



UvA-DARE (Digital Academic Repository)

Between empowerment and manipulation

The ethics and regulation of for-profit health apps

Sax, M.

Publication date

2021

Document Version

Final published version

[Link to publication](#)

Citation for published version (APA):

Sax, M. (2021). *Between empowerment and manipulation: The ethics and regulation of for-profit health apps*.

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.



Between Empowerment and Manipulation

**The Ethics and Regulation
of For-Profit Health Apps**

Marijn Sax

Between Empowerment and Manipulation

The Ethics and Regulation of For-Profit Health Apps

Marijn Sax

© Marijn Sax, 2021

ISBN: 978-94-93184-78-7

Coverontwerp: Paul Middelkoop

Layout: Guus Gijben | Proefschrift-aio.nl

Drukkerij: Proefschrift-aio.nl

Between Empowerment and Manipulation

The Ethics and Regulation of For-Profit Health Apps

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor

aan de Universiteit van Amsterdam

op gezag van de Rector Magnificus

prof. dr. ir. K.I.J. Maex

ten overstaan van een door het College voor Promoties ingestelde commissie,

in het openbaar te verdedigen

op vrijdag 26 maart 2021, te 16:00 uur

door Marijn Sax

geboren te Utrecht

Promotiecommissie

Promotores:	Prof. dr. N.H. Helberger	Universiteit van Amsterdam
	Prof. dr. B. Roessler	Universiteit van Amsterdam
Copromotor:	Dr. N. Bol	Tilburg University
Overige leden:	Prof. dr. H. Nissenbaum	Cornell Tech
	Prof. dr. H.-W. Micklitz	European University Institute
	Prof. dr. T. Blanke	Universiteit van Amsterdam
	Prof. dr. J.C.M. van Weert	Universiteit van Amsterdam
	Prof. dr. M.M.M. van Eechoud	Universiteit van Amsterdam
	Dr. T.R.V. Nys	Universiteit van Amsterdam

Faculteit der Rechtsgeleerdheid

Table of Contents

Acknowledgements	11
Introduction	15
1. Health Apps and the Digital Society	15
2. Outline of the Dissertation: On the Ethical Evaluation and Legal Regulation of For-Profit Health Apps	19
3. Why This Is Not a Dissertation about Privacy	22
4. Methodology: On Combining Ethics and Law	23
Chapter 1 - Health (Apps) for Everyone	27
1. Introduction	27
2. One Simple But Important Observation: Most Popular Health Apps Are Commercial Services	31
3. Health Apps and the (Freemium) App Economy	34
4. The Design of Health Apps: Where Technology Meets Business	38
4.1. How Silicon Valley Became Addicted to Behavioral Economics	38
4.2. More Data, Both in Terms of Quantity and Quality	42
4.3. Analyzing and Using All Those Data	44
4.4. Example of Proto-Hypernudging: Garmin Connect	48
5. Conclusion	52
Chapter 2 - 'Apps That Will Make You Happier and Healthier:' The Strategic Framing of Health in Health Apps	55
1. Introduction	55
2. Health as a Contested Concept	57
2.1. Health as Biological Normality	57
2.2. Nordenfelt's Holistic-Humanistic Perspective	58
3. The Contemporary Culture of Health	61
4. Health, Wellness, and Healthy Lifestyles According to Popular Health Apps	68
4.1. Fitbit: Every Moment Matters and Every Bit Makes a Big Impact	68
4.2. MyFitnessPal: How Even a 'Simple' Calorie Counter Is Secretly about Everything You Do	71
5. Health Is Like Sex, It Sells	77
6. Conclusion	79

**Chapter 3 - Autonomy, Vulnerability, Trust, and Manipulation: Key Concepts
for the Ethical Evaluation of For-Profit Health Apps** 83

1. Introduction	83
2. Theorizing Autonomy	86
2.1. Autonomy as an Uncontroversial Starting Point in Liberal Democracies	86
2.2. Autonomy: A Brief Conceptual Sketch	87
2.3. From Ideal Theory to Normative Reconstruction: An Example from the Health App Context	89
2.4. Relational Autonomy	91
2.5. Reconciling Procedural and Relational Autonomy	93
2.6. Conclusion	95
3. Vulnerability and Trust	96
3.1. Vulnerability and Autonomy	97
3.2. A Taxonomy of Vulnerability	97
3.2.1. Sources of Vulnerability	98
3.2.2. States of Vulnerability	100
3.3. Vulnerability and Trust	101
3.4. Taking Stock Before We Move On to Manipulation	106
4. Manipulation	107
4.1. Introduction	107
4.2. Moralized Versus Non-Moralized Conceptions of Manipulation	108
4.3. Conceptualizing Manipulation	109
4.4. Two Possible Objections: The Means of Manipulation and Hiddenness	113
4.4.1. Means of Manipulation	113
4.4.2. Hiddenness of Manipulation	113
4.5. Manipulation as a Success Concept: Individual Acts of Manipulation versus Manipulative Environments	115
5. Manipulation and Autonomy (, and Vulnerability and Trust)	117
6. Power	118

Chapter 4 - The Vague But Important Line between Supportive and Manipulative Digital Health Environments	123
1. Introduction	123
2. The Question of Autonomy in the Context of Health Apps	125
3. The Helpful Side of the Vague Line: How Autonomy, Trust, and Attention or Vulnerabilities Ideally Come Together	128
3.1. Vulnerabilities	129
3.2. Trust	131
3.3. Conclusion	133
4. The Other Side of the Vague Line: Health Apps as Manipulative Digital Environments	134
4.1. Conditions for Manipulation in Health Apps	135
4.1.1. Helping the Bottom Line by Helping the User?	136
4.1.2. Relationship-building with Users: Targeting and Exploitation of Vulnerabilities	140
4.1.3. Relationship-building with Users: Trust, Trustworthiness, and Hidden Influences	144
4.2. Conclusion: The Manipulative Potential of For-Profit Health Apps	148
5. Manipulative Health Apps and User Autonomy	149
Chapter 5 - Health Apps and Unfair Commercial Practices	153
1. Introduction	153
2. The Unfair Commercial Practices Directive	155
2.1. General Aims: Autonomy & Fairness	155
2.2. Structure of the Directive	158
2.3. Scope of the Directive: Are Health Apps Covered?	160
2.4. The Average Consumer & Vulnerabilities – The Average Health App User as a Vulnerable Consumer?	165
2.5. Misleading Commercial Practices	173
2.5.1. How Applicable Is the Current Concept of Misleading Commercial Practices to the Health App Context?	177
2.5.2. Native Advertising in Health Apps	177
2.5.3. Health Apps as Free Services & Lack of Transparency Concerning Commercial Intent	180
2.5.4. Towards a More Ambitious Interpretation of the Concept of Misleading Commercial Practices	183
2.6. Aggressive Commercial Practices	189

2.6.1 Using the Concept of Undue Influence to Address the Manipulative Potential of Health Apps	194
2.7. Annex I – The Blacklist	200
2.8. The General Clause on Unfair Commercial Practices	203
3. Conclusion	206
Conclusion	211
1. The Importance of the Relational Perspective	211
2. Manipulation as a Key Concept to Understand and Evaluate Digital (Health) Environments	212
3. Health Apps: Between Empowerment and Manipulation	213
4. Addressing Manipulative Health App Practices	214
5. Suggestions for Improvement	215
References	219
Summary	233
Samenvatting	239

Acknowledgements

When asked, every researcher claims to love doing interdisciplinary research. In practice, however, interdisciplinary research requires hard work and turns out to be pretty challenging. So when doing a PhD, as a philosopher, at an information law institute, with supervisors from Law, Philosophy, and Communication Science, one cannot simply assume that the process of writing a PhD under those conditions will be smooth sailing. A plethora of methodologies, approaches, perspectives, questions, writing styles, and vocabularies have to somehow be combined into one coherent whole. In the end, interdisciplinary research only works when the people involved are willing to put in the effort, listen to each other, and learn from each other.

I am therefore lucky to have had you, Natali, Beate, and Nadine, as my supervisors. From the very start, all of you made a genuine effort to explore – together with me – how to combine different approaches into one PhD dissertation. At times it seemed like each of you was *more* interested in the other supervisors' respective disciplines than in your own discipline. Our discussions were never focused on the guarding of boundaries associated with particular academic disciplines, but only on one thing: how to craft the best argument, period. I am deeply grateful for this truly interdisciplinary research experience, made possible by you, Natali, Beate, and Nadine.

This brings me to the Institute for Information Law – IViR – which has been my academic home for the past four years. To an outsider, doing a PhD as a philosopher at IViR might seem a bit odd. But to me it never felt odd. IViR is much more than an institute for information law. Not only do the non-lawyers (almost) outnumber the lawyers; IViR lawyers are also the type of lawyers who always want to collaborate with non-lawyers. Who would have thought lawyers could be so much fun? I want to thank *everyone* at IViR from the bottom of my heart for providing me with such a warm, stimulating, exciting, fun workplace.

Also a big shout-out to Tom, Brahim, Joanna, Judith, and Claes, my stable ASCoR connection over the years and my Personalized Communication collaborators 'from the other side.' Apart from being good company, you also showed tremendous patience in explaining communication science research to a philosopher.

The Philosophy and Public Affairs (PPA) Capacity Group of the UvA continued to feel like home, even though my desk was at the Law Faculty. As much as I have come to like ‘the lawyers,’ I always enjoyed attending the PPA reading groups and PPA colloquium to discuss philosophy, which is and will always remain the academic discipline I love most.

I am also very grateful to Helen Nissenbaum for hosting me as a visiting researcher at The Digital Life Initiative at Cornell Tech and for giving me such a warm welcome in New York City. I had a fabulous time, which would not have been possible without Michael, Eran, Jake, Erica, Elizabeth, Yvonne, Lauren, and Jessie.

Marjolein, my partner in crime in the academic conference business. Our back-to-back trips to Copenhagen and Washington, DC were pretty educational, eh? Which city has good hipster micro-breweries, uh, I mean, which city has the most relevant academic conference next year?

Jef and Paddy, my IViR neighbors on the other side of the paper-thin walls which enclose our offices. I heard laughter coming from your room around 47 times a day. Both of you enriched my vocabulary. Jef, thanks for teaching me Flemish: Frigo, isomo, vogelenpik, valies, supappeke. Paddy: grimmige shit, grote speler.

Philosophy is important and beautiful, but not as important and beautiful as friendship. Paul, Yorian, Mark, Anne, Isabel, Martijn, Anoeek, Evelien, Sietske, Isa, Erik, Joris, Florean, Berend, Bas, Margot, Lorenzo, Cassandra, Eduard, Igor, Wessel, Steven, Maurits, and Max, thanks for making my life outside of academia fun and exciting.

Thanks also to my parents, for always stimulating me to be interested in the world, ask questions, and for telling me to pursue whatever career I wanted to.

Jitske, thank you for being you, *vreg* and all. You always call bullshit when I launch into a semi-coherent monologue on whatever. Bickering with you is the best thing in the world and a better exercise in critical thinking than any kind of academic education could ever provide.

Introduction

1. Health Apps and the Digital Society

This is a dissertation about health apps. At the same time, it is a dissertation about a larger, more general development of which health apps are but a symptom. I want to start this introduction by briefly sketching this larger, more general development, before I explain how health apps fit into this larger picture.

We are currently living in a – or: the – digital society. It is hard to tell when the analogue society stopped being analogue and started being digital. I am not interested in putting a precise date on the start of the *digital* society. If we stop for a moment and examine our own lives, we find enough evidence to support the claim that right now – in the summer of 2020 – we are living in a digital society. Many of our daily activities are enabled by, or even take place in, digital environments: our communications with friends, family, colleagues, strangers, commercial entities, and government agencies; our shopping; our news and media consumption; the planning of and navigation during transportation; the managing of our agendas; our search for advice or instructions; and so on.

One important feature of the digital society, in which an increasing amount of daily activities (partly) take place in digital environments, is the fact that the designers and operators of these digital environments become increasingly powerful actors. They are *choice architects* who design and operate *choice architectures*. Thaler and Sunstein (2008) popularized these terms, as well as the accompanying proposal that based on insights from behavioral economics and psychology, choice architectures can be designed with behavioral change in mind. By understanding and anticipating the various biases that people tend to exhibit, choice architects can design their environments to tap into these biases to steer or shape the behavior of people. As a result, choice architects are powerful actors because they can direct people's decisions and behavior towards particular ends or outcomes, often without the 'inhabitants' of choice architectures noticing that their behavior is deliberately and intentionally being shaped by other actors. This power can, of course, be used responsibly – and it often is. But this power also requires careful and continuous scrutiny, especially in the digital society.

In the old days, choice architects had to study and understand human behavior to design a physical choice architecture that was then put in place, permanently. In the digital society, things are different. As people live their lives in and through digital environments,¹ all their interactions with(in) those environments are turned into data. These user data can be used to analyze the functioning of digital environments *in real time*: which type of user tends to interact with which elements of the digital environment, and how do different users respond to different design choices or stimuli? The resulting insights can be used to – sometimes dynamically and in real time – adjust the digital environment to see whether the desired patterns of behavior emerge. Put simply: digital choice architectures can in principle be run as continuous experiments in behavior adjustment or optimization. These experiments are often barely transparent to the people who serve as the ‘test subjects.’ Many people know at an abstract level that the digital environments they use study their behavior and change their environments in response to the generated insights. But even to the trained eye, it often remains unclear *what* digital environments know about their users, *how* they use that knowledge, and for *which* purpose.

The potential to shape behavior through digital choice architectures has received an increasing amount of attention in recent years. The challenge has been framed as one of “dark patterns” (Gray *et al.* 2018) or as the engineering of human behavior and, even more extreme, the engineering of human nature (Frischmann & Selinger 2018). Zuboff (2015, 2019) has developed a line of research around the concept of ‘surveillance capitalism’ and its logic which leads to large-scale data collection and data-driven attempts to perfectly predict the behavior of consumers by *creating* said behavior. Lastly, Susser, Roessler, and Nissenbaum (2019a, 2019b) have offered an elaborate reappraisal and interpretation of the concept of manipulation for the ethical evaluation of the influences digital choice architectures can exert on people’s behavior. As different as these approaches are, they all stem from the shared concern that digital choice architectures can influence and shape human behavior in a highly effective yet undesirable manner.

This dissertation is motivated by the same concern and is inspired by (among others) the perspectives mentioned above. I take the potential for behavior modification through digital choice architectures by a less-than-transparent means to be a key challenge in the digital society. More specifically, the key challenge for me as an ethicist is to help evaluate when and how behavior should be allowed to be influenced in the digital society where the digital choice architectures available

1 I use the terms ‘digital environments’ and ‘digital choice architectures’ interchangeably in the rest of this dissertation.

can exert increasingly subtle, difficult-to-understand yet efficacious influences on people. After all, influencing the behavior of other people is not good or bad per se. What we need, then, is a set of ethical norms or a framework to help us differentiate the legitimate attempts to influence behavior through digital choice architectures from the illegitimate ones. Moreover, how can and should the law respond to the normative tensions introduced by digital choice architectures? How can ethical approaches to digital choice environments be translated into *legal* solutions?

To address this abstract challenge of the ethics and regulation of behavior modification in the digital society, I choose to focus on a more specific phenomenon, namely health apps. Focusing on health apps is interesting for at least two reasons. First, by focusing on *one* phenomenon, I can describe and analyze in much more detail how the larger challenge I seek to address plays out on the ground, in the real world. This will allow me to offer a more grounded, and, hopefully, more relevant analysis. Second, and maybe even more importantly, health apps are incredibly interesting in their own right. As a value, health is universally desired and health is – to varying degrees of course – a concern to every human being. Health, in short, *matters to everyone*. It follows that focusing on digital *health* environments, which seek to influence the *health*(-related) behavior of persons, matters too. Moreover, health apps come with a promise of *empowerment*. As smartphones have become mainstream consumer products, health apps are seen as easily accessible tools that put health-related behavior advice and tracking in the hands of many people. The fact that health apps are installed on one's *personal* device and can collect *personal* information also allows for the *personalization* of the services offered by health apps. As a result, their services and functions can be catered to specific individuals, making them potentially more relevant. It should come as no surprise then that the European Commission (2012, 2014) and the World Health Organization (2011, 2015, 2016a, 2016b, 2018a, 2018b, 2019a, 2019b, 2020) have emphasized the role mobile health technology can play in securing positive health outcomes for both individuals and societies at large. In this respect, the promise of health apps mirrors the promise of digital choice architectures more generally: the provision of smart, dynamically adjustable digital environments that can deliver to their users the most 'relevant' suggestions, information, and services.

Still, health apps also embody the larger challenges posed by digital choice architectures in the digital society. To help me tease out the tensions introduced by data-driven dynamically adjustable digital health environments, I choose to focus on a subset of health apps: popular for-profit health apps. This subset covers nearly all widely used health apps. The choice for popular for-profit health apps is not an arbitrary one. As will become clear in the following chapters, for-profit health apps

introduce a difficult but interesting tension, which also sheds light on the larger, more general challenge of digital choice architectures in the digital society. On the one hand, popular for-profit health apps offer advice and support on something that is important and universally desired: health. In this regard, they can be helpful or even empowering. On the other hand, these apps also seek to exploit² the users' desire and need for health for their own financial gain. In order to do so, popular for-profit health apps tend to offer advanced – i.e., data-driven (dynamically) adjustable – digital health environments within which they seek to build ongoing, *profitable relationships* with users. Health sells, but to what extent can appeals to health be used to shape the (economic) behavior of people in ways that benefit the health app providers before such a practice becomes problematically exploitative?

Health apps direct our attention to another important perspective on the issue of choice and behavior modification in the digital society. Providers of commercial digital choice environments (such as health apps) want users to keep coming back to these environments. If people keep coming back, the providers of such environments can 'get to know' these returning users, which can make it easier to steer their behaviors, for example, by developing persuasion profiles (Kaptein 2015, Kaptein *et al.* 2015). Another way of putting this is to say that commercial digital choice environments want to build ongoing *relationships* with their users, and these ongoing relationships – rather than individual, separate interactions of a user with a digital choice environment – should be the focus of our analysis. We need a structural – or: relational – approach to choice in digital environments.

Admittedly, the term 'relationship' may feel a bit strong in this context. Can we really say that people can become involved in an ongoing *relationship* with an app, in the same way they become involved in a relationship with their partner or their friends? It should be clear that developing a relationship with another person is not the same as developing a relationship with a data-driven service. I use the term 'relationship' to refer to the fact that as people keep coming back to an app or a digital service, *something* grows over time. Put differently, there is a difference between a user returning (again and again) to an app, and a completely new user. Precisely because digital choice architectures can, and often do, collect user data, they can also learn about returning users. Returning users, on their side, 'store value' in an app and can, depending on their experiences, have changing motivations and expectations (Eyal 2014). The intricate tracking of users over time, coupled with the possibility to *personalize* the environment based on knowledge

2 I use the word exploit in the neutral sense here. In later chapters (especially chapters 4 and 5) we will consider whether and when this 'exploitation' is in fact problematic.

gained from tracking, is a game changer. In what follows, I choose to use the term ‘relationship’ to refer to this progressive dynamic between users and health apps and in the digital society.

My overall aims in this dissertation are to: 1) develop a conceptual framework that helps me evaluate – ethically – when the influences digital health environments seek to exert on their users should be considered legitimate, and when such attempts cross a line and should be considered illegitimate, and 2) to use that framework to explore how current regulation – and more specifically, European unfair commercial practice law – could be (re)interpreted in order to address ethically illegitimate commercial health app practices. I will argue that popular for-profit health apps, and more generally digital environments like health apps, can harbor a manipulative potential. I want to emphasize the word ‘potential,’ because I will not be arguing that popular for-profit health apps are always or necessarily manipulative. I will argue, however, that there is a *potential* for health apps to exploit their users’ desire for health in a manipulative manner, rendering the practices of such health apps ethically illegitimate. Moreover, I will argue that the European Union’s Unfair Commercial Practices Directive (UCPD), when informed by my ethical analysis, offers promising legal solutions to manipulation worries in consumer environments such as popular for-profit health apps.

In the remainder of this introduction I will do three things. First, I will provide an outline of the five chapters that follow to summarize the main argument of this dissertation. Second, I will briefly explain why I chose not to focus on privacy, which is a ‘usual suspect’ in this context. Third and last, I will discuss my methodology and the relation between my ethical and legal analysis.

2. Outline of the Dissertation: On the Ethical Evaluation and Legal Regulation of For-Profit Health Apps

Before I can develop my ethical framework, I use **chapter 1** to introduce the phenomenon of health apps. Here I discuss the technological capabilities of health apps as well as the larger app economy of which they are a part. We need to understand the dynamics of the app economy to understand how the technological capabilities of health apps – and the digital environments they offer to users – are designed and operated. The contemporary app economy incentivizes app developers to build *ongoing relationships* with users and monetize those relationships *over time*. To successfully monetize one’s userbase, there is an incentive for health apps to design their digital health environment to optimize user retention, user

engagement, and user conversion. Health apps, like other apps, can use their near complete power over the design and operation of the digital environments they offer to try to give shape to the user-app relationships. For example, health apps can collect user data in order to learn more about their users. These insights can be used to constantly experiment with and update the digital environment.

In **chapter 2**, I explore the notion of ‘health’ and how it is used in the health app context. It is clear that health is universally desired. When given the choice, everyone prefers health over illness. However, when we ask what it means to be healthy, it turns out that health is a contested concept. This conceptual fluidity might be regrettable to philosophers aiming for conceptual rigidity, but it can also be put to strategic use by health apps: it affords them the opportunity to allude to health’s universally desired nature while remaining vague regarding their contributions to their users’ health. In this chapter, I argue that health apps construct and propagate particular health discourses and narratives which serve their business interests. So to understand why health apps are designed and run the way they are, we must not only understand their business models and technological affordances, but also their portrayal of health.

In **chapter 3**, I provide the conceptual foundation for the ethical framework I use to evaluate the commercial practices of for-profit health apps. I build on four interrelated concepts: autonomy, vulnerability, trust, and manipulation. Why these four concepts? Put simply, autonomy is the value that is at stake in digital environments that shape human behavior. The concept of manipulation is promising because it can help explain *how* and *why* autonomy can be at stake in the context of health apps. Lastly, vulnerability and trust serve as important conceptual links between autonomy and manipulation. I mainly build on Susser, Roessler, and Nissenbaum (2019a, 2019b) to argue that, in essence, manipulation is about the intentional infiltration of a manipulee’s decision-making in order to dispose the manipulee to the manipulator’s ends, by way of targeting known or inferred weaknesses of said manipulee in a manner that is not meant to be (fully) transparent to the manipulee. Put simply, when a person is influenced in a manipulative manner, their autonomy is threatened. Theories of vulnerability can help us get a better, more nuanced understanding of how people’s weaknesses and biases can be exploited by manipulative practices. Attention for the role of trust is important because health apps (and digital for-profit services more generally) seek to build trust with users, since it helps to build profitable relationships with them. Although trust in technology can be helpful, it can also be exploited in a manipulative manner.

In **chapter 4**, I use the conceptual framework from chapter 3 to provide an ethical evaluation of the (commercial) practices of health apps. To understand the manipulative potential of popular for-profit health apps and to understand how the autonomy of health app users can be at stake, we must not only focus on isolated one-off interactions between a health app and its users. It is equally important to analyze how health apps use their control over the digital health environment they offer to build *ongoing* relationships with their users over time. In the first part of the chapter I explore how health apps could *scaffold* their users' autonomy (Heath & Anderson 2010). The collection and analysis of user data to learn about users' vulnerabilities can, for instance, be used to better anticipate what individual users need. Moreover, insights into vulnerabilities and preferences can be used to engineer feelings of trust towards a health app, making it easier for the app in question to help users. We can thus imagine that under (somewhat) ideal conditions, health apps can empower their users. In the second part of the chapter I explore the manipulative potential of health apps. The properties of health apps discussed in the first part of the chapter, can, under different conditions, turn from being empowering into being manipulative. An important source for health apps' manipulative potential is the freemium app economy's emphasis on optimizing user retention, user engagement, and user conversion. When a health app engineers its digital health environment to identify and target vulnerabilities of users for the purpose of engineering profitable relationships which will optimize retention, engagement, and conversion rates, most ingredients for manipulation are present. In practice, the line between empowerment and manipulation can be difficult to draw. Those properties that make health apps potentially useful also put them in a position of power vis-à-vis their users, and that position of power can also be *misused* for commercial practices that are manipulative in nature.

Chapter 5 is devoted to the question of how we can address – legally – the difficult tension between empowerment and manipulation. I focus on the UCPD and present an interpretation of the Directive that is informed by my ethical analysis. At its core, the UCPD is concerned with protecting the possibilities for autonomous decision-making of consumers. I try to show how a more relational understanding of autonomy, coupled with a more nuanced understanding of vulnerability, can and should inform our reading of the Directive. Theoretically, the concept of misleading commercial practices seems to align rather well with the challenges posed by manipulative practices; in cases of manipulation, manipulees are often (but not necessarily) “led astray” (Noggle 1996: 44) or, in other words, misled. While the concept of misleading practices and the corresponding call for better transparency requirements are helpful, I argue that the concept of aggressive commercial practices – and more specifically the concept of undue influence – seems (even)

more promising. By focusing on the position of power of health app providers and the possible *exploitation* of that position of power by way of targeting exploitable characteristics of users, we can get much closer to the *heart* of the manipulation problem. The overall conclusion of the chapter is that if we are willing to inform our interpretations of key concepts in the Directive with my ethical analysis, the UCPD can be a promising legal tool to address manipulative practices in digital (health) environments.

3. Why This Is Not a Dissertation about Privacy

If we understand privacy to be about the control people have over access to things – i.e., places, information, decisions – that are important to them, then it is easy to see why health apps naturally give rise to privacy concerns.³ Health apps tend to collect large amounts of data, and some of those data may even encode sensitive information that is related to people’s health. So it is only natural to ask how health app users can control access to their (health-related) data, as well as to their decisional sphere. Even more so when we consider how respect for privacy is constitutive of people’s personal autonomy (Benn 1984, Reiman 1994, Roessler 2005). By affording people the ability to control access to those things that are important to them, people are also given the opportunity to exercise control over how they present themselves and how they make their own decisions (Roessler 2005). Privacy, in short, is important, and in no way do I mean to suggest that questions concerning privacy should not be asked in the health app context. Still, I have made the decision not to base my main argument on the function and value of privacy.

As I explained above, I want to focus on the potential of digital choice architectures to – often subtly – manipulate our behavior. If we frame this problematic in terms of privacy, we would fail to get to the *heart* of the problem. To be sure, the problem of privacy in the digital society is important in this context, but it is only a secondary problem. Data-driven digital choice architectures need data to learn about their users and, if required, make adjustments to steer the behavior of users. Privacy, in this context, is important because it can function as a sort of shield that can limit the flow of data, thereby depriving digital choice architectures of

3 Nissenbaum (2010) has developed an influential critique on control-access theories of privacy, proposing the ‘contextual integrity’ approach as a better way to capture the challenges posed to privacy by technological developments. From a contextual integrity perspective, data-hungry for-profit health apps also give rise to privacy concerns to the extent that the logics and norms of the market and the health(care) domain clash.

the ‘fuel’ they need to learn about their users and make adjustments (Calo 2016). But notice how by focusing on this – important – role of privacy, we leave the real problem unaddressed. Privacy helps us understand and manage one of the major *preconditions* – i.e., availability of personal data – for manipulation, but it does not address the question of manipulation in the digital society *itself*.⁴

Much to the credit of privacy advocates worldwide, privacy has become a widely shared (although not universal) public and political concern. This increased attention for privacy has also come with what I want to call a ‘privacy reflex:’ as soon as a service collects, processes, or stores data, the immediate response is to ask about the implications for privacy. The same has already happened for health apps.⁵ This is a good thing, because privacy is important. Privacy, however, is not the only useful perspective to address challenges posed by technology. Because the ‘privacy reflex’ already secures ample attention for the privacy implications of health apps, I choose to focus on slightly different perspectives to, in the end, contribute to a more complete understanding of the impact of technology on our lives.

4. Methodology: On Combining Ethics and Law

In this dissertation I combine ethical and legal approaches. The ethical approach serves as the foundation of this dissertation. I do not only want to describe and explain the phenomenon of health apps (and digital choice architectures more generally); I also want to offer the reader guidance on how to *evaluate* them. To evaluate a phenomenon, one needs *standards of evaluation*. This is where ethics – or more generally put: normative theory – comes in. In chapters 3, 4, and 5 I formulate and use conceptions of autonomy and manipulation.⁶ These concepts provide us with criteria to evaluate when we should consider (commercial) health app practices as legitimate, because they respect or even promote autonomy, and when we should consider them illegitimate, because they are manipulative in nature and undermine (the possibilities for practicing) autonomy.

4 Lanzing (2018, 2019) has argued, correctly I think, that we can use the concept of *decisional* privacy to begin to address manipulation concerns. Still, by combining theories of manipulation and autonomy, as I do in this dissertation, we can develop a much more comprehensive approach to behavior modification through digital choice architectures in the digital society.

5 See, e.g., Huckvale *et al.* 2015, Huckvale, Torous & Larsen 2019, Mulder 2019, O’Loughlin *et al.* 2019.

6 I also work with the concepts of trust and vulnerability. Although these concepts are essential to my story, they do not necessarily serve as evaluative standards themselves. I mainly use these concepts to better understand and describe how autonomy and manipulation are at stake in the digital society.

In the field of normative philosophy, the distinction between ideal theory and non-ideal theory describes two types of methodologies. Running the risk of oversimplifying, the distinction can be understood as follows. When doing ideal theory, one starts from abstract principles, ideals, and values to formulate one's theory, while bracketing actual socio-economic conditions. Once one has formulated one's – ideal – theory, the theory is, in a next step, applied to the *real* – i.e., non-ideal – world. Rawls (1971) famously applied this methodology, and has received criticism for it (see, e.g., Mills 2005). An argument in favor of ideal theory is that we (sometimes) first need to formulate the values and ideals we want to strive for, before we can see how our non-ideal reality can be changed for the better. Non-ideal theory, on the other hand, holds that we should take note of the non-ideal circumstances of our societies from the very start, because only then can we devise theories that help us understand and address the challenges we *actually* face. Running the risk of oversimplifying again, ideal theory can be considered more of a top-down approach where theory-building comes first and application to empirical reality comes second, whereas non-ideal theory can be characterized as a bottom-up approach, where empirical reality *informs* theory-building from the very start.

Ideal theory and non-ideal theory can themselves be considered to be ideal types; that is, not every method of doing normative philosophy falls neatly into either of these categories. But as ideal types, they do help to typify one's approaches. When looking at the approach I adopt in this dissertation, I try to avoid doing ideal theory by taking into consideration, from the start, how health apps function and are used in the real world. I discuss characteristics and practices of real health apps and use those to tease out normative tensions that follow from the real-world design and use of health apps. These normative tensions I observe in real-life cases inform my conceptualization and use of concepts like autonomy, manipulation, vulnerability, and trust.

The attentive reader may object that such an approach can never be *completely* bottom-up, for how can we recognize normative tensions in the first place without having *some* preconceived notion of which values and principles are at stake? I can only bite the bullet and acknowledge that I come to this topic with *some* preconceived notions of what, for example, autonomy and manipulation are about, roughly speaking. Seen from this perspective, my methodology is not one of *pure*⁷ non-ideal theory. I do not regard this as a problem, if it can be called a problem at all. In the end, my aim is simply to be mindful of the empirical reality I seek to capture and evaluate with my ethical approach. In practice, this means that I do

7 To non-ideal theorists, the notion of *pure* non-ideal theory will, I suspect, sound peculiar.

come to this topic with some theoretical and normative preconceptions, but my aim is to *test* and, when necessary, *readjust* them in the face of the challenges the digital society confronts us with. For example, in chapter 3 I start with a relatively basic, uncontroversial conception of autonomy to then, over the course of the chapter, explore how this basic conception should be adjusted to be able to account for the types of challenges health apps give rise to.

The next methodological question to address is how my ethical approach interacts with the legal objectives of this dissertation. The perspective I adopt is an *external* one. I build on both the text of the UCDP and secondary literature to capture the general aims and structure of the Directive. I am not, however, interested in describing the law as accurately and elaborately as possible, or to provide an exhaustive overview of relevant case law. As an ethicist, my aim is to explore how my external ethical perspective can inform our interpretations of legal concepts in the UCPD. This implies that I see legislation – and the system of law as a whole – as a set of rules that require interpretations, and there are no ‘necessarily correct’ interpretations that are embedded in the rules themselves, just waiting to be captured. However, even though the practice of interpretation is, in essence, open-ended, it does not follow that legal interpretation is a practice where anything goes. Interpretations admit to degrees of plausibility. For instance, an interpretation that destabilizes the coherency of an entire area of law, or, more specifically, a Directive or Regulation, is generally thought to be less plausible than an interpretation that does not upset the entire constitution of an area of law. Seen from this perspective, the practice of interpretation resembles Quine’s (1992) theory of underdetermination of evidence and theory.

My objective, then, is to suggest interpretations of key concepts in UCPD that I think can best address the ethical concerns we face. But while doing so, I do try to be mindful of the structure, language, and objective of the Directive because I want my ethically inspired interpretations to be at the very least minimally plausible to legal professionals.

Chapter 1

Health (Apps) for Everyone

1. Introduction

Health is important. It always has been, and it always will be. Health, one could say, is what enables a long and prosperous life, and is universally desired. The way we think about health can differ, however, between both times and cultures. In our current Western societies, the culture of health takes a particular form. As Devisch (2013: 21) explains, one used to be healthy until one fell ill, meaning that in the absence of a clear illness one was seen as healthy. Now, Devisch (2013: 21) argues, this logic has been reversed: we are presumed to be *unhealthy* until we have *proven* that we are healthy. As a result, health is increasingly seen as something that can and should be self-managed and optimized (Cederström & Spicer 2015). Health is a process, a journey, which requires the right *lifestyle*. Seen from this perspective, the popularity of health apps is understandable. They are seen as easy to use, accessible technological tools of *empowerment* that will help people manage and improve their own health and lifestyle. Moreover, health apps are seen as promising technology which can help to achieve positive societal outcomes, for instance by lowering the public costs of health care⁸ due to lower rates of illness and, relatedly, by ensuring a more productive workforce.

As a result, we see that health apps are propagated as a promising technology by many different types of actors. On the policy level, for instance, the European Commission (2012, 2014) and the World Health Organization (2011) have emphasized the role mobile health technology can play in securing positive health outcomes for both individuals and societies at large. The WHO has even published a series of handbooks called ‘Be He@lthy Be Mobile’ on how to design and implement mobile health solutions (WHO 2015, 2016a, 2016b, 2018a, 2018b, 2019a, 2019b, 2020). The European Commission argues that health apps can streamline prevention and allow for a stronger focus on quality of life solutions,

⁸ See, e.g., IQVIA 2017

can make health care more efficient and sustainable, and can empower people by allowing for “a more participative role while enhancing their responsibility over their own health” (European Commission 2014: 5).

We also see a strong interest in health apps among employers. Many popular health apps offer corporate wellness programs, such as Fitbit’s ‘Health Solutions’ for employers⁹ and Garmin’s ‘Corporate Wellness’ programs.¹⁰ One company that has adopted a corporate wellness program is the Dutch bank ING. Employees receive a smartwatch which allows them to keep track of their health and wellness progress. The data that the smartwatch generates is fed into an app that calculates a personalized ‘Wellness Quotient.’ The program has even become a part of the formal annual evaluation of employees: “At its core, the project is aimed at making staff more energetic and happier to produce a more effective sales force. ING salespeople do not have revenue targets and some of their Key Performance Indicators solely focus on personal goals, such as the wellness quotient” (Business Insider 2017). Put shortly: health and wellness effectively become Key Performance Indicators on the basis of which employees are evaluated. To be a good employee, you must be both happy and healthy.

Health insurers also show a great interest in health apps. In the Netherlands, major health insurer Zilveren Kruis has developed a mobile health platform and accompanying app, both called *Actify*.¹¹ Interestingly, both the online platform and the app have undergone major changes over the years, as Zilveren Kruis continues to search for the best and most appropriate way to deliver health and lifestyle advice to the Dutch population. Actify started out as an app that had to be linked with a commercial fitness wearable in order to participate in ‘health challenges’ to earn points. These points could be redeemed in the Actify Shop, which contained some products (vaguely) related to one’s health, but mostly products not clearly related to health, such as suitcases, car sat navs, headphones, frying pans, and high-quality 100% leather aprons. In 2018, Zilveren Kruis closed the Actify Shop because the insurer no longer deemed this reward scheme “appropriate”, due to the uneasiness many users felt over the mixture of commerce and health.¹² (I

9 <https://healthsolutions.fitbit.com/employers/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/mLZq7gz>.

10 <https://www.garmin.com/en-US/health/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/UtJPaCD>.

11 <https://www.actify.nl/> (last accessed 22 September 2020).

12 The Actify shop was been taken down by after Zilveren Kruis in August 2018, after Zilveren Kruis reconsidered the “appropriateness” of this reward scheme (See: <https://www.smarthealth.nl/2018/08/23/webshop-actify-sluit-liever-lopen-voor-het-goede-doel-dan-voor-korting-op-een-waterkoker/>, last accessed 22 September 2020). Screenshots of the now closed webshop are available here: <https://imgur.com/a/e2x38YU>.

suggested back in 2017 that Actify’s reward scheme was not entirely unproblematic, see Sax 2017). Zilveren Kruis then rebranded Actify as a mindfulness platform and offered an (new) accompanying Actify Mindfulness¹³ app, which it claimed offered scientifically proven mindfulness solutions.¹⁴ This new mindfulness app was later taken down as well. After yet another round of rebranding, Actify is now (in the summer of 2020) available as a ‘vitality coach app’, which is presented as an all-purpose solution to help its users adopt a healthier lifestyle. For this new direction of the Actify platform, they teamed up with renowned professor in psychology and behavioral economics Dan Ariely to offer behavior change solutions that are proven to work.¹⁵

Why should this brief history of a Dutch health insurer’s attempts to develop a health platform with accompanying health app be of interest of anyone? Actify itself is not necessarily interesting to an international audience, but the brief history of Actify does exemplify many of the more general questions and challenges that I want to address in this dissertation. For example, Actify’s development from simple step counter into a mindfulness app and finally into an all-purpose healthy lifestyle app, shows the various ways in which the value, concept, and function of health can be understood: health can be about one’s level of physical activity and fitness, about one’s mental well-being, and about one’s (even) more general patterns of living. The abandonment of Actify’s first project (i.e., of rewarding users for steps by offering discounts on products from partner brands predominantly unrelated to health) shows the sometimes uneasy combination of pursuing commercial interests and, at the same time, offering health-related advice and services to people. Lastly, Zilveren Kruis’ persistence in trying to find new approaches and technological solutions to support people in their pursuit of health clearly shows the potential they – and with them, many other actors – see in health apps.

The promise of user empowerment through health apps of course also coincides with technological developments. Almost everyone owns a smartphone, apps can be installed in a matter of seconds, and smartphones contain many sensors which

13 The Actify Mindfulness app was later taken down by Zilveren Kruis. Screenshots of the app’s description in Google’s Play Store can be found here: <https://imgur.com/a/YOj9a9R>.

14 “Would you like to relax more? Reduce stress and be more aware of yourself? Actify Mindfulness takes you on a journey to a more mindful life. Multiple studies show the effect of mindfulness. The professors who conducted these studies have also developed the mindfulness and meditation exercises. So you learn from the masters!” (Translated from Dutch by me). Screenshot available here: <https://imgur.com/a/YOj9a9R>.

15 The message describing the collaboration with Dan Ariely is currently (June 2020) on their homepage (<https://www.actify.nl/>, last accessed 22 September 2020). A screenshot of the message is available here: <https://imgur.com/a/kSRZkj6>.

can collect useful data. Since people tend to use their smartphone rather intensively throughout the day, smartphones offer many opportunities to learn about (the habits of) people and, moreover, offer many opportunities to reach and influence people. These technological circumstances further explain why many different actors understandably see health apps as the perfect tools to empower people in the adoption of ‘healthy lifestyles.’

Despite the promises of empowerment, critical questions concerning health apps have already been raised. Most of the questions commonly asked tend to focus on the *effectiveness*¹⁶ of health apps (“Do they *really* help people to live healthier?”) or on the *data practices*¹⁷ of health apps (“What types of sensitive data do health apps collect, what do they use these data for, and is doing so legally permissible?”). In many of these discussions, the health apps themselves are presumed to be (largely) neutral technological tools; health apps may or may not be *effective* in actually achieving positive health outcomes, but at the very least they are often presumed to be tools that only try to achieve positive health outcomes. In this dissertation, I will question this presumed neutrality of health apps as tools that simply aim to achieve positive health outcomes. More specifically, I emphasize that all the popular, widely used health apps are commercial services. It is only when we understand and analyze health apps as commercial services that we are able to see certain *tensions* they give rise to. On the one hand, popular for-profit health apps can serve as helpful tools of user empowerment. On the other hand, the technological affordances of health apps *combined* with the dictates and incentives of the app economy also introduce a manipulative potential as the value and importance of health can be *misused*. To better understand and evaluate the manipulative potential of health apps, the remainder of this chapter explores the (health) app economy and the technological affordances of health apps. In later chapters I will build on these insights into the app economy and technological affordance of health apps to address the tension between empowerment and manipulation found in popular for-profit health apps (and commercial digital choice architectures more generally).

16 See, e.g., Murnane, Huffaker & Kossinets 2015, Wang *et al.* 2016, Whitehead & Seaton 2016, Wang, Varma & Proserpi 2018, Wisniewski *et al.* 2019,

17 Huckvale *et al.* 2015, Huckvale, Torous & Larsen 2019, Mulder 2019, O’Loughlin *et al.* 2019.

2. One Simple But Important Observation: Most Popular Health Apps Are Commercial Services

Nearly all consumer-facing health apps in existence can be downloaded from Apple's App Store or Google's Play Store. Only a very small number of health apps – typically those that are designed to be used in a very specific *medical* context – are not acquired in one of the two major app stores. Such *medical* health apps usually do not have many users, because they are aimed at a very specific population and are, moreover, not distributed through publically accessible channels. In this dissertation, I will focus on consumer-facing health apps that can be acquired in app stores and brand themselves as being about, or in important ways related to, health. I do so for two reasons. First, most of the people that use health apps use consumer-facing health apps. Second, and maybe even more importantly, consumer-facing for-profit health apps introduce specific tensions and raise specific questions that I think need to be addressed.

App stores are thoroughly commercial contexts. All of the popular health apps that are used by millions of people (e.g., Fitbit, MyFitnessPal, Strava, Garmin Connect, Lifesum, Headspace, Calm, Runkeeper, Endomondo, Runtastic) are commercial services, built on the basis of business models that ultimately serve the purpose of generating profits. The users of these health apps should thus be considered *consumers*. These popular for-profit health apps do not just try to influence the health-related behavior of their users, but also, at the same time, try to influence their economic behavior (Sax, Helberger & Bol 2018). Moreover, there is often no clear delineation between the health content and motives, and commercial content and motives in popular health apps – health and commerce are deliberately merged in subtle ways. The reason for this is simple: Health, like sex, sells.¹⁸ It is universally desired and at the very same time is shrouded in an aura of innocence – who could possibly be *against* health?!

18 See, for instance, this article called ‘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)’ Last accessed 22 September 2020 at <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/>, screenshot available here: <https://imgur.com/a/aUL9qDJ>. “If you're going to make just 1 New Year's Resolution this year, then it HAS to involve feeding 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?”

To see how popular health apps are operated as commercial services, consider the example of MyFitnessPal. Under Armour acquired this very popular health app together with Endomondo and MapMyFitness for around \$710 million.¹⁹ In July 2018, their jobs page listed a vacancy for a new Account Executive at MyFitnessPal. Among the responsibilities of the new account executive are:

- “Proactively prospect, qualify, grow, and maintain a national account list that includes CPG [customer packaged goods] and QSR [quick service restaurants] companies and regional agencies”;
- “Sell current portfolio of digital ad solutions across MapMyFitness, MyFitnessPal and Endomondo – including mobile and web display media, and custom Challenge activation”;
- “Develop thoughtful and tailored POVs [points of view] for clients relative to UA’s [Under Armour’s] leadership in health and wellness”;
- “Meet and exceed revenue and sales targets”;
- “Coordinate with pre and post-sales teams at Under Armour to maximize revenue potential with optimal efficiency”;
- “Negotiate pricing and contractual agreements with clients and agencies”; and
- “Manage activity using Salesforce.com.”²⁰

This list of responsibilities clearly shows that a health app like MyFitnessPal is actively run like a company – which is unsurprising because it *is* a company. They develop portfolios with “ad solutions” to best match their health content to those who are interested in buying ads. They actively scout the food and customer packaged goods markets to see whether there are companies they can target with their current portfolio – and if not, how the portfolio should be modified to attract those companies to these ad solutions. Another example is the app Headspace, which is a mindfulness and meditation app with “one mission: to make the world a happier and healthier place.”²¹ They looked to hire a new Senior Product Analyst that could “[d]rive projects to identify key levers for new user growth, retention, and engagement”, “[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.”²² Examples like these particular

19 <http://fortune.com/2016/04/05/under-armour-apps-healthier/> (last accessed 22 September 2020).

20 Job posting has since been removed. Screenshot available here: <https://imgur.com/a/IVRfx8e>.

21 <https://www.headspace.com/about-us> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/7rjYNLV>.

22 Job posting has since been removed. Screenshot available here: <https://imgur.com/a/uIwpvds>.

job positions clearly show that some apps are focused on maximizing revenue by exploring how apps can be optimized for that purpose through continuously testing and implementing small, profitable adjustments.

It is precisely because popular health apps like MyFitnessPal and Headspace are run like companies (because they are companies), that one should also try to understand and analyze them as such. It is only from this perspective that one is able to understand the rationales behind the commercial practices employed.

The reason I want to emphasize the commercial practices of health apps is that they give rise to an important challenge. Through their commercial practices as profit-seeking services, health apps run the risk of *misusing* the importance of health to people and society to influence – or maybe we should say *manipulate* – users' (economic) behavior in problematic ways. My proposal thus is to focus on the commercial practices of health apps and to ask where to draw the line between, on the one hand, 'business as usual' in the health app economy, and, on the other hand, ethically illegitimate commercial health app practices which target the consumers' concern for their health in an exploitative manner. It is important to emphasize at the outset of this dissertation that I am *not* going to argue that any type of commercial practice by a health app is problematic per se. The health app is not the first phenomenon in human history that introduces or enables the (further) commercialization of health (see, e.g., Navarro 1976, Barsky 1988, Hesse-Biber 1996, Mol 2008). Health has been commercialized for decades and it is not the phenomenon of the commercialization of health which I seek to investigate. What I do want to call attention to is the particular ways in which health apps can – and sometimes do – engage in commercial practices within a health context. Some – but certainly not all – of these commercial health app practices are ethically illegitimate.

In order to provide more structure to this inquiry, I propose we first look into the (health) app economy, the popular business models in the app economy, and the commercial practices those business models give rise to.

3. Health Apps and the (Freemium) App Economy

Every popular health app is built – and continuously tweaked – on the basis of an underlying business model. Interestingly, both Apple²³ and Google²⁴ offer ‘how-to’ guides to prospective app developers.²⁵ They both emphasize how essential it is to develop one’s app around a business model. For example, on its developer page Apple writes: “[t]he right business model for your app will balance your goals with your target market’s expectations. Consider choosing a business model before you start developing your app so that you can build it seamlessly into the user experience.”²⁶ Google suggests that “[t]o maximize your revenue, consider multiple monetization models for your app.”²⁷ So, which types of monetization models – the core of app business models – can one consider? Both Apple and Google provide a list, which largely overlap and cover roughly the same monetization models.²⁸

First of all, the most simple monetization model is to ask users to pay a one-time fee before being allowed to download and use the app. Secondly, an app developer can choose to sell space for advertising that is shown to the user within the app. Google helpfully suggests to consider “native ads that allow you to match ads to your app’s look and feel.”²⁹ Thirdly, in-app purchases can be offered to users, for instance to unlock additional content or features. Physical or digital items can of course also be offered in an in-app shop. The option of in-app purchases is often present in apps that are free to download. Apple calls this the freemium model: “[t]his model allows users to get something great at no cost and have the option to pay if they want to

23 <https://developer.apple.com/app-store/business-models/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/26l9COh>.

24 <https://developer.android.com/distribute/best-practices/earn/monetization-options> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/27MHERZ>.

25 I got the idea to look at Apple’s and Google’s pages for developers from Fahy, Van Hoboken & Van Eijk 2018.

26 <https://developer.apple.com/app-store/business-models/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/26l9COh>.

27 <https://developer.android.com/distribute/best-practices/earn/monetization-options> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/27MHERZ>.

28 Apple calls these models the free model, freemium model, subscription model, paid model, and paymium model. Google lists five types of potential income, namely in-app purchases, subscriptions, advertising, paid apps, and e-commerce.

29 <https://developer.android.com/distribute/best-practices/earn/monetization-options> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/27MHERZ>.

enhance their experience or engage more deeply.”³⁰ Fourthly, app developers can opt for a subscription model where users pay a recurring monthly (or weekly, or daily) fee in order to gain and retain access to either the entire app, or to some additional content or features.

The four monetization options mentioned above are not necessarily mutually exclusive – Google even suggests combining them, as we saw above. For example, an app could contain both (native) advertising and the option of in-app purchases. Or users could have to pay a fee to download an app that also contains (native) advertising or in-app purchases, or both. Note, however, that not all combinations of monetization options are as likely as others. Usually, an app that costs money in an app store does not show advertising and does not require additional in-app purchases or subscriptions to gain access to all content or features. Reversely, apps that can be downloaded free of charge usually do contain (native) advertising, in-app purchases, or both.

Although different monetization strategies can inform the business models of health apps, some general incentives and related strategies will be present regardless of the particular monetization strategies that are adopted. A simple observation which is true for all business models is that it is all about making sure one has as many 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, or selling subscriptions, in all cases more users means more potential to generate income.

Virtually all of the widely used health apps are freemium apps (e.g., Fitbit, MyFitnessPal, Strava, Garmin Connect, Lifesum, Headspace, Calm, Runkeeper, Endomondo, Runtastic). The freemium model comes with its own peculiarities, which are captured by Apple’s insistence on engagement as “the path to monetization” in the freemium model. Apple explicitly advises freemium app developers to “prepare for the long term” because the freemium model works best when one is able to engage users for extended periods of time.³¹ Successful freemium apps thus try to *build profitable on-going relationships* with their users. A freemium app developer tries to build apps to which users keep coming back and on which they spend as much time as possible. Over time, as users keep coming back to an app, the likelihood increases that their

30 <https://developer.apple.com/app-store/business-models/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/26l9COh>. A more in-depth advice page on monetization in the freemium model can be found here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKlNlD>.

31 See Apple’s advice page here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKlNlD>.

behavior can be influenced in ways that profit the app provider.³² As an app developer quoted by Apple explains, “[w]e don’t just look at short-term retention, we look at very, very long-term retention.” People must, in the words of Eyal (2014), get hooked on an app.³³ So when using the freemium model, amassing large numbers of users is just the start. For freemium apps, it is equally important to *persuade* users to display the kinds of behavior that are profitable to the app provider, such as: making users come back to the app to build an audience to serve (native) advertising to; persuading users to make use of the in-app purchase options (either for single transactions, or for subscriptions); and persuading users to use integrated social features so more potential consumers can be reached (who, in turn, may be persuaded to display behaviors that are profitable for the app provider). Apple prefers to use the term ‘engagement’ rather than ‘persuasion’: “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.”³⁴

Apple also advises that “[s]uccessful freemium apps have analytics built into the experience so that developers can understand user preferences and continually improve the apps.”³⁵ Put differently, the freemium monetization model comes with an incentive to collect data about one’s users, generate insights from those user data, and use those insights to engineer continuous engagement by *targeting* users and by (dynamically) adjusting different elements of the digital environment offered to users. Google also emphasizes the need to implement and use tools for analytics because “different audiences have different preferences.”³⁶ It is essential to *know* one’s users to understand how the ‘user experience’ can be ‘optimized’ for every type of user or even for every individual user. One of Google’s “best practices for optimizing your monetization strategy” is therefore to “[a]nalyze statistics in the Google Play Console. Gain an understanding of your users’ purchasing patterns, including statistics on average revenue per paying user and new vs.

32 Remember, users do not actually need to spend money to be considered ‘profitable’ to the app provider. When a health app sells advertising space and/or monetizes user data, the sheer fact that people keep coming back to spend time on an app serves the app’s bottom line.

33 There are many self-help-style books and articles for (aspiring) app developers that explain how to build profitable apps. Nearly all of those books and articles emphasize the potential and importance of the freemium model (see, e.g., Wooldridge & Schneider 2011, Kumar 2014, Seufert 2014)

34 <https://developer.apple.com/app-store/freemium-business-model/>

(last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/ahnBwmv>.

35 See Apple’s advice page here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKINID>.

36 <https://developer.android.com/distribute/best-practices/earn/monetization-options>

(last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/27MHERZ>.

returning buyers.”³⁷ Other suggestions include to “[r]un pricing experiments. Use in-app A/B testing to experiment with different prices and encourage users to take up subscriptions or purchase in-app products.”³⁸

We can summarize the central insights from this section as follows: 1) Nearly all popular for-profit health apps use a freemium business model which necessitates the careful development and nurturing of the *commercial relationship* between the user and the app; and 2) health app providers should use their *control* over the digital choice architectures they offer to users to *shape* the commercial relationships with their users in the manner they desire (i.e., optimize for user retention, engagement, and conversion).

I want to focus on two different but equally important ways in which health apps seek to carefully develop their commercial relationships with their users. First of all, I will focus on health app designs and features, and the technological capabilities and data practices (which are often seen as ‘fueling’ these technological capabilities) of health apps. Here one can think of, for instance, the application of insights from behavioral economics, personalized targeting and persuasion profiling (Kaptein *et al.* 2015), as well as practices that could be called ‘hypernudging’ (Yeung 2017). This ‘technological element’ of health app business models will be elaborated upon in this chapter.

Secondly, I will address an element of health app business models that is often overlooked but no less important, namely the carefully crafted health discourses one can encounter in health apps, and how these discourses connect with prevailing cultures of health. Health is framed in very *particular* ways in popular health apps, and I will argue that strategic, commercial considerations inform these deliberate framing practices. I devote an entire chapter (chapter 2) to this ‘discourse element’ of health app business models, because this perspective receives little attention in the literature and therefore requires an elaborate and careful analysis.

Together, the ‘technology perspective’ and the ‘health discourse perspective’ allow for a comprehensive understanding of health app business models. Due to this combination of perspectives, the technological affordances which structure the commercial practices of health apps can, at the very same time, be understood within their socio-cultural context.

37 <https://developer.android.com/distribute/best-practices/earn/monetization-options> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/27MHERZ>.

38 <https://developer.android.com/distribute/best-practices/earn/monetization-options> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/27MHERZ>.

4. The Design of Health Apps: Where Technology Meets Business

In this section, I will discuss those technological affordances of health apps that are most relevant to my argument. More specifically, I will discuss how the availability of (1) increasing amounts of data, (2) analysis and targeting technology, and (3) insights from behavioral economics, can together explain why and how health apps can persuasively influence their users – now and even more so in the near future. I will start by discussing how insights from behavioral economics increasingly get transposed to health apps (and digital services more generally). Applying these insights to the general design and function of health apps can make apps much more engaging and persuasive. Next, I will explain how the availability and clever usage of data can ‘supercharge’ these general insights into engaging app designs and features, by using data to profile users, and personalize targeting, content, and design.

4.1. How Silicon Valley Became Addicted to Behavioral Economics

Before we dive into the world of data collection, data analytics, persuasion profiling, and hypernudging, it is important to emphasize that the shaping of the commercial relationship between user and health app is not solely informed by sophisticated applications of technology. An equally important factor is the application of recent insights from the literature on behavioral economics which can already inform app designs and features at a very general level. Authors like Thaler and Sunstein (2008)³⁹ and Kahneman (2011) have popularized behavioral economics and have inspired many app developers and others working in Silicon Valley to apply insights into how and why people “predictably err” (Thaler & Sunstein 2008: 7) in app designs (e.g., Berman 2015, Lohr 2016, Conick 2017, Lewis 2017, Shaw 2017).

Thaler and Sunstein build on Kahneman’s dual process theory, which claims that people have two cognitive ‘systems’, often simply called System 1 and System 2 (‘The Automatic System’ and ‘The Reflective System’ in Thaler and Sunstein (2008: 19)). They explain that although we often feel like we make deliberate, rational, calculated choices (System 2, the Reflective System), in reality much of what we end up doing is the result of System 1 – the Automatic System – which Thaler and Sunstein (2008: 20) describe as the uncontrolled, effortless, associative, fast, unconscious, skilled system, that allows us to do many things in an intuitive manner without having to think a lot (or at all). This Automatic System is very

39 Both Thaler and Sunstein have published many works on nudging. See, e.g., Sunstein & Thaler 2003, Thaler & Sunstein 2008, Sunstein 2014, Sunstein 2015a, Sunstein 2015b, Sunstein 2016.

useful, because we simply do not have the cognitive resources to always carefully deliberate about *everything*. To help us navigate the world without having to think about everything, System 1 makes use of a wide range of heuristics that get the job done most of the time. However, precisely because System 1 functions in an automatic fashion and through the use of heuristics, it also makes us susceptible to various biases (e.g., status quo bias, availability bias, self-serving bias). As a result, we “predictably err” (Thaler & Sunstein 2008: 7) because these mental shortcuts we use intuitively or automatically do not always deliver the results we, upon closer critical inspection, prefer or endorse.

Thaler and Sunstein couple these insights from behavioral economics with a specific policy proposal which they famously called ‘nudging:’ the adaptation of choice architectures based on knowledge of how System 1 functions, in order to get people to act in their own best interest by exploiting System 1’s heuristics and biases, without explicitly *forcing* them to act in a particular way. There is, one could say, a *descriptive* and a *normative* side to the nudging story as told by Thaler and Sunstein. The descriptive part of their work focuses on dual process theory and how, based on that knowledge, choice architectures can be adapted to achieve, or be more likely to achieve, particular outcomes. It is this descriptive part I am mainly interested in and when I use the term ‘nudging.’ The normative part of their nudging story focuses on the question of whether nudges have a special normative status given their properties. It is especially this normative part of Thaler and Sunstein’s story that has sparked lively philosophical debates on the ethics of nudging (see, e.g., Bovens 2009, Anderson 2010, Yeung 2012, Hansen & Jespersen 2013, Saghai 2013, Wilkinson 2013, Rebonato 2014, Nys & Engelen 2017, Engelen & Nys 2019). As interesting as these debates are, I will not (explicitly) engage with them. In what follows, I am interested in the descriptive part of the nudge story and will consider nudging to be just one of the many ways in which people’s behavior can be influenced or steered. I will not assume that influencing people with the use of nudges has – by definition – special normative implications.

For our purposes, it suffices to emphasize that behavioral economics has become a great inspiration for Silicon Valley.⁴⁰ The central insight from Thaler and Sunstein’s and Kahneman’s work is that choice architects can play a decisive role in influencing and steering the behavior that will be displayed within the choice architectures that they design. Developers developing a health app are literal choice architects

40 See Dow Schüll 2014 for a good example of an analysis of carefully crafted technological designs aimed at ensuring as much (profitable) consumer engagement as possible.

designing choice architectures. Moreover, they are very *powerful* choice architects since they can manipulate nearly all aspects of the design of, and features in, the app (i.e., the choice architecture).

The book *Hooked* by Eyal (2014) is a great example of how insights from behavioral economics are translated to app designs. *Hooked* is written as a self-help book for app developers and explains how one can use behavioral economics to design apps that users get ‘hooked’ on. Most of the techniques Eyal describes are not based on technologically advanced, big data-fueled analytics, but rather on careful thinking about how to get app users into “The Habit Zone” through the use of very general design principles that exploit general knowledge about the heuristics and biases of System 1 (Eyal 2014: 29-31). Eyal develops what he calls “The Hook Model”. Put shortly, Eyal suggests that a well-designed app that ‘hooks’ users is based on a cyclical design that provides triggers that motivate users to perform actions. These actions should be linked to *variable* rewards. Actions are thus performed *in anticipation of* a reward, and this reward should be variable so as to ensure that users keep on being interested in performing the action – an app that offers very predictable rewards soon becomes dull. Next, users should be seduced to invest time, energy, and – ultimately – money in the app. When users do this, they are – in Eyal’s words – “storing value” into an app, which “increases the likelihood they will use it again in the future” (Eyal 2014: 145). Lastly, this stored value can be used to “load the next trigger” which can start the ‘trigger, action, variable reward, investment’ cycle again (Eyal 2014: 154). As The Hook Model shows, it is not just *technological* sophistication that can determine the persuasiveness of apps’ attempts to influence user behavior. Using insights from behavioral economics at the level of the general app design already explains part of the potential persuasiveness of apps.

Another example can be found on a popular website for ‘user experience’ (UX) professionals. It features an article⁴¹ called ‘Top 5 Behavioral Economics Principles for Designers’, describing the insights gained from a multi-day workshop with Dan Ariely, to create better (i.e., more engaging) app designs. Ariely is a big star in the world of behavioral economics and has done much to popularize the field.⁴² One recommendation is to focus on defaulting, which also happens to be one of Thaler and Sunstein’s favorite examples. “The combination of loss aversion with mindless

41 <https://uxplanet.org/the-top-5-behavioural-economics-principles-for-designers-ea22a16a4020> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/eNbWjUP>.

42 Ariely has written popular books such as *Predictably Irrational* (2008), *The Upside of Irrationality* (2010), and *The (Honest) Truth About Dishonesty* (2012). He also is the co-founder and director of the Center for Advanced Hindsight and has a special on Netflix called *(Dis)Honesty: The Truth About Lies*.

choosing implies that if an option is designated as the “default,” it will attract a large market share. [...] In many contexts defaults have some extra nudging power because consumers may feel, rightly or wrongly, that the default options come with an implicit endorsement from the default setter” (2008: 35). Predictably, the UX advice is to “Set the default to the favoured outcome you want the user to take.”⁴³ The mindfulness app Headspace is a good example. When you start using the app, the default program that is placed right at the top of your screen (the most salient place) is a 10-day program. Given the way monetization in health apps works, this is understandable. By making a multiday program the default, the app immediately tries to draw you in, letting you store value, which increases the chance you (and your eyeballs) will return to the app.

Another recommendation revolves around ‘friction costs.’ The idea is, again, simple: the more friction there is, the more effort it requires to get something done. Therefore, the deliberate elimination or introduction of friction can steer users’ behavior. “Direct users or help them complete a task by removing small barriers. Conversely, add small barriers to hinder undesirable behavior.”⁴⁴ The calorie counter (and recipe and lifestyle) app MyFitnessPal makes use of this principle. Users of calorie counting apps will tell you how difficult it can be to insert all the right information for home cooked meals.⁴⁵ That is why MyFitnessPal offers a lot of recipes through its blog, which is prominently integrated into the app. These recipes often come with a ‘Log It’ button, which will automatically record all the correct caloric information in your diary, thereby taking away a lot of ‘friction’ in the food registration process. MyFitnessPal is thus actively steering users towards its blog and the recipes on there. This makes great sense from a business perspective since the blog contains a lot of native advertising, as do the recipes which are often offered in collaboration with commercial third parties such as popular food blogs. More users looking at the blog means more revenue.

The widely prevalent share button also functions to lower the friction costs of sharing content with others. Be it a recipe you like, a workout you have just completed, or a meditation tip you value, there will always be a share button conveniently present to help you to quickly inform your friends or followers.

43 <https://uxplanet.org/the-top-5-behavioural-economics-principles-for-designers-ea22a16a4020> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/eNbWjUP>.

44 <https://uxplanet.org/the-top-5-behavioural-economics-principles-for-designers-ea22a16a4020> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/eNbWjUP>.

45 Say you cooked pasta for 4 people and used some mushrooms, zucchini, bell pepper, onion, garlic, tomatoes, and some parmesan. It would be difficult to determine how much of each of these separate ingredients you have personally eaten.

From the perspective of business models this is, again, understandable since all monetization works through eyeballs, and the more content is shared, the more eyeballs will be attracted to your app's content. Moreover, the share button also fits the use of "social nudges" (Thaler & Sunstein 2008: 53-71) by making it easier to involve others in our own health app experiences. With others involved, health apps emphasize how many people have seen your progress, achievement or activity. Health apps also use your friends and their health endeavors as "primes" (Thaler & Sunstein 2008: 69) to influence the likeliness of you displaying similar behavior. Strava, for instance, will by default inform you of your friends' running or cycling activities. Social sharing features can also introduce the spotlight effect – one's feeling "that others are closely paying attention to what they are doing" (Thaler & Sunstein 2008: 60-62). Given the social status that being in great health and living healthily bestow on individuals, the spotlight effect can be employed to increase the chance that people feel that they have to 'perform' and 'prove' their health in the correct manner (more on this in chapter 2). The health app with all its social sharing features is a convenient vehicle to do so.

The above-mentioned examples of insights from behavioral economics being applied to health apps at the general design level are not meant to form a comprehensive list. The examples show the many ways in which health app designs and features can be carefully tweaked to steer behavior. Moreover, it has hopefully become clear by now that these insights are very popular among – and easily available to – health app developers. In the next section, I will explain how these behavioral economics principles can be *supercharged* with intensive data collection and subsequent clever applications of data. First, I briefly address how increasing amounts of data can be collected, before turning to the usage of those data.

4.2. More Data, Both in Terms of Quantity and Quality

The new smartphone models you can buy right now have more different types of sensors in them than the smartphone models you could buy five years ago. Take, for example, the Global Positioning System (GPS) functionality. GPS used to be an exciting new feature, something to show off to friends. Today it is (close to) impossible to buy a smartphone without GPS functionality. A more recent example is the heartrate monitor. At the time of writing – 2020 – a built-in heartrate monitor is still a rather unique smartphone feature. It might very well be the case, however, that in a few years from now, many more smartphones will come equipped with such a sensor.

As smartphones come equipped with ever more sensors, they can collect both *more* data (quantity) and a *wider variety* of data (quality). This development is further exacerbated by other factors. First of all, the increased storage capabilities allow more data to be stored. Not only do phones come with more internal storage, but there is also an increasing availability of cloud storage solutions which means the user is no longer bound by the internal storage limit of a device. Secondly, we tend to carry and use our smartphones (nearly) everywhere, while making use of all kinds of different apps. As a result, there are simply more opportunities for apps to collect data on their users in different kinds of locations and situations. Thirdly, Bluetooth (yet another smartphone functionality that used to be cutting-edge but that is a default nowadays) can be used to connect other devices to smartphones, such as wearables. Wearables that are connected to the smartphone (usually via a special app) add new sensors to the smartphone ecosystem such as accelerometers, barometric altimeters, and heartrate monitors. Many fitness bracelets are even designed to be worn while sleeping and taking a bath or shower. On its 'Wear & Care' webpage, Fitbit explicitly states that "Our products are meant to be worn all day and night."⁴⁶ So the introduction of wearables also allows for (even more) continuous collection of data.

Put simply, the smartphone and wearable technology that is currently available makes for a perfect infrastructure to collect and store a wide range of different, continuously updated (user) data. Because mobile technology like smartphones and wearables now blend so naturally into almost all parts of our everyday life, it should also be emphasized that data on nearly all of our *behaviors* can and will be collected. The importance of the datafication of many of our behaviors should not be underestimated, as we will see in the next section where I discuss what can be done with all these data (Van Dijck 2014, Mau 2019). Before we turn to the next section, it is insightful to briefly consider the fact that health apps themselves are also acutely aware of the enormous amounts of data they (can) collect. On their jobs page MyFitnessPal emphasizes how exciting it is to work for them as a data scientist by stating that "MyFitnessPal has the largest database of human eating habits in the world. The opportunities for a data scientist here are almost endless."⁴⁷ It is to these 'almost endless opportunities' that I now turn.

46 <https://www.fitbit.com/productcare> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/CU5jvxz>.

47 <https://www.myfitnesspal.com/jobs> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/KVo0EWc>.

4.3. Analyzing and Using All Those Data

Collecting and storing endless amounts of data is one thing, but making use of those data is something else. With sensors to *collect* different types of data, smartphones are also increasingly capable when it comes to *analyzing* data. The computing power of modern smartphones enables increasingly sophisticated computational tasks.⁴⁸ From the perspective of analyzing data, smartphones' ability to connect to other devices and servers via 3/4/5G, Wi-Fi, or Bluetooth is interesting as well. Data can be sent to and received from cloud services. With these cloud services, data can either be stored and/or *analyzed* elsewhere (namely 'in the cloud'). As powerful as smartphones have become, cloud computing services offer many more possibilities for data analytics than smartphones themselves do. First of all, cloud computing services simply offer more computing power. Secondly, this additional computing power is interesting because 'in the cloud' data can be easily recombined with other available data. This is where the promise⁴⁹ of big data enters the scene. By being able to draw on enormous datasets consisting of – often – recombined data, the promise is that new, non-trivial insights can be extracted (or generated) out of the data (Rubinstein 2013: 76). Or as Tene and Polonetsky (2013: 259) put it: “the big data business model is antithetical to data minimization. It incentivizes collection of more data for longer periods of time. It is aimed precisely at those unanticipated secondary uses, the “crown jewels” of big data.” Since popular health apps have literally millions of users, all of which generate data, the promise is that many new, interesting, profitable insights can be extracted out of the available (health) data. Or as MyFitnessPal put it on their aforementioned jobs page: “the opportunities for a data scientist here are almost endless.”⁵⁰

So, we have established that lots of data can be collected and that there is an abundance of computing power available, even for relatively simple health apps. What is more interesting, however, is that these data can be 'put to work' in increasingly interesting ways. Put shortly, these data can be used to devise, test, and optimize increasingly persuasive digital choice architectures. Let me explain this in more detail below.

48 To use an example from personal experience: At the time of writing, I am using the OnePlus 3 as my phone. This phone contains a Qualcomm Snapdragon 820 quadcore Central Processing Unit (CPU) of 2.2 GHz, and 6 Gigabytes of Random Access Memory (RAM). To put this into perspective: the desktop computer I wrote both my bachelor theses on (in 2010 and 2011) had a slower dual core CPU and only 2 GB of RAM. It also was many times bigger than my current smartphone.

49 See Sax 2016 for a critical analysis of this promise.

50 <https://www.myfitnesspal.com/jobs> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/KVo0EWc>.

To start, a simple fact should be observed. The fact that people take their smartphone wherever they go, and wear their wearable at all times, not only means that there are ample opportunities to collect data. It also means that people can be *reached* nearly everywhere and all the time via their smartphones and wearables. Due to this simple fact, the persuasive potential of health apps dramatically increases. Kaptein *et al.* (2015: 38) observe that “[t]o be effective, persuasive systems should deliver the right message, at the right time, in the right way.” There are more opportunities to ‘get it right’ when you can reach your target at nearly all times.

Getting it right, however, requires much more. The reason health apps (and other apps, and websites, and social media platforms, and so on) are so data-hungry, is that data can be used to infer all kinds of information about both individual users and groups of users (see, e.g., Turow 2011). With the help of data, one can find out what makes users ‘tick’ (Kaptein 2015). And of course, the more insights one can gain about what makes individual users or particular (sub)groups ‘tick,’ the more effective the profiling can be. Kaptein *et al.* (2015: 38), for instance, discuss adaptive persuasion profiles, which are: “systems that adapt the message, the timing, and the persuasive approach to the situation at hand.” The more data one has available to profile users, the better the persuasive approach can be made to ‘fit’ the individual or (sub)group one is trying to influence.

To test whether a persuasive approach is a good ‘fit,’ app developers often make use of A/B testing to test whether a new feature has the desired effect. Big data analytics can (but does not have to) be used to suggest interesting correlations between, for instance, user characteristics and responsiveness to design features that warrant further testing. Such testing follows the big data mindset of ‘if it works, it works’ (Mayer-Schönberger & Cukier 2013: 13). Usually, the goal is not to understand *why* a particular design feature works or not. App developers simply test a lot of different things and whatever seems most effective is implemented.

One potential worry is that the huge amount of data available, coupled with analytic capacities, will also allow health app providers to find out about – and then target – specific weaknesses or vulnerabilities of individuals or subgroups of the population. Consider the following example. In May 2017, a confidential Facebook document was leaked (The Australian 2017). This document detailed how Facebook was working on algorithms that “can determine, and allow advertisers to pinpoint, ‘moments when young people need a confidence boost’” (Machkovech 2017). Moreover, the document showed “a particular interest in helping advertisers target moments in which young users are interested in ‘looking good and body confidence’ or ‘working out and losing weight’” (Machkovech 2017). Put shortly: Facebook was,

according to the internal documents, building a system that allowed it to identify temporarily experienced vulnerabilities *in real time* and allow advertisers to target users, based on those vulnerabilities, *in real time* as well. As one can imagine, advertisers would be very keen on using such a system, since presenting individual users with a message or offer that taps directly into the insecurities, fears, or doubts they experience during these moments of vulnerabilities, will certainly be much more persuasive than a generic message or offer.

The Facebook example can be considered quite shocking, partly because it involves the explicit targeting of young people's vulnerabilities. It should be noted, however, that the general idea behind this Facebook experiment, and the mechanism through which it works, does not differ substantially from the types of persuasive systems that many apps and other platforms (try to) build. Yeung (2017) coins the term 'hypernudging' to refer to systems like the one in the Facebook example, or the type of systems Kaptein (2015) and Kaptein *et al.* (2015) write about. As Yeung explains: "[...] unlike the static Nudges popularized by Thaler and Sunstein (2008) such as placing a salad in front of the lasagna to encourage healthy eating, Big Data analytics nudges are extremely powerful and potent due to their networked, continuously updated, dynamic and pervasive nature (hence 'hypernudge')" (Yeung 2017: 118). She continues: "Big Data-driven nudges make it possible for automatic enforcement to take place *dynamically* (Degli Eposti 2014), with both the standard and its execution being continuously updated and refined within a networked environment that enables real-time data feed which, crucially, can be used to *personalize* algorithmic outputs (Rieder 2015)" (Yeung 2017: 122).

The possibility of hypernudging (or behavioral targeting practices that are not quite hypernudging yet but already come close to it) is also further enabled by the ecosystem within which health apps operate and which offers built-in optimization tools. Take the already mentioned Android developers platform. Google offers Google Play Analytics⁵¹ and Google Analytics for Firebase⁵² to app developers to help them analyze how users interact with their apps.⁵³ On their page detailing the features of Google Analytics for Firebase, Google explains that:

51 <https://support.google.com/googleplay/android-developer/answer/139628>

(last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/JA4CHNM>.

52 <https://developer.android.com/distribute/best-practices/earn/improve-conversions>

(last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/uQcgtO>.

53 <https://developer.android.com/distribute/best-practices/earn/monetization-options>

(last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/27MHERZ>.

getting users to complete key actions, or *conversions*, is your main success metric. By identifying whether users are encountering roadblocks before converting, or which app areas have high exit rates, you can identify opportunities to improve your app experience and increase conversion rates.⁵⁴

Google also offers some best practices, some of which come very close to the suggestion to implement hypernudge-like features:

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.

[...]

Improve your acquisition workflow. Use integration with AdWords 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.

[...]

View and analyze your Google Analytics for Firebase data in real time with BigQuery. Combine all of your raw, unsampled data with custom data for BigQuery for additional analysis.⁵⁵

These very explicit and public suggestions made by Google show that futuristic sounding concepts like ‘persuasion profiling’ and ‘hypernudging’ are becoming an easily accessible option for app developers. (Google now also offers Firebase as an affordable off-the-shelf tool for any business seeking to datafy their operations.)⁵⁶ The interactions of every single user with an app can be tracked and analyzed – in real time – by the app provider. Tweaks to the app that are either generic, fully personalized, or something in between (e.g., tweaks to the app experience of a

54 <https://developer.android.com/distribute/best-practices/earn/improve-conversions> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/uQcgtoO>.

55 <https://developer.android.com/distribute/best-practices/earn/improve-conversions> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/uQcgtoO>.

56 See <https://firebase.google.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/2QuVdxU>.

group of users that share characteristics X, Y, and Z) can be suggested based on the real-time analytics. As a result, app providers have the tools to test *for every single user* how the app experience should be customized to allow for optimal ‘conversion.’ Moreover, tools are available to app providers to dynamically and in real time change the app experience for every individual user, based on what they predict will lead to optimal conversion.

It is important to differentiate between two ways of using (real-time) analysis of user behavior. First, (real-time) analysis of user behavior can be used to measure the effects on a personal and group level of (very) small tweaks of the app. Think of the A/B testing of font sizes, colors used, subtle differences in language or prices, and so on. Such very small tweaks can lead to increases in conversion, even though it may not always be clear why. (This is a typical example of the big data mindset: explore, test a lot, and just see if it works *for some reason*). These are usually the types of practices that are encouraged and enabled by Google Firebase and Apple’s developers’ built-in tools and which do not require much independent programming on the part of the app developers. Second, there are the more elaborate, app-specific features which depend on personalized feedback to users. These come closer to what Yeung has in mind when she talks about hypernudging. Now, we can question whether actual hypernudging of the kind Yeung describes is already happening within apps. What we already can observe, however, are proto-hypernudging systems that already come very close to it. Let me present one such example.

4.4. Example of Proto-Hypernudging: Garmin Connect

Shortly after starting working on my PhD project I got myself a fitness bracelet (the Garmin vivosmart HR) in order to get acquainted with the phenomenon I was researching. This wearable connects to my smartphone with an app called ‘Garmin Connect.’ In this app all the data that are collected about me are visualized in various graphs and tables. The wearable continuously records my heartrate, the amount of steps I take, the amount of stairs I ascend and descend, the amount of calories I burn, and the intensity of my activities, so there are a lot of data that can be visualized. Interestingly, the app does not merely visualize the data, but it also actively analyzes my data to establish my patterns of behavior, and it uses these inferred patterns to support me in living healthy. The function called ‘Insights’ is especially revealing in this regard. The app describes this function in the following way:

Welcome to Insights

Insights: Go Beyond Tracking

You know it's important to track your health data, but how does it all tie together? Insights give you real-time analysis of your patterns and help you get where you want to go. You'll see how you compare to your own past routines and to other people like you. Insights will appear periodically when Garmin Connect detects a pattern that might be useful to you.

shows a sample chart

More Data Means Better Insights

It's no surprise; the more information you have, the easier it is to find patterns. To get the best picture of your health, sleep and movement, wear your Garmin device all day, every day... and night. With more data, you'll get more meaningful Insights to help you identify areas you'd like to improve and meet goals you're working toward.

The Garmin Connect app is a good example of an app that makes use of many of the technological capabilities I described above. It facilitates the internet of things-like interaction between smartphone, wearable, and cloud. It makes use of all kinds of sensors – heartrate monitor, accelerometer, altimeter, GPS – to collect data. It can perform rather advanced tasks like inferring patterns from the data it collects, and translating those patterns into predictive analyses. It can relate these patterns both to the patterns of other users and to the current activities of a user. It can make sense of patterns and then provide *personalized* messages and recommendations throughout the day, based on patterns of previous activity and current activity. In sum, the app functions as a tracking tool that is fed with data of everything you do, all the time, and uses cloud-based data analytics to personalize the recommendations it gives you.

To illustrate this, let's take a closer look at some of the 'Insights' I have been sent by Garmin Connect.

*A step in the right direction*⁵⁷

Based on your activity pattern thus far, we predict that you will only take 5600 steps today. Go for a walk in the coming hours to make up your lack of steps.

*You are on the right track for meeting your goal*⁵⁸

If you follow your regular activity pattern today, we predict you will meet today's step goal. Well done!

*Do you sleep enough?*⁵⁹

It looks like you get less hours of sleep than might be good for you. Last month, 67% of the persons of your age and sex have slept more than you have. Sleep is very important for the health of your heart and helps your brain process information, so make sure you get enough hours of sleep.


The messages quoted directly above contained just a health or lifestyle tip. The app also sends personalized messages that contain native advertising.⁶⁰ To show how the native advertising is integrated into the health advice, I insert a screenshot from such a personalized message.

57 A screenshot of the full message can be found here: <https://imgur.com/a/LmmkesB>.

58 A screenshot of the full message can be found here: <https://imgur.com/a/g9I8cmZ>.

59 A screenshot of the full message can be found here: <https://imgur.com/a/53V3gwf>.

60 Both of the short 'Read More' messages have exactly the same style and feel as the advice that could be clicked on, but lead to different webpages with a "daily wellness tip" from the Cleveland Clinic Wellness Institute. At the very top of the webpages one can see a bright orange banner reading "Holiday Gift Guide: Fun and Healthy Gifts They'll Love!" that directs visitors to the web shop (where, for instance, a "16oz Stainless Steel Colored Mug" reading "Cleveland Clinic" can be bought, reduced from \$16.99 to \$13.59). Although it is not indicated anywhere, I think it is fair to presume that these Cleveland Clinic messages are in fact sponsored content, i.e., native advertisement. The Cleveland Clinic Wellness Institute is mentioned as a partner on the Garmin website, on a page called "Partner with us". The page shows a header that reads "Garmin Wellness API" and continues, under the header, to read "When you partner up with Garmin, we'll provide you with the programs and tools to make integration quick and easy. Using our flexible Wellness API (Application Programming Interface), and family of our products, you can validate a wide variety of activities" (<http://sites.garmin.com/en-US/wellness/partners/>, the webpage has since been taken down, screenshot available here: <https://imgur.com/a/F2f85HK>). The exact nature of the relationship between Garmin and the Cleveland Clinic Wellness Institute is not clearly stated, but since Garmin is a commercial company that most likely will not give others access to its API for free, some kind of commercial transaction between Garmin and the Cleveland Clinic Wellness Institute is probably involved.



All these Insights messages were accompanied with a graph to visualize the messages. As these messages clearly state, the Garmin Connect app has learned what my 'normal' activity patterns look like. It can use this knowledge to send me highly personalized messages during the day, by relating *my* patterns to *my* current levels of activity. It should also be noted that the goals that the app shows are personalized and constantly adjusted as well. The app does not disclose how it adjusts daily step goals, but it does mention that it has something to do with my past movement patterns and predicted future movement patterns. Through dynamic goal adjustment and Insights notifications, the app tries to influence my behavior by providing me with my own personal and unique nudges. Different people would receive different nudges at different moments, due to the simple fact that their activity patterns differ (slightly) from mine.

The Insights function of the Garmin Connect app is pretty sophisticated. It is not hard to imagine, however, how it could evolve into an even more persuasive experience with a few developments that can take place in the *very* near future. I am not trying to suggest that I personally think that such developments are a good idea. My claim is simply that given the current state of technological development, these would be natural next steps in the ‘optimization’ of the app.

First of all, the *content* of the nudges is currently partly personalized – different users will receive different recommendations – but the look, feel, and tone of the recommendations appears to be static. By making use of persuasion profiling in combination with tracking of the conversion (e.g., time spent in app, click through rates, engagement with advertising or social features or shop), Garmin Connect could experiment with differences in layout, tone of voice, and so on, for every user. In a similar vein, the native advertising shown directly beneath the recommendations seems to be adapted to the rather general content of the recommendation: so a recommendation on the amount of steps is accompanied with a native ad on shoes and walking. This could be improved (from the perspective of Garmin) by engaging in much more elaborate profiling and native ad matching to those profiles. Because Garmin Connect collects many different types of data continuously, it can also try to leverage those data to find out *when* users will be most susceptible to *what* types of native advertising packaged as health nudges.

Time will tell whether Garmin will end up implementing such features in the near future. Based on what is technologically possible and the desire to monetize the relationship with their users, I predict that the current proto-hypernudges will become (even) better able to persuade users to display behaviors that are profitable to the service provider.

5. Conclusion

I started this chapter with the very simple observation that popular health apps are commercial services that are run as companies and engage in a range of commercial practices to monetize their userbase. Many easily accessible technological tools exist nowadays to run health apps as successful companies. Health apps can now conduct constant, large-scale optimization experiments – where ‘optimization’ should be read as ‘optimal conversion’, and optimal conversion, in turn, should be read as ‘making sure the relationship between app and user is structured and developed in a way that is as profitable as possible.’ These experiments are further informed by the popularity of behavioral economics in Silicon Valley. Many of

the suggestions made by authors like Thaler and Sunstein (2008), and Eyal (2014) can now be implemented differently for different users, tested in real time, and dynamically adjusted (in a personalized manner) in real time. The commercial relationship between health app and user is thus one where technologically savvy analysts now have the tools to leverage all available data to shape the relationship in such a way that the user is steered towards the most profitable behavior.

Based solely on the technological developments described in this chapter, it is already possible to see the potential for unethical attempts to influence the users of health apps. For example, the ability to find out about insecurities, biases and other types of exploitable characteristics – whether they are generally present or only immediately felt at very particular moments that could be identified as well (see aforementioned Facebook study) – offers great potential to target people in ways that may manipulate them into displaying behavior they do not really endorse, especially if one understands such targeting practices in the context of peoples' general need and desire for health.

Before I provide a more elaborate ethical analysis of commercial health app practices, I will first turn in the next chapter to the current culture(s) of health and the strategically crafted health discourses one can find in health apps. It is only when both the technological and the socio-cultural aspects of the health app phenomenon are on the table, that I can provide a convincing ethical analysis.

Chapter 2

‘Apps That Will Make You Happier and Healthier:’ The Strategic Framing of Health in Health Apps

1. Introduction

In the previous chapter, I argued that health apps are – and should be analyzed as – commercial services with corresponding commercial interests and strategies. In order to get a critical analysis of health apps off the ground, I propose to first look into the technology and business models behind health apps. Understanding the business model behind an app provides important insights into the design choices that have been made in designing and (continuously) developing the health app in question.

However, to understand how digital (health) environments can end up shaping or steering the behavior of their users, we need to broaden the scope of the enquiry beyond the technical specifications of health apps and their underlying business models. Health app providers also try to influence how users understand the experience offered by an app, as well as their own relation to their health, through the propagation of a particular health discourse. So we also need to analyze the health discourses (or: rhetoric) one encounters in popular for-profit health apps to understand and evaluate how the behavior of the user can be steered. Qua commercial services, health apps have an interest in strategically presenting health in a particular way. How health is framed by the propagation of a particular discourse (partly) decides what type of users are attracted to the app, how they (will start to) understand their own health, what aspects of health are made salient to users, what type of health suggestions they are responsive to, and so on. The discourse employed in health apps matters, because it suggests to users what

kind of health norms should be observed. It gives shape to the semantic ‘field of understanding’ within which the users of health apps come to understand what health entails and how they should understand – and act upon – their own health.

The discourses on health in health apps should be understood in conjunction with the business models and technology behind apps. Especially with the popular for-profit health apps, there is a logic behind the specific ways in which business model, technology, and framing of health come together.

It is important to emphasize that different discourses on health can successfully frame health in different ways precisely because ‘health’ is itself a fuzzy, fluid concept. At first sight, it might seem that there is a conceptual rigor to the term ‘health.’ In our daily speech, the usage of the term health does not very often lead to confusion. Most of us agree that exercising is healthy, and eating excessive amounts of junk food is unhealthy. Most of us intuitively grasp that health has ‘something’ to do with the body, how well it functions, and how it is or can be affected by diseases and injuries. Upon closer inspection, however, it turns out to be very hard to pin down what exactly ‘health’ is supposed to mean. This fuzziness aides the strategic usage of the term ‘health’ for commercial purposes.

In what follows, I will first briefly explain how the concept of health is discussed in the philosophy of health. This section aims to show that despite a long tradition of conceptual analysis, the question of what ‘health’ means remains a fundamentally open-ended one. There is not one objective definition of health that is the correct one as a matter of fact. One’s understanding of health depends on the purpose one wants a concept of health to serve. This, in turn, explains why health is indeed a fuzzy concept and why health discourse can be subject to strategic commercial choices.

I then turn to our contemporary culture of health, which is the socio-cultural context within which health apps operate. I discuss the contemporary logic of health – the idea that one has to constantly prove that one is (still) healthy – and the fact that health is turning into the much broader and vaguer notion of wellness. This particular development of contemporary health culture is unsurprising, since the particular notion of health it propagates is one which is most commercially attractive: everyone *needs* and *should want* health and wellness products or services, and *everything* can be considered a health and wellness product or service.

In a next step, I show how these general features of contemporary health culture are used, reproduced, and exploited by health apps. My argument here is that the adopted health discourse is an important part of the business model of health apps. All of the important monetization strategies available to health apps incentivize the propagation of a health discourse which emphasizes how nearly everything is about one's health and how everyone can and should be preoccupied with their own health.

2. Health as a Contested Concept

In this section, I provide a brief overview of the most important philosophical discussions on the concept of health. In fact, significant and persistent disagreements regarding health's nature, meaning, and definition exist in the philosophy of health: health is a contested concept. I discuss the theories by Christopher Boorse and Lennart Nordenfelt, which represent the two main opposing views of what health consists of, to briefly show why health is a contested concept. In the next sections, it will become clear that precisely *because* health is contested concept, actors can choose to present health in a particular manner for strategic reasons.

2.1. Health as Biological Normality

Boorse defends a theory of health that others have called “atomistic-biological” (Nordenfelt 1986: 281) and which he himself calls a “functional account” (Boorse 1975: 58, 64). The starting point of Boorse's theory is the observation that in a very basic sense “health is normality” (Boorse 1975: 50). But which norms are we to use in order to determine what is normal? Boorse proposes to look at the human as a biological entity consisting of “systems and subsystems” (Boorse 1975: 57) which are oriented towards certain goals. What these goals are depends on the biological design of the species we are discussing, since every species, *qua* species, is biologically designed to be able to do certain things. Humans, for instance, have legs that are shaped in such a way that they can achieve the goal of walking upright on two legs, while bees have wings that allow them to achieve the goal of hovering above flowers. This, in turn, means that all the parts of the human physique have a “natural function,” which should be understood as “a standard causal contribution to the goal actually pursued by the organism” (Boorse 1975: 57). To return to the example, the natural function of the legs is to allow humans to walk and run upright on two legs. If we look back at the idea of “health as normality,” we now see that normality means nothing else than a part behaving in accordance with its natural function. This is why Boorse says that “the root idea of this account is that normal is the natural” (Boorse 1975: 57).

What does this imply for Boorse's definition of health? Boorse defends a derivative definition of health, which defines health as the absence of disease. Notice that this means that for Boorse, disease – and not health – is the primary unit of analysis, since health is completely defined in terms of disease. So, what counts as a disease? A disease is nothing else than a (sub)part not working in accordance with its natural function. A person, then, is healthy to the extent that she does not suffer from disease, which is another way of saying that a person is healthy to the extent that all of her (sub)parts are functioning in accordance with their natural functions.

Boorse's theory of health has an important implication. In his functional account, health is an entirely descriptive concept. To make a health judgment is to *describe* to what extent (sub)parts are functioning normally, i.e. help to achieve goals that are in line with their species biological design. "If diseases are deviations from the species biological design, their recognition is a matter of natural science, not evaluative decision" (Boorse 1977: 543). Determining whether someone is healthy thus requires, in this view, scientific empirical knowledge of the biology and physiology of a species. Moreover, to judge that someone is healthy does not mean that it is *good* or *desirable* to be healthy, or *bad* or *undesirable* to have a disease. To return to the example again, to conclude that a broken leg is an unhealthy leg is merely to conclude that it cannot perform its natural function. Such a health judgment implies nothing about the desirability of the condition.⁶¹

2.2. Nordenfelt's Holistic-Humanistic Perspective

Nordenfelt defends a theory of health that he calls a "holistic-humanistic perspective" (Nordenfelt 1986: 282). Where Boorse primarily focuses on an individual's actual physiology which he then analyses in terms of objectively determinable natural functions and biological species design, Nordenfelt starts with a socially situated person that wants to be able to do things in order to live a fulfilling life. For Nordenfelt, the primary unit of analysis is health. The basic idea of Nordenfelt's conception is easy to understand: health is about a person's ability to do what she wants to do in order to be happy. To the extent that a person is able to do so, she is healthy.

61 To be sure, Boorse is not committed to saying that certain diseases (such as broken legs) cannot be very unpleasant, or are not, in fact, *not* desired by people. He is merely committed to saying that, in relation to disease, there is nothing inherently desirable or undesirable about the condition in question. To still be able to account for the (un)desirability of certain conditions, he proposes a distinction between 'disease' and 'illness.' The notion of illness "*does* have a negative evaluation built into it" (Boorse 1975: 61) and can be understood as applying to those diseases that are experienced as being undesirable.

Let me first provide a more formal definition, which can then be unpacked more precisely in order to appreciate the important implications of this definition. Nordenfelt defines health as follows: “A is in health (sic) if, and only if, A has the ability, given standard circumstances, to realize his vital goals, i.e. the set of goals which are necessary and together sufficient for his minimal happiness” (Nordenfelt 1987: 90).

The general structure of the theory is that health consists of a person having the *ability* to realize those *goals* that are necessary for their *minimal happiness*. Let us first look at ‘ability.’ The first thing to notice is that Nordenfelt does not provide a specification of this term. This is a deliberate choice because it allows Nordenfelt to move beyond Boorse’s narrow focus on human biology and physiology and, instead, also incorporate the importance of social conditions for health. *Anything* that supports or enhances a person’s ability to realize her vital goals is conducive to her health. For example, if a person with a broken leg wants to visit a museum (because it contributes to realizing her vital goals), then a wheelchair being pushed by a friend can help her achieve this goal. The social support of a friend thus helps the person with a broken leg to live a healthy life, by enabling the realization of vital goals. Notice, moreover, that Nordenfelt speaks of the ‘ability to’ and not ‘achievement of’. The theory does not claim that a person is only healthy when she actually succeeds in being (minimally) happy.

Next, notice that the focus on the realization of “*his* vital goals” introduces inherent subjectivity in this definition of health, since different people will have different sets of goals that are important for their happiness. To see the importance of this, consider two young adults who are identical from a physiological perspective, except for one condition. One of them has a very particular knee injury, which only acts up while running. For Boorse, this is all we need to know in order to judge that the person with the knee problem is less healthy. After all, the natural function of the knee is – among other things – to allow the owner of the knee to run. Under Nordenfelt’s definition, however, both persons could be considered equally healthy, depending on their vital goals. It might be the case that the person with the knee problem has always hated running anyway and therefore does not need the ability to run in order to realize her vital goals. If this is indeed the case, then the physiological difference between both persons does not lead to different evaluations of their health status. Health is thus always relative to a particular person’s subjective vital goals. (If, on the other hand, the person with the knee injury is also the person who is really fond of running, then the knee injury might indeed be construed as making the owner of the knee less healthy).

The subjectivity of a person’s vital goals brings us to the last important element of Nordenfelt’s theory of health: happiness. In the end, Nordenfelt’s theory of health is about happiness, since the vital goals that we need the ability to achieve are defined

in terms of happiness. Here we encounter the root of the inherent subjectivity of the theory: what makes one happy differs from person to person. In effect, what conditions – whether they be social or physiological in nature – make a person healthy will differ accordingly, based on different subjective requirements for happiness. At the same time, the centrality of happiness also explains why health is an evaluative notion for Nordenfelt. If one accepts the proposition that happiness is desirable – and who wouldn't? – then it automatically follows that health is desirable, since health is a person's ability to achieve minimal happiness. In order to avoid confusion, it is important to emphasize that Nordenfelt speaks of *minimal* happiness. Health does not require perfect, heavenly, ecstatic happiness. All that is required for health is that people have the *ability* to be reasonably happy – happy *enough*.

So, what are we to make of these two contrasting theories of health? Clearly, people will differ as to which theory they find most plausible. My aim, however, is not to argue that one theory is more plausible than the other. Neither of the theories is *necessarily* correct – the correctness of a particular understanding of health cannot be evaluated on the basis of its conformity to empirical reality, because empirical reality does not provide us with facts that make any definition of health necessarily true.

What we can do is observe that although no understanding of health is necessarily true, there are plausible pragmatic reasons for adopting a particular understanding of health in a particular context. For example, an orthopedic surgeon inspecting an injured knee will probably adopt a largely Boorsean perspective on health. She will inspect the knee by comparing its current injured state to an ideal 'natural' state and, next, determine what steps can be taken to alleviate the current condition by restoring the 'natural' state of the knee. In this example of the orthopedic surgeon we can clearly see that a functionalist Boorsean understanding of health helps the surgeon to perform the task she is expected to do.

For the health app context, things are not quite so clear. There is not one conception of health which is obviously the 'correct' one – whatever 'correct' may mean here – for the health app context. Above all, it makes sense to emphasize the fact that precisely because it is fundamentally unclear what health means that its meaning is up for grabs in contexts that are not as formalized and institutionalized as the context of the hospital. Given the fact that the health app context is a thoroughly *commercial* context, it should not surprise us that very pragmatic and strategic reasoning informs the conceptions of health that are supported and propagated. If something can be successfully sold as a product that is relevant to one's health, then it will be framed as a health product aimed at improving one's health.

In what follows, I focus on the way in which the strategic employment of a particular health discourse can be seen as an important part of the business models and monetization strategies of health apps. First, I discuss the general culture of health within which health apps operate. Second, I discuss how health apps specifically fit within this general culture of health, gratefully making use of, perpetuating, and exploiting this culture of health for commercial gain.

3. The Contemporary Culture of Health

Health has become a dominant aspect of our contemporary culture.⁶² Consider, for example, the Dutch festival *Healthy Fest* – “a break for your mind and a boost for your body”.⁶³ Just like music festivals celebrate people’s love for music, this festival celebrates people’s love for a healthy lifestyle. In the Q&A section of the website, the organizers provide an answer to the question “What is Healthy Fest?”

Healthy Fest is an inspiring and challenging festival weekend where you can concern yourself with sport, mind, food, and friends in a fun and responsible manner. A three day line-up consisting of: professional trainers, serious workouts, energizing relaxation, and, of course, healthy food, all while having plenty of time to chat with each other in a relaxing way. You will return home with more energy than when you left. In sum: a real break for your mind and a boost for your body.⁶⁴

62 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 my argument certainly has a Foucauldian flavor, I do not incorporate this literature into my arguments in a systematic manner because this literature comes with its own particular vocabulary and methodologies. Within the overall structure of this dissertation, the introduction of the Foucauldian perspective of biopower, as well as its relation to digital health technologies, would distract too much from my own line of argument which already builds on various other theories, concepts, and disciplines.

63 <https://legacy.healthyfest.nl/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/0o0OQBI>.

64 <https://legacy.healthyfest.nl/veelgestelde-vragen/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/1WVm0jm>. Original quote is in Dutch, translation by me. “Healthy Fest is een inspirerend en uitdagend festivalweekend waarin je op een leuke en verantwoorde manier bezig bent met sport, mind, food en je friends. Een line-up van drie dagen lang: professionele trainers, serious workouts, energizing ontspanning én natuurlijk gezond eten waarbij er genoeg tijd is om te ontspannen en bij te praten met elkaar! Je komt met meer energie thuis dan dat je bent gekomen. Kortom; een echte break for your mind and boost for your body.”

Healthy Fest combines the idea that health is hard work with the idea that health should be actively pursued, performed for the world to see, and celebrated. In the description of the festival, the organizers are quite frank: besides inspiring, the festival is also going to be *challenging*, both physically and mentally. Still, people are more than willing to pay around €250 for 3 days of healthy lifestyle tips, together with like-minded people. *Healthy Fest* is a perfect metaphor for our time: people pay good money to be told by the commercial sponsors of a festival what the new health trends are and thus what a healthy lifestyle should look like; to work tirelessly to achieve such a perfect healthy lifestyle; and to broadcast the pursuit of health to the world by showing everyone (preferably via social media) that they are the type of person that goes to *Healthy Fest* for #healthinspiration.

The current culture of health is a rather complex phenomenon to dissect for it is made up of many different, interlocking elements. To start unpacking the different elements of contemporary health culture, let us start with its underlying logic. Devisch speaks of a “better than well” culture (2013: 21) based on an “infinite logic” of health (2013: 22).

Health is a task that we never know if we have fulfilled in a satisfactory manner. When are we healthy? The most logical answer would be: *if* we are not ill. But when are we healthy or fit enough in order not to be considered ill? Because we are emphatically asked to pursue our own health, we are no longer in a position to say that we are healthy *as long as* we are not ill. Since health is a norm that we are supposed to adhere to, the reverse is closer to the truth nowadays: we are *ill as long as* we are not fully healthy (Devisch 2013: 21).⁶⁵

According to Devisch, it used to be the case that someone was declared healthy in the absence of (obvious) disease or illness. Nowadays, health is something which must be actively achieved and we must constantly *prove* that we are (still) healthy. Devisch speaks of an “infinite logic” because the process of proving one’s health is a never-ending process with no logical conclusion to it. Even if one is in supreme health, one has to keep proving in the (near) future that one is *still* in

65 Translated from Dutch by me. The original quote: “Gezondheid is een opgave waarvan we nooit weten of we de taak naar behoren hebben volbracht. Wanneer zijn we gezond? Het meest logische antwoord luidt natuurlijk: *als* we niet ziek zijn. Maar wanneer zijn we gezond of fit genoeg om niet langer ziek te zijn? Omdat we nadrukkelijk onze gezondheid (moeten) najagen, gaat het niet langer op te zeggen dat we gezond zijn *tot zolang* we niet ziek zijn. Doordat gezondheid een norm is die we horen na te streven, sluit het omgekeerde dicht bij de dagelijkse realiteit aan: we zijn *zolang* ziek *tot* we gezond zijn.”

supreme health. We encounter a similar argument in Mau (2019: 68) who argues that quantified health scores in “personal health monitoring” tools “motivate us not ‘merely’ to be healthy, but to keep striving for improvement. Now part of a chain of association with activity and fitness, health is becoming a commodity we can never get enough of.” (As we will see below, health’s contemporary logic is also infinite because the very norm of what counts as healthy keeps shifting and becoming more demanding. Increasingly, many activities are being labeled as ‘relevant to one’s health,’ meaning one has to pursue and optimize an ever increasing number of aspects of one’s health).

Health apps and related products such as activity trackers find a natural place within this new culture. These technologies give individual users the tools to actively engage with their health and prove their health. Lupton (2013: 397) observes that advocates of self-tracking technologies such as health apps and activity trackers “tend to place emphasis on the potential for the ‘empowerment’ of lay people by these technologies and the importance of ‘taking responsibility’ with one’s health”. Since “the common assumption is that health must be achieved,” (Crawford 2006: 402) and since we now actually ‘have the tools’ to achieve and prove our own health, a process of *responsibilization* emerges (Cederström & Spicer 2015: 4-8; Lupton 2013: 397-398). We are being made responsible for observing and managing our own health.

With the responsibilization of health also comes what could be called the *moralization* of health. Now that we are responsible for our health, and now that health must constantly be achieved and maintained, we are *expected* to take care of our health and to adopt a *healthy lifestyle*. The moral and social demand to be healthy also perpetuates a culture of ‘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” (Lupton 2013: 397). The phenomenon of healthism also explains why the phrase “a healthy lifestyle” is becoming commonplace. If we are expected to live healthy, then it should come as no surprise that health is advocated as a norm according to which people should evaluate and organize their lives. A ‘lifestyle’ is quite literally the style of one’s life, and thus refers to one’s values and reasons on the basis of which one gives shape to one’s life. In a culture of healthism, health is one of the most important values providing one with reasons for shaping one’s life in a particular way.

Thus far, I have discussed the *logic* of the contemporary culture of health, which revolves around the need and the demand to constantly prove one's health and organize one's life(style) in a 'healthy' way. However, understanding the logic of health does not yet explain what health *itself* means – or, put differently, what *counts* as health and a healthy lifestyle. This is where it becomes difficult to come up with a clear, coherent discussion of health culture. Why? Because so many activities, products, services, habits, practices, and so on are being drawn into the health sphere. Nearly everything is health-related today. To still be able to understand what health itself means nowadays, I suggest we follow Cederström and Spicer (2015), who argue that the notion of health has morphed into a broader, vaguer notion of *wellness*.

Wellness refers to the general idea of 'doing well' and 'feeling good,' which, in turn, can refer to anything and everything in a person's life. There is of course not one single, objective conception of wellness that everyone in the commercial health context agrees upon. What is clear, however, is that nearly everything can be – and is – framed as a matter of wellness (which is, coincidentally, precisely why it is hard to formulate one clear conception of wellness as it is found in contemporary health culture).

The incredible broadness and vagueness of the term 'wellness' is unsurprising. Commercial actors in the health and wellness business are amongst the strongest advocates of both the importance of wellness and an incredibly broad understanding of wellness: the more expansive the notion of wellness, the more can be advertised as being – in some (farfetched) sense – relevant to wellness. As a result, there is a very wide range of products, services, and activities in contemporary health culture aimed at wellness and that are associated with a healthy lifestyle. One of the most dramatic examples I encountered is this interview⁶⁶ about news apps framing themselves as wellness apps that are part of a healthy lifestyle. Brodesser-Akner (2018) accurately captures the wide range of things that are supposedly about wellness:

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 hands and bow as a way of saying thank

66 "Wellness apps, but for news: Can Neva Labs build a news reading experience that feels healthy?". Last accessed 22 September 2020 at <http://www.niemanlab.org/2018/04/wellness-apps-but-for-news-can-neva-labs-build-a-news-reading-experience-that-feels-healthy/>, screenshot available here: <https://imgur.com/a/GYcvVEc>.

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.

Because wellness can potentially mean anything, there is a constant flux of new healthy wellness trends that are extensively discussed on popular and trendy lifestyle platforms such as *Mind Body Green*⁶⁷, *Well+Good*⁶⁸, *goop*⁶⁹, and *Elle*⁷⁰, but also *The Guardian*.⁷¹ (Interestingly, discussions of these wellness trends are often coupled with commercial suggestions by ‘partners’ or native advertising – more on this below).

With the meaning of health being morphed into the much broader notion of wellness, it also becomes much easier to incorporate a ‘happiness’ discourse into the health and wellness content. Many health apps effortlessly link the terms ‘health’ and ‘happiness,’ promising that their apps will make you *happier and healthier*. For example, Headspace “has one mission: to improve the health and happiness of the world.”⁷² Lifesum has the “vision to make it simple for people everywhere to form habits that build towards healthier, happier lives.”⁷³ MyFitnessPal write that they are “helping millions of people across the world to live healthier and happier.”⁷⁴ The popular Dutch health and wellness website The Green Happiness advertises its subscription for exclusive content with the slogan “A happy & healthy you.”⁷⁵

67 ‘11 Wellness Trends to Watch’ (<https://www.mindbodygreen.com/articles/wellness-trends-2018>, last accessed 22 September 2020, screenshot available here: <https://imgur.com/a/RORG1oo>).

68 ‘The Top 18 Trends of 2018 Are Here’ (<https://www.wellandgood.com/fitness-wellness-trends/>, last accessed 22 September 2020, screenshot available here: <https://imgur.com/a/4E52oMz>).

69 <https://goop.com/wellness/>, (last accessed 22 September 2020) screenshot available here: <https://imgur.com/a/BPnWbxI>.

70 ‘30 of the Most Popular Wellness Trends of 2017’ (<https://www.elle.com/culture/g14444727/health-wellness-trends-2017/>, last accessed 22 September 2020, screenshot available here: <https://imgur.com/a/Xwhgah3>).

71 ‘From Psychoactive drugs to Silent Spas: 2018’s Weirdest Wellness Trends’ (<https://www.theguardian.com/lifeandstyle/2018/jan/05/2018-wellness-trends-silent-spas>, last accessed 22 September 2020).

72 <https://www.headspace.com/about-us> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/7rjYNLV>.

73 <http://jobs.lifesum.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/d8vt2wM>.

74 <https://www.myfitnesspal.com/jobs> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/KVo0EWc>.

75 <https://mygreenhappiness.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/cGpzdUO>.

On top of making you *happier*, some health apps also say they will simply make you *better* – which I interpret as meaning a better human, or someone living a better life. Under Armour Record writes “Improve your health and fitness with our family of apps. From sleep, fitness, step activity and nutrition, these apps share one mission: to make you better.”⁷⁶ Fitbit invites visitors of its website to “See how Fitbit can help you exercise, eat, sleep & live better.”⁷⁷ So besides the fact that health has come to mean wellness in contemporary health culture, there is also a constant insistence that being preoccupied with one’s health in the manner suggested by health websites, blogs, apps and so on makes one’s life happier and better. Health is thus framed as an essential component of a successful life.

In many of these wellness promotions, advertisements, and tips for a healthy lifestyle, the language and meaning of wellness, health, and happiness are merged – just like Cederström and Spicer (2015) suggest. Take, for example, the popular website *goop*, which offers its own line of *goop wellness* products and which are promoted through many native advertising articles on their websites. In the promotion of these wellness products *goop* constantly emphasizes that their products will make you feel and look beautiful and good, and *thus* are also healthy. Take, for instance, their *goopglow Morning Skin Superpowder* (promoted via dedicated native advertising articles dressed up as legitimate health/wellness tips⁷⁸), which the company claims you should drink “as part of your healthy morning ritual” to “drink your way to a glowing skin.”⁷⁹ The message here is that health is about having a glowing skin. *Goopglow* is just one example of the way in which nearly everything is now being framed as a *healthy* wellness product or service. In the next section, I will discuss multiple examples of health apps eagerly working with the ever broader, vaguer, all-encompassing notions of health as wellness.

At this point, it is important to recall that the *logic* of health demands that health is constantly pursued and proven to oneself and the world. This logic, coupled with the fact that health becomes wellness, offers the perfect conditions for the commercial exploitation of health. Because of wellness’ expansiveness as a term, all kinds of trendy, hip, beautiful lifestyle products can be (re)branded as healthy lifestyle products. And of course, because they are now considered ‘healthy,’ these

76 <https://record.underarmour.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/YLOO77j>.

77 <https://www.fitbit.com/home> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/JZ1lpnd>.

78 ‘Drink Up and Glow Up’. Available at <https://goop.com/beauty/skin/drink-up-and-glow/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/KFnmtOk>.

79 <https://shop.goop.com/shop/products/goopglow> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/A8f50RL>.

products can be sold at what could be called a ‘health premium.’ The *goopglow Morning Skin Superpowder* costs \$60 for a pack of 30, meaning daily use of this supplement which is claimed to be so good for your health will cost over \$720 a year (in 2019). If yoga gives you a better body-mind balance which helps you deal with your work stress, then surely that one pair of \$100 yoga pants are a health product, right? You cannot put a price on health, so it must be worth it.

As a result, health magazines, blogs, platforms, social media accounts, and apps are eagerly promoting and making use of the contemporary logic of health, because it affords them with the opportunity to tell their (prospective) audiences what a curated healthy lifestyle should look like and which products and services fit such a lifestyle. “As consumers, we are required to curate a lifestyle aimed at maximizing our wellbeing [...] wellness has wormed itself into every aspect of our lives” (Cederström & Spicer 2015: 3).

Thus understood, it is entirely unsurprising that we experience such an explosion of healthy lifestyle and wellness content. As the online advertising company *NativeAdBuzz* explains in a New Year’s resolutions post:

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?⁸⁰

These remarks from *NativeAdBuzz* serve as a good reminder that much of the health content on websites, blogs, social media, and in health apps is not in essence about health. Above all, health is an incredibly convenient and useful vehicle for commercial practices, since everyone agrees that health is important and valuable: who can possibly be *against* health? Framing one’s commercial practices in terms of health provides these companies with a sense of innocence. This is why the contemporary culture of health – and the types of health discourse it allows for – is so central to properly understanding the commercial practices of health apps.

80 <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/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/aUL9qdJ>.

Health apps ride the waves – which they themselves partly co-create and reinforce – of a culture where everything matters to a healthy lifestyle. The business models of health apps thus also depend on what ‘counts’ as health and what can be sold as ‘relevant to health.’

In the following section I will discuss two concrete examples from the health app context to show how the contemporary general culture of health is reflected in the health discourse one can encounter in health apps. By focusing on concrete examples, it should also become clear how the use of a particular health discourse can structure the user experience of an app. The ways in which users interact with an app is not just determined by its technical specifications, but also by the presentation and framing of the content. Put shortly, my argument is that it is a deliberate business strategy (1) to use expansive ‘health as wellness’ notions of health, and (2) to use discourses on health that communicate the need to constantly pursue one’s health, since this offers the highest potential for monetization.

4. Health, Wellness, and Healthy Lifestyles According to Popular Health Apps

4.1. Fitbit: Every Moment Matters and Every Bit Makes a Big Impact

When you go to the [fitbit.com](https://www.fitbit.com) homepage, you are greeted by the aforementioned message “See how Fitbit can help you exercise, eat, sleep and live better.”⁸¹ When clicking on this message (which is accompanied by a large image of very fit and able-bodied woman in full sportswear looking at her Fitbit device), you are directed to a new page that shows the following text, projected over a video playing in the background which shows people working out, dancing at a party, competing in a race, walking down the street in sportswear, and cooking with lots of vegetables:

“On the walk to work, at the weight room or in the last mile.

Somewhere between first tries and finish lines.

Pillow fights and pushing limits. That’s where you find fitness.

81 <https://www.fitbit.com/nl/home> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/JZ1lpnd>.

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.

How you spend your day determines when you reach your goals. And seeing your progress helps you see what's possible.

*Seek it, crave it, live it*⁸²

This statement could not communicate the message more clearly that, according to Fitbit, your fitness is about *everything* you do. “It’s all the time.” If you scroll down, you see the “Meet Fitbit” screen, which tells you that “Fitbit tracks every part of your day – including activity, exercise, food, weight and sleep.” Even further down the page you read that “Fitbit’s devices, dashboard and app work together to deliver a complete, connected and fun experience that’s 24/7.”

Since Fitbit manufactures wearables that should be used in conjunction with their health app (simply called Fitbit as well), it makes sense for the company to propagate a discourse that makes everything about health. Their wearables are designed (and advertised) to be worn all the time. The term wearable is, in a sense, too generic, for it refers to any type of device that can be worn by persons. More specifically, Fitbit calls its wearables ‘activity trackers’ aimed at tracking all your activity.

The catch-all nature of Fitbit’s health discourse is also clearly illustrated by the “Real People, Real Stories” section on the homepage.⁸³ One story is about Sandile M. who became partly paralyzed after a motorcycle accident, but “returned to the gym to regain confidence in his body, relying on his Fitbit Charge 2 as a coach and companion to take on the challenge.”⁸⁴ The story of Hayley S. (“Can Positive Thinking Lead to Weight Loss? It Did for This Woman”⁸⁵) is, in a way, much less spectacular. To the outside world, she appeared to be someone who had it all. “But on the inside I was very unhappy. I wasn’t eating well or exercising, and I didn’t like the reflection I saw in the mirror.” The text on the Fitbit website continues:

82 <https://www.fitbit.com/nl/whyfitbit> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/qw8nb5u>.

83 <https://stories.fitbit.com/nl/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/N7ql8oB>.

84 <https://stories.fitbit.com/nl/featured/sandile/> (last accessed 22 September 200), screenshot available here: <https://imgur.com/a/GqNPtRE>.

85 <https://stories.fitbit.com/nl/more-stories/hayley/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/e3RR7ee>.

“Still, it was hard for Haley [sic] to commit to improving her health. “I had the worst mindset,” she says. “I would think, ‘Well, most of America is obese, so what’s the difference?’” Long story short: Hayley started using a Fitbit Flex to track her activity and that not only helped her to set (and achieve!) goals, but also, along the way, helped her gain a more positive mindset. Yet another ‘Real Story’ introduces professional musician Kiyoshi M.⁸⁶ who overcame his stage fright by using a Fitbit to calm down through guided breathing.

I want to make three observations here. First, notice how *different* these stories are. One is about someone who became disabled, one is about a person that is struggling with issues many people face and the power of positive thinking, and one is about someone with a relatively rare profession that comes with very specific challenges (i.e., stage fright before a concert). What this wide range of stories illustrates, and is *meant* to illustrate, is that the Fitbit is for *everyone*. Whatever your life looks like, and whatever your goals or ambitions are, Fitbit can help you live healthier and happier. Second, when focusing on the specific activities and goals Fitbit says it helps support and improve by tracking with a Fitbit activity tracker, it turns out that *everything* is worth tracking. The stories are about sleep, working out, relaxing, feeling positive, weight loss, breathing, and so on. In all those cases, we are made to believe that Fitbit is helpful in achieving improvement. Third, Fitbit makes use of the term health in a very broad, loose sense, largely in line with the ‘health as wellness’ discourse I discussed earlier in this chapter.

Another telling message on the Fitbit website is that “[e]veryone’s approach to fitness is different. One-size-fits-all doesn’t always fit you. That’s why we created a family of products that work seamlessly with each other, your budget, your favorite apps and your goals.”⁸⁷ Precisely because health and fitness can mean everything, it can also mean different things to different people. As a result, it of course makes sense to offer a wide range of different products (ranging from very simple bare bone trackers to high performance sports watches, and from smart scales to a variety of fashionable wristbands for one’s tracker) that can be used in many different contexts and during many different kinds of activities. Fitbit thus explicitly recognizes that health can mean everything to everyone – and it gladly provides products to support people in whatever health endeavors that may want to undertake.

86 <https://stories.fitbit.com/nl/featured/kiyoshi/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/SZQMKJA>.

87 <https://www.fitbit.com/nl/whyfitbit> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/VQ7IY0A>.

These stories, together with all of Fitbit's other (marketing) texts, propagate a very particular understanding of health which mirrors and reproduces more general principles of the contemporary culture of health. Health is presented as a *lifestyle* that concerns not only the body, but also the mind. Everything you do contributes – either directly or indirectly – to your health and fitness (or: wellness). The particular language Fitbit uses illustrates this point; many terms are used almost interchangeably in a rather nonchalant fashion: fitness, happiness, well-being, weight, confidence, workout, running, mindfulness, stress, sleep, relaxing, ambition, activity, nutrition, vegetables, eating, gym, pushing limits, motivation, steps, goals, food, community, freedom, training, evolving, journey, maximization, body fat, cardio, breathing sessions, BMI, lean mass, fun, smile, empowerment, experiences, moving, heart rate. All of these terms and many more are used in describing why Fitbit matters for you and your health. In the Fitbit universe, health thus is about your general wellness, which allows them to design commercial practices which be presented as relevant to everything you do.

4.2. MyFitnessPal: How Even a ‘Simple’ Calorie Counter Is Secretly about Everything You Do

MyFitnessPal is another example of a health app that propagates a health discourse aiming to make nearly everything you (can) do about health (or, again: wellness). MyFitnessPal has over 150 million users.⁸⁸ Before installing the app, the app store description suggests it is a simple calorie counter app: “Lose weight with MyFitnessPal, the fastest and easiest to use calorie counter.” After installing the app, it immediately becomes clear that the app is much more than a calorie counter. The home screen shows, all the way at the top, two lines worth of calorie counting information. Directly below this, the majority of the home screen (around 90%) is filled with MyFitnessPal blog posts you can scroll through infinitely (because they keep on loading automatically once you reach the bottom of the screen). These blog posts feature stories on health in the broad sense: the blog clearly relies on the

⁸⁸ It is difficult to find out exactly how many users health apps have. Most of the popular apps do not release precise numbers and Google's Play Store and Apple's App Store only show rough estimates. For example, Google's Play Store reports “50,000,000+ downloads” for MyFitnessPal, which could mean 50,000,001 users, but just as well 76,519,567 users. Another difficulty is the fact that different language versions of app stores show different app descriptions to users. For example, the Danish App Store description of MyFitnessPal reports that “over 100 million people use MyFitnessPal”; while the English language version of the same App Store does not provide this information in the description. An interview with a MyFitnessPal employee that can be found on LinkedIn reports at least 165 million users (<https://www.linkedin.com/pulse/interview-how-myfitnesspal-app-got-165-million-users-david-jones>, last accessed 22 September 2020, screenshot available here: <https://imgur.com/a/0bCwCsq>) but that number is, for reasons mentioned above, difficult to verify.

broad wellness frame. One can find inspirational blogs by runners, healthy recipes by dietitians, chefs, and popular bloggers, blogs about yoga and mindfulness, blogs about weight loss, blogs about sleeping habits, blogs about people's (positive) experiences with self-tracking, blogs about people going for strolls during their lunch breaks, and so on.

Most of these blogs are clearly written to serve as *suggestions* to users: why not try “[t]his Thai-inspired vegetarian pumpkin curry from Dietitian Debbie [hyperlink to her personal website] [...] [p]acked with antioxidants and 12 grams of fiber, it also contains 21 grams of protein, which is not easy to find in a vegan dish.”⁸⁹ Or how about a “5-Pose Yoga Fix: Yoga to Ease Cold Symptoms”⁹⁰ or a “5-Pose Yoga Fix: Energize Your Morning”⁹¹? Why not read “How Jorge Made Health and Fitness a Family Affair.”⁹²

If you spend some time using the MyFitnessPal app, it becomes clear how much more it is than just a calorie counter app. Calorie counting is only your very first step in the health world of MyFitnessPal. From the outset, a majority of your attention is being drawn to aspects of your health other than calorie counting. The very prominently placed blog posts mentioned above are only one example. In the main sidebar, there are over 15 tabs that lead the user to different parts of the app. There are, for instance, tabs for ‘Recipes, Meals & Food,’ for other ‘Apps & Devices,’ for ‘Steps,’ for ‘Community,’ for ‘Shop Fitness Gear,’ and for the aforementioned ‘Blog.’ There is, in other words, a wealth of health-related things to do in the app.

To better understand the significance of all these functionalities of the app, it is important to know that MyFitnessPal – originally founded in 2005 and later funded by venture capital firm Kleiner Perkins Caufield & Byers for \$18 million⁹³ – was acquired by sports apparel brand Under Armour in 2015. In an effort to become “the world’s

89 <http://blog.myfitnesspal.com/vegan-pumpkin-curry-recipe/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/AHZo38q>.

90 <http://blog.myfitnesspal.com/5-pose-yoga-fix-yoga-ease-cold-symptoms/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/xqXstqb>.

91 <http://blog.myfitnesspal.com/stretches-energize-morning-5-pose-yoga-fix/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/MC1JR9U>.

92 <https://blog.myfitnesspal.com/watch/what-it-took-for-this-obese-doctor-to-take-his-own-health-advice/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/CTe0p3h>.

93 <https://www.crunchbase.com/organization/myfitnesspal> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/niXPjXV> & <https://techcrunch.com/2013/08/11/myfitnesspal-funding/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/SpStMUv>.

first 24/7 connected fitness system,⁹⁴ Under Armour acquired multiple popular health apps, with MyFitnessPal as its most expensive purchase for a price of \$475 million.⁹⁵ With the purchase of Endomondo (a workout app) and the MapMy app family (consisting of various kinds of activity tracking apps),⁹⁶ the total investments made by Under Armour add up to at least \$710 million.⁹⁷

As a result, MyFitnessPal has become part of a larger effort by Under Armour to interconnect many different aspects of health into one “24/7 connected fitness system.”⁹⁸ From the perspective of Under Armour, users ideally start using Under Armour Record, a health app that functions as a dashboard that brings all the different acquired apps together within one app (advertised with again roughly the same text: “the world’s first 24/7 connected health and fitness system.”⁹⁹) However, even if a user only uses MyFitnessPal, without installing any of the other apps or Under Armour Record, the other apps still enter the user experience. For example, after I installed the MyFitnessPal app, the very first message on the home screen that greeted me was a video called ‘Motivation with Kym Nonstop.’ In the one-minute video, Kym greets the viewer (in full Under Armour sport apparel) by saying “Put on those shoes, get to the gym, walk out your door, take a walk, take a job. Do the first step right now. Literally as soon as you finish watching this, I need you to go.” The embedded video is accompanied by another message: “Get Connected. More than 30 apps and devices work with MyFitnessPal.” Directly below it there is a big blue button (blue is the main color of the MyFitnessPal layout) that says “See Apps & Devices”, and leads you to the same page as the sidebar tab ‘Apps & Devices.’ On this page the user sees suggested connected apps, with at the top of the list UA Record, MapMyRide, MapMyFitness, MapMyRun, UA Scale, UA Healthbox, and UA Band (which are all apps owned by Under Armour).

Without wanting to discuss all these different apps, I do want to emphasize that they affect the ‘health experience’ within the MyFitnessPal app. From a business perspective, it is understandable that Under Armour wants a maximum number of users engaging with as many Under Armour health apps as possible. Practically speaking, it means that

94 <https://blog.underarmour.com/>, screenshot available here: <https://imgur.com/a/EQnWsfC>

95 <https://www.wsj.com/articles/under-armour-to-acquire-myfitnesspal-for-475-million-1423086478> (last accessed 22 September 2020).

96 Namely MapMyRide, MapMyFitness, MapMyRun, and MapMyWalk.

97 <http://fortune.com/2015/04/21/under-armour-connected-fitness/> (last accessed 22 September 2020)

98 <https://record.underarmour.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/YLOO77j>.

99 <https://record.underarmour.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/YLOO77j>.

MyFitnessPal, as a ‘simple’ calorie counter, explicitly and persistently tries to nudge its users to interact with a wide range of apps that focus on altogether *different* aspects of the user’s health.

In terms of the health discourse found in the app, the propagation of all the different app functionalities in combination with the blog posts that cover a very wide range of stories, makes, like with Fitbit, for a very broad catch-all discourse on health as wellness. Due to its prominence in the app and the content it features, the MyFitnessPal blog is a suitable function of the app to examine for its health discourse. It should be emphasized that MyFitnessPal is constantly creating new content for the blog, resulting in multiple new stories per day which are presented to all users on the app’s homescreen. If someone uses the app for, say, a month, they will thus have been presented with over a hundred blog posts about the very wide range of activities and products that are framed and presented as health-related.

The willingness of the app to frame nearly everything as a ‘health thing’ can be clearly seen in the many recipes that are shared on the blog and which often contain native advertising. Take for instance a recipe for “Roasted Fall Vegetables Salad With Maple Orange Cinnamon Dressing”¹⁰⁰ which is provided by the food blog The Roasted Root.¹⁰¹ At the bottom of the recipe as it is displayed on the MyFitnessPal blog, there is a blue highlighted ‘about’ section describing The Roasted Root. “Julia Mueller writes the food blog The Roasted Root and is author of “Delicious Probiotic Drinks” and “Let Them Eat Kale!” [...] She enjoys developing recipes that are nutrient-dense.” Directly at the start of the actual recipe as it is displayed on the MyFitnessPal blog, there is a bold, bright, blue hyperlink to the same recipe on The Roasted Root food blog. Users are thus encouraged to read the recipe on The Roasted Root. In the recipe, Julia Mueller inserted an additional paragraph reading “Yesterday, I posted Chocolate Peanut Butter Chia Seed Smoothie along with a Health Warrior Chia Bar giveaway (which you should go enter...RIGHT NOW!). I have been adding the white chia seeds to just about everything under the sun, and it turns out they’re great on salads.”¹⁰² The references to both the posts come as bold purple hyperlinks, asking to be clicked. As it turns out, both posts are sponsored content, with the second hyperlink leading directly to an Amazon product page.

100 <http://blog.myfitnesspal.com/roasted-fall-vegetable-salad-with-maple-orange-cinnamon-dressing/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/gHHLN2Q>.

101 <https://www.theroastedroot.net/roasted-fall-vegetables-salad/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/sGtLRfy>.

102 <https://www.theroastedroot.net/roasted-fall-vegetables-salad/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/sGtLRfy>.

One could address (and object to) the sleek native advertisement strategies that MyFitnessPal is explicitly endorsing by associating itself with a food blog that is built around a native advertising business model. I have critiqued such practices elsewhere (Sax, Helberger & Bol 2018), and will also do so in later chapters. For now, I want to focus on what such a recipe means in terms of the health discourse that is perpetuated by the MyFitnessPal app. First of all, the recipe is posted on the MyFitnessPal blog and features directly on the homepage of the app. The recipe comes with a blue MyFitnessPal button ‘Log It,’ which allows the user to directly import the nutritional information to the calorie counting function. So even these clearly commercially inspired references to the natively advertised Health Warrior Chia Bar are presented and incorporated into the overall design and experience of the app as a legitimate health-related product. Health, apparently, is also about Health Warrior Chia Bars.

Besides recipes, there are many other things that are framed as relevant to users’ health and which are often accompanied with suggested products or services. A few examples are as follows:

- **Yoga** is frequently discussed in articles such as “Yoga for Sciatica Relief,”¹⁰³ “Yoga for Optimal Digestion,”¹⁰⁴ and “A Day in the Life of a Yoga Teacher.”¹⁰⁵ All these articles come with direct links to the Under Armour webshop to “gear up for your next yoga session.” The “A Day in the Life of a Yoga Teacher” article also features native advertising as the yoga teacher in question – Liz Arch – writes in the middle of the article that “Skincare comes next. I’m obsessed with a non-toxic, plant-based, biodegradable and compostable skincare line called Yuni [embedded hyperlink to product’s website]. I wash my face with Yuni Concentrated Cleansing Crème, then hydrate my skin with the line’s moisturizer.” Here we see the classic health-as-wellness discourse in action, which ticks many of the trendy wellness boxes by referring to the product as “non-toxic, plant-based, biodegradable, and compostable.” Moreover, we see activities like ‘skin care’ being dragged into the health domain and being merged with other vague health tips such as taking a “curcumin supplement to help promote overall health and reduce inflammation,” and “an omega-3

103 <https://blog.myfitnesspal.com/5-pose-yoga-fix-yoga-sciatica-relief/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/d4nqXcR>.

104 <https://blog.myfitnesspal.com/5-pose-yoga-fix-yoga-optimal-digestion/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/sNOJlKC>.

105 <https://blog.myfitnesspal.com/day-life-yoga-teacher/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/KijSKaz>.

supplement and a B-complex supplement for brain health and nervous system support,” or eating “vegan bibimbap [...] which is full of probiotics and helps promote a healthy gut.”

- **On-ear headphones** that “can actually withstand your workouts.”¹⁰⁶ Health and fitness writer Aleisha Fetters shares her headphone story with the MyFitnessPal blog, since wearing the right headphones during your workout is an important health issue, it claims. The article – which in its entirety is a piece of native advertising – discusses how good the “UA Sport Wireless Train on-ear Headphones Engineered by JBL” are.
- **Music playlists** that are “built to make you better”¹⁰⁷ and picked by “certified strength and conditioning specialist [...] [and] fitness expert” Anthony J. Yeung. With these playlists, Anthony explains, you can perform better: “If you’re doing fast-paced circuit, listen to fast-paced music – it can help you boost your performance.” This article, like the last one, is in fact a piece of native advertising and part of the same campaign for the Under Armour JBL headphones.
- **Family vacation destinations** are also framed as a health issue, with the article suggesting “some out-of-the-box options” for “those who like to mix it up.”¹⁰⁸ The article appears to be merely about ways in which one can plan a healthy, active family vacation. In reality, the article is native advertising for very specific destinations.
- **The alarm clock** is an important health topic because “When you wake up to a loud alarm clock, your body releases cortisol – a reaction to stress of being startled awake. Your heart beats faster and blood pressure rises, just like in a flight-or-fight response.”¹⁰⁹ The article start by listing some more (quasi-) academic studies on sleep disorders before introducing “certified sleep science

106 <https://blog.myfitnesspal.com/these-on-ear-headphones-can-actually-withstand-your-workouts/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/9wBLM2g>.

107 <https://blog.myfitnesspal.com/these-playlists-were-built-to-make-you-better/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/lbvc8sJ>.

108 <https://blog.myfitnesspal.com/best-u-s-outdoor-family-vacations/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jPUTUgW>.

109 <https://blog.myfitnesspal.com/why-and-how-you-should-nix-an-alarm-clock/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/1qy1faF>.

coach” Chris Brantner from SleepZoo,¹¹⁰ a mattress shop (that itself tries very hard to look like a sleep information website). Later in the article the reader is also referred to Tuck.com,¹¹¹ a sleep information and review website that also benefits financially¹¹² from users being directed there from the MyFitnessPal blog.

This list could be much, much longer, but these few examples hopefully show how diverse the activities, products, and services are that are framed as being relevant to one’s health in the MyFitnessPal app.

In this section we have seen how well health apps and their health discourse fit in – and eagerly make use of – the contemporary health culture. The next section will briefly discuss how exactly the used health discourse can serve the business models of health apps.

Understanding this connection between discourse and business models will help to critically analyze the various commercial practices of health apps.

5. Health Is Like Sex, It Sells

As I explained in chapter 1, health app business models can be based on any of the following monetization principles, or a combination thereof:

1. One-time fee, often to be paid before the user can install and use the app;
2. (Native) advertising, to be shown in one or more places in the app;
3. In-app purchases, either for products offered by the app developer itself, or in a store that also features (paid-for, promoted) products by others;
4. Subscriptions, where users pay a recurring monthly (or weekly, or daily) fee in order to gain and retain access to either the entire app, or to some additional content or features.

110 <https://www.sleepzoo.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/IGg5Pnf>.

111 <https://www.tuck.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/fQsl0aw>.

112 <https://www.tuck.com/disclosure/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/bhATUan>.

The health discourse discussed above serves each of these monetization principles. As I observed in chapter 1, all of these monetization principles flourish through the engagement of as many users as possible. When nearly everything in a person's life can be framed as relevant to that person's health, and if, moreover, the general culture of health makes the pursuit and (public) performance of proving one's health very important, then everyone can plausibly be addressed as a potential user of a health app. Moreover, the moralization of health, by which health (or: wellness) can become a (social) demand, is also helpful in getting users to engage with health apps and the commercial practices that come with health app usage. As a result of the prevalent health discourse, the potential for user engagement is much higher than if health apps more narrowly focused on health as, for instance, the absence of (serious) illness or disease. This line of reasoning also holds for each of the separate (but often combined) monetization principles.

1. When monetization occurs through one-time installation fees, the potential for monetization obviously increases when as many prospective users as possible see the health app in question as relevant to their health.
2. Especially native advertising works best when the advertised content blends together naturally with the health content. So if nearly everything one does is (supposedly) relevant to one's health, then the pool of possible buyers of native advertising space increases. On top of this, there will simply be more opportunities for native health-related advertising when everything is a health issue. The same holds for regular advertising: the more expansive the health discourse, the more products and services can be advertised as relevant to the user.
3. In-app purchases are served by a broad, vague health discourse in a similar fashion: the broader and vaguer the health discourse, the more in-app purchases can be presented to the user as relevant. For example, a workout app can offer a set of free workouts, but suggest that by paying an additional fee, the user will also gain access to guided yoga sessions to help achieve some wellness goal. The MyFitnessPal app also shows that a broader conception of health allows the app to suggest fitness gear (ranging from fitness apparel to headphones) from their own in-app shop to users.
4. If you want to sell subscriptions for either the usage of an app or for additional content and features, you need to make sure potential users see your product as useful. Emphasizing that working on your health is important and that

“every bit makes a big impact”¹¹³ obviously helps to sell subscriptions. Again, emphasizing that nearly everything people do is health-related broadens the potential audience of the app. The moralization of health and people’s internalized sense of responsibility for ‘taking control’ of their health also helps to make people consider paying for the use of an app. What, after all, is a €10 a month subscription if this helps you live happier, healthier, and better?

We thus see that health apps cleverly latch on to the general culture of health and incorporate (and reproduce) this culture into their own health discourse. The deliberate and careful propagation of this health discourse serves their business models especially well, since it helps to produce and sustain circumstances which provide many opportunities for monetization.

6. Conclusion

In this chapter, I have argued that health apps should be understood and analyzed as commercial services that are built on the basis of business models which, in turn, inform and shape their commercial practices. Furthermore, I have suggested that there are, roughly speaking, two elements of health apps we should understand before we can move on to a critical ethical analysis of the commercial practices of health apps: (1) the technological capabilities and data practices of these apps; and (2) the carefully crafted health discourse the apps employ.

This chapter has been an attempt to capture the health discourses one can encounter in health apps and to explain how they serve the business models and commercial practices of these apps. The user experience of a health app is, for an important part, determined by the health discourses that are used to frame the app’s content and functionalities. Although health is desired by everyone, it is very hard to pin down what exactly health means. The persistent philosophical disagreements on what health means is a case in point. Because health is universally desired yet difficult to define, it lends itself perfectly to *strategic* (re)formulations that frame it in a very *particular* manner to serve particular goals and interests. Health apps do precisely this.

¹¹³ <https://www.fitbit.com/nl/whyfitbit> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/qw8nb5u>.

Health, I have argued, has come to mean *wellness*: a rather vague and broad term used to refer to products, activities, and services that are ‘good for you’ and ‘make you feel good.’ Everything someone does or desires can – and is – framed as a health matter. Health, moreover, has become *healthism*: the achievement and continual maintenance of health has become an important social and moral demand, meaning people are expected to take care of and (self-)manage their health. The underlying contemporary logic of health prescribes that one has to continually *prove* that one is healthy. Within such a culture, which is mirrored and eagerly reproduced by health apps, commercial health practices in health apps can flourish.

If we combine the insights from this chapter with the insights from the first chapter, we are left with the following conclusion. As commercial services, health apps can – and in fact do – try to influence us and our behavior in a variety of subtle, mutually enforcing ways. Through the careful crafting and perpetuation of a very particular health discourse, these apps try to influence how we understand health and how we understand ourselves in relation to our own health. On top of this, health apps are acutely aware of the potential of applying insights from behavioral economics to app design – especially when ‘supercharged’ with (user) data. Health apps, like many other popular apps, are essentially big optimization experiments, where ‘optimization’ should be understood as optimizing user engagement and then conversion.

In the next chapters, I will turn to an ethical and legal analysis of the commercial practices of health apps. My main focus will be the question of whether health apps are manipulative and, in effect, violate the autonomy of users in problematic ways. Based on the descriptive work done in this chapter and the previous one, there is reason to suspect that health apps sometimes try to manipulate the desires and/or behavior of users under the (apparently innocent) guise of ‘being concerned with the user’s health.’

Chapter 3

Autonomy, Vulnerability, Trust, and Manipulation: Key Concepts for the Ethical Evaluation of For-Profit Health Apps

1. Introduction

In chapters 1 and 2 I introduced the phenomenon of health apps. I explained what health apps are, how they are designed and function, and what the scope of this dissertation is, namely for-profit health apps. For-profit health apps are of special interest because they need to monetize their userbase while offering ‘solutions’ for a variety of health, lifestyle, and wellness challenges. In doing so, for-profit health apps have to strike a balance between helping users and profiting from users. Those two goals do not have to conflict; a for-profit health app can try to design a commercial service that helps users as well as it possibly can, thus creating satisfied customers. Health apps can thus be used to help promote the universally desired and needed value of health.

There is, however, also a risk of for-profit health apps *misusing* people’s desire and need for health. As I explained in the previous chapters, most popular for-profit health apps are *freemium* apps that rely on the creation and nurturing of long-term relationships with their users in order to monetize their userbase. Health apps are in control of the digital environment within which these relationships are developed, so they are also able to (try to) *shape* these relationships to a significant degree. Health apps can also collect data on users (data which encode their characteristics, behavioral tendencies, as well as actual interactions with the app), which can inform the design and management of the digital environment. Put simply, there can be an

incentive for for-profit health apps to target and exploit people's desire and need for health and to use their technological capabilities to *create* profitable relationships with users. There is a risk that users misconstrue the nature of the relationship, or do not notice or understand the extent to which the relationship is engineered to take a particular form or shape.

To draw the line between, on the one hand, legitimate commercial health app practices, and, on the other hand, illegitimate ways of influencing health app users for commercial gain, I need a normative framework. In this chapter I provide the conceptual foundation for this normative framework, by discussing the key concepts needed to evaluate the commercial practices of for-profit health apps. I build on four interrelated concepts: autonomy, vulnerability, trust, and manipulation. Why these four concepts? Put simply, autonomy is the value that is at stake in digital environments that shape human behavior. The concept of manipulation is promising because it can help explain *how* and *why* autonomy can be at stake in the context of health apps. Lastly, vulnerability and trust serve as important conceptual links between autonomy and manipulation. I will mainly build on Susser, Roessler, and Nissenbaum (2019a, 2019b) to argue that, in essence, manipulation is about the intentional infiltration of a manipulee's decision-making in order to dispose the manipulee to the manipulator's ends by way of targeting known or inferred weaknesses of the manipulee, in a manner that is meant not to be (fully) transparent to the manipulee. Put simply, when a person is influenced in a manipulative manner, her autonomy is threatened. Theories of vulnerability can help us get a better, more nuanced understanding of how people's weaknesses and biases can be exploited by manipulative practices. Attention for the role of trust is important because health apps (and digital for-profit services more generally) seek to build trust with users, as it helps to build profitable relationships with them. Although trust in technology can be helpful, it can also be exploited in a manipulative manner.

This chapter is structured as follows. First, I develop an account of personal autonomy that should help analyze how (commercial) health app practices can impact user autonomy. My aim is to formulate an uncontroversial account of autonomy that can help us identify and think through some of the important normative tensions that come with the introduction and everyday use of popular for-profit health apps. To arrive at such an account, I will focus on the everyday, non-ideal circumstances under which our autonomy is shaped and practiced, both in general and in the context of health apps. From the start, I acknowledge that autonomy is never perfect or complete, and always both empowering and restrictive.

I conclude that we can rely on procedural notions of autonomy to *conceptually define* autonomy, but that we should emphasize autonomy's inherent *relationality* to understand and appreciate how autonomy figures in our everyday lives.

Second, I elaborate on the concepts of vulnerability and trust. I further develop the conceptual entwinement between autonomy and vulnerability. Moreover, I explain how the concept of vulnerability can be understood in a differentiated manner (following Mackenzie, Rogers & Dodds 2014 and Mackenzie 2016), while still acknowledging that “ontological or universal vulnerability is simply a shared universal *fact* that characterizes human life” (Straehle 2017: 1). I highlight how different people can experience different vulnerabilities under different circumstances. Such a concept of vulnerability is helpful in analyzing how manipulative influences work, why they are problematic, and how they threaten to undermine autonomy. I then turn to the concept of trust to explore how it can help enable agency in the face of vulnerabilities, but how, at the same time, the reliance on trust *itself* makes us vulnerable to others.

Third, I develop a conception of manipulation. I argue in favor of a moralized conception of manipulation and emphasize that it is an intentional process of instrumentalizing others by targeting and exploiting vulnerabilities or weaknesses while not drawing explicit attention to the manipulative nature of the attempt. I also explain that not only individual interactions between health apps and a user can be manipulative, but that also entire digital environments such as (parts of) health apps can be manipulative.

Fourth, I address the relation between autonomy, vulnerability and trust, and manipulation. I argue that one should not only focus on individual interactions and decisions to check whether those individual decisions are the result of manipulation and thus non-autonomous. It is equally important to consider how environments that are designed to be manipulative can undermine social conditions that help develop and sustain autonomy.

Lastly, I briefly discuss the concept of power, which figures in the background of my analysis. As choice architects designing and operating data-driven dynamically adjustable choice architectures, health app providers are in a position of power. To properly appreciate their position of power, we need a sufficiently broad concept of power.

2. Theorizing Autonomy

2.1. Autonomy as an Uncontroversial Starting Point in Liberal Democracies

Autonomy is both a fact of life and an uncontested value in liberal democracies. To be sure, there is persistent disagreement over the precise meaning and definition of autonomy, but as an abstract value it is and will remain a normative anchor point for the liberal democratic tradition. Mill's foundational insight that people are generally speaking the best judges of their own lives and should (while observing the harm principle) be afforded the freedom to, as Raz (1986: 380) put it, "make their own lives," is still at the heart of our liberal democracies. Autonomy, moreover, is not only valuable to individuals because it affords them authorship over their own lives. In liberal democracies, the democratic process also depends on the existence of a citizenry composed of autonomous persons, both to make the democratic process meaningful in the first place and to ensure the quality of the deliberative process. Without autonomous citizens, a democracy cannot exist.

Now, against the background of this generally accepted and uncontroversial understanding of the meaning of autonomy, it is tempting to provide a conceptual analysis of what autonomy, under ideal conditions, looks like, and to then, in the next chapter, apply this conception of autonomy to the case of health apps. Although such an approach is certainly not absurd or necessarily unproductive, it is not a very attractive strategy for developing my argument. If I were to develop a rather idealized conception of autonomy, the predictable conclusion that health apps do not conform (perfectly) to this idealized conception of autonomy would surely follow. Such a conclusion would be unhelpful, because autonomy is *never* practiced under ideal conditions and, maybe even more importantly, no one can ever be 'perfectly' autonomous. The complexity of life, both in the digital and the pre digital society, simply does not allow for 'perfect' autonomy (Roessler 2017, 2021).

So instead of first formulating an idealized conception of autonomy, and then applying it to health apps in order to evaluate them in terms of autonomy, it makes more sense to adopt a bottom-up approach and to embrace, from the start, the fact that the development and practicing of autonomy in everyday life happens under non-ideal conditions and is, put simply, messy. For my analysis this means that I will start with only a very minimal conceptual sketch of what autonomy, at its core, is about. This minimal conceptual sketch, which only serves as a start to my discussion of autonomy, should be plausible to as many readers as possible. I will then, in a next step, spend much more time on theorizing what autonomy under non-ideal conditions looks like. It is already at this stage that I will incorporate

examples from the health app context, to help me explore how autonomy and technologies that seek to have some influence on our behavior (such as health apps) can be both constraining and empowering.

My method, then, is not one of doing ideal theory, but rather one of ‘normative reconstruction.’ I start with a very minimal and (hopefully) widely acceptable conception of autonomy and incorporate real cases from the very start to tease out the normative tensions that emerge when real people use real health apps and that my conception of autonomy must be able to capture.

2.2. Autonomy: A Brief Conceptual Sketch

Let me start by briefly sketching a minimal conception of what autonomy, at heart, is about. This brief sketch should be seen as a starting point, providing us with a conception that we can then, in a next step, *refine* by exploring the non-ideal circumstances under which we develop and practice our autonomy. I will discuss the standard procedural account of autonomy to serve as our starting point. Such an account is called ‘procedural’ because it essentially outlines what the procedure to make up one’s mind and make decisions should look like to be considered autonomous. Procedural accounts are often contrasted with *substantive* accounts of autonomy, which are generally considered to be (much) more demanding or even perfectionist.¹¹⁴ Because I want to start with a not-too-demanding and minimal conceptual sketch that should be acceptable to most readers, I start with the procedural account of autonomy which captures what autonomy – in its essence – is about.

According to Christman (2004: 148, emphasis added), “[a]utonomy has come to be understood in this literature as *competence* in reflection and decision making and (on some views) *authenticity* of values, desires, and so on that constitute the person and motivate choice.” Elsewhere, Christman (1989: 3) summarizes the concept of

114 Where procedural accounts describe what the procedures of decision-making should look like to declare people autonomous, substantivists emphasize that the *content* of our desires, intentions, and eventual decisions also matter, irrespective of whether the ‘correct’ procedures have been followed. Oshana (1998) is one of the most prominent substantive theorists. By discussing counterexamples such as the ‘contented slave’ who – in a manner that satisfied all procedural requirements – sells himself into slavery, substantive theorists seek to show that certain decisions can be non-autonomous *qua their content* (Oshana 1998: 86-89). Moreover, strong substantive theories such as Oshana’s also require a range of socio-economic conditions to be in place for the people (and the decisions they make) to be called autonomous. Substantivists like Oshana thus argue that the presence of the right socio-economic conditions and relational embeddings of agents are *part of the conceptual definition of autonomy*, rather than ‘merely’ causal contributors to the development and practicing of autonomy. Substantive theories, especially the stronger ones such as Oshana’s, are generally considered to be rather demanding.

autonomy as referring to “an authentic and independent self.” Competency and independence refer roughly to the same thing. Both terms are meant to capture that for people to be autonomous, they require certain competencies which allow them to act as competent decision-makers who can reason and act in a sufficiently independent fashion. Authenticity refers to the requirement that autonomous people act on the basis of desires, ideas, and values that are truly their own. Let me elaborate on these two main building blocks of procedural accounts of autonomy.

Starting with the second element – independence or competency – we can say that a person is autonomous to the extent that she herself, and not others (hence the independence), is in control of her deliberations. Autonomy thus requires a particular ‘competency’ in deliberation and decision-making. Valdman (2010: 765, emphasis added) elaborates on this required competency when he writes that “autonomous agents must exercise a kind of *managerial control* over the motivating elements of their psychology; they must shape their lives by deciding which elements to act on, which to shed, and which to ignore.” An autonomous person decides independently on the basis of which of her values, desires, and goals she acts. A person acting autonomously thus deliberates with herself, as it were, taking notice of, on the one hand, the information available and the options that are open to her and, on the other hand, the “motivating elements of her psychology”. She should, then, be the one who decides how her values, desires, and goals inform her intentions for acting and, thus, constitute reasons for (not) doing something. Consider the fact that in real life we often ask people for their reasons for having done something. For example, if Jenn has decided to install the free health app Strava and purchase the additional premium service it offers, she will be able to explain that she chose to do so because 1) she loves to track her workout sessions because it motivates her, 2) she has done some research and has seen that Strava offers good tracking options and has a big community, and 3) she trusts Strava to treat her and her data respectfully, and that taken together, these gave her sufficient reason to install Strava and purchase the premium features.

We should add to this the requirement of authenticity. Authenticity refers to a person’s relationship to her own values, desires, and goals. The idea behind this notion of authenticity is that it is one thing to have ‘managerial control’ over the way the values, desires, and goals *you happen to have* inform your decisions, actions, and lifestyle, but it is quite another thing for those values, desires, and goals to be *truly your own*. In other words, we do not want to declare people autonomous who are stuck with values, desires, and goals they do not consider to be their own, for instance because they were manipulated into forming certain desires. More concretely then, what is required for autonomy is that the person’s

relationship to her own values, goals, and desires should take *a certain form*. Often, this ‘certain form’ is referred to as *identification* – the requirement of identifying with one’s own values, desires, and goals in some sense. Different interpretations of what identification should entail exist. Some argue that a person should be able to reflect critically on one’s own values, desires, and goals as they are, and should, then, endorse them (Dworkin 1989; Frankfurt 1971). Others, most prominently Christman (1991), focus on the historical process through which the desires, values, and goals developed: “what matters is what the agent thinks about the *process* of coming to have the desires, and whether she resists that process when (or if) given the chance” (Christman 1991: 10). Let us look at Jenn again. As we have already seen, her values, goals, and desires constituted compelling enough reasons for her to start using Strava. But what if we find out that Strava correctly identified her problematically compulsive personality which leads her to track everything, including her workouts, compulsively. If Strava targeted that part of her personality with personalized offers to try to get her to install a tracking app that is going to exacerbate her compulsive behavior, and if Strava used that knowledge to engage in personalized upselling to get her to purchase the premium services, then we can question whether Jenn’s decisions were sufficiently autonomous. Even though Jenn is an intelligent and capable person who can think coherently and clearly about her decisions, she is also stuck with some desires and impulses that she does *not* identify with – they are not authentically hers. When precisely those desires and impulses are identified, targeted, and leveraged against her, they can ‘taint’ the procedure to the extent that they figure in the decision-making procedure. Put simply, for a decision-making procedure to be truly autonomous, the ‘input’ needs to be sufficiently authentic.

The joint requirements of independence (or competency) and authenticity make for a standard procedural account of autonomy. The procedural account captures the basic intuition that autonomy is about people’s ability to think for themselves, make their own decisions, and give shape to their own life, informed by what they themselves truly care about.

2.3. From Ideal Theory to Normative Reconstruction: An Example from the Health App Context

The conceptual sketch of the procedural account gives us a good idea of what autonomy is – in essence – about. But it is unsuitable to serve as an evaluative standard for health apps because it is too abstract and ideal. To see why, it can be helpful to briefly consider an example to help us draw out the kind of normative tensions we want our conception of autonomy to address.

Consider the already mentioned health app Headspace. Headspace is a freemium mental wellness app that offers meditation sessions to its users. The app can be installed for free and offers a few meditation ‘packs’ for free. Access to additional meditation packs must be purchased. Headspace’s business model is to attract users to the app with its limited set of free meditation exercises and to try to build an ongoing relationship with these users to retain them as users and to, hopefully, get them engaged enough with the service that they purchase additional meditation packs or subscriptions. As articles on Headspace as well as the job description accompanying some of their vacancies clearly indicate,¹¹⁵ the app is designed and run as an advanced digital environment that can continuously measure how users interact with the environment in order to respond in real time and in a personalized manner to those user interactions. Headspace’s business model, in combination with its technological capabilities, gives rise to a potentially difficult tension. On the one hand, we can assume that Headspace wants to offer a service that is helpful to its users and it can use the user interaction data it collects to make its services more helpful to users. On the other hand, Headspace is a for-profit enterprise with an incentive to design and operate its digital environment to shape the relationships with its users to be as profitable as possible. However, designing a digital environment that shapes *profitable* relationships with users is not necessarily the same as designing a digital environment that aims to be as helpful as possible. For example, with all the user data Headspace collects, they can build a pretty good profile of individual users’ anxieties, fears, and insecurities. Such information can be used to sell more services, but tapping into anxieties, fears, and insecurities to sell meditation sessions that (indirectly, maybe) address those anxieties, fears, or insecurities is not always the best course of action from a mental health perspective. Here, the fine line between designing helpful or exploitative digital health environments clearly presents itself.

This is not the place to draw such lines. But as the example hopefully shows, the autonomy of the user vis-à-vis apps such as Headspace is at stake in complex and maybe even ambiguous ways. It clearly makes no sense to declare a health app such as Headspace to be *either* autonomy-enhancing *or* autonomy-undermining. As users start and continue to use such an app, a dynamic relationship with the app emerges *within* the digital environment that is designed and controlled by the app. To properly understand how the autonomy of the user is at stake here, we need more than a conception that outlines, in an abstract manner, when and how

115 See, for instance, Chaykowski (2017) for an article on Headspace’s data-driven business model. For an example of one of Headspace’s job postings detailing the data intensive operations of the app, see: <https://imgur.com/a/uIwpvds>.

individual, separate decisions of users are either autonomous or non-autonomous. What we also have to capture is how the *dynamic relationship* between the user and the app (and the digital environment that can be adjusted differently in response to different users) develops *over time* and how that can both enable and constrain a person's autonomy.

To do so, I propose we turn to theories of relational autonomy to help us theorize how our social and relational embedding into the world both enable and constrain our autonomy.

2.4. Relational Autonomy

Procedural accounts of autonomy are said to implicitly presuppose that idealized agents that are perfectly rational, individualistic, and “unencumbered” (Sandel 1984) are the persons capable of living an autonomous life. For example, the independence condition as I discussed it might suggest that *the more* independent someone is – and, thus, *the more* someone can exercise *perfect* control over all the motivation elements of her psychology – *the more* this person can be seen as being autonomous. Similarly, the authenticity condition might suggest that those people who *only* act on the basis of their *own* individual, atomistic desires, values, and goals, are truly autonomous.

The central objection that is articulated by the literature on relational autonomy is that such an idealized conception of the person neglects the inherent “intersubjective and social dimensions of selfhood and identity for individual autonomy and moral and political agency” (Mackenzie & Stoljar 2000: 4). This negligence of the intersubjective and social dimensions of autonomy is important to address for at least two reasons. Firstly, a purely individualistic, atomistic, rationalistic conception of the person and the related conception of autonomy are hardly attainable (if possible at all), and brush over the fact that we need social relations to be(come) autonomous. Secondly, even if such an idealized conception of autonomy were attainable, it would be unattractive for (nearly?) everyone to live such a ‘perfectly autonomous’ life. Let me elaborate on these reasons.

Relational autonomy scholars argue that the conventional literature on autonomy silently presupposes an “[a]utonomous man [who] is – and should be – self-sufficient, independent, and self-reliant, a self-realizing individual who directs his efforts towards maximizing his personal gains” (Code 1991: 78). This silent presupposition does not do justice to the fact that humans are fundamentally vulnerable and dependent on others. Interdependency and vulnerabilities are, from the perspective of procedural theories of autonomy, a threat to autonomy and,

therefore, undesirable. In reality, however, interdependencies and vulnerabilities are – in the right amount and constellations – an enrichment of peoples’ lives, as well as unavoidable parts of them. By explicitly theorizing the “intersubjective and social dimension of selfhood and identity” (Mackenzie & Stoljar 2000: 4), the literature on relational autonomy provides an important and essential correction to often overly idealized conceptions of autonomy.

Consider the fact that we all need or have needed our parents, our friends, our significant other(s), etc., to live meaningful and autonomous lives. They have raised us, or listened to us when we needed to vent, or tried to make helpful suggestions when we were in doubt, or consoled us when we felt down, and so on. Moreover, we need not just those close to us, but other people in general. Society could not function if there were *no* interdependencies between people. We all need (anonymous) others to, for example, produce the goods and provide services we need to live our lives – people can never be completely “self-sufficient, independent, self-reliant, and self-realizing” (Code 1991: 78). Importantly, many of the relationships we have with other people are actually constitutive of our autonomy, since we need to (learn to) develop our capacity for autonomy. “If we ask ourselves what actually enables people to be autonomous, the answer is not isolation, but relationships – with parents, teachers, friends, loved ones – that provide the support and guidance necessary for the development and experience of autonomy” (Nedelsky 1989: 12).

Besides the fact that interdependencies (which also render us vulnerable) are unavoidable and at the very same time constitutive of autonomy, it would be equally *unattractive* to live a completely self-sufficient, independent, self-reliant, and self-realizing life, even if we could. The ideal of ‘perfect autonomy’ is, for many people, itself normatively undesirable.¹¹⁶ For example, being involved in loving relationships also means that one has to take the values, goals, and desires of others into consideration, which is formally speaking a constraint on a person’s self-sufficiency, independence, and self-reliance (and thus autonomy). We do, however, willingly embrace such constraints (if we can call them ‘constraints’ at all) because self-sufficiency, independence, and self-reliance are not absolute ideals.

116 It is important to note that given my commitment to a procedural account of autonomy (to *define* what, in essence, autonomy is about), I accept the possibility that there are persons who do have the authentic desire to live ‘as autonomously as possible’, understood in individualistic terms. Think of a hermit, living in a remote location in a completely self-sufficient manner. Such a hermit should be considered fully autonomous if the hermit *truly* wants to live that way. As a matter of empirical fact, however, I do believe that the relational autonomy literature is correct in claiming that engaging in (the right sort of) relationships contributes *causally* to the development and practicing of autonomy for almost everyone.

If they were, we would not make ourselves vulnerable to others. But we do just that. Loving relationships are an enrichment of our lives, even though they make us vulnerable; we make our happiness partly dependent on others, without any absolute guarantees.

Relational autonomy theorists thus emphasize that interdependency and vulnerability are social facts of life that should figure in a theory of autonomy. We *need* social relations to live autonomously, even though they render us vulnerable to others. Moreover, we also *want* to live our lives with others in the various social contexts we have to navigate. Later in this chapter I will say more about the ‘entwinement’ of autonomy and vulnerability. For now, I want to emphasize that by theorizing the inherent relational dimension of autonomy, we are also better equipped to understand and evaluate how the relationship between a user and an app that develops over time can impact the user’s autonomy. In the digital society, other people *as well as technology* are part of the (social) contexts in which we develop and practice our autonomy. An important theoretical question that still needs to be addressed is how the minimal procedural account I started out with relates to the relational dimension of autonomy I just highlighted. The next section addresses this question.

2.5. Reconciling Procedural and Relational Autonomy

How to reconcile a commitment to proceduralism while still acknowledging the important insights from relational theories of autonomy? What we want to hold on to from procedural theories of autonomy is the central idea that autonomy requires a specific kind of *competence* which allows for (relatively) independent decision-making. At the same time we want to hold on to the idea that autonomy is only learned, practiced, and improved within social contexts, together with other people.

Christman (2005: 155) helpfully suggests to focus on competency conditions to see how procedural autonomy and relational autonomy can come together: “it is *competency* conditions in proceduralist views of autonomy that are problematic insofar as they do not include or make room for the wide variety of capacities for care, intimacy, social interaction, and the like that will be crucial for socially embedded persons to flourish.” If we focus on: 1) the *kind* of competency required for autonomy, and 2) the social conditions under which we learn and practice that competency, we can preserve the central insights from both the types of theories of autonomy.

An answer along these lines can still take on different – i.e., more and less demanding – forms. The demandingness of different answers will be determined by the types of social (and economic) conditions one requires to be in place for the development and practicing of one’s autonomy competences to be possible. Demanding answers can be rather attractive in at least one sense – by providing elaborate and precise accounts of the conditions that need to be in place for autonomy to be developed and practiced, they are rather informative. In other words: such more demanding theories do not only tell us that socio-economic circumstances matter, but also *how*. But this quality of being informative comes at the price of being demanding. The more precisely one describes all the socio-economic conditions that need to be met, the more people will be declared non-autonomous (or at least *less* autonomous) because they live under non-ideal conditions.

Consider, for instance, Oshana’s (1998) substantive theory of autonomy. What *conceptually defines* autonomy according to Oshana is not only being an agent that decides and reflects independently and authentically (i.e., competently), but, moreover, “autonomy only obtains when social conditions surrounding an individual live up to certain standards” (Christman 2004: 150). When social conditions are sufficiently suboptimal, causing them to fall short of the substantive standards defined, persons living under these social conditions are *by definition* lacking autonomy. Another example is Anderson and Honneth’s (2005) relational account of autonomy, informed by Honneth’s (1996) theory of recognition, to provide a conceptual framework in order to help explain why personal autonomy can only be achieved *interpersonally*. They explain how the social contexts we inhabit must provide us with three modes of recognition – namely self-trust,¹¹⁷ self-respect, and self-esteem – for us to be able to (come to) *understand* and *trust* ourselves as persons capable and worthy of autonomous decision-making (Honneth 1996: 129; Anderson & Honneth 2005: 131; Anderson 2014: 141).

Precisely because of their explanatory richness, I do want to build on substantive theories such as Anderson and Honneth’s (2005). But I want to build on such theories in a specific way, namely as a kind of interpretational horizon which helps us understand how autonomy can be developed, practiced, as well as threatened in various contexts. Such substantive theories can be very helpful without there being a need to distill from them a list of *necessary* conditions for the practicing and development of autonomy. In the next chapters, I mainly want to argue that a more structural, relational approach to digital choice architectures (such as health apps) is needed to understand and evaluate both their promises and risks. To do

117 (Basic) ‘self-confidence’ in Honneth 1996.

so, the ‘simple’ insight that autonomy is relational is already very helpful and more substantive theories can help explain *how* autonomy’s relationality takes shape. I just do not want or need to translate the explanatory richness of substantive theories into a list of demanding necessary conditions.

I want to follow Roessler (2019, 2021) who offers a significantly less substantive account of relational autonomy. As persons, we are always involved in various social contexts. It is in those social contexts that we can (and have to) develop and practice our autonomy. In this sense, autonomy is *necessarily* relational; there is no way to live one’s life – and, in effect, to develop and practice one’s autonomy – outside of social contexts. Now, it is clear that some social contexts tend to be more conducive of (the development and practicing of) autonomy than others. For example, growing up in an abusive home in a country that is ruled by an authoritarian dictator is not exactly helpful when it comes to growing up to be an autonomous individual. Theories such as Oshana’s and Anderson and Honneth’s can certainly help us to see why that is the case. But it does not follow that we must conclude our hypothetical person simply *lacks* the conditions or resources to live an autonomous life. It will be more challenging to develop the necessary competency to act autonomously, but not impossible. Our hypothetical person will still have to navigate various social contexts and it would be counterintuitive to argue that those social contexts simply cannot offer possibilities to develop and practice one’s autonomy. Theorizing the relational nature of autonomy thus does not (exclusively) come down to the formulation of precise and strict *requirements* for the social contexts we navigate. It is, certainly for my purposes in later chapters, more helpful to emphasize the fact that autonomy is necessarily relational in nature, without committing to substantive theories of *how* social contexts must (necessarily) be arranged to be conducive of autonomy. “It is one thing to say that models of autonomy must acknowledge how we are all deeply related; it is another to say that we are autonomous *only* if related in certain idealized ways” (Christman 2004: 151).

2.6. Conclusion

In these sections on autonomy I have been searching for a concept of autonomy that can help me tease out and evaluate the precise normative tensions that can arise when real people use real, popular for-profit health apps. So let me briefly revisit the Headspace example to see whether the conception I have been developing can capture the normative tensions that such health apps can give rise to. The core tension that needs to be explored in the next chapter is how we should understand and evaluate the dynamic *relationship* that develops between a user and an app within the digital environments that the apps themselves design and control.

These digital environments provide opportunities to ‘scaffold’¹¹⁸ the autonomy of users, but we can also identify commercially-inspired incentives to design the digital environments in ways that are not aimed at supporting – and moreover risk undermining – the user’s autonomy. Our conception of autonomy should thus allow us to look beyond the individual and the separate, isolated decisions the user makes and also be attentive to the ways in which digitally mediated relationships shape our possibilities for developing and practicing autonomy over time. To meet this requirement, I build on an uncontroversial, minimal conception of autonomy which I interpret in a procedural as well as relational manner. Autonomy is *defined* in procedural terms, but to understand *how* autonomy can be developed, practiced, and threatened, we need to build on theories of relational autonomy which provide an indispensable interpretational horizon. My relational interpretation of autonomy is, however, not very substantive; i.e., I do not use theories of relational autonomy to introduce additional *necessary* requirements for autonomy. By emphasizing autonomy’s relationality in this manner, I should be able to address the important normative tensions introduced by health apps in the next two chapters.

3. Vulnerability and Trust

Before I continue to discuss the concept of manipulation, I first want to discuss the concepts of vulnerability and trust because they help us to better understand not only autonomy and manipulation, but also the connection between autonomy and manipulation. Autonomy, as we will see below, is *entwined* with vulnerability. At the same time, as I will argue in the next section, manipulation typically involves the targeting of vulnerabilities of manipulees. So if we want to understand how the digital environments provided by health apps can both enable and undermine autonomy, and how manipulative influences can be exerted, we should also understand how vulnerabilities can exist and manifest themselves. Additionally, I spend some time exploring the relation between trust and vulnerability. As I argue throughout this dissertation (and more specifically and elaborately in the next two chapters), popular for-profit health apps seek to create on-going relationships with their users. These relationships may sometimes resemble trusting relationships – or actually constitute trusting relationships – since users use these apps to manage something of great importance to them: their health. By exploring how vulnerability and trust are, much like autonomy and vulnerability, entwined, we also have more conceptual tools at our disposal to analyze the manipulative potential of health apps.

118 I borrow this term from Heath and Anderson (2010).

3.1. Vulnerability and Autonomy

As a start, it helps to begin with a broad and general understanding of vulnerability. Anderson (2014: 135) helpfully suggests that “a person is *vulnerable* to the extent to which she is not in a position to prevent occurrences that would undermine what she takes to be important to her.” Vulnerability thus is about one’s relation to the world, the forces (social, physical) in the world that can affect anything one deems important, and one’s (lack of) control or power over those forces. Starting from this general understanding of vulnerability, we can see how autonomy and vulnerability are entwined. As I argued above, the relational nature of autonomy bears emphasis. We are continuously thrown into various social contexts we inhabit together with other people and technology, which means we also develop and practice our autonomy together with others. These social contexts – over which we lack complete control – can, however, also change in a manner that might affect our (possibilities for) autonomy. So because we have no complete control over those social conditions that causally contribute to our autonomy, we are always vulnerable to influences that might hamper our autonomy. Consider the relationships one has with one’s parents, close friends, or partner. Such relationships can give one the support needed to make autonomous choices. We look to others for inspiration and we turn to others for advice or reassurance. It is also clear, however, that those relationships may change for reasons (partly) beyond our control and in manners that undermine, rather than support our autonomy. As a result, our dependence on others makes us vulnerable.

Being vulnerable to others is not necessarily a bad thing: to the extent that we want to live truly autonomous lives, a certain degree of vulnerability to autonomy-undermining changes (generally put) cannot be completely avoided. This is why autonomy and vulnerability can be said to be *entwined* on a conceptual as well as practical level (Anderson 2014). Vulnerability is not something that can in principle, were the circumstances more ideal, be eliminated from an autonomous life; the very conditions for autonomy *necessarily* introduce vulnerability. The fact that autonomy and vulnerability are entwined does not, however, mean that *every* type or source of vulnerability is acceptable. For example, vulnerabilities that are the result of unequal treatment on the basis of categories such as gender or race are problematic.

3.2. A Taxonomy of Vulnerability

Vulnerability then is not intrinsically problematic or undesirable, although certain vulnerabilities certainly can be problematic. So how are we to think of the normative status of vulnerability? In what follows, I will *not* discuss or develop a free-standing theory on the normativity of vulnerabilities. I will, instead, discuss a

taxonomy of vulnerability (Rogers, Mackenzie & Dodds 2012, Mackenzie, Rogers & Dodds 2014) which explicates different sources and states of vulnerability. Although this taxonomy itself does not provide *direct* answers to questions on the normative status of vulnerability, it does provide a more refined understanding of vulnerability, which, in turn, helps us to discuss the role vulnerabilities play in manipulative digital environments. By focusing on the role vulnerabilities play in manipulation and autonomy, we can *indirectly* discuss the normative significance of vulnerabilities through the concepts of manipulation and autonomy.

3.2.1. Sources of Vulnerability

Mackenzie, Rogers, and Dodds (2014) make a difference between *sources* of vulnerability and *states* of vulnerability. Starting with the sources of vulnerability, we can differentiate between ‘inherent’ and ‘situational’ sources. Inherent vulnerabilities are those vulnerabilities we experience as humans simply because we are humans. Inherent vulnerabilities are “intrinsic to the human condition” and “arise from our corporeality, our neediness, our dependence on others, and our affective and social natures” (Mackenzie, Rogers & Dodds 2014: 7). So to the extent that we agree that all humans are social beings who strive to live (relatively) autonomous lives, we can understand the entwinement of autonomy and vulnerability as a good example of *inherent* vulnerability.

Situational sources of vulnerability refer to vulnerabilities that can arise in particular contexts or situations. As Mackenzie, Rogers, and Dodds (2014: 7) emphasize, a wide variety of influences of a different nature – “personal, social, political, or environmental” – can cause or exacerbate vulnerabilities. For example, natural disasters can make people vulnerable to having their house destroyed, and (unjust) social policies can make people vulnerable to unfair treatment. Moreover, the impact these kinds of influences have on individuals or groups are also mediated by the characteristics of their specific situation. Not every person responds, or can respond, in the same manner to the same influence. People live in different social communities, have different psychological characteristics, different levels of affluence, and so on, all of which can influence how one is (or is not) rendered vulnerable to particular threats.

Mackenzie, Rogers, and Dodds (2014) also highlight the temporal dimension of situational vulnerabilities. Contrary to inherent vulnerabilities (which are intrinsic to the human condition and therefore always present), situational vulnerabilities can “be short term, intermittent, or enduring” (Mackenzie, Rogers & Dodds 2014: 7). For example, a person can go through a very bad break up and, as a result, be especially vulnerable to bad social influences for some time (say a few weeks or

months) as a result of being emotionally devastated. We can also imagine situations where a person is systematically vulnerable to a range of influences as a result of a social situation that is nearly impossible to change. Think of systemic racism, which can render people permanently vulnerable to a wide range of unjust practices.

This last example of systemic racism also reveals that the distinction between inherent and situational vulnerabilities is not always easy to maintain. A situation of systemic racism is not, in a metaphysical sense, a necessary feature of the human condition and therefore falls outside of the category of ‘inherent vulnerability.’ In practice, however, systemic racism can very much feel like an inescapable, *inherent* feature of one’s situation or life in general. In such cases, the formal distinction between inherent and situational vulnerabilities can be difficult to maintain. Moreover, the fact that some vulnerabilities – such as the vulnerability of one’s body to injury – are inherent to the human condition sometimes invites sentimentalist commentary. Inherent vulnerabilities are sometimes presented as an ‘equalizer’ between humans: irrespective of socio-economic or cultural circumstances everyone is, deep down, equally vulnerable *qua human being*, or so the story goes. Although this sentiment is understandable and, at a purely conceptual level, not untrue, I want to caution against such a sentimentalist understandings of (especially) inherent vulnerabilities because it can make us lose sight of the many ways in which vulnerabilities are *not* evenly and justly distributed. For now, I just want to flag that I am aware of these issues and will be cautious not to rely on such sentiments.¹¹⁹

In the context of health apps specifically we can also expect interesting connections between inherent and situational vulnerabilities to arise. One of the defining aspects of the human conditions is people’s inherent vulnerability to diminishing mental or physical health, either through injuries or steady, expected decline. This inherent vulnerability interacts with situational social aspects that can impact one’s health (e.g., one’s upbringing, socio-economic status, availability of ‘healthy’

119 I am finalizing this chapter during the 2020 corona crisis. Interestingly, this crisis has come with a lot of attention for the role (the concept of) vulnerability plays in it. Some have argued that because the virus does not discriminate, it reminds us that as humans we are all equally vulnerable in some respects. Nancy (2020) made such a suggestion when he wrote that: “In fact, the virus actually communizes us. It essentially puts us on a basis of equality, bringing us together in the need to make a common stand.” Others, however, have argued that coronavirus clearly exposes why we need a concept of “differential vulnerability” (Lorenzini 2020). Although people are *qua human beings* equally vulnerable to the virus, we also clearly see that people living in less privileged communities tend to be hit harder by the virus. For example, the choice to self-quarantine at home to avoid infection also presupposes an *ability* to do so. One’s employment situation, home situation, and financial situation all influences one’s ability to stay home or not.

options) as well as with situational characteristics related to one's personality (e.g., one's psychological and affective biases, one's insecurities) that can affect one's susceptibility to having one's health-related behavior influenced by health apps.

3.2.2. *States of Vulnerability*

Besides different sources of vulnerability, we can also distinguish between different *states* of vulnerability. Mackenzie, Rogers, and Dodds (2014: 8-9) use the categories of 'dispositional' and 'occurrent' to describe states of vulnerability. Both inherent and situational vulnerabilities can be dispositional and occurrent. Dispositional vulnerabilities can roughly be translated to potential vulnerabilities; this category describes vulnerabilities that are not present yet, but that can manifest themselves given the inherent or situational features of the person or situation involved. We speak of an occurrent vulnerability when a dispositional vulnerability actually occurs. Mackenzie, Rogers, and Dodds (2014: 8) provide a helpful example:

[A]ll fertile women of childbearing age are dispositionally vulnerable to life-threatening complications in childbirth. But whether or not a pregnant woman is occurrently vulnerable to such complications will depend on a range of factors, both inherent and situational, such as her physical health, medical history, socioeconomic status, geographical location, access to health care, and cultural norms relating to pregnancy and childbirth.

The distinction is helpful because we do not only want to focus on vulnerabilities that have already manifested themselves, but also want to focus on vulnerabilities that could potentially come into existence – or be 'activated' – under particular circumstances. In the context of health apps, the distinction also promises to be helpful. Popular for-profit health apps offer sophisticated digital environments to users that can target groups or individuals dynamically and can personalize content or interventions dynamically. In such a digital environment, we should not only focus on those vulnerabilities that are occurrent, but also on dispositional vulnerabilities that can become relevant in response to the dynamically changing digital environment. The occurrent-dispositional distinction can thus also be action-guiding if we require actors to not only be attentive to occurrent and more

easily observable vulnerabilities,¹²⁰ but also require them to take into account dispositional – not yet present but foreseeable – vulnerabilities that might manifest themselves as a result of the actions of another actor. For example, we do not only want health apps to be cautious when targeting occurrent vulnerabilities related to someone’s age or present health status. Precisely because digital (health) environments are becoming increasingly good at discovering (or even *creating*) more situational and dispositional vulnerabilities,¹²¹ we might also want to evaluate the ways in which health apps deal with such ‘dynamic’ vulnerabilities.

3.3. Vulnerability and Trust

Vulnerability and trust are, much like vulnerability and autonomy, entwined (Wiesemann 2017). This entwinement can be understood in two related ways. Firstly, trust is an important social practice or institution that helps us navigate a social world where vulnerabilities are always (dispositionally or occurrently) present. Secondly, to the extent that we actually get involved in relationships of trust, those relationships themselves introduce vulnerabilities.

As we have already seen, vulnerability in its most basic sense could be understood as “the extent to which [someone] is not in a position to prevent occurrences that would undermine what [that someone] takes to be important” (Anderson 2014: 135). Vulnerability has its roots in a world where other people, as well as social and natural forces, can exert influences on us that are (partly) beyond our control. To still get things done in such a world of uncertainty (with the vulnerabilities that result from that uncertainty) we have to rely on a range of social practices and institutions. One obvious example is the enforceable legal contract where two parties agree on terms and make themselves accountable to each other. Another equally important but less formal social practice or institution that structures our social world in the face of uncertainty is *trust*.

Although philosophers disagree on the precise meaning, nature, and significance of trust (see, e.g., McLeod 2015 for an overview of different philosophical discussions on trust), we can point to some uncontroversial core features. First of all, as Hardin

120 It should be noted that it is not always true that occurrent vulnerabilities are easier to detect than dispositional vulnerabilities. For example, a person may be struggling with an eating disorder (making this person occurrently vulnerable) without this being immediately clear and observable to the outside world. A healthy woman that is 9 months pregnant, on the other hand, is clearly dispositionally vulnerable to medical complications; everybody can see she is pregnant and everybody understands that pregnancies come with the risk of medical complications.

121 Think, for instance, of a health app that is able to discover that a perfectly healthy, young, affluent, successful person tends to experience brief but intense moments of insecurity concerning his/her body at very predictable times.

(2002: 9)¹²² points out, “a characteristic of trusting relationships [...] is that trust is generally a three-part relation: A trusts B to do X.” So usually a trust relationship involves a party (A) that has to ‘place one’s trust’ in *another* party (B) because there is something (X) of value that A needs to trust B with. This ‘something of value’ can be nearly anything: I can trust a health app to keep my data safe, as if my data were an object that I handed over for safekeeping. I can also trust the health app to only *use* my data for particular purposes, such as providing me with the best possible health advice and not, for instance, to build a persuasion profile (Kaptein 2015) for me to target my insecurities. Secondly, trust relationships are only possible, and at the same time useful, under conditions of *uncertainty*. When A trusts B to do X, A cannot be certain that B will indeed do X. In this sense, trust necessarily involves ‘a leap of faith;’ trust comes without absolute guarantees, for trust with absolute guarantees ceases to be trust.

That is not to say that people generally display ‘blind’ trust towards other people or institutions. When entering into a trusting relationship, we usually try to consider the *trustworthiness* of the other party.¹²³ What it is that makes actors trustworthy is subject to philosophical debate. Baier (1986: 235) famously argued that “[w]hen I trust another, I depend on her good will toward me.” What makes others trustworthy in this account is our estimation/assumption that the other party will act out of good will towards us. Another famous account is Hardin’s (2002: 13) trust-as-encapsulated-interest account: “[y]our trust turns not directly on your own interests but rather on whether these are encapsulated in the interests of the trusted. You trust someone if you believe it will be in her interest to be trustworthy in the relevant way at the relevant time, and it will be in her interest because she wishes to maintain her relationship with you.” So what makes other people trustworthy in this account is our rational belief that it is in the other party’s interests to also take our interests into account.

It bears emphasis that “unlike some other values, trust is a psychological state that represents the trusted person or object as being trustworthy, and this may or may not actually be the case” (Nickel 2015: 552). There is, in other words, no *necessary* connection between a person’s, system’s, or object’s (inherent) trustworthiness and the trust people have in said person, system or object. It follows that to secure someone’s trust, one can 1) try to be as trustworthy as possible in order to ‘earn’ other people’s trust; or 2) try to ‘engineer’ circumstances, interactions or appearances

122 Baier (1986: 236) proposes a similar three-part relation to model trust relationships.

123 Hardin (1996: 28) even argues that most philosophical discussions of trust are, in essence, discussions of trustworthiness: “Surprisingly, much of the literature on trust hardly mentions trustworthiness even though much of it is primarily about trustworthiness, not trust.”

that are likely to induce the psychological state of trust, without there (necessarily) being an underlying source of real trustworthiness to warrant the engineered trust: “Showing that people trust (within) a design does not imply that it is trustworthy, nor the other way around” (Nickel 2015: 559). Contemporary digital environments that can collect large amounts of user data and can dynamically personalize the user experience seem to be especially well-suited to “induce” psychological states of trust in users (Nickel 2015: 553). This ability to engineer trust can be helpful in contexts where the existence of trust enables helpful/valuable interactions *and* where the parties involved are, in fact, trustworthy. There is, however, also a risk of digital technologies focusing predominantly/exclusively on providing all the right cues to induce psychological states of trust, without attention for or investments in *actual* trustworthiness. In the context of health apps we also want to be able to differentiate between, on the one hand, apps that, for strategic reasons, focus on engineering the right circumstances to induce trust in order to exploit that trust, and, on the other hand, health apps that really are trustworthy and therefore *deserve* their users’ trust.

Even though we might disagree on what it is that makes people trustworthy – I will not try to answer that difficult question here – it is evident that *whatever it is that makes people trustworthy* matters, because our estimation of the trustworthiness of a person, system, institution, or app is an important predictor of our ability or willingness to enter into relationships of trust with another party. Moreover, all accounts of trust and trustworthiness agree that entering into a relationship of trust necessarily makes the person placing one’s trust in another party dispositionally vulnerable to have that trust betrayed. Whatever account of trust and trustworthiness one prefers, in the end all accounts have to explain how *something* that is, strictly speaking, uncertain, can provide us with enough reason(s) to place our trust in the other party. Whatever this *something* is, basing our trust on it can backfire. Baier (1986: 235), for instance, write that “[w]here one depends on another’s good will, one is necessarily vulnerable to the limits of that good will.”

Precisely because trust is entwined with vulnerability, the *betrayal* of trust (and voluntarily accepted vulnerability) is often seen as a social or even moral transgression. As Baier (1986: 239) observes, “[i]f part of what the truster entrusts to the trusted are discretionary powers, then the truster risks abuse of those and the successful disguise of such abuse. The special vulnerability which trust involves is vulnerability to not yet noticed harm, or to disguised ill will.” The betrayal of trust is thus also, in a way, an abuse of power exercised to take advantage of the truster. Taking advantage of the ‘special vulnerability’ that is constitutive of trust can help explain why a betrayal of trust raises moral questions.

That is not to say that the betrayal of trust is *intrinsically* morally problematic and impermissible. Think for instance of a terrible dictatorship where an opponent of the regime infiltrates the highest ranks of the dictatorship in order to earn the trust of the dictator. Once she has earned this trust, she takes advantage of it to help overthrow the dictatorship. Clearly, this particular breach of trust is not, all things considered, problematic or impermissible. Considered in isolation, a breach of trust can probably be considered a *prima facie* moral wrong, but once we consider other elements of the case and the other values and interests involved, our evaluation of the case may change.

Even if we agree that, in principle, a particular breach of trust is morally problematic, we can still ask *how* problematic it is. Depending on what it is we trust others with, our evaluation of the severity of the transgression can vary. If I trust someone to hold my sandwich while I get my coat and the person I entrust with my sandwich eats it or throws it on the floor, then the betrayal of my trust certainly affects our personal relationship. But because the valued thing – a sandwich – which was the object of the trust relationship is not all that important to me, the harm can be said to be comparatively minor. If, however, I trust someone with my deepest fears, insecurities, or secrets and the person goes on to betray my trust, then the betrayal will feel much more severe because the betrayal revolves around something that is very important (and intimate) to me. In short, the importance of the object of trust determines for the most part how severe the possible betrayal of trust is.¹²⁴

It follows that in the context of health apps, the importance of health to people should figure in our evaluation of popular for-profit health apps and their attempts to build trust(-like) relationships with users. Such relationships between a health app and their users are not problematic by their very nature; quite to the contrary, trust can serve an important social function and the potential usefulness of health apps can be explained partly by their ability to build trust(-like) relationships with their userbase. However, when that trust is betrayed or misused, we can speak of a serious moral transgression precisely because health is so important to people. Here we also see one of the reasons why health apps are different compared to other digital technologies and environments. Health apps are certainly not the only technology aimed at building trust(-like) relationships with their users, but out

124 The nature of the object of trust is not necessarily the only factor determining the severity of a possible betrayal of trust. For example, *how* people act in order to mislead us, resulting in a betrayal of trust, can also determine how we evaluate the betrayal. When someone develops elaborate, sneaky plans to take advantage of a trust relationship, we can certainly feel a deeper sense of betrayal than in cases where someone impulsively betrays our trust, or when someone betrays our trust out of desperation.

of all those widely available consumer-facing digital technologies, few deal with a value as central and important to people's lives as health. So if, as I have argued, trust also introduces a "special vulnerability" (Baier 1986: 239) to have one's trust misused, and if the misuse of trust is especially problematic in cases where something important or of great value is at stake, then trust(-like) relationships in the commercial context of health apps certainly deserve our scrutiny.

At this point it could be asked whether 'trust' is really the right concept to use in this context. Can we really say that people can (come to) *trust* a health app? Or is trust too 'heavy' a concept to use and would it be better to state that people use a health app because they have certain interests, expectations or hopes, without necessarily *trusting* the app in question? To start, it is important to acknowledge that not every user of a (health) app will relate to the app in a similar fashion; some users may take a more cynical stance, not expecting too much from the app to begin with, while others may, over time, experience the development of a somewhat 'deeper' relation to the app. Still, the question remains whether the concept of trust is the correct concept to help us interpret and evaluate such user-app relations.

There are two reasons to retain the concept of trust. First, app developers themselves think – albeit mostly for strategic reasons – in terms of trust and explicitly state that they build their apps with the aim of developing trusting relationships with their users. On the earlier mentioned Apple Developer's page, multiple app developers are quoted on the pages that describe how to build a successful freemium app. They all emphasize – as we have seen before – that freemium apps need to build *ongoing* relations with their users. One app developer is quoted saying that "[f]rom their first session, we're trying to develop trust with the users. And we do that by demonstrating value without asking for anything from them first. Over time, as they see value in the service and the tools we provide, many are going to want to pay for those more robust features."¹²⁵ Another app developer also emphasizes the need to establish trust with the users: "[o]ne of the ways we establish trust is that when you get a trial or some free content, you're getting the same quality that you would when you pay," says Joe Ghazal, Chief Technical Officer at Originator."¹²⁶ To be sure, the trust that app developers speak about does not necessarily live up to the more substantive philosophical accounts of trust discussed above. But I do want to insist that there is something telling about app developers insisting on the importance of trust.

125 See Apple's advice page here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKlNID>.

126 See Apple's advice page here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKlNID>.

Second, the concept of trust as I discussed it still offers a useful model to understand the dynamic of the (potential) relations between users and apps, even when we use a slightly more ‘relaxed’ notion of trust. To trust one’s best friend is obviously not the same as trusting a health app, but this rather obvious difference should not concern us. My aim is not to argue that people can trust a health app *like they trust their best friend*; all I want to argue is that the concept of trust can help us to understand the *structure* of the relationship between (some) users and a health app. The fact remains that when people decide to use an app for a longer period of time, and decide to volunteer (user) data as well as access to their decisional sphere, they do so in the absence of guarantees that the app will (try to) help them and will not (try to) manipulate them. So in that basic sense, users do have to trust that the health app they use does not violate this ‘basic’ trust. Moreover, as we have seen in the previous paragraph, health app providers *know* that establishing trust is a necessary precondition for (financial) success, which means that health apps also actively try to convince users to place their trust in them. So even though trust comes in different shapes and forms, I maintain that using the concept of trust can help us understand the user-app relation better. When I discuss trust in the remainder of this dissertation, I refer to it in this wider sense, which encompasses both the trust people can have in, say, their best friend, and the trust people can have in a health app. To accentuate the difference between these different ‘shades’ of trust, I will sometimes speak of ‘trust(-like)’ relationships when I discuss user-app relations.

3.4. Taking Stock Before We Move On to Manipulation

Let me briefly summarize where we stand now, before I continue with an analysis of the concept of manipulation. The aim of this chapter is to develop a conceptual apparatus to evaluate the digital environments that are provided to users by popular for-profit health apps and the relationships between user and app that are cultivated in those digital environments. So far, I have discussed the concepts of autonomy, vulnerability, and trust. Simply put, the concept of autonomy provides us with an important perspective to see what is at stake; we want health apps to help people to be engaged with their health in ways that truly matters to *them* and that gives *them* the possibility to make decisions that are important to them. To become autonomous and to practice one’s autonomy, however, one is always partly dependent on other actors; autonomy is relational. This explains why health apps are so promising – health apps can be this other actor scaffolding the user’s autonomy – but it also, at the same time, explains why there is always a threat of others (people, institutions, companies) maneuvering themselves into our decisional sphere to influence us in ways that are *not* supportive of our autonomy. To help decide how popular for-profit health apps fit on this scale, the concepts of vulnerability and trust are helpful. Precisely because of autonomy’s relational nature, we are vulnerable (inherently and

situationally, in either a dispositional or occurrent manner) to others, affecting our autonomy adversely. Most popular for-profit health apps want their users to use their app for an extended period of time, which means that the apps try to build trust(-like) relationships with their users. Because of the *entwinement* of autonomy and vulnerability, trust can be helpful in overcoming people's reservations around seeking help or cooperating. But trust can also *betrayed*. In the case of health apps, there is a risk of health apps misusing the trust users place in an app by, for instance, exploiting people's health-related fears, insecurities, or desires to steer their behavior in ways that undermine rather than scaffold their autonomy. That is why in the last part of this chapter I turn to the concept of manipulation, which can help to 1) explain *how* digital environments can be designed and operate in ways that undermine autonomy by, for instance, exploiting trust(-like) relationships and vulnerabilities; and 2) evaluate why such conduct is problematic.

4. Manipulation

4.1. Introduction

Philosophical treatments of manipulation often start by emphasizing how difficult it is to define manipulation in a coherent manner. The philosophical difficulties surrounding manipulation may come as a surprise since the term manipulation is often used quite effortlessly in everyday conversation. The “commonness of manipulation in everyday life” is, however, precisely one of the reasons for it being so difficult to define (Noggle 1996: 43).

To illustrate the widespread use of the term manipulation, consider the wide range of practices we (can) call manipulation in everyday conversations. We say that we manipulate *objects* such as levers and juggling balls when we move them around at our will. We also say that we manipulate *persons* such as our neighbors, friends, and lovers. Lastly, we can say that *more abstract entities* such as institutions are manipulated (e.g., elections), or that we ourselves are manipulated by them (e.g., when the dating app Tinder sells premium ‘pay-to-win’ features to users by exploiting their fear of being put at a disadvantage vis-à-vis their competitors who are already using premium features to boost their chances).

Given the very wide range of phenomena that are so commonly said to involve manipulation, the task to come up with a definition of manipulation that covers all of our everyday intuitions about what manipulation is, is rather challenging. Such a definition would have to be very general and broad for it to be able to capture all the colloquial usages of the term manipulation (if indeed this is possible at all).

Before we move on and try to actually define manipulation, we should therefore ask ourselves what purpose we would like a definition to serve. Trying to define manipulation in such a way that it captures all (or at least as many as possible) of the phenomena we call manipulation in our everyday life is but one possible purpose of a definition.

Another option is to try to come up with what could be called a ‘technical definition.’ Such a technical definition does not try to account for all the possible colloquial uses of the term, but rather tries to serve a more specific purpose. I would like to suggest that the purpose a definition of manipulation could serve is to capture a particular *pro tanto* moral wrong that warrants special attention. Such a technical definition may not capture *all* of our colloquial uses of the term, but we need not be worried by this, as long as we clearly specify what the purpose of our definition is.

4.2. Moralized Versus Non-Moralized Conceptions of Manipulation

A good way to start a discussion on the conceptualization of manipulation is to address its moral status. In its colloquial use, the term is used to refer both to morally innocent and to morally wrong behaviors and actions. There is, for instance, nothing wrong with manipulating one’s juggling balls when juggling, but there is (we can assume for the sake of argument) something wrong with manipulating another person to do something that that person does not want to do and which conflicts with that person’s true informed interests.

In philosophical discourse, there is persistent disagreement concerning manipulation’s moral status. Some authors claim that manipulation is a morally neutral term that refers to a wide range of practices, none of which are *necessarily* morally wrong simply by virtue of being manipulation (e.g., Buss 2005, Wood 2014). (This of course still leaves open the possibility that instances that are labeled manipulation turn out to be morally wrong, but in that case there would have to be some other wrong-making feature of the situation that explains the wrongness of it.) Other authors use manipulation as a *moralized* term which always denotes a (*pro tanto*) moral wrong (e.g., Noggle 1996, Susser, Roessler & Nissenbaum 2019a, 2019b).

Notice that these disagreements on manipulation’s moral status can be (partly) explained by the different kinds of projects philosophers are engaged in. For example, if one’s project is to provide a conceptual analysis of the phenomenon of manipulation as it can be encountered in our everyday social and linguistic practices, then one is likely to end up with a conceptual understanding of manipulation that does not moralize the term. If, however, one’s project is to provide an ethical

analysis of a particular practice or technology, then a moralized conception of the term manipulation can aid one's ethical analysis since a moralized conception can provide a standard of evaluation.

For my purpose of ethically analyzing the commercial practices of health apps, it is most useful to develop a technical conception of manipulation which denotes a particular kind of (*pro tanto*) moral wrong. With a fully developed, moralized conception of manipulation in place, I can ethically evaluate commercial practices of health apps with reference to the moralized conception of manipulation.

4.3. 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 – in other words, manipulation *instrumentalizes* the manipulee. Noggle (1996: 44), for instance, writes that “[t]he 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 driver 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 clearly 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 that 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 “[t]hough 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, on 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 their 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. It should be emphasized at the outset that these are *cumulative* criteria: all need to be present to be able to speak 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 – at least not in the sense that the manipulator seeks to respect those. 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 about 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 a while, 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 (*pro tanto*) 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* characteristic of a person's psychology that can be exploited to steer that person's behavior in the (self-serving) directions that the manipulator desires. Put differently, manipulators are always looking for exploitable vulnerabilities in their targets, where 'vulnerability' should be understood in a broad manner to encompass any characteristics of a person that can be exploited by the manipulator. In the context of manipulation, information often equals power. The more one knows about the potential target, the better one will generally be able to identify

vulnerabilities.¹²⁷ In the absence of perfect information, manipulators will target presumed or inferred vulnerabilities. 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.¹²⁸ Manipulators always attempt to target vulnerabilities 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 rising 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.) All four characteristics of manipulation need to be present to speak of manipulation. When, say, two of the four characteristics are present it does not follow that a little bit of manipulation occurs.

127 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*.”

128 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.”

4.4. 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.

4.4.1. 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, 2019b). 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 the 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 ‘manipulation techniques,’ but rather starts with the ends and looks for *any* means that can help her achieve those ends. What typifies a manipulator is a mindset of seeing others as instruments that can be used to effectuate the desired outcomes. Consider, for instance, the app Headspace I discussed earlier. In their job postings they made clear that they are looking for persons to run large-scale behavioral experiments to identify *any* exploitable characteristics their users have may have and that can be targeted and used to personalize the user experience and promotional offers. In cases of manipulation, the means follow the ends, not the other way around.

4.4.2. 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, 2019b) 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. 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, 2019b) would *not* argue that our fictive 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 disposes 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 continuously identify how every individuals’ buttons can be pushed as effectively as possible in order to make them useful pawns in someone else’s scheme. That is the core intuition I would 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.

We could also ask what a ‘hidden influence’ means in the first place. What has to be hidden to be able to speak of a hidden influence? Must the very presence of the influence be hidden? Must the specific manipulation techniques used be hidden to the target? Or must the ends which the manipulator seeks to achieve be hidden to the manipulee? Because my account does not require the manipulative influence to be completely hidden, but merely states that manipulators attempt to keep the manipulation hidden, it does not have to grapple with the difficult question of what precisely constitutes ‘hiddenness’ in the context of manipulation.

4.5. Manipulation as a Success Concept: Individual Acts of Manipulation versus Manipulative Environments

Manipulation is a success concept. As Susser, Roessler, and Nissenbaum (2019a: 27) explain, “the claim that someone has been manipulated refers not only to the strategies employed by the influencer but also to the effects of those strategies on the influenced.” To say that someone has been manipulated is to say that the intended manipulation actually succeeded. Manipulation as a success concept is easy to understand and apply in simple interpersonal cases of manipulation where A manipulates B with action Y to do X. Such cases are simple because we can focus on a neatly circumscribed isolated interaction between two actors. There is a clear manipulative act we can focus on to establish whether manipulation actually occurred.

In digital environments like health apps, focusing on manipulative *acts* is not a very productive approach. As users use a health app for an extended period of time, their use of the app is made up of many separate interactions. When asking whether a health app is manipulating its users, one could look at all the separate interactions

between the user and app and check for all of those separate interactions whether there is a manipulative intent and whether manipulation has actually occurred. However, performing such a manipulation check is not only highly impractical – for popular apps with millions of users, one would have to look at hundreds of millions of interactions – but also rather uninformative. In the freemium app economy (see chapter 1) digital environments are built to create relationships with users *over time*. It can be difficult to point to separate, isolated manipulative actions in order to explain the *overall* manipulative potential of such digital environments and the relationships they seek to create with their users. To get this more *structural dynamic* of the interactions between app and user into focus, it is more helpful to concentrate on structural design and operations of those digital environments and ask how they – over time – seek to influence their users. Or, put differently, it is helpful to focus *less* on individual acts of manipulation and more on the intentional development and deployment of manipulative strategies by and through digital environments. I suggest we follow Susser, Roessler and Nissenbaum (2019a: 27) who introduce the concept of manipulative practices: “we focus [...] on the concept of *manipulative practices* – strategies that a reasonable person should expect to result in manipulation – and not on the success concept of manipulation, in toto.”

Hence, we can call a digital environment (such as a popular for-profit health app) manipulative if it is designed and operated in such a manner that we can be almost certain that at least *some* users will be manipulated. Reversely, it follows that even if we can provide counterexamples of some users that were *not* manipulated by interactions happening within the digital environment, we can still claim that that digital environment as such is manipulative. What matters is whether a digital environment is designed and operated in a manner that betrays a manipulative intent and employs strategies that can be considered manipulative, not whether *every single user* will *always* be manipulated when interacting with or within the digital environment.

Also notice that the focus on digital environments has important implications for the question of intentionality: it is not necessary to establish whether there was an explicit intention on the part of the health app provider to manipulate every single user. All I need to establish is that the environment is intentionally designed in such a manner that the health app provider can be reasonably sure that at least some users will be manipulated by it.

5. Manipulation and Autonomy (, and Vulnerability and Trust)

The relation between manipulation and autonomy is usually thought of as being very straightforward. Let us start with an isolated manipulative interaction between a manipulator and a manipulee. When a person is successfully manipulated, that person is not (fully) autonomous. It is easy to see why, especially when one adopts (as I do) a procedural conception of autonomy, which requires the process by which persons make up their minds about their preferences, values, and ultimately decisions to meet certain requirements. When a manipulator targets vulnerabilities or weaknesses to exploit them with the aim of getting the manipulee to do something the manipulator is reasonably sure the manipulee does not want to do, both the independence and authenticity of the target's decision-making are not respected. So if we focus on individual interactions between a manipulator and a manipulee, we can see how a particular decision that follows from the manipulative interaction will not be (fully) autonomous.

I would like to point to another important relation between autonomy and manipulation. In the first part of this chapter I drew on relational accounts of autonomy to theorize the social conditions for (procedural) autonomy. In the second part of this chapter, I introduced a distinction (following Susser, Roessler, and Nissenbaum 2019a) between manipulative individual acts and manipulative practices or environments. Combining these two insights, it follows that manipulative environments can not only threaten autonomous decision-making right at the moment when a decision is made, but that more generally speaking manipulative environments can also undermine the social conditions necessary for developing and practicing autonomy.

For example, consider the different modes of recognition that help persons develop an understanding of themselves as competent, autonomous agents. One could ask whether (potentially) manipulative commercial health app environments that are designed to optimize conversion are the type of environment that helps people recognize themselves as competent autonomous agents. Does such an environment stimulate a relation to one's own body and health that, from the perspective of autonomy, is valuable? Also consider the fact that theories of relational autonomy emphasize the importance of relationships for autonomous agency. I argued earlier that as health apps are used over time, a (commercial) relationship between the user and the app develops over time as well. One could thus also ask how the development of this (commercial) relationship impacts the users' possibilities for practicing autonomy.

We can gain an even better understanding of the relation between manipulation and autonomy – especially in the context of health apps – by incorporating the concepts of vulnerability and trust (as discussed above) into the analysis. As I have argued, manipulation involves the targeting of vulnerabilities, so a more refined understanding of vulnerabilities also helps to analyze the manipulative potential of health apps in a more refined manner. If we take autonomy’s relational nature seriously, as I have argued we should, we should also acknowledge how the development and practicing of autonomy necessarily comes with vulnerabilities. So precisely because autonomy and vulnerability are entwined, the conditions for autonomy also, at the same time, provide a necessary (but not sufficient) condition for manipulation, namely the existence of vulnerabilities that can be targeted and exploited by manipulators. It is also from this perspective that we should analyze the role trust can play in manipulative digital environments. Trust(-like) relationships provide a *source* of vulnerabilities that can be exploited by manipulators, especially when they revolve around a value that is important to people, such as health.

6. Power

Before proceeding to the next chapter, I want to briefly discuss the concept of power. As I have argued in earlier chapters, health app providers should be considered choice architects. By designing digital choice environments, health app providers can shape and/or steer the behavior of the users of those environments, which puts them in a position of power. This position should be properly understood, for it is precisely this position of power which can – but does not have to – enable the introduction of manipulative practices. My aim in this section is not to provide a full-fledged theory of power in the digital society. I do, however, want to offer some conceptual refinements in order to better understand how power can operate and be exercised in digital environments.

So how should power be understood? In its most basic, intuitive sense, the exercise of power is often associated with the ability to *interfere* with other people’s decisions and behavior. Dahl (1957: 202-203), for instance, writes that “my intuitive idea of power, then, is something like this: A has power over B to the extent that he can get B to do something that B would not do otherwise.” Although this intuitive and commonsensical way of understanding power is certainly not wrong, it does require more nuance. Lukes (1974: 11-15) famously argued that power has three dimensions and called the ‘intuitive’ idea of power as formulated by Dahl “one-dimensional.” Power can be exercised in many intricate ways which go beyond the simple, direct interference with someone else’s actions. For example, Lukes argued that power is

not just about observable behavior in situations where there is overt conflict over particular decisions and where one party has the power to compel a particular outcome. Following Bachrach & Baratz (1962, 1963, 1970) Lukes discusses how there is also power in agenda-setting and controlling what counts as a legitimate issue or question in the first place. Moreover, power can also be exercised by way of shaping/changing the preferences and desires of others, causing them to actually desire that which is in the interest of the actor in the position of power. We do not have to accept Lukes' particular framework to conclude that power can function in a variety of intricate ways which go beyond direct interference.

One perspective that seems especially relevant in the context of adaptive digital (health) environments is the republican perspective on power (see Pettit 1997, 2012, 2018). As the literature on the republican conception of power teaches us, power should not only be associated with actual interference. The mere fact that an actor is in a position where she *can* interfere¹²⁹ with someone or something already gives her power over others, irrespective of whether she actually chooses to interfere. So to understand the power dynamics at play in a digital (health) environment, we should not just focus on instances of actual interference, but also on the ways in which the (affordances of) digital environments grant their owners the ability to interfere with the users of the digital environment. More concretely, one can think of the collection of data about one's users to learn about their personal characteristics or circumstances. Such information gives the owner of the digital environment power over its users because such data can be used to target users with (personalized) messages or nudges (Helberger 2016: 155). As we have already seen in the sections on manipulation, acquiring information about one's targets also renders those targets more vulnerable to manipulation – the more one knows about the target, the higher the chance of finding out which buttons to push and which strings to pull.

Even though my aim has not been to offer a full-fledged theory of power in the digital society, I do hope that these brief elaborations on the concept of power help the reader to see and understand how questions and issues of power operate in the background of the ethical and legal analyses I offer in this dissertation. Choice architects are powerful actors in the digital society, who can exercise their power not only by actually interfering with decisions to *compel* particular outcomes, but also by *shaping* preferences and desires, or by manipulating others. Moreover, their

129 'Interfere' should be understood in a wide sense, in line with Lukes' (1974) argument. So interference does not just refer to *direct* interference to *compel* particular outcomes, but can also refer to more subtle or even manipulative practices aimed at securing particular outcomes or cooperation.

power consists not only in their actual attempts to influence people, but also, more importantly, in their very position as choice architects who *can* – if they so desire – interfere with users of their digital environments.

Keeping these remarks on power in mind, it is time to move on to the next chapter, in which I will build on the conceptual work done in this chapter to provide an ethical evaluation of health apps. I will explore the tension between user empowerment and user manipulation in for-profit health apps.

Chapter 4

The Vague But Important Line between Supportive and Manipulative Digital Health Environments

1. Introduction

In the previous chapters I have discussed the design of the digital health environments provided by health apps, the commercial practices they engage in, the health discourse one can find in popular for-profit health apps, as well as the conceptual framework I will use. The key question to be addressed in this chapter is where to draw the line – ethically – between legitimate and illegitimate commercial health app practices. Or, put differently, where does legitimate ‘business as usual’ end and where does illegitimate manipulative commercial activity begin in the context of for-profit health apps?

Autonomy will serve as the guiding principle in this chapter; ultimately, we should want and strive for health apps that enable their users to be preoccupied with and work on their health in an autonomous fashion. The concepts of trust, vulnerability, and manipulation will be used to explore under what conditions health apps do, or do not, promote users’ ability to practice autonomous health-related behavior.

Before I proceed to my actual analysis, I would like to use this introduction to explain what *kind* of analysis I will be providing. Most for-profit health apps are freemium apps, which means that their monetization strategy is based on recruiting users which they then try to retain and engage for longer periods of time (see chapter 1). For-profit health apps thus try to build ongoing *relationships* with their users. These relationships are situated within – and shaped by – the *digital environments* that are provided by health apps. So to understand and be

able to scrutinize the ways in which users interact with, and are influenced by, for-profit health apps, we need to focus on 1) the structural design features of these digital health app environments, within and through which 2) dynamic, ongoing relationships between users and apps take shape.

By adopting a focus on structural design features of digital environments and the resulting dynamic, ongoing relationships between those environments and their users, I want to move away from the type of analysis that only focuses on separate, individual interactions between an app and a user, asking, for example, whether *this particular individual intervention* by an app is ethically problematic or not. While the analysis of separate, individual interactions can be valuable, I argue that such analyses ignore the larger context within which these interactions take place. Many of the digital environments we navigate – including for-profit health apps – are *not* designed with (just) separate, isolated interactions in mind; rather, they operate at a higher level, constantly analyzing (and experimenting) how the *entire environment* of the app or service can be optimized when taking account of the fact that people will use (or can be made to use) an app or service for an extended period of time. This reality should be reflected in normative analyses of these apps and services (see Frischmann & Selinger 2018 for a similar argument).

It follows that if we want to ask under what conditions for-profit health apps can support the autonomy of their users when dealing with their health, we should not only focus on the question of whether a particular individual interaction between user and app respects the user's autonomy *at the moment of the interaction*. It is equally important to ask whether the way in which the entire digital environment is set up, and the way in which the health app – through that environment – tries to shape and structure the relationship with the user *over time*, shows sufficient respect for that user's autonomy (Frischmann & Selinger 2018; Susser, Roessler & Nissenbaum 2019a, 2019b; Zuboff 2019). This is why I argue in favor of a conception of autonomy that not only takes the standard notions of authenticity and competency on board, but also *relationality*. The ways in which our relational embedding makes us understand ourselves, and enables and constrains our practicing of autonomy, matters.

I develop my analysis in this chapter in the following way. I start by providing a brief schematic sketch of the different ways in which the autonomy of health apps users can be affected by the influences exerted by health apps. Put simply, health apps can both enhance and undermine their users' autonomy in different ways. As will become clear throughout this chapter (and throughout the entire dissertation for that matter), there is a *vague line* between helpful and manipulative digital

(health) environments. Now, precisely because we are dealing with a vague line, it will be incredibly difficult, if not impossible, to focus on the line itself and explain to the reader what the line looks like and where exactly it will be located in real-life cases. In reality, moreover, we are often dealing with a continuum. The digital environments provided by health apps are rarely either purely helpful or purely manipulative; specific functionalities can be manipulative in nature, or specific functionalities can have a manipulative influence on a subset of users. So instead of focusing on the line *itself*, I will focus on both sides of the line. This allows me to describe in slightly more general terms what it takes for digital (health) environments – or parts thereof – to fall on either side of the line, without having to describe the line itself.

After I have briefly indicated how the question of autonomy presents itself in the contexts of health apps, I provide an overview of the conditions that need to be place for health apps to serve as genuine ‘scaffolding’ (Heath & Anderson 2010) for their users’ autonomy. Here I also discuss how trust(-like) relationships between users and health apps, as well as a respectful attention for vulnerabilities¹³⁰ of health app users, can be important contributors to the potential of health apps to scaffold users’ autonomous dealings with their health. By focusing on the role trust and vulnerabilities can play in the user-app relationship, we can, however, also start to see how trust and vulnerabilities can be *misused* by health apps to promote ends that can clash with or even undermine the ends of users. In a next step, I will discuss under what conditions for-profit health apps can turn into manipulative digital environments. As much as trust and careful attention for vulnerabilities can be sources of *helpful* relationship-building between users and health apps, they can also easily turn into sources of manipulative potential that can undermine, rather than enhance, the autonomy of health app users.

2. The Question of Autonomy in the Context of Health Apps

When a person’s behavior is influenced by an external source, it makes intuitive sense to ask how that external influence impacts this person’s (possibilities for practicing) autonomy. Since health apps exert – or at the very least attempt to

¹³⁰ I use the term ‘vulnerability’ in the wide ethical sense developed in chapter 3, basically referring to a wide range of ‘exploitable characteristics’ of persons. In chapter 5, I will discuss the *legal* concept of vulnerability in the context of the UCPD. In the legal context, ‘vulnerability’ is typically understood in a much narrower sense. In chapter 5 I will argue that the ethical concept of vulnerability can and should inform our interpretation of the (much narrower) legal concept of vulnerability.

exert – an influence on their users’ behavior, they automatically give rise to ‘the question of user autonomy.’ It bears emphasizing that even the most basic health apps exert an influence on the user’s behavior, however minor, non-persuasive, and even indirect the influence might be. Think, for instance, of a calorie counting app that only offers the functionality of inserting caloric intake information, and then showing this information to the user in an organized manner without making any recommendations or performing any analyses. The simple act of showing self-reported data in an organized manner can exert an influence on a user’s behavior; it might give her a better understanding of her eating patterns and might prompt her to adjust her eating behavior. Contrary to the hypothetical, very simple app I just described, most of the more popular health apps I discuss try to *actively* shape and steer the user’s behavior by, for instance, making recommendations, personalizing content, implementing persuasive UX designs, and so on.

The formulation of a potential “impact” on the autonomy of a user of a health app needs to be refined. An external influence on a person’s behavior can, roughly speaking, preserve or enhance, or threaten to undermine that person’s autonomy. Before spending the last part of this chapter speaking about the potential threats to autonomy, I also want to discuss how health apps can *preserve* or *enhance* the (possibilities for practicing) autonomy of users. Doing so does not only give us a more complete story, but also – and more importantly – provides us with a positive image to which we can refer. It helps us to better understand both sides of the ‘vague line’ I spoke of earlier.

So how can health apps preserve or enhance the autonomy of their users? Let us start with the most straightforward situation where a person has, in a sufficiently autonomous fashion, decided that she wants to pursue goal X. She then uses a health app as a mere instrument to help her pursue goal X more successfully or efficiently. If the health app does indeed function as such a useful instrument, and if, moreover, the use of the health app does not introduce any new influences or side effects that may threaten the user’s personal autonomy, then the health app does indeed help the user to decide and behave in an autonomous fashion. An example could be a person that wants to keep track of her diet for health reasons. Until now, she has been doing this with a custom spreadsheet on her computer, which helps her to calculate manually how much calories, carbs, sugars, etc. she eats on a given day. This process is relatively time-consuming and cumbersome. To ease this process, she decides to use a calorie counter, which automatically does a lot of the calculations for the user and also comes with a large database of food

products and produce with preprogrammed information. As a result, this user's autonomous intention to keep track of her health can be executed more efficiently and accurately.

A slightly more complicated situation is one where (proactive) suggestions of a health app evoke desires which then inform a person's intentions or projects. Here the health app does not function as a mere tool to help execute or follow through with a settled, sufficiently autonomous intention, but rather suggests it. This does not necessarily undermine the user's autonomy of course, since desires and the resulting intentions do not need to spontaneously originate in a person to be autonomous. In this case, what is crucial is the process through which the desire and resulting intentions are adopted by the user.¹³¹ Depending on the precise conception of autonomy one accepts, one would give a different answer to the question of what conditions need to be fulfilled for intentions to count as sufficiently autonomous. Roughly speaking, however, most people would agree that a person needs to engage in some critical reflection to decide what to do with the suggestion by the app (see also chapter 3). If, after critical reflection, the user embraces, identifies with, or does not resist the newly found desire and resulting intentions sparked by the suggestion, we can declare the new intention sufficiently autonomous. It could even be said that a health app can make a person *more* autonomous through its suggestions, if the suggestions make a person consider courses of actions, goals, or projects the person had not considered before, but, after critical reflection, regards as valuable new additions to her life. Think, for instance, of a person who looks for a new diet as an answer to some health concern. A health app might suggest an entirely new cuisine to her that she did not know about, but which offers delicious food that fits perfectly with her dietary restrictions. The proactive suggestion might spark her joy for cooking (providing her with an exciting new hobby) while also helping her achieve her health or lifestyle goals in a pleasant manner.

The last type of situation I want to discuss is one where a user uses a health app to *restrict* some of her own behaviors. Intuitively, using an app in such a way might be construed as restricting one's own autonomy, but that conclusion would be too quick. Consider the phenomenon of weakness of the will. A person can have a serious, autonomous intention to lose weight while also experiencing moments where the weakness of the will manifests itself, for instance when presented with snacks. In such a situation a more immediately felt desire – wanting the unhealthy

131 In chapter 3 I already discussed Christman (1991) who argues that the process through which a person comes to adopt/accept a desire, rather than the source of the desire *per se*, matters most to autonomy.

snacks – can compete successfully with the desire that the person finds (all things considered) more important, which is losing weight. Or put differently, in Frankfurtian terms, a person may often find herself having a first-order desire to eat an unhealthy snack, while also having a stable second-order desire to lose weight (Frankfurt 1971). When giving in to the immediate desire to snack, such a person might experience herself as lacking autonomy, even though she can fulfill an immediate desire she has. In order to pursue her autonomous intention to lose weight, she might try to organize her surroundings in such a way that it becomes easier to resist desires for snacks – she can, put differently, try to ‘scaffold’ her autonomy through a slight reorganization of her environment (Heath & Anderson 2010). For example, she might install an app on her phone that not only tracks her food intake and weight, but also sends her notifications around times when she usually feels the urge to snack. Although this person is, in a way, *restricting* her own freedom to act by using the app in this way, there is still a case to be made that she is enhancing her autonomy in doing so. Heath and Anderson (2010: 235) use the term “extended will” to explain how we can try to offload some of the ways in which our will manifests itself to the environment or to technology. In this case, the person with the truly autonomous intention to lose weight by eating fewer snacks can ‘extend’ her will by autonomously imposing self-made restrictions on herself in anticipation of moments where she experiences the weakness of the will she does not autonomously *want* to experience. If the decision to extend our will in such a manner is sufficiently autonomous, it can be concluded that the self-imposed restrictions should be considered to be autonomy preserving – or even enhancing – as well.

3. The Helpful Side of the Vague Line: How Autonomy, Trust, and Attention for Vulnerabilities Ideally Come Together

It should be clear then that health apps can respect or even enhance the autonomy of their users. So what conditions should – ideally – be in place for a health app to preserve or enhance its users’ autonomy? To start with, a very simple observation: an app should seek to support its users to make health-related decisions that are in accordance with their autonomous intentions. It is of course possible for an app to *accidentally* help some of its users to act upon their autonomous intentions. But if an app wants to claim that it explicitly *aims* to support or enhance people’s

autonomous health behavior, it should be designed and organized in a manner that prioritizes the identification and tracking of people's authentic, autonomous health-related intentions.¹³²

An autonomy preserving or enhancing health app should, moreover, also be transparent (enough) about the influences it exerts on its users. As we have seen above, a health app can restrict or steer a person's will and still be autonomy-preserving or -enhancing, but this can only be the case if the person being influenced has been able to *understand* and *autonomously endorse* current or future restrictions and steering influences. So the transparency should not necessarily be absolute, but there should be enough transparency for users to make a truly informed decision. We can, for instance, imagine an app that offers behavior-restricting interventions. Such an app can be perfectly transparent about the fact that some of its behavior-restricting interventions will function in ways that are themselves not entirely transparent to the user. As long as this information about future, deliberate opaqueness is communicated in a transparent manner, the user has enough information at her disposal to make an autonomous choice.

Let me now move on to ways in which attention for the vulnerabilities of users as well as the building of trusting relationships with users can serve as important conditions for offering autonomy-enhancing digital health environments.

3.1. Vulnerabilities

Health apps can identify and target various vulnerabilities in ways that can help users to deal with their health while respecting and supporting those users' autonomy. In the health app context, we can differentiate between two types or levels of vulnerabilities: 1) vulnerabilities pertaining to the user's health, and 2) vulnerabilities pertaining to the user's cognitive and affective biases. The first type of vulnerabilities refers to the ways in which particular persons are vulnerable to have their physical or mental health or wellness affected by any conceivable influence. For example, some people may have a weak knee and risk knee injuries if they do not carefully plan and observe their exercise regimes, while other people might risk sliding (back) into a depression if they do not adhere to particular routines. The second type of vulnerabilities refers to the various cognitive and affective biases people are susceptible to and which can be targeted by health apps (and digital choice architectures more generally) to influence behavior. Although

132 Notice that I am not arguing that health apps can only engage in acceptable commercial practices if they *only* serve the interests of the user while neglecting their own interests. My aim here is to sketch how health apps can – ideally – support or enhance the autonomy of the user, not to prescribe a standard that every app should meet in order to be seen as ethically acceptable.

many of the known cognitive and affective biases are widely shared among people, we can still expect some variations on an individual level. For example, some people might respond more strongly to ‘social’ biases that emphasize what friends or ‘people like you’ do, while other people might be more susceptible to ‘sticky defaults’ (Thaler & Sunstein 2008: 83-87).

Now, although the phrasing ‘targeting vulnerabilities’ might sound rather negative or even exploitative, identifying (or inferring) and targeting vulnerabilities can in fact *help* people practice their autonomy vis-à-vis their health. Let us start by considering the first type of vulnerabilities, which pertain to the health of the user. Put very simply, the more a health app knows about a person’s specific vulnerabilities, the better the health app should (in principle, but not necessarily) be able to offer personalized suggestions to this person. Knowledge about the specific circumstances of a user creates the possibility to offer solutions that cater to those specific circumstances. Consider a health app that allows users to track running sessions and which also offers users a feature to create training schedules. If the app also has information about injury histories and other relevant physiological data of its users, it can personalize the recommendations for training schedules to minimize the risk of injuries. So by knowing more about our specific, personal health-related vulnerabilities, health apps can help us to better address those vulnerabilities.¹³³

The second type of vulnerabilities can, broadly speaking, be identified and targeted in order to help people change their health and wellness-related behaviors and lifestyles. As the literature on behavioral economics has taught us, behavior and lifestyle change can be difficult to achieve, but when our cognitive and affective biases are leveraged in the correct manner, we can be helped by choice architects to achieve our desired outcomes. It follows that our cognitive and affective biases (i.e., the second type of vulnerabilities) can be leveraged to help us change our *health*-related behavior and lifestyle. Because the behavioral sciences have described a wide range of cognitive and affective biases,¹³⁴ we can think of a similarly wide range of strategies for targeting known or inferred cognitive and affective

133 Also consider our institutionalized medical practices, which are also founded on the sharing of information to allow for better treatments. Doctors typically question and examine a patient to get a clear picture of what is going on with this *particular* patient and what treatment would benefit this *particular* patient.

134 See Thaler & Sunstein (2008) for an expansive overview of such biases. They discuss, for instance, sticky defaults (pp. 83-87), the above average effect (p. 32), the effectiveness of framing (pp. 36-37), the mere measurement effect (p. 70), people’s sensitivity to group norms and susceptibility to peer pressure (p. 57), loss aversion (pp. 33-34), inertia (p. 35), the planning fallacy (p. 7), and so on.

vulnerabilities aimed at supporting people's autonomous health behavior through health apps. Health apps can, for instance, personalize their content, visuals, tone of voice, or timing of notifications based on what the app thinks will lead to more motivational or persuasive interactions. Because phenomena such as (temporary) lapses of motivation to achieve desired outcomes or (temporary) weakness of the will manifest themselves differently for different people, an intimate knowledge of the psychological and behavioral tendencies of individuals can in principle be used to help people achieve what they, all things considered, value most. This line of thinking also accurately captures the often-touted potential for self-empowerment through the proliferation and use of health apps. We, of course, often encounter this promise of self-empowerment in health apps themselves (see chapter 2) in order to advertise their services. That these promises of personalization and self-empowerment are not unfounded is shown by a growing body of research which looks into the various ways in which personalization of different features of online health services and health apps can improve health outcomes by influencing health behavior and health decisions (see, e.g., Kroeze *et al.* 2006, Noar *et al.* 2007, Krebs *et al.* 2010, Lustria *et al.* 2013, Kaptein *et al.* 2015, Bol *et al.* 2019).

In sum, health apps can leverage knowledge of health and wellness-related vulnerabilities as well as cognitive and affective vulnerabilities to help people deal with their health in a more efficient, effective, and – under the right circumstances – autonomous manner. To fully understand how influences exerted by health apps can steer people's behavior in an autonomy-respecting manner, I now turn to the role trust can play in user-app relationships.

3.2. Trust

Trust can play an important role in understanding and explaining how health apps can help people work on their health in autonomy-preserving or even autonomy-enhancing ways. In the previous two sections on autonomy and vulnerabilities, most of the potential positive effects we hope health apps can have are premised on the possibility of building ongoing relationships between apps and their users. As people use apps for extended periods of time, the apps can 'get to know the users,' meaning the app can – ideally – learn about vulnerabilities (of either type) and about the considered autonomous intentions of their users. It should be clear that such fruitful relationships between health apps and users can only develop when people *trust* – or at the very least do not actively *distrust* – the health apps they are using. Both mental and physical health are usually considered to be rather sensitive or intimate topics. Letting an anonymous digital service that you use through your smartphone in on your health-related vulnerabilities does not just happen automatically. You need to *trust* this service with (potentially) sensitive or intimate

information in at least two senses. First, you need to trust the health app to keep your information safe; you must trust the app to take adequate measures to protect your health-related information against (digital) theft and misuse. Moreover, you must trust the health app not to voluntarily disclose your health-related information to other actors you do not condone of. Second, as you let health apps within your *decisional sphere* to allow them to exert influences on your (health-related) behavior, you also need to trust those apps to use their insights into your vulnerabilities in a proper manner. In sum, trust serves as a necessary background condition which must be met to facilitate the positive effects that health apps can achieve.

Users do not only need to trust health apps for them to be able to perform a useful role at all; there are also use cases where health apps may be uniquely positioned to offer health solutions which users trust. In recent years, the quality of artificial conversational agents (or, in simpler terms, chatbots) seems to have seen a significant increase. Chatbots have become increasingly good at mimicking ‘real’ conversations and offering meaningful responses to questions posed by people interacting with the chatbot. Against the background of these general improvements in artificial conversational agents, there has been an increased interest in the development and deployment of mental health and digital counseling chatbots (see, e.g., Cameron *et al.* 2017, Oh *et al.* 2017, Lee *et al.* 2019). Many of these chatbots are founded on the idea that people may find it easier to share and discuss intimate mental health questions with AI rather than with a *real* person; people do not feel shame towards AI, potentially making them more inclined to seek help or disclose intimate information (Hasler, Tuchman & Friedman 2013).

As I discussed in chapter 3, “unlike some other values, trust is a psychological state that represents the trusted person or object as being trustworthy, and this may or may not actually be the case” (Nickel 2015: 552). Trust can be engineered or induced independent of the actual trustworthiness of the trusted actor or object. Generally speaking, there is a wide range of factors that can contribute to the trust that people have in health apps. Source credibility, for instance, can be important to people: they might be more likely to trust Apple’s Health app because they trust Apple as a company. Other people might be more likely to trust certain apps based on their friends’ experiences. For health apps to be helpful health and wellness companions, it can of course be beneficial if they are *designed* specifically with the building of trust in mind. The chatbots discussed in the previous paragraph only work well when they manage to instill a sense of trust in their users. Some of the cognitive and affective biases that have been discussed in several places in this dissertation can also be leveraged to induce trust. However, for health apps to be *truly* helpful

in a respectful, autonomy-supporting or autonomy-enhancing manner, there needs to be an actual corresponding trustworthiness that also *justifies* the trust they seek to induce. For example, a chatbot can be designed to induce a great sense of trust in its users, making it more likely that the users share intimate information to receive mental health support. But if the chatbot then uses the information about its users' vulnerabilities for marketing purposes, the cleverly engineered trust has been betrayed by the untrustworthy behavior.¹³⁵

3.3. Conclusion

When we combine the insights from the previous sections on autonomy, vulnerability, and trust, the following picture emerges: *under the right conditions* health apps can serve as especially helpful tools when they identify and target vulnerabilities and when they manage to build trust(-like) relationships with their users. So what are those 'right conditions'? The gist of the answer is deceptively simple: targeting vulnerabilities and the attempt to build trust relationships can, generally speaking, be considered unproblematic or even helpful in a context where it is the health app's one and only goal to help the user deal with her health in an autonomous fashion. Put differently, the design and operation of the health app must be geared towards the genuine attempt to understand what the users *really* need from a health perspective, and how corresponding features of the app can be best designed and used to respond to those real health needs. The interest of the app must be to serve the interests of the user. That is not to say that as long as the only interest of an app is to serve the *true, authentic* interests of the user, all the resulting interactions are by definition unproblematic. Between the high-level intention to only serve the true, authentic interests of the user and the actual execution of that high-level intention a lot has to happen and, as a result, a lot can go wrong.

I have called this answer 'deceptively simple' rather than just 'simple' because in the real world there are few (if any) popular health apps that are designed and operated *purely and only* for the sake of serving the interests of their users. Most, if not all, of the popular health apps that are used by large audiences are *for-profit* services. And so in the messier real world, interests and incentives are not as clearly aligned and organized as in the case of the hypothetical health app that meets the right conditions and *only* seeks to serve the *real, authentic* interests of the user. As soon as these 'right conditions' are not fully met, health apps that identify and target vulnerabilities and that try to build trust relationships can quickly turn from helpful into problematic services.

135 Here I am assuming, for the sake of argument, that the chatbot service did not (clearly) disclose to users that they are also using user data for marketing and targeting purposes.

There turns out to be a fine line between the helpful use of vulnerabilities and (the development of) trust and the *misuse* of vulnerabilities and (the development of) trust.

In the remainder of this chapter, it is precisely this other side of the vague line between use and misuse of vulnerabilities and trust that I am interested in. As we will see, the potential misuse of vulnerabilities and trust prove to be a fertile ground for manipulative practices. So in what follows, I will build on my concept of manipulation (developed in chapter 3) to explore how digital health environments that also seek to serve their own interests can cross this vague line and become manipulative. My aim is, above all, to sketch the general conditions that make a health app manipulative. I do not provide an exhaustive overview of the extent to which all popular for-profit health apps do, in fact, meet these general conditions for being manipulative digital health environments.

4. The Other Side of the Vague Line: Health Apps as Manipulative Digital Environments

Manipulation offers a useful perspective to explore the fine line between, on the one hand, helpful support by health apps, and, on the other hand, problematic interferences with our autonomy (in the wide sense as I developed the concept in chapter 3). To help us understand the manipulative potential of popular for-profit health apps, we need to briefly revisit chapter 1. In chapter 1 I explained how most of the popular for-profit health apps adopt a freemium business model. The freemium business model comes with a set of relatively fixed incentives and corresponding design choices. Freemium (health) apps typically attempt to build on-going relationships with their users (user retention) and will try to keep their users *engaged* with the app. The freemium philosophy is built on the idea that when users keep coming back to the app and remain engaged with the app, they will, in the end, be more willing to spend time and money on the app. The initial free offering of the service is thus used as a kind of 'bait' to lure new users. Once the users have installed the app, the real work of trying to optimize the retention, engagement, and ultimately conversion of the users begins. One of the core arguments of this section is that the app designs and commercial practices that flow from the freemium business model are, in the context of *health* apps, precisely the types of design choices and business practices that risk making the health app in question a manipulative digital health environment.

For ease of the analysis, let me repeat my conception of manipulation. 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.)

Let me emphasize again that we can only speak of manipulation (or of a manipulative practice) when all four elements of manipulation are present. So when, say, three of the four elements are present, it does *not* follow that *some* manipulation is going on. This should be kept in mind because in what follows I cannot avoid discussing the four elements of manipulation separately (although I will of course also relate the four elements to each other to paint the bigger picture). When I discuss the different elements of manipulation, I provide examples of how those separate elements are sometimes present in health apps. My aim is not to single out particular apps, but to *illustrate* the more general argument I am developing and which concerns the conditions which make a health app either a helpful or a manipulative tool.

4.1. Conditions for Manipulation in Health Apps

To provide the reader with some guidance, let me start by briefly summarizing the argument I am going to develop in this section. For-profit health apps operate in a commercial context where it can be difficult to disentangle the app provider's interests and the users' interests. This is exemplified by the incentives that flow from the freemium business model: in order to optimize the metrics of user retention, user engagement, and conversion, health apps have a strong incentive to track users and experiment with various methods to keep users coming back to, and engaging with, the app. In such a situation it can become increasingly unclear whether users interact with a health app because they really *want* or *need* to, or because the health apps cleverly leveraged vulnerabilities and the users' trust in the app for its own gain. It can be difficult for users to know what is what in such health apps, because most popular for-profit health apps are deliberately designed to look and feel like unobtrusive self-help instruments which direct attention away from their often advanced retention, engagement, and conversion optimization efforts. The result is a digital health environment where monetization of the userbase can take place through manipulative means.

It is tempting to conclude that I am essentially arguing that commercial activities such as advertising and sales promotion are manipulative and thus illegitimate and that, by logical extension, such commercial activities are also manipulative and

thus illegitimate *when they occur in health apps*. The question whether commercial activities such as advertising and sales promotion are mostly – or even always – manipulative is not one I seek to answer. They certainly can be, and maybe most of them even are manipulative to a certain degree.¹³⁶ But my argument is not focused on the commercial exploitation of health *per se*. My argument must be seen, and placed, within the *technological* context of adaptive, personalized digital environments with targeting capabilities. It is one thing to design a billboard or a television ad which, in a very broad and general sense, ‘targets’ an ‘exploitable characteristic’ (say, the desire to be more successful). It is another thing to build a digital environment which can sniff out exploitable characteristics of *individual users* and *related to something as essential as one’s health* and to use personalized targeting capabilities to exploit those privileged insights into the psychology of *individual users*. That is not to say that all popular for-profit health apps do in fact operate in such a manner, but it is precisely this *potential* of such apps that should be scrutinized; and that is what I will do in the remainder of this chapter.

4.1.1. *Helping the Bottom Line by Helping the User?*

Let me now flesh out the different parts of my argument in more detail. I want to start by focusing on the interests that are at play in the freemium health app economy. In the first part of this chapter it became clear that health apps *can* be of great help to users, but only when they meet some important conditions. One of the most important conditions is that the health apps collect data and design their app in order to (make a genuine attempt to) serve the real, authentic health-related interests of their users. In a commercial context, where health apps both seek to help users *and* want to make a profit from doing so, we should ask how the aim to profit from health advice can be reconciled with a respectful treatment of the users’ decisional health sphere.

Let me begin by stating the obvious: there can be no doubt that for-profit health apps seek to further their own interests. The health apps that are the focus of this dissertation are for-profit; their aim is – by definition – to generate profits. Besides being for-profit services, many of the popular health apps are funded by venture capital firms that demand steady and quick return on investment.¹³⁷ In order to

136 Some have argued that this is the case: see for instance Santilli 1983 and Crisp 1987.

137 For example, Headspace was funded through four funding rounds, raising \$75 million (<https://www.crunchbase.com/organization/headspace#section-investors>, last accessed 22 September 2020). MyFitnessPal also received funding from venture capital firms (<https://www.crunchbase.com/organization/myfitnesspal#section-investors>, last accessed 22 September 2020) and was later acquired by Under Armour for \$475 million (Olson 2015). Fitbit also saw four funding rounds raising \$66 million in venture capital (<https://www.crunchbase.com/organization/fitbit#section-investors>, last accessed 22 September 2020) before being acquired by Google, and Strava raised \$41.9 million over six funding rounds (<https://www.crunchbase.com/organization/strava#section-funding-rounds>, last accessed 22 September 2020).

provide the required return on investments, health apps have to grow and generate profits. To be sure, I am not going to argue that *because* for-profit health apps seek to (also) serve their own ends they must also be manipulative. It would be unfair and unrealistic to place on health apps the *requirement* to *always only* try to track and serve the interests of their users. Such a position would deny the possibility of market activity in the health sector taking place in a legitimate manner, which is too strong a claim to defend. However, the fact that profiting from health apps is not, in principle, illegitimate, does not mean that we should not require for-profit health app providers to observe precautionary principles. Even when a for-profit health app ‘simply’ seeks to maximize profits without intentionally seeking to harm their users, the mere fact that a health app tends to *disregard* the interests of their users can (but does not have to) lead to *foreseeable* illegitimate treatments of (parts of) its userbase.

Now, the fact that health apps profit from their activities does of course not necessarily imply that they cannot also serve the interests of their users. A free market optimist may even propose that competitive market forces will force for-profit health apps to track and serve the real interests of their users as well as possible because a failure to do so would result in customers taking their business elsewhere.

Although there is certainly some truth to the free market optimist’s position, we cannot rely on that position to *guarantee* to us that the interests of the health app provider and the interests of the health apps users will necessarily coincide. First of all, accepting the free market argument implies that for-profit health apps only need to *convince* users they are serving their interests, not that they actually need to serve their true interests.¹³⁸ Secondly, one of the most interesting and challenging aspects of health apps (and similar technologies) is the fact that the digital environments they provide can exert very subtle yet efficacious influences on their users to shape their behavior (Susser, Roessler & Nissenbaum 2019a, 2019b; Zuboff 2019). As a result, we cannot simply assume that consumers will have enough information about, and insight into, the ways in which they are influenced by the app to make a fully informed choice. It follows that health app users may not always be able to judge whether the health app they are using is in fact concerned with their real, authentic health concerns.

138 This is also in line with the earlier discussions of trust and trustworthiness. Trust is, in essence, a psychological state that does not necessarily correspond to an object or actor that can be considered trustworthy.

Still, rejecting the idea of a necessary alignment of interests also does not imply that such an alignment cannot occur. Since we are dealing with *health* apps, it could be suggested that people being preoccupied with their health – a universally desired good – with the help of an app can be rather beneficial to them, even when the provider of the app may ultimately seek to further its own interests. At this point we should briefly revisit my conception of manipulation. One important criterion for manipulation (or manipulative practices) is *not* that a health app should actively try to *undermine* the interests of their users. Rather, we should be alert to situations where (potential) manipulators show a *disregard* for the interests of the people they deal with. So instead of just focusing on the question of whether a digital environment undermines the interests of its users, we should also investigate situations where the design and organization of a digital health environment betrays a disregard for the interests of its users.

To debate this suggestion, we need to briefly go back to chapter 1 where I discussed the more technical side of things. There, I argued that in the contemporary app economy where freemium is the standard, the engineering of user *retention* and *engagement* is essential to the commercial success of one's app. Engaged users keep coming back to an app, keep 'storing value' in it (Eyal 2014), and have a higher chance of interacting with revenue-generating content and features. Because popular for-profit health apps are just another example of the way the freemium app economy works, they are just as any other freemium app organized as optimization systems (Overdorf *et al.* 2018) that optimize for the retention, engagement, and conversion metrics.

Now, it could be argued that optimizing engagement with content in and features of a *health* app translates directly into benefits for the users precisely because they are engaging with *health*-related content and features. Although it makes intuitive sense to equate optimization of engagement with health content and features with optimization of health, we should resist the temptation to accept such a simplistic equation. Health is universally desired and being healthy is a good thing for every person. It does not follow, however, that to optimize one's health, one should optimize one's engagement with (the partly commercial content in) health apps.

Consider, for instance, the calorie counting app MyFitnessPal and its use by users with (a history of) eating disorders. Trying to get users to engage with a calorie counter such as MyFitnessPal as much possible can be highly problematic for this group because it can trigger problematic eating habits, even though it may also be helpful to some of these users (Eikey & Reddy 2017, Eikey *et al.* 2017).

One might object at this point that people with a specific medical condition (e.g. eating disorders) should be considered an extreme case that requires a separate analysis. I would not agree with this objection, because eating disorders are also partially *caused* by health ideals that can be propagated and circulated by health apps.¹³⁹ But even if we exclude those people who suffer from (mental) health afflictions from the analysis, the same argument still holds: maximizing engagement with an app such as MyFitnessPal does *not* necessarily lead to healthier behavior or a happier life. Lupton (2018: 1) argues on the basis of qualitative research conducted with Australian women that food tracking apps such as MyFitnessPal (which is explicitly discussed in the study) can lead to “frustration, disappointment, the fear of becoming too controlled, and annoyance or guilt evoked by the demands of the app.” This is very much in line with my analysis in chapter 2 where I discussed how the MyFitnessPal experience is structured around a healthy lifestyle blog that propagates a very particular and idealized lifestyle which it couples with a very wide range of native advertising posts. Claiming that optimizing engagement with this content is exactly the same as aiming to help users with their health to the best of the app’s abilities is clearly implausible. Getting users to interact with as much native advertising posts which feature idealized images of a healthy lifestyle as possible could, for instance, make people feel insecure about their body or fitness level (Lupton 2018), or it could help create false desires for particular ‘health’ products or services.¹⁴⁰ Put simply, the connection between optimizing engagement and optimizing health is ultimately an accidental rather than a necessary one.

So when a for-profit health app aims to optimize user engagement with health content and features, it cannot be argued that it *therefore* automatically acts in line with the true interests of its users. It could still be asked, however, whether one can speak of a real *disregard* for the users’ real interests in such cases. I think we can. To see why, the concept of ‘encapsulation’ (Hardin 2002, Nys 2016), which I also discussed in chapter 3, can be helpful. Health apps eagerly attempt to *encapsulate* people’s general desire for a healthier and happier life: they want to put people’s general desire for health to strategic use. Health apps want their users to see and understand their service as one that – obviously – helps them improve their health.

139 For example, a review study found that “Indeed, media exposure, perceived pressure to be thin, thin-ideal internalization, and thinness expectancies have all been shown to prospectively predict increased levels of disordered eating cognitions and behaviors (e.g., body dissatisfaction, dieting, bulimic symptoms) in adolescent and young adult females” (Culbert, Racine & Klump 2015: 1145).

140 I am not arguing that the app’s attempt to maximize engagement with native advertising posts will *always* lead to users developing a negative self-image or false needs. All I want – and need – to argue is that optimizing for engagement is *not* exactly the same as being concerned with the user’s health.

The carefully constructed health discourses discussed in chapter 2 are testament to this encapsulation. By emphasizing how important health is, how a healthy life means a happy life, how nearly everything is health-related, and how health should be practiced and optimized, health apps try to communicate that they are there to improve people's health. Such discourses can also help a health app to build a trust(-like) relationship with users. When users trust a health app, it can be easier to encapsulate their interests.

It should be noted that even when a health app encapsulates the interests of its users, the health of users can of course still be improved when using the app. But, and this is crucial, the encapsulation is strategic in the sense that it is *not primarily* focused on health improvement. Encapsulating the health interests of users is a means to serve the commercial aims of the for-profit health apps. Had the real and primary aim of for-profit health apps been to provide optimal health support, they would not aim to *always* and *indiscriminately* optimize retention and engagement because, as we have seen, optimizing engagement with a health app is simply not the same as aiming to provide optimal health support.

So to the extent that health apps aim for indiscriminate engagement (and conversion) optimization, they do in fact disregard the real interests of their users, even though some users may still benefit from using the app. As I explained at the beginning of this section, we cannot require for-profit health apps to ignore their own interests and *only* try to *fully* track and serve the interests of their users. However, when a for-profit health app serves its own interests while (partly) disregarding the interests of its users, we should require such health apps to be cautious not to exploit their users in an illegitimate manipulative fashion. To understand how manipulative practices could take place in this context, I will now turn to the ways in which vulnerabilities can be exploited to render people (more) susceptible to manipulative practices.

4.1.2. Relationship-building with Users: Targeting and Exploitation of Vulnerabilities

It is against the background of the difficult tension between the interests of health app providers themselves and the interests of their users that we should evaluate the ways in which popular for-profit apps try to build relationships with their users. Data analytics as well as the use of the right health discourse play an important role in this process. To build relationships with users in a way that optimizes retention and engagement rates, there is a constant pressure to 'get to know' one's users by building persuasion profiles (Kaptein 2015, Kaptein *et al.*) in order to learn how the app should be tweaked to optimize the desired metrics.

Freemium apps in general, and popular freemium health apps in particular, tend to adopt a rather ‘instrumental’ approach: many of them search for *anything* that can boost their retention, engagement, and conversion rates. This is precisely the reason why popular for-profit (health) apps tend to be data-hungry and find it important to experiment constantly with features and the (visual) design of their app.¹⁴¹ Insights into the users and how they interact with the app can inform the (real-time) personalization of the digital environment provided by the app (Yeung 2017, Lanzing 2018, Lanzing 2019). 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 maximize their engagement with the app.

It is precisely in this constant search for what it is that makes the users ‘tick’, and what the users respond to, that the conditions for *helpful support* and *manipulation* can come together. In the health app context, the pressure to get to know one’s users basically comes down to a pressure to – indirectly – collect information about one’s users’ vulnerabilities.¹⁴² These kinds of optimization systems are basically organized to infer or predict dispositional vulnerabilities of users, and, in a next step, infer or predict when or under what conditions those dispositional vulnerabilities can ‘materialize’ and become actual occurrent vulnerabilities. (Remember that earlier in this chapter I made a difference between vulnerabilities pertaining to a person’s health status, desires, or worries, and vulnerabilities that pertain to a person’s cognitive and affective vulnerabilities that can be leveraged by digital environments to influence behavior.)

As we have already seen in the previous sections, the collection of information about, and the targeting of, vulnerabilities can be used to create helpful mobile health solutions. But when information about vulnerabilities (either health-related, or cognitive and affective biases-related) is collected and used to target users or to (dynamically) adjust aspects of the digital environment *for the sake of optimizing retention and engagement*, we quickly veer into manipulation territory.

141 In chapter 1 I discussed multiple job descriptions for jobs at popular health apps. They all shared a strong focus on data analytics and emphasized the need for constant experimentation with every aspect of the app to make sure it is as engaging as possible to as many users as possible.

142 As I also wrote in the previous chapter, I am finalizing this chapter during the 2020 ‘corona crisis’ which, interestingly, has come with an increased attention for (the concept of) vulnerabilities in relation to our health. See, for instance, the *European Journal of Psychoanalysis* which featured a collection of brief response by philosophers to the global coronavirus crisis, many of which involved reflections (both explicitly and implicitly) on vulnerability (<http://www.journal-psychoanalysis.eu/coronavirus-and-philosophers/>, last accessed 2 June 2020).

A clear example of how this can happen comes from Facebook. In 2017, Australian journalists got their hands on internal Facebook documents which detailed the development of:

algorithms [that] can determine, and allow advertisers to pinpoint, “moments when young people need a confidence boost.” If that phrase isn’t clear enough, Facebook’s document offers a litany of teen emotional states that the company claims it can estimate based on how teens use the service, including “worthless,” “insecure,” “defeated,” “anxious,” “silly,” “useless,” “stupid,” “overwhelmed,” “stressed,” and “a failure.” (Machkovech 2017)

In their search for (almost fully automated) advertising systems that can deliver the highest return on investment, Facebook ended up building a system that was literally designed to identify and target exploitable vulnerabilities in an already vulnerable population. I do not mean to suggest that most (or even some) popular for-profit health apps operate in a similar manner. But the Facebook example should at the very least serve as a cautionary tale of what *can* happen when one builds digital environments that are designed to optimize a limited set of metrics such as retention, engagement, and conversion.

To see how for-profit health apps can – intentionally or inadvertently¹⁴³ – harbor a manipulative potential by identifying and targeting vulnerabilities, it is important to emphasize, again, that optimizing for retention, engagement, and conversion is *not* the same as offering a health service that is, first and foremost, designed to improve people’s health. So when a health app’s dominant design and organization goals are the optimization of these metrics, they are – as argued in the previous section – serving their own interests in the first place, while displaying a (at least partial) disregard for the interests of their users. If these self-serving interests are pursued *by means of targeting vulnerabilities*, we have to conclude that most (but not yet all) criteria for manipulation have been met. To fully see the manipulative potential of such digital health environments, remember that a manipulator will

143 The attentive reader may object that I have argued that manipulation is – by definition – intentional, so how can a digital environment *inadvertently* harbor a manipulative potential? What I am claiming here is that for-profit health apps going about their business of making the app a financial success may end up identifying and targeting specific vulnerabilities without necessarily aiming to identify and target *those specific vulnerabilities*. So even though the identification of the *specific vulnerabilities that end up being targeted* is not intentional per se, the entire design and operation of the digital environment is of course a thoroughly intentional process and the fact that the design of digital environment is going to lead to the targeting of *some* vulnerabilities is entirely foreseeable.

look for ‘buttons to push’ and ‘strings to pull,’ i.e., the manipulator will try to find out what kind of targeted treatments her targets might be responsive to in order to serve the ends of the manipulator. Information about one’s targets is important here, because it can be used to find out or estimate what the targets are responsive to. Rudinow (1978: 346 emphasis mine), for instance, argues that “the manipulator’s behavior is normally either deceptive or *predicated on some privileged insight into the personality of his intended manipulee.*” Manipulation, in other words, typically works best when one has access to persuasion profiles (Kaptein 2015, Kaptein *et al.* 2015). What digital environments like health apps offer, is the ability to test *all* known buttons and pull *all* known strings by continuously iterating through different experiments. Their ability to continuously run experiments and receive real-time feedback on the experiments down to the level of individual users also means that, in theory at least, real-time personalized targeting based on an individual’s unique exploitable characteristics is possible (Kaptein *et al.* 2015, Yeung 2017).

There is one last aspect of vulnerabilities that should be discussed in this context. There is increasing attention for the fact that digital environments like health apps cannot only identify and target *preexisting* vulnerabilities, but might also be able to *create* or *exacerbate* them (Hanson & Kysar 1999, Calo 2014). Not only technological features (e.g., data analytics, targeting) but also communication strategies and adopted health discourse should be considered in this regard. As I explained earlier (chapter 2), 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, Devisch 2013, Lupton 2013). Health apps can ride this healthism 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* put it in a piece aimed at other advertisers.¹⁴⁴ 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 communicate that nearly everything in a person’s life is relevant to one’s

144 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/> (last accessed April 16, 2019). Screenshot available here: <https://imgur.com/a/aUL9qdj>

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 more receptive to the engagement optimization practices.

Put differently, one could say that 1) through the propagation of their health discourses, and 2) through their personalization and targeting capabilities, popular for-profit health apps can *create* or *exacerbate* vulnerabilities in their users which can then be targeted and exploited. As far back as 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."

4.1.3. *Relationship-building with Users: Trust, Trustworthiness, and Hidden Influences*

So far we have established that 1) for-profit health apps tend to function as optimization systems which optimize retention, engagement, and conversion, and 2) that the prioritization of these metrics incentivizes such apps to constantly search for vulnerabilities of their users and to constantly experiment with different ways to exploit those vulnerabilities. However, to speak of a *problematic* and potentially *manipulative* relationship between such apps and their users, we need to consider more criteria for manipulation. After all, it might be entirely clear to health app users that they are involved in an instrumental relationship with the app, in which they use the app and the app uses them.

One of the criteria for manipulation is that manipulators attempt to keep their manipulative influences hidden. When the motives and functioning of a for-profit health app are (almost entirely) transparent to the user and the user consciously decides to accept the package deal, we could conclude (tentatively) that there is no harm and no foul.¹⁴⁵ But we should also consider the possibility that the deal users are making when they start and continue to use a health app is not, in fact, always as clear as it may seem. To see why, this section will focus on the ways in which

¹⁴⁵ Of course, not every deal in which the contents are entirely transparent to the chooser is, by definition, unproblematic *because* the terms of the deal are transparent. Think for instance of extortion cases, which are usually transparent but the situations in which the deals are struck are problematic.

health apps try to instill trust in their users, and how trust and trustworthiness can come apart when health apps also try to obfuscate or hide certain aspects of their conduct.

The core tension which we should analyze is the following. The ability of an app to successfully build trust can be very beneficial *to the users* of an app. When users trust an app – as a result of which they feel comfortable volunteering data and allowing the app inside their decisional sphere – they also allow the app to help them better. However, in the freemium app economy, popular for-profit health apps also have a clear interest in engineering trust because the better the app is able to instill the mental state of trust in their users, the better it will be able to optimize retention and engagement. So the question becomes to what extent the building of trust is a *helpful feature* of the app, and to what extent it is a clever *business practice* that is not necessarily backed by a corresponding underlying trustworthiness.

To see how trust and trustworthiness can potentially come apart, we should consider the ways in which popular for-profit health apps sometimes attempt to obfuscate, hide, or direct attention away from some of their intentions, design choices, and so on. The main idea here is that when health apps try to engineer trust, but at the same time also attempt to hide or obfuscate particular elements or dimensions of their apps, there are good reasons to question whether the engineered trust is commendable. To bring out some of the ways in which health apps can attempt to impose hidden influences on their users, I mainly focus on the health discourses in apps, as well as their layout and communications to their users.

One common strategy is for health apps to craft and use a health discourse that helps to distract users from the *commercial* nature of the relationship between themselves and the app (Sax, Helberger & Bol 2018). Put differently, to be commercially successful, it is in a health app's interest to be perceived as a *helpful health service* rather than a *profit-driven service*. This is of course not unique to health apps; even internet service providers want to sell their service by presenting the choice for a particular internet service provider as a lifestyle choice.¹⁴⁶ What makes health apps unique, however, is the fact that they deal with a universally desired value – health – which affords them with lots of possibilities to frame their activities in a manner that can sound appealing to nearly everyone.

146 For example, the Dutch internet service provider KPN presents itself as selling services which connects families, closeness with loved ones, and meaningful emotional moments: <https://www.youtube.com/watch?v=2D5KSIgD7N0>

It is therefore no accident that one encounters carefully constructed health narratives in popular for-profit health apps. As I explained in chapter 2, a health app can attempt to (co)shape how users understand the app and the app's relationship to their health, as well as the users' understanding of their own health, through the productive power of a particular health discourse. The propagation of a particular health discourse can thus be used to draw attention away from the (potentially) manipulative elements of the app – i.e., the fact that popular for-profit health apps primarily aim to optimize engagement and conversion through the continuous identification and targeting of vulnerabilities. Such an attempt to keep manipulative elements of an app out of the users' sight can *itself* be understood as a (further) manipulative element of an app.

Looking at real examples, popular for-profit health apps tend to emphasize (1) how important health is in general; (2) how important this particular health app is when it comes to securing the offered health benefit(s); and (3) how the app is concerned about *you* and *your* health. Moreover, apps tend to have visual designs that communicate how fun, caring, and/or simple the app is. Take, for example, the Fitbit app I discussed earlier. Fitbit emphasizes how important health is and how Fitbit helps you pursue this important goal by declaring 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.”¹⁴⁷ Fitbit also stresses that it is focused on “how *you* spend *your* days” because that is what determines “when *you* reach *your* goals” and “seeing *your* progress helps *you* see what’s possible.” Fitbit helps “*you* find *your* fit,” and supports “*your* path.”¹⁴⁸ The message is clear: Fitbit is merely a tool that exists to help you get in shape to improve your health.

Headspace is another good example because the app and the accompanying website clearly display all the health discourse elements I just listed. When visiting the homepage of Headspace, one is greeted by the following text: “Your guide to health and happiness. Live a healthier, happier, more well-rested life with Headspace.” Directly below this welcoming text, the page continues by stressing how this app is going to make your life healthier and happier. Here, Headspace clearly communicates that users who decide to use the app will enter into a relationship with the app that is focused on helping *you* as a user get healthier and happier – which is important to do because meditation offers so many benefits. Moreover,

147 <https://www.fitbit.com/whyfitbit> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/qw8nb5u>

148 <https://www.fitbit.com/whyfitbit>, emphases added (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/qw8nb5u>

Headspace has a very consistent visual style that is built around very simplistic cartoon-like characters that smile or show other basic emotions. The colors used throughout the interface are all soft and calming and the design is very minimalistic. When using Headspace, the entire user experience is designed to make the user feel she is using a very simple service that just offers some helpful meditation sessions. The advanced data analytics and engagement optimization efforts are carefully kept under the user's radar. Health apps like Headspace and Fitbit thus try to create a user experience that feels like a genuine *health* experience to make the commercial nature of the user-app relationship less apparent.

Attempts to sell a particular idea of health and, moreover, to sell the image that the app in question is there *just* to help you with your health, can be understood as attempts to evoke trust in users. When such attempts *distract* people from the *commercial* nature of their relationship with the app they can contribute to the manipulative potential of a digital environment.

Another clear example of attempts to keep (potentially) manipulative influences hidden is the frequent use of native advertising. As I explained earlier, native advertising, “also called sponsored content, [...] is a term used to describe any paid advertising that takes the specific form and appearance of editorial content from the publisher itself” (Wojdyski & Evans 2016: 157). In the case of health apps, this usually means that advertising content is intentionally made to look and feel like genuine health content. 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). The use of native advertising is thus a clear example of health app providers deliberately attempting to hide the commercial nature of some of the content they try to get users to engage with. In chapter 2 I explained how the user experience of MyFitnessPal is built around the MyFitnessPal blog, which features lots of native advertising.

Besides the role health discourse can play in obfuscating the *commercial* nature of the relationship between app and user, we should also consider the role insights from behavioral economics play. Because most popular for-profit health apps are designed and run as engagement optimization systems, they also seek to identify and target vulnerabilities of their users, such as the cognitive and affective biases that have been elaborately described by authors such as Thaler and Sunstein (2008) and Kahneman (2011). When cognitive and/or affective biases are targeted, there is

a serious chance that people's behavior is influenced without them (fully) realizing that it is happening or how it is happening (Susser, Roessler & Nissenbaum 2019a, 2019b). This is precisely why targeting cognitive and affective biases as popularized by the behavioral economics movement is so popular in Silicon Valley and among policy makers: when targeted correctly, one can evoke behavior change without having to *convince* people. Biases operate (semi-)automatically and can thus also be leveraged to influence people's behavior (partly) behind their back. Importantly, data analytics and targeting informed by behavioral economics can also be used to find out which personalized treatments work best to instill trust in users. When such insights are used to *engineer* trust by exploiting cognitive and affective biases, subjectively experienced trust can easily become detached from (more) objective trustworthiness.

In sum, for-profit health apps have a commercial interest in making their users trust their app. Health apps can use a variety of techniques to instill or evoke trust. Some of these techniques, such as the ones discussed above, are designed to exert – or at the very least risk exerting – a hidden influence on their users. To the extent that such hidden influences are exerted to *hide* or *distract attention* from information that can be important to users – such as the fact that the relationship between app and user is (also) a commercial one – we can speak of a manipulative potential.

4.2. Conclusion: The Manipulative Potential of For-Profit Health Apps

As I already explained in chapter 3, we can make a distinction (following Susser, Roessler & Nissenbaum 2019a, 2019b) between individual acts of manipulation and manipulative environments or practices. The potential of digital health app environments manipulating users into making *specific decisions* that users would not, all things considered, necessarily agree with or feel comfortable with if they critically reflected on them, should be taken seriously. Users of health apps making decisions they would not autonomously want to make, as a result of a manipulative interference, constitutes a clear, identifiable harm. However, as I explained in the introduction to this chapter, I am mainly interested in analyzing and evaluating manipulative *environments* because this allows me to address more *structural* issues of power and information asymmetries and the influence digital environments have on shaping our self-understanding and behavior.

The perspective of manipulative *environments* – as opposed to manipulative *acts* – thus helps us see the more structural ways in which health apps can shape and steer our behavior. For-profit health apps do not only seek to influence individual decisions in a manner that is 'optimal' from the perspective of conversion.

It is at least as, if not more, important to influence their users in a manner that helps to give shape to a *long-term relationship* between user and app that serves the commercial goals of the respective app. The overall goal is to build a digital environment that can *create* and *foster* relationships with the userbase that are optimal from a business perspective. In practice, this means that health apps will, for instance, try to make users feel that the app in question is important to their quest to improve or maintain (an aspect of) their health (in the broadest sense of the word) and lifestyle. We can thus see the interactions between the user and the app as a kind of *habituation process* where the app tries to shape the users' perception and understanding of their health in a manner that suits the app.

The manipulateness of a digital health environment, then, can certainly *not* be reduced to its commercial nature. The (hyper)commercialization of health in health apps is an important *precondition* for the building of manipulative digital health environments, but it is certainly not a *sufficient* condition. The real manipulateness takes shape when the commercial nature of health apps leads to the development and implementation of systems that seek to optimize metrics that are not necessarily beneficial to users, by employing tactics and techniques (e.g., personalized targeting) that are intentionally designed not to be transparent and systematically seek to sniff out characteristics of groups and individuals that can be exploited.

5. Manipulative Health Apps and User Autonomy

Having established that popular for-profit health apps can under certain conditions be considered as manipulative environments, the last question to answer is what this means for the autonomy of health app users. In chapter 3 I argued that a basic procedural account of autonomy should be reconciled with the insights from the literature on relational autonomy. Autonomy is, in essence, about competent, authentic decision-making. We are, however, already continually thrown into social contexts together with other people (or technology) and it is within these social contexts that we need to develop and practice our autonomy. Autonomy is thus also deeply relational. In the digital society, we therefore need to ask how the (affordance of the) digital environments we navigate structure our possibilities for developing and practicing autonomy, just like we would ask for offline contexts. In the case of health apps the question thus becomes whether the types of digital environments they offer, as well as the user-app relationships that develop within those digital environments, support or undermine autonomy.

Let us start, briefly, with individual acts of manipulation where the relation between autonomy and manipulation is most straightforward. A decision that is the result of manipulation is not an autonomous decision. Autonomy requires a person to critically reflect – to some degree – on her reasons for acting, her desires that inform her reasons, and her options available, to come to a decision she can identify with and consider her own.¹⁴⁹ Put simply, “manipulation challenges both conditions of autonomy – competency and authenticity” (Susser, Roessler & Nissenbaum 2019a: 38).

When we shift our focus from isolated acts of manipulation to the manipulateness of environments, the connection between autonomy and manipulation becomes less straightforward – but not less significant. For single acts of manipulation, we can point to a clear moment or instance where autonomy is undermined. However, when adopting the perspective of manipulative *environments*, the (main) aim is not to point to single, isolated acts of manipulation, but rather to the structural tendency and ability to shape not only particular decisions but also longer-lasting behavior patterns. To still be able to explain how such manipulative digital environments can undermine user autonomy without having to refer to single, isolated ‘incidents,’ we can rely on autonomy’s inherent *relationality*.

As we have seen, popular for-profit health apps try to build long-term *relationships* with their users, within a sophisticated digital environment they themselves control. From the perspective of autonomy, we should then ask how these relationships impact the possibilities for practicing autonomy. Are the digital environments that these health apps offer the kinds of environments that help people to become (or remain) competent decision-makers when it comes to their health, and do they stimulate people to make health-related decisions that are informed by authentic desires?

If one wanted to design a digital health app environment that was, from a relational perspective, supportive of the user’s autonomy, one would want that environment to be designed to attend to the specific needs of different users. Such an environment would try to identify what different users *genuinely* (or: authentically) want and, equally important, would stimulate users to try to find out what they genuinely

149 As I discussed elaborately in chapter 3, we do not have to – and should not – require a person to make an autonomous decision in complete isolation. Other people (e.g., friends, lovers, family) can shape intentions for acting, or give one the necessary confidence for considering oneself a sufficiently autonomous person. Such ‘relational influences’ do not undermine – but tend to *shape* – one’s autonomy to the extent that they, in the end, help one to develop into a person that feels competent and worthy to make her own decisions.

want – e.g., by stimulating critical reflection and careful experimentation, or by trying to provide a digital environment that fostered self-respect, self-trust, and self-esteem (Anderson & Honneth 2005). Moreover, such a hypothetical maximally supportive digital health environment would try to engage users with the health services it provided only to the extent that they actually contributed – or were expected to contribute – to users’ health and lifestyle decisions. Given the current state of technology, which has been discussed throughout this dissertation, we can actually imagine there being health apps that (try to) provide such digital environments. The great promise of empowerment through health apps that can provide personalized health and lifestyle solutions is, from this perspective, entirely understandable.

Popular for-profit health apps do not always fulfill these promises of empowerment. Their digital environments are not designed to meet the above-described optimal conditions for the scaffolding of the autonomy of their users. Quite to the contrary, they tend to be run as cynical optimization systems, just like any other type of for-profit freemium app. Their environments are built to recruit as many users as possible and to get them to spend as much time as possible in their app in order to have them interact with as many revenue-generating content, features, and services as possible. To be sure, such digital environments do not *preclude* the practicing of autonomy; it is still possible for users to use the apps in a way that does not (significantly) undermine their autonomy. But it should also be acknowledged that the digital environments they provide are *not designed with just the goal of stimulating their users to deal with their health and lifestyle in an autonomous fashion in mind*. Rather, the types of relationships they seek to build and nurture with their users are relationships of a colder, more calculative, profit-driven nature. This is not a surprising conclusion. It is, however, a conclusion that bears repetition to not be lulled to sleep by the overly rosy and often propagated discourse on self-empowerment through health apps.

Chapter 5

Health Apps and Unfair Commercial Practices

1. Introduction

In the previous chapter I provided an *ethical* analysis of for-profit health apps. I explored the vague but important line between helpful and potentially manipulative digital health environments. I concluded that digital environments can, under the right conditions, be very helpful tools. However, in the context of a freemium app economy that tends to prioritize the recruitment, retention, engagement, and conversion of users, there is also a clear potential for for-profit health apps to design and offer manipulative digital health environments to their users. One key insight from the previous chapters is that to truly understand the influence and impact of digital (health) environments, we should not just look at single, isolated interactions between users and said environments, but also, more importantly, at the relationship-like dynamic that develops *over time* between a digital environment and its users. It is in this relationship-building enabled by dynamically adjustable, personalizable choice architectures that most of the manipulative potential is to be located. Although this more structural perspective on for-profit health apps and their users helps identify manipulative potential, it does not allow for brief, clear, and snappy conclusions. Most apps cannot be ruled to be manipulative in their entirety. More often than not, there are certain features or element of apps that can be considered manipulative and therefore ethically problematic, while other features or elements of the same app are not necessarily ethically problematic. The resulting conclusions can feel a bit messy. This, to me, is not a bug, but a feature; the for-profit health app context simply *is* a bit messy, precisely because of the uneasy tension between developing business practices to profit from people's pursuit of health and the simultaneous attempt to (really) help people get healthier.

This empirical and normative messiness carries over to this *legal* chapter, where I explore the possibilities of legally addressing the manipulative potential of health apps. I will not pretend to be able to present an exhaustive legislative agenda to tackle all issues for-profit health apps (potentially) give rise to. What I will offer is an *exploration* of the Unfair Commercial Practices Directive (UCPD) of the European Union to search for possible legal angles to address the challenging tensions I have presented in chapter 4. My reason for focusing on the UCPD is straightforward: the types of challenges this Directive seeks to address – namely the undermining of consumers’ autonomy as a result of influences that can be considered unfair – align rather well with the ethical analysis I have provided. I do not argue that the UCPD offers a complete, or the best, answer to ethically illegitimate health app practices. For example, health apps also give rise to obvious and important privacy and data protection questions, which are better addressed by the General Data Protection Regulation (GDPR). In the digital society, many of the legal questions posed by online services (including health apps) are best addressed by *combining* different legal approaches. Helberger, Zuiderveen Borgesius and Reyna (2017) as well as Van Eijk, Hoofnagle and Kannekens (2017) have argued convincingly that European data protection law and European consumer law can, and should, be considered as highly complementary. My focus on the UCPD is thus pragmatic. The type of issues I focus on align better with the aims of the UCPD, seeing that the GDPR focuses on the *processing* of personal data, but not explicitly¹⁵⁰ on the subsequent *use* of those data for, for instance, manipulative practices. Moreover, the GDPR has already received a fair amount of attention in the context of health apps since it is more of a ‘usual suspect.’¹⁵¹ So given the underexplored nature of the UCPD in the context of health apps and the promising alignment between my ethical analysis and the aims of the UCPD, I choose to focus on the UCPD.

In what follows, I provide an introduction to and explanation of the UCPD’s most important aims, definitions, and mechanisms. The perspective I adopt here is an ‘external’ one, in the sense that I rely not only on the text of the UCPD itself, but also on secondary literature which discusses and analyzes the UCPD. As I discuss the Directive, I also address the interpretational questions the Directive may raise vis-à-vis popular for-profit health apps. First of all, it should be asked whether and to what extent the commercial practices of popular for-profit health apps are covered by the UCPD. I will argue that many of the potentially manipulative

150 See, for instance, Clifford and Ausloos (2018) who argue that the principle of fairness is one of the master values (largely implicitly) underlying the GDPR.

151 Since health apps rely on the collection, storing, analyzing, and disseminating of personal data they naturally capture the attention of data protection lawyers (see, e.g., Gostin, Halabi & Wilson 2018, Muchagata & Ferreira 2018, Mulder 2019, Mulder & Tudorica 2019).

features of the digital health environments provided by popular for-profit health apps can be addressed by the UCPD. The second question is what the UCPD can achieve, given the fact that it should cover many commercial health app practices. In addressing these interpretational questions, I will rely on my ethical analyses from the preceding chapters to explore how ethics can inform our understanding and use of the UCPD. In doing so, I will always attempt to offer interpretations that should at least be *plausible* to readers with a background in law. But, and this bears emphasis, my aim is *not* to offer a detailed doctrinal analysis of the UCPD.

This also brings me to one important caveat about the following legal analysis. Unfair commercial practice law in Europe is very much an example of ‘judge-made law;’ in the end, national judges are the ones who have to decide on the interpretative questions discussed below and who, through their judgments, give shape to unfair commercial practice law. My legal analysis will not, however, be based on an exhaustive, detailed analysis of national UCPD judgments. I will mostly rely on secondary literature which, of course, discusses and is informed by national case law. My aim is not to stay *as close as possible* to the current developments in case law, but rather to ‘import’ my ethical insights into the Directive by way of volunteering new interpretations.

2. The Unfair Commercial Practices Directive

2.1. General Aims: Autonomy & Fairness

Influencing consumers is an essential feature of the marketplace. As a commercial party you can only be successful if you manage to reach consumers and get them to become a customer of your product or service. A wide variety of common market practices exemplifies this. For example, advertising is carefully designed for and targeted at particular groups or even individuals; discounts are offered; customer loyalty programs are carefully managed. Although many commercial practices which aim to influence consumers are seen as unproblematic and ‘part of the game,’ there are also practices that clearly are problematic. Such practices include the use of outright lies to advertise goods and services, or the use of blackmail to ‘persuade’ consumers. Between the ‘business as usual’ cases and the clearly problematic cases lies a large grey area of commercial practices of which it is more difficult to determine whether, and if so why, they are problematic. In the health app context, we can refer back to what I have called the ‘vague line’ between helpful and manipulative commercial health app practices. What is clear though, despite the existence of this vague line, is the fact that we want to be able to draw the line

somewhere. The Unfair Commercial Practices Directive is an answer to the question of where to draw the line between legitimate and illegitimate attempts to influence consumers.

At its core, the UCPD is concerned with “[t]he autonomy of the consumer’s decision-making process” (Howells 2006a: 174). It aims to “keep and maintain the consumer’s autonomy” (Micklitz 2006: 104). The central idea is that if consumers have sufficient and correct information at their disposal and if their capacities for autonomous decision-making are not otherwise undermined by the (commercial) influences exerted on them, they are in a position to make economic decisions in the marketplace in an unproblematic fashion. To that end, the Directive focuses on business-to-consumer commercial practices “directly related to influencing consumers’ transactional decisions in relation to products”, and lays down norms to determine when such influences are *distorting* consumers’ economic decision-making (Rec. 7). To understand when an influence *distorts* the consumers’ decision-making, one should ask how the influence impacts a consumer’s autonomy and her ability to make an informed decision. “The true challenge,” Micklitz (2006: 104) writes, “is to give shape to the consumer’s autonomy, which lies at the heart of the Directive.” The task of the Directive, then, is to formulate legal norms that help capture when a commercial practice is unfair because it threatens the consumers’ autonomy in an illegitimate fashion.

However, Micklitz (2006: 104) also points out that “[o]n its own, the concept of autonomy in the Unfair Commercial Practices Directive, even in the entire European unfair trading law, seems largely incomprehensible. In essence, it is a legal term; however, it needs a reference point outside the legal system.” So although it is clear that autonomy is one of the guiding principles, if not *the* guiding principle, of the UCPD, it is not all that clear how autonomy should be understood. An important anchor point for understanding the significance of the concept of autonomy in European law is the landmark decision by the German Constitutional Court in 1983, which established the concept of “informational self-determination.”¹⁵² In this landmark decision the German Constitutional Court clearly established that persons have a general right to their personality and that this fundamental right “guarantees the authority conferred on individuals to, in principle, decide for themselves on the disclosure and use of their personal data.”¹⁵³ The decision mainly

152 BverfG, 1 BvR 209/83, 15 December 1983, last accessed 22 September 2020 at http://www.bverfg.de/e/rs19831215_1bvr020983.html

153 Abstract of the official translation of BVerfG, 1 BvR 209/83, last accessed 22 September 2020 at https://www.bundesverfassungsgericht.de/SharedDocs/Entscheidungen/EN/1983/12/rs19831215_1bvr020983en.html

focuses on privacy and the right people have to decide who gets to collect, access, and disseminate their personal data. This is understandable because privacy is *constitutive* of personal autonomy.¹⁵⁴ We can, however, also see the contours of a (implicit) theory of autonomy here. In discussions on the difference between the rationale for privacy protections in the European Union and the United States, different authors have emphasized that in Europe such protections are grounded in the dignity and inviolability of the person and her right to determine for herself how to act and present herself (Zwick & Dholakia 2001, Kang & Buchner 2004, Kobrin 2004, Whitman 2004). Informational self-determination as a right thus enables and protects a person's right and ability to practice her autonomy. Although this general background to the value of autonomy in European law helps us situate and contextualize the role of autonomy in the UCPD, it falls short of providing a full-blown theory of autonomy. In what follows, I will also build on the *ethical* concept of autonomy I myself have been developing throughout this dissertation to inform my (largely) *legal* analyses in this chapter. My extra-legal conception of autonomy should thus be seen as the "reference point outside the legal system" that Micklitz (2006: 104) refers to.¹⁵⁵

So let me briefly revisit my conception of autonomy and explain how it informs the rest of this chapter. My conception of autonomy should not be seen as an alternative to the legal conception of autonomy associated with the literature on informational self-determination. The idea of informational self-determination is perfectly compatible with my conception of autonomy, which should rather be seen as a further elaboration of the foundational ideas enshrined in the concept of informational self-determination.

Put simply, autonomy is about people being able to give shape to, and have authorship over, their own lives through the decisions they make. Put in slightly more technical terms, autonomy is about 1) the competency of decision-makers to make decisions that are their own and that they can understand as their own; and 2) about the authenticity of desires that inform (competent) decision-making. As emphasized throughout this dissertation, these requirements should be interpreted in the light of realistic, everyday circumstances. In practice, this means that a person does not have to be a maximally rational, calculated, ideal typical decision-making machine. What is required is a *sufficient* (but not necessarily a maximal) amount of critical reflection on, and identification with, one's desires and decisions. Adopting

154 See Roessler 2005 for a detailed analysis of how privacy is constitutive of autonomy and what that implies for privacy's value.

155 See Sax, Helberger & Bol 2018 for a similar argument. In that article we did not develop autonomy's relationality because we had a limited amount of space at our disposal.

autonomy as a guiding principle does not require us to conceptualize autonomy in an overly ideal fashion (see chapters 3 and 4). Moreover, it is especially important in the context of unfair commercial practice law to emphasize autonomy's *relational nature*. The environments within which persons reside and operate co-determine not only how they can become autonomous – e.g., by providing sources of self-respect, self-trust, and self-esteem (Anderson & Honneth 2005) – but also their possibilities for actually practicing their autonomy. Because for-profit health apps are essentially sophisticatedly engineered digital environments, it is important to look further than the individual person and her competency as a decision-maker, and incorporate the ways in which features of the digital environment – and the commercial relationships that are developed within those environments – help (or do not help) with the development and practicing of autonomy. If we go back then to the question of what autonomy as the guiding principle for unfair commercial practice law should entail, it turns out that a narrow focus on the question whether separate, individual decisions by consumers are autonomous is an incomplete and outdated focus. The autonomy of the consumer can only be properly understood and safeguarded if one acknowledges autonomy's relationality and the fact that digital commercial environments (try to) build relationships over time. This is of course not to say that there is a moral or legal obligation to only build commercial digital environments that *promote* autonomy as much as possible. The absence of such a positive obligation does not, however, imply a permission to build digital health environments that are manipulative.

There is one further, related, perspective that can help us understand the general aim of the UCPD: a fair marketplace. Seen along the lines of fairness, the UCPD mainly deals with fairness by focusing on the asymmetry of power and information between traders and consumers. Because an asymmetry of power and information translates into an asymmetry of proverbial 'weapons' at the disposal of traders and consumers in the marketplace, the fairness approach in the UCPD can be understood as the aim to 'soften' the asymmetry by curtailing the (ability to exercise) power of traders. The Directive tries to do so by banning commercial practices that exploit said asymmetry in an unfair manner.

2.2. Structure of the Directive

The general legal framework to define the unfairness of commercial practices has a "threefold structure" (Micklitz 2006: 86). At the first level, there is the general clause on unfair practices (Art. 5) which rules all commercial practices that are considered unfair to be prohibited. According to this general clause, a commercial practice is unfair if "it is contrary to the requirements of professional diligence [of the trader]" (Art. 5(2)(a)) *and* "it materially distorts or is likely to materially

distort the economic behavior with regard to the product of the average consumer whom it reaches or to whom it is addressed, or of the average member of the group when a commercial practice is directed to a particular group of consumers” (Art. 5(2)(b)). At the second level, the Directive defines two types of practices that are *particularly* unfair, namely commercial practices that are *misleading* (Arts. 6-7) or *aggressive* (Arts. 8-9). At the third level we find Annex I (the ‘blacklist’), which lists 31 specific practices which are always prohibited. By moving from level one to level three, one moves from the most general to the most specific rules.

Given these three levels, it should be asked what the internal *logic* of the Directive is: how do these three levels relate to each other? I follow Micklitz (2006: 86) who argues that “the threefold approach of the Directive must be turned upside down. Any application would have to start with the test of whether the practice is prohibited under Annex I and, if not, whether it is misleading or aggressive.” The general clause would then only come into play as “a remedy of last resort” in cases where the more specific clause cannot capture the unfairness of a commercial practice, yet there still is a strong suspicion that the practice is particularly unfair.¹⁵⁶

The 31 prohibited commercial practices listed in Annex I are so specific that they will rarely be directly applicable, meaning one will have to move on to the clauses on misleading and aggressive practices in nearly all cases. So under this proposed interpretation, the clauses on misleading (Arts. 6-7) and aggressive (Arts. 8-9) commercial practices will, in practice, be the most important ones for judging the

¹⁵⁶ This interpretation is also in line with the case law of the Court of Justice of the European Union (CJEU). To see why, one must remember that the general clause (the first level) is formulated in a *cumulative* manner: a practice is unfair if it is contrary to the requirements of professional diligence *and* if it is (likely to) distort the economic behavior of the average consumer of the targeted group. In *CHS Tour Services* (C-435/11) the CJEU was confronted with a case where the trader “cannot be criticised for any breach of the duty of diligence,” yet also, through no fault of their own, had misled consumers, thereby distorting their economic behavior. Now, the question the CJEU had to answer was whether the conduct of the trader should be considered unfair because consumers were misled, even though the trader met all the requirements of professional diligence. If the general clause with the cumulative criteria (first level) is considered to be *logically prior* to the clauses on misleading and aggressive practices (second level) and the Annex I (third level), it should follow that a commercial practice cannot be considered unfair and prohibited as long as the vendor has met all requirements of professional diligence. As the CJEU made clear, however, a commercial practice can be considered misleading *and therefore unfair*, without having to establish whether a trader acted contrary to the requirements of professional diligence: [I]f a commercial practice satisfied all the criteria specified in Article 6(1) of that directive for being categorized as a misleading practice in relation to the consumer, it is not necessary to determine whether such a practice is also contrary to the requirements of professional diligence as referred to in Article 5(2)(a) of the directive in order for it legitimately to be regarded as unfair and, therefore, prohibited in accordance with Article 5(1) of the directive (C-435/11, para 49).

unfairness of commercial practices. The general clause then only comes into play as a “remedy of last resort” as Micklitz (2006: 86) puts it. I will therefore spend most of my time explaining when commercial practices count as misleading or aggressive. Before I do so, I will first briefly discuss the scope of the Directive: when does a practice qualify as a *commercial* practice, and what counts as “influencing consumers’ transactional decisions in relation to products” (Art. 5(2)(b))?

2.3. Scope of the Directive: Are Health Apps Covered?

To understand the scope of the UCPD, and to see whether the commercial health app practices I discuss fall within the scope of the Directive, there are two main concepts that need to be discussed: (1) the related concepts of a ‘transactional decision’ and ‘economic behavior’ of the consumer, and (2) the concept of a ‘commercial practice’ itself.

At first glance, the UCPD might seem to have a relatively narrow scope since it only applies to ‘transactional decisions’ and the ‘economic behavior’ of consumers. Moreover, Article 1, which lays down the purpose of the Directive, speaks of protection against “unfair commercial practices *harming* consumers’ *economic* interests” [emphasis added], which can, when interpreted stringently as the need to prove actual significant monetary damage, be a difficult standard to meet.

In practice, however, these terms get such broad interpretations that the UCPD “covers any of the vast range of transactions that consumers might enter into, as long as there is some commercial flavor” (Willett 2010: 249). Wilhelmsson (2006: 49) also emphasizes that the Directive has “a very broad scope.” Given this broad understanding of what constitutes a ‘transactional decision’ or ‘economic behavior,’ it is also not difficult to argue that many of the practices popular for-profit health apps engage in may have an impact on the economic behavior of consumers.

The broad scope of the Directive can be understood along different lines. First, the Directive covers a wide range of practices, as Article 2(d) of the Directive makes clear: it covers “any act, omission, course of conduct or representation, commercial communication, including advertising and marketing, by a trader, directly connected with promotion, sale or supply of a product.” Second, the Directive has a “wide temporal scope” (Wilhelmsson 2006: 57), as Article 3(1) specifies: “This Directive shall apply to unfair business-to-consumer practices [...] *before, during* and *after* a commercial transaction in relation to a product” [emphasis added]. Third, when considering the phrase ‘harming consumers’ interests,’ Wilhelmsson (2006: 58) observes that “the importance of the phrase at hand relates more to the word ‘interests’ rather than ‘economic,’ since the latter has to be understood very

broadly anyway.” The Directive’s main purpose is explicitly *not* to protect individuals against significant economic harm. Rather, the Directive focuses on the *collective* protection of consumers. “One of the purposes of collective protection of consumers is to create a remedy when the actual losses of individual consumers are so small that individual mechanisms for redress are not worthwhile to use” (Wilhelmsson 2006: 58). Helberger (2016: 153) also emphasizes the significance of the Directive’s focus on the *collective* protection of citizens: “It [the UCPD] also touches upon broader, more conceptual questions about the kind of information economy we would like to live in, and the values that shape it. Fairness, privacy, autonomous choice may be important rights or entitlements of individual consumers/citizens, but they are also the quintessential building blocks of a free digital society.”

So put simply, the Directive covers everything a trader does not only to *actually* make consumers buy products or services, but also to promote them, stimulate consumers to continue to use the product or service, and so on. So long as the trader’s conduct vis-à-vis consumers has “some commercial flavor” (Willett 2010: 249) it will be covered. For a practice to be possibly considered unfair, the Directive also does not require the practice to be an egregious distortion of consumers’ behavior; very minor distortions – or even *likely* minor distortions – potentially fall within the scope of the Directive. The scope of the Directive can thus rightfully be called very broad (Wilhelmsson 2006; Willett 2010; Sax, Helberger & Bol 2018). But is it – in this respect – broad enough to cover the commercial health app practices I seek to address in this dissertation? Remember, I am not just interested in individual, isolated interventions or offers by a health app that can be considered manipulative, but also, and more importantly, in the more structural ways in which digital health environments are designed in (sometimes) manipulative ways to give shape to commercial relations with health app users. For their monetization efforts, popular for-profit health apps (most of which are freemium apps) *depend* on the continuous retention and engagement of users. The mere fact that an app manages to retain and engage users is, from the perspective of the app, economically significant. The freemium business models works by engaging users to let them store value in the app, which increases the likelihood they will engage with revenue-generating features. So it seems to follow that when 1) a health app provider intentionally designs (features of) an app aimed at retention and engagement, and 2) as a result of those efforts users (continue) to use a health app, which 3) directly contributes to the business interests of the app, the health app has influenced the economic behavior of the users. Put differently, when users decide to start using a health app, or decide to *continue* to use a health app and volunteer (user) data to the app, such decisions can be considered transactional decisions – even if consumers don’t spend any money and are using, or continuing to use, a health app that is, in principle, ‘free.’

So far I have established that the Directive's broad interpretation of 'transactional decision' and 'economic behavior' means that most commercial health app practices should fall comfortably within the scope of the Directive. One further argument needs to be developed, however, to determine whether the commercial health app practice I address in this dissertation fall within the scope of the Directive. The concept of 'commercial practice' itself is defined in Article 2(d), and we should also ask how this definition relates to commercial health app practices. The definition provided in Article 2(d) reads as follows:

Business-to-consumer commercial practices (hereinafter also referred to as commercial practices) means any act, omission, course of conduct or representation, commercial communication including advertising and marketing, by a trader, directly connected with the promotion, sale or supply of a product to consumers.

The analysis provided by Willett (2010: 249) on something being covered "as long as there is some commercial flavor" carries over to this formal definition of commercial practices, which mentions "any act, omission, course of conduct or representation, commercial communication including advertising and marketing." There is, however, still a further question to be answered concerning the Directive's scope. The definition in Article 2(d) speaks of commercial practices being "*directly* connected with the promotion, sale or supply of a product to consumers." For single, isolated interactions between a user and a health app – e.g., a piece of native advertising shown to a user – it is clear that it can be considered *directly* related to the "promotion, sale or supply of a product to consumers." But as I have emphasized throughout this dissertation, I do not just want to address single, isolated interactions between a consumer and an app; it is much more important to address the more structural concern of how digital health environments can be designed in such a way as to give a particular commercial shape to the longitudinal relationships with the users of apps. It could be asked whether the structural features of a digital health environment can be considered a commercial practice since it is difficult to relate such structural features *directly* to "the promotion, sale or supply of a product to consumers."

I want to present two arguments here. First of all, I do think there is an argument to be made for including the more structural features of digital health environments into the scope of the Directive. Secondly, even if the first argument succeeds, it is still important to (re)consider the concept of 'commercial practices' as laid down in the UCPD because it is outdated and does not align well with the structure and impact of sophisticated digital environments.

So how can one argue that the structural design features and operation of a digital health environment are, in a relevant sense, directly related to “the promotion, sale or supply of a product to consumers”? The health apps under investigation in this dissertation can, and probably will, try to argue that only the direct offering of products and services to consumers, as well as showing advertising, is truly *directly* related to the “promotion, sale or supply of a product to consumers.” All the structural ‘environmental’ aspects I have discussed are background chatter that is at best *indirectly* related to the commercial interactions with consumers; or so they could argue. Such an answer would be understandable (from a strategic perspective), but also inaccurate. As I have emphasized throughout this dissertation, nearly all popular for-profit health apps offer highly sophisticated digital environments where every part of the environment is carefully optimized. What these apps optimize for is retention, engagement, and conversion, which are all metrics that when optimized should play out in favor of the health app *in the long run*. On its Developer pages, Apple explicitly advises app developers to “prepare for the long term” because the freemium model works best when one is able to engage users for extended periods of time.¹⁵⁷ An app developer quoted by Apple on the Developer page says that “[w]e don’t just look at short-term retention, we look at very, very long-term retention.” Apple also advises that “[s]uccessful freemium apps have analytics built into the experience so that developers can understand user preferences and continually improve the apps.”¹⁵⁸ Put differently, the monetization model is one of (trying to) engineer continuous engagement, which, in turn, makes other monetization options (e.g., selling freemium options, selling more advertising, collecting more user data to monetize) more feasible. So when a popular for-profit health app sends a personalized push notification to, for instance, notify me that a friend liked a particular healthy lifestyle blog post, one could try to argue that there is no *direct* commercial purpose behind the push notification. The app is not trying to sell something directly to you. But such an answer would be overly simplistic, because such push messages are an *integral part* of – and directly related to – the monetization efforts of freemium apps such as most of the popular for-profit health apps I discuss in this dissertation. The engineering of engagement is *the* central pillar of the monetization philosophy in freemium apps (as evidenced by Apple’s own

157 See Apple’s advice page here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKlNID>.

158 See Apple’s advice page here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKlNID>.

advice to app developers¹⁵⁹). So when apps design a digital health environment that continuously keeps track of user interactions in order to analyze how to, for instance, best personalize the experience and target users with content, notifications or offers, I would argue that *all* those efforts are in fact *directly* related to the commercial activities of the app.

This brings me to my second argument. The health app case shows how the concept of ‘commercial practices’ is a somewhat uneasy fit with the current developments in user monetization efforts in online environments. The UCPD appears to mainly focus on clearly identifiable, isolated commercial interactions between users and traders (e.g., Is this particular communication misleading? Did this advertisement improperly target vulnerable consumers?). From an UCPD perspective, the marketplace looks like a large chain, or collection, of discrete commercial practices that are related, but also clearly discernable as individual practices. As my analysis of health apps shows, more structural dynamics in the contemporary economy of freemium services risk escaping the scope of the UCPD. App developers design their digital environments, in the word of Apple, “for the long term;”¹⁶⁰ a web of interlocking design choices which are meant to help develop economically feasible *relationships* with users. Because service providers such as health apps tend to have (almost) complete control over the design and operation of the digital environments they offer, they also have a significant amount of power. One of the most important questions for contemporary unfair commercial practice law should be whether and how this structural power to design choice architectures can and cannot be used to shape people’s self-understanding and possibilities for decision-making. To properly address such questions, a concept of commercial practices that is more sensitive to the ‘structural’ or ‘environmental’ nature of online market interactions is necessary.

Such a more ‘structural’ or ‘environmental’ interpretation of commercial practices also makes sense from the perspective of autonomy, which is the master principle informing unfair commercial practice law. If we accept that autonomy is not just about an individual’s capacities for decision-making but also about the ways

159 See Apple’s advice page here: <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKINID>. Apple writes that “In most cases, providing a great experience to all users regardless of whether they choose to spend is an integral aspect of the freemium model. 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.”

160 <https://developer.apple.com/app-store/freemium-business-model/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/jdKINID>

in which (digital) environments shape our possibilities for developing and practicing autonomy, then it follows that a commitment to autonomy also implies a commitment to a more 'structural' or 'environmental' interpretation of the UCPD.

In sum, the broad interpretation of what constitutes a 'transactional decision' combined with the argument that intentionally designed structural features of a digital health environment can be interpreted as directly related to "the promotion, sale or supply of a product to consumers" lead to the conclusion that many commercial health app practices fall within the scope of the Directive.

In the next section, I discuss one of the UCPD's most important evaluative standards, the average consumer standard.

2.4. The Average Consumer & Vulnerabilities – The Average Health App User as a Vulnerable Consumer?

The average consumer standard serves as an essential evaluative standard when analyzing the unfairness of commercial practices; "The average consumer is the measure of all things" (Micklitz 2006: 12). When analyzing the potential unfairness of a commercial practice, it should be asked whether the practice is unfair *from the perspective of the average consumer*. If a commercial practice is aimed at a particular group, then the average consumer of that particular group serves as the evaluative standard. More specifically, the average consumer standard entails that when considering the unfairness of commercial practices, it should be asked whether the practice leads – or is likely to lead – the average consumer to make a transactional decision or display economic behavior she would not have made or displayed otherwise.

In order to avoid confusion, let me point out at the start of this analysis of the average consumer standard that the standard should not be seen as a separate 'test' or 'step' when analyzing the possible unfairness of commercial practices. The average consumer standard is always already part of every UCPD analysis; the standard serves as a kind of 'interpretational horizon' throughout the Directive. So even though I discuss the standard separately in this section, it should not be seen as a separate part of unfair commercial practice analyses. In the later sections on misleading and aggressive commercial practices, the annex, and the general clause, I will continue to include the average consumer standard in my analysis.

Who, then, is the average consumer? The Directive takes as a benchmark an average consumer "who is reasonably well-informed and reasonably observant and circumspect, taking into account social, cultural and linguistic factors" (Recital

18), which was established by the CJEU in *Gut Springenheiden GmbH*.¹⁶¹ It is emphasized that the “average consumer standard is not a statistical test,” as it is ultimately up to courts to determine what the average consumer standard entails in a particular case (Recital 18).

As Recital 18 makes clear, the concept of the average consumer serves the purpose of not imposing unreasonable burdens on traders: “In line with the principle of proportionality [...] this Directive takes as a benchmark the average consumer.” In its Guidance document on the implementation and application of the UCPD, the European Commission (2016a: 42) underlines the fact that “The UCPD adopted this notion to strike the right balance between the need to protect consumers and the promotion of free trade in an openly competitive market.” Put simply, it would be unreasonable to require traders to actively anticipate the effect their influences might have on *every single possible consumer* – especially given the fact that unpredictable ‘outliers’ are difficult to anticipate. Imagine a trader that advertises rubber ducks on the Internet. This trader is not required to anticipate the possible existence of that one consumer in the entire population who has an irrational fear of rubber ducks, causing this consumer to experience an immediate stress reaction when confronted with the image of a rubber duck and thereby distorting this consumer’s capacities for autonomous decision-making. (If, however, there were a nation state in which a widely feared guerilla faction used a rubber duck as its symbol, then the (mis)use of this widely present fear by rubber duck traders might very well be problematic.)

Although the Directive clearly does not want to impose overly demanding requirements on traders, there is an additional ‘vulnerability’ clause which focuses on specific *groups* of consumers who are “particularly vulnerable to the practice or underlying product because of their mental or physical infirmity, age or credulity in a way which the trader could reasonably be expected to foresee” (Article 5(3)). Even when such groups of “particularly vulnerable” consumers are not the intended audience, a trader should still take the effect a commercial practice can have on such groups into account. In other words, the average consumer standard cannot be used as an excuse to engage in commercial practices that are likely to treat members of vulnerable groups unfairly, *even if the trader never explicitly targeted these vulnerable groups*. The only possible excuse for a trader would be to insist that there was no way she “could reasonably be expected to foresee” the potential impact of a commercial practice on vulnerable groups.

161 C-210/96.

Although the attention for particularly vulnerable groups is welcome, the Directive also betrays a rather narrow understanding of what constitutes vulnerability in the marketplace. The Directive seems to opt for an approach to vulnerability where particular *static* groups are designated as vulnerable on the basis of (*semi-*)*permanent* physical or mental characteristics that prohibit them from acting in a fully autonomous fashion in the marketplace. Put bluntly, the Directive seems to draw a very hard line between, on the one hand, a few *fixed* groups that are vulnerable and warrant special protection (i.e., children, the elderly, and disabled people), and, on the other hand, all the ‘regular’ consumers who are not children, very old, and/or disabled, and are therefore considered average and not particularly vulnerable. In the UCPD, “[t]he average consumer is the rule, whereas the particularly vulnerable consumer is the exception” (Micklitz 2006: 112). One should, however, not make the mistake of thinking that the average consumer is a perfectly autonomous ‘superhuman’ who cannot be misled or influenced adversely. The average consumer is, in the end, only human. Still, from a *legal* perspective there is a sharp distinction between the average consumer and the vulnerable consumer.

Over the years, this supposed hard line between ‘the vulnerable’ and ‘the average’ consumer has increasingly been called into question (e.g., Incardona & Poncibò 2007; Duivenvoorde 2013; Sax, Helberger & Bol 2018.) Advances in, and the popularization of (e.g., Kahneman 2011, Thaler & Sunstein 2008), behavioral sciences have made it clear that the ‘average consumer’ portrayed in the Directive (and European consumer law at large) is more of an ideal type than an empirical reality. Incardona and Poncibò (2007: 35) conclude that:

The average consumer test reflects the economists’ idealistic paradigm of a rational consumer in an effective marketplace. This notion may be useful for economists’ calculations and projections, but departs from the unpredictable realities of human behavior and is hardly an appropriate standard for legislative or juridical sanctions.

The *real* average consumer is one whose behavior is shaped by various cognitive and affective biases, the effects of which are often not entirely transparent to herself. Instead of designating certain groups as *inherently* vulnerable as a result of some (semi-)permanent characteristic all members share (e.g., age), it seems more appropriate to acknowledge the fact that *all* people are at least susceptible – or: dispositionally vulnerable – to manipulative commercial practice; not necessarily always, but at least *sometimes*, under *some* circumstances, or as a result of personalized targeting practices. Every person has a persuasion profile

(Kaptein 2015, Kaptein *et al.* 2015). Consider an able-bodied, educated, employed, upper middle class, heterosexual, white male – the prime example of a privileged and supposedly non-vulnerable consumer. It is possible to conceive of the possibility that he could experience brief moments of temporary, but still serious, vulnerability. For example, a trader could target him specifically at Friday, around 7 pm, when he gets home from an emotionally draining week of hard work, ready to collapse on the couch, all alone because he has not managed to date due to his strenuous schedule, of which he is painfully aware on Friday evenings. If this trader targets these brief, highly specific moments in which our male experiences intense loneliness to push a product related to this loneliness, we might want to say that a vulnerability – however brief and temporary – has been unfairly exploited.¹⁶² Whether this hypothetical example is a good example of a (type of) vulnerability that should be covered by the UCPD remains to be seen.

Recent insights from the field of behavioral economics are not the only reason to argue for a more nuanced understanding of both the average consumer standard and our understanding of vulnerabilities in the legal context. As I discussed in chapter 3, there is also an ethical literature on vulnerability – usually working with a broader notion of vulnerability than found in the legal contexts – which highlights the importance of a refined conceptualization of vulnerability (see Anderson 2014, Mackenzie 2014; 2016, Mackenzie, Rogers & Dodds 2014, Straehle 2017, Jacobs 2020). From an ethical perspective, a proper understanding of (and resulting ability to identify) vulnerabilities is important, because the fact that a person can be considered vulnerable may¹⁶³ give rise to (additional) responsibilities and duties on others to take those vulnerabilities into account, or try to alleviate them. In their taxonomy of vulnerability, Mackenzie, Rogers and Dodds (2014: 7-9) identify different sources of vulnerability and different states of vulnerability, thereby “acknowledg[ing] the ontological vulnerability that is inherent in the human condition while at the same time enabling the identification of context-specific forms of vulnerability” (Mackenzie, Rogers & Dodds 2014: 7-8). As *sources* of vulnerability, they identify 1) inherent vulnerabilities, which are “sources of vulnerability that are intrinsic to the human condition;” and 2) situational sources,

162 Also consider the Facebook example discussed in section 4.3 of chapter 1 and section 4.1.2 of chapter 4 in this regard.

163 I use the word ‘may’ here because there is no consensus on the normative status of vulnerabilities. Some would argue that ‘vulnerability’ is a normative concept that always implies that the vulnerability in question is bad and generates responsibilities (on others, or the community) to cater to the responsibility. Others hold that vulnerability is a descriptive rather than a normative concept which *may*, but does not necessarily generate responsibilities.

“by which we mean vulnerability that is context specific.”¹⁶⁴ These sources can manifest themselves in two different *states*, namely dispositional or occurrent. Put simply, occurrent vulnerabilities are vulnerabilities that, due to either inherent or situational sources (or a combination thereof), have actually materialized. Dispositional vulnerabilities are vulnerabilities that *can* materialize (i.e., become occurrent) because a person either has some inherent features or finds herself in a particular situation. So “the dispositional-occurrent distinction refers to states of potential versus actual vulnerability” (Mackenzie, Rogers & Dodds 2014: 8).

Thus from an ethical perspective, there are also good reasons to opt for a more nuanced understanding of vulnerability. The taxonomy by Mackenzie, Rogers and Dodds highlights, correctly in my view, how people are of course always vulnerable already but also, more importantly, can constantly move in and out of *states* of vulnerability due to a constant interaction between their personal (semi-static) features and the environments within which they – either temporarily or permanently – reside. For now, it suffices to say that the strict conceptual distinction between ‘the average consumer’ and ‘the vulnerable consumer’ is difficult to maintain, either on empirical or normative grounds. But even if we agree that the current sharp distinction between the average consumer and the vulnerable consumer is problematic, we can still disagree about the available strategies to address this challenge. One option, which feels intuitively most plausible to me as an ethicist, is to insist that we simply need to *broaden* the legal concept of vulnerability, for example along the lines of my discussion of the ethical concept of vulnerability. The current legal concept of vulnerability is, however, well-established, so it might be difficult to propose a complete overhaul. Another option is to direct most of our efforts to nuancing the concept of the average consumer and to try to make the average consumer more human. According to this proposal, the legal concept of vulnerability would retain (most of) its current meaning, while there would be more attention for the various ways in which the average consumer is susceptible to manipulative commercial practices, without this average consumer necessarily being *vulnerable*.

When we return to the actual text of the Directive, we can find some interpretative room for a broader understanding of the concept of vulnerability. The Guidance document of the European Commission (2016a: 47) points out that although Article 5(3) “appears to exclusively qualify consumers as vulnerable because of their ‘*mental or physical infirmity, age or credulity*’, Recital 19 provides a non-

¹⁶⁴ Mackenzie, Rogers and Dodds (2014: 9) also identify *pathogenic* vulnerabilities, which are “a subset of situational vulnerabilities that are particularly ethically troubling.”

exhaustive list of characteristics that make a consumer ‘particularly vulnerable.’” Recital 19 refers to “certain characteristics *such as* age, physical or mental infirmity or credulity” that make a consumer “*particularly susceptible* to a commercial practice” [emphases mine]. The fact that the Commission emphasizes the open-ended formulation of Recital 19 over the narrow, closed formulation of Article 5(3) is telling; clearly the Commission did not intend for the vulnerability clause to be as narrow as the language of Article 5(3) might suggest. Moreover, in its Guidance document the Commission explicitly mentions, and quotes from, its own study *Consumer Vulnerability across Key Markets in the European Union* (2016b). This study also concludes that – especially in the face of recent insights from behavioral sciences – overly static conceptions of vulnerability will not provide adequate levels of consumer protection, since “[c]onsumers may move in and out of states of vulnerability and they may be vulnerable in respect to some categories of transactions but not others [...] vulnerability is best seen as a spectrum rather than a binary state” (European Commission 2016b: xvii-xviii). Opting for a broad rather than a narrow understanding of vulnerability thus seems to be very much in line with the (policy) goals of the European Commission, as well as with its preferred interpretation of the UCPD.

Let us take stock for a moment. So far, I have used this section to argue that 1) the average consumer standard and the notion of vulnerability together function as an important ‘interpretational horizon’ in the UCPD; 2) there are both empirical and normative reasons to adopt a broader, more fluid understanding of the average consumer and vulnerability; and 3) there is some interpretational room already present in the Directive to adopt this broader approach to the average consumer and vulnerability. Having established all this, the next important question to answer is how the average consumer standard and the concept of vulnerability should be understood *in the health app context*.

The most obvious defining feature of health apps is of course that they are organized around the value of *health*. As I have argued before, health is, in a rather literal sense, a central concern in every person’s life – every person needs a certain degree of health to continue living. Moreover, it is safe to say that being healthy is an important precondition for living a happy, fulfilling life.¹⁶⁵ So as a start, it makes

165 We can of course think of cases where being healthy is not a *necessary* precondition for living a happy, fulfilling life. Think, for instance, of a person living with a serious chronic disease. It would be absurd to say that such a person would be incapable of living a good life. It does seem plausible to me, however, to say that *other things being equal*, people will generally prefer being healthy to being ill because it makes for a more pleasant life. See the debate between Boorse (1975, 1977) and Nordenfelt (1986, 1987) on the normative status of health and its relation to well-being.

intuitive sense to say that the average consumer of health apps has a significant interest in her health. However, the fact that health is an important *precondition* for living a good life does not necessarily imply that people are therefore also necessarily actively concerned about their health, or actively involved in managing their health. One could argue that as long as a sufficient basic level of health is achieved, the *concern* for health could cease being a central concern in a person's life – when a person already has what she needs (i.e., health) she does not need to worry about it.

As I argued in chapter 2, such an argument is overly formalistic and neglects the ways in which health figures in our culture and society. Health has real social significance (Crawford 2006). Being healthy affords social status, whereas being unhealthy or living (what is perceived to be) an unhealthy life are seen as socially unacceptable. For example, Lupton (2018) shows how there still exists serious social stigma around 'being fat,' which is, for instance, seen as a symptom of character flaws and as being a burden to society. Given health's social significance, it is also important to understand *how* the imperative of health is given shape in contemporary culture. What 'counts' as health, how should people understand themselves in relation to their health, and how should they be preoccupied with their health?

As my discussion (see chapter 2) of the contemporary health discourse – which is eagerly reproduced in health apps – shows, health is not just understood as the absence of (serious) illness. Nowadays, health is portrayed as something that should be optimized (Devisch 2013). Even when you are already healthy by most standards, your health is still something that should be continuously monitored and maintained. Health is not something you achieve, but something you *do*. One is supposed to adopt a 'healthy lifestyle,' regardless of whether one is on its way to become healthy or is already considered to be (very) healthy. Cederström and Spicer (2015: 4-8) as well as Lupton (2013: 397-398) also observe a *responsibilization*¹⁶⁶ in the health context, fueled partly by the wide availability of health apps that 'give people the tools needed' to deal with their health responsibly.

166 There is a larger literature on health responsibilization. Rossiter (2012: 193), for example, argues that "contemporary public health messages, which are couched in the language of risk and statistical fact, in fact operate as morality tales that serve a political purpose by individualizing health responsibility and turning attention away from systemic health problems. Here, I have argued that public health messages become effectual through the mechanism of predictive fantasy, and that these fantasies have profound effects on our bodily practices and routines. Specifically, I argue that these bodily practices are rooted in a kind of anxiety that occurs when taking personal responsibility for health is indicative of a moral orientation towards health. Health status, I have argued, is represented by public health discourse not as an indication of what kind of a life an individual has chosen to live, but rather as a sign of one's moral proclivities." Other examples are Shamir 2008, Gray 2009, Brown & Baker 2012, and Sparke 2016.

It is precisely against this background of a culture of health optimization and responsabilization, as well as the social status implications of health, that we should understand the ‘average health app consumer.’ The average health app consumer lives in a society where health is very much a part of her life, and where health demands are real and omnipresent. We thus live in societies where health is made an important concern for most people. Because the contemporary health culture mandates *everyone* to be preoccupied with their health, it also follows that everyone, in principle, could be vulnerable to being manipulated on the basis of health concerns. The traditional hard distinction between the average consumer on the one hand, and the vulnerable consumer on the other, breaks down here. It no longer makes sense to point to some fixed set of characteristics that makes a fixed subset of citizens vulnerable to unfair commercial practices in the digital health context. Under the traditional distinction (which requires reconsideration), one would probably divide the digital health world into two camps: on one side, we would find children, the elderly, and everyone with serious health issues (i.e., the vulnerable), while the other side would be made up of all the able-bodied adults under the age of 80 (i.e., the non-vulnerable). The first camp would then be seen as at risk of being manipulated by health apps, whereas the second camp of ‘normal,’ ‘average’ people would not be considered to be particularly susceptible to manipulation. This picture, as I have argued, does not allow us to properly understand vulnerability in the digital health context. *Every* person is potentially vulnerable to manipulative commercial practices in the digital health context because health is an important concern to *everyone*. The digital health environments provided by popular for-profit health apps are built to optimize retention, engagement, and conversion, as a result of which there is a *structural* need to identify health-related characteristics, desires, and concerns, and to construct persuasion profiles (Kaptein 2015, Kaptein *et al.* 2015). Consider the young, successful, able-bodied, highly intelligent, affluent student or young professional using MyFitnessPal on a daily basis. The health discourse that is (re)produced and propagated by MyFitnessPal is very much catered to people like her. She is, as much as any user, vulnerable to manipulation by an environment that seeks to ‘optimize’ its interactions with her to keep her coming back to the app, so she engages with the content in the app as much as possible. As the earlier discussed studies by Culbert, Racine and Klump (2015), Eikey and Reddy (2017), and Eikey *et al.* (2017) show, our hypothetical MyFitnessPal user is maybe even *especially* vulnerable to have her self-understanding and behavior influenced (or manipulated) in ways that undermine rather than scaffold her autonomy.

It is therefore important to follow the European Commission’s (2016b) lead and embrace a more *dynamic* understanding of vulnerability, which acknowledges empirical insights from behavioral economics and normative insights from

philosophy. Put simply, people are not either vulnerable or not vulnerable, full stop. Rather, different people can be vulnerable under different conditions, or at different times. People can experience short episodes of vulnerability (e.g., for years, months, days, or even hours), and these episodes might only occur under particular conditions (e.g., when being tired, when visiting a particular place, when being targeted in a particular manner). This is *especially* the case in the health app context where *every* person has a fundamental interest in their own health and where the popular for-profit health apps have *structural* reasons to identify, target, and exploit every user's feelings, desires, and concerns regarding their own health for reasons of optimization.

The upshot of this argument then is not that every consumer should always be considered vulnerable. Rather, my argument is that for every single consumer who regularly inhabits a digital health environment, it is very much a serious *open question* whether and when she can and should be considered vulnerable. Acknowledging the *inherent potential* for every consumer to be rendered (temporarily) vulnerable in the digital health context is essential to make the UCPD best serve its purpose of creating fair market environments that respect people's autonomy.

In the next two sections, I deal with the concepts of misleading commercial practices and aggressive commercial practices; arguably the two most important concepts in the UCPD to evaluate whether a commercial practice should be considered unfair or not. The average consumer standard as well as the critical remarks on vulnerability should be kept in mind as an important 'interpretational horizon' when reading the following sections.

2.5. Misleading Commercial Practices

This section will proceed in three steps. First, I will introduce the UCPD's Articles on misleading commercial practices and discuss to what extent they can capture some of the ethical tensions I have outlined in chapter 4. Second, I will discuss the usefulness of the UCPD's concept of misleadingness when it comes to addressing manipulation worries. Third and last, I will discuss how my ethical analysis can inform the concept of misleadingness to allow for an interpretation of the concept that is better able to address manipulation worries in the digital society.

As I will argue below, the concept of misleading practices can be readily applied to some questionable, rather specific, health app practices (i.e., native advertising in the health context and lack of transparency concerning commercial intent). As a whole, however, the category of misleading commercial practices as it is found

in the UCPD will do little to *directly* address the more structural questions on the power relations between health apps and their users and the manipulative potential of digital health environments. That does not mean, however, that the concept of misleading commercial practices should not be explored further. Manipulative practices, as I have defined them in chapters 3 and 4, are often misleading in nature, which means that there is, at a theoretical level at least, a natural connection between manipulation and misleading commercial practices. It is no accident that many of the philosophers that have tried to conceptualize manipulation have considered the question of whether manipulation is necessarily deceptive (e.g., Rudinow 1978, Noggle 1996, Baron 2003, Greenspan 2003, Cohen 2018, Susser, Roessler & Nissenbaum 2019a, 2019b). Noggle (1996: 44 [emphasis added]), for instance, writes that one core intuition we have when analyzing cases of alleged manipulation is that “in many of them the victim has been *led astray* in some way.”

Let me first explain the UCPD’s concept of misleading commercial practices before I analyze its applicability, usefulness and shortcomings, and room for improvement. The Directive specifies what constitutes a misleading commercial practice in Article 6, which deals with misleading actions, and Article 7, which deals with misleading omissions. I start by providing the two Articles, which I then use as a basis to discuss the Directive’s stance on misleading practices. Article 6(1) reads as follows:

A commercial practice shall be regarded as misleading if it contains false information and is therefore untruthful or in any way, including overall presentation, deceives or is likely to deceive the average consumer, even if the information is factually correct, in relation to one of the following elements, and in either case causes or is likely to cause him to make a transactional decision that he would not have taken otherwise.

Article 6(1)(a-g) then goes on to list “elements” that should be used to consider whether information, either in content or presentation, can mislead consumers: the nature of the product (a); the main characteristics of the product (b); the extent of the trader’s commitment, the motives of the commercial practice and the nature of the sales process (c); the price or the manner in which the price is calculated (d); the need for a service, part, replacement or repair (e); the nature, attributes and rights of the trader or his agent (f); and the consumer’s rights (g).

Article 7(1) explains what constitutes a misleading omission:

A commercial practice shall be regarded as misleading if, in its factual context, taking account of all its features and circumstances and the limitations of the communication medium, it omits material information that the average consumer needs, according to the context, to take an informed transactional decision and thereby causes or is likely to cause the average consumer to take a transactional decision that he would not have taken otherwise.

Articles 7(2-5) provide further specifications that are not immediately relevant to my argument so I will not discuss them more here.

As both Articles 6 and 7 make clear, the Directive employs a heavy focus on the types of *information* a consumer needs to make an informed decision, and the ways in which *information* is communicated to consumers to establish the misleading nature of commercial practices. This should be kept in mind when reading the following discussion on the interpretation of the Directive's clauses on misleadingness.

Much in the same spirit of the Directive's overall wide scope, the clauses on what constitutes a misleading (and therefore unfair and prohibited) commercial practice should also be interpreted in a broad manner. Let me start with the 'materiality condition.' For a commercial practice to qualify as being misleading it is not enough to consider the content and presentation of relevant information in isolation. For a practice to be unfair and prohibited, it must cause, or must be likely to cause, the consumer to make a transactional decision she would not have made without the (allegedly) misleading influence. The phrase 'likely to cause' is highly significant here, since it means that no *actual* distorted transactional decision is required to have taken place. Moreover, Wilhelmsson (2006: 137-138) speaks of a "relaxed attitude towards the proving of materiality. No empirical proof concerning the influence of certain information on the actual transactional decisions of consumers is required." So for the materiality condition to be met, a judge only has to believe that a commercial practice *may* mislead consumers. The 'relaxed attitude' is even more relaxed due to the fact that the term 'transactional decision' has a rather broad meaning. As Article 2(k) makes clear, 'transactional decision' refers not just to the actual decision to buy a product or service, but also to, for example, questions of payment, and the question of whether to retain or dispose of a product or service. In the section on the scope of the Directive, I have already argued that given the freemium business models of most popular for-profit health apps

and, most importantly, the resulting intentionally designed features of the digital health environments, nearly all user-app interactions in these health apps can be considered transactional decisions. All in all, the materiality condition can be considered a “low hurdle” (Wilhelmsson 2006: 137).

Another important feature of the clauses on misleadingness is the fact that the Directive does not refer to the actual intentions of traders. The possible misleadingness of a commercial practice is determined by analyzing whether the practice does, or is likely to have, a significant influence on the consumer’s economic behavior (in the broad sense explained above). There is no need to determine whether a trader proceeded from a plan to mislead, or failed to take necessary precautions to avoid misleading consumers. It is all about the (likely) effects.

As I mentioned above, Articles 6 and 7 are really focused on the role information plays in the consumer’s economic behavior. Information does not have to be incorrect or even untruthful to be (potentially) misleading. Article 6(1) explicitly mentions the importance of considering the “overall presentation” of a commercial practice, since a collection of truthful pieces of information can, when combined in a specific manner, still be misleading. “Half-truths, or combinations of given facts with information being omitted, may result in a picture that is misleading, even though one cannot point to any single concrete piece of information that would actually be false” (Wilhelmsson 2006: 130). ‘Overall presentation’ should also include the more ‘visual elements’ of presenting information. As the nudge movement popularized by Thaler and Sunstein (2008) has taught us, the actual visual organization and presentation of information is just as important as the actual *content* of the information when it comes to influencing the behavior of people. The explicit reference to the ‘overall presentation’ of information thus also seems to cover the role of, for instance, UX design (and its “dark patterns” (Gray *et al.* 2018)) in influencing consumers’ behavior when considering the unfairness of a commercial practice.

Concerning misleading *omissions*, I would like to highlight parts of Article 7(2) that are especially relevant in the context of for-profit health apps. In this further specification of Article 7(1), we read that “it shall also be regarded as a misleading omission when [...] a trader [...] fails to identify the commercial intent of the commercial practice if not already apparent from the context.” To see the potential significance of this phrasing, remember that what constitutes a commercial practice is defined very broadly: I already established that as long as there is “some commercial flavor” it is probably covered by the Directive.

2.5.1. How Applicable Is the Current Concept of Misleading Commercial Practices to the Health App Context?

Based on my analysis in the previous chapters, there are two commercial health app practices that I want to address as being ‘misleading’ in the Directive’s current sense of the word. First, there is the issue of poorly labeled (or even completely unlabeled) native advertising which, especially in the *health context*, can be considered problematically misleading. Second, there is the slightly more abstract issue of health app providers’ (mis)characterizations of the functioning of, and the commercial intent behind, the digital health environments they provide to users. Popular for-profit health apps tend to emphasize that they are free services, while de-emphasizing the *economic* nature of the relationship between the user and the app. This can be considered misleading.

2.5.2. Native Advertising in Health Apps

So let us first look at native advertising. In the previous chapters I have discussed various examples of native advertising in health apps. The very reason for native advertising to ‘blend’ a sponsored message into the non-advertising content, by matching its look and feel, is to make sure the advertising does not, in fact, look like advertising. So in essence, native advertising is designed to mislead. That does not mean, however, that native advertising is always forbidden. Nearly all advertising is – to some extent – misleading in its attempt to lure, persuade, or tempt consumers to consider a product. We can and should expect from consumers that they know and understand that there are some misleading qualities to (nearly all) advertising. Here the importance of the average consumer standard reveals itself. One important line that can be drawn between advertising that should be considered ‘business as usual’ or ‘part of the game’ and illegitimate instances of advertising is the consumer’s ability to recognize that she is being served a piece of advertising in the first place. So for native advertising to be fair game, it should at the very least be clear to consumers that they are looking at advertising. This requirement, of course, introduces a tension, since the idea behind native advertising is to make the advertising look like organic content. In practice, the key question becomes whether a disclosure accompanying a piece of native advertising is good *enough*.

To start, there are easy cases where native advertising is not labeled at all. For example, the earlier discussed blog post called ‘A Day in the Life of a Yoga Teacher’¹⁶⁷ was featured in the main feed of the MyFitnessPal app that is, by default, shown to all users of the app. The post clearly intends to look like regular editorial content

¹⁶⁷ <https://blog.myfitnesspal.com/day-life-yoga-teacher/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/KijSKaz>

(i.e., an informative piece about what it is like to be yoga teacher) while blending in advertising messages for skin care products. Because there is a complete absence of labels such as ‘advertising’ or ‘advertorial’ and no mention of the commercial intent behind the message, it is clearly misleading. Article 7 on misleading omissions could apply here, since the complete absence of any information which communicates the commercial nature of the advertorial is likely to have a material impact on the economic decision-making of the health app user.¹⁶⁸

In most cases, however, native advertising in health apps contains at least some labels or disclosures to flag the commercial nature of the communication. Some interpretative work needs to be done before such cases can be judged as either native advertising that is misleading and therefore prohibited, or as native advertising that is not misleading (in the UCPD sense of the word). In these cases the ‘good enough’ question presents itself. To answer that question for the context of health apps, I want to draw on my analyses from the previous chapters. In short, I would argue that most native advertising in health apps should qualify as being misleading. A few considerations are important here.¹⁶⁹ First of all, remember that the average consumer benchmark is used to judge whether the piece of native advertising is likely to distort economic decision-making. The fact that native advertising takes place *in a health app* is significant from the perspective of the average consumer benchmark. Second, it should be remembered that we are, ultimately, concerned with the consumer’s ability to make autonomous decisions. In the context of native advertising, I want to propose that we draw on a more relational notion of autonomy (as developed in chapters 3 and 4) to properly assess the potentially misleading nature of native advertising.

The average consumer is notoriously bad at recognizing native advertising (Wojdyski *et al.* 2017, Hyman *et al.* 2018). Moreover, I have already argued that health is an important concern to the average consumer. So when native advertising

168 Alternatively, item 11 of the blacklist, which directly addresses native advertising, possibly applies here.

169 When considering whether native advertising in health apps should be (more stringently) regulated, it should also be emphasized – again – that the UCPD has an especially wide material scope. As discussed earlier, it is not just the actual spending of money which is to be counted as ‘displaying economic behavior’ in the freemium app context. When a user reads a piece of native advertising (such as the A Day in the Life of a Yoga Teacher advertisement by MyFitnessPal) without purchasing the advertised product, the user’s engagement with the native advertising can still be economically beneficial to the app provider (engagement metrics are a core feature of the freemium app economy). So the mere fact that a health app provider manages to get a user to engage with a piece of native advertising within a health app could be considered a display of economic behavior.

is integrated into content on a topic – i.e., health – that is of great importance to the (average) consumer, we should be especially vigilant because there is a greater potential for manipulation. To see why, we should not just focus on the moment a piece of native advertising is served to a user. We should also, again, take the process of relationship building between the app and the user into account. For-profit health apps will often seek to instill trust in users to build lasting relationships with their users, and these relationships (and the relationship-building efforts) are often framed in terms of the importance of health, and, more specifically, the importance of the user's health. It is against this background that we should consider the question of native advertising in health apps and its misleadingness in relation to the user's autonomy. A significant portion of users will encounter native advertising after they have already spent a considerable amount of time 'with' the app in question and with which they have (sometimes) developed a trust-like relationship. Given autonomy's relational nature, we should consider the autonomy of the user's behavior vis-à-vis the native advertising within the larger context of the relationships that develop between an app and its users over time. So instead of just asking how an isolated piece of native advertising may influence a random health app user, we should ask how the deliberate relationship-building efforts influence the users' ability to practice their autonomy. The average health app user that encounters a piece of native advertising is often *already* involved in an ongoing trust-like relationship with the app. When an app intentionally designs its digital health environment to develop ongoing trust-like relationships with its users while insisting on the importance of health, the app in question can and should expect that its users may be more inclined to engage with content that is *presented to them* as being genuine health content. Offering native advertising to users under such circumstances seems especially misleading, because the relationships that the app tries to build with the users are not set up to help the user recognize and critically reflect on the native advertising offered to them. Quite to the contrary, in fact; the relationship-building efforts of the apps undermine rather than enhance the users' ability to critically question the commercial nature of what is presented to them.

Very clear and explicit labels can go *some* way to help make the advertising less misleading. But considering the importance of the average consumer benchmark as well as the relational nature of autonomy, we should be especially vigilant vis-à-vis native advertising in health apps. Labels can barely be clear and explicit enough in health apps, and in this context regulators and judges should be very critical of native advertising labeling practices that try to meet disclosure duties in a minimally obtrusive manner. Moreover, the same research that finds that people do a very poor job of recognizing native advertising also tends to find that disclosures have a rather moderate effect, with some types of disclosures obviously

producing better results in terms of consumer recognition than others (see, e.g., Wojdnyski & Evans 2016, Boerman & Van Reijmersdal 2016, Hyman *et al.* 2017). In sum, progress can be made in mandating the use of better native advertising disclosures, but one should not expect miracles. Transparency measures in the form of disclosures are not a silver bullet against misleading native advertising practices.

2.5.3. *Health Apps as Free Services & Lack of Transparency Concerning Commercial Intent*

Nearly all the popular for-profit health apps are promoted as being free to use, with the possibility of buying premium features if the user desires to do so. For example, the homepage of MyFitnessPal greets people with a distinctive blue “Start for free” button¹⁷⁰ and Headspace’s homepage contains a distinctive orange button saying “Sign up for free.”¹⁷¹

Such health apps thus effectively communicate – and often eagerly emphasize – that a user can use the app completely free of charge (as long as the user does not purchase access to premium features, which is not required). In reality, the business models of health apps are of course more complicated, as I discussed in chapter 1. As, for instance, the many job openings at health apps show, the constant and dynamic engineering of the digital environment to optimize user retention, user engagement and user conversion is a key objective of most for-profit freemium health apps. The more users can be retained for longer periods of time, the more users can be served with advertising, the more user data can be collected (either to resell or to further optimize the service), and the more users can be made into ambassadors for the app to their (social media) friends. In short: users are not just useful parts of the health app economy when they spend actual dollars or euros – the mere fact that they keep engaging with a ‘free’ app can also make them valuable to the app provider.

It follows that when health apps engage in practices that are aimed at building profitable relationships with their userbase, these practices have a clear commercial intent *even if they are not aimed at directly selling goods or services to the users*. Consider the following examples. The adoption and propagation of a health-as-wellness discourse which emphasizes the need to always optimize and curate one’s lifestyle – ideally with the help of health apps – serves the clear commercial purpose of convincing people to use *and continue to use* health apps. The constant search

170 <https://www.myfitnesspal.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/1vaMPOr>

171 <https://www.headspace.com/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/Q3zID37>

by data scientists and UX designers to either change or personalize the ‘look and feel’ of (features of) the app to optimize the metrics of engagement, retention, and optimization, is aimed at creating a digital environment that nurtures relationships with users in a way that is as profitable as possible. Yet at the same time, these health apps often clearly and intentionally emphasize to users that they are *free* services without drawing *any* explicit/specific attention to the fact that user and app are – from the start – involved in an *economic* relationship.

As the discussion of the UCPD made clear, the failure to communicate the *commercial intent* behind one’s commercial practices in a transparent, understandable and honest manner violates the Directive because it is considered to be misleading.¹⁷² Put simply, promising users a free health app and, after installation, providing them with a digital health experience that never explains or communicates how all the interactions with the app are part of a commercial strategy to monetize the users’ attention, data, and social connections clearly constitutes a failure to inform users about commercial intent. The Italian consumer authority has already made such an argument when it fined Facebook in November 2018 for promoting itself as a free service while not informing consumers about the commercial use of the data it collects on its users, effectively rendering Facebook a service where users ‘pay with data’ rather than with money.¹⁷³

In this context it is also important to acknowledge that it is of course not accidental that health apps do not communicate to users the value to their business of ‘capturing’ these users in their digital health universes. Ignorance of commercial intent makes it easier to pursue the desired economic outcomes. Studies on native advertising show time and again that consumers are really bad at recognizing native advertising, which explains the popularity of the practice (see, e.g., Hyman *et al.* 2018). Forcing health apps to be more transparent about their commercial intentions behind various elements and practices of their apps could go some way towards leveling the informational playing field (see, e.g., Wojdyski *et al.* 2017). However, more transparency and information alone is obviously not enough. Since manipulation typically works through the targeting of known cognitive or affective vulnerabilities that people often cannot help but have, mere knowledge of the commercial intent to leverage those vulnerabilities will not always be enough to render the manipulative influence ineffective.

¹⁷² As we will see later in this chapter, item 20 of Annex I (the blacklist) maybe also applies here.

¹⁷³ Decision PS11112. See also: <https://www.theguardian.com/technology/2018/dec/07/italian-regulator-fines-facebook-89m-for-misleading-users> (last accessed 22 September 2020).

One defining feature of my conception of manipulation is that manipulators (attempt to) manipulate *intentionally*. From an ethical perspective this makes sense, for it helps capture the *wrongness* of manipulation: there is something deeply insulting about people deliberately trying to make other people serve their ends by sneakily finding out how to push their buttons. One practical worry, however, that emerges as a result of the intentionality requirement is the fact that it also seems to introduce a high burden of proof; for manipulation to occur, it has to be made plausible that someone *intended* the manipulation. Especially in the legal context, this can be a complicating factor. As I have demonstrated in the first part of this chapter, however, a commercial practice can be judged to be misleading irrespective of the intentions of the trader.¹⁷⁴ What matters is the (likely) material impact the practice has on the consumer, not what (likely) impact the trader *intended* to have.

So far I have considered the potential misleadingness of offering a for-profit health app as if it is a completely free service. There is, however, another potential source of misleadingness; namely presenting a health app to consumers as being a purely health-oriented¹⁷⁵ service while obfuscating the profit-driven nature of the app. Here, the potential to mislead consumers originates in the allusions to (the importance of) *health*. In chapter 2 I discussed the importance of health in our contemporary culture. Health affords social status (Crawford 2006, Lupton 2018) and should constantly be achieved and performed (Devisch 2013, Lupton 2013, Cederström & Spicer 2015). The result is a culture of health optimization and responsabilization, which creates a social pressure to be preoccupied with one's health (Shamir 2008, Gray 2009, Brown & Baker 2012, Rossiter 2012, Sparke 2016). Health apps are a part of this culture of health optimization and responsabilization because they provide consumers with easily accessible tools to manage (some aspects of) their health. Unsurprisingly, health apps eagerly emphasize how their sole purpose is to support the user's health goals without mentioning or explaining how the relationships they seek to build with users are profit-oriented. To be sure, (nearly) all types of for-profit apps offer their services to consumers without *explicitly* emphasizing or communicating their commercial intentions. One could thus argue that the average consumer should be aware of the fact that apps may have a commercial intent. But given the societal context I just sketched – i.e., the culture of health optimization and responsabilization – the exclusive focus on health promises by health apps can be considered especially misleading. Health app providers happen to work in a digital market that deals in something – health –

174 Remember *CHS Tour Services*, C-435/11, where a trader had met all the requirements of professional diligence but was still judged to have engaged in an unfair commercial practice.

175 Health in the wide sense of the word, also including 'wellness and lifestyle' (see chapter 2).

that is very important to people. When their health is on the line or in question, consumers may be more willing to have their (economic) behavior steered in ways that *seem* to be aimed at supporting their health, but that in reality are (also) motivated by commercial goals. In sum, precisely because health is so important to everyone, even the average consumer may be more likely to be misled by health apps that hide their commercial intent behind enthusiastic health-related promises.

2.5.4. *Towards a More Ambitious Interpretation of the Concept of Misleading Commercial Practices*

At the start of this section, I already indicated that *on a theoretical level*, there is a rather natural, straightforward relation between manipulation concerns and the UCPD's focus on misleading commercial practices. As the analysis above proves, it does not require too much of an interpretational stretch to indicate how some aspects of health apps that can contribute to their manipulative potential can be considered as misleading and therefore unfair. Still, at a *practical level*, we can question whether this natural connection leads to useful solutions. One central aspect of the Articles on misleading commercial practices is their focus on the provision of information and transparency. So when we identify a commercial practice that is misleading, the natural solution to alleviate the misleading nature of the practice is to explore different or better ways to provide information to consumers. This, however, results in a somewhat underwhelming repertoire of solutions; one of the key insights informing this entire dissertation is the fact that cleverly designed choice architectures can shape behavior in increasingly subtle yet persuasive ways that are not easily countered by simply providing more or different information.¹⁷⁶

A more ambitious interpretation of what it *should* mean for a commercial practice to be considered misleading is needed. If we take the manipulative potential of adaptive digital choice architectures such as health apps seriously, we must look

¹⁷⁶ The attentive reader may interject at this point and say that one of the key elements of manipulation discussed in chapter 3 is the manipulator's attempt to keep the manipulation hidden. More transparency thus sounds like a natural – and potentially effective – solution to disarm manipulative practices. However, another key element of manipulation is the manipulator's targeting of weaknesses that can be exploited to steer the manipulee's behavior in the desired direction. We need to take this additional key element into account when discussing transparency measures. When a manipulator does a good job at identifying exploitable weaknesses, the targeting of those weaknesses can be rather efficacious *even if the attempt is (partly) transparent* to the user. Transparency can certainly be part of a solution against manipulative practices, but by itself it is rarely an effective solution when the underlying (and maybe even more pernicious) practice of systematically identifying and targeting weaknesses remains unaddressed.

beyond simple transparency requirements concerning native advertising. The theory of manipulation I developed in chapter 3 (following Susser, Roessler, and Nissenbaum 2019a, 2019b) can be used to broaden the scope of the concept of misleadingness. Instead of just focusing on the question of how consumers should be *informed* about, for instance, native advertising, it makes more sense to focus on the question whether – and if so, when – digital health environments should be considered misleading *as such*. Put simply, it can be argued that offering digital (health) environments that enable manipulative practices should be understood to be misleading. In the end, the very design of manipulative digital choice architectures should be seen as misleading, not the (lack of) information about them.

The four cumulative criteria for an act or practice to be considered manipulative (see chapter 3) can offer guidance. Remember, 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.) I will not consider the question of intentionality here, because the UCPD is not so much concerned with intentionality as it is with the *effects* of commercial practices. Focusing on the last three requirements, manipulation is about the (intentional) steering of a person's goals and ends towards those of the manipulator by way of targeting weaknesses in a manner that is not meant to be transparent to the user. Manipulative digital health environments are thus misleading in nature, since they seek to exert influences on people that are not meant to be (fully) transparent to them.¹⁷⁷ So to the extent that a digital (health) environment is designed in a manner that introduces a serious manipulative potential, we can also claim that the digital environment is *designed to be misleading*.

Still, even a broader interpretation – inspired by my ethical analysis – of what it should mean for digital health environments to be considered misleading, must, in the end, fit the logic and language of the Directive in order to be helpful from

¹⁷⁷ In chapter 3 I argued that manipulators will always *attempt* to influence their targets behind their backs, but that we can also think of scenarios where the targets find out or understand the manipulative influences exerted on them only to still be influenced in the way desired by the manipulator. For this reason, I cannot argue that manipulation is always necessarily misleading. I can, however, argue that manipulation is always *intended* to be misleading and often ends up misleading manipulees.

a *legal* perspective. So for the argument that the overall design of a digital health environment is misleading to work *in the context of the UCPD*, we still need to take notice of some key elements of the Articles on misleading practices. This can of course introduce a tension between, on the one hand, my more ambitious suggestions inspired by ethics, and, on the other hand, the limited ‘interpretative room’ provided to us by the language of the Directive. There is no way – and no need – to ‘resolve’ this tension. It is precisely the point of my analysis to explore the boundaries of the room for interpretation when we bring an ethical perspective to the UCPD. The following two key aspects of the UCPD need to be considered when analyzing my proposal for a broader interpretation of what constitutes a misleading practice.

First, the *overall design* of a digital health environment must be characterized as a commercial practice, for the Directive focuses on commercial practices. Building on my argument in section 2.3 of this chapter, there are good reasons to use a much broader notion of ‘commercial practices.’ As I argued in that section, monetization in freemium apps (and digital choice environments more generally) happens through a holistic approach to digital choice environments. Building ongoing relationships with users is the path to monetization and so digital choice environments are geared towards that goal. For example, the continuous collection of user data is an essential part of most for-profit health apps, because it allows them to learn about their users and adjust the digital choice environment accordingly. So within this context, virtually *all* user behavior that takes place within digital choice environments can be considered *economic* behavior, for the simple fact that the tracking of all user behavior is an essential element of the business practices of for-profit health apps.

Second, Articles 6 and 7 focus (almost) exclusively on the provision and/or omission of *information*. So to the extent that the overall design of a digital environment seems misleading from an ethical perspective *and* can be seen as a commercial practice, we still need to somehow tie the alleged misleading design of a digital environment to the provision and/or omission of information. Without a clear link to the provision and/or omission of information, my proposal will not fit the scope of the Directive.

With these two key elements in minds, let us get back to the question of how an ethical perspective can inform our understanding of misleading commercial practices. If we focus on the very design of digital (health) environments to determine their potential misleadingness, it also follows that protecting consumers against misleading commercial practices must start by addressing the very

constitution of the digital environments that are designed to (sometimes) mislead them. We can use the separate cumulative criteria for manipulation to direct our attention towards those aspects of digital choice architectures that can and should be addressed to make them less misleading. Digital health environments offered by health apps are misleading if 1) they place their own commercial interests above the interests of the user, by 2) targeting the user's exploitable characteristics (which are often inferred by collecting and analyzing user data), 3) without making this transparent to the user. It bears emphasizing once again that it is the *collective* presence of these elements that establishes the manipulative nature of the digital environment and which can have a misleading effect on the consumer. So the mere fact that an app provider prioritizes its own ends above the user's ends does not constitute manipulation with a (potentially) misleading influence on consumers. Similarly, the collection of user data to infer exploitable characteristics is, by itself, not a manipulative practice with the potential to mislead consumers.

If we follow the three elements of manipulation directly above, we can try to further specify how the misleading nature of manipulative digital health environments can be addressed. To start, it does not make much sense to spend time discussing the first element – i.e., the fact that for-profit health apps place their own commercial interests above the interests of their users. All profit-driven vendors pursue their own interests and there is nothing inherently wrong with this. Moreover, the goal of the UCPD is not to outlaw commercial practices, but to make sure that commercial practices are fair. If we focus on the collection of user data to find exploitable characteristics of users, we could try to argue that health apps should *refrain* from collecting user data for the purpose of finding exploitable characteristics that can be used to target users in order to dispose them to the ends of the health app (rather than to the users' own ends). Such an approach would not really address the potentially misleading nature of health app practices *itself*, but would rather be focused on taking away one of the important *preconditions* for building digital environments whose overall nature could be considered misleading as such. It would be difficult to make such an approach work within the framework offered by the Directive and exploring the possibility to rely on the GDPR to address questions of data collection and processing might be a better bet. The GDPR and UCPD could be complementary in this regard (see also Helberger, Zuiderveen Borgesius and Reyna 2017, as well as Van Eijk, Hoofnagle and Kannekens 2017 for such a proposal).

When we shift our focus back to the constraints of the Directive itself, with its focus on the provision and/or omission of information, our best bet to address the misleading *overall design* of health apps is to impose transparency requirements that inform users *about* the overall design of the health apps. If we rely on Articles

6 and 7, we end up with transparency requirements which demand that health app providers not only reveal their commercial intent but also, more specifically, their strategies and/or techniques to pursue their commercial agenda. Here, the tension between my ethical perspective on the Directive and the language of the Directive clearly presents itself: from an ethical perspective, one would want to address *the very constitution* of manipulative digital health environments, while the Articles on misleading commercial practices only allow us to address information about them.

It should not be forgotten that broader and more demanding transparency requirements do seem to address the last key element of manipulation, i.e., the manipulator's attempt to keep the manipulative practices hidden. Intuitively, focusing on transparency requirements can still offer a more or less direct method to alleviate potentially misleading influences; when a consumer is informed about what is happening to her, she cannot – strictly speaking – be misled about it. However, this rather formalistic line of reasoning also clearly shows the *limits* of transparency requirements. First of all, if transparency measures make it more difficult to mislead consumers, it is still possible for consumers to be influenced adversely in different ways. In chapter 3, I argued that manipulators will typically *try* to keep their manipulative practices hidden, but that manipulative influences can still successfully dispose manipulees to the manipulator's ends when the manipulative intent becomes transparent. The reason for this, I argued, is that manipulators target exploitable characteristics of people, and if the targeting is done right – i.e., if the manipulator manages to identify and target 'the right buttons to push' – the manipulator may still successfully dispose her targets to her own ends. Put simply, consumers can still be manipulated without actively being misled. Secondly, the question about *which* information must be provided *in what form* does not have an obvious answer. Designing precise transparency requirements goes beyond the scope of this dissertation. But it does seem clear that good transparency requirements should take their inspiration from, for instance, research by communication scientists which could help us understand how persuasion knowledge can be activated (Friestad & Wright 1994, Friestad & Wright 1995, Boerman & Van Reijmersdal 2016, Wojdyski *et al.* 2016).

So if we want to use the UCDP to address what – from an ethical perspective – can be labeled the overall misleading nature of some digital health environments, the language of Articles 6 and 7 forces us to rely on solutions focusing on the provision of information and transparency requirements. There is, however, one final argument I would like to briefly consider. Article 6(2) contains a subtle suggestion that there can be more to misleadingness than the provision and/or omission of information. In Article 6(2) we read that:

A commercial practice shall also be regarded as misleading if, in its factual context, taking account of all its features and circumstances, it causes or is likely to cause the average consumer to take a transactional decision that he would not have taken otherwise, and it involves: [...] non-compliance by the trader with commitments contained in codes of conduct by which the trader has undertaken to be bound, where [...] the trader indicates in a commercial practice that he is bound by the code.

Notice how in this Article the misleading nature of a commercial practice is not primarily defined in terms of the provision and/or omission of information. To be sure, information plays a role because the trader must have indicated that she is bound by a particular code of conduct. But it also seems to be *very the breaching of* the code of conduct to adversely influence the transactional decision of the consumer that constitutes the misleadingness in this case. A commercial practice can, it seems, *in principle* be misleading in more respects than just the provision and/or omission of information. It is just this particular insight that I am interested in, because it opens up the – admittedly – very abstract possibility to think about other sources of misleadingness within the scope of the UCPD. This would also potentially open up more interpretative room to align my argument that the *overall design* of digital health environments can be misleading with the concept of misleading commercial practices. As things stand now, however, the language of the Directive clearly wants us to focus on the provision and/or omission of information in the context of misleading commercial practices.

Let me briefly summarize this section. I have suggested that if we take the manipulative potential of health apps seriously, we should also acknowledge that, from an ethical perspective, *the very constitution* of such digital environments can be misleading in nature. If, however, we want to address the potential misleadingness of health apps within the constraints of the UCPD, we end up with a call for more stringent transparency requirements. Such requirements can strengthen the position of the consumer, but cannot resolve the complete manipulative potential of digital choice environments by themselves. In the next section I turn to aggressive commercial practices. Since there is a less direct ‘natural’ connection between the manipulative potential of health apps and the legal concept of aggressive commercial practices, it will require some more interpretational and argumentative work to show how the Articles on aggressive commercial practices can be helpful. The concept of aggressive commercial practices is, however, rather promising because it does not focus on transparency but instead goes (more) to the

heart of the issue of how consumers can be approached and targeted. The concept of aggressiveness thus has the potential to address the deeper, more structural challenges posed to the autonomy of the health app user.

2.6. Aggressive Commercial Practices

Aggressive commercial practices are defined in Articles 8 and 9 of the UCPD. Where misleading commercial practices focus on the role *information* plays in influencing consumers, aggressive commercial practices deal with situations where, roughly speaking, traders – often enabled by their positions of relative power vis-à-vis consumers – exert problematic *pressure* on consumers either by threatening them or by exploiting circumstances or characteristics of consumers. Put more generally, the Articles on aggressive commercial practices are mainly concerned with the misuse of positions of power by traders.

In what follows, I will first discuss the two Articles on aggressive practices and the role they can play – in their current form – in addressing manipulative health app practices. Then, in a next step, I will explore the question of how my ethical analysis of manipulation in digital (health) environments can inform the legal concept of aggressive commercial practices.

Let me start by providing Article 8 and 9 in full to serve as a basis to discuss the finer details of aggressive commercial practices. Article 8, which defines what constitutes an aggressive commercial practice, reads as follows:

A commercial practice shall be regarded as aggressive if, in its factual context, taking account of all its features and circumstances, by harassment, coercion, including the use of physical force, or undue influence, it significantly impairs or is likely to significantly impair the average consumer's freedom of choice or conduct with regard to the product and thereby causes him or is likely to cause him to take a transactional decision that he would not have taken otherwise.

Article 9 provides five broadly defined types of considerations that should be taken into account when assessing whether a commercial practice uses harassment, coercion, or undue influence:

In determining whether a commercial practice uses harassment, coercion, including the use of physical force, or undue influence, account shall be taken of:

- (a) its timing, location, nature or persistence;
- (b) the use of threatening or abusive language or behavior;
- (c) the exploitation by the trader of any specific misfortune or circumstance of such gravity as to impair the consumer's judgment, of which the trader is aware, to influence the consumer's decision with regard to the product;
- (d) any onerous or disproportionate non-contractual barriers imposed by the trader where a consumer wishes to exercise rights under the contract including rights to terminate a contract or to switch to another product or another trader;
- (e) any threat to take any action that cannot legally be taken.

Article 8 introduces three requirements. For a commercial practice to count as aggressive, it must be (1) a case of harassment, coercion, and/or undue influence; (2) impair or be likely to impair the average consumer's freedom of choice or conduct; and (3) cause or be likely to cause the consumer to take a transactional decision she would not have taken otherwise. The first part basically breaks the concept of 'aggressive' down into the three concepts of harassment, coercion, and undue influence. The second and third part together introduce, just like in Articles 6 and 7 on misleading actions and omissions, a materiality condition: a practice is only aggressive and therefore unfair if the practice impairs or is likely to impair the transactional decision of a consumer. Let me first briefly address the materiality condition before moving on to the concepts of harassment, coercion, and undue influence.

Here, like for Articles 6 and 7, the materiality condition should be considered a low hurdle. Article 8, like Articles 6 and 7, explicitly mentions that a practice that is *likely* to impair choice or conduct – as opposed to actually impairing it – can be considered aggressive and so here, again, no actual empirical test is necessary. The term 'transactional decision' should also be interpreted in the broad sense I discussed earlier: it not only refers to the actual decision to buy a service or product, but to all decisions (potentially) leading up to a purchase, and decisions

after a purchase. Moreover, as I argued in section 2.3 of this chapter, nearly all behavior displayed by users in digital commercial environments that optimize for retention, engagement, and conversion should be considered ‘economic behavior.’ In short: as long as there is a “commercial flavor” to the practice and as long as a judge considers it likely that the practice can lead to a transactional decision the consumer would not have taken without the aggressive influence, the materiality condition is met (Willett 2010: 249).

With the materiality condition posing a relatively low hurdle, much comes down to the interpretation of the concepts of harassment, coercion, and undue influence; even more so because “the term [aggressive] is not a developed concept within most Member States of European Law” (Howells 2006a: 167). So in the absence of a well-developed meaning of the term ‘aggressive,’ the way we understand the concept into which aggressive is broken down – i.e. harassment, coercion, undue influence – is essential to understanding what constitutes aggressiveness. It should, first of all, be noted that a commercial practice does not have to fall neatly within the boundaries of one of these three concepts. “Aggressive practices can fall into one, two or all of those concepts and there is no need to allocate it to a particular head. It is the overall impact that counts. Rogue traders will often use a variety of techniques” (Howells 2006a: 173). Still, for analytical purposes it is important to attempt to understand how the three concepts should be interpreted separately.

So when does a commercial practice actually qualify as harassment, coercion, or undue influence? Howells (2006a: 168-169) as well as Schulze and Schulte-Nölke (2003: 36) point out how difficult it is to draw a line between legitimate hard selling and illegitimate aggressive selling. The Directive is not particularly helpful in drawing such a line. It only defines the term ‘undue influence,’ leaving the terms ‘harassment’ and ‘coercion’ undefined. Article 2(j) states that “‘undue influence’ means exploiting a position of power in relation to the consumer so as to apply pressure, even without using or threatening to use physical force, in a way which significantly limits the consumer’s ability to make an informed decision.” So what makes an influence an *undue* influence? Micklitz *et al.* (2010:147) write that “the fundamental questions about the undue influence concept are: what is meant by a position of power and when does conduct amount to exploitation of that position?” A first important observation to make is the fact that the formulation of Article 2(j) suggests that an undue influence can, but does not have to, involve explicit, direct threats or pressures. This seems to set undue influence apart from harassment and coercion. Without formal legal definitions of harassment and coercion,¹⁷⁸ it seems

¹⁷⁸ Remember, the Directive leaves these two concepts undefined.

intuitively evident that both harassment and coercion *necessarily* involve explicit, direct pressure on a consumer, for the simple fact that coercion and harassment cannot work if it is not abundantly clear to the target that she is being pressured.¹⁷⁹

Schulze and Schulte-Nölke (2003: 37) see a similar distinction between, on the one hand, harassment and coercion, and, on the other hand, undue influence: “While harassment and coercion are the most blatant forms of aggressive selling techniques, the notion of undue influence designates more subtle and therefore more effective methods of unfairly influencing the consumer’s transactional decision.” So one way to distinguish undue influences from harassment and coercion is to differentiate between more subtle types of influences and more directly threatening influences. (That still leaves the question of whether there are more directly threatening influences that do *not* qualify as harassment and coercion but that *do* result from an exploitation of a position of power.) In the absence of legal definitions of the terms ‘coercion’ and ‘harassment’ I follow Schulze and Schulte-Nölke’s (2003) commonsensical distinction between, on the one hand, coercion and harassment as blatant aggressive practices which involve direct forms of pressure, and, on the other hand, undue influence as (mainly, but not necessarily exclusively) more subtle and even sneaky ways of influencing consumers. For the purposes of my argument, it thus makes most sense to work with the concept of undue influence. I do not focus on the aggressive use of lies, threats, blackmail and so on; I am much more interested in the manipulative potential that is sometimes present in digital health environments that can shape long-term commercial user-app relationships.

So if we use the concept of undue influences to address the less direct, subtler ways of exerting pressure on consumers, how should we understand the position of power of the health app provider? In considering what constitutes a position of power, we should start by asking how to understand the term ‘power.’ In chapter 3 I already elaborated on the concept of power, highlighting the importance of understanding power in a wider sense than just an actor’s ability to compel another actor to decide or behave in a particular manner. More specifically, the republican perspective on power is important in this context (Pettit 1997, 2012, 2018). We should not just consider *actual* interferences with behavior, but also the *mere ability* to interfere, and which results from an actor’s privileged position. By virtue of their position as choice architects in the digital society, health app providers are in a position

179 You can only coerce a person when that person understands that she is being coerced and actually feels the pressure resulting from the coercive actions. Similarly, harassment only works when the target actually perceives the harassment and experiences a pressure resulting from the harassment. Put differently, no one attempting to coerce or harass someone can do so without also attempting to exert a pressure on the targets *that is noticed and felt by the targets*.

of power. As designers and operators of data-driven digital choice architectures, health app providers can learn about their users' preferences, desires, and patterns of behavior to build persuasion profiles (Kaptein 2015, Kaptein *et al.* 2015). Moreover, these types of information concerning (health-related) preferences, desires, and patterns of behavior also render health app users – in principle at least – more manipulable; the more one knows about a person, the higher the chance one can figure out which buttons to push or strings to pull.

Going back to the Directive, we need to look at Article 9 for more guidance on what can or should be taken into account when evaluating potential instances of coercion, harassment, or undue influence. The article provides a list of additional considerations that should be taken into account when deciding whether a commercial practice uses harassment, coercion, or undue influence. The five sections which make up Article 9 do not relate specifically to either undue influence or coercion and harassment. Rather, all the mentioned considerations are potentially relevant to all types of aggressive practices. Sections (a) and (c) are most relevant for my argument on health apps, because when taken together they seem to address situations in which traders (subtly) target and seek to exploit particular characteristics or circumstances of consumers. The other sections (i.e., (b), (d), and (e)) mainly focus on the more direct and blatant ways of threatening consumers, as exemplified by the references to threatening or abusive language, threats to take illegal actions, and explicit attempts to prevent consumers from exercising their rights.

The phrasing of sections (a) and (c) betrays the need to interpret 'aggressive commercial practice' in a broad way. Section (c) speaks of "*any* specific misfortune or *circumstance*" which is of enough "gravity" to impair decision-making. (Remember, because the materiality condition poses a low hurdle, it is relatively easy to qualify an influence on decision-making as impairment.) Moreover, section (a) resembles the insistence of Articles 6 and 7 on the importance of the 'overall presentation' of information when considering whether a commercial practice is misleading. In this case, by mentioning "its timing, location, nature, or persistence," it is clear that the Directive wants us to consider a wide range of elements that can (probably) make a commercial influence capable of exploiting weaknesses to impair decision-making. Here we can again think of the practice of collecting and analyzing user data, and the attempt to build trusting relationships with users, as elements that can contribute to the exploitation of a position of power (Helberger 2016: 155).

2.6.1 Using the Concept of Undue Influence to Address the Manipulative Potential of Health Apps

So given the way aggressiveness is operationalized in the Directive, what role can it play in addressing the manipulative potential of health apps in order to safeguard users' autonomy (in its wider relational sense)? I want to propose that despite its relative vagueness, the concept of undue influence offers a lot of potential to address manipulative practices. Articles 2(j), 8 and 9 contain enough 'material' to develop the argument that manipulative practices as I have defined them are a form of undue influence and should thus be covered by the UCPD.¹⁸⁰ Undue influence originates in the misuse of a position of power. Digital health environments such as health apps can afford a position of power to their designers and operators (more in this below). Moreover, data-driven, dynamically adjustable choice architectures afford the designers and operators of those architectures the means to *exploit* their position of power. The ethical concept of manipulation can help explain what constitutes an exploitation of power. Manipulative practices are characterized by the search for weaknesses or biases in potential targets that can be exploited to subtly shape or steer the behavior of targets in the way the manipulator desires. In what follows, I want to offer an interpretation of the concept of undue influence that is inspired by my ethical analyses from the preceding chapters. By doing so, I hope to convince the reader that the concept of undue influence can play an important role in addressing the manipulative potential of health apps and digital choice architectures more generally.

In exploring the usefulness of the concept of undue influence, we need to consider 1) the position of power of health app providers and 2) the possible exploitation of such a position of power.¹⁸¹ I want to argue that to the extent that a digital environment harbors a manipulative potential, the provider of that environment is also necessarily in a position of power. When the provider actually abuses

180 The fact that manipulation can be seen as an undue influence does not imply that every undue influence is a form of manipulation.

181 Strictly speaking we also need to consider the materiality condition built into the concept of undue influence. I will only do so very briefly because I have already established elsewhere that the materiality condition should be considered a "low hurdle" (Wilhelmsson 2006: 137). The exploitation of a position of power is only relevant to the UCPD to the extent that it can adversely influence the *economic behavior* of a consumer. I want to refer back to the earlier discussion on the scope of the Directive and, more precisely, what counts as a 'transactional decision' and 'economic behavior.' I have already established that in the freemium app economy, nearly everything that happens within a deliberately designed digital (health) environment should be considered 'economic behavior.' In the freemium app economy, just trying to *retain* users to *engage* them with your service in the long run is essential to the (financial) success of the app. So when a position of power is used to, for instance, retain and engage users, we can in fact say that the app influences the economic behavior of their users.

that power to shape or steer the behavior of consumers in a manner that can be characterized as being manipulative, we can speak of an *exploitation* of that position of power.

So let me first address the potential position of power of popular for-profit health apps. As I have argued throughout this dissertation, we should adopt a *structural* and *relational* perspective when analyzing popular for-profit health apps. The possible position of power of health app providers should also be understood from this more structural and relational perspective. In the previous section I briefly hinted at the importance of data collection and (personalized) targeting capabilities for the health app's position of power. Helberger (2016) has already argued that data collection aimed at acquiring knowledge about users so as to be better able to target them and personalize services can plausibly be construed as – or at the very least contribute to – being in a position of power. This argument also applies to most popular for-profit health apps because nearly all of them rely on (in their own words: advanced) data analytics to better cater their services to their users. For example, Headspace claims that it can “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.”¹⁸² Even though the shroud of corporate secrecy typical of Silicon Valley companies makes it almost impossible to find out whether Headspace can actually deliver on those promises, we can at the very least conclude that they *themselves* understand their own position of power. Headspace, like most other popular for-profit health apps, clearly presents itself as having a competitive edge over other services precisely because of their clever exploitation of user data. Such a data-driven approach also affords health apps like Headspace power over their users because it increases the *potential* for manipulation (Pettit 2018). The more one knows about one's users, the higher the chances are that one can find out which ‘buttons to push’ or ‘strings to pull’ to shape their behavior in the desired manner (see, e.g., Rudinow 1978: 346). Put simply, information equals manipulation power. So the more the provider of a digital environment knows about its users, the greater the manipulative potential.

The role of trust is also important to understanding the possible position of power of health apps. In chapter 4, I argued that the deliberate attempt to create trusting or trust-like relationships can, on the one hand, contribute to a health app's mission to help its users achieve beneficial health outcomes, but can, on the other hand, also help introduce a manipulative potential. When a user comes to trust a health app over time – either because the user (thinks she) recognizes the trustworthiness

182 Job posting has since been removed. Screenshot available here: <https://imgur.com/a/uIwpvds>.

of the app, or because the app has cleverly *instilled* a sense of trust in the user – she can become less likely to always, in every interaction, perform her role as a critical, circumspect consumer. Moreover, the more a user trusts an app, the more we can expect a user to be willing to volunteer data to the app and provide the app with access to her decisional sphere.¹⁸³ Since manipulation tends to work best when a manipulator has information about her targets, the increased willingness to volunteer data and heightened access to one’s decisions also increases the manipulative potential of the health app in question.¹⁸⁴ This manipulative potential which comes with the development of trust(-like) relationships between apps and users translates to a stronger relative position of power of the health app provider in question. To fully appreciate this argument, one should also remember how trust is entwined with vulnerability: when you trust an actor, you also render yourself vulnerable to having that trust betrayed by that actor. So when you trust an actor, that actor is placed in a position of power since that actor now has the power to betray your trust. It follows that when the providers of health apps try to develop trusting relationships with their users, it contributes to their position of power.

There is an additional argument to be made for the claim that the providers of health apps are potentially in a position of power. This argument builds on the observation that health is a central concern in the life of every person. When a health app manages to convince a user that using the app is going to contribute to, for instance, a healthy lifestyle, the app can also maneuver itself into (the direction of being in) a position of power to the extent that the user can feel at least partly dependent on the app for the pursuit, maintenance, or optimization of something as important as her health.

Now, if we accept that there are good reasons for considering popular for-profit health apps to be in a position of power when they build ongoing trusting relationships with users within data-driven digital health environments, the next question becomes when we can consider such a position of power to be *exploited*. Here, influences that rise to the level of manipulation should be considered an exploitation of power. Remember that manipulation involves the targeting of

183 See, for instance, Patterson (2013) for an analysis of people’s trust in a health app, their willingness to disclose information, and associated privacy risks. She suggests that if users trust a health app, they will indeed be more inclined to share (health) data. She does, however, argue that users’ reasons for trusting a health app and being willing to share (health) data are mostly contextual in nature. Ostherr *et al.* (2017: 9) also find that “[t]he rhetoric of personalization and sharing appears to facilitate an understanding of user-generated health data that deemphasizes the risk of exploitation in favor of loosely defined benefits to individual and social well-being.”

184 To be sure, the presence of a manipulative potential does not mean that actual manipulative acts or practices will materialize.

exploitable characteristics of persons to (attempt to) shape or steer their behavior behind their backs in the direction desired by the manipulator. Put simply, I propose that using one's position of power to identify and target exploitable characteristics for the purpose of subtly¹⁸⁵ steering behavior of health app users in directions that first and foremost benefit the health app provider should be considered an undue influence. In practice, this means that we have to deal with the challenge – already elaborately discussed in chapter 4 – of drawing the line between legitimate, helpful health app practices which involve some degree of targeting of weaknesses and biases (i.e., for the purpose of helping, supporting, or empowering the user), and illegitimate manipulative ways of targeting weaknesses and biases (i.e., to steer their behavior in directions that do *not* benefit the user, but (only) the health app provider).

However difficult this practice of line drawing is, the conceptual refinements I proposed in chapter 3 for the concept of vulnerability are of help here. First a word of warning is in order. Ethical and legal conceptions of vulnerability differ, so when ethical and legal conceptions of vulnerability are combined, misunderstandings can occur. In the legal context, the concept of 'vulnerability' still has a very narrow meaning (as we have seen in section 2.4 of this chapter), referring to a fixed set of personal characteristics – i.e. “mental or physical infirmity, age or credulity” (Article 5(3)). I have already argued in chapters 3 and 4, as well as in section 2.4 of this chapter, that we need a broader understanding of vulnerability inspired by ethical theories of vulnerability to make the concept fit for the digital society. In what follows, I am interested in formulating how the targeting of weaknesses and biases of consumers can constitute an undue influence, and what ethical theories of vulnerability can teach us in that regard. So when I use the term 'vulnerability' in the remainder of this section I refer to the *ethical* notion of vulnerability and when I speak of 'weaknesses and biases' I do not mean to suggest that those are currently (completely) covered by the narrow *legal* notion of vulnerability.

Now, I cannot propose a general formula which will always provide a precise answer to the question of when a commercial practice involving some form of weakness-and-bias-based targeting will constitute an undue influence and, in effect, should be considered an aggressive commercial practice. The answer to such a question depends on the particularities of each case. What I can do is draw on the previous chapters to emphasize how the broader ethical notion of vulnerability can help us become more attentive to how the targeting of weaknesses and biases

185 I.e., attempting to shape or steer behavior in a manner that is not meant to be fully transparent to one's targets.

can render consumers susceptible to manipulative undue influences. To determine which practices involving the identification and targeting of weakness and biases are to be considered a form of undue influence, it is of the utmost importance to consider the different sources and states of vulnerabilities (see chapters 3 and 4). Moreover, it should be kept in mind that there are, broadly speaking, two types of vulnerabilities that are important in the health app context, namely, on the one hand, vulnerabilities pertaining to people's health status, and, on the other hand, people's general cognitive and affective biases, both of which can be targeted to influence users' interactions with(in) health apps. Taking these perspectives into account should be considered to be in line with Article 9(c)'s insistence on "taking into account the exploitation by the trader of any specific misfortune or circumstance of such gravity as to impair the consumer's judgment, of which the trader is aware, to influence the consumer's decision with regard to the product."

So to understand the proper scope of the concept of undue influence, we start from the premise that the average consumer is a consumer who can be rendered manipulable by the targeting of specific weaknesses and biases. Moreover, to see the potential for undue influences, we should focus most of our attention on the optimization logics that are a part of the (health) app economy. When popular for-profit health apps aim to optimize for retention, engagement, and conversion, there is a need to be especially vigilant because such optimization logics can give rise to manipulative practices in the health app context. In practice, this means we need to focus on the various techniques health apps can use to build profitable relationships with users, as well as how those various techniques can leverage weaknesses and biases.

Let us look at the already discussed app called Headspace for a moment. This app offers a limited number of meditation sessions for free to attract users to the app and then tries to retain and engage users to increase the chance that they buy a subscription or additional meditation packs. One of the meditation packs on offer is 'Coping with Cancer.'¹⁸⁶ For a company that prides itself on being good at using data to get to know and serve its customers (see previous chapters), the first obvious question that presents itself is whether, and if so how, the targeting of people dealing with a cancer diagnosis (either because they themselves suffer from the disease or because they know someone who does) can happen. Offering meditation packs to people dealing with cancer is not inherently problematic; it is certainly possible that meditation can offer benefits to those people. However, when

186 <https://www.headspace.com/blog/2014/12/02/coping-with-cancer/> (last accessed 22 September 2020), screenshot available here: <https://imgur.com/a/wReVJOa>. Screenshot of the meditation pack in the app available here: <https://imgur.com/a/Ez20ejC>.

such people are targeted with paid-for meditation packs based on the (inferred or actually established) information that they themselves or people in their surroundings are dealing with cancer, the line between helpful and exploitative offers is easily crossed. Headspace is also a perfect example of an app that tries to build trusting relationships with its users, by emphasizing how much they care about the health and well-being of their users, and by using friendly, cozy, soothing colors and illustrations to make users feel at home. This can contribute to the app provider's position of power. Targeting people based on a disease they (or a loved one) suffers from, with the promise they will be better able to cope with the disease, can be seen as the targeting of a weakness that is directly related to health status. Because cancer is such a widespread and terrible disease, it will seem intuitively clear to most readers that targeting based on health status can be considered a problematic exploitation of a position of power – and thus aggressive – in this context. The less 'serious' (either in public perception of subjectively experienced) the targeted health status becomes, the less clear it is when the line between helpful and exploitative is crossed.

Another factor to consider is the targeting of cognitive and affective biases in order to influence the health-related and economic behavior of users. To stick to the example: it is not just the health status of patients that can be used to target them as prospective customers. Maybe the app has already successfully attracted them to the app with free meditation packs and then uses data analytics to choose between different nudges for different users to maximize the number of subscriptions or additional meditation packs sold. For example, some new users may simply receive a personalized offer with a price that is determined based on what the data tell the user will respond to.¹⁸⁷ Other users may receive personalized social messages ('your friends benefitted from...') or personalized messages that emphasize what tends to work for 'people like you.' Now, when such targeting of cognitive and/or affective weaknesses/biases is done in the context of a gaming app that tries to sell virtual coins, we may be inclined to say that such targeting practices are just business as usual in the app economy. But in a health context, we may want to evaluate such practices differently. Not only is the targeting done to promote a health-related service such as the Coping with Cancer meditation pack; apps like Headspace also tend to try to build trusting relationship with their users by emphasizing the *health* aspect of the relationship (Sax, Helberger & Bol 2018). In such a context,

187 Zuiderveen Borgesius and Poort (2017: 349) write that "personalized pricing seems to be relatively rare." One can, however, make an educated guess that personalized pricing will increase rather than decrease in the years to come. Zuiderveen Borgesius and Poort (2017: 355), referring to Turow *et al.* 2005, also note that "many people regard personalized pricing generally as unfair and manipulative."

where the meaning of the user-app relationship is framed in terms of health (in the wide sense of the word), even the targeting of weaknesses and biases that are not *directly* related to someone's health status can still sometimes be considered a form of undue influence.

To fully appreciate the scope of this argument, it is also important to consider the fact that weaknesses and biases such as the ones mentioned in the previous paragraphs need not be permanent or semi-permanent. As the taxonomy of vulnerability (Mackenzie, Rogers & Dodds 2014) clearly shows, the vulnerabilities (in the wider ethical sense of the word so as to encompass the 'weaknesses and biases' I have been speaking of) we must take into account can be latently present and only materialize under certain circumstances. Moreover, weaknesses and biases need not be permanent and can be experienced for shorter periods of time, or only under certain conditions. For example, a health app may target the insecurities a person has vis-à-vis her health. Those insecurities do not need to be permanent in order for them to count as relevant weaknesses or biases. Even if the insecurities only 'surface' when triggered by particular messages or circumstances (e.g., seeing particular people, products, or stories), they can still be relevant to consider as weaknesses or biases that can be exploited by health apps that occupy a relative position of power.

In conclusion to this section, let me circle back to the principle of autonomy, which informs the Directive as a whole. In the end, the question we should ask ourselves is how the Directive helps to both safeguard and scaffold the health app user's autonomy in the wider relational sense I have advocated throughout this dissertation. The concept of aggressive commercial practices – with a special emphasis on the concept of undue influence – is promising in this regard. By bringing the concept of undue influence in line with the ethical concepts of manipulation and vulnerability, undue influence can be used to address those commercial practices in the digital society that are specifically designed *not* to respect the autonomy of the citizen-consumer. Manipulation, with its (intentional) targeting of weaknesses and biases to adversely influence the behavior of persons, should be seen as the epitome of an undue influence.

2.7. Annex I – The Blacklist

The blacklist contains 31 commercial practices, 23 of which are labeled as misleading commercial practices and 8 of which are labeled as aggressive commercial practices. Although the subdivision of the blacklist into misleading and aggressive practices provides some structure, the blacklist in its entirety has been described as “a rather rag bag collection of unfair practices [...] it is not even

clear that all practices are under the appropriate heading” (Howells 2006b: 22). Moreover, nothing important hinges on the labels ‘misleading’ and ‘aggressive’ in this context because items on the list are always forbidden, without the need to explain why or how they are deemed to be misleading or aggressive. As long as a commercial practice is judged to be sufficiently similar to one of the blacklisted commercial practices, it is forbidden, period.

As mentioned earlier, the 31 commercial practices listed by the blacklist are so specific that they will not apply in most cases. Still, it makes sense to highlight a few items on the list which may be relevant in the context of for-profit health apps. The first item on the blacklist to consider is number 11:

Using editorial content in the media to promote a product where a trader has paid for the promotion without making that clear in the content or by the image or sounds clearly identifiable by the consumer (advertorial).

This item on the blacklist basically deals with native advertising and might be relevant because for-profit health apps feature a lot of native advertising. In practice, it can be difficult to rely on this blacklisted practice because native advertising comes in many forms and shapes and not everyone will agree that what I call native advertising is in fact native advertising according to the UCPD. As I already discussed in the section on misleading commercial practices, much comes down to the question of when a disclosure counts as good *enough* (also see section 2.5.2 of this chapter). For example, is the way in which professional athletes use Strava’s social and podcast features to draw attention to their personal fashion and lifestyle brands native advertising? I would argue that such practices qualify as native advertising. But based on the phrasing of this particular blacklist item, it is not at all obvious that it applies in a straightforward manner to native advertising practices such as the one mentioned in the Strava example.

Blacklist item number 17 might also be relevant in the health app context:

Falsely claiming that a product is able to cure illnesses, dysfunction or malformation.

Health apps can of course make health claims, and to the extent they do, this practice on the blacklist is relevant. However, popular for-profit health apps often take great care *not* to make explicit *medical* promises and claims regarding the

health of their users. Most for-profit health apps position themselves as ‘healthy lifestyle’ apps rather than health apps in the (more) medical sense of the word in order not to be regulated as medical devices or services.

Blacklist item number 20 can potentially be relevant to any service that is part of the contemporary data economy where ‘free’ services can be offered because users of the service ‘pay’ in other ways (such as with their data or with their attention). For-profit health apps are very much part of this contemporary data economy. Blacklist item number 20 reads as follows:

Describing a product as ‘gratis’, ‘free’, ‘without charge’ or similar if the consumer has to pay anything other than the unavoidable cost of responding to the commercial practice and collecting or paying for the delivery of the item.

In very abstract terms, it certainly makes sense to say that health apps that advertise themselves as a free service, but whose business models depend on users ‘paying’ with other means such as user data or attention, are in breach of this blacklisted commercial practice to the extent that they describe their product as free. As things stand now, it has not been definitively established that consumers that have to ‘pay’ with their attention or data are not legitimately enjoying a *completely* ‘free’ product. The argument does have legal purchase though as the European Commission’s (2016a: 97) Guidance on the UCPD explicitly acknowledges that “products presented as ‘free’ are especially common in the online sector. However, many such services are only accessible for consumers on the condition that they provide personal data.” The European Commission (2016a: 97) further points to the “increasing awareness of the economic value of information related to consumers’ preferences, personal data and other user-generated content” in order to conclude that a commercial practice could indeed be considered unfair (because it is misleading) if a service is promoted as ‘free’ without explicitly mentioning the costs to consumers of having to volunteer their data or attention. No. 20 of the Annex could thus apply to health apps that advertise themselves as being ‘free’ while their business models depend on users volunteering their data and/or attention.

Apart from considering the applicability of the existing items on the blacklist, we could also ask whether new items should be added to the blacklist to help address unfair commercial practices in health apps. One possibility to explore is a ban on the targeting of insecurities, fears, or concerns that people have concerning their health for the purpose of influencing the economic behavior of health app users in ways that (solely) benefit the health app provider. Such a proposal would basically

turn my arguments on undue influence and the exploitation of “specific misfortunes or circumstances” (Art 9(c)) into a separate blacklist item specifically for the health (app) context.

Although this proposal is certainly worth exploring, it also raises some obvious questions regarding its practical implementation and use. It will be difficult to determine what counts as a ‘health concern.’ As was discussed in chapter 2, the meaning of the term ‘health’ is fundamentally contested, and it is precisely this conceptual fluidity that is used by popular for-profit health apps to allude to the user’s health without explicitly handing out health advice in the medical sense of the word. It will, in effect, be a challenge to capture, legally, what counts as the targeting of a *health* concern without being either overinclusive or underinclusive. Another challenge, related to the previous one, is to draw a line between influences that aim to, on the one hand, help the user with her health (or: empower the user) and, on the other hand, influences that are (solely) aimed at influencing the economic behavior of users in ways that serve the interests of the health app provider. Because this line between empowerment and manipulative exploitation is so difficult to draw¹⁸⁸ – especially in the health context where it can be easy to suggest that an app has some positive effect on the user by alluding to the user’s health – the blacklist with its brief and specific items is maybe not the best instrument to legally address these challenging questions.

2.8. The General Clause on Unfair Commercial Practices

Although the General Clause on unfair commercial practices (laid down in Article 5) has been described as a “remedy of last resort which only applies to particularly unusual circumstances” (Micklitz 2006: 86), it is worth exploring its somewhat underestimated potential to address unfair commercial practices. Remember that the General Clause holds that “a commercial practice shall be unfair if (a) it is contrary to the requirements of professional diligence, and (b) it materially distorts or is likely to materially distort the economic behavior with regards to the product of the average consumer whom it reaches or to whom it is addressed, or of the average consumer of the group when a commercial practice is directed to a particular group of consumers” (Art 5(2)(a)(b)). The last part (b) of Article 5(2) does not need to be discussed here, since it describes the standard materiality condition already discussed multiple times in this chapter. To explore the potential of the General Clause in addressing unfair commercial practices, especially in the

188 The entire fourth chapter is devoted to the challenge of differentiating between claims of empowerment and (the potential for) manipulative exploitation of users.

context of *health* apps – as opposed to any other type of app – we should focus on the first part (a) of Article 5(2). The central question to ask is how to interpret ‘professional diligence’ and, importantly, the *requirements* that come with it.

The term ‘professional diligence’ is defined in Article 2(h) as “the standard of special skill and care which a trader may reasonably be expected to exercise towards consumers, commensurate with honest market practices and/or the general principle of good faith in the trader’s field of activity.” This definition is still rather broad and leaves important interpretative questions open; for example, what constitutes an ‘honest market practice’? That is not to say, however, that there is no way to understand ‘professional diligence’ in a concrete, practical manner. Codes of conduct could perform a helpful role. If a trader is subject to a code of conduct but at the same time acts contrary to the requirements of the code of conduct, it is not difficult to argue that the trader has acted contrary to the requirements of professional diligence. At the time of writing there are, however, no codes of conduct for health apps that could perform such a role. The European Commission started facilitating the process of drafting a code of conduct for privacy in health apps in 2014, but with the arrival of the GDPR the code of conduct was not approved because it did not meet all the necessary GDPR requirements.¹⁸⁹ So right now, there is no code of conduct that could be given a prominent role in addressing unfair commercial practices via the General Clause of the UCPD. If a code of conduct does get approved in the near future, the General Clause could become a (much) more relevant and prominent instrument for addressing unfair commercial practices.¹⁹⁰

Apart from the relatively straightforward suggestion of checking whether a code of conduct has been breached, it is worth spending some time reflecting on what, ideally, meeting the requirements of professional diligence *should* mean in the context of health apps. Here, we can draw on my analyses from earlier chapters. To determine which requirements of professional diligence we would like for-profit health apps to meet, we need to recount a few important aspects of health apps and their users. First, health apps deal with an important and relatively sensitive issue, namely the health of their users. Second, health apps, in their attempts to build (profitable) *relationships* with their users *over time*, will also – at least in some cases – build trust-like relationships with their users. Lastly, in their attempts to retain and engage users, for-profit health apps will often – intentionally or

189 <https://ec.europa.eu/digital-single-market/en/privacy-code-conduct-mobile-health-apps> (last accessed 22 September 2020).

190 Although the usefulness of the General Clause would then still be determined by the actual content of the hypothetical code of conduct.

unintentionally – end up identifying and targeting the weaknesses and/or biases of their users. It is the *combination* of these elements that make health apps the kind of service providers for which the requirements of professional diligence should be understood (almost) as a duty of care. In engaging with data-hungry health apps over longer periods of time, users make themselves especially vulnerable and – in effect – manipulable to the health apps they use. The fact that people choose to use health apps does not absolve those apps from the responsibility not to *misuse* the trust users put in an app, as well as the privileged insights into users' health-related behaviors, insecurities, and desires that the apps can infer. As I argued in the section on aggressive commercial practices, the trust users put in health apps, combined with the privileged access to the behavioral health sphere of users *and* the access to information about users' health-related worries, desires, and insecurities makes for a position of power on the part of the app. With such a position of power should come strict requirements of professional diligence.

It could still be asked what these requirements should, ideally, look like. It is clearly implausible to demand that for-profit health apps should *only* be allowed to act as if they were health professionals that *only* try to help their clients with their health to the best of their abilities. Here we encounter, once again, the vague but important line between illegitimately profiting from health in manipulative ways and trying to grow a legitimate business that profits from health and lifestyle advice. Short of forbidding any kind of economic activity in the health app context, we can think of some requirements of professional diligence that should, ideally, be met.

When health apps collect data which encode weaknesses and biases, the use of those data *purely for the sake of optimizing retention, engagement, and conversion* should be seen as a failure to meet the requirements of professional diligence in the health app context. An important principle of precaution should be observed in this context: when you collect large amounts of data *in a health context* and which are directly or indirectly related to the weaknesses and biases of users, and when you use those data to pursue optimization of business metrics, you can and should foresee that you can end up exploiting your privileged position of power and treat users unfairly.

Furthermore, I want to re-emphasize the distinction I introduced in chapter 3 (following Susser, Roessler and Nissenbaum 2019a) between an intended act of manipulation and manipulative practices. Instead of focusing on all the separate, individual interactions between health apps and users, and to ask for all those separate interactions when determining whether an intended act of manipulation was successfully executed, it makes more sense to take a more structural perspective

and ask whether the intentional overall design and operation of a digital health environment tends to produce ample opportunities for manipulative influences on the behavior of users. Similarly, the requirements of professional diligence should be evaluated from the perspective of the overall design and operation of the digital health environment offered to users.

Maybe we should go even one step further and suggest that *the very fact* that one offers a data-driven, dynamically adjustable choice environment to users should come with a corresponding professional duty not to manipulate persons. If all the preceding chapters made one thing clear, it is the fact that in the digital society, digital choice architectures harbor a serious manipulative potential that can shape and steer the behavior of large groups of persons. The cliché that ‘with great power comes responsibility’ should inform our idea of professional diligence and, equally important, the associated professional duties. The cumulative criteria for manipulation formulated in chapter 3 could function as the starting point for a more specific understanding of the professional duties of choice architects in the digital society. For example, one should only be allowed to collect significant amounts of user data if one does *not* use those user data (and persuasion profiles) for the purpose of discovering exploitable characteristics which are then targeted to steer the behavior of people in ways that only benefits the choice architect.

3. Conclusion

Having discussed the UCPD and the various instruments it offers to address potentially unfair commercial practices in health apps, how and to what extent can the Directive actually help to safeguard the autonomy of health app users? As my analysis in this chapter shows, the UCPD offers many legal tools for addressing manipulative health app practices if one is willing to actively look for them. In many instances, the phrasing of the Directive might suggest a rather narrow scope, but if one embraces the overall philosophy behind the Directive – i.e., safeguarding and promoting the autonomy of consumers – and appreciates the European Commission’s (2016a) more recent Guidance document, it becomes clear that the Directive can and should be interpreted in a broad rather than narrow manner.

The UCPD dates back to 2005. At that time, the (freemium) app economy did not yet exist. It was only in 2007 that Apple’s first iPhone, which was the first broadly used consumer smartphone offering the ability to install ‘apps’ as we know them today, was released. It is therefore unsurprising that it requires considerable interpretational effort to show how commercial practices in the (freemium) app

economy (and more specifically in the health app context) are still captured by a Directive which predates this new part of the economy. The (freemium) app economy is built on the philosophy of building long-term relationships with users in data-driven digital environments, rather than focusing on optimizing financial gains in one-off interactions. To make the UCPD fit for the future, we need to embrace broader, more dynamic concepts of what constitutes a commercial practice and a transactional decision. In the contemporary app economy, nearly everything a (health) app provider does to attract users to one's app, to keep them engaged, to stimulate users to volunteer their time and data, and to sometimes even get users to spend money, can be considered (part of) a commercial practice. Similarly, nearly all behaviors of consumers within such digital environments can be considered transactional decisions (or, put more generally, economic behavior). If we do not rely on these broader understandings of what constitutes a commercial practice or a transactional decision, the UCPD will be(come) largely irrelevant in the digital society.

When we ask how the UCPD can help safeguard or even scaffold the autonomy of the health app user, it is essential to keep autonomy's relational nature in mind. The (freemium) app economy incentivizes the building of relationships with users, in order to monetize these relationships *over time* in various ways. Moreover, these relationships take shape within the data-driven and dynamically adjustable digital environments that the health apps themselves largely control. Throughout this dissertation, I have localized the manipulative potential of popular for-profit health apps in their control over the conditions under which ongoing relationships are shaped.

To address this manipulative potential, the concept of misleading practices may – at first glance – seem like the best candidate. Manipulation often (though not necessarily) involves an attempt to mislead manipulees, making for an intuitive connection between manipulation and misleading commercial practices. In practice, however, the concept of misleading commercial practices is of limited use due to its strong emphasis on the provision of (the right kinds of) information to consumers. The concept of misleading commercial practices can be used to address some significant but ultimately derivative problems such as native advertising and the misrepresentation of the 'free' nature of health apps. Misleading commercial practices do not get us to the heart of the problem, namely the position of power – and possible exploitation thereof – health apps are in vis-à-vis their users due to their control over the digital health environment within which ongoing relationships are shaped. This is precisely where the concept of aggressive commercial practices comes in and is much more promising. More specifically, the concept of undue

influence with an emphasis on the *exploitation* of a *position of power* does a good job of capturing the core challenge of differentiating between legitimately helpful and illegitimately manipulative health app practices.

In thinking through the question of how a position of power can be established and exploited, we must also appreciate the important role vulnerabilities (in the wider ethical sense of the term) play. To properly capture and address the challenges posed by the (health) app economy – and being a consumer in the digital society in general – the sharp distinction between, on the one hand, the average consumer, and, on the other hand, the vulnerable consumer, should be overcome. As this dissertation – in line with the literature on digital market manipulation (Hanson & Kysar 1999, Calo 2014) – has made clear, vulnerability should not be considered a ‘state of exception,’ but rather a permanent aspect of the digital market economy. That is not to say that all consumers are always vulnerable. However, every person can move in and out of states of vulnerability, and different vulnerabilities can materialize (often temporarily) under different conditions (Mackenzie, Rogers & Dodds 2014). Popular for-profit health apps, which tend to collect user data concerning health concerns and behavior, can use the generated insights to target or even *evoke* – intentionally or unintentionally – the vulnerabilities of their users. Accepting a more dynamic concept of vulnerability which helps describe how *most* consumers are dispositionally vulnerable to manipulative practices is essential to the UCPD’s ability to safeguard user autonomy, especially in the context of aggressive commercial practices where determining whether a position of power has been exploited is the core question to be answered.

The often-underappreciated blacklist and general clause also have a role to play. Because we are dealing with apps that operate in the *health* domain – even if we accept that most apps do not deal with health in the medical sense of the word – we need to consider the need for additional blacklisted practices as well as an insistence on the professional diligence of app providers. Precisely because health is such an inescapable and fundamental need and concern of everyone, even the ostensibly very healthy part of the population, we should consider assigning a duty of care to health app providers; or at the very least something resembling a duty of care as part of their professional diligence.

In conclusion, the UCPD offers ample opportunities to address manipulative health app practices, but only if one is willing to opt for broad interpretations of key terms in the Directive. By building on ethical insights into autonomy, vulnerability, trust, and manipulation, I have offered what I hope to be a plausible reading and interpretation of the UCPD which can help to truly safeguard the autonomy of

consumers in the face of increasingly advanced digital choice architectures. The increased power that data-driven dynamically adjustable digital choice architectures afford to their designers should be matched by our willingness to adopt broader, future-proof interpretations of key concepts in the UCDP.

Conclusion

1. The Importance of the Relational Perspective

In the introduction, I explained that this dissertation is not only about health apps, but also about the role digital choice architectures plays in the digital society. So I want to start this conclusion by briefly discussing my main findings concerning the ethical evaluation of digital environments.

One of the defining features of contemporary digital choice architectures is that they can learn about their users. What they learn can be used to adjust and personalize the experience of users. Another way of putting this is to say that a kind of relationship between a user and a digital environment can develop over time. Much like 'real' – i.e., interpersonal – relationships, these digital relationships can be pleasurable, helpful, or useful, but they also render one *vulnerable* to the other party. When someone, be it a person or a digital environment, knows you, they can better anticipate your needs; this can be helpful. At the same time, granting someone a privileged insight into your psychology and behavior also makes you vulnerable to the other party, because they can exploit those privileged insights to, for instance, manipulate you for their own benefit. This, in essence, is the core tension between empowerment and manipulation that I have been dealing with throughout this dissertation. Those properties of digital (health) environments that explain their potential to *help* and *support* us are also precisely the same properties that explain their manipulative potential. It also follows that digital choice architectures do not raise a very straightforward challenge to our personal autonomy. Digital environments such as health apps can respect, support, or scaffold our autonomy by collecting data about us, learning about us, and personalizing aspects of their environment for us. But, again, those same capabilities can end up undermining our personal autonomy when we reside in digital environments of a manipulative nature.

Now, to better be able to capture and evaluate this tension, I have emphasized the importance of what can be called the 'relational perspective.' To understand the impact of digital (health) environments on our autonomy, we need to draw on theories of *relational* autonomy. When asking how digital (health) environments impact our personal autonomy, it is tempting to focus on individual, separate

interactions between an environment and a user and ask whether a *particular influence* respects the autonomy of the user, or whether a *particular decision* can be considered (sufficiently) autonomous. If we were to adopt such an approach, we would fail to see the bigger picture: contemporary digital choice architectures can build relationships with us *over time* and, relatedly, attempt to give a particular shape or direction to those relationships to try to shape our behavior over time. To be able to ask the question of autonomy in a sensible manner, we thus need to ask how the relationship-building efforts of digital (health) environments impact our possibilities for developing and practicing our autonomy over time. After all, autonomy is something we learn and develop *through* our relationships with others and with technology, and our ability to practice our autonomy is co-determined by the support we receive from our (technological) surroundings. Since we increasingly tend to inhabit *digital* environments where we develop (partly) *digital* relationships, autonomy's relational nature has never been more important.

2. Manipulation as a Key Concept to Understand and Evaluate Digital (Health) Environments

Following Susser, Roessler, and Nissenbaum (2019a, 2019b) I have argued that the concept of 'manipulation' is especially useful to both understand and evaluate how contemporary digital choice architectures can challenge our autonomy. I have defined manipulation as 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.) This concept of manipulation helps us to understand when the building of commercial relationships within digital health environments can become problematic. Put simply: we do not want health apps to *exploit* their privileged access to people's (health-related) concerns and behaviors to learn how to exploit their weaknesses and biases in order to shape their behavior and decisions in ways that benefit the health app provider, without taking into account – and acting on – the interests and concerns of those same people.

Although purely instrumental relations between service providers and consumers are a common feature of the marketplace, the manipulative potential of health apps – and digital choice architectures in general – does deserve special scrutiny. It would be naive to insist that what I call 'manipulation' is always simply 'business

as usual.’ In an analogue world, outside of digital choice architectures, this might be true – to a certain extent. For example, a vendor in a physical television store may target people’s general desire for success with a clever advertising campaign that, somehow, links owning a particular television with living a successful life. In doing so, the vendor seeks – in a way – to target a generic exploitable characteristic of (most) people to make them serve her own ends (i.e., selling televisions) while disregarding¹⁹¹ the interests of her prospective customers. Running such an advertising campaign can indeed be seen as ‘business as usual’ and is not particularly nefarious. This largely innocent nature of offline advertising campaigns does not, however, necessarily carry over to the way things are done in the digital society. In the digital society, we deal with *data-driven* and *adaptive* digital environments that can learn about individuals (e.g., their patterns of behavior, their cognitive and affective biases, their concerns, and their desires) and *target* those in a *personalized* fashion with interventions catered to *a specific individual*. This difference between the analogue society and the digital society is a game changer, and my arguments concerning manipulation are meant to address manipulation *in the digital society*.

3. Health Apps: Between Empowerment and Manipulation

If we focus our attention on popular, for-profit health apps, we see that the technological affordances that make them possible sources of user empowerment are, at the very same time, the possible source of their manipulative potential. Put differently, the technology *itself* is not inherently empowering or manipulative; it comes down to the specific implementations of technology to determine whether the technology serves an empowering or manipulative role. The data-driven and adaptive nature of digital health environments can be used to empower users. For example, collecting data on people’s behavioral and psychological tendencies and learning about their specific circumstances and desires can allow for *personalized* solutions. This technological potential for empowerment must, however, be situated within the (demands of the) contemporary app economy. Nearly all popular health apps are for-profit, and many of them operate on the basis of a freemium business model. To be commercially successful, health apps must – just like any other type of app – use the technological means available to them to design their digital environments for the optimization of user retention, user engagement, and user conversion. It is here that we encounter the often vague but important line between empowerment and manipulation. Do you use the data you

191 To be sure, there may exist television vendors who truly care about their customers and only want to sell televisions their prospective customers *truly* need. But such a vendor would be the exception. I think it is safe to say that most television vendors just want to sell televisions and do not *always* and *truly* care about the interests of their customers.

collect on your users' behavioral and psychological tendencies as well as their (health-related) concerns, insecurities, and desires to *help* and *support* your users? Or do you use those same data-driven insights to build a digital environment that stimulates behaviors and self-understandings that are most profitable to you, the app provider? In practice, there is never a sharp, hard distinction between these two options. Health apps that design their digital health environments with (mainly) the development of *profitable* relationships with their users in mind need not necessarily undermine the autonomy of their users. But – and this is the important point – when health apps operate just like any other type of for-profit app and, above all, optimize their digital environments for user retention, engagement, and conversion, they knowingly and willingly accept the risk of exploiting their position of power vis-à-vis at least some of their users. Precisely because health is universally desired and needed, the promise to support people in their pursuit of health places the health app provider in a position of power which should come with enhanced responsibility. So when a health app provider tries to instill trust in their users and learns about their vulnerabilities, the *exploitation* of that position of power for the deployment of commercial practices that are manipulative in nature is worrying.

4. Addressing Manipulative Health App Practices

The Unfair Commercial Practices Directive (UCPD) is a promising legal instrument to address many of the manipulation worries health apps – and commercial digital environments more generally – give rise to. However, for the Directive to be 'fit for the future,' we need interpretations of key concepts in the Directive that are sensitive to contemporary developments in the digital society. There should be more attention for the fact that in the digital society, service providers seek to develop ongoing relationships with consumers and try to monetize these relationships over time. The concept of what constitutes a 'commercial practice' should reflect this reality: consumers no longer engage in a series of clearly separable commercial interactions with vendors or service providers when they use digital services. More often than not, the 'boundaries' between separate commercial practices become blurry. In a similar vein, what constitutes a 'transactional decision' or 'economic behavior' should be seen in light of the freemium app and digital services economy. As business models gravitate towards the optimization of user retention and engagement, efforts by apps and digital service providers to stimulate behaviors which contribute to the optimization of these metrics should also be interpreted as invitations to display economic behavior. Money, data, and attention are all valuable in the contemporary digital society and should all figure into our interpretations of what 'counts' as a transactional decision or as economic behavior.

Another important area of attention is our understanding of vulnerability and how vulnerability relates to the average consumer standard. Building on broader, ethical theories of vulnerability, the legal concept of vulnerability should be understood in a broader and more nuanced manner. There are many sources and states of vulnerability, and vulnerabilities can be permanent or temporary. It is increasingly difficult to maintain a hard distinction between, on the one hand, the average consumer, and, on the other hand, the particularly vulnerable consumer; everyone is potentially vulnerable sometimes, under some conditions, to manipulative practices. Such a broader understanding of vulnerability is especially important in the digital society. One of the defining properties of data-driven, dynamically adjustable digital choice architectures is the fact that they can *identify* and *target* various vulnerabilities (and they may even be able to *evoke* or *create* vulnerabilities).

To address the manipulative potential of health apps and digital (health) environments more generally, it is essential to focus on the position of power such apps or environments occupy. For health apps, the fact that they deal in health advice, often collect data pertaining (directly or indirectly) to vulnerabilities, and, moreover, tend to attempt to establish trust-like relationships with their users, all contribute to their position of power. When we adopt a broader, more nuanced conception of vulnerability in the digital society and understand the position of power apps can occupy, the concept of aggressive commercial practices – with a specific focus on the concept of undue influence – can become a promising regulatory tool.

5. Suggestions for Improvement

This dissertation is essentially a warning about the manipulative potential of digital environments, such as health apps. I have, however, also emphasized the potential for useful, empowering applications of the same technologies. So how can we strive for empowering digital environments while trying to temper manipulative tendencies?

In the health app context, there is one clear tension that should and can be addressed. The incentive structure that comes with, and the design choices that flow from, the freemium business model do not go well together with *health* apps. To be successful financially, freemium apps need to build long-lasting, intense relationships with their users. To do so, there is a structural pressure to utilize the technological capabilities offered by contemporary apps which allow for the collection and analysis of user data to identify exploitable characteristics

that can be targeted to boost user retention, engagement, and conversion. In a context where something as important and often sensitive as health is at stake, the freemium model can quickly turn manipulative. For health apps, other business models which are less focused on the constant collection, analysis, and exploitation of user data *for commercial purposes* seem more appropriate. Think, for instance, of subscription models. When a subscription for a longer period of time has been secured, there should be less pressure to constantly search for vulnerabilities which can be targeted to engineer ongoing engagement.¹⁹²

Enforcement efforts by national consumer protection agencies – for instance on the basis of aggressive commercial practices and undue influence – can also help to draw some lines in the sand. The line between empowerment and manipulation can be vague, but that does not mean that the line is not important. What is clear though, is the fact that the line has to be drawn *somewhere* and enforcement efforts can help to do this. These regulatory line-drawing exercises are certainly going to be subject to criticism and debate. This is not a problem and should actually be considered a blessing in disguise: we *need* there to be more debate about the at times uneasy tension between helping people improve their health and companies profiting from those same efforts.

192 An emphasis on subscription models is, of course, not a perfect solution. Apps that depend on subscription fees also need to recruit users and convince them to purchase a subscription, which can still incentivize the collection and exploitation of user data to help achieve those ends. Moreover, a health app can still depend on manipulative commercial practices to get users to buy subscriptions.



References

- Ajana, B. (2017). Digital Health and the Biopolitics of the Quantified Self. *Digital Health*, 3, 1-18.
- Anderson, J. (2010). *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Thaler, Richard H. and Sunstein, Cass R.. Yale University Press, 2008. x+ 293 pages.[Paperback edition, Penguin, 2009, 320 pages.]. *Economics & Philosophy*, 26(3), 369-376.
- 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 & J. H. Honneth (Eds.), *Autonomy and the Challenges to Liberalism*. Cambridge University Press (pp. 127-149).
- Ariely, D. (2008). *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. New York: HarperCollins.
- Ariely, D. (2010). *The Upside of Irrationality: The Unexpected Benefits of Defying Logic at Work and at Home*. New York: HarperCollins.
- Ariely, D. (2012). *The (Honest) Truth about Dishonesty*. New York: HarperCollins.
- Armstrong, D. (1995). The Rise of Surveillance Medicine. *Sociology of Health & Illness*, 17(3), 393-404.
- Bachrach, P., & Baratz, M. S. (1962). The Two Faces of Power. *American Political Science Review*, 56, 947-952.
- Bachrach, P., & Baratz, M. S. (1963). Decisions and Nondecisions: An Analytical Framework. *American Political Science Review*, 57, 641-651.
- Bachrach, P. & Baratz, M. S. (1970). *Power and Poverty. Theory and Practice*. New York: Oxford University Press.
- Baier, A. (1986). Trust and Antitrust. *Ethics*, 96(2), 231-260.
- Barksy, A. J. (1988). The Paradox of Health. *New England Journal of Medicine*, 318(7), 414-418.
- Baron, M. (2003). Manipulativeness. *Proceedings and Addresses of the American Philosophical Association*, 77(2), 37-54.
- Benn, S. I. (1984). Privacy, Freedom, and Respect for Persons. In F. D. Schoeman (Ed.), *Philosophical Dimensions of Privacy: An Anthology* (pp. 223-244). Cambridge: Cambridge University Press.
- Berman, K. (2015). How Startups Should Use Behavioral Economics. *Techcrunch*, 7 December. Last accessed 22 September 2020 at <https://techcrunch.com/2015/12/07/how-startups-should-use-behavioral-economics/>
- Boerman, S. C., & Van Reijmersdal, E. A. (2016). Informing Consumers about 'Hidden' Advertising. A Literature Review of the Effect of Disclosing Sponsored Content. In P. De Pelsmacker (Ed.), *Advertising in New Formats and Media: Current Research and Implications for Marketers* (pp. 115-146). Bingley: Emerald.

- Bol, N., Helberger, N., & Van Weert, J. C. M. (2018). Differences in Mobile Health App Use: A Source of New Digital Inequalities. *The Information Society*, 34(3), 183-193.
- Bol, N., Høie, N. M., Nguyen, M. H., & Smit, E. S. (2019). Customization in Mobile Health Apps: Explaining Effects on Physical Activity Intentions by the Need for Autonomy. *Digital Health*, 5, 1-12. <http://dx.doi.org/10.1177/2055207619888074>
- 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. T. Grüne-Yanoff & S.O. Hansson (eds.) *Preference Change: Approaches from Philosophy, Economics and Psychology*, pp. 207-219. Dordrecht: 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. Last accessed 22 September 2020 at <https://www.nytimes.com/2018/07/25/magazine/big-business-gwyneth-paltrow-wellness.html>
- Brown, B. J., & Baker, S. (2012). *Responsible Citizens: Individuals, Health and Policy under Neoliberalism*. London: Anthem Press.
- Business Insider (2017). Why Dutch Bank ING Swapped Client Lunches for Gym Trips and Uses Smart Watches to Track Staff Well-Being. Last accessed 30 September 2020 at <https://www.businessinsider.nl/ing-wellbeing-quotient-tomtomb-program-2017-4/>.
- 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.
- Calo, R. (2016). Privacy, Vulnerability, and Affordances. *DePaul Law Review*, 66(2), 591-604.
- Cameron, G., Cameron, D., Megaw, G., Bond, R., Mulvenna, M., O'Neill, S., Armour, C., & McTear, M. (2017) Towards a Chatbot for Digital Counselling. *Proceedings of the 31st British Computer Society Human Computer Interaction Conference*. <http://dx.doi.org/10.14236/ewic/HCI2017.24>
- 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. Last accessed 22 September 2020 at <https://www.forbes.com/sites/kathleenchaykowski/2017/01/08/meet-headspace-the-app-that-made-meditation-a-250-million-business/#7641f4f81f1b>.
- Christman, J. (Ed.) (1989). *The Inner Citadel*. Oxford: Oxford University Press.
- 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.
- Clifford, D., & Ausloos, J. (2018). Data Protection and the Role of Fairness. *Yearbook of European Law*, 37, 130-187.

- 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.
- Conick, H. (2017). Read This Story to Learn How Behavioral Economics Can Improve Marketing. *American Marketing Association* website, 31 December. Last accessed 22 September 2020 at <https://www.ama.org/publications/MarketingNews/Pages/read-story-learn-how-behavioral-economics-can-improve-marketing.aspx>
- 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.
- Dahl, R. A. (1957). The Concept of Power. *Behavioral Science*, 2(3), 201-215.
- Devisch, I. (2013). Proloog [Prologue] in I. Devisch (Ed.), *Ziek van gezondheid: voor elk probleem een pil?* [Sick of Health: A Pill for Every Problem?], pp. 15-26. Amsterdam: De Bezige Bij.
- Degli Esposti, S. (2014). When Big Data Meets Dataveillance: The Hidden Side of Analytics. *Surveillance & Society*, 12, 209-225.
- Dijck, J. van (2014). Datafication, Dataism and Dataveillance: Big Data between Scientific Paradigm and Ideology. *Surveillance & Society*, 12(2), 197-208.
- Dow Schüll, N. (2014). *Addiction by Design. Machine Gambling in Vegas*. Princeton University Press.
- Duivenvoorde, B. (2013). The Protection of Vulnerable Consumers under the Unfair Commercial Practices Directive. *Journal of European Consumer and Market Law*, 2(2), 69-79.
- Dworkin, G. (1989). The Concept of Autonomy. In J. Christman (Ed.), *The Inner Citadel* (pp. 54-62). Oxford: Oxford University Press.
- Eijk, N. van, Hoofnagle, C. J., Kannekens, E. (2017). Unfair Commercial Practices: A Complementary Approach to Privacy Protections. *European Data Protection Law Review*, 3(3), 325-337.
- 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 of the 2017 CHI Conference on Human Factors in Computing Systems*, 624-654.
- 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.
- Engelen, B., & Nys, T. R. V. (2019). Nudging and Autonomy: Analyzing and Alleviating the Worries. *Review of Philosophy and Psychology*, 11, 137-156.
- 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). Last accessed 22 September 2020 at <https://ec.europa.eu/digital-single-market/en/news/green-paper-mobile-health-mhealth>
- European Commission (2016a). *Guidance on the Implementation/Application of Directive 2005/29/EC on Unfair Commercial Practices*. Brussels. Last accessed 22 September 2020 at <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52016SC0163>
- European Commission (2016b). *Consumer Vulnerability across Key Markets in the European Union*. Report written by London Economics, VVA Consulting & Ipsos Mori, Brussels. Last accessed 22 September 2020 at https://ec.europa.eu/info/sites/info/files/consumers-approved-report_en.pdf
- 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.
- Fotopoulou, A., & O’Riordan, 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.
- Frankfurt, H. G. (1971). Freedom of the Will and the Concept of a Person. *The Journal of Philosophy*, 68(1), 5-20.
- Friestad, M., Wright, P. (1994). The Persuasion Knowledge Model: How People Cope with Persuasion Attempts. *Journal of Consumer Research*, 21(1), 1-31.
- Friestad, M., & Wright, P. (1995). Persuasion Knowledge: Lay People’s and Researchers’ Beliefs and Psychology of Advertising. *Journal of Consumer Research*, 22(1), 62-74.
- Frischmann, B., & Selinger, E. (2018). *Re-engineering Humanity*. Cambridge: Cambridge University Press.
- Gostin, L. O., Halabi, S. F., & Wilson, K. (2018). Health Data and Privacy in the Digital Era. *The Journal of the American Medical Association*, 320(8), 233-234.
- Gray, G. C. (2009). The Responsibilization Strategy of Health and Safety: Neo-liberalism and the Reconfiguration of Individual Responsibility for Risk. *The British Journal of Criminology*, 49(3), 326-342.
- Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The Dark (Patterns) Side of UX Design. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.
- Greenspan, P. (2003). The Problem with Manipulation. *American Philosophical Quarterly*, 40(2), 155-164.
- Habermas, J. (1987 [1981]). *The Theory of Communicative Actions. Volume 2: Lifeworld and System*, translated by T. McCarthy. Boston: Beacon Press.
- Hansen, P. G., & Jespersen, A. M. (2013). Nudging and the Manipulation of Choice: A Framework for the Responsible Use of the Nudge Approach to Behavior Change in Public Policy. *European Journal of Risk Regulation*, 4(1), 3-28.
- 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. (1996). Trustworthiness. *Ethics*, 107(1), 26-42.

- Hardin, R. (2002). *Trust and Trustworthiness*. New York: Russell Sage Foundation.
- Hasler, B. S., Tuchman, P., & Friedman, D. (2013). Virtual Research Assistants: Replacing Human Interviewers by Automated Avatars in Virtual Worlds. *Computers in Human Behavior, 29*, 1608-1616.
- Hayek, F. A. (2006 [1960]). Freedom and Coercion. In D. L. Miller (Ed.), *The Liberty Reader* (pp. 80-99). New York: Routledge.
- Heath, J. & Anderson, J. (2010). Procrastination and the Extended Will. In C. Andreou & M. D. White (Eds.), *The Thief of Time: Philosophical Essays on Procrastination* (pp. 233-252). Oxford: Oxford University Press.
- Helberger, N. (2016). Profiling and Targeting Consumers in the Internet of Things – A New Challenge for Consumer Law. In R. Schulze & D. Staudenmayer (Eds.), *Digital Revolution: Challenges for Contract Law in Practice* (pp. 135-162). Baden-Baden: Nomos.
- Helberger, N., Zuiderveen Borgesius, F., & Reyna, A. (2017). The Perfect Match? A Closer Look at the Relationship between EU Consumer Law and Data Protection Law. *Common Market Law Review, 54*(5), 1427-1465.
- Hesse-Biber, S. J. (1996). *Am I Thin Enough? The Cult of Thinness and the Commercialization of Identity*. New York, NY: Oxford University Press.
- Howells, G. (2006a). Aggressive Commercial Practices. In G. Howells, H.-W. Micklitz & T. Wilhelmsson (Eds.), *European Fair Trading Law: The Unfair Commercial Practices Directive* (pp. 167-194). London: Ashgate.
- Howells, G. (2006b). Introduction. In G. Howells, H.-W. Micklitz & T. Wilhelmsson (Eds.), *European Fair Trading Law: The Unfair Commercial Practices Directive* (pp. 1-26). London: Ashgate.
- Huckvale, K., Tomás Prieto, J., Tilney, M., Benghozi, P.-J., & Car, J. (2015). Unaddressed Privacy Risks in Accredited Health and Wellness Apps: A Cross-Sectional Systematic Assessment. *BMC Medicine, 13*(1), 214.
- Huckvale, K., Torous, J., & Larsen, M. E. (2019). Assessment of the Data Sharing Practices and Privacy Practices of Smartphone Apps for Depression and Smoking Cessation. *JAMA Network Open, 2*(4), e192542.
- Hurka, T. (1993). *Perfectionism*. Oxford: Oxford University Press.
- Hyman, D. A., Franklyn, D., Yee, C., & Rahmati, M. (2018). Going Native: Can Consumers Recognize Native Advertising? Does It Matter? *Yale Journal of Law and Technology, 19*(1), 77-112.
- Incardona, R., & Poncibò, C. (2007). The Average Consumer, the Unfair Commercial Practices Directive, and the Cognitive Revolution. *Journal of Consumer Policy, 30*(1), 21-38.
- IQVIA (2017). *The Growing Value of Digital Health. Evidence and Impact on Human Health and the Healthcare System*. Last accessed 22 September 2020 at <https://www.iqvia.com/insights/the-iqvia-institute/reports/the-growing-value-of-digital-health>.
- Jacobs, N. (2020). Two Ethical Concerns about the Use of Persuasive Technology for Vulnerable People. *Bioethics, 34*(5), 519-526.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Macmillan.
- Kang, J., & Buchner, B. (2004). Privacy in Atlantis. *Harvard Journal of Law and Technology, 18*(1), 229-267.

- Kaptein, M. (2015) *Persuasion Profiling: How the Internet Knows What Makes You Tick*. Amsterdam: Business Contact Publishers.
- Kaptein, M., Markopoulos, P., De Ruyter, B., & Aarts, E. (2015). Personalizing Persuasive Technologies: Explicit and Implicit Personalization Using Persuasion Profiles. *International Journal of Human-Computer Studies*, 77, 38–51.
- Kobrin, S. J. (2004). Safe Harbours are Hard to Find: The Trans-Atlantic Data Privacy Dispute, Territorial Jurisdiction and Global Governance. *Review of International Studies*, 30(1), 111-131.
- Krebs, P., Prochaska, J. O., & Rossi, J. S. (2010). A Meta-Analysis of Computer-Tailored Interventions for Health Behaviour Change. *Preventive Medicine*, 51(3), 214–221.
- Kroeze, W., Werkman, A., & Brug, J. (2006). A Systematic Review of Randomized Trials on the Effectiveness of Computer-Tailored Education on Physical Activity and Dietary Behaviors. *Annals of Behavioural Medicine*, 31(3), 205–223.
- Kumar, V. (2014). Making “Freemium” Work. *Harvard Business Review*, May 2014 Issue. Last accessed 22 September 2020 at <https://hbr.org/2014/05/making-freemium-work>.
- Lanzing, M. (2018). “Strongly Recommended”: Revisiting Decisional Privacy to Judge Hypernudging in Self-Tracking Technologies. *Philosophy & Technology*, 32(3), 549-568.
- Lanzing, M. (2019). *The Transparent Self: A Normative Investigation of Changing Selves and Relationships in the Age of the Quantified Self*. PhD Dissertation, Eindhoven University of Technology.
- Lee, M., Ackermans, S., Van As, N., Chang, H., Lucas, E., & IJsselsteijn, W. (2019). Caring for Vincent: A Chatbot for Self-Compassion. *CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3290605.3300932>
- Lewis, P. (2017). ‘Our Minds Can Be Hijacked’: The Tech Insiders Who Fear a Smartphone Dystopia. *The Guardian*, 6 October. Last accessed 22 September 2020 at <https://www.theguardian.com/technology/2017/oct/05/smartphone-addiction-silicon-valley-dystopia>
- Lohr, S. (2016). Goodbye, Ivory Tower. Hello, Silicon Valley Candy Store. *New York Times*, 3 September 2016. Last accessed 22 September 2020 at <https://www.nytimes.com/2016/09/04/technology/goodbye-ivory-tower-hello-silicon-valley-candy-store.html>
- Lorenzini, D. (2020). Biopolitics in the Time of Coronavirus. *Critical Inquiry*. Last accessed 2 June 2020 at <https://critinq.wordpress.com/2020/04/02/biopolitics-in-the-time-of-coronavirus/>.
- Lukes, S. (1974). *Power: A Radical View*. New York: Macmillan.
- 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). *Fat*. London: Routledge.
- Lupton, D. (2018). ‘I Just Want It to Be Done, Done, Done!’ Food Tracking Apps, Affects, and Agential Capacities. *Multimodal Technologies and Interaction*, 2(2). <https://doi.org/10.3390/mti2020029>.
- Lustria, M. L. A., Noar, S. M., Cortese, J., Van Stee, S. K., Glueckauf, R. L., & Lee, J. A. (2013). A meta-analysis of web-delivered tailored health behavior change interventions. *Journal of Health Communication*, 18(9), 1039–1069.

- Machkovech, S. (2017). Report: Facebook Helped Advertisers Target Teens Who Feel “Worthless”. *ArsTechnica*, 1 May 2017. Last accessed 19 May 2020 at <https://arstechnica.com/information-technology/2017/05/facebook-helped-advertisers-target-teens-who-feel-worthless/>
- Mackenzie, C. (2014). The Importance of Relational Autonomy and Capabilities for an Ethics of Vulnerability. In C. Mackenzie, W. Rogers & S. Dodds (Eds.), *Vulnerability: New Essays in Ethics and Feminist Philosophy* (pp. 33-59). Oxford: Oxford University Press.
- Mackenzie, C. (2016). Vulnerability, Needs and Moral Obligation. In: C. Straehle (Ed.), *Vulnerability, Autonomy, and Applied Ethics* (pp. 83-100). London: Routledge.
- Mackenzie, C. & N. Stoljar (Eds.) (2000a). *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*. Oxford: Oxford University Press.
- Mackenzie, C. & N. Stoljar (2000b). ‘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.
- Mackenzie, C., Rogers, W., & Dodds, S. (2014) Introduction: What Is Vulnerability, and Why Does It Matter for Moral Theory? In: C. Mackenzie, W. Rogers & S. Dodds (Eds.), *Vulnerability: New Essays in Ethics and Feminist Philosophy* (pp. 1-29). Oxford: Oxford University Press.
- Marwick, A. E. (2013). *Status Update: Celebrity, Publicity, and Branding in the Age of Social Media*. New Haven, CT: Yale University Press.
- Mau, S. (2019). *The Metric Society*. Cambridge: Polity Press.
- Mayer-Schönberger, V., & Cukier, K. (2013). *Big Data. A Revolution That Will Transform How We Live, Work, and Think*. New York: Houghton Mifflin Harcourt.
- Mayes, C. (2015). *The Biopolitics of Lifestyle: Foucault, Ethics and Healthy Choices*. London: Routledge.
- McLeod, C. (2015). Trust. In: E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Last accessed 22 September 2020 at <https://plato.stanford.edu/archives/fall2015/entries/trust/>
- McStay, A. (2016). Emphatic Media and Advertising: Industry, Policy, Legal and Citizen Perspectives (The Case for Intimacy). *Big Data & Society*, 3(2), 1-11.
- Micklitz, H.-W. (2006). The General Clause on Unfair Practices. In G. Howells, H.-W. Micklitz & T. Wilhelmsson (Eds.), *European Fair Trading Law: The Unfair Commercial Practices Directive* (pp. 83-122). London: Ashgate.
- Micklitz, H.-W., Sibony, A.-L., & Esposito, F. (Eds.) (2018). *Research Methods in Consumer Law: A Handbook*. Cheltenham: Edward Elgar Publishing.
- Micklitz, H.-W., Stuyck, J., & Terry, E. (Eds.) (2010). *Cases, Materials and Text on Consumer Law*. Oxford: Hart Publishing.
- Mill, J. S. (1859). *On Liberty*.
- Mills, C. W. (2005). “Ideal Theory” as Ideology. *Hypatia*, 20(3), 165-184.
- Mol, A. (2008). *The Logic of Care: Health and the Problem of Patient Choice*. London: Routledge.
- Muchagata, J., & Ferreira, A. (2018). Translating GDPR into the mHealth Practice. *2018 International Carnahan Conference on Security Technology (ICCST)*. IEEE. Last accessed 22 September 2020 at <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8585546>.

- Mulder, T. (2019). Health Apps, Their Privacy Policies and the GDPR. *European Journal of Law and Technology*, 10(1). Last accessed 22 September 2020 at <http://ejlt.org/article/view/667>.
- Mulder, T., & Tudorica, M. (2019). Privacy Policies, Cross-Border Health Data and the GDPR. *Information & Communications Technology Law*, 28(3), 261-274.
- Murnane, E. L., Huffaker, D., & Kossinets, G. (2015). Mobile Health Apps: Adoption, Adherence, and Abandonment. *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*, 261-264.
- Nancy, J. L. (2020). Communovirus. *Verso Books Blog*. Last accessed 2 June 2020 at <https://www.versobooks.com/blogs/4626-communovirus>.
- Navarro, V. (1976). *Medicine under Capitalism*. New York: Prodist.
- Nedelsky, J. (1989). Reconceiving Autonomy: Sources, Thoughts and Possibilities. *Yale Journal of Law and Feminism*, 1, 7-36.
- Nickel, P. J. (2015). Designing for the Value of Trust. In J. van der Hoven, P.E. Vermaas & I. Van der Poel (Eds.), *Handbook of Ethics, Values, and Technological Design* (pp. 551-567). Dordrecht: Springer.
- Nissenbaum, H. (2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford University Press.
- Noar, S. M., Benac, C. N., & Harris, M. S. (2007). Does tailoring matter? Meta-analytic review of tailored print health behavior change interventions. *Psychological Bulletin*, 133(4), 673-693.
- 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*, 40(4), 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.
- Nys, T. R. V., & Engelen, B. (2017). Judging Nudging: Answering the Manipulation Objection. *Political Studies*, 65(1), 199-214.
- O'Loughlin, K., Neary, M., Adkins, E. C., & Schueller, S. M. (2019). Reviewing the Data Security and Privacy Policies of Mobile Health Apps for Depression. *Internet Interventions*, 15, 110-115.
- Oh, K.-J., Lee, D., Ko, B., & Choi, H.-J. (2017). A Chatbot for Psychiatric Counseling in Mental Healthcare Services Based on Emotional Dialogue Analysis and Sentence Generation. *2017 18th IEEE International Conference on Mobile Data Management (MDM)*. DOI: 10.1109/MDM.2017.64.
- Olson, P. (2015). Under Armour Buys Health-Tracking App MyFitnessPal for \$475 Million. *Forbes*, February 4, 2015. Last accessed 22 September 2020 at <https://www.forbes.com/sites/parmyolson/2015/02/04/myfitnesspal-acquisition-under-armour/#1a75350c6935> (last accessed April 16, 2019).
- Oshana, M. A. L. (1998). Personal Autonomy and Society. *Journal of Social Philosophy*, 29(1), 81-102.

- Ostherr, K., Borodina, S., Conrad Bracken, R., Lotterman, C., Storer, E., & Williams, B. (2017). Trust and Privacy in the Context of User-Generated Health Data. *Big Data & Society*, 4(1), 1-11.
- Overdorf, R., Kulynych, 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. *TPRC 41: The 41st Research Conference on Communication, Information and Internet Policy*. Last accessed 22 September 2020 at <https://ssrn.com/abstract=2242144>
- Petersen, A., & Bunton, R. (Eds.) (1997). *Foucault, Health, and Medicine*. London: Routledge.
- Pettit, P. (1997). *Republicanism: A Theory of Freedom and Government*. Oxford: Clarendon Press.
- Pettit, P. (2012). *On the People's Terms: A Republican Theory and Model of Democracy*. Cambridge: Cambridge University Press.
- Pettit, P. (2018). Is Facebook Making Us Less Free? *iai*, 26 March 2018. Last accessed 3 June 2020 at <https://iai.tv/articles/the-big-brotherhood-of-digital-giants-is-taking-away-our-freedom-auid-884>.
- Quine, W. V. (1992). *Pursuit of Truth*, Revised Edition. Cambridge, MA: Harvard University Press.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Raz, J. (1986). *The Morality of Freedom*. Oxford: Clarendon Press.
- Rebonato, R. (2014). A Critical Assessment of Libertarian Paternalism. *Journal of Consumer Policy*, 37(3), 357-396.
- Reiman, J. (1994). Driving to the Panopticon: A Philosophical Exploration of the Risk to Privacy Posed by the Highway Technology of the Future. *Santa Clara High Technology Law Journal*, 11(1), 27-44.
- Rieder, B. (2015). On the Diversity of the Accountability Problem – Machine Learning and Knowledge Capitalism. Paper presented at meeting of the NYU Information Law Institute, Algorithms and Accountability Conference, New York City, NY.
- Roessler, B. (2005). *The Value of Privacy*. Cambridge: Polity Press.
- Roessler, B. (2017). *Autonomie. Ein Versuch über das gelungene Leben*. Berlin: Suhrkamp.
- Roessler, B. (2021). *Autonomy. An Essay on the Life Well-Lived*. Cambridge: Polity Press.
- Rogers, W. A., Mackenzie, C., & Dodds, S. (2012). Why Bioethics needs a Concept of Vulnerability. *International Journal of Feminist Approaches to Bioethics*, 5(2), 11-38.
- Rossiter, K. (2012). Talking Turkey: Anxiety, Public Health Stories, and the Responsibilization of Health. *Journal of Canadian Studies*, 46(2), 178-195.
- Rudinow, J. (1978). Manipulation. *Ethics*, 88(4), 338-347.
- Saghai, Y. (2013). Salvaging the Concept of Nudge. *Journal of Medical Ethics*, 39(8), 487-493.
- Sandel, M. J. (1984). The Procedural Republic and the Unencumbered Self. *Political Theory*, 12(1), 81-96.
- 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. (2016). Big Data: Finders Keepers, Losers Weepers? *Ethics and Information Technology*, 18(1), 25-31.

- Sax, M. (2017). Let op de subtiele vermenging van commercie en adviezen in gezondheidsapps. *De Volkskrant*, 25 March 2017. Last accessed 22 September 2020 at <https://www.volkskrant.nl/columns-opinie/let-op-de-subtiele-vermenging-van-commercie-en-adviezen-in-gezondheidsapps~b2fe4387/>.
- 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.
- Schulze, R., & Schulte-Nölke, H. (2003). Analysis of national fairness laws aimed at protecting consumers in relation to commercial practices. Brussels: European Commission. Last accessed 22 September 2020 at https://lirias.kuleuven.be/bitstream/123456789/204413/1/unfair_practices_en.pdf.
- Seufert, E. B. (2014). *Freemium Economics: Leveraging Analytics and User Segmentation to Drive Revenue*. Waltman, MA: Morgan Kaufmann.
- Shamir, R. (2008). The Age of Responsibilization: On Market-Embedded Morality. *Economy and Society*, 37(1), 1-19.
- Shaw, T. (2017). Invisible Manipulators of Your Mind. *The New York Review of Books*, April 20. Last accessed 22 September 2020 at <https://www.nybooks.com/articles/2017/04/20/kahneman-tversky-invisible-mind-manipulators/>
- Sparke, M. (2016). Health and the Embodiment of Neoliberalism: Pathologies of Political Economy from Climate Change and Austerity to Personal Responsibilization. In S. Springer, K. Birch & J. MacLeavy (Eds.), *The Handbook of Neoliberalism* (pp. 237-251). New York: Routledge.
- Spencer, S. B. (2019). The Problem of Online Manipulation. Last accessed 22 September 2020 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.
- Straehle, C. (Ed.) (2017). *Vulnerability, Autonomy and Applied Ethics*. London: Routledge.
- Straehle, C. (2017). Introduction: Vulnerability, Autonomy and Applied Ethics. In C. Straehle (Ed.), *Vulnerability, Autonomy and Applied Ethics* (pp. 1-10). London; Routledge.
- Sunstein, C. R. (2014). *Why Nudge? The Politics of Libertarian Paternalism*. New Haven, NJ: Yale University Press.
- Sunstein, C. R. (2015a). *Choosing Not to Choose: Understanding the Value of Choice*. Oxford: Oxford University Press.
- Sunstein, C. R. (2015b). Nudges, Agency, and Abstraction: A Reply to Critics. *Review of Philosophy and Psychology*, 6(3), 511-529.
- Sunstein, C. R. (2016). *The Ethics of Influence: Government in the Age of Behavioral Science*. Cambridge: Cambridge University Press.
- Sunstein, C. R., & Thaler, R. H. (2003). Libertarian Paternalism Is Not an Oxymoron. *The University of Chicago Law Review*, 70(4), 1159-1202.
- 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). DOI: 10.14763/2019.2.1410.

- Tene, O., & Polonetsky, J. (2013). Big Data for All: Privacy and User Control in the Age of Analytics. *Northwestern Journal of Technology and Intellectual Property*, 11(5), 239-273.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth and Happiness*. New Haven, CT: Yale University Press.
- The Australian (2017). Facebook Targets Insecure Young People to Sell Ads. Last accessed 22 September 2020 at <http://www.theaustralian.com.au/business/media/digital/facebook-targets-insecure-young-people-to-sell-ads/news-story/a89949ad016eee7d7a61c3c30c909fa6>.
- Turow, J. (2011). *The Daily You: How the New Advertising Industry is Defining Your Identity and Your Worth*. New Haven: Yale University Press.
- Turow, J., Feldman, L., & Meltzer, K. (2005). *Open to Exploitation: America's Shoppers Online and Offline*. Report from the Annenberg Public Policy Center of the University of Pennsylvania. Last accessed 26 September 2020 at http://repository.upenn.edu/asc_papers/35.
- Wang, K., Varma, D.S., & Prospero, M. (2018). A Systematic Review of the Effectiveness of Mobile Apps for Monitoring and Management of Mental Health Symptoms or Disorders. *Journal of Psychiatric Research*, 107, 73-78.
- Wang, Q., Egelandstad, B., Amdam, G. V., Almlil, V. L., & Oostindjer, M. (2016). Diet and Physical Activity Apps: Perceived Effectiveness by App Users. *JMIR Mhealth Uhealth*, 4(2), e33.
- Whitehead, L., & Seaton, P. (2016). The Effectiveness of Self-Management Mobile Phone and Tablet Apps in Long-term Condition Management: A Systematic Review. *Journal of Medical Internet Research*, 18(5), e97.
- Whitman, J. Q. (2004). Two Western Cultures of Privacy: Dignity versus Liberty. *The Yale Law Journal*, 113, 1151-1221.
- Wiesemann, C. (2017). On the Interrelationship of Vulnerability and Trust. In C. Straehle (Ed.), *Vulnerability, Autonomy and Applied Ethics* (pp. 157-170). London; Routledge.
- Wilhelmsson, T. (2006). Scope of the Directive. In G. Howells, H.-W. Micklitz & T. Wilhelmsson (Eds.), *European Fair Trading Law: The Unfair Commercial Practices Directive* (pp. 49-81). London: Ashgate.
- Wilkinson, T. M. (2013). Nudging and Manipulation. *Political Studies*, 61(2), 341-355.
- Willett, C. (2010). Fairness and Consumer Decision Making under the Unfair Commercial Practices Directive. *Journal of Consumer Policy*, 33(3), 247-273.
- Wisniewski, H., Liu, G., Henson, P., Vaidyam, A., Hajratalli, N.R., Onnela, J.-P., & Torous, J. (2019). Understanding the Quality, Effectiveness and Attributes of Top-Rated Smartphone Health Apps. *Evidence-Based Mental Health*, 22, 4-9.
- 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.
- Wooldrige, D., & Schneider, M. (2011). *The Business of iPhone and iPad App Development: Making and Marketing Apps That Succeed*, Second Edition. New York: Springer Science+Business Media.

- World Health Organization (WHO) (2011). *mHealth. New Horizons for Health through Mobile Technologies*. Geneva: World Health Organization.
- World Health Organization (2015). *Be He@lthy Be Mobile: A Handbook on How to Implement mTobaccoCessation*. Geneva: World Health Organization. Last accessed 22 September 2020 at [https://www.who.int/publications/i/item/mobile-health-for-tobacco-cessation-\(mtobacco cessation\)](https://www.who.int/publications/i/item/mobile-health-for-tobacco-cessation-(mtobacco cessation)).
- World Health Organization (2016a). *Be He@lthy Be Mobile: A Handbook on How to Implement mDiabetes*. Geneva: World Health Organization. Last accessed 22 September 2020 at [https://www.who.int/publications/i/item/mobile-health-for-diabetes-prevention-and-management-\(mdiabetes\)](https://www.who.int/publications/i/item/mobile-health-for-diabetes-prevention-and-management-(mdiabetes)).
- World Health Organization (2016b). *Be He@lthy Be Mobile: A Handbook on How to Implement mCervicalCancer*. Geneva: World Health Organization. Last accessed 22 September 2020 at [https://www.who.int/publications/i/item/mobile-health-for-cervical-cancer-\(mcervicalcancer\)](https://www.who.int/publications/i/item/mobile-health-for-cervical-cancer-(mcervicalcancer)).
- World Health Organization (2018a). *Be He@lthy Be Mobile: A Handbook on How to Implement mBreatheFreely*. Geneva: World Health Organization. Last accessed 22 September 2020 at [https://www.who.int/publications/i/item/mobile-health-for-asthma-and-chronic-obstructive-respiratory-disease-\(mbreathefreely\)](https://www.who.int/publications/i/item/mobile-health-for-asthma-and-chronic-obstructive-respiratory-disease-(mbreathefreely)).
- World Health Organization (2018b). *Be He@lthy Be Mobile: A Handbook on How to Implement mAgeing*. Geneva: World Health Organization. Last accessed 22 September 2020 at <https://www.who.int/publications/i/item/be-healthy-be-mobile-a-handbook-on-how-to-implement-mageing>.
- World Health Organization (2019a). *Be He@lthy Be Mobile: A Handbook on How to Implement mTB-Tobacco*. Geneva: World Health Organization. Last accessed 22 September 2020 at <https://www.who.int/publications/i/item/a-handbook-on-how-to-implement-mtb-tobacco>.
- World Health Organization (2019b). *Be He@lthy Be Mobile: Personas Toolkit*. Geneva: World Health Organization. Last accessed 22 September 2020 at <https://www.who.int/publications/i/item/be-healthy-be-mobile-personas-toolkit>.
- World Health Organization (2020). *Be He@lthy Be Mobile: A Handbook on How to Implement mHypertension*. Geneva: World Health Organization. Last accessed 22 September 2020 at <https://www.who.int/publications/i/item/be-healthy-be-mobile-a-handbook-on-how-to-implement-mhypertension>.
- Yeung, K. (2012). Nudge as Fudge. *The Modern Law Review*, 75(1), 122-148.
- 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 Prospects 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.
- Zuiderveen Borgesius, F. J., & Poort, J. (2017). Online Price Discrimination and EU Data Privacy Law. *Journal of Consumer Policy*, 40(3), 347-366.
- Zwick, D., & Dholakia, N. (2001). Contrasting European and American Approaches to Privacy in Electronic Markets: Property Right versus Civil Right. *Electronic Markets*, 11(2), 116-120.

Summary

This is a dissertation about health apps. At the same time, it is a dissertation about a larger, more general development of which health apps are but a symptom. We are currently living in a – or: the – digital society. One important feature of the digital society, in which an increasing amount of daily activities (partly) take place in digital environments, is the fact that the designers and operators of these digital environments become increasingly powerful actors. They are *choice architects* who design and operate *choice architectures*. The digital environments they design can be very helpful and empowering. However, all the features of these digital environments which make them potentially helpful and empowering also contain a potential for *manipulation*. In this dissertation, I explore the – often difficult to draw – boundaries between empowering and manipulative (features of) digital choice environments.

To address this abstract challenge, I choose to focus on a more specific phenomenon, namely for-profit health apps. Focusing on for-profit health apps is interesting for at least two reasons. First, by focusing on *one* phenomenon, I can describe and analyze in much more detail how the larger challenge I seek to address plays out on the ground, in the real world. This will allow me to offer a more grounded, and, hopefully, more relevant analysis. Second, and maybe even more importantly, for-profit health apps are incredibly interesting in their own right. As a value, health is universally desired and health is – to varying degrees of course – a concern to every human being. Health, in short, *matters to everyone*. It follows that focusing on digital *health* environments, which seek to influence the *health*(-related) behavior of persons, matters too. Moreover, health apps come with a promise of *empowerment*, but also, at the same time, with a potential for manipulation.

To explore the boundaries between empowerment and manipulation in the digital health environments that for-profit health apps provide to their users, I first *describe* how such digital health environments are designed and operated. I then, in a next step, develop an ethical framework which should help with the ethical *evaluation* of digital health environments. In a last step, I use the insights from my ethical evaluation to explore the possibilities of addressing manipulation worries in the health app context through European consumer law.

In **chapter 1** I introduce the phenomenon of health apps. Here I discuss the technological capabilities of health apps as well as the larger app economy of which they are a part. We need to understand the dynamics of the app economy to understand how the technological capabilities of health apps – and the digital environments they offer to users – are designed and operated. The contemporary app economy incentivizes app developers to build *ongoing relationships* with users and monetize those relationships *over time*. To successfully monetize one's userbase, there is an incentive for health apps to design their digital health environment to optimize user retention, user engagement, and user conversion. Health apps, like other apps, can use their near complete power over the design and operation of the digital environments they offer to try to give shape to the user-app relationships. For example, health apps can collect user data in order to learn more about their users. These insights can be used to constantly experiment with and update the digital environment.

In **chapter 2**, I explore the notion of 'health' and how it is used in the health app context. It is clear that health is universally desired. When given the choice, everyone prefers health over illness. However, when we ask what it means to be healthy, it turns out that health is a contested concept. This conceptual fluidity might be regrettable to philosophers aiming for conceptual rigidity, but it can also be put to strategic use by health apps: it affords them the opportunity to allude to health's universally desired nature while remaining vague regarding their contributions to their users' health. In this chapter, I argue that health apps construct and propagate particular health discourses and narratives which serve their business interests. So to understand why health apps are designed and run the way they are, we must not only understand their business models and technological affordances, but also their portrayal of health.

In **chapter 3**, I provide the conceptual foundation for the ethical framework I use to evaluate the commercial practices of for-profit health apps. I build on four interrelated concepts: autonomy, vulnerability, trust, and manipulation. Why these four concepts? Put simply, autonomy is the value that is at stake in digital environments that shape human behavior. The concept of manipulation is promising because it can help explain *how* and *why* autonomy can be at stake in the context of health apps. Lastly, vulnerability and trust serve as important conceptual links between autonomy and manipulation. I mainly build on Susser, Roessler, and Nissenbaum (2019a, 2019b) to argue that, in essence, manipulation is about the intentional infiltration of a manipulee's decision-making in order to dispose the manipulee to the manipulator's ends, by way of targeting known or inferred weaknesses of said manipulee in a manner that is not meant to be

(fully) transparent to the manipulee. Put simply, when a person is influenced in a manipulative manner, their autonomy is threatened. Theories of vulnerability can help us get a better, more nuanced understanding of how people's weaknesses and biases can be exploited by manipulative practices. Attention for the role of trust is important because health apps (and digital for-profit services more generally) seek to build trust with users, since it helps to build profitable relationships with them. Although trust in technology can be helpful, it can also be exploited in a manipulative manner.

In **chapter 4**, I use the conceptual framework from chapter 3 to provide an ethical evaluation of the (commercial) practices of health apps. To understand the manipulative potential of popular for-profit health apps and to understand how the autonomy of health app users can be at stake, we must not only focus on isolated one-off interactions between a health app and its users. It is equally important to analyze how health apps use their control over the digital health environment they offer to build *ongoing* relationships with their users over time. In the first part of the chapter I explore how health apps could *scaffold* their users' autonomy (Heath & Anderson 2010). The collection and analysis of user data to learn about users' vulnerabilities can, for instance, be used to better anticipate what individual users need. Moreover, insights into vulnerabilities and preferences can be used to engineer feelings of trust towards a health app, making it easier for the app in question to help users. We can thus imagine that under (somewhat) ideal conditions, health apps can empower their users. In the second part of the chapter I explore the manipulative potential of health apps. The properties of health apps discussed in the first part of the chapter, can, under different conditions, turn from being empowering into being manipulative. An important source for health apps' manipulative potential is the freemium app economy's emphasis on optimizing user retention, user engagement, and user conversion. When a health app engineers its digital health environment to identify and target vulnerabilities of users for the purpose of engineering profitable relationships which will optimize retention, engagement, and conversion rates, most ingredients for manipulation are present. In practice, the line between empowerment and manipulation can be difficult to draw. Those properties that make health apps potentially useful also put them in a position of power vis-à-vis their users, and that position of power can also be *misused* for commercial practices that are manipulative in nature.

Chapter 5 is devoted to the question of how we can address – legally – the difficult tension between empowerment and manipulation. I focus on the UCPD and present an interpretation of the Directive that is informed by my ethical analysis. At its core, the UCPD is concerned with protecting the possibilities for autonomous

decision-making of consumers. I try to show how a more relational understanding of autonomy, coupled with a more nuanced understanding of vulnerability, can and should inform our reading of the Directive. Theoretically, the concept of misleading commercial practices seems to align rather well with the challenges posed by manipulative practices; in cases of manipulation, manipulees are often (but not necessarily) “led astray” (Noggle 1996: 44) or, in other words, misled. While the concept of misleading practices and the corresponding call for better transparency requirements are helpful, I argue that the concept of aggressive commercial practices – and more specifically the concept of undue influence – seems (even) more promising. By focusing on the position of power of health app providers and the possible *exploitation* of that position of power by way of targeting exploitable characteristics of users, we can get much closer to the *heart* of the manipulation problem. The overall conclusion of the chapter is that if we are willing to inform our interpretations of key concepts in the Directive with my ethical analysis, the UCPD can be a promising legal tool to address manipulative practices in digital (health) environments.



Samenvatting

Dit proefschrift gaat over gezondheidsapps. Tegelijkertijd gaat dit proefschrift ook over een grotere en meer algemene ontwikkeling waar gezondheidsapps slechts een onderdeel van zijn. We leven momenteel in een – of: de – digitale samenleving. Een belangrijk aspect van onze digitale samenleving waarin steeds meer dagelijkse bezigheden (deels) plaatsvinden in digitale omgevingen, is de groeiende macht van de vormgevers van die digitale omgevingen. Zij zijn zogenaamde *keuze-architecten* die *keuze-architecturen* inrichten en exploiteren. De digitale omgevingen die zij inrichten kunnen erg behulpzaam en *empowering* zijn. Echter, precies die eigenschappen van digitale omgevingen die hen mogelijk behulpzaam en *empowering* maken, dragen ook al een potentie tot manipulatie in zich. In dit proefschrift verken ik de vaak lastig te trekken grens tussen behulpzame en manipulatieve (eigenschappen van) digitale keuze-architecturen.

Om deze abstracte uitdaging bij de hoorns te vatten, richt ik mij op één specifiek fenomeen: for-profit gezondheidsapps. Een focus op for-profit gezondheidsapps is om ten minste twee redenen interessant. Ten eerste, door mij op één fenomeen te richten kan ik in meer detail en met meer precisie onderzoeken hoe een algemene ontwikkeling in de praktijk, in de echte wereld, zijn uitwerking heeft. Hierdoor kan ik een meer ‘geaarde’ en, hopelijk, bruikbaarere analyse uitvoeren. Ten tweede, en misschien wel belangrijker, zijn for-profit gezondheidsapps ook al van zichzelf razend interessant. Gezondheid is een universeel nagestreefde waarde en speelt – zij het in verschillende mate – een belangrijke rol in elk mensenleven. Gezondheid is voor *iedereen* van belang. Hieruit volgt dat een analyse met een focus op digitale *gezondheidsomgevingen* die aan *gezondheid* gerelateerde gedragspatronen proberen te beïnvloeden, ook van groot belang is. Bovendien vindt men in gezondheidsapps ook zowel de belofte van *empowerment* als een duidelijk manipulatiepotentieel.

Om deze spanning tussen *empowerment* en manipulatie die men ook aantreft in gezondheidsapps te onderzoeken, beschrijf ik in dit proefschrift eerst hoe zulke digitale gezondheidsomgevingen zijn vormgegeven en hoe zij ingezet kunnen worden om gedrag(spatronen) vorm te geven. Daarna, in een volgende stap, ontwikkel ik een ethisch raamwerk aan de hand waarvan digitale gezondheidsomgevingen geëvalueerd kunnen worden. In een laatste stap gebruik ik de inzichten van de ethische evaluatie om te verkennen hoe, en zo ja in welke mate, manipulatie-*risico's* aangepakt kunnen worden door het Europese consumentenrecht.

In **hoofdstuk 1** introduceer ik het fenomeen gezondheidsapps. Ik bespreek de technologische mogelijkheden van gezondheidsapp alsook de bredere app-economie waar gezondheidsapps onderdeel van uitmaken. Een goed begrip van de dynamiek van de app-economie is belangrijk, omdat de technologische mogelijkheden van gezondheidsapps – en de digitale omgevingen die zij aan gebruikers aanbieden – mede worden *gevormd* door de app-economie. De huidige app-economie spoort app-ontwikkelaars aan om doorlopende relaties met gebruikers op te bouwen en die doorlopende relaties door de tijd heen te gelde te maken. Om het gebruikersbestand van een app op succesvolle wijze te gelde te maken, worden app-ontwikkelaars aangespoord om hun digitale gezondheidsomgeving zo in te richten dat gebruikers*engagement*, gebruikersretentie en gebruikersconversie geoptimaliseerd worden. Gezondheidsapps kunnen, net als elk ander type app, hun bijna volledige macht over de inrichting van de digitale gezondheidsomgevingen die ze aanbieden, gebruiken om de relatie tussen app en gebruiker vorm te geven. Neem bijvoorbeeld de gebruikersdata die gezondheidsapps kunnen verzamelen en analyseren om meer over hun gebruikers te weten te komen. De verkregen inzichten kunnen gebruikt worden om verdere, constant doorlopende experimenten te voeden, waarvan de uitkomsten op hun beurt weer gebruikt kunnen worden om de digitale omgeving (subtiel) aan te passen.

In **hoofdstuk 2** verken ik hoe het begrip ‘gezondheid’ wordt gebruikt in de gezondheidsappcontext. Het moge duidelijk zijn dat iedereen verlangt een gezond leven te leiden. Als men zou kunnen kiezen, zou met altijd gezondheid boven ziekte verkiezen. Dat betekent echter niet dat ‘gezondheid’ een helder begrip is. Sterker nog, ‘gezondheid’ is een klassiek voorbeeld van een *contested concept*. De conceptuele fluiditeit die hiermee gepaard gaat is lastig voor filosofen die conceptuele rigiditeit nastreven. Maar voor de aanbieders van gezondheidsapps biedt deze conceptuele fluiditeit ook een uitkomst: het stelt hen in staat om tegenover hun (toekomstige) gebruikers te refereren aan het vage maar aantrekkelijke containerbegrip van gezondheid, zonder direct heel precies te hoeven uitleggen *hoe* hun service nu precies bijdraagt aan ‘de gezondheid’. In dit hoofdstuk beargumenteer ik dat gezondheidsapps vaak specifieke gezondheidsdiscoursen en -narratieven creëren en propageren met als voornaamste doel om hun bedrijfsdoelen te ondersteunen. Als we dus willen begrijpen waarom gezondheidsapps zo ontworpen en ingezet worden als nu het geval is, dan moeten we niet alleen naar de onderliggende technologische mogelijkheden kijken, maar ook naar de manier waarop ze het begrip ‘gezondheid’ vormgeven.

In **hoofdstuk 3** formuleer ik het conceptuele fundament van het ethische raamwerk dat ik later gebruik om de handelspraktijken van for-profit gezondheidsapps te evalueren. Ik gebruik (voornamelijk) vier aan elkaar gerelateerde concepten,

namelijk: autonomie, kwetsbaarheid, vertrouwen en manipulatie. Waarom deze vier concepten? Simpel gezegd is autonomie de waarde die op het spel staat in digitale omgevingen die menselijk gedrag proberen vorm te geven. Het concept manipulatie is veelbelovend omdat het goed kan vatten en uitleggen *hoe* en *waarom* autonomie in de gezondheidsappcontext op het spel staat. Kwetsbaarheid en vertrouwen, tot slot, zijn belangrijke conceptuele verbindingen tussen autonomie en manipulatie. Ik borduur vooral voort op het werk van Susser, Roessler en Nissenbaum om te beargumenteren dat manipulatie in essentie draait om het intentioneel infiltreren van het besluitvormingsproces van andere mensen teneinde hen in de richting van jouw eigen belangen te sturen, op een manier die zich kenmerkt door het op verkapte wijze uitbuiten van bekende of veronderstelde uitbuitbare zwakheden. Wanneer iemands gedrag op manipulatieve wijze beïnvloed wordt, zou men simpel gezegd kunnen stellen dat zijn of haar autonomie in het geding is. Theorieën van kwetsbaarheid helpen ons om beter in beeld te krijgen hoe de kwetsbaarheden en biases van mensen uitgebuit kunnen worden door manipulatieve praktijken. Aandacht voor de rol die vertrouwen speelt is belangrijk omdat gezondheidsapps (en meer algemeen gesteld digitale for-profit diensten) vaak proberen om een vertrouwensband met hun gebruikers op te bouwen omdat dat hen beter in staat stelt winstgevende relaties met hun gebruikers aan te gaan. Daarmee is niet gezegd dat vertrouwen in technologie onbelangrijk of schadelijk is, maar het is wel belangrijk om in het oog te houden dat vertrouwen ook op manipulatieve wijze misbruikt kan worden.

In **hoofdstuk 4** maak ik gebruik van het conceptuele raamwerk beschreven in hoofdstuk 3 om een ethische analyse van de (handels-) praktijken van gezondheidsapps te formuleren. Om het manipulatiepotentieel van for-profit gezondheidsapps te begrijpen en om te zien hoe de autonomie van gezondheidsappgebruikers op het spel staat, is het essentieel om niet alleen te kijken naar eenmalige interacties tussen een gebruiker en een app. Het is even belangrijk om na te gaan hoe gezondheidsapps hun controle over de digitale omgevingen die ze aanbieden gebruiken om *doorlopende* relaties met hun gebruikers op te bouwen. In het eerste deel van dit hoofdstuk onderzoek ik hoe gezondheidsapps de autonomie van hun gebruikers, om met Heath en Anderson te spreken, kunnen *scaffolden*, oftewel ‘in de steigers kunnen zetten’. Om een voorbeeld te noemen: het verzamelen en analyseren van gebruikersdata om zo de kwetsbaarheden van gebruikers te leren kennen, kan gebruikt worden om de individuele behoeften van gebruikers beter te bedienen. Bovendien kunnen inzichten in kwetsbaarheden en voorkeuren gebruikt worden om (gevoelens van) vertrouwen in de gebruiker op te wekken, wat het makkelijker kan maken om de gebruiker te helpen. Het is dus niet moeilijk om je voor te stellen dat een gezondheidsapp onder (ietwat) geïdealiseerde omstandigheden de gebruikers kan *empoweren*. In het tweede gedeelte van dit hoofdstuk verken in het manipulatiepotentieel van gezondheidsapps.

Alle eigenschappen van gezondheidsapps die ik in het eerste deel van het hoofdstuk heb besproken, kunnen onder andere omstandigheden van *empowering* omslaan in manipulatief. Een belangrijke bron van het manipulatiepotentieel van gezondheidsapps is de sterke nadruk die de *freemium*-economie legt op het optimaliseren van gebruikersretentie, gebruikers*engagement* en gebruikersconversie. Wanneer een gezondheidsapp bewust de digitale omgeving inricht op een manier die erop gericht is kwetsbaarheden van gebruikers te identificeren en uit te buiten, met als doel om winstgevende relaties met gebruikers op te bouwen waarmee gebruikersretentie, gebruikers*engagement* en gebruikersconversie worden geoptimaliseerd, dan zijn de meeste ingrediënten voor manipulatie aanwezig. In de praktijk blijft het altijd lastig om een duidelijke lijn te trekken tussen *empowerment* en manipulatie. Precies die eigenschappen van gezondheidsapps die ze potentieel zeer behulpzaam maken, zijn dezelfde eigenschappen die gezondheidsapps in een machtspositie ten opzichte van hun gebruikers kunnen plaatsen. En het is die machtspositie die ook *misbruikt* kan worden voor manipulatieve handelspraktijken.

Hoofdstuk 5 is gewijd aan de vraag hoe de spanning tussen *empowerment* en manipulatie juridisch begrepen en geadresseerd kan worden. Ik richt mij op de Europese Richtlijn oneerlijke handelspraktijken (EROH) en presenteer een door mijn ethische analyse geïnspireerde interpretatie van de Richtlijn. In de kern draait de EROH om het beschermen van de mogelijkheid tot het maken van autonome beslissingen door consumenten. Ik probeer te laten zien hoe een meer relationeel begrip van autonomie en een meer genuanceerde opvatting van kwetsbaarheid kan leiden tot een meer bruikbare lezing van de EROH. Vanuit theoretisch perspectief lijkt het concept van ‘misleidende handelspraktijken’ het best aan te sluiten op de manipulatiezorgen die ik in eerdere hoofdstukken beschrijf; in gevallen van manipulatie worden de gemanipuleerden vaak (maar niet altijd) “om de tuin geleid”, of, in andere woorden, misleid. En hoewel het begrip van misleidende handelspraktijken en de daarbij horende roep om meer transparantievereisten nuttig zijn, beargumenteer ik dat het concept van ‘agressieve handelspraktijken’ – en meer in het bijzonder het concept van ‘ongepaste beïnvloeding’ – waarschijnlijk meer gewicht in de schaal kan leggen. Door aandacht te besteden aan de machtspositie van gezondheidsapp-aanbieders en de mogelijke *uitbuiting* van die machtspositie door kwetsbaarheden van gebruikers te identificeren en uit te buiten, komen we veel dichterbij de kern van het manipulatieprobleem. De algehele conclusie van dit hoofdstuk luidt dan ook dat de EROH in potentie zeer toepasselijke juridische middelen aanreikt om manipulatie in de digitale samenleving te adresseren, maar dat daarvoor wel nieuwe interpretaties – in lijn met mijn ethische analyse – van kernconcepten in de Richtlijn noodzakelijk zijn.

