

“Strongly recommended” revisiting decisional privacy to judge hypernudging in self-tracking technologies

Citation for published version (APA):

Lanzing, M. (2019). “Strongly recommended” revisiting decisional privacy to judge hypernudging in self-tracking technologies. *Philosophy & Technology*, 32(3), 549-568. <https://doi.org/10.1007/s13347-018-0316-4>

DOI:

[10.1007/s13347-018-0316-4](https://doi.org/10.1007/s13347-018-0316-4)

Document status and date:

Published: 15/09/2019

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the “Taverne” license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

“Strongly Recommended” Revisiting Decisional Privacy to Judge Hypernudging in Self-Tracking Technologies

Marjolein Lanzing¹

Received: 4 October 2017 / Accepted: 22 May 2018 / Published online: 6 June 2018

© The Author(s) 2018

Abstract This paper explores and rehabilitates the value of decisional privacy as a conceptual tool, complementary to informational privacy, for critiquing personalized choice architectures employed by self-tracking technologies. Self-tracking technologies are promoted and used as a means to self-improvement. Based on large aggregates of personal data and the data of other users, self-tracking technologies offer personalized feedback that nudges the user into behavioral change. The real-time personalization of choice architectures requires continuous surveillance and is a very powerful technology, recently coined as “hypernudging.” While users celebrate the increased personalization of their coaching devices, “hypernudging” technologies raise concerns about manipulation. This paper addresses that intuition by claiming that decisional privacy is at stake. It thus counters the trend to solely focus on informational privacy when evaluating information and communication technologies. It proposes that decisional privacy and informational privacy are often part of a mutually reinforcing dynamic. Hypernudging is used as a key example to illustrate that the two dimensions should not be treated separately. Hypernudging self-tracking technologies compromise autonomy because they violate informational and decisional privacy. In order to effectively judge whether technologies that use hypernudges empower users, we need both privacy dimensions as conceptual tools.

Keywords Ethics of privacy · Decisional privacy · Informational privacy · Autonomy · Hypernudging · Self-tracking · Personalized feedback · Big data · Surveillance

✉ Marjolein Lanzing
m.lanzing@tue.nl

¹ University of Technology Eindhoven, Het Eeuwsel 57, IPO 1.13, 5600 MB, 5612 AZ Eindhoven, The Netherlands

“What if your Fitbit knew exactly what to say on a particular day to motivate you to get off the couch and run a 5K?”

- Persado, Schwab 2017

1 Introduction

New technologies that use our data in order to steer our behavior are often accompanied by worries about (mass)-manipulation. Uber’s (offline) collection of real-time data in order to predict your next ride and tailor on-on-the-go recommendations (sushi or noodles?) based on one’s location and past choices makes us uneasy (Schlosser 2016). The Facebook experiment, in which tampered newsfeeds influenced the behavior of users, sparked outrage (Rushe 2014) and visualizing a FitBit that uses personalized nudges to coach the user into “healthy” behavior is met with suspicion (Schwab 2017). Moreover, the recent “fake news” controversy surrounding Facebook and Cambridge Analytica reignited the debate about the manipulative aspects of data-driven personalized communication and behavioral targeting in the online realm (Citron and Pasquale 2014; Hildebrandt 2008; Pariser 2011; Turov 2011; Zittrain 2014; Zuiderveen Borgesius et al. 2016; Zuboff 2015). Yet, drawing in Big Data to nudge individuals with personalized feedback to change their behavior, or “hypernudging,” is the latest feature of many new technologies. The new frontier, and the subject of this article, is self-tracking (Danaher 2016: 3–4; Galic et al., 2017: 30). Fuelled by real-time data, algorithms create personalized online choice architectures that aim to nudge individual users to effectively change their behavior. The question arises to what extent the data-driven personalized recommendations of coaching technologies are in fact empowering.

In this paper, I criticize the (potential) use of hypernudging in the field of self-tracking. I focus on self-tracking technologies because most people wear them precisely because of the personalized feedback they offer. If my critique succeeds, it follows that information and communication technologies (ICTs) that hypernudge users without their knowledge, such as Facebook, are ethically problematic too. I criticize self-tracking technologies from an informational privacy and decisional privacy perspective. The aim is to explore and rehabilitate the importance of decisional privacy as a conceptual tool to carry out this critique and to counter the trend to focus solely on informational privacy when evaluating ICTs by emphasizing the relationship between surveillance and decisional interference (Roessler 2005; Koops et al. 2017).

The claim of this paper is that hypernudges compromise autonomy because they violate both informational and decisional privacy as complementary dimensions. I support this claim in four steps. First, I argue that the type of personalized feedback offered by self-tracking technologies should be interpreted as hypernudging. Building on Karen Yeung (2017), who coined the term, I define hypernudging and distinguish it from “regular” nudging. Subsequently, I show how its features of extensive surveillance, hiddenness, and predictive capacity increase its potential for unjustified interference by going beyond the safeguards of “good” nudging. Secondly, I explore the concept of decisional privacy as a complementary dimension to informational privacy

and its value as a conceptual tool for evaluating hypernudging by drawing on research by Beate Roessler (2005), Jean Cohen (2002), and Bert-Jaap Koops et al. (2017). Third, I argue that in order to address our intuitions about the manipulation involved in hypernudging and evaluate whether hypernudging compromises autonomy, we should interpret this phenomenon from *both* an informational and decisional privacy perspective. Informational and decisional privacy are part of a mutually reinforcing dynamic—rather than separate types. Fourth, I raise and counter three potential objections to my argument. Finally, I conclude that self-tracking technologies that use hypernudging compromise a user’s autonomy, because they violate both informational and decisional privacy. Interestingly, it seems that while self-tracking technologies promise to empower users, they simultaneously compromise their autonomy in a different way. Moreover, I conclude that there is value in decisional privacy as a conceptual tool for assessing whether hypernudging compromises or strengthens autonomy.

1.1 1.1 Self-Tracking: the New Frontier for Hypernudging

Self-tracking, also referred to as life-logging, quantified self, personal analytics, and personal informatics, is the practice of quantifying behavior through extensive self-surveillance for the purpose of behavioral change. Users can record their behavior through wearing digital devices (such as a clip-on, a wristband, headband, or ring) or monitor their actions with applications on their smartphones. These technologies are often connected to external online platforms where the data is pooled, analyzed, shared, and compared (Lupton 2016: 22–23). Deborah Lupton interprets social media as self-tracking technologies, and vice versa, because of the increasing interconnectedness between different personal devices, applications, and social media. For instance, Strava is a self-tracking technology but a social network for athletes at the same time. Especially tracking one’s health and fitness data has become an increasingly popular practice. Fitness apps and devices that enable the user to improve their athletic performances and overall fitness such as FitBit, Runkeeper, and Strava are well known examples. Other examples include DrinkLess (reduce your alcohol intake), SleepCycle (improve your sleeping patterns), Lose it (aimed at weight loss), SexPositive (to track sexual activity), and What to Expect (pregnancy). Medical apps meant for diagnosing symptoms (23andMe) and apps that track specific medical data, such as glucose levels by diabetics (MySugr), are becoming common and even recommended (Van Dijck and Poell 2016: 2).¹

The main attraction and promise of self-tracking is self-improvement through personalized feedback (Danaher 2016: 17). Personalized feedback is valuable because it is an effective tool for behavioral change. Tailoring and personalization are powerful strategies of persuasion associated with more effective online health behavior change interventions, because users experience tailored feedback as more relevant to their person and situation (Fogg 2003; Krebs et al., 2010; Smit et al. 2015).

If tailoring is persuasive, then what about Big Data-driven personalized choice architectures? Choice architectures are designs in which options are presented to users or consumers (Hausman and Welch 2010: 124). The design can shape the decision-making processes of users significantly by presenting options in a particular way, by

¹ The difference between fitness and medical data is vague. One can make assumptions about a user’s health based on fitness data and vice versa.

offering a certain number of options or by implementing a “default” option. Big Data has enabled “personalized” choice architectures: choice architectures that are designed according to user data feedback. Personalized feedback in self-tracking is based on the analysis of large aggregates of (personal) information or “Big Data,” also referred to as personal analytics. The analysis aims to identify patterns and interesting correlations in the data. Based on the analysis, many devices, and apps make suggestions to their users about how they can change or improve their behavior, what choices to make. For instance, your FitBit can tell you to increase your steps based on the individual user performance it has measured and based on the performances of other users or “peers” (Lupton 2016: 24–26).² Another example would be an energy app that compares your personal energy data to the data of the neighborhood population and encourages the user to make “green” choices.³ Normative interventions are common in self-tracking. Most apps offer feedback with regard to the performance of a user based on the average for his or her age and sex, personal goals, or on a standard set by for instance the World Health Association.

Most self-tracking technologies are still at an early stage of development. Nevertheless, their potential with regard to behavioral change and steering choices is growing along with the rapid progress that is made in real-time data processing, predictive analytics, and Big Data-driven (automated or guided) decision-making processes. The potential of behavioral change through self-tracking lies in highly personalized online choice architectures enabled by smart algorithms that learn from and adapt to the behavior of the user (Michie et al. 2017). MyBehavior, a self-tracking app recently designed by Cornell, is promoted as “the Netflix for you health behavior” and fine-tunes the algorithmic recommendations for personalized feedback for behavior change that “sticks” (Metz 2015; Rabbi et al. 2015).⁴ One can imagine that this has attracted the attention of policy makers who are interested in battling national health issues like obesity, of employers who would like to keep their employees healthy and productive and of companies that see the monetary value in aggregate collections of health data. Because of the current trend in insurance, policy, and employment, in which self-tracking technologies are imposed on or donated to clients, citizens, and employees, it is worth evaluating Big Data-driven behavioral steering that self-tracking technologies may be capable of in the future. For the purpose of this paper, I criticize self-tracking technologies that use Big Data-driven decision-making processes and are hosted by corporations and governmental institutions.

1.2 Features of Nudging

Personalized feedback offered by self-tracking technologies could be interpreted as harmless “nudges,” as ways to scaffold a user’s autonomy by offering “a form of choice architecture that changes the behaviour of people in a predictable way without forbidding any other options or changing their economic incentives” (Thaler and

² Lupton (2016) lists other examples of self-tracking devices that use “coveillance” and pool the data of a particular group in order to monitor behavior. Work Time allows employers and employees to track and encourage each other’s progress. Virgin Pulse tracks the fitness, diet, weight, sleep, and work commitment of employees and compares the aggregated data for employers.

³ This app is used in one of Amsterdam’s living labs: <http://oud.amsterdamsmartcity.com/projects/detail/id/85/slug/city-zen-testliving-lab>.

⁴ For the project website of Rabbi et al. 2015: <http://idl.cornell.edu/projects/mybehavior/>

Sunstein 2008: 6). Ideally, nudges do not compromise your freedom. In fact, according to Thaler and Sunstein’s theory of libertarian paternalism, you can change people’s choice in such a way that they will choose what is best for them and what they would have chosen themselves, had they not been limited by their human flaws such as weakness of will. Importantly, nudges do not reduce or eliminate options, but simply order your choice architecture in a way that favors specific options. The main criticism of nudging is its potential for manipulation (Hausman and Welch 2010: 128; Wilkinson 2013: 347). Manipulation, as I understand it here, refers to the intentional but “hidden” steering of people’s choices by promoting and shaping decision-making processes that persons generally would not use for making rational decisions (Wilkinson 2013: 347; Goodin 1980: 17). For instance, shaping a choice architecture so that a person will only perceive one option and will subsequently choose that option would be manipulative.

Nudges use psychological mechanisms in order to steer decision-making. For instance, bright red arrows pointing towards a staircase will prompt people to take the stairs instead of the elevator. The critique is that “nudges” are not fully in charge with respect to their behavior. Someone else steers their decisions based on psychological mechanisms instead of rational deliberation and argumentation (Nys and Engelen 2016: 4). Moreover, because nudges are “physically” unobtrusive (otherwise they would not work) and its intentions are generally also not transparent, they are potentially manipulative.

Then, in order to ensure “good nudges” Thaler developed three criteria. First, all nudging should be transparent and never misleading. Users should be able to “see” the nudge and to hold the choice architects, the engineers of corporations, or (governmental) institutions who structure the environment in such a way as to encourage a certain type of action, accountable. Secondly, it should be as easy as possible to opt out of the nudge, preferably with as little as one mouse click. Thirdly, there should be good reason to believe that the encouraged behavior will improve the welfare of the nudgee (Thaler 2015). Now, let us assume for a moment that if we would adhere to Thaler’s criteria, we could tolerate nudging. What happens when nudges become powered by Big Data?

1.3 Features of “Hypernudging”

The rise of Big Data practices adds more worries to the nudging debate. Yeung has recently coined and defined “hypernudging” as the algorithmic real-time personalization and reconfiguration of choice architectures based on large aggregates of (personal) data. Yeung stresses that the hyper personalization of a user’s digital choice-environment based on Big Data is incredibly potent and possibly manipulative.

By constantly (re)-configuring and thereby personalizing the user’s informational choice context, typically through algorithmic analysis of data streams from multiple sources claiming to offer predictive insights concerning the habits, preferences, and interests of targeted individuals, these nudges channel users choices in directions preferred by the choice architect through processes that are subtle, unobtrusive, yet extraordinarily powerful (Yeung 2017: 119).

Hypernudges are also known as “Big Data driven decision-guidance processes” or “recommender systems.” Contrary to automated decision-making processes, decision-guidance processes allow the user to make the final decision. A hypernudge “merely”

steers or optimizes someone's decision-making process via algorithms that offer a personalized selection to the "targeted" individual based on the profile constructed from (personal) information. This is also the reason why it is referred to as a type of "nudge."

Hypernudges process past and real-time information from many sources within the networked environment. Hypernudging is therefore based on live data streams as well as a user's personal data history. Moreover, it is not only the data of the individual user that provides feedback, but also all the data of other users. Recommender systems use collaborative filtering, which means that choice selections are optimized based on "people like you" or people who choose and behave like you. These profiles are then often informed by and mixed with individual informational input, in which the individual user can insert information about certain options (by liking or accepting certain options). The choice architect can then provide feedback not only based on the individual's behavior, but also based on and compared to an entire population.

Think about the personalized advertisements a Facebook user receives: these recommendations are constructed real-time based on your own behavior but also on the behavior of people that share similar political views, lifestyles, or music interests. Another well-known example is "GoogleMaps" that updates and suggests one's itinerary real-time by collecting the (GPS) information of other users and traffic information. Self-tracking technology Strava uses a similar technique by combining GPS data and comparing athletic performances among users.

In sum, hypernudges use personalized recommendation to steer behavior. The effectiveness of their interventions is powered by surveillance. The refinement of a target's choice environment requires continuous (corporate) large-scale data collection about people's decisions in order to specify data profiles of targets—which is stored and can of course be used for other applications (Yeung 2017: 122).

1.4 Nudges Versus Hypernudges

Hypernudges are more sophisticated, intrusive, and powerful than Thaler and Sunstein's "nudge." Thaler's criteria for "good nudges" are difficult to meet because of three features that also distinguish hypernudging from regular nudging.

The first feature of hypernudging is *dynamism* or the real time, personalized feedback dynamic. This feature is enabled by the *networked quality* or "surveillance" of hypernudging: the unobtrusive, real time collection, combination and analysis of Big Data, drawn from multiple sources. This feature is powerful because of its one-to-many capacity and personalization. Based on real time data, it can change the choice architectures of millions of users in one mouse click. Moreover, it can offer each and every one of those users a personalized set of options. A regular nudge is aimed at a general public rather than directed at a specific, targeted individual and can only offer a "one size fits all" option.

The second is its *predictive capacity*, which is constituted by smart algorithms that "learn" from the collected data and make behavioral predictions that inform the constant reconfiguration of an individuals' choice architecture. While nudges may be adjusted, this is a time-consuming enterprise. Hypernudges receive immediate feedback about the effectiveness of their interventions.

The final, overarching, and most important feature is the *hiddenness and hidden intentions* of hypernudges. While nudges are also often not immediately detectable, they are and should be “visible” in the physical world (we can “see” the red arrow pointing to the staircase). Hypernudges are hidden in a more complicated and sophisticated way. First, most users are not aware of hypernudges because they are unobtrusively integrated in most of our online informational environments. Furthermore, they are also not aware that the choice architects behind hypernudges are corporations with economic incentives. Google, Facebook, or FitBit may deliberately steer users in a certain direction without their knowledge of the underlying intentions. Of course, this is also a problem in regular nudging, if it does not apply to Thaler’s criteria, but in hypernudging this hiddenness is inherent to the technology.

All three features problematize meeting Thaler’s criteria for “good nudges.” The hiddenness of hypernudges compromises both Thaler’s transparency and welfare criteria. Because hypernudges are unobtrusive, they can be misleading, unjustified, but powerful interferences with decision-making processes. Moreover, because of the corporations behind many hypernudges—after all, data is the currency that makes most online services and technologies commercially viable—we cannot be certain (or have a way to find out) that the intentions and reasons behind hypernudging are legitimate and are guaranteed to improve the welfare of the user (in the future).⁵ As I will emphasize in the third section, the pre-selection of choices offered by the algorithm to the self-tracker may be more aligned with the interest of the actor that controls the technology than with the user.

Furthermore, all three features make it difficult to meet Thaler’s second criterion that it should be easy to opt out of a hypernudge. For one, the level of persuasion increases as choice architectures become more personalized due to real-time surveillance and predictive capacity. Also, opting out is problematized by the hiddenness and unobtrusiveness of these systems, and, many hypernudges cannot be opted out from without quitting the service altogether. For instance, not showing women the same high-paid job advertisements as men entails unjustified interference with someone’s choices and opportunities (Gibbs 2015). Choice architects are responsible for whether people can see their options and can opt out. If they are reckless or negligent, for instance by employing hidden hypernudges, then this could be an unjustified interference with someone’s decision-making process.

The emerging picture is that self-tracking technologies that use hypernudging, interfere with users’ decision-making processes by using real-time, continuous surveillance. While “regular nudging” is often the subject of worry, the features of its Big Brother, hypernudging, make it incredibly difficult—if not impossible—to meet safeguards that should prevent nudges from becoming *unjustified* interferences with decision-making processes. In the next part I will argue that these interferences can

⁵ Often, algorithms are corporate secrets, creating more barriers in understanding why one receive particular feedback. Moreover, to complicate matters further, in some cases, hypernudges use inherently complex machine learning algorithms (Burrell 2016: 3–5). The inner workings of algorithms are “black boxes” and cannot be (easily) explained (Pasquale 2015). Even expert choice architects often do not understand or can explain how “deep learning” algorithms work and have to rely on outsider feedback for mistakes made by faulty machine learning (Byrnes 2016; O’Neil 2016: 154). The “right to explanation” has become an important topic in recent debates about algorithms and machine learning. The renewed European General Data Protection Regulation is claimed to protect this right, but its feasibility is contested (Wachter et al., 2017).

be specified as violations of both decisional and informational privacy. Moreover, because these dimensions are constitutive of our autonomy, hypernudging is worrisome from an autonomy perspective. Interestingly, while self-tracking technologies promise to scaffold one's autonomy, they compromise one's autonomy at the same time (Lanzing 2016).

1.5 Two Complementary Dimensions: Informational and Decisional Privacy

Informational privacy has become the most widely used concept to evaluate the use of data by ICTs. Informational privacy entails the ability to control who has access to one's personal information and to what extent (Westin 1967). Informational privacy is therefore bound up with the concept of reasonable expectations: it is reasonable to expect that the information shared with one's physician will not be shared with a health insurance agency for instance. These expectations about sharing and withholding information are dynamic and context dependent (Nissenbaum 2010). They constitute social norms that mediate and shape our social relationships. Scholars have found informational privacy useful for explaining the harmful aspects of online data collection by third parties that cannot reasonably be expected to have access to that information. For instance, Yeung states that the right most clearly implicated by hypernudges is the right to informational privacy, given the continuous monitoring of individuals and the collection and algorithmic processing of personal digital data that it entails (Yeung 2017: 124).

Although the focus has been on the dimension of informational privacy, a more general typology of privacy should involve many other dimensions (Koops et al. 2017: 2). Scholars have identified dimensions such as privacy of the body, privacy of behavior, privacy of thoughts, local privacy, and decisional privacy. Informational privacy is usually presented as a separate type of privacy that exists alongside these other types of privacy (Roessler 2005; Allen 1988). Yet, recently it has been argued that this may be a misrepresentation. Other dimensions, like decisional privacy, are historically and conceptually related to informational privacy and should be considered complementary concepts (DeCew 2016; Koops et al. 2017).

Decisional privacy is broadly defined as the right against unwanted access such as unwanted interference in our decisions and actions (Allen 1988: 97; Roessler 2005: 9). Roughly, "being interfered with" means that (un)known actors or entities have access to one's behavior and decisions, which allows them to comment upon, interpret, or change one's behavior and steer one's decisions, while this access does not fall under the reasonable expectations of the user or subject or was not granted in the first place.

In the literature, there exist either very narrow or very broad accounts of decisional privacy. On the one hand, decisional privacy is often narrowly associated with "nongovernmental decision-making," intimate choices including (same sex) marriage and childrearing and the right to reproductive liberties. This stems from the USn jurisprudence that grounded reproductive liberties in the right to decisional privacy preceded by *Roe v Wade* [1973] (Allen 1988: 97).⁶ On the other hand, some descriptions are broad, encompassing not only fundamental decisions about one's life projects, such as religion or relationships, but also actions, modes of behavior, and ways of life

⁶ *Roe v Wade* [410 U.S. 113 1973]

or lifestyles (Roessler 2005: 14–15, 79). I will not commit to a particular view, but I will assume, for the purpose of this paper, the broader description.

Privacy is a social negotiation. Whereas informational privacy would regulate access between people to certain information, decisional privacy regulates the access of others in the form of interpretation, objection, commenting, and other forms of intervention in the way you live your life. Of course, the more significant certain behavior, actions, and choices are, the more salient the need for calling them “private” in the sense that they are, quite literally in the case of data mining, none of anyone’s business.

Decisional privacy provides the necessary breathing space to carry out one’s chosen life unhindered in different social contexts, which is important for leading a self-determined life and so, for autonomy (Roessler 2005: 80). In addition, according to Ruth Gavison, it protects you from the interference of others, from the “chilling effect”: conforming your actions to perceived social norms out of fear for (social) sanctions.

Privacy thus prevents interference, pressures to conform, ridicule, punishment, unfavorable decisions, and other forms of hostile reactions. To the extent that privacy does this, it functions to promote liberty of action, removing the unpleasant consequences of certain actions, and thus increasing the liberty to perform them (Gavison 1980: 448).

With regard to ICTs that use hypernudging, while informational privacy can capture the wrong in collecting information, decisional privacy can explain the distinctive wrong involved in using that information to subsequently interfere with a person’s (or group’s) decision-making process. For instance, hypernudges compromise decisional privacy, because users may not know or expect by whom their decisions will be interfered with *based* on their collected information. Decisional privacy concerns should therefore not be reduced to or merely understood in terms of informational privacy concerns. Decisional privacy could be a promising complementary conceptual tool for criticizing hypernudging.

1.6 Privacy and Autonomy

A right to decisional privacy aims to protect freedom from intrusions and interference of the mind and the freedom to exercise autonomous (personal) decision-making. Although decisional privacy does not feature as a concept in the European legal tradition, article 8 of the European Convention on Human Rights does acknowledge the function of privacy as a right to personal development and autonomy as its underlying value. Therefore, the right to choosing one’s own way of living, to self-determination and the right to make one’s own choices about one’s body are mentioned in article 8. Koops et al. argue that decisional privacy is “a distinct type of privacy, which protects the autonomy of persons to make decisions about their body or other aspects of their private life” (Koops et al. 2017: 40). This echoes Cohen’s conceptualization of the right to privacy as protecting decisional autonomy (Cohen 2002: 44).

Decisional privacy thus resonates strongly with the liberal ideal of autonomous decision-making. Moreover, we need decisional privacy in order to ensure decisional autonomy (Cohen 2002). It protects, more or less, the freedom to select our own behavior, actions, and ways of life without interference, as long as we do not harm others—even though others may not agree with our choices because they consider them to be “foolish, perverse, unhealthy, wrong or abnormal” (Mill 2006 [1858]: 19).

Decisional privacy is rooted in and closely related to autonomy. The two are sometimes difficult to distinguish. Nevertheless, I focus on decisional privacy rather than the more general concept of autonomy because it does distinctive work in explaining the particular wrong at stake in hypernudging. As I argue below, its explanatory power lies in the fact that it focuses on the contextual justification of interferences with (and influencing of) decision-procedures. Importantly, it explains why interference with certain decisions is off-limits for particular parties. Expectations about decisional interference are context dependent. For instance, we do not reasonably expect interference with decisions about one's personal health or fitness by commercial enterprises. Decisional privacy protects these expectations. I will return to this point shortly.

All dimensions of privacy protect aspects of autonomy, for instance autonomous decision-making, self-development, or self-presentation. Without privacy, these aspects of autonomy cannot be developed or exercised. Like informational privacy, decisional privacy has a functional relationship with the concept of autonomy. Privacy cannot be reduced to another value such as autonomy, yet the reason why we value privacy is rooted in autonomy (Roessler 2005: 67).

Furthermore, as I announced, there is a difference between autonomy and decisional privacy. While the latter may be violated, this does not entail an immediate loss of autonomy. Roessler argues that decisional privacy protects autonomous authorship with regard to one's own, unique biography: a life free from interpretation and comments, from people whom one does not want to grant this kind of interpretative power (Roessler 2005: 84). Of course, this does not mean that we make our decisions as isolated individuals. "Privacy as decisional autonomy" does not deny that we are embedded, interdependent individuals (Cohen 2002: 47). Our social relations constitute our autonomy. They influence our decision-making processes, control parts of our lives, and determine or provide many of the values, beliefs, and reasons we identify with.

What it means is that the level or kind of decisional interference we accept depends on the social norms per social context. It may be wrong for your boss to comment on the way you raise your children (non essential for the relationship), but not for your partner (essential for the relationship). The kind of privacy that protects us against interference by contexts that we did not reasonably expect (or grant) to interfere with our decisions is decisional privacy. This is exactly what allows a wide variety and plurality of self-chosen ways of life, behavior, actions, and choices, while the reasons that underlie these decisions may very well be rooted in one's community values or discussions with one's partner (Cohen 2002: 48).

To return to the point under discussion, remarks, or advice on one's exercising pattern, partner of choice, friends, religious expression, eating habits or career choices, and (false) inferences or interpretations about one's sexuality, music preferences, or political decisions can be very serious interferences and intrusions even if one's autonomy is not immediately lost. For instance, when an algorithm offers a user feedback on reducing her alcohol intake, including three self-help literature recommendations, the suggestion to visit a physician—including three nearby choices—or several recommendations for insurance policies, users can still plan their lives as they please. However, it is a violation of decisional privacy. The decision-procedure of the user is influenced and interfered with. This is not necessarily problematic of course. What

makes this interference problematic is that this interference is done by a commercial choice architect with economic incentives and that the object of decisional interference is “health,” which is a domain we do not reasonably expect commercially driven interferences with our decision-making process. Moreover, users are not aware of the algorithm and its underlying incentives and intentions.

To summarize, decisional privacy is a precondition for autonomy that enables a person to pursue certain lifestyles and life projects as she pleases without others interfering. Decisional privacy seems to grasp something particular about the moment that a person decides to choose or do something that he or she wants to choose or do and who is reasonably expected to interfere with one’s decision-making process at that moment in that particular social context. As Judith Wagner DeCew argues, many cases concerned with autonomous decision-making about one’s body, intimate relationships, and lifestyle are in fact privacy cases (DeCew 2016: 40).

Stanley Benn argues that in order to respect people as persons, we should regard them as agents: as persons who are capable of making autonomous choices.⁷ To interfere with their decision-making processes through surveillance is to violate both their informational and decisional privacy. Without these privacies, users can no longer be certain whether they are acting based on their own reasons, reasons they selected themselves and identify with, or those of the manipulator (Benn 1971). In that vein, I want to conclude this section by stating that approaching the problem of hypernudging from the perspective of informational and decisional privacy is useful because it tells us something specific about the distinctive wrong of decisional interference enabled by hypernudging. When hypernudges are hidden and when the intentions behind it are commercially driven, they entail a kind of decisional interference that is enabled by and powered by pervasive and pernicious surveillance. This is a violation of both our decisional and information privacy and a potential threat for our autonomy. In the next section, I will use hypernudging as a key example of how these two dimensions are intertwined and therefore should both be used as criteria for evaluating hypernudging.

1.7 Dynamic Dimensions

We have clarified informational and decisional privacy and have established that the two dimensions are distinct, but related to and protective of autonomy. We can now proceed with the suggestion that these dimensions should not be treated as separate types but as part of a mutually reinforcing dynamic. Koops et al. (2017) argue that informational dimension is strongly connected with other kinds of privacies such as decisional privacy and vice versa. A violation of your decisional privacy often implies a violation of your informational privacy (Koops et al. 2017: 56). For example, decisional privacy is not limited to controlling and limiting interference with one’s decisions, but also involves controlling and limiting *information about those decisions* (Koops et al. 2017: 68). In other words, it is not only about determining who can actually interfere with your choices (who has decisional access), but also about determining who knows (who has informational access) about these decisions—and to what extent.

⁷ This argument is similar to Onora O’Neill’s argument about treating people as agents. See: Kerstein 2009.

Likewise, when informational privacy is violated, this may have consequences for your decisional privacy. Informational privacy is not limited to controlling information, but also involves controlling and limiting *interference with one's decisions based on that information*. In other words, it does not only mean that you can determine who or what has access to your information and to what extent, but also about determining who can interfere with your decisions based on that information (who has decisional access based on information).

Similarly, DeCew argues that there are important connections between these dimensions and that they are not very adequate by themselves. DeCew describes the link between surveillance and our ability to make our own decisions so clearly that I hope the reader allows me to quote in full:

(...) because the interests that justify the screen on information include the interest in being free to decide and make choices about family, marriage and lifestyle absent the threat of the same problematic consequences that accompany an information leak. In other words, it is plausible to maintain that worries about what information others have about me are often due to worries about social control by government or others. What one can do to me, or what I can do free of the threat of scrutiny, judgment and pressure to conform, may often depend on what information (personal or not) an individual, state or others have about me. Clearly my behaviour is also affected by the extent to which I can make my own choices. Therefore, both the threat of an information leak and the threat of decreased control over decision making can have a chilling effect on my behaviour (DeCew 2016: 42).

The rehabilitation of decisional privacy and emphasis on the relationship between informational and decisional privacy could be helpful in clarifying some of our ethical intuitions regarding new technologies that use our data to steer our decisions. When we say that hypernudging interferes with our decision-making processes, we can specify this as a violation of decisional and informational privacy.

Hypernudges in self-tracking technologies could serve as an example to illustrate Koops' thesis that informational and decisional privacy are part of a mutually reinforcing dynamic (see Fig. 1). Hypernudges can also be described as a feedback loop in which data about decisions and actions is used for interference with decisions. The resulting data about these decisions is then again fed into the system in order to further tailor and personalize the choice architectures, thus increasing the persuasiveness of interference.

If we want to explore in what way self-tracking devices that use hypernudging could violate decisional and informational privacy simultaneously, we should focus on two sorts of violations that are also deeply intertwined. First, we should look for aspects that violate the user's controlling and limiting abilities with regard to data collection about their decisions. Secondly, we should look for aspects that violate user's abilities to control interference with one's decisions based on data collection. I present several self-tracking examples that illustrate how the two dimensions are compromised in self-tracking as a hypernudging technology. I will conclude this section with a discussion of three potential objections against my argument.




Fig. 1. Hypernudging

1.8 Controlling Access to Information about Decisions

There are several aspects about hypernudges that violate the user’s ability to control and limit data collection about their decisions. Hypernudges are networked and collect data from many different sources. Hypernudges have a recursive nature. They process information in order to shape the user’s decision-making process. Subsequently, information about decisions is used to shape future decisions. Information about decision-making behavior, including the decisions, is therefore of crucial importance to hypernudging. However, users virtually have no ability to control this informational access to their decisions. A good example is the recent scandal involving the WeVibe, a vibrator that tracked the sex life of its users through a corresponding app. The app collected usage information and connected it to user’s e-mail addresses and customer accounts without their knowledge or consent (Domonoske 2017).

Decisions about our sex lives, our political preferences, and health are considered “private,” in the sense that they are not anybody’s business in many social contexts. However, Big Data-driven technologies are slowly blurring the boundaries between contexts, granting access to parties that formerly did not have access to information about your decisions. Commercial enterprises such as FitBit, Facebook, Amazon, and Google now know when users decide to quit smoking, lose weight, vote Democrat, to start a family, or to change careers. More importantly, based on user information, they know what decisions users are *likely* to make in the future. This is problematic for several reasons. The pervasive data collection and surveillance results in a loss of control over the information about one’s decisions, because third parties are able to access this information and use it for ends that the users could not foresee (Brey 2006: 161). The parties that control hypernudges are usually actors with commercial interests. In order to create self-tracking technologies, they should be commercially viable. Users pay for the services with their data and the choice architects use the data for their commercial ends. Examples are health insurance agencies such as Vitality or Aetna that adjust health premiums in return for using a self-tracking technology and achieving certain health goals, but also employers who use self-tracking technologies to increase productivity by monitoring the actions and choices of employees (Boyd 2017 and Dato 2014). Information that was formerly not accessible by, for instance, one’s employer, like one’s geo-location, fitness, or calorie-intake, becomes accessible and

subject to evaluation, interpretation, and, ultimately, interference. Data that is currently collected can be used to make predictions about groups and individuals and steer how they will behave in the future. For instance, information about your lifestyle decisions may become an excuse to interfere with lifestyle choices (“Your geo-location indicates you were moving in Amsterdam last Monday at 02:00AM, while we had an important meeting on Tuesday morning. Can you explain your behaviour?”). Data can be retroactively used for purposes beyond our current imagination. Being aware of this can cause a “chilling” effect on our behavior.

Finally, it is difficult to control or expect what decisional information is used in order to receive unbiased and “accurate” hypernudges. Hypernudges feed past decisional data into choice architectures. This can create a feedback loop that results in a “self-fulfilling prophecy” (O’Neil 2016: 144–146). If you chose thrillers on Netflix in the past, Netflix may recommend you thrillers in the future, which you will subsequently choose, thus reinforcing the feedback loop. This hampers serendipitous encounters, creating Pariser’s notorious filter bubble. But, there is a more serious filtering problem: collaborative filtering. Hypernudges collect data from “people like you”. You may subsequently receive options that are “personalized,” but not only based on your personal data but the data of an entire population. This can lead to unjust and discriminatory choice architectures. In order to make decision-making processes more efficient, ICTs use profiles: stereotypical (social) categorizations of data patterns that can be powered by surveillance. Profiles “produce new forms of vulnerability” (Ball et al., 2009: 352). They hide or remove social contexts and relationships by reducing bodies and behavior to data: data that can be easily controlled and manipulated and is enshrined in a misleading aura of technological objectivity and neutrality (Monahan 2009: 291). For instance, because a user’s geo location indicates that one runs laps in an area where the zip code corresponds with “bad” public health, one may be excluded a health insurance or you may be targeted with alarming health warnings or pricy health products (O’Neil 2016). Hypernudges may perform actions that do not correspond to the needs or intentions of its user, because they made incorrect inferences and/or because the results are unjust (Brey 2006).

In sum, who is in charge of the data, how the data is used (in the future), how the data is interpreted and shared, how long it is saved, and what the social or individual consequences are are all beyond the control of the user (Mittelstadt and Floridi 2016: 319).

1.9 Controlling Interference with Decisions Based on Information

Another cluster of issues revolves around the fact that hypernudges can interfere with a user’s decision-making process based on Big Data collection. Several aspects of hypernudging violate the user’s ability to control interference with one’s decisions based on their information.

First, the main worry about hypernudges is that they meddle with our private lives by interfering with our decision-making processes. Moreover, the worry is that rationales and reasons behind the choices they offer us will remain hidden to us and that we will not even be aware of the fact that we are steered because we simply consume the defaults offered by the “seamless” informational environments we increasingly live in. Having choices, being able to identify with one’s choices, and being able to provide

reasons for these choices are fundamental aspects of being a competent decision-maker. Yet, hypernudges are hidden. Why one is offered a higher health insurance premium, a certain recommended exercise program, or particular career prospects is unclear. Based on the extensive surveillance and the—inherent—selling of personal data, we can be interfered with significant choices in our lives in hidden and unobtrusive ways by (un)known third parties—including companies and governments.

Second, contrary to nudging, hypernudging is robbed of its soothing blanket of libertarianism. One of the ground rules of nudging is that all options remain available to the user (Sunstein and Thaler 2003). The options are re-ordered, but none are taken away. In recommender systems, such as hypernudging, options are taken away based on the user profile. From every 100 posts, a user may only see 10, selected by the smart algorithm. Users would never know what the other options were. Also, a large portion of the selection that users see on Facebook or Google is visible because someone paid to make it visible. Moreover, nudges should always be in the best interest of the *nudgee*. The problem is that the pre-selection of choices offered by the algorithm to the self-tracker may be more aligned with the interest of the actor that controls the technology than with the user. The user may then be steered in a certain commercial or political direction (Owens and Cribb 2017: 12–14). As stated before, a user profile does not only represent the needs of the user, but also those of third parties. The devil is in the default. A hypernudge can “tell us what to choose” because it will require several actions to negotiate and correct the default options or to uncover the options that are not shown to us (Brey 2006). This severely limits the user control for restricting interference with her decisions based on her information and it could be argued that this even constitutes a case of coercion (Raz 1986: 377–378).

Third, these systems are without visible, responsible agents. Users often do not know who interferes with their lifestyle, career choices, or political affiliation based on personal data or why. This is a violation of the user’s ability to control interference with one’s decisions based on their information. Part of having decisional privacy means being able to have reasonable expectations about who can and cannot interfere and to be able to hold actors accountable for (the consequences of) transgression of these boundaries.

Fourth, the hiddenness and unobtrusiveness of hypernudges is particularly risky in light of its scope and structural cumulative effect. Hypernudges can influence many people at the same time, with a tailored menu and create a constant barrage of intrusions and therefore amount to structural interference in many different domains in a user’s life, which is very difficult to control.

1.10 Three Objections

This section discusses three potential objections to my argument. One objection could be that collecting and sharing one’s data are part of the trade-off that offers the benefit of scaffolding one’s autonomy in return.

My response is that I do not argue that it is wrong or impossible to scaffold one’s autonomy by using technology that collects one’s data. I do however argue that a technology can never truly scaffold a user’s autonomy, when it violates informational and decisional privacy. This is the case when the choice architect is a commercial or

other third party that should not reasonably be expected to interfere with the object of one's decision-making process and when it is unclear how and why parts of our decision-making process are interfered with (Van den Berg 2016: 186–188). In practice, however, monetization of data is part of the business and development plan of most viable self-tracking technologies. The widespread use of non-commercial self-tracking technologies is unlikely.

As a second, and related, objection, one could say that using self-tracking technologies implies consent with their operation and how this affects the user.

My response is that the problem is that the meaningfulness of consent is challenged in the age of Big Data. Meaningful, *informed* consent requires awareness and knowledge of the practices of the technology, which is problematized when these practices are hidden. When we consent to what happens to our data, we usually do so based on experience and a combination of contextual legal and social norms. Before, those expectations we consented to were relatively clear and easy to enforce when transgressed. However, emerging networked technologies blur informational boundaries and corresponding norms and expectations. With the widespread and ubiquitous collection of data, individuals often are unaware that their data is gathered and unable to review all the data processes their data is involved in. Moreover, they are unable to assess whether these are lawful, let alone to go to court if they are not (Van der Sloot 2017: 77). Also, users often rely on informational boundaries and norms that pre-date the Big Data era when they use a new technology (Patterson 2013). Furthermore, because the practice of these technologies changes rapidly, the risks change as well. Corporations change their policies, take over formerly idealistic start-ups, and draw up incomprehensible or deceptive terms and agreements (Turow et al. 2007: 747). Data may be used for other (harmful) purposes in the future, which we cannot currently foresee (Van der Sloot 2017: 76).

A third objection to my argument would concern what might appear to be the most obvious resistance strategy, namely, simply not using self-tracking technologies. One might wonder why this is not a viable strategy. Why would we simply not stop using the technology altogether?

There are several reasons why I do not think that this is a viable strategy. First of all, in order to avoid the harms of self-tracking, one should be able to afford not to use the technology or have the resources to buy or use other technologies that do not violate one's privacy. This is unfortunately dependent on one's privileges (Prainsack 2017: 121–122).

Furthermore, the use of self-tracking technologies is becoming increasingly institutionalized or “pushed.” Employers, health insurances, and even NGO's “offer” them to their employees, clients, and target-groups (Lupton 2014: 7).⁸

But most importantly, resisting the harmful aspects of technology by advising users to stop using the technology shifts the responsibility for a political and social problem, with regard to how we (should) handle data and protect the value of privacy, to the individual.

⁸ For Unicef's self-tracking initiative see: <https://unicefkidpower.org>. For a health insurance example that cooperates with FitBit or an Apple Watch see <https://www.investopedia.com/news/fitbit-healthcare-deal-unitedhealth/> or <https://www.vitality.co.uk>. For an example from the workplace see: <https://www.theverge.com/2018/2/1/16958918/amazon-patents-trackable-wristband-warehouse-employees>

2 Conclusion

In this paper, I aimed to both provide an evaluation of hypernudging and an exploration of decisional privacy as a helpful conceptual tool for evaluation hypernudging in self-tracking. In closing, I want to draw two corresponding conclusions. First, I evaluated hypernudging and argued that under certain conditions, hypernudges may violate both informational and decisional privacy. Hypernudging is a key example in which informational and decisional privacy are closely linked and in which they might both be threatened. Big data-driven decision-guiding processes collect and interpret data about our decisions on an unprecedented scale, with unprecedented scope, across multiple contexts, and from multiple sources. This real-time surveillance allows for real-time (re)configuration and further personalization of choice architectures. This makes the technology highly appealing but also a very powerful technology (Fogg 2003). Moreover, hypernudges are hidden. We have to be careful to allow hidden technologies into the fabric of our online environment and our decision-making processes. Hiddenness complicates the fact that corporations produce most self-tracking technologies. Collected lifestyle and health data can be used for steering users into making "profitable" decisions, to act on certain offers, services, or products. This makes users vulnerable to unwanted, profit-driven, interference, and intrusion in health and lifestyle-related decision-making processes.

Secondly, I argued that decisional privacy has value as a conceptual tool for evaluating hypernudges in self-tracking. As we are increasingly adopting more self-tracking technologies that use data in order to steer our behavior, it is helpful for understanding and criticizing hypernudges if we conceptualize hypernudging not only as an informational, but also as a decisional privacy issue. The two concepts are part of a mutually reinforcing dynamic. Intuitions about the manipulative aspects of hypernudging technologies can be specified as violations of informational and decisional privacy. Moreover, now that I have argued that self-tracking technologies that use hypernudging are ethically problematic, it follows that other technologies, like Facebook, are problematic too.

In sum, self-tracking technologies that use hypernudges are potentially powerful means of behavioral change. Although self-tracking technologies are intended to support user autonomy, they might compromise autonomy on a different level. Self-tracking technologies can interfere with influence and steer our decisions with regard to our behavior based on extensive data collection about our decisions. Self-tracking technologies promise to empower users but violate informational and decisional privacy when commercial parties are involved in hidden, extensive surveillance, and interference with decision-making processes that they should not reasonably be expected to be. Since informational and decisional privacy protect autonomy, autonomy is under threat. Self-tracking technologies that violate informational and decisional privacy are therefore problematic from an autonomy perspective.

Acknowledgments For helpful comments on earlier presentations and drafts of this paper I would like to thank the participants at the 2016 Conference of the Dutch Research School of Philosophy (University of Groningen), the 2017 Conference E-Coaching for Health and Well Being (De Bazel, Amsterdam), the Open Minded Session at ELaw (Leiden University) and the workshop participants and commenters at the Privacy Law Scholars Conference 2017 (Berkeley Law, University of California). Special thanks to Beate Roessler, Daniel Susser, Philip Nickel, Anthonie Meijers, Mark Alfano and Sven Nyholm for their constructive suggestions.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Allen, A. L. (1988). *Uneasy access: privacy for women in a free society*. Totowa: Rowman and Littlefield.
- Ball, K., Green, N., Koskela, H., & Phillips, D. (2009). Editorial: surveillance studies needs gender and sexuality. *Surveillance and Society*, 6(4), 352–355.
- Benn, S. I. (1971). Privacy, freedom and respect for persons. In F. Schoeman (Ed.), *Philosophical dimensions of privacy: An anthology*. Cambridge: Cambridge University Press.
- Boyd, A. (2017) Could Your Fitbit Data Be Used to Deny You Health Insurance? February 20th 2017. *The Observer*. Accessed on March 15th 2017 at: <http://observer.com/2017/02/could-your-fitbit-data-be-used-to-deny-you-health-insurance/>.
- Brey, P. (2006). Freedom and privacy in ambient intelligence. *Ethics and Information Technology*, 7, 157–166.
- Byrnes, N. (2016) Why we should expect algorithms to be biased. *MIT Technology Review*. June 24th 2016. Accessed on December 6th 2016 at: <https://www.technologyreview.com/s/601775/why-we-should-expect-algorithms-to-be-biased/>.
- Burrell, J. (2016). How the machine ‘Thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society* 1–12.
- Citron, D. K., & Pasquale, F. (2014). The scored society: due process for automated predictions. *Washington Law Review*, 89, 1–33.
- Cohen, J. L. (2002). *Regulating Intimacy: a new legal paradigm*. Princeton: Princeton University Press.
- Danaher, J. (2016). The threat of algocracy: reality, resistance and accommodation. *Philosophy and Technology*, 29(3), 245–268.
- Dattoo, S. (2014) These companies are tracking the fitness of their employees. March 17th 2014. *The Guardian*. Accessed on December 6th 2016 at: <https://www.theguardian.com/technology/2014/mar/17/why-companies-are-tracking-the-fitness-of-their-employees>.
- DeCew, J. (2016). Connecting informational, fourth amendment and constitutional privacy. In A. D. Moore (Ed.), *Privacy, security and accountability: ethics, law and policy*. London & New York: Rowman and Littlefield International.
- Domonoske, C. (2017) Vibrator maker to pay millions over claims it secretly tracked use. *National Public Radio*. March 14th 2017. Accessed at March 14th 2017 at: <http://www.npr.org/sections/thetwo-way/2017/03/14/520123490/vibrator-maker-to-pay-millions-over-claims-it-secretly-tracked-use>.
- Fogg, B. J. (2003). *Persuasive technology. Using computers to change what we think and do*. San Francisco: Morgan Kaufman Publishers.
- Galic, M., Timan, T., & Koops, B.-J. (2017). Bentham, Deleuze and beyond: an overview of surveillance theories from the panopticon to participation. *Philosophy and Technology*, 30(1), 9–37.
- Gavison, R. (1980). Privacy and the limits of the law. *The Yale Law Journal*, 89(3), 421–471.
- Gibbs, S. (2015) Women less likely to be shown ads for high-paid jobs on Google, study shows. *The Guardian*. July 8th 2015. Accessed December 6th 2016 at: <https://www.theguardian.com/technology/2015/jul/08/women-less-likely-ads-high-paid-jobs-google-study>.
- Goodin, R. (1980). *Manipulatory politics*. New Haven: Yale University Press.
- Hausman, D. M., & Welch, B. (2010). Debate: to nudge or not to nudge. *The Journal of Political Philosophy*, 18(1), 123–136.
- Hildebrandt, M. (2008). Defining profiling: a new type of knowledge? In M. Hildebrandt & S. Gutwirth (Eds.), *Profiling the European citizen: cross-disciplinary perspectives* (pp. 17–145). Dordrecht: Springer.
- Kerstein, S. (2009). Treating others merely as means. *Utilitas*, 21(2).
- Koops, B.-J., et al. (2017). A typology of privacy. *University of Pennsylvania Journal of International Law*, 38(2), 483–575.
- Krebs, P., Prochaska, J. O., & Rossi, J. S. (2010). A meta-analysis of computer-tailored interventions for health behavior change. *Preventive Medicine*, 51(3), 214–221.
- Lanzing, M. (2016). The transparent self. *Ethics and Information Technology*, 18(1), 9–16.
- Lupton, D. (2014). Self-tracking modes: Reflexive self-monitoring and data practices. Available at SSRN: <http://ssrn.com/abstract=2483549> or <https://doi.org/10.2139/ssrn.2483549>.

- Lupton, D. (2016). *The quantified self*. Cambridge: Polity Press.
- Metz, R. (2015) A health-tracking App you might actually stick with. *MIT Technology Review*. July 28th 2015. Accessed at September 5th 2017 at: <https://www.technologyreview.com/s/539721/a-health-tracking-app-you-might-actually-stick-with/>.
- Michie, S. et al. (2017) Developing and evaluating digital interventions to promote behavior change in health and health care: recommendations resulting from an international workshop. *Journal of Medical Internet Research*, 19 (6).
- Mill, J.S. (2006) [1859] On liberty. In: *On Liberty and the Subjection of Women* (Ryan, A. Ed.) London: Penguin Classics.
- Mittelstadt, B., & Floridi, L. (2016). The ethics of big data: current and foreseeable issues in biomedical contexts. *Science and Engineering Ethics*, 22, 303–341.
- Monahan, T. (2009). Dreams of control at a distance: gender, surveillance, and social control. *Cultural Studies Critical Methodologies*, 9(2), 286–305.
- Nissenbaum, H. (2010). *Privacy in context: technology, policy, and the integrity of social life*. Stanford Palo Alto: Stanford University Press.
- Nys, T., & Engelen, B. (2016). Judging nudging: answering the manipulation objection. *Political Studies*, 65(1), 1–16.
- O'Neil, C. (2016). *Weapons of math destruction: how big data increases inequality and threatens democracy*. New York: Crown Publishing Group.
- Owens, J., & Cribb, A. (2017). 'My Fitbit thinks I can do better!' Do health promoting wearable technologies support personal autonomy? *Philosophy and Technology*. <https://doi.org/10.1007/s13347-017-0266-2>.
- Pariser, E. (2011). *The filter bubble: what the internet is hiding from you*. London: Penguin Books.
- Pasquale, F. (2015). *The black box society: the secret algorithms that control money and information*. Cambridge: Harvard University Press.
- 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*. SSRN: <http://ssrn.com/abstract=2242144> or <https://doi.org/10.2139/ssrn.2242144>.
- Prainsack, B. (2017). *Personalized medicine: empowered patients in the 21st century?* New York: New York University Press.
- Rabbi, M., Aung, M.H., Zhang, M. & Choudhury, T. (2015). MyBehavior: automated personalized health feedback from user behavior and preference using smartphones. *The 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing UbiComp 2015*. Accessed at September 5th at: <https://pdfs.semanticscholar.org/4610/744f6410035292e7856c2c949346588bceb9.pdf>.
- Raz, J. (1986). *The morality of freedom*. Oxford: Clarendon Press.
- Roessler, B. (2005). *The value of privacy*. Cambridge: Polity Press.
- Rushe, D. (2014). Facebook sorry – almost – for secret psychological experiment on users. *The Guardian*. October 2nd 2014. Accessed at December 6th 2016 at: <https://www.theguardian.com/technology/2014/oct/02/facebook-sorry-secret-psychological-experiment-users>.
- Schlosser, K. (2016) Uber redesigns app to predict where riders are headed and give them more to do in the car. November 2nd 2016. *GeekWire*. Accessed on May 5th 2017 at: <https://www.geekwire.com/2016/uber-redesigns-app-predict-riders-headed-give-car/>.
- Schwab, K. (2017) Made you click: meet the AI lurking in your inbox. March 8th 2017. *FastCoDesign*. Accessed at July 17th 2017 at: <https://www.fastcodesign.com/3068766/made-you-click-meet-the-ai-lurking-in-your-inbox>.
- Smit, E.S., Linn, A.J., & van Weert, J.C.M. (2015) Taking online computer-tailoring forward: the potential of tailoring the message frame and delivery mode of online health behaviour change interventions. *The European Health Psychologist*, 17 (1).
- Sunstein, C.R., & Thaler, R.H. (2003). Libertarian paternalism is not an oxymoron. *The University of Chicago Law Review*, 70(4), 1159–1202.
- Thaler, R.H. (2015) The power of nudges, for good and bad. *The New York Times*, October 31st 2015. Accessed at December 6th 2016 at: https://www.nytimes.com/2015/11/01/upshot/the-power-of-nudges-for-good-and-bad.html?_r=0.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven & London: Yale University Press.
- Turow, J. (2011). *The daily you: how the new advertising industry is defining your identity and worth*. New Haven: Yale University Press.
- Turow, J., Hoofnagle, C.J., Mulligan, D.K., Good, N. & Jens Grossklags. (2007). The federal trade commission and consumer privacy in the coming decade, 3(3) I/S: J. L. & Pol'y for Info. Soc'y 723, 724.

- Van den Berg, B. (2016). Coping with information underload. In M. Hildebrandt & B. Van den Berg (Eds.), *Information, freedom and property* (pp. 173–198). New York: Routledge.
- Van der Sloot, B. (2017). *Privacy as virtue. moving beyond the individual in the age of Big Data*. School of Human Rights Research Series, Volume 81.
- Van Dijck, J. & Poell, T. (2016). Understanding the promises and premises of online health platforms. *Big data & Society*, January–June:1–11.
- Wachter, S., Mittelstadt, B. & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation (December 28, 2016). *International Data Privacy Law*, 2017. Available at SSRN: <https://ssrn.com/abstract=2903469> or <https://doi.org/10.2139/ssrn.2903469>.
- Westin, A. (1967). *Privacy and freedom*. New York: Atheneum.
- Wilkinson, T. M. (2013). Nudging and manipulation. *Political Studies*, 61(2), 341–355.
- Yeung, K. (2017). ‘Hypermudge’: Big data as a mode of regulation by design. *Information, Communication & Society*, 20(1), 118–136.
- Zittrain, J. (2014). Engineering an election. *Harvard Law Review Forum*, 127, 335.
- Zuboff. (2015). Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30, 75–89.
- Zuiderveen Borgesius, F. & Trilling, D. & Möller, J. & Bodó, B. & de Vreese, C. & Helberger, N. (2016). Should we worry about filter bubbles? *Internet Policy Review*, 5(1).