University of Twente

Manipulation, Aggravation, Autonomy and IoT

Developing an understanding of manipulation by IoT devices

Jan Kaptijn (s2362864)

Master thesis
University of Twente, Faculty of Behavioural, Management, and Social Sciences, Enschede, the Netherlands
MSc Philosophy of Science, Technology and Society - PSTS

Supervisor: dr. A. Henschke
Second reader: dr. J.P. Bergen

2-19-2025
Wordcount: 23994 (Page 4 - 47)

Contents

Acknowledgements

The only thing you know is that you know nothing. This is the first lesson every philosopher learns. It turns out that writing a master's thesis is quite challenging when knowing nothing. Therefore, in the past year, I pretended to know at least some things. I want to thank everyone who supported me during this time; without all of you, I would have never finished.

First, I want to thank Adam for our countless meetings and his invaluable advice on both personal and academic topics. I could not have wished for a better supervisor. Your light-hearted ways of giving advice and amazing sense of humour are things I'll remember for a long time.

Second, I want to thank Mirjam. At the start of this thesis, you were my girlfriend. And now, at the end, you are my wife. Thank you for all the emotional support you gave, especially during the darker days, which were many. Thank you for believing in me when I was on the verge of quitting. Without your relentless support, I would not have been able to finish this thesis. Also, thank you for showing me, even when I did not believe so myself, that I was working hard enough.

Jan-Peter, thanks for your enthusiasm. Your way of doing philosophy is inspiring. Also, if you ever compare me to Jordan Peterson again, I know where to find you...

Niels, thank you for scouting PSTS for me a year in advance. Also, thank you for letting me vent about how much I hated writing this thesis and then proceeding to give me advice I needed to hear but not always wanted to.

I want to thank Mathijs and Camilla for always checking up on me and ensuring I stayed buffed during this thesis.

I want to thank all members of the cult we call HBB for staying interested in me and my thesis, even though I was (and will be) constantly complaining about how technology will be the end of us and everything we hold dear.

Thank you to the boys from Urk for always being there with cold beer, an abundance of cigarettes, and snarky comments whenever I came to visit.

I want to thank my parents (in-law) and the rest of my family for believing in me even though it was sometimes hard to follow what I was doing (trust me, half the time, I do not even know myself).

I want to thank the PSTS cohort of 2022-2024 for all the great lectures, drinks, discussions, and for the abundance of existential crisis. In the last two years, it is you from whom I learned the most.

A special thanks to Marianne, Jedidja, and Matthijs for providing valuable feedback on my drafts.

Summary

This thesis provides an answer to the question of how individual IoT users in a neoliberal society can develop an understanding of manipulation by concrete IoT devices that negatively impact their autonomy. The answer is divided into three chapters.

The first chapter establishes a foundational understanding of digital manipulation through conceptual analysis. By answering the question, “What differentiates digital manipulation by IoT devices from non-digital manipulation?”. It defines manipulation as "an attempt to change an individual's behaviour by undermining their decision-making process by using means of deception or playing on a vulnerability the individual has." By conducting secondary research based on the work of Jongepier and Klenk (Jongepier & Klenk, 2022a), the chapter argues that digital manipulation is distinguished from non-digital manipulation by four aggravating factors that increase its potency: personalization, flow, opacity, and lack of user control. The analysis then extends Jongepier and Klenk’s framework by adding a fifth factor: omnipresence.

The second chapter argues that it is important to generate an understanding of manipulation since it can impact our autonomy in the structural, temporal, and relational dimensions. By answering the question: “How should we understand the relationship between IoT's aggravating factors and their impact on different dimensions of personal autonomy?” Through a definitional analysis of autonomy and an examination of Klenk and Hancock's (2019) work, the chapter establishes a conceptual connection between manipulation and autonomy loss. It then demonstrates that all the previously stated aggravating factors that increase the manipulative potency of digital technologies are present in IoT devices and how, through these factors, IoT devices can diminish specific dimensions of autonomy.

The third chapter synthesizes these insights and introduces the concept of “aggravating webs” as a tool for developing an understanding of the interconnected role of aggravating factors in manipulation by IoT devices. It answers the question: “How can the interplay between aggravating factors in IoT devices be analysed to safeguard personal autonomy against manipulation?” It answers this question. Through examination of aggravating webs by conducting a case study regarding smartwatch users and aggravating factors.

In answer to the previously stated question, this thesis concludes that individual IoT users in a neoliberal society can develop an understanding of manipulation by concrete IoT devices that negatively impact their autonomy by using aggravating webs to analyse the manipulative impact of specific IoT technologies.

Introduction

We used to visit the internet by sitting behind grey boxes we called computers. However, when computers transformed from big grey boxes to chips the size of a postage stamp, visiting the internet became obsolete. Today, we no longer visit the internet, but we can access it everywhere we want (*Global Coverage By floLIVE*, n.d.). Small computers are embedded in everyday objects, allowing for internet access through devices we would not describe as computers in the first place. Computers have become ubiquitous. Phones, watches, lightbulbs, fridges, and shoes now contain tiny computers that allow these objects to connect to the Internet (Apple, n.d.; *The Best Smart LED Light Bulbs for 2024*, n.d.; Samsung, 2024; Harish, 2017), enabling the integration of the Internet into our everyday lives. This integration allows users to navigate cities, monitor vitals, remotely control lights, automatically order milk and track walking patterns. The widespread adoption of the Internet of Things shattered the border between the internet and the physical world. We used to go surfing the web, but now we are caught in a tsunami of internet connectivity.

Throughout this thesis, I focus on the concept of The Internet of Things. The Internet of Things (IoT) refers to everyday objects that are connected to the internet (Gokhale et al., 2018). This internet connection enables these objects to send and receive data over the internet or to remotely activate features of a device (commonly called actuators). IoT-enabled products are often marketed as 'smart', indicating their internet connectivity to the Internet. Examples include smart fridges, smart homes, smart locks, smart water bottles, and smartphones[1] (Google, n.d.-a, n.d.-b; Samsung, 2024; WATERH, n.d.).

IoT devices enable "cyberphysical interactions with users" (J. Zhang & Tao, 2021). IoT devices can alter aspects of the physical world based on digital events and generate digital data based on events in the physical world. The ability for cyberphysical interactions stems from three integral components: sensors, actuators and internet connectivity. Actuators change things in the physical world, like moving a lock, turning on a lamp or adjusting home temperature. For example, the actuator in smart locks (Google, n.d.-b) is the small electrical motor that moves the lock in a door.

Sensors in IoT devices collect data regarding the device itself or its environment, like house temperature, car speed, or the position of the lock on your door. Returning to the smart lock, the sensors in the lock are present in its ability to detect whether someone is trying to break in (Google, n.d.-b).

[1] I want to specifically state that I regard smartphones as IoT-enabled devices as well. First, because they too contain sensors, actuators and are connected to the internet, and second because most smart features of smartphones are dependent on the internet.

Figure 1 Nest x Yale Lock (retrieved from https://store.google.com/us/product/nest_x_yale_lock?hl=en-US)

Without internet connectivity, a smart lock would be an ordinary electrical lock. However, its internet connectivity brings the actuators and sensors together to become an IoT device. The Internet connectivity of the Yale smart lock, for example, allows users to remotely control the actuators and monitor the door's status through an app.

In some cases, IoT devices only contain sensors or actuators instead of both. The sensors in smart cities (Mitton et al., 2012) often lack actuators. They only collect data and send it over the Internet. Conversely, Smart light bulbs (*The Best Smart LED Light Bulbs for 2024*, n.d.) often only contain actuators in the form of LEDs, which can be remotely controlled over the internet. Yet both can be considered IoT devices since they are connected to the internet, allowing them to work in a network in which the sensor of one device can control the actuator of another device.

By themselves, these smart devices might not seem remarkable. However, their internet connectivity enables them to work together in a network. Just like the internet connects computers, the IoT connects things, allowing them to work together. In a smart home, for example, the smart lock can indicate to the lamps that they should turn on because the user is home. However, users are not the only party that is able to access this network. Manufacturers often have their own interest in participating.

Alongside the development of IoT, which started to take off in the 2010s, Big Tech companies made a discovery. They discovered that data collected on users by IoT devices can be used to modify their behaviour. This discovery plays a central role in the system Shoshana Zuboff calls *surveillance capitalism* (Zuboff, 2019). According to Zuboff, surveillance capitalism is an economic system in neo-liberal societies in which capitalist players employ surveillance technology to collect data. This data can then be used to manipulate users in two ways: to generate more capital and to ensure that users are willing to keep interacting with the surveillance devices. Thereby creating a self-reinforcing cycle of manipulation that is fuelled by technology.

IoT devices play a fundamental role in surveillance capitalism, especially when it comes to manipulating users, for two reasons. First, IoT technology enables unprecedented data collection by transforming everyday objects into surveillance tools, thereby creating a ubiquitous network that tracks people's activities everywhere they go. Second, actuators in IoT technology allow for the remote control of the user's environment, thereby allowing for effective means of controlling their user. This manipulative influence of IoT devices deployed by surveillance capitalist companies can have destructive effects on users' autonomy.

If IoT is ubiquitous and surveillance capitalists are using it to manipulate us, what can individual users do to fight back and break free of this manipulative cycle? Rather than focusing on changing companies, this thesis is focused on what individuals themselves can do to safeguard from manipulation by technology companies. This approach empowers users to take control and resist the impulse that might draw them back into the manipulative cycle of surveillance capitalism. However, for users to resist IoT devices' manipulative effects, they must first understand how these devices manipulate them. Developing this understanding is the main purpose of this thesis.

This thesis addresses the vital question: "**How can individual IoT users in a neoliberal society develop an understanding of manipulation by concrete IoT devices that negatively impact their autonomy?**"

Answering this question can provide individual IoT users with the conceptual tools they need to understand the manipulation of the IoT devices they encounter in everyday life. This is the first step in breaking free of the manipulative cycle of surveillance capitalism since understanding how manipulation by IoT devices impacts autonomy can help users analyse potential countermeasures in terms of the positive impact these countermeasures have on autonomy. Additionally, answering this question enables users to maintain their autonomy while not fully rejecting the use of IoT devices altogether. Since they can use their newfound understanding, it provides users with conceptual tools that they can use to their own liking. In this sense, this thesis is not about how users can avoid using IoT technology but about providing an understanding of IoT's manipulative properties that can help users find new ways to engage with these technologies. Some might argue that it would be simpler for individuals to reject IoT technology instead of developing an understanding of manipulative effects. However, this position overlooks several important factors. First, IoT is becoming an integral part of our lives; moreover, without the conceptual tools to think about what individuals can do, they remain subject to the power of others. Second, users will only grasp the severity of this manipulation once they understand how IoT technologies are manipulating them.

This thesis shifts the focus of the debate regarding manipulation and surveillance from examining brought systems like Zuboff's surveillance capitalism to focusing on concrete technologies for specific individuals, allowing users to see the impact of these systems in their personal everyday lives. The concept of surveillance capitalism underlines the importance of our research by giving a detailed account of the manipulative power big technology companies have and by showing how widespread

and invasive the manipulative power of these companies is. However, Zuboff's work does not give us all the tools we need to answer our main research question. I argue that it falls short in two aspects. First, the concept of surveillance capitalism does not focus on individuals in this system. Instead, surveillance capitalism is, as the name implies, focused on an economic system. While this focus serves as a fitting lens to criticise systems of power, it steers us away from focusing on the individuals who live in this system. This thesis fills this gap. Second, Zuboff's work lacks conceptual clarity regarding the concept of manipulation. While Zuboff provides a detailed explanation of this system and many concrete examples of companies trying to steer user behaviour, she fails to provide a concrete definition of manipulation. This lack of conceptual clarity proves problematic for our question since understanding manipulation requires defining it. Developing this definition is the focus of the first sub-question of this thesis.

I answer the main research question by dividing it into three chapters. In the first chapter, I answer the question, “what differentiates digital manipulation by IoT devices from non-digital manipulation”? This chapter provides us with conceptual clarity regarding manipulation, which is what we need to answer our main research question. It also provides the basis for distinguishing manipulation by IoT devices from non-digital types of manipulation. The first chapter provides an answer to its question by first offering a conceptual analysis of the concept of op manipulation by drawing on literature regarding this concept. Second, this chapter introduces a distinction between manipulation and digital manipulation by conducting secondary research based on the work of Jongepier and Klenk (2022a).

The second chapter asks the following question: “How should we understand the relationship between IoT's aggravating factors and their impact on different dimensions of personal autonomy?” Answering this question allows us to see whether the assumption of a connection between manipulation by IoT and loss of autonomy present in the main research question holds. However, before establishing this connection, it is necessary to define the concept of autonomy. Therefore, the second chapter has the following structure: first, we define autonomy by drawing from the literature regarding the concept of autonomy in the context of technology. Second, the importance of autonomy is argued for by discussing its relevance in different aspects that are important in a neo-liberal society. Third, the chapter investigates the connection between autonomy loss and manipulation by examining an argument by Klenk and Hancock (2019) regarding this connection. Fourth, the aggravating factors that differentiate digital manipulation are further investigated by examining each factor individually and analysing how these factors decrease autonomy. This allows us to make the connection between manipulation by IoT devices and autonomy loss explicit.

The final chapter tests the conceptual work of the first two chapters by applying these concepts to a case study regarding a smartwatch; by using insights from this case study, it answers the question, “How can the interplay between aggravating factors in IoT devices be analysed to safeguard personal autonomy against manipulation?”. In answering, it develops a way of understanding the factors and their relationship to each other.

1. What is digital manipulation?

1.1. Introduction

Zuboff's book "The Age of Surveillance Capitalism" (Zuboff, 2019) sparked discussion regarding companies using digital devices that shape our behaviour and how this might be unwanted. However, companies have been trying to manipulate us for a long time by non-digital means. Traditionally, companies have tried to change people's behaviour by informing them about facts regarding the product (Curtis, 2016). However, ever since the introduction of Freud's psychoanalysis in marketing at the end of World War I, the way companies tried to change people's behaviour became more manipulative (Curtis, 2016). Products were no longer marketed as a necessity but as a way to quench an unfulfilled desire. We could see Google as just another marketing company that uses clever ways to manipulate our desires so we will buy more goods and services. However, there seems to be something novel about the digital means of manipulation that they employ versus the traditional means of manipulation.

So, if companies have been trying to shape our behaviour for a long time, what, if anything, is the difference when they try to use digital means for this? More precisely stated, we could ask the following question: "What differentiates digital manipulation by IoT devices from non-digital manipulation?". We answer this question in this chapter.

This is important because if we want to find out how individuals can develop an understanding of digital manipulative practices, we should first see how this type of manipulation is different from non-digital manipulation. Demarcating digital and non-digital manipulation allows us to focus on the relevant features of digital manipulation.

This chapter asks the question, "What differentiates digital manipulation by IoT devices from non-digital manipulation?" and provides an answer in three steps. First, by offering a conceptual analysis of the concept of influence. Second, by conducting a conceptual analysis of the concept of manipulation, which serves as the basis for understanding the difference between regular and digital manipulation. And last, by discussing factors that differentiate digital manipulation from non-digital manipulation by using the works of Jongepier and Klenk (Jongepier & Klenk, 2022a). After taking these three steps, we should have all the puzzle pieces to formulate a satisfying answer to the question: "What differentiates digital manipulation by IoT devices from non-digital manipulation?"

1.2 Influence

Manipulation is a form of influence that tries to change our behaviour. However, not all forms of influence are the same, and certainly, not all forms of influence are manipulation. Let me illustrate some examples of the many different kinds of behaviour changes we encounter on a daily basis to further illustrate the concepts of influence and behaviour change. Advertisers, for example, try to influence us to buy their clients' products, which is evident in the sheer amount of advertisements we encounter on billboards, television, and social media. In essence, the goal of advertisers is to influence us. Normally, we would go through our day unaware of many products or services. It is the advertiser's job to create advertisements that make us aware of these offerings and present them in such a way that we are inclined to buy them. If the advertiser succeeds in this task, they have effectively shaped our behaviour since we would not have purchased the product without the advertisement. In this example, the change in human actions is illustrated in the person buying the product, and the outside factor is the advertisement.

Parents actively try to influence their children's behaviour as part of raising them, steering them in specific directions. For example, a child might not be inclined to say "thank you" after receiving a piece

of candy from an old lady. This is not because the child is unthankful but because the thought simply does not arise in the child. By telling the child to say “thank you”, parents try to influence the child’s behaviour to be more compliant with the social norm of thanking people when they give you something.

The previous two examples are quite visible forms of influence. Influence, however, can also be more mundane. For example, whenever someone asks you to pass the salt during dinner, they shape your behaviour. Before the request, you might not even have been aware of the existence of the salt on the table, but now you are actively interacting with it.

Furthermore, things[2] constantly influence us in various ways. A red traffic light, for example, ‘tries’ to influence us to stop our car, a locked door prevents us from entering a room, and the pavement steers us to walk there instead of on the road. Digital things shape us, too; our phone influences us to pick it up when it rings, and our navigation makes us drive a certain route. Additionally, software can shape our behaviour, too. Instagram draws us to scroll while we are bored, and Microsoft Word urges us to click the word with the red squiggly line to correct the spelling.

There are many differences in the ways we can define influence. One difference, for example, could be the goals behind a specific form of influence. Marketers and parents both try to assert influence, but their goals differ. For marketers, the goal is to make money, while parents try to improve their children's lives by teaching them to conform to social norms. The same is true for Instagram and Microsoft Word. Instagram's goal is to make its users spend as much time as possible on their platform (Bhargava & Velasquez, 2021), while Word's goal[3] is to help users write without spelling mistakes.

Another difference lies in the means by which the influence is asserted. Instagram, for example, influence us to spend more time on their platform by using specific psychological hooks or triggers (Lukyanchikova et al., 2023) to reel us in and keep us invested. In contrast, Word uses visual cues to indicate something is wrong with the spelling of a specific word and make us change it.

We can think of many more differences when comparing these different types of influence. This raises questions about which factor or factors we should focus on when deciding whether a specific form of influence counts as manipulation. In the literature regarding manipulation, this problem is called the demarcation problem (Jongepier & Klenk, 2022a). So, in conclusion, there are many different ways in which outside factors can influence our behaviour, and these differences impact whether we regard such influence as manipulation. Now, let us shift our focus to the concept of manipulation as a specific form of behaviour change.

1.3 Manipulation

If we want to find out what is novel about digital manipulation based on IoT devices, as opposed to non-digital manipulation, we should first find out what manipulation is to begin with. If we do not do this, it becomes a challenge to compare the manipulation present in IoT with the manipulation present in marketing. Without defining manipulation, it might even be the case that we are investigating two different forms of manipulation that are not comparable to begin with. Therefore, this section will elaborate on the concept of manipulation by offering a conceptual analysis based on the work of Fleur

[2] Of course, it questionable whether it are the objects themselves that try to influence or if it are the designers of these objects. We return to this question in a later part of the chapter.

[3] We could also argue that both Word and Instagram have the goal to generate profit, however our focus here is on the impact the have on the behaviour of their users.

Jongepier and Michael Klenk. I chose their work since their chapter in the book "The Philosophy of Online Manipulation" (Jongepier & Klenk, 2022b) offers a comprehensive overview of the concept of manipulation. By elaborating on the concept of manipulation, we provide a basis for a definition of digital manipulation, which is crucial to define if we want to answer our main research question.

A tentative definition of manipulation we can use for now is *a form of influence in which one party (the manipulator) tries to change the behaviour of another party (the manipulatee) in a way that is not coercive or persuasive.* However, literature shows that there are many different answers to how we should interpret this tentative definition (Jongepier & Klenk, 2022a)**.** The differences between definitions of manipulation often lie in one of the following two categories. Firstly, definitions demarcate manipulation differently. That is, they differ in defining where manipulation stops and becomes another type of behaviour change. Second, most definitions agree that there must be some form of intentionality present in manipulation, but there are different ways of defining this intentionality. Let me further elaborate on these three aspects so we can figure out what relevant aspects we should define for our definition of manipulation.

1.3.1 The demarcation problem

When defining manipulation, it is essential to delineate where manipulation stops, and other forms of influence begin. Jongepier and Klenk describe this as the demarcation problem (Jongepier & Klenk, 2022a). Demarcating manipulation from other types of influence is challenging. Most scholars seem to agree that for influence to be manipulative, the manipulated should experience a loss of freedom in some form (Jongepier & Klenk, 2022a). However, it is challenging to reach a consensus on what being “not free” means or to what extent one’s freedom should be limited for influence to count as manipulation. Influence exists on a spectrum, with coercion on one side, persuasion on the other and manipulation somewhere in between. However, the distinction between manipulation and these concepts often falls into a grey area rather than being a clear-cut separation.

First, let us demarcate manipulation from coercion. Coercion involves changing someone’s behaviour by communicating how choosing a specific option will yield negative consequences inflicted by the coercer (S. Anderson, 2023). For example, if I say, “I will stab you with a knife if you do not give me money,” I am coercing you to give me money. Here, I do not take your options away; I only connect a negative consequence, stabbing, to one of the options. A relevant feature of coercion, as argued by Rundiwo (1978), is that it can be used as a reason when justifying your actions. For instance, if your spouse asks, “Why did you give away your money?” referring to the threat of being stabbed would serve as a viable reason. However, not all cases of coercion are this clear-cut. For example, some digital devices will not continue to function if you refuse to install an update issued by the manufacturer (Sunstein, 2015). Here, it is not necessarily clear whether stopping the device’s functionality is a threat or whether this threat bears enough negative consequences to count as coercion. Moreover, most device manufacturers would argue that updating your device is a free service that they offer instead of a burden a user should be coerced to perform.

Second, let us demarcate manipulation from persuasion. Persuasion involves trying to assert influence by offering some form of information regarding the choice at hand (Susser et al., 2019). Persuasion differs from coercion in that it does not connect negative consequences to a specific action that the persuader will impose. However, the distinction between persuasion and manipulation is also not always clear-cut. For example, giving specific information at a specific time can be manipulative in itself. Sunstein (Sunstein, 2015) shows that presenting the same textual information can have different effects based on the circumstances in which this information is presented. This raises the question of

whether optimising your message to assert maximal influence is manipulative. Moreover, the existence of frames shows that presenting information is seldom or never neutral (Halffman, 2019).

Lastly, it seems fitting to briefly demarcate manipulation from nudging as a separate category to demarcate manipulation from. Since nudges are widely discussed in both academia and policymaking (Pedwell, 2017), in their influential book, Thaler and Sunstein (2021) introduce the concept of a nudge. They define a nudge as "any aspect of the choice architecture that alters people's behaviour in a predictable way without forbidding any options or significantly changing their economic incentives" (Thaler & Sunstein, 2021). Like manipulation, nudges are a form of "non-coercive influence" (Jongepier & Klenk, 2022a). Unlike manipulation, nudges only exist in a decision-making context, as evidenced by the use of "choice architecture" in the definition. There is a broad spectrum of what people call a nudge in current discourse (Hummel & Maedche, 2019), which might be one reason why people have questioned the effectiveness of nudges (Porter, 2023). Research shows that the most promising nudges are default options and that other types of nudges might be less promising (Hummel & Maedche, 2019). Again, the line between nudging and manipulation is not clear-cut. Susser, for example, argues that some forms of nudges are manipulative while others are not, making the distinction based on whether the intentions of the nudger are known (Susser et al., 2019). I agree with Susser that some nudges are manipulative while others are not. However, here, too, the demarcation problem prevents us from deciding which nudges are manipulative and which are not.

So, coercion and persuasion are clearly different from manipulation, but how they are different and based on what difference we should demarcate the concept is still unclear. Let us, therefore, regard ways in which we can make this distinction clearer.

1.3.2 Solving the demarcation problem

To differentiate manipulation from other behaviour-shaping concepts, scholars have proposed various solutions that often fall into one of three categories (Jongepier & Klenk, 2022a).

Outcome approach to manipulation: Harms to self-interest

The first category is concerned with outcome-based accounts of manipulation. These accounts determine whether something is manipulative based on the result of the action. On these accounts, for manipulation to occur, the manipulated individual's self-interest must be harmed in some way (Jongepier & Klenk, 2022a). For example, an algorithm is considered manipulative only if it yields negative consequences for the user. In an outcome-based approach, the YouTube algorithm, designed to keep people engaged for extended periods, is not manipulative if it fails to draw me in for longer than I intended. Outcome-based accounts can help distinguish manipulation from persuasion, as the informational nature of persuasion typically does not harm the user's self-interest since providing information in itself is often harmless. Thus, persuasion and manipulation can be differentiated based on the harm to self-interest. However, the outcome-based approach fails to distinguish manipulation from coercion, as both manipulation and coercion can harm the user's self-interest.

A counterexample to outcome-based accounts is found in nudges or parenting. There seem to be benevolent forms of manipulation that do not harm the self-interest of the manipulated individual but instead promote it. For instance, nudging someone to eat healthier could be seen as a form of manipulation that benefits the person.

Process-based approaches to manipulation: Covert Influences and Bypassing Rationality

Second, the process approach describes manipulation based on the methods used to exert influence. As Jongepier and Klenk state, "process views of manipulation interpret manipulation in terms of characteristic processes or modes of influence that lead to a given behaviour or action"(2022a). Daniel

Susser et al. (2019) offer a widely referenced process-oriented account based on covert influence. According to Susser, influence counts as manipulation when the influence is intentionally hidden. This view helps demarcate manipulation from coercion and persuasion, as the latter two are "forthright" forms of influence while manipulation is hidden (Susser et al., 2019).

Additionally, Susser's account can help classify whether a nudge is manipulative. He states that a nudge is only manipulative when it is hidden (2019). For instance, nutrition labels are not manipulative, while placing certain products at eye level is manipulative (Susser et al., 2019). The difference lies in the former's clear intention—to make users reflect on their food consumption—while the latter's purpose is not immediately apparent.

A counterexample to the process approach can be found in the YouTube algorithm. Even if I know that the algorithm is designed to keep me on the platform longer, that knowledge does not prevent me from falling for it. Here, the influence is not hidden since I am aware of the algorithm's purpose. However, I could still argue that the algorithm manipulates me because I end up spending more time than I intended. A possible rebuttal to this counterexample is to focus on the intended covertness rather than the actual covertness. In this case, it is not the actual process of manipulation that defines whether an influence is manipulative but the intended workings of this process.

Another example of a process-based approach is based on bypassing rationality (Jongepier & Klenk, 2022a). In this approach, influence is considered manipulative if it bypasses our rational decision-making capacities. This approach helps differentiate manipulation from coercion and persuasion, as both require the manipulated individual to engage in rational processes. To coerce someone, they must understand the negative consequences of taking a specific action. Similarly, for persuasion, the individual needs to comprehend the persuasive message to change their actions.

So, process-based approaches demarcate manipulation from coercion and persuasion by focusing on the process by which influence is asserted.

Norm-Based Views on Manipulation: Violating values

Lastly, norm-based views consider influence as manipulative if it 1) violates a norm of the manipulated individual or 2) is asserted in a manner that does not conform to social norms (Jongepier & Klenk, 2022a). For example, Gorin (2014) proposes a norm-based view that focuses on the manipulator breaking acceptable norms of influence (Jongepier & Klenk, 2022a). For instance, we might accept an opt-out strategy for a paid subscription since we are used to it. However, if the opt-out process is tedious, we might not find it acceptable. Thus, if the opt-out process is tedious, it would count as manipulation according to this account.

Noggle (1996) offers another norm-based view that is based on making the manipulated individual violate a personal norm (Jongepier & Klenk, 2022a). For example, I might have a personal norm to limit my screen time to one hour. However, the addictive nature of the YouTube algorithm can manipulate me to violate this personal norm. Norm-based views aid the demarcation process by focusing on norms, allowing for ethical considerations relevant to manipulation instead of merely focusing on the exact type of asserted influence.

Norm-based approaches, however, fail to make a clear-cut line between coercion and manipulation. One counterexample to the norm-based view would be threatening someone with a knife. While this is not an acceptable norm for asserting influence, the influence here seems more like coercion than manipulation. Thus, violating a norm alone is not enough for an influence to count as manipulation.

Choosing an approach

To answer our question regarding what individuals can do against manipulative practices, I think it is most fruitful to adopt a process-based approach to the demarcation problem manipulation since it allows us to analyse the specific role IoT devices play in the manipulative process. Focusing on the role of IoT devices in the process of influence allows us to make more general claims about how IoT devices impact manipulation. Furthermore, I think choosing an outcome- or norm-based approach would limit our analysis.

Outcome-based approaches fall short because they focus too much on specific individuals. An outcome-based approach decides whether a specific individual is manipulated based on the consequences of a behaviour-shaping intervention. However, a specific intervention can yield different consequences for different individuals. Therefore, an outcome-based approach will not help our investigation if we want to make claims that work for individuals in general since what counts as manipulation for one individual might count as persuasion for another.

Norm-based approaches for demarcating manipulation suffer from a similar problem. Norms are different from person to person and are therefore unfit for deriving a definition of manipulation that yields a general connection to IoT technology. Moreover, an individual's norms can also be the result of previous manipulation, thereby creating the possibility of an infinite regress when trying to identify whether an IoT device was asserting manipulative influence. Simply put, when an IoT device influences me in a way that is not against my norm, I could still claim that my norm is the result of different influences that were asserted in ways that are against my norm.

Thus, instead of a norm-based or outcome-based approach, we want to use a process-based approach for demarcating manipulation since it allows us to make general claims regarding the role of IoT devices in this process. Now, based on this definition of manipulation, let us consider whether it is even possible for an IoT device to manipulate individuals.

1.3.3 Manipulation, Intention, and Mental states

In the previous sections, we assumed that technologies can manipulate us. If we want to investigate *how* digital IoT devices can manipulate us, we must first define manipulation in a way that can be applied to technologies in order to confirm our assumption. Additionally, addressing this assumption will allow us to better understand the role of technology in the process of manipulation.

In her analysis regarding manipulative machines, Pepp (2022) shows how some scholars, like Barron (2014), define manipulation as requiring the manipulator's intent, which can only arise from a specific mental state or 'mens rea'. Others, like Wood (2014), argue that manipulation can occur through objects and institutions without mental states, focusing instead on the properties of the asserted influence.

To ascribe manipulative properties to digital technology, we must thus deal with the fact that they do not have a mental state. This can be approached in two ways. First, we could argue that machines have intentions and manipulate us without needing a mental state. If we accept that intentionality is needed for manipulation, we must show how digital technology can possess these intentions. While I do not argue that digital technology can form cognitive states resulting in intentions, I take the position that designers can embed their intentions into technology, allowing it to make users act in line with the designer's intentions (Norman, 1999). Technology can be created to afford specific uses, making certain actions more likely (Norman, 1999). For instance, the glass on my desk can be used as a paperweight, but its hollow shape makes it more likely to be used for drinking. However, one might argue that just because a technology has specific affordances, it is not manipulative, as these

affordances are not hidden and do not exploit our decision-making. It would be nonsensical to claim that the designer of the glass on my desk manipulated me to put water in it simply because it was designed with a hollow shape. However, design can be manipulative in some cases. Take dark patterns, for example. Dark patterns are specific design choices in user interfaces that lead users to act against their interests (Luguri & Strahilevitz, 2021). End user license agreements (EULAs) are an example of such a dark pattern. EULAs are often deliberately made long and hard to understand, causing people to click “accept” without reading them. Thereby making people act in line with the designer's intention but potentially against their own interests. Thus, even if technology cannot form intentions by itself, designers can create technology that influences people in ways that are not beneficial to them by deceiving them or playing on their vulnerability, thereby manipulating users through the technology.

Second, we can resolve the issue by arguing for a definition of manipulation where intentionality is not required and mental states are unnecessary. These definitions often focus on the type of influence rather than the manipulator’s intentions (Pepp, 2022). In this view, both people and objects can influence us, and depending on the process by which the influence is asserted, we can determine whether the influence was manipulative. This process often involves circumventing decision-making processes, leading a person to make choices against their own best interests (Pepp, 2022). For example, Wood (2014) argues that systems like capitalism encourage immediate desire satisfaction instead of considering long-term importance. Advertisements often target our immediate desires (or system 1 thinking) (Sunstein, 2015), causing us to focus on the present and ignore the consequences of our purchases (Wood, 2014). Here, the appeal to our immediate desires and subsequential bypassing of more reflective deliberation makes the system manipulative rather than a specific intention.

In conclusion, IoT technology does not need to have cognitive states to manipulate people. This is because, first, designers can embed intentions, and second, we can formulate definitions of manipulation that focus on the way influence is asserted instead of intention. Therefore, the lack of a cognitive state does not prevent digital technologies from being seen as manipulative.

1.3.5 Defining manipulation

In the previous sections, we concluded that a definition of manipulation should have the following two features. First, our definition should focus on the process by which influence is asserted. Second, the influence does not necessarily have to be asserted by an individual. Now, let us consider a definition that incorporates these three features.

A definition that incorporates these three features would be as follows**.** An influence should be regarded as manipulation when it is ***an attempt to change an individual's behaviour by undermining their decision-making process through means of deception or playing on a vulnerability the individual has*.** Let us discuss the different concepts in this definition to further explain this definition.

First, let us focus on the word "attempt." An “attempt to change” means it does not need to succeed. So, a specific influence does not necessarily need to be successful in changing behaviour for it to be manipulative. This aspect arises from demarcating manipulation using a process-based approach instead of an outcome-based approach. By focusing on the process, it becomes more relevant how the influence is asserted than what the result was of said influence. For example, a specific deceptive advertisement might try to convince me to buy a product by providing incomplete information. In our definition, this would count as manipulation since the process by which the influence was asserted is by means of deception. It is not relevant whether this deceptive advertisement succeeds in making me buy said product. Additionally, this attempt does not necessarily need to be ascribed to a person. As we discussed previously, an attempt can also be asserted through technology.

Second, our account focuses on the individual. While it is possible to manipulate groups through propaganda (Ellul, 1965), these groups will ultimately consist of individuals. Also, our focus on the individual is more fitting for answering our research question since the research question focuses on how individuals can develop an understanding of manipulation.

Third, the influence is focused on a person's decision-making process. By a decision-making process, I mean the choices individuals make. Based on these choices, people act in a specific way. These choices can be either conscious or unconscious or, put differently, based on fast or slow thinking (Stanovich & West, 2000), meaning that I do not necessarily have to be aware of the choices I make while making them.

I deliberately refrain from strictly defining decision-making, which means that I will not choose one singular concept of what specifically constitutes a decision. Further analysing the concept “decision” might benefit our understanding of how specific influences impact an individual. However, this is beyond the scope of this thesis. Additionally, not strictly defining decision-making allows us to apply our concept of manipulation to a wider field of theories of behaviour and cognition instead of having to carefully define whether these theories are compatible with our specific definition of manipulation.

The focus on specific means is a direct consequence of choosing a process-based approach. Our definition contains two aspects: deception and playing on a vulnerability. At least one of these aspects must be present in the process for an influence to be counted as manipulative. I arrived at these two aspects by combining the definition of manipulation by both Sher (2011) and Susser (2019).

I have chosen deception based on Susser's account of manipulation. According to Susser, manipulation should be understood as “hidden or covert influence” (Susser et al., 2019). Covert influence can be understood in the following way. When an individual does not know that they are being influenced, it immediately impacts their ability to make decisions for themselves (Susser et al., 2019) since they cannot take into account the fact they are being influenced to make specific decisions.

Our account of manipulation differs from Susser's since an influence does not necessarily need to be covert or hidden; it can also try to circumvent our decision-making process by other means. For Susser, 'forthright' influence would not count as manipulation. For example, he argues that because seduction is not covert, it is, therefore, not manipulative (Susser et al., 2019). He argues that because of the non-covert nature of seduction, we can still take this form of influence into account in our decision-making processes. However, I think that the ability to take into account a specific form of influence is not enough for this influence to be non-manipulative. That is because some forms of seduction can be so powerful that the ability to reflect is impaired.

For example, I might decide that I want to reduce my social media use. When I receive a push notification to open my social media feed, I am aware of the fact that this is done to make me scroll on that app for an extended time, and therefore, I decide to ignore the notification. However, if the specific app deliberately sends the notification during moments of vulnerability, like when I am bored and tired, my diminished willpower makes it much harder to resist clicking it, and I might end up clicking it anyway. I argue that in this example, I was still manipulated since the app engaged with a vulnerability that I am aware of and still caused me to change my behaviour since the seduction was too strong. Therefore, only focusing on covertness is not enough. We should also take into account if the influence uses a specific vulnerability.

I have chosen vulnerability based on Sher's account of manipulative marketing. According to Sher, vulnerability is a specific aspect that exists in a person that can be used to steer their decision-making process (Sher, 2011). It is, therefore, not necessarily something negative or fault of the person. For

this reason, Sher deliberately uses the word vulnerability instead of weakness to show that manipulation is not the fault of the manipulated.

Let me explain what a vulnerability is by using an example. I might like the colour pink and therefore might be more inclined to buy pink products. In itself, preferring pink products is not something negative. However, if someone learns I prefer pink, they might use this information to make me buy more products than I actually want or need because they have been coloured pink. Suddenly, my innocent preference has become something that can be exploited to steer my behaviour. My preference has changed due to vulnerability.

Now that we have defined manipulation, we can restate our main research question by replacing the word manipulation with our new definition. This results in the following question: How can individual IoT users in a neoliberal society develop an understanding of *attempts* by specific IoT devices *to change their behaviour by undermining their decision-making process by using means of deception or playing on a vulnerability,* which will negatively impact their autonomy? Restating our question like this adds much conceptual clarity; however, one question remains. Namely: "What differentiates digital manipulation from non-digital manipulation?"

1.5 Let's get digital: Digital manipulation

Now that we have defined manipulation and argued that technology can, in fact, manipulate us, we are still left with one question: what differentiates manipulation by an IoT device from other forms of manipulation? To answer this question, I will use the concept of aggravating factors as introduced by Fleur Jongepier and Michael Klenk. They introduce this concept in the first chapter of the book "The Philosophy of Online Manipulation" (Jongepier & Klenk, 2022b).[4]

Jongepier and Klenk define aggravating factors as follows:

> "An aggravating factor is a factor that sometimes or typically either (a) makes manipulation more effective, its effects worse or morally wrong, or (b) makes it harder for individuals to avoid or contest manipulative practices and technologies." (Jongepier & Klenk, 2022a).

Both characteristics (a) and (b) are relevant to the question we are trying to answer. (a) since more effective manipulation will increase the risk for an individual to be manipulated, and (b) since it is regarded with individuals' capabilities to mitigate and safeguard themselves against manipulative effects. In their work, Jongepier and Klenk (2022a) provide four aggravating factors. In this section, I will briefly go over them and introduce a fifth factor. For each factor, I explain what it is and how it increases the process of manipulation.

1.5.1 Personalisation

The first aggravating factor Jongepier & Klenk (2022a) describe is *personalisation*. By personalisation, they refer to the fact that digital technologies can adjust their specific manipulative tactics based on data collected about the individual in question. All of this is made possible by the changing and adaptive nature of technologies.

Personalisation might be the most publicly known aggravating factor. The Cambridge Analytica scandal (Greenfield, 2018) brought to light how manipulative tactics were successfully employed to exploit the vulnerability of voters. During a presentation (Concordia, 2016), Alexander Nix, the CEO of Cambridge

[4] Jongepier and Klenk use the term online manipulation to refer to technologies that manipulate us. I use the term digital manipulation since it is more akin to IoT devices. The word online has a connotation of actively visiting a website. By using the word digital instead of online I hope to indicate that one does need to visit a website or open an app to be manipulated by a technology.

Analytica, claimed that they had a psychographic model of every adult in the US and that they changed their influence based on this model.

Personalisation aggravates manipulation by enabling technologies to target specific vulnerabilities in individuals. For example, I might like pink while my spouse hates it. Showing me an advertisement for a pink product will, therefore, be more likely to influence me into buying than showing that same advertisement to my spouse. Not only can digital technologies change their influence based on our preferences and personality, but IoT devices have been demonstrated to be able to adapt themselves based on how we feel (Lee et al., 2007).

We must keep in mind that personalisation is not the only aggravating factor that can be employed to manipulate us; moreover, we can also be manipulated without any personalisation at all, for example, by using *flow*.

1.5.2 Flow

The second factor Jongepier and Klenk describe is Flow (2022a). Flow can describe many different aspects of a technology. Jongepier an Klenk do not provide a strictly defined definition of flow. Instead, they show different ways in which one can interpret the concept of flow.

The first description of flow they provide is the ease of use of a specific technology (Jongepier & Klenk, 2022a). The less effort it takes to use a technology, the more we use it. I call this user flow. User flow can increase the effectiveness of manipulation by making users less reflective when using a technology. This lack of reflection causes users to be less critical and, therefore, more susceptible to manipulation. For example, when wearing a smartwatch, we are mostly unaware that it is there. This makes us less reflective of the fact that we could be collecting data while going through our day.

Second, flow can also be related to the flow of information (Jongepier & Klenk, 2022a). IoT devices allow information to flow from one place to another without the need for user intervention. I call this information flow. Information flow makes it easier for manipulating parties to collect user data. Information flow can make it harder for users to escape manipulation and make manipulation more effective since having information about a person makes it easier to deceive them or to play on their vulnerabilities. Smartwatches are an example of a technology that offers a lot of user and information flow. Smartwatches can, among many other things, keep track of a user's heart rate. Since smartwatches are connected to the internet, they can send this data to other parties without a user's knowledge. One might argue, however, that a user can find out what happens to their data by reading the terms of service. However, this might be tedious or impossible.

Third, flow is often used to describe an experience "during which individuals are fully involved in the present moment" (Nakamura & Csikszentmihalyi, 2014). For lack of a better term, let us call this engaged flow. So, instead of thinking about something else, one fully engages in one's activity. This notion of user flow bears resemblances to the postphenomenological notion of *transparency* (Verbeek, 2015). Transparency, in this context, "refers to the degree to which a device (or an aspect of that device) fades into the background of a user's awareness as it is used" (Verbeek, 2015). When using a smartphone, for example, one is not concerned with looking at a pane of glass but with the content that is presented through the glass. In this case, the glass is both literally and figuratively transparent. Instead of drawing attention to them, IoT devices often want to move to the background so a user can focus on the task at hand. A smart (IoT-enabled) coffee machine, for example, can automatically make coffee at specific times, allowing users to get their coffee without focusing on the process of coffee making, thereby drawing the users' attention away from the machine.

1.5.3 Opacity

Third, let us discuss the aggravating factor of opacity. According to Jongepier and Klenk (2022a), opacity can also be conceptualized as a lack of transparency[5] or the ability to discern whether one is being manipulated. Thus, something is transparent when it is not Opaque and Vice-versa. There are different dimensions in which technology can be transparent or opaque. The next chapter discusses these dimensions. For now, it is sufficient to state that it can be challenging to determine whether digital technology is attempting to manipulate us because of its opaqueness.

Further to this, a lack of transparency can make manipulation more effective by hiding how things work, making it easier to be deceived. This is not unlike a magic trick: once you know how it works, you are less likely to fall for it. The same is true with manipulation in technology. Consider, for example, the YouTube application. The app might send a push notification prompting one to open it. Upon opening, one is immediately greeted with short-form videos. These videos are often addictive(X. Zhang et al., 2019), leading users to spend more time on the app. However, if one knows that the YouTube app will show addictive content first, it might lead that person to engage with the app differently, thereby circumventing manipulation.

Opacity can thus make it easier to manipulate people since they are unaware of specific asserted influence.

1.5.4 Lack of user control

Now, let us consider the aggravating factor of "lack of user control." Lack of user control refers to the inability of users to control the technology they are interacting with. For example, users are often unable to disable specific data collection features of devices without losing specific functionality, or website visitors are unable to control the way in which they receive cookie pop-ups. If we regard IoT devices as a form of influence, a lack of user control amounts to the inability to "influence the influence" (Jongepier & Klenk, 2022a).

This inability to control the way we are influenced increases the risk of manipulation since users lack the ability to disable how they are controlled. For example, if I know I am susceptible to specific kinds of manipulation by a device, I want the ability to shield myself from this kind of manipulation. However, since I lack user control, I am unable to disable this form of manipulation. Additionally, a lack of user control can increase the risk of manipulation indirectly by preventing users from disabling the data collection features of a device. While data collection in itself is not manipulative, it can be used to further personalize manipulative strategies. A smartwatch user, for example, might be unable to disable features that collect their location. Here, knowing the location of a user in itself is not manipulative, but it can be used to increase the effectiveness of personalization tactics, for example, by showing advertisements for restaurants in their city.

1.5.5 Omnipresence

As a fifth and final aggravating factor, I would like to add omnipresence. Jongepier and Klenk (2022a) specifically state that the list of aggravating factors they provide is non-exhaustive. This means that, according to them, there are more factors that differentiate digital manipulation from the four they present. I argue that omnipresence is a helpful addition to further differentiate digital manipulation from non-digital manipulation.

I define omnipresence as the unavoidable interaction with or sensing by digital technology. For individuals in Western neo-liberal societies, it is almost impossible to live their lives without interacting

[5] Not to be confused with the earlier discussed concept op postphenomenological transparency.

with digital technologies. Most people own smartphones, and even those without are still reliant on computers in their everyday lives, as was exemplified by the societal disruption following the recent CrowdStrike (Pilkington & Aratani, 2024) incident where a substantial percentage of Windows computers stopped working shows our reliance on digital technology.

The constant exposure and reliance on digital aggravate its manipulative influence. First, since we either directly or indirectly interact with digital technology for the most part of our waking hours, these technologies can assert their influence whenever they want. Second, our reliance on digital technology prevents us from taking a step back to flee their manipulative influence.

In the second chapter of this thesis, I will elaborate further on the claims made in this paragraph and show how omnipresence is specifically important in analysing IoT's manipulative influence. For now, it is enough to note that omnipresence is one of the factors that aggravates the severity of manipulation by digital technology.

1.5.6 Digital vs non-digital

So, digital manipulation is a form of manipulation with a more potent ability to deceive or play on a vulnerability, resulting from some combination of the aggravating factors that are present in digital technology. What differentiates digital manipulation from non-digital manipulation, then, is the extent to which digital manipulation encompasses these aggravating factors. Digital manipulation encompasses aggravating factors to a greater extent. For example, let us compare the extent to which a billboard, as a case of non-digital manipulation, and a smartwatch, as a case of digital manipulation, can personalize their influence. In both cases, an advertiser might try to sell a pair of running shoes. When using a billboard, the advertiser is limited in the extent to which they can personalize their influence to accommodate deceiving or playing on a vulnerability. The advertiser could, for example, place the billboard in a specific area where they know there might be an audience that is more likely to buy running shoes. Placing the billboard next to a gym might be more effective than placing it next to a restaurant since gym visitors might be more likely to purchase running shoes than people who are dining. Here, the advertiser has some ability to personalise manipulation; however, not nearly as much as he would have with the smartwatch.

Using the smartwatch, the advertiser can collect data regarding people's activity levels and send the advertisement only to those who have a level of activity. Moreover, since the smartwatch is able to communicate with other services as well, the advertiser might learn a favourite colour and combine information about the person to show them an advertisement for a colour of running shoes while running if we compare the extent to which a manipulative influence can be personalised between a smartwatch and a billboard. The smartwatch allows for more personalisation and is, therefore, more effective in playing on vulnerabilities. The same holds for all the other factors.

This increased potency in digital manipulation to deceive or play on vulnerabilities does not only hold for personalisation but also for the other four aggravating factors. This is exemplified in the previous considerations of individual factors.1.6 Conclusion

In this chapter, I provided an answer to the question: what differentiates digital manipulation by IoT devices from non-digital manipulation? The answer I derived is that digital manipulation and non-digital manipulation can be differentiated by at least five aggravating factors: flow, opacity, lack of user control, and omnipresence, which *make digital manipulation more effective than non-digital manipulation.* I arrived at this answer by first defining manipulation as a specific form of influence. More specifically, manipulation is a form of influence that attempts to change an individual's behaviour by undermining their decision-making process by using means of deception or playing on a

vulnerability the individual has. This definition focuses on the process by which the influence is asserted, namely by "means of deception or playing on a vulnerability". It is precisely in this process that we can find the difference between digital and non-digital manipulation. *Digital* manipulation is more effective than 'traditional' modes of manipulation in deceiving people or playing on vulnerabilities because it contains the aggravating factors of personalisation, flow, opacity, lack of user control, and omnipresence that are inherent to digital technology. Thus, in the end, we can conclude that digital manipulation and non-digital manipulation can be differentiated by at least five aggravating factors that make digital manipulation more effective than non-digital manipulation.

2. Digital manipulation and autonomy

2.1 Introduction

In the previous chapter, I argued that digital manipulation by IoT devices can be differentiated from non-digital manipulation because of aggravating factors that increase digital manipulation's potency.

The question I answer in this chapter is: "How should we understand the relationship between IoT's aggravating factors and their impact on different dimensions of personal autonomy?". I argue that digital manipulation augmented by aggravating factors poses a threat to autonomy, thereby illustrating why it is relevant to consider digital manipulation. So that afterwards, in the last chapter, we can explore solutions to this threat to our autonomy.

To arrive at this conclusion, I first define autonomy by considering personal autonomy. So, we can come up with a definition of autonomy that is relevant for digital manipulation by IoT devices. Then, I argue why autonomy is important so that our main claim will remain relevant. I do so by examining the importance of autonomy in moral deliberation and democracy, creating meaning, and achieving our ends. Third, I explain how digital manipulation threatens autonomy by arguing that more effective manipulation also results in more autonomy loss. I do this by examining if there is a conceptual link between manipulation and autonomy. Last, I show how each of the previously defined aggravating factors of personalisation, flow, opacity, user control, and omnipresence all make IoT's manipulative capabilities more effective, thereby increasing manipulation's threat to autonomy and show how each of these factors is related to autonomy loss. I do so by further defining what these factors entail and looking at real-world examples of how IoT impacts manipulation and autonomy.

2.2 What is autonomy

To argue that manipulation threatens autonomy, we must first define what autonomy is. The word autonomy comes from two Greek words, 'autos' and 'nomos'. 'autos' means 'self', and nomos means 'law'. so put together, they mean self-law or self-governing (*Autonomy | Etymology of Autonomy by Etymonline*, n.d.). However, defining autonomy simply as self-law does not capture the implications of this concept. It does not explain what it means to make law for oneself or when someone has autonomy. Also, the concept does not tell us when something or someone has autonomy. In other words, it does not clarify the degree of control one must exert over oneself to qualify as autonomous.

A widely discussed viewpoint in autonomy literature is personal or individual autonomy. Generally speaking, personal autonomy refers to the ability to make one's own life (Buss & Westlund, 2018). For instance, Raz defines it as "the ability to make one's own life in both existential and mundane choices". (Raz, 1988). This definition fits our investigation as it allows us to explore how manipulation by IoT-enabled technology can affect this ability. However, defining autonomy as the ability to make one's own life leaves room for interpretation regarding what constitutes "one's own life" or "mundane and existential" choices.

2.2.1 Own: Basic, Ideal And Structural dimension of autonomy

One such interpretation of personal autonomy as "the ability to make our own life" is the extent to which "own" in "own choices" should be considered in relation to outside influence. The concept of basic autonomy regards actions as one's own when a person can be held responsible for such actions (Christman, 2020). In contrast, ideal autonomy considers actions autonomous when there is no external influence whatsoever (Christman, 2020). A problem with this view, as discussed in the previous chapter, is the difficulty in defining when something is an influence. For example, imagine using an IoT-enabled coffee machine that automatically decides the strength of your coffee based on your calendar; the busier you are, the stronger your coffee; if you want to manually adjust the strength

of your coffee, you can do so using the app. Additionally, the device can send you a push notification that you can click to automatically make coffee for you. Whether making coffee using this device is autonomous would be answered differently in both cases. Adhering to basic autonomy would arguably mean that the action is my own since I can reasonably be held accountable for the strength of my coffee. In contrast, in light of ideal autonomy, we would probably conclude that this action is not my own as there is likely an outside factor influencing me, such as having a busy week or the tediousness of opening an app just adjust your coffee strength.

Furthermore, for a person to make their “own” life, it should be clear to them what they desire and that they can identify with these desires. This ability to identify with our desires is what Pérez-Verdugo calls the structural dimension of autonomy (2023). Again, this degree can be traced back to deliberation on what "own" in one's own life means. Frankfurt relates to the structural dimension of autonomy in terms of first and second-order desires (1988). A first-order desire is immediate, like the desire to get caffeinated. A second-order desire is a desire about a desire (Frankfurt, 1988), like the desire to not consume too much caffeine. My life then is only “my own” or autonomous when I can live according to my second-order desires. For example, I might desire to drink coffee (first-order), but my desire to not consume too much caffeine (second-order) might stop me from acting on my desire to drink coffee. The structural dimension then allows us to ask the following question in light of our case: “Do I desire my desire to drink coffee?” If the answer is no, but I’m still clicking the push notification to make coffee, I am not autonomous.

2.2.3 Life: Temporal dimension of autonomy

The notion of “life” in the personal definition of autonomy (the ability to make one's own life) also allows us to ask questions. Life, or existence, is temporal, meaning that when I reflect on my life, I can consider at least three temporal directions: forward, backward, and present. These directions constitute what Pérez and Barandiaran call the temporal domain of autonomy (2023).

The forward temporal direction considers whether actions are in line with what I want to be in the future (Pérez-Verdugo & Barandiaran, 2023). In this sense, life is about what is to come. In regard to autonomy, we can ask ourselves if our actions align with our future plans (Bratman, 2000). Returning to our example, we can investigate it in a forward direction in the following way: assuming that I want to sleep at 20:00, I could ask whether drinking coffee now is in line with my desire to sleep at 20:00. It could be if drinking coffee allows me to finish work before a deadline and helps me sleep because of a lack of stress. However, if drinking coffee ends up keeping me awake for longer than I expected, it might not be.

We can also look back on our lives, which Pérez and Barandiaran explain as the backward temporal direction (2023). This direction is concerned with asking whether, in retrospect, I would have resisted my desires and if I agree with their source (Pérez-Verdugo & Barandiaran, 2023). In light of our case, the backward temporal direction makes us ask: Do I resent drinking coffee in the past? If I can easily fall asleep that night, I might not resent it. On the contrary, I might regret my extra cup of coffee if I cannot fall asleep.

Both the forward and backward temporal dimensions are concerned with a more global and existential view of autonomy. In contrast, the (what I call) immediate temporal direction allows us to examine actions in the here and now. This is also referred to as local autonomy (Oshana, 2016). In this case, deciding whether one is autonomous depends on whether one is currently being controlled or restricted in the actions one wants to take. Some scholars define this lack of restriction as freedom (Pérez-Verdugo & Barandiaran, 2023). Therefore, we could restate the autonomy condition for immediate temporal autonomy as whether one is free enough to take the actions one wants to take.

Looking at our case, nobody is forcing us to drink coffee when receiving a notification. In that sense, I am autonomous in the immediate direction when pressing the notification to make coffee.

However, the line between being free from influence and being influenced is not always clear-cut. For example, we could argue that the ease of making coffee by only pushing a notification influences me to drink more coffee. Given the omnipresence of the IoT in our world, it is this line between person and environment that the relational dimension is concerned with.

2.2.4 Making: Relational dimension of autonomy

The notion of “Making” in “Making one’s own life” raises even more questions regarding the barrier of “one's own life”. We rarely make choices in a vacuum outside of any influence, and we always stand in relation to other people and our environment. This is what Pérez-Verdugo and Barandiaran call the Relational domain of autonomy (Pérez-Verdugo & Barandiaran, 2023). Feminist critiques have shown us that an agent must always interact with its environment. So you are never free from influence. (Stoljar, 2024). Instead of defining "own" here the absence of outside forces, we should regard it as our “own” position towards these forces. In the relational dimension, autonomy is seen as having relevant or realistic (Oshana, 2016) control over one's actions in relation to the environment in which one finds oneself. In this sense, autonomy can be seen as a power struggle between oneself and the environment, where the environment asserts control over us while we try to assert control over it.

Returning to our example, the notification on our phone might cause us to drink coffee when we first did not think about it. Whether this would be a breach of autonomy is dependent on whether the control I have is both realistic and relevant. On the surface, the ability to choose whether to take my phone out might seem both relevant and realistic. However, this might not be the case, as the coffee machine, which is part of the environment, contains features that will put one on the losing end of the power struggle between the coffee machine and me. Like the invasive nature of push notifications, which can draw our attention whenever they please. Or the vast amount of information that IoT devices collect about their users, which can be used more effectively to play on our vulnerabilities.

2.2.5 Three dimensions of autonomy and their implications

In conclusion, as I stated previously, the aim of this investigation was not to find a new definition of personal autonomy. Instead, I have shown how multifaceted the concept of personal autonomy is. Coming back, I think that the definition of “the ability to make one’s own life” is still helpful in our case. The rich degrees of interpretation in this definition allow for a wide investigation into how manipulative technologies threaten autonomy. When using this definition, we must keep in mind that autonomy is a multidimensional concept. The multidimensional nature of personal autonomy is evident in three key aspects: First, the concept of life inherent in this definition has a temporal dimension. Second, the notion of "own" should be understood as authentic, yet one should acknowledge that one is never entirely free from external influences. Lastly, making implies meaningful control. Meaningful control implies an ongoing negotiation between the individual and their environment. These three elements collectively form what I would like to call a multidimensional understanding of personal autonomy. As I will argue later, all three dimensions of autonomy are, to some degree, negatively impacted by IoT devices. Let us first regard why we should be concerned about the negative impact of IoT devices.

2.3 The importance of autonomy

When arguing that manipulation threatens autonomy, it is necessary to argue why autonomy is important. Otherwise, we risk our argument being valid and sound while remaining irrelevant. Therefore, this section argues for the importance of autonomy in three fields. First, I show how autonomy is essential in two major moral frameworks. Second, I illustrate why autonomy is crucial to

a neo-liberal democracy. Third, I show autonomy's importance by highlighting psychological insights into why people need autonomy to live a good life.

First, there are moral reasons to consider when arguing for the importance of autonomy. We examine these reasons to understand why robbing people of their autonomy is immoral. As discussed earlier, we can (at least) distinguish between two moral frameworks: deontology and consequentialism.

In Kant's deontological moral philosophy, autonomy serves as the basis for moral deliberation. Autonomy here acts as the basis for performing moral actions (Christman, 2020). It is because we can decide for ourselves, based on our own thoughts and rational deliberations, that moral questions on how to act become relevant (Christman, 2020). Therefore, if we adhere to a deontological moral framework and want to act morally, we must conclude that it is important to have autonomy. Of course, Kant's notion of autonomy is not necessarily in line with the personal notion of autonomy that we discussed in the previous section. Still, I think Kants' insights are valuable, even when adhering to a personal notion. If we regard ourselves as having personal autonomy, we must also acknowledge others' personal autonomy. This acknowledgement can guide us in acting morally towards others since they, like us, possess personal autonomy.

In Mill's consequentialist moral philosophy, autonomy itself is an object of value that we should optimise (Christman, 2020). Mill describes a lack of autonomy in the following way: "He who lets the world, or his own portion of it, choose his plan of life for him, has no need of any other faculty than the ape-like one of imitation" (Mill, 1859, Chapter 3). According to Mill, personal autonomy is mandatory for a valuable life. Without the ability to deliberate on why we do things, the value of human life is diminished, and we become no more than machines (Mill, 1859, Chapter 3). In turn, it is then unethical to diminish someone's autonomy because it would diminish the value of their life.

Thus, because autonomy plays a vital role in at least two major moral philosophical frameworks, we can conclude that it is important to cherish it. However, autonomy is important, not only for moral reasons but also for political reasons.

Second, autonomy is essential for liberal democracy since liberal theories operate under the condition that individuals are autonomous (Christman, 2020). Thus, we must conclude that, at least in liberal democracy, autonomy is valuable. However, we must remember the nuances of autonomy, as discussed previously. Not all political theorists hold the same conception of autonomy or consider it important for the same reasons. Let us briefly contrast two perspectives.

Anderson and Honneth believe that safeguarding personal autonomy is one of the most important duties of a liberal democracy (Anderson & Honneth, 2005). The state should, therefore, act in such a way that protects and increases the autonomy of the individuals under it. Further to this, according to Rawls, only the approval of autonomous individuals can make a state's deliberations legitimate; without autonomy, this approval becomes superficial (Rawls, 2005). I must note that more needs to be said about the connection between liberalism and autonomy (Christman, 2020). However, these deliberations are outside the scope of our investigation. Although they might differ, these two examples show that liberal theories still operate under the condition that individuals are autonomous.

Third, autonomy plays a vital role in everyday life. Autonomy is needed for human well-being (Deci & Ryan, 2013; Nyholm, 2022). Drawing from self-determination Theory, Deci & Ryen (2013) show that autonomy is necessary for motivation. A lack of autonomy leads to passivity and decreased motivation, hindering our ability to actively pursue our goals. Also, a lack of agency is linked to decreased well-being (Deci & Ryan, 2013).

Conversely, making intrinsic goals is connected to greater human well-being (Deci & Ryan, 2013). However, reaching these goals requires autonomy. Without the ability to make our own decisions in life, it becomes hard to live up to our goals. Nyholm even argues that "it is plausible that personal autonomy is a key component of a meaningful life." (Nyholm, 2022). This further stresses the importance of autonomy as a necessary component of a meaningful life.

Thus, we have moral, political, and existential reasons to believe autonomy is important. This further stresses why we should be concerned about IoT devices diminishing our autonomy. However, thus far, we have only discussed that IoT devices can manipulate us, not that this manipulation causes a loss of autonomy. Therefore, let us continue by investigating the relationship between manipulation and autonomy to see if manipulation actually leads to autonomy loss.

2.4 How autonomy and Manipulation are related

As noted, our argument currently contains the implicit premise that manipulation leads to autonomy loss. This section aims to make the connection between these two concepts explicit. First, I examine whether a conceptual link exists between manipulation and autonomy loss. Second, I connect our definition of manipulation to the previously established concept of personal autonomy by introducing the concept of conceptual connection.

First, let us examine whether there is a conceptual link between autonomy and manipulation. A conceptual link between two concepts allows us to deductively claim some consequence by proving only the first concept as true. Concretely, a conceptual link between manipulation and autonomy loss lets us *a priori* know that manipulation always results in autonomy loss. In the same way that when we know a bachelor is always unmarried. In this case, one must only prove that a person is a bachelor to show that this person is unmarried.

Klenk and Hancock (2019) pick up on the implicit premise in popular digital manipulation literature (Susser et al., 2019; Zuboff, 2019) that autonomy and manipulation are conceptually linked. This connection is illustrated in Susser's (Susser et al., 2019) account of manipulation as a covert influence. Susser describes the relationship between manipulation and autonomy loss, stating that manipulation "can lead them to act toward ends they haven't chosen, and second, it can lead them to act for reasons not authentically their own." (Susser et al., 2019). This description clearly illustrates how manipulation can result in a loss of autonomy.

However, Klenk and Hancock (2019) argue that one can be manipulated following Susser's definition of autonomy while still acting towards one's own ends and for reasons that are originally one's own. Therefore, there is no conceptual link between manipulation and autonomy loss.

They make this argument by showing that one can deliberately choose to manipulate oneself into achieving ends of one's own choosing, like using an app that can covertly manipulate one to eat healthily. Looking at this counterargument with a multidimensional understanding of autonomy, we can draw a conclusion regarding the temporal dimension. Autonomy loss in the present can potentially lead to increased autonomy in the future, thereby increasing total autonomy over time. Therefore, we can conclude that in our multidimensional understanding of autonomy, there are at least some dimensions where the conceptual link between manipulation and autonomy does not hold.

While agreeing with Klenk that autonomy lacks a direct conceptual link to manipulation, I argue that a conceptual connection still exists between the two. This distinction is important because there may not be an inherent or necessary relationship (a link), but there is still a meaningful association (a connection). For this reason, the lack of a conceptual link is not necessarily problematic for our argument. Let me illustrate what I mean by conceptual connection with an example. Take a

sledgehammer as an analogy for manipulation and destroying a house as an analogy for a decrease in autonomy; a sledgehammer does not necessarily destroy houses. However, it can destroy the bricks that constitute the house. Here, there is no link between a sledgehammer and destroying houses. However, because of its ability to destroy bricks, a sledgehammer likely will be used to destroy houses. This connection between the ability to destroy bricks and destroying houses is what I call a conceptual connection.

Similarly, there is a conceptual connection between our definition of manipulation and autonomy loss. In our definition of manipulation, influence is asserted by means of deception or by playing on a vulnerability. I argue that this process of asserting influence diminishes the necessary conditions to act autonomously in almost, if not all, dimensions of autonomy, thereby resulting in autonomy loss. Just as a sledgehammer possesses specific properties that predispose it to destroying houses. Manipulation, in our definition, holds specific properties that tend to diminish autonomy. Take, for example, the structural dimension of autonomy. Deception can cause our actions to be inauthentic by impacting our ability to critically reflect on whether our desires are truly our own. The relational dimension of autonomy is also likely to be affected by manipulation. By exploiting our vulnerability, manipulation can decrease our ability to make independent choices. This exploitation puts us at the losing end of the power struggle between us and the environment.

Lastly, it is important to note that when someone's autonomy is diminished, this does not mean that that person does not have any autonomy at all, as autonomy is not a binary concept but, as argued, should be understood in terms of dimensions.

So, in conclusion, while there might not be a conceptual link between our conceptualisations of manipulation and autonomy, there is still a conceptual connection. Therefore, the implicit premise present in the works of Susser and Zuboff, that manipulation causes autonomy loss, is still plausible. More importantly, our premise that manipulation causes autonomy loss still stands. The upcoming section will further explore how manipulation through IoT devices causes this loss of autonomy.

2.5 Aggravating factors and autonomy loss

This section returns to the previously introduced aggravating factors and investigates these factors in more detail by first further elaborating on the factors themselves. Second, it provides real-world cases of IoT devices that exemplify the impact of aggravating factors in increasing manipulative potential and diminishing autonomy. Thereby showing how IoT devices in our everyday lives are manipulating us and thereby, in most cases, diminish our autonomy.

2.5.1 Personalisation

The first aggravating factor I discuss is personalisation. This concept focuses on the fact that digital content is personalised for individuals (Jongepier & Klenk, 2022a). Instead of sharing the same content with everyone, each person receives their own customised content. What form of content the person sees is dependent on specific information about the individual. For personalisation to work effectively, data about each person is needed. The data collection capabilities of many IoT devices make them perfect tools for personalisation. Furthermore, internet-connected technology can be adapted on the fly, allowing for more fluid personalisation. The effectiveness of personalisation of adjustments has been questioned, but it turns out it is effective (Zarouali et al., 2022).

Several techniques are used for personalisation, including A/B testing, micro-targeting, and proactive steering. A/B testing involves showing different versions of content to users and analysing which version performs better according to a chosen metric (Tamburrelli & Margara, 2014). Click-through rate is a prevalent example of such a metric (Chae et al., 2018). Micro-targeting works by making a

user's psychological profile and tailoring advertisements to the impact of the advertisement (Borgesius et al., 2018). As mentioned in chapter I, Alexander Nix, former CEO of Cambridge Analytica, claimed in a presentation that messages can be optimised for people based on their psychological profile (Concordia, 2016). Furthermore, through personalisation, people can be proactively steered into behaving in a specific way (Biddle, 2018). By predicting what people will likely do, advertisers can proactively target these people by making this prediction a reality or stopping it.

Let us turn to an example of IoT and personalisation: Google's advertising model. Based on location data from your smartphone, Google tracks past whereabouts. Or as Google's privacy statement puts it. "For example, if you search for where to buy milk nearby on Google, you may see ads for grocery stores in the general area where you frequently browse Google Search while waiting for your bus or train." (*How Google Uses Location Information – Privacy & Terms – Google*, n.d.). This allows Google to show advertisements specifically tailored to individual users based on data derived from their physical location. Google frames this as making advertisements more 'relevant'. However, another way to interpret this is that the manipulation becomes more effective.

As we have seen, data can be used to personalise manipulation strategies and make them more effective. Knowing more about a person increases the effectiveness of manipulation, as it can be crafted specifically for a given individual. IoT devices are particularly effective for personalisation since not only can their sensors collect data about the physical world, as opposed to the digital one, they can even personalize their influence by changing aspects of our physical reality.

2.5.2 Flow

Second, let us have a more in-depth look at how the aggravating factor of flow impacts manipulation and, thereby, autonomy. As discussed, Jongepier and Klenk do not strictly define flow but provide three different ways to conceptualize it. User flow is regarded as a frictionless user experience (Jongepier & Klenk, 2022a). Information flow is regarded as the free flow of data between different parties (Jongepier & Klenk, 2022a). And Existing in a flow state in which users get absorbed in what they are doing (Jongepier & Klenk, 2022a). I will not discuss this last conceptualization of flow as it does not seem to be relevant for IoT devices. Since the purpose of most IoT devices seems to be the opposite, their goal is to make our interactions with technology less obvious. The goal of a smart LED, for example, is to draw our attention away from controlling the light (*The Best Smart LED Light Bulbs for 2024*, n.d.). Now, let us first consider user and information flow in more detail. In doing so, I hope to show that these concepts are not only theoretical but grounded in everyday life.

An example of an IoT device with both high user- and information flow is the Amazon dash button. This is a button that can be placed somewhere in one's house. When one presses the button, it automatically purchases specific products without the need for any further interaction (*Amazon.Com: Amazon Dash Button - Frequently Asked Questions: Amazon Devices & Accessories*, n.d.). User flow is present because the steps for buying a product are significantly decreased when using this button. Instead of opening an app and searching for the product one needs, one only needs to push a button. Information flow is present in the amount of data needed to place an order. With the push of a button, data regarding one's banking information, address, and type of product are sent to Amazon's servers to place an order.

For the Dash button, the flow can increase the effectiveness of manipulation in the following way: by making purchasing a product as easy as pressing a button, users become less reflective about buying, thereby steering them into making more purchases effortlessly. The specific vulnerability played on here is our tendency to do what is easy. Also, the dash button only works for specific products, like specific brands of toilet paper, making us less likely to purchase other products. It is not hard to see

how manipulation here makes us less autonomous. By making us opt for the easy choice, the Dash button makes us less reflective, resulting in less authentic choices.

2.5.3 Opacity and Transparency

Opacity is a lack of transparency (Jongepier & Klenk, 2022a). There are different ways in which technology can be transparent. I call these "dimensions of transparency."[6].

The organisational dimension of transparency concerns how the organisation behind a technology is open about its practices (Jongepier & Klenk, 2022a). In this dimension, an organisation is transparent when it is forthcoming about why it uses technology in a specific way. Take, for example, using an Android smartphone. Google is open to using historical location data of Android smartphones to offer advertisements. Thus, they are transparent in the organisational dimension.

The active outreach dimension of transparency concerns organisations that proactively communicate why their technology works in a specific way (Jongepier & Klenk, 2022a). This is often a strategy to differentiate their technology from competitors. Take the messenger app Signal, for example (*Signal Messenger*, n.d.). Signal is notably transparent about the fact that they do not use trackers or store metadata about chats. This is how Signal uses active outreach transparency to differentiate itself from competitors like WhatsApp.

The factive transparency dimension refers to when users know that a technology is attempting to manipulate them (Jongepier & Klenk, 2022a). Fitness apps provide a fitting example. These apps are designed to make you work out more regularly by employing gamification methods to engage users. (Lister et al., 2014). Users might download these types of apps specifically for their ability to motivate them.

However, while the facts about a technology's manipulative practices might be available, this does not entail that they are easily accessible. This is what I call “factive opaqueness” since the facts about manipulation are there, but the way they are presented is opaque. For example, End User Licence Agreements (EULAs) are often hard to read (Prichard & Hayden, 2008). While they might contain all the relevant information regarding manipulative practices, their complexity and length often discourage users from reading them. In this case, the factual information is present, but it is presented in such a way that users will likely not engage with it.

Jongepier and Klenk (2022a) also introduce the notion of engaged transparency. Engaged transparency is the same as factive. However, in Engaged transparency, the users do not make use of the access to their knowledge about manipulation (Jongepier & Klenk, 2022a); this might be for different reasons. Take YouTube shorts (*Introducing YouTube Shorts*, n.d.), for example. A user might not want to watch short videos but end up watching these addictive (X. Zhang et al., 2019) videos despite knowing they are addictive. Here, the addictive nature causes engagement that results in the user not caring about the fact that he/she is manipulated.

Last, I propose an additional form of opaqueness: technical opaqueness. This refers to the ability to examine the inner workings of a technology. By understanding how it functions, users can determine exactly what it does and whether it is covertly attempting to manipulate them.

Again, Signal offers an excellent example. The inner workings of Signal are fully exposed since the app's source code is openly available (*Signalapp/Signal-Desktop: A Private Messenger for Windows, macOS, and Linux.*, n.d.). By exposing the inner workings, users can confidently use the app without worrying

[6] Instead of offering one example of a concrete IoT device, I will provide an example for each dimension.

about covertly collected data. Of course, reading source code is not a skill most people possess. Thereby, there is still some opaqueness in the form of a technical barrier.

However, a caveat to this dimension is that knowing how something works does not necessarily equate to knowing why it works like that. Signal could, for example, choose a specific colour scheme to manipulate people into using the app longer. In this case, the technical workings do not explain why a specific colour scheme is used.

A prime example of an opaque IoT technology is the Google Home (Google, n.d.-a), as the software for this device is closed source and the EULA is long. Leaving users without the ability to understand how the device uses its many sensors.

2.5.4 Lack of user control

Another prevalent factor that significantly increases IoT's manipulative capabilities is the lack of control users have over these systems. This lack of control is remarkable since IoT systems can influence us greatly, yet we have little influence on the systems that influence us (Jongepier & Klenk, 2022a).

As users, we lack control over the IoT technology we use. In this case, a lack of user control amounts to the inability to use a purchased product in the way we want. Let us consider the case of iRobot's Roomba to exemplify further what a lack of user control means for consumers[7].

A Roomba is an IoT-enabled autonomous vacuum-cleaning robot. The Roomba can not only clean your house, but it can also make a detailed floor plan of your living space (*Guide to Imprint™ Smart Maps*, n.d.). The floor plan can help the Roomba to clean the floor even better. However, these floor plans are also valuable for data-collecting parties, which allows iRobot to make a great profit by selling them (Wen, 2017). Unsurprisingly, not all customers like the fact that their floorplan is sold. Therefore, iRobot offers the option to opt out of the data selling at the cost of not being able to use the floor map feature. However, opting out not only means that the user cannot use the floor map function, but It also disables most of the smart features on the Roomba, Like the ability to schedule a clearing or receive software updates (Zuboff, 2019). This significantly degrades customers' purchases by taking away control over their products.

While the lack of user control in the Roomba example relates to the control over a device with regard to its functionality, it illustrates a wider problem with IoT devices. Most IoT devices contain hardware that is perfectly capable of executing software that is different from the software provided by the manufacturer (Beinsteiner, 2020). The reason why this is problematic is that the closed nature of IoT devices prevents users from creating alternatives that are less manipulative.

Concluding, a lack of user control robs users of the means to adjust or not comply with IoT's manipulative practices. This is true not only for free services like Google Maps but also for paid goods like Roombas. By limiting user control and reducing the ability to defend against data collection practices, manipulation becomes more efficient, as these practices are essential for effective manipulation. Furthermore, lack of user control also directly impacts autonomy by decreasing the choices a user can make.

2.5.5 Omnipresence

As mentioned in the previous chapter, we used to visit the internet. Now, the internet comes to us. We are surrounded by it because things surround us. With the internet finding its way into objects, it can be everywhere. Even mundane items like water bottles (WATERH, n.d.) or shoes (Harish, 2017) can be

[7] Originally this example is from Zuboff in the age of surveillance capitalism.

connected to the internet. In other words, the internet is omnipresent. I believe this omnipresence can significantly increase the effectiveness of manipulation.

This is why I would like to add omnipresence as an extra aggravating factor to the list provided by Klenk and Jongepier(Jongepier & Klenk, 2022a). In the first chapter, I defined the omnipresence of digital technology as the unavoidable interaction with or sensing by digital technology. This definition specifically fits IoT devices. Let us look at three cases to exemplify this.

Ever since the introduction of the iPhone, consumers have been walking around with fully fledged internet-connected computers called smartphones. Smartphones, too, are omnipresent. Most people in the developed world now own a smartphone (*How Many People Own Smartphones?*, 2021). Their multifaceted nature offers many uses. They have increasingly replaced other devices we used to carry, such as MP3 players, watches, and wallets, making them indispensable for navigating everyday life.

Second, ordinary objects embedded with internet connectivity are becoming increasingly omnipresent. Items like doorbells, household appliances, cars, and speakers can now be connected to the internet, further integrating it into our everyday lives. Smart doorbells, for example, are gaining popularity. As these omnipresent devices film public spaces, it becomes increasingly difficult to walk on the street without being detected. Some brands have been shown to actively create facial profiles of the people they record (Moore, 2022). This allows companies to potentially track individuals' movements.

Third, some companies behind IoT devices are becoming increasingly omnipresent. It seems almost impossible to go through a day without interacting with Google, Microsoft, or Facebook in one way or another. People use Google to navigate, search for products, send emails, browse the web, use their Android phones with proprietary Google services, and watch videos on YouTube. Similarly, Microsoft is hard to evade. People use computers running Windows, send emails through Outlook, or write essays using Microsoft Word. Facebook also seems omnipresent wherever we try to connect with others, including WhatsApp and Instagram.

These examples are all, to some degree, visible. However, we also interact with these companies without noticing it. They apply trackers to monitor the websites you visit, which work even if you disable cookies (Bogna, 2021). Public transport systems could rely on Microsoft Azure cloud infrastructure to keep trains running on time, and Google might index your address and Wi-Fi network by using a Wi-Fi map (Leiteritz, 2010). This all happens without any active interaction with users.

Lastly, the emergence of Smart Cities makes it increasingly difficult to move around in urban areas without being sensed by IoT devices. Smart cities collect data from people within them using IoT devices (Gaur et al., 2015). This data can then be used to actively manage various aspects of the city, such as traffic flow or waste management. In a city full of invisible IoT-enabled sensors, it's hard to evade being constantly detected.

To conclude, let us consider how omnipresence makes manipulation more effective and harder to evade. Omnipresence allows for uncannily accurate personalisation. We have become so accustomed to the idea of omnipresent surveillance that people believe the devices around them are listening to their conversations to show advertisements. However, these devices are probably not actively listening, as it would not be technologically feasible to do so (Kröger & Raschke, 2019). Instead, the different data gathered by your smartphone, such as location and acceleration, allow for the creation of uncannily accurate advertisement profiles that increase the efficiency of manipulation.

Furthermore, the omnipresence of IoT devices makes it harder to evade their manipulative influence simply because they seem to be everywhere. Moreover, they are in places that we both have to use and cannot control. This increase in efficiency and difficulty in evading impacts autonomy by making the power struggle between individuals and their environment constant and more intense, thereby decreasing autonomy in the relational dimensions. Everywhere you go, there is a possibility for your autonomy to be breached. Additionally, the omnipresence of companies impacts autonomy by stripping away options to use technology from other companies.

2.5.6 Aggravating factors and autonomy.

All aggravating factors can increase the potency of manipulation by IoT devices since these factors increase the potency of a specific influence to play on a weakness or to make this influence more covert. The question arises as to whether this increased potential for manipulation also increases the potential for losing autonomy. I argue it does.

The example of a sledgehammer indicated the conceptual connection between manipulation and autonomy loss by showing how a sledgehammer can be connected to destroying houses by its ability to destroy bricks. Just as a sledgehammer can destroy houses through its ability to break bricks, manipulation can diminish autonomy through its influence on decision-making. A more destructive hammer does not guarantee more destroyed houses; however, when it is used to destroy a house, its impact will be greater. Simply put, a better sledgehammer is also better at destroying houses. Similarly, more potent manipulation through aggravating factors will also increase the potential harm to autonomy.

We can apply this example to manipulation made more potent by aggravating factors since a more potent form of manipulation through these will likely increase the harm done to autonomy when it is used.

Moving away from a binary conceptualisation of autonomy by regarding it in dimensions allows us to see how autonomy can be impacted by aggravating factors at different levels. Let me illustrate this by examining how each of the aggravating factors can impact different dimensions of autonomy.[8]

Increased personalisation, for example, provides IoT technologies with a benefit in the power struggle between users and their environment, thereby increasing its potency to diminish autonomy in the relational dimensions. Additionally, high personalisation can lead users to question whether an action they are performing is even their own. For example, if Google knows that I like Italian food and I use google maps to find a restaurant, it might show me primarily places of Italian restaurants. Choosing a restaurant based on a selection presented by google might seem autonomous, but given the fact that google showed only specific restaurants, it can leave a user wondering whether the choice they made was their own.

The seamless interaction with IoT devices resulting from user flow can harm autonomy in the forward and backward temporal dimensions by making users less reflective when engaging with the devices. Thereby increasing the risk of engaging in behaviour that they might not endorse when reflecting on it later. For the same reason, increased manipulation can impact structural dimensions. Increased user flow can cause users to make choices that do not align with their goals or second-order desires.

[8] This list is not exhaustive meaning that there are more ways in which factors can impact lead to a negative impact to the user's autonomy.

Opacity has a clear connection with the immediate temporal dimension of autonomy. As a user, I might want to understand how my IoT device works and what data it does with my data. Opacity, however, prevents users from getting this information, thereby limiting what a user can do and negatively impacting the immediate temporal dimension. Additionally, opacity impacts the relational dimension of autonomy by preventing users from understanding their environment, especially in combination with omnipresence this aspect is pressing. Being surrounded by devices you do not understand will put users at the losing end in the power struggle between them and their environment.

The lack of user control present in IoT devices can impact the relational and the direct-temporal dimensions of autonomy. In the power struggle between the user and their environment, users want to assert control over the IoT devices that try to control them. This is exemplified in the language that the hacker community uses when forcing user control over a device; forcing user controlled is known as Jailbreaking (Masjedi, 2024). Hackers perceive their devices a locked in jail and have to break it free from restrictions imposed by the manufacturers of these devices. Often, attempts to jailbreak a device are themselves a result of a lack of immediate temporal autonomy, like the inability to install apps from outside specific app stores (Masjedi, 2024).

Last, the omnipresence of IoT devices impacts the structural and relational dimensions of autonomy. The structural dimension is impacted by the fact that it becomes increasingly harder not to use a specific device once more people start to use it. If more people use smart doorbells, for example, others will expect that you will use one as well. Additionally, omnipresence impacts the relational dimension by constantly exposing users to IoT devices, thereby making the environment more hostile. Once IoT devices become more omnipresent, it becomes increasingly hard not to interact with them, which limits a user's power over their environment.

So, all aggravating factors have the potential to impact different dimensions of autonomy. This is why it is important to mitigate the manipulative effects that result from these factors.

2.6 Conclusion

In this chapter, I asked the question: “How should we understand the relationship between IoT's aggravating factors and their impact on different dimensions of personal autonomy?”. In answering, I provided the argument that digital manipulation, augmented by aggravating factors, poses a threat to autonomy, thereby illustrating why it is relevant to consider digital manipulation. Autonomy in this context should be understood through the personal notion of this concept, namely, the ability to make one’s own life. However, defining autonomy as the ability to make one’s own life still leaves room to interpret the concepts of own, life, and making. Leaving room for degrees of interpretation of these three concepts allows for a multidimensional understanding of autonomy, in which there are at least three dimensions that determine whether an individual is autonomous.

First, the structural dimension, the concept of “own”, allows us to contemplate when life is really one own life. It allows us to ask questions as to how we should relate to outside influence and whether our desires are authentic. Second, the concept of “life” makes us think about the temporal dimension of autonomy since the life of an individual exists in the past, present and future. Third, the concept of making allows us to focus on an individual’s ability to make and how this ability to make can be restricted by our environment. This gives rise to the relational dimension, in which we can conceptualize autonomy as a power struggle between the individual and their environment.

This multidimensional conception of autonomy allows us to go past understanding autonomy as a binary, which aids us in understanding the impact of manipulative influences by IoT devices. Digital-

manipulative influence augmented by aggravating factors in IoT, therefore, does not fully diminish the manipulatee's autonomy but instead negatively impacts specific dimensions to a certain extent. Considering autonomy as a dimensional concept allows us to overcome the lack of a conceptual connection between manipulation and autonomy loss, thereby leaving room for our argument.

It is important to mitigate these negative impacts on specific dimensions since autonomy is relevant for a multitude of reasons. First, it is immoral to diminish someone's autonomy. Second, a neo-liberal democracy cannot function without its subject' demonstrating autonomy. And third, autonomy is necessary to lead a meaningful life.

By providing a more in-depth analysis of aggravating factors and considering concrete examples of manipulative IoT technologies, I showed how an increased threat of manipulation also increases a threat to various dimensions of autonomy. This provides evidence for our claim that manipulation by IoT devices augmented by aggravating factors poses a threat to autonomy. Now, let us turn to how individual users in a neo-liberal society can develop an understanding regarding this increased threat to autonomy.

3. Developing an understanding of manipulation by IoT devices

3.1 Introduction

In this chapter, I consider a solution that can mitigate the manipulative and autonomy-diminishing impact of the aggravating factors.

If we put the goal of this chapter into the concepts we discussed in the previous chapters, we arrive at the following. In the first chapter, we defined manipulation as *an attempt to change an individual's behaviour by undermining their decision-making process by using means of deception or playing on a vulnerability the individual has*. Limiting manipulation as a result of aggravating factors in IoT devices should, therefore, focus on limiting the ability of these devices to deceive us or play on our vulnerabilities. In the second chapter, we defined personal autonomy as "the ability to make our own life" while keeping in mind that this definition can be interpreted across at least three different dimensions: temporality, authenticity, and power.

Thus, joining the findings of the previous two chapters will result in the following goal for this chapter. We want to find a way to limit aggravating factors in deceiving us and play on our vulnerability so that we can ensure our actions align with what we truly want, both now and in the future, and ensure our choices come from within ourselves and reflect our own desires. In the remainder of this chapter, I will call this "safeguarding" our personal autonomy.

This chapter answers the question, "How can the interplay between aggravating factors in IoT devices be analysed to safeguard personal autonomy against manipulation?" and does so by first conducting a case study to see how aggravating factors can impact someone in everyday life. Second, the case study is analysed to discover how aggravating factors can be mitigated.

3.2 Case study

To discover how we can safeguard ourselves from the manipulative influence of IoT devices, let us consider the case of a smartwatch user.

A smartwatch is fitting for such a case since it is a prime example of an IoT technology. First and foremost, a smartwatch is an IoT device because it is a physical thing that is connected to the internet. This device can collect data using sensors and send it over the internet. For example, it can track a user's location, body temperature, voice and heart rate. By using this data, smartwatches can infer information about the user's health, such as the quality of their sleep or whether they have been drinking alcohol (*Apple Introduces Groundbreaking Health Features*, 2024).

However, one aspect that would make smartwatches a less fitting example for IoT devices is their lack of large visible actuators. A smart coffee machine or smart lock, for example, both have visible actuators that change the environment around them by heating water or locking a door. However, I do not think that a lack of visible actuators will pose a problem for using smartwatches as an example for IoT devices since smartwatches are increasingly used to control devices with large actuators, exemplifying their connectivity between different devices (Google, n.d.-b).

3.2.1 Larry, the smartwatch user

For our case, let us imagine a smartwatch user named Larry. Larry has a specific schedule. Every day after work, he eats dinner and drinks one beer. After dinner, Larry takes a 20-minute walk to the

supermarket to get groceries for the next day. Then Larry gets home, watches television for two hours and goes to sleep. One particular evening, Larry decided to buy a smartwatch because his mother had said she liked hers. The day after, he wears it to work, and the smartwatch starts tracking his behaviour. When Larry comes home from work, he starts his usual routine. He gets home, eats, drinks a beer, walks to the shop, watches television and goes to sleep. The day after, he repeats his schedule. At the end of the day, the health rings on his smartwatch indicate that 20 minutes of walking is not enough to meet the daily step goal that the watch has set for him. Also, he receives a notification that his mother has just reached her health goal for the day (Apple, n.d.).

Figure 2: Activity ring on an Apple watch from (Levin, 2020)

The next day, Larry decides that he wants to reach the fitness goal because he wants his mother to know that he is healthy. After eating dinner and drinking a beer, he goes to the shop and takes a different route that is twice as long. Coming home, Larry happily notices that his ring is filled, and a notification is sent to his mother to show that Larry, too, has finished his goal for this day. At the end of the week, the smartwatch generates a report about Larry's health and indicates that Larry's heart rate was not high enough during his extended walks. For this reason, Larry decides that in the upcoming week, he will jog to the shop instead of enjoying his normal walking pace. That same evening, a third party sends Larry an advertisement on his phone about a running shoe sale. Since, the smartwatch indicated to them that Larry's behaviour corresponded to that of potential runners. After seeing the advertisement, Larry decides to buy the shoes.

3.2.2 Is Larry autonomous?

This story about Larry asks the following question: Was Larry autonomous when changing his schedule and buying the shoes, or did the smartwatch manipulate him? First, we must conclude that the smartwatch shaped Larry's behaviour since before he had the smartwatch, the idea of jogging to the shop or buying running shoes had not occurred to him. However, not all forms of influence are manipulation and not all manipulation results in a loss of autonomy. As discussed in the previous chapter, critics like Klenk and Hancock (2019) could argue that we can manipulate ourselves to reach our own goals and thereby increase our autonomy. However, this does not seem to be the case in this example since Larry did not deliberately use the smartwatch to achieve a specific goal. He bought it because his mother said she liked hers. Moreover, if we ask where Larry's drive to buy the sports shoes originates from, we must conclude that it comes from the smartwatch. The smartwatch showed him

rings about his behaviour and mentioned Larry's mother had filled her ring, which drove him to change his behaviour. Thus, following the discussion from Chapter 2, Larry's desire is not authentic because his drive to change originated from the smartwatch and not from a second-order desire. Moreover, the behaviour shaping in this case meets the requirements for our previously defined definition of manipulation. The way in which the watch tries to make Larry more active clearly plays on vulnerability by showing how his mother has reached her goal already. This plays into the human bias to do things they know other people do (Burchell et al., 2013). This can be conceptualized as a negative effect on the relational dimension of autonomy since the smartwatch, as part of Larry's environment, tries to assert dominance over Larry, thereby limiting Larry's ability to make his own choices.

So, in the end, we can conclude that some degrees of Larry's autonomy were diminished because the smartwatch manipulated him. Now, what could Larry do in such a situation to keep his autonomy? As discussed earlier, the aggravating factors make digital manipulation more effective. Therefore, I suggest that we decrease the manipulative effects of the smartwatch by limiting the impact of aggravating factors. However, before we can limit the impact, we must first identify how the aggravating factors are present in Larry's case.

3.2.3 Aggravating factors in the smartwatch

Let us analyse this case to see whether our findings in the previous chapter regarding aggravating factors, manipulation and autonomy hold for this case. The smartwatch, in our case, encompasses all the aggravating factors that can increase manipulation by strengthening its ability to play on a vulnerability or make it more covert, which leads to a decrease in autonomy. Let us briefly discuss how each of the factors is present in our example so we can then start discussing ways to mitigate the manipulative and autonomy-diminishing effects of these factors.

Personalisation increases the ability of the smartwatch to play Larry's vulnerabilities. This is exemplified in multiple aspects of the case. The most obvious is the personalised advertisement Larry receives. This advertisement would have been way less successful in steering Larry's behaviour if he had received the advertisement before buying the smartwatch. It is because the smartwatch shared personal information about Larry that the advertisement became effective. Also, the smartwatch succeeded in pushing Larry to run more by sending him a notification about how his mother already achieved her goal. The push notification would have been way less effective if it had not been personalised and mentioned that a mere stranger had completed their goal.

Flow, too, is present in our example and helps the smartwatch to increase its covert influence. User flow (ease of use) causes Larry to not focus on the smartwatch when it is collecting data. For most parts of the day, Larry is unaware that he is wearing the watch since it just sits there comfortably on his wrist. This user flow enables the watch to collect data that can be used to manipulate Larry without Larry noticing it. Additionally, information flow is also present in this example. Larry is unaware of the different parties that receive his personal data. Just going through his day, the smartwatch is busy collecting, sharing, and analysing data about Larry's health and behaviour.

Opacity increases the covertness at which the smartwatch can assert influence. Larry does not know what is done with all the data from smartwatches and would be unable to find out exactly how the smartwatch handles his data because of the closed-source nature of the device. By reading the terms and services(*Fitbit Terms of Service - Help*, n.d.), for example, he might be able to infer that his data is 'shared' with third parties. However, it does not tell how these parties use his data. The opacity robs Larry of the ability to guard himself against possible misuse of his data. Some might argue that this opacity is not present in all smartwatches. Apple, for example, the creator of the popular Apple Watch, states explicitly on its website that all data is processed locally on the smartwatch and not shared with

other parties (*Privacy*, n.d.). However, this is just one dimension of transparency, and it does not succeed in making the technology fully transparent since the closed-source nature of Apple devices prevents users from actually checking whether these claims are valid. Furthermore, it is unclear what Apple does with this information.

A lack of user control is present in the fact that smartwatches do not allow users to uninstall specific apps or functions. This inability asserts the smartwatch's capability to play on a vulnerability for all users. Even though it technically should be feasible to delete such apps. Apple, for example, does not allow users to turn off activity rings (*How Do I Turn off the Activity Rings? - Apple Community*, n.d.). When using this device, users have no choice but to be constantly reminded of their movement. Similarly, Google smartwatches need the Fitbit app to function. What is problematic about this is that the Fitbit app is known to share data with third parties (**Privacy Not Included Review*, n.d.). This data is a valuable resource for manipulating users. Both these examples directly impact autonomy in that a user is unable to choose how they want to use the device. They increase the manipulative abilities of the smartwatch by preventing the user from turning off the manipulative features or features that can increase the risk of being manipulated.

Lastly, omnipresence is present in multiple aspects of the smartwatch. First, the watch is omnipresent for Larry since he always wears it, even when he sleeps. This allows the watch to constantly collect data and strike with its manipulative influence whenever it wants. Second, smartwatches are becoming more and more ubiquitous. When more people wear a smartwatch, it allows companies to collect more data that can be used to devise more elaborate manipulative strategies.

In conclusion, Larry's choice to change his route to the supermarket and buy running shoes is the result of manipulation by the smartwatch. The reason why this specific behaviour change that resulted from using the smartwatch counts as manipulation is because the smartwatch used Larry's vulnerabilities to steer his behaviour. The manipulation is particularly effective because of the five aggravating factors that increase the effectiveness of the watch's manipulation, thereby demising Larry's autonomy. Now, let us explore how we can diminish the effect of these factors so that Larry might maintain his autonomy in the future.

3.3 Mitigating aggravating factors

3.3.1 One factor at a time

I argue that we should consider how different aggravating factors are related if we want to safeguard our autonomy by mitigating their manipulative effects. For this, I provide two reasons.

First, we should consider the relationship between aggravating factors because we cannot reduce the increase in manipulative effect to a singular factor. Ascribing the increase in manipulative effect to one factor is what I call a reductionist approach since it tries to reduce the problematic aspects of the aggravating factors to a singular factor. For example, a reductionist approach could claim that personalisation is *the* underlying factor that causes Larry's loss of autonomy. Going for a reductionist approach might be tempting since some factors are more visible than others. Personalisation, for example, is a factor that receives much attention. It is widely discussed in many different fields (Cavdar Aksoy et al., 2021), and users have become accustomed to the idea that their technology is adapting itself to them (Kröger & Raschke, 2019).

Personalisation is reliant on data (Jongepier & Klenk, 2022a), meaning that more data about someone makes it easier to personalise manipulative practices towards that person. Given the importance of

data in the privacy debate and for personalisation, it might be tempting to reduce the other factors to just another way to increase data collection. However, the other factors by themselves are more than just that, meaning they do not have to be understood in terms of personalisation. Take the omnipresence in Larry's case, for example. It is possible to describe the omnipresence of his smartwatch as a particularly effective way to collect more data, which will ultimately allow for more accurate personalisation. However, as discussed, the omnipresence does not only increase the amount of data that can be collected. It also increases the risk of manipulation by being constantly present to assert influence since the smartwatch can always send a notification to Larry[9].

I propose that we should see data as an aspect that connects the two factors of personalisation and omnipresence. This connection does not entail that omnipresence should be understood in terms of personalisation. This means that if we want to safeguard our autonomy, we should focus on *both* factors that can separately increase manipulation while keeping in mind that there are aspects that will connect them. This connection between personalisation and omnipresence through data illustrates how different aggravating factors can relate to one another without being reduced to one of both factors. Not only is this true for personalisation and omnipresence, but other factors are interrelated as well.

Second, we should consider the relation between different factors when safeguarding our autonomy because their interconnectedness prevents us from taking countermeasures against a singular factor. Take again the personalisation in Larry's case. A lack of user control prevents Larry from turning off the Activity rings on his watch, leaving Larry no choice but to look at them when he wants to see the time. Opacity, too, prevents Larry from dealing with personalisation. While the rings are a prominent and visible way of personalisation, there might also be more covert ways in which the watch is personalising itself. Opacity prevents Larry from discovering these hidden forms of personalisation, thereby rendering him unable to mitigate the manipulative effects of personalisation.

This example shows how two factors, lack of user control and opacity, prevent Larry from safeguarding his autonomy from another factor, personalisation. Thereby further stressing that we should not focus on a singular factor if we want to safeguard our autonomy.

Thus, in conclusion, if we want to safeguard our autonomy, we should not focus on a singular factor since, first, manipulative effects result from the interplay between different factors, and second, different factors prevent diminishing other factors, thereby making it necessary to understand the relation between factors.

[9] For simplicity's sake this part only discusses the connection between personalisation and omnipresence to illustrate how different aggravating factors can be connected to each other. In reality all five factors can impact each other simultaneously.

3.3.2 Aggravation and relations

Before looking into methods to analyse the connection between different factors, let us briefly reflect on the concept of aggravation itself to see if it can help us with finding such methods.

First, it might be fruitful to investigate how aggravation is used in a legal context since here, the term is widely used to describe how different seemingly unrelated factors together can form a whole that is greater than the sum of its parts (Ashworth, 2005). In legal contexts, aggravation describes how circumstances beyond the core crime itself can increase both its severity and the resulting sentence (Ashworth, 2005). For example, a judge might decide I have to pay a small fine because I stole candy from someone. However, the same judge might decide to increase the fine because I stole the candy from a defenceless child. Here, the fact that I stole from a child instead of from another adult aggravates the severity of my crime and thereby allows the judge to increase my fine. If it then turns out I am wealthy and could easily have paid for candy myself, while the child used the little money they had to buy the candy, the judge might increase my sentence even more. Still, my original crime was stealing, but the severity of the crime is aggravated by the fact that I stole from a child and did not need to steal in the first place. Thus, the severity of my crime increased due to two seemingly unrelated factors.

This legal definition does not exactly match the definition of aggravation that we used in the previous chapters. It falls short in accounting for the interplay, or synergy, between the different factors. The fact that I stole from a child does not directly impact the fact that I could have easily paid for the candy myself. Instead, these are two separate factors that both impact the severity of my crime by themselves. On the contrary, the aggravating factors in the context of digital manipulation are connected to each other. For example, manipulation in Larry's case is not effective only because it involves personalisation and a lack of user control. It is effective because the manipulation and personalisation together prevent a user from turning off personalisation features. So, merely looking at aggravating factors using the legal concept will not aid us in coming up with a fitting way to safeguard our autonomy because it does not see the factors as interrelated.

A different approach that might help is seeing different factors as *aggregating* instead of aggravating.[10] In the context of privacy, Solove (2004a) introduces the concept of aggregation. He argues that individual pieces of information do not reveal much about us. However, taken together, these loose pieces of information can tell more about a person than their individual parts alone. This interaction between different pieces of information is what he calls the "aggregating effect" (Solove, 2004b), meaning that different individual pieces of information can add up to more than the sum of individual parts. Using the aggregating effect as an approach appears to be more helpful in finding a way to safeguard our autonomy since it takes into account the relation between different factors.

Still, the concept of the aggregating effect misses one key feature. While it does indicate that different factors are connected, it does not show *how* they are connected. Take Larry's case again: here, the concept of aggregation will allow us to indicate that, for example, flow and omnipresence together result in more potent manipulation, as different factors impact each other in a way that increases the overall manipulative effect of the smartwatch. However, it does not show us why this is true. Simply stating that flow and omnipresence together increase the manipulative potential of the smartwatch does not tell us why they do so.

[10] The words aggregating and aggravating are linguistically similar, but they have two different meanings.

If we want to find a way to safeguard our autonomy, it is essential to understand how different factors relate. So, let us now focus on finding ways to uncover these relations.

3.3.3 Aggravating webs

One way in which we can analyse the relations between aggravating is by analysing how they impact each other in a sequential way. Factor X impacts factor Y, which in turn impacts Z. I propose we call this method “aggravating chains”. An example of such a chain that corresponds with Larry's case is presented in Figure 4, with the five aggravating factors arranged in a causal chain.

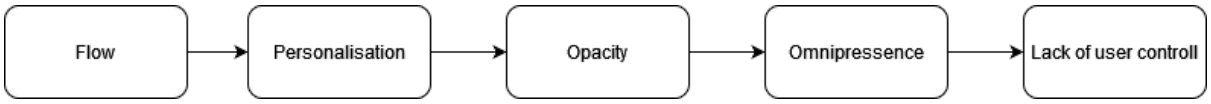

Figure 3: aggravating chain

Let us go through it factor by factor. High flow can lead to more personalisation because high flow makes Larry less aware of the data his watch collects, which will result in more data that can be used to personalise influence more accurately. The increased personalisation can, in turn, lead to more opacity in how the device works since the smartwatch might function differently for Larry than for his mother, making it harder for him to understand it. Increased opacity can make the device more omnipresent since Larry's lack of understanding will probably make him less sceptical about the data that is collected about him, thereby causing him to wear it more. Ultimately, omnipresence can increase the lack of user control, which can result in a lack of user control. Thus, this causal chain results in Larry’s autonomy being significantly diminished, something we might have missed if we had only focused on a single factor.

However, one major flaw of aggravating chains is that they deliberately choose one route in which these factors impact one another instead of acknowledging that the impact factors have on each other is probably not linear. Flow, for example, does not exclusively result in increased personalisation; it can also impact opacity by making Larry less aware of the fact he is using his smartwatch. This would result in a different chain where flow is followed by opacity instead of personalisation. Also, if a chain is as strong as its weakest link, an analysis would fall apart if the connection between two factors is weak or does not exist at all. Additionally, chains assume a form of linearity in the impact factors have on each other, which does not correspond with how these factors work in reality. Flow, for example, can impact personalisation and vice versa. Personalisation can impact flow. Thus, chains will not provide Larry with a sufficient understanding of the manipulation by his smartwatch since they offer too many shortcomings.

Another way for Larry to analyse his smartwatch, which might overcome the shortcomings of aggravating chains and help Larry to develop an understanding of the impact of aggravating factors, while keeping in mind they are connected, is by focusing on how each individual factor impacts every other factor. Let us call this the focused approach since it focuses on each aggravating factor individually. For example, we could analyse how flow impacts personalisation, opacity, omnipresence, and lack of user control. A focused approach overcomes the linearity of aggravating chains by acknowledging that each individual factor has the potential to impact every other factor.

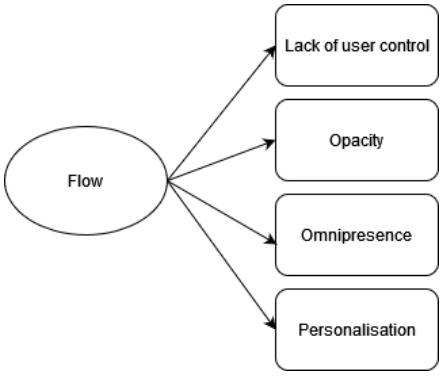

Figure 4: focused approach

Conducting such an analysis might be beneficial for finding tactics to safeguard autonomy. However, I think it is impractical for users to conduct since it would require each factor in relation to the other four, requiring that 20 different interactions need to be analysed. Such an analysis would be overly lengthy since it would lead to 20 different analyses that need to be considered, rendering it less useful for Larry to analyse his smartwatch since the focused approach does not provide a hierarchy in which relations are important to keep in mind. Having 20 equally important analyses will fail to illuminate those connections that generate specifically problematic results in relation to autonomy. This problem will get exponentially worse once for every new factor that is introduced. Furthermore, not all aggravating factors increase each other's impact. Omnipresence, for example, does not necessarily increase the flow of a device, so acting as if such a connection exists will only further obfuscate how aggravating factors really impact one another.

As a solution to the shortcomings of both aggravating chains and the individual approach, I propose we combine relevant features of the two into what I call aggravating webs. By combining chains with the focused approaches, we can create a method that does not follow a singular linear flow in which factors impact each other. Instead, it creates a dynamic approach in which one factor multiple other factors and in which a factor can impact their preceding factors. Also, aggravating webs do not assume that one factor necessarily impacts every other factor. An example of such an aggravating web that corresponds with Larry's example is present in Figure 5. In this figure, we see that personalisation can lead to more omnipresence and user flow, while user flow can also lead to more personalisation and omnipresence.

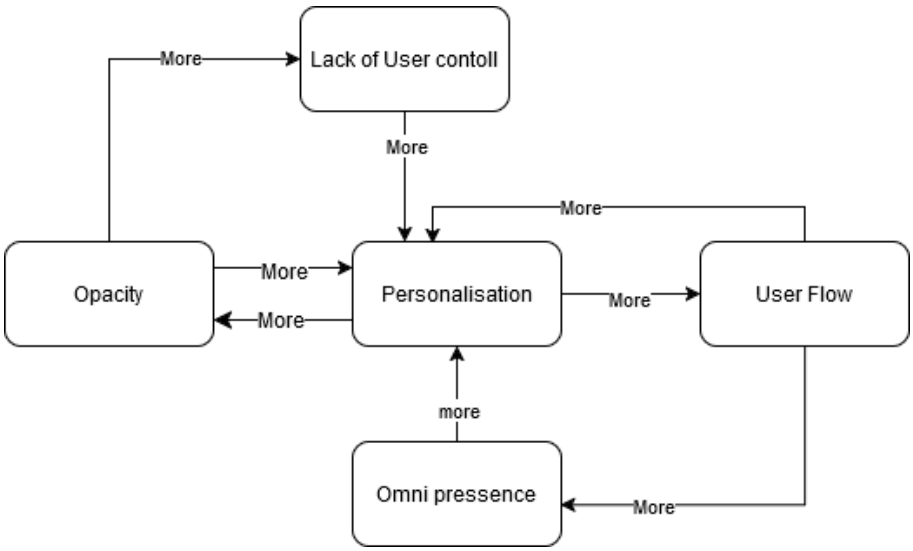

Figure 5: aggravating web

Let us explain the aggravating web by beginning with the aggravating factor of opacity. In the web, we see that the opacity of Larry's smartwatch increases the potency of two other factors: personalisation and user control. As stated, the opacity is present in the smartwatch since Larry is prevented from understanding the workings of his smartwatch because of its closed-source nature. This lack of understanding, in turn, decreases Larry's ability to control his device by preventing him from using the hardware in a way he sees fit. Additionally, opacity increases the risk of personalisation since Larry is unable to see how data collected by the smartwatch is used. The inverse is also true. Personalisation leads to more opacity since when the workings of the smartwatch differ from person to person, it becomes harder to infer how the smartwatch works, thereby increasing opacity.

The web of Larry's case provides us with a valuable understanding of the manipulation by his smartwatch that can result in his autonomy loss: that all other aggravating factors for Larry's smartwatch lead to increased potency for personalisation. This is illustrated by all the arrows that connect back to personalisation. The insights gained from this web can be the first step in analysing the effect of potential countermeasures. For example, the connection of other factors to personalisation helps us to see that it will not be effective for Larry to just focus on limiting the aggravating factor of personalisation by itself since it is amplified by every other factor.

While aggravating webs are a useful way of providing an understanding of aggravating factors, we must acknowledge one shortcoming. Aggravating webs as a concept do not offer the user a specific method by which should constitute a web. Therefore, webs run the risk of being ambiguous since different users might come up with different webs for the same technology. Despite this shortcoming, I still believe that aggravating webs is a valuable way of creating an understanding of manipulation by specific IoT technologies since the ambiguity does not prevent users from creating a web that describes their understanding of the situation. Additionally, it is only when comparing maps with other individuals that the ambiguity becomes apparent. This ambiguity might even spark a fruitful discussion regarding the way different individuals perceive the manipulation by a specific IoT device.

3.4 Conclusion: How to mitigate aggravating factors

In conclusion, by analysing the case of Larry, the smartwatch user, we have discovered that our theory regarding aggravating factors can be applied to real-world scenarios. The manipulative and autonomy-diminishing powers of these factors do not originate from a singular factor. It is through the interplay between different aggravating factors that IoT devices gain their manipulative and autonomy-diminishing powers. Therefore, it is not adequate to safeguard autonomy by simply limiting the effects of a singular factor since limiting the impact of an arbitrary factor does guarantee a sufficient decrease in manipulate power. Instead, we should analyse how the different factors strengthen each other so we can devise a strategy that might help us decrease the manipulative powers of specific IoT technologies.

I proposed aggravating webs as a method for analysing the interplay between the aggravating factors. Aggravating webs consider how different aggravating factors impact each other by focusing on whether the existence of a specific factor in a concrete technology increases the effectiveness of other factors in that technology. By mapping out the relation between aggravating factors in this fashion, it becomes possible to analyse the impact of specific countermeasures regarding a singular factor. Aggravating webs look different for different technologies.

Using aggravating webs to understand the relation between aggravating factors for specific IoT technologies can be the first step for individuals to mitigate the manipulative effects these technologies have in their everyday lives. Since the relations sketched out in aggravating webs, allow users to contemplate the effectiveness of specific countermeasures against a singular aggravating factor in relation to the other factors. This offers a more effective way to take countermeasures against manipulation by IoT devices than focusing on a singular issue without any knowledge regarding the impact of limiting that specific factor.

Conclusion

This thesis answered the following question: "**How can individual IoT users in a neoliberal society develop an understanding of manipulation by concrete IoT devices that negatively impact their autonomy?**" The response to this question consists of two steps:

> First, individual users in a neoliberal society can develop an understanding of IoT devices' manipulation that negatively impacts their autonomy by first conceptualizing manipulation by IoT devices as digital manipulation—a form of manipulation that has increased potency through the aggravating factors of personalization, flow, opacity, lack of user control and omnipresence.
>
> And second, by using aggravating webs as a way to map out the interplay between the aggravating factors for specific technologies, users can develop an understanding of how these specific factors together increase manipulative potency, which is the first step in safeguarding their autonomy.

This conclusion emerged by answering different sub-questions. Let us briefly examine these questions to understand how answering each of them has led to our final answer to the main research question.

The first chapter addressed the question, "What differentiates digital manipulation by IoT devices from non-digital manipulation?". The answer to this question contains two steps. First, through conceptual analysis of manipulation, I derived a foundational definition of manipulation. Second, the concept of aggravating factors was applied to this definition to differentiate manipulation from digital manipulation. The conceptual analysis resulted in the following definition. An influence should be regarded as manipulation when it is an attempt to change an individual's behaviour by undermining their decision-making process by using means of deception or playing on a vulnerability the individual has.

This definition emphasizes the specific process by which the influence is trying to undermine an individual's decision-making process, namely by using means of deception or playing on a vulnerability. It is the effectiveness of achieving these means that differentiates digital and non-digital manipulation. Drawing upon Jongepier & Klenk (2022a), I argued that digital manipulation should be differentiated from non-digital manipulation by five aggravating factors that increase its effectiveness of achieving these means—namely, flow, personalization, opacity, lack of user control, and omnipresence. In light of our definition of manipulation, it is these aggravating factors that increase the effectiveness of the specific means employed to assert influence. However, claiming that digital manipulation by IoT devices can be more effective than non-digital manipulation does not explain the significance of digital manipulation by IoT devices, nor does it show a connection autonomy loss. These two aspects are discussed in the second chapter.

The second chapter provided answers to why we should be concerned about autonomy loss and argued that manipulation by IoT devices is likely to result in autonomy loss. To make this argument, I adopted a personal notion of autonomy (Buss & Westlund, 2018), which is 'the ability to make one's own life'. This notion of autonomy should be understood as a multi-dimensional concept and can (at least) be differentiated into three different dimensions. The structural dimension concerns whether an individual is able to follow their own desires; The temporal concerns how an individual reflects their actions as autonomous in the past, present and future; and the Relational concerned with the power

struggle between an individual and their environment; Having established this framework, I provided several arguments for why autonomy itself is valuable.

Drawing on this conception of autonomy, I provided an argument that digital manipulation is likely to result in a loss of autonomy. Lastly, I elaborated on each of the aggravating factors by analysing a real IoT device to show how every factor in real IoT technology contributes to increased manipulative potential.

In light of this understanding of autonomy and its value and the relation between manipulation and autonomy, it becomes crucial to develop an understanding of these aggravations to accommodate potential solutions for safeguarding autonomy. This understanding is further developed in the third chapter. Building on the conceptual work of the first two chapters. The third chapter answers the main question by developing a method of understanding the relations between aggravating factors and testing this method using a case study.

The case study illustrated first that all aggravating factors can be present in a concrete IoT technology, in our case, a smartwatch. And second, the aggravating factors together are able to diminish autonomy in different dimensions. Then, by drawing on the case study, the chapter argues that mitigating the aggravating factors cannot be done by simply limiting a single factor. Instead, aggravating factors should be understood as a connected whole. Generating this understanding should not be done by considering how one factor impacts one other factor, nor should every possible connection between every possible factor be considered. As a middle ground between these two approaches, aggravating webs are a way to develop an understanding of aggravating factors as a connected whole.

Aggravating webs help users to see what specific aspects of an IoT device increase its manipulative capabilities and also show how these features together form a manipulative whole that is greater than the sum of its parts. Aggravating webs are a way of asking oneself how aggravating factors are present in specific devices and how they are related. Creating an aggravating web can function as the first step in mitigating the aggravating factors that increase the manipulative effects of an IoT device. In extension, they function as the first step to safeguard autonomy since they offer the necessary understanding to analyse the impact of specific countermeasures against an aggravating factor.

Now, do aggravating webs as a way of understanding the digital manipulative influence of IoT devices offer a sufficient understanding to start developing methods for mitigating this influence? My answer would be yes. Using aggravating webs can help individuals consider their relation with concrete IoT technologies in such a way that they can evaluate the impact of specific countermeasures against the manipulative effect of the technology, thereby increasing their chances of safeguarding their autonomy. Especially in the Relational dimension of autonomy, aggravating webs can be useful; they can function as a tool to protect oneself in the power struggle between an individual and the environment. By providing individuals with a better understanding of that environment.

This understanding is the basis for my recommendation to individuals who want to safeguard their autonomy against manipulation by a specific IoT technology.

> First, identify how and to what extent each of the aggravating factors is present in the technology.
>
> Second, once the aggravating factors are identified, use an aggravating web to determine how the factors work together to exploit your vulnerabilities further.
>
> Third, use this newfound understanding to analyse the impact of specific countermeasures against aggravating measures.

I note here that this third step remains under-explored in this thesis. This could be further expanded on in further research. This could be done by letting actual users come up with countermeasures against manipulation by a technology of their choice based on aggravating webs they themselves created.

In finishing, it is important to address two limitations in the answer provided to the main research question. First, the list of aggravating factors I provided could be expanded. Second, the concept of Aggravating webs requires further exploration.

The implications for these two limitations are as follows.

Regarding the first limitation, Jongepier and Klenk (2022a) specifically state that the list they provide is not exhaustive. In my research, I acknowledged this by adding omnipresence as an extra aggravating factor. However, further expansion remains possible. I chose only to add one extra factor to limit the scope of this thesis. As a possible extra factor, I suggest dependence. The concept of dependence might help us understand how the dependence on specific IoT devices might improve their ability to deceive us or play on vulnerabilities we have. However, expanding the list of aggravating factors might impact the practical application of aggravating webs. Adding an extra factor might increase their complexity, which might make it harder for them to be created for specific technology and further increase the risk of ambiguity.

The second limitation concerns the practical effectiveness of aggravating webs in decreasing manipulative potential. This aspect remains unexplored. However, it is for this reason that the thesis is focused on understanding manipulation by IoT devices as a prerequisite for deriving countermeasures instead of focusing on concrete solutions. Empirical research should apply the findings of this thesis regarding aggravating webs to real-world cases to determine its usefulness in predicting the effectiveness of concrete countermeasures against aggravating factors. Only through the practical implementation of aggravating webs can we discover their full utility.

My concrete advice for IoT users who wish to develop an understanding of manipulation by specific IoT devices in their everyday lives is to identify how aggravating factors are present in the device and to map out how these factors relate by using an aggravating web. The understanding that results from this web can then be used to test the impact of countermeasures to mitigate aggravating factors.

Bibliography

Amazon.com: Amazon Dash Button—Frequently Asked Questions: Amazon Devices & Accessories. (n.d.). Retrieved July 17, 2024, from https://www.amazon.com/b?node=17437623011

Anderson, J., & Honneth, A. (2005). Autonomy, Vulnerability, Recognition, and Justice. In J. Anderson & J. Christman (Eds.), *Autonomy and the Challenges to Liberalism: New Essays* (pp. 127–149). Cambridge University Press. https://doi.org/10.1017/CBO9780511610325.008

Anderson, S. (2023). Coercion. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Spring 2023). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2023/entries/coercion/

Apple. (n.d.). *Apple Watch—Close Your Rings*. Apple. Retrieved October 28, 2024, from https://www.apple.com/watch/close-your-rings/

Apple introduces groundbreaking health features. (n.d.). Apple Newsroom. Retrieved December 3, 2024, from https://www.apple.com/newsroom/2024/09/apple-introduces-groundbreaking-health-features/

Ashworth, A. (2005). Aggravation and mitigation. In *Sentencing and Criminal Justice* (pp. 151–181). Cambridge University Press.

Autonomy | Etymology of autonomy by etymonline. (n.d.). Retrieved July 1, 2024, from https://www.etymonline.com/word/autonomy

Baron, M. (2014). The Mens Rea and Moral Status of Manipulation. In C. Coons & M. Weber (Eds.), *Manipulation: Theory and Practice*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199338207.003.0005

Beinsteiner, A. (2020). Conviviality, the Internet, and AI. Ivan Illich, Bernard Stiegler, and the Question Concerning Information-technological Self-limitation. *Open Cultural Studies*, *4*(1), 131–142. https://doi.org/10.1515/culture-2020-0013

Bhargava, V. R., & Velasquez, M. (2021). Ethics of the Attention Economy: The Problem of Social Media Addiction. *Business Ethics Quarterly*, *31*(3), 321–359. https://doi.org/10.1017/beq.2020.32

Biddle, S. (2018, April 13). *Facebook Uses Artificial Intelligence to Predict Your Future Actions for Advertisers, Says Confidential Document*. The Intercept. https://theintercept.com/2018/04/13/facebook-advertising-data-artificial-intelligence-ai/

Bogna, J. (2021, April 28). *What Is Google's FLoC, and How Will It Track You Online?* How-To Geek. https://www.howtogeek.com/724441/what-is-googles-floc-and-how-will-it-track-you-online/

Borgesius, F. J. Z., Möller, J., Kruikemeier, S., Fathaigh, R., Irion, K., Dobber, T., Bodo, B., & Vreese, C. de. (2018). Online political microtargeting: Promises and threats for democracy. *Utrecht Law Review*, *14*(1), Article 1. https://doi.org/10.18352/ulr.420

Bratman, M. E. (2000). Reflection, Planning, and Temporally Extended Agency. *The Philosophical Review*, *109*(1), 35–61. https://doi.org/10.2307/2693554

Burchell, K., Rettie, R., & Patel, K. (2013). Marketing social norms: Social marketing and the 'social norm approach.' *Journal of Consumer Behaviour*, *12*(1), 1–9. https://doi.org/10.1002/cb.1395

Buss, S., & Westlund, A. (2018). Personal Autonomy. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2018). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2018/entries/personal-autonomy/

Cavdar Aksoy, N., Tumer Kabadayi, E., Yilmaz, C., & Kocak Alan, A. (2021). A typology of personalisation practices in marketing in the digital age. *Journal of Marketing Management*, *37*(11–12), 1091–1122. https://doi.org/10.1080/0267257X.2020.1866647

Chae, Y., Nakazawa, M., & Stenger, B. (2018). Enhancing Product Images for Click-Through Rate Improvement. *2018 25th IEEE International Conference on Image Processing (ICIP)*, 1428–1432. https://doi.org/10.1109/ICIP.2018.8451513

Christman, J. (2020). Autonomy in Moral and Political Philosophy. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2020). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2020/entries/autonomy-moral/

Concordia (Director). (2016, September 27). *The Power of Big Data and Psychographics | 2016 Concordia Annual Summit* [Video recording]. https://www.youtube.com/watch?v=n8Dd5aVXLCc

Curtis, A. (Director). (2016, June 20). *The Century of the Self - Part 1: "Happiness Machines"* [Video recording]. https://www.youtube.com/watch?v=DnPmg0R1M04

Deci, E. L., & Ryan, R. M. (2013). The Importance of Autonomy for Development and Well-Being. In B. W. Sokol, F. M. E. Grouzet, & U. Müller (Eds.), *Self-Regulation and Autonomy: Social and Developmental Dimensions of Human Conduct* (pp. 19–46). Cambridge University Press. https://doi.org/10.1017/CBO9781139152198.005

Ellul, J. (1965). *Propaganda: The Formation of Men's Attitudes* (1973rd ed.). Vintage books.

Frankfurt, H. G. (1988). Freedom of the Will and the Concept of a Person. In M. F. Goodman (Ed.), *What Is a Person?* (pp. 127–144). Humana Press. https://doi.org/10.1007/978-1-4612-3950-5_6

Gaur, A., Scotney, B., Parr, G., & McClean, S. (2015). Smart City Architecture and its Applications Based on IoT. *Procedia Computer Science*, *52*, 1089–1094. https://doi.org/10.1016/j.procs.2015.05.122

Global Coverage By floLIVE. (n.d.). floLIVE. Retrieved February 18, 2025, from https://flolive.net/coverage/

Gokhale, P., Bhat, O., & Bhat, S. (2018). Introduction to IOT. *International Advanced Research Journal in Science, Engineering and Technology*, *5*(1), 41–44. https://doi.org/10.17148/IARJSET.2018.517

Google. (n.d.-a). *Google Home*. Retrieved December 2, 2024, from https://home.google.com/welcome/

Google. (n.d.-b). *Yale x Nest*. Retrieved October 11, 2024, from https://store.google.com/us/product/nest_x_yale_lock?hl=en-US

Gorin, M. (2014). Do Manipulators Always Threaten Rationality? *American Philosophical Quarterly*, *51*(1), 51–61.

Greenfield, P. (2018, March 25). The Cambridge Analytica files: The story so far. *The Guardian*. https://www.theguardian.com/news/2018/mar/26/the-cambridge-analytica-files-the-story-so-far

Guide to Imprint™ Smart Maps. (n.d.). Retrieved July 22, 2024, from https://homesupport.irobot.com/s/article/64102

Halffman, W. (2019). Frames: Beyond Facts Versus Values. In E. Turnhout, W. Halffman, & W. Tuinstra (Eds.), *Environmental Expertise: Connecting Science, Policy and Society* (pp. 36–57). Cambridge University Press. https://doi.org/10.1017/9781316162514.004

Harish, J. (2017, July 13). *Xiaomi Mi Smart Shoes review: Obsessively track your workouts*. Android Central. https://www.androidcentral.com/xiaomi-mi-smart-shoes-review

How do I turn off the activity rings? - Apple Community. (n.d.). Retrieved October 29, 2024, from https://discussions.apple.com/thread/252372267

How Google uses location information – Privacy & Terms – Google. (n.d.). Retrieved July 17, 2024, from https://policies.google.com/technologies/location-data?hl=en&gl=en

How Many People Own Smartphones? (2024-2029). (2021, November 19). Exploding Topics. https://explodingtopics.com/blog/smartphone-stats

Hummel, D., & Maedche, A. (2019). How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics*, *80*, 47–58. https://doi.org/10.1016/j.socec.2019.03.005

Introducing YouTube Shorts. (n.d.). [Video recording]. Retrieved July 18, 2024, from https://www.youtube.com/shorts/J38Yq85ZoyY

Jongepier, F., & Klenk, M. (2022a). Online manipulation: Charting the field. In *The Philosophy of Online Manipulation* (pp. 15–48). Routledge.

Jongepier, F., & Klenk, M. (Eds.). (2022b). *The Philosophy of Online Manipulation*. Routledge. https://doi.org/10.4324/9781003205425

Klenk, M., & Hancock, J. (2019). Autonomy and online Manipulation. *Internet Policy Review*, *1*, 1–11.

Kröger, J. L., & Raschke, P. (2019). Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping. In S. N. Foley (Ed.), *Data and Applications Security and Privacy XXXIII* (pp. 102–120). Springer International Publishing. https://doi.org/10.1007/978-3-030-22479-0_6

Lee, C.-H. J., Chang, C., Chung, H., Dickie, C., & Selker, T. (2007). Emotionally reactive television. *Proceedings of the 12th International Conference on Intelligent User Interfaces*, 329–332. https://doi.org/10.1145/1216295.1216359

Leiteritz, R. (2010, April 27). *Copy of Google's submission to several national data protection authorities on vehicle-based collection of wifi data for use in Google location based services*. Google. http://static.googleusercontent.com/media/www.google.com/en//googleblogs/pdfs/google_submission_dpas_wifi_collection.pdf

Levin, B. (2020, February 28). *Apple Watch Activity: Your guide to fitness tracking on Apple's smartwatch*. CNN Underscored. https://www.cnn.com/cnn-underscored/electronics/apple-watch-activity

Lister, C., West, J. H., Cannon, B., Sax, T., & Brodegard, D. (2014). Just a fad? Gamification in health and fitness apps. *JMIR Serious Games*, *2*(2), e9. https://doi.org/10.2196/games.3413

Luguri, J., & Strahilevitz, L. J. (2021). Shining a Light on Dark Patterns. *Journal of Legal Analysis*, *13*(1), 43–109. https://doi.org/10.1093/jla/laaa006

Lukyanchikova, E., Askarbekuly, N., Aslam, H., & Mazzara, M. (2023). A Case Study on Applications of the Hook Model in Software Products. *Software*, *2*(2), 292–309. https://doi.org/10.3390/software2020014

Masjedi, Y. (2024, February 7). *What Does Jailbreaking an iPhone Do? (Risks and Benefits)*. https://www.aura.com/learn/what-does-jailbreaking-an-iphone-do

Mill, J. S. (1859). *On liberty*. Oxford University Press.

Mitton, N., Papavassiliou, S., Puliafito, A., & Trivedi, K. S. (2012). Combining Cloud and sensors in a smart city environment. *EURASIP Journal on Wireless Communications and Networking*, *2012*(1), 247. https://doi.org/10.1186/1687-1499-2012-247

Moore, P. (Director). (2022, November 23). *Eufy leaking your "private" images/faces & names... To the cloud.* [Video recording]. https://www.youtube.com/watch?v=qOjiCbxP5Lc

Nakamura, J., & Csikszentmihalyi, M. (2014). The Concept of Flow. In M. Csikszentmihalyi (Ed.), *Flow and the Foundations of Positive Psychology: The Collected Works of Mihaly Csikszentmihalyi* (pp. 239–263). Springer Netherlands. https://doi.org/10.1007/978-94-017-9088-8_16

Noggle, R. (1996). Manipulative Actions: A Conceptual and Moral Analysis. *American Philosophical Quarterly*, *33*(1), 43–55.

Norman, D. (1999). Affordance, conventions, and design. *Interactions*, *6*, 38–42. https://doi.org/10.1145/301153.301168

Nyholm, S. (2022). Technological Manipulation and Threats to Meaning in Life. In *The Philosophy of Online Manipulation* (pp. 235–252). Routledge.

Oshana, M. (2016). *Personal Autonomy in Society* (1st ed.). Routledge. https://doi.org/10.4324/9781315247076

Pedwell, C. (2017). Habit and the Politics of Social Change: A Comparison of Nudge Theory and Pragmatist Philosophy. *Body & Society*, *23*(4), 59–94. https://doi.org/10.1177/1357034X17734619

Pepp, J., Sterken, R., McKeever, M., & Michaelson, E. (2022). Manipulative machines. In F. Jongepier & M. Klenk (Eds.), *The Philosophy of Online Manipulation* (1st ed., pp. 91–107). https://doi.org/10.4324/9781003205425-6

Pérez-Verdugo, M., & Barandiaran, X. E. (2023). Personal Autonomy and (Digital) Technology: An Enactive Sensorimotor Framework. *Philosophy & Technology*, *36*(4), Article 84. https://doi.org/10.1007/s13347-023-00683-y

Pilkington, E., & Aratani, L. (2024, July 19). US transportation, police and hospital systems stricken by global CrowdStrike IT outage. *The Guardian*. https://www.theguardian.com/technology/article/2024/jul/19/crowdstrike-microsoft-outage

Porter, E. (2023, November 7). Opinion | Economists loved so-called nudge thinking. But it's a dud. *Washington Post*. https://www.washingtonpost.com/opinions/2023/11/07/nudge-thinking-economics-policy-flawed/

Prichard, J., & Hayden, M. (2008). Assessing the readability of freeware end-user licensing agreements. *Issues in Information Systems*, *9*(2), 452–459.

Privacy. (n.d.). Apple. Retrieved December 3, 2024, from https://www.apple.com/privacy/

**Privacy Not Included review: Fitbit Charge 6*. (n.d.). Mozilla Foundation. Retrieved December 3, 2024, from https://foundation.mozilla.org/en/privacynotincluded/fitbit-charge-6/

Rawls, J. (2005). *Political Liberalism: Expanded Edition* (2nd ed.). Columbia University Press. https://www.jstor.org/stable/10.7312/rawl13088

Raz, J. (1988). *The Morality of Freedom*. Oxford University Press. https://doi.org/10.1093/0198248075.001.0001

Rudinow, J. (1978). Manipulation. *Ethics*, *88*(4), 338–347. https://doi.org/10.1086/292086

Samsung. (2024). *Family Hub*. Samsung Electronics America. https://www.samsung.com/us/support/answer/ANS00049761/

Sher, S. (2011). A Framework for Assessing Immorally Manipulative Marketing Tactics. *Journal of Business Ethics*, *102*(1), 97–118. https://doi.org/10.1007/s10551-011-0802-4

Signal Messenger: Speak Freely. (n.d.). Signal Messenger. Retrieved July 18, 2024, from https://signal.org/

signalapp/Signal-Desktop: A private messenger for Windows, macOS, and Linux. (n.d.). Retrieved July 18, 2024, from https://github.com/signalapp/Signal-Desktop

Solove, D. J. (2004a). *The Digital Person: Technology and Privacy in the Information Age*. New York University Pres. https://papers.ssrn.com/abstract=2899131

Solove, D. J. (2004b). *The Digital Person: Technology and Privacy in the Information Age (full text of book)* (SSRN Scholarly Paper 2899131). https://papers.ssrn.com/abstract=2899131

Stanovich, K. E., & West, R. F. (2000). Advancing the rationality debate. *Behavioral and Brain Sciences*, *23*(5), 701–717. https://doi.org/10.1017/S0140525X00623439

Stoljar, N. (2024). Feminist Perspectives on Autonomy. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Summer 2024). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2024/entries/feminism-autonomy/

Sunstein, C. (2015). Fifty Shades of Manipulation. *Journal of Marketing Behavior*, *1*. https://doi.org/10.1561/107.00000014

Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, autonomy, and manipulation. *Internet Policy Review*, *8*(2). https://policyreview.info/articles/analysis/technology-autonomy-and-manipulation

Tamburrelli, G., & Margara, A. (2014). Towards Automated A/B Testing. In C. Le Goues & S. Yoo (Eds.), *Search-Based Software Engineering* (Vol. 8636, pp. 184–198). Springer International Publishing. https://doi.org/10.1007/978-3-319-09940-8_13

Thaler, R. H., & Sunstein, C. R. (2021). *Nudge: The Final Edition*. Yale University Press.

The Best Smart LED Light Bulbs for 2024. (n.d.). PCMAG. Retrieved November 12, 2024, from https://www.pcmag.com/picks/the-best-smart-light-bulbs

Verbeek, P.-P. (2015). A Field Guide to Postphenomenology. In R. Rosenberger & P.-P. Verbeek (Eds.), *Postphenomenological Investigations: Essays on Human-Technology Relations* (pp. 9–41). Lexington Books.

WATERH. (n.d.). *WaterH BOOST*. WaterH™. Retrieved July 22, 2024, from https://www.waterh.com/products/waterh-boost-smart-water-bottle

Wen, M. (2017, July 26). iRobot shares surge on strong sales of Roomba vacuum cleaners. *Reuters*. https://www.reuters.com/article/technology/irobot-shares-surge-on-strong-sales-of-roomba-vacuum-cleaners-idUSKBN1AB2QU/

Wood, A. W. (2014). Coercion, Manipulation, Exploitation. In C. Coons & M. Weber (Eds.), *Manipulation: Theory and Practice* (pp. 17–50). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199338207.003.0002

Zarouali, B., Dobber, T., De Pauw, G., & de Vreese, C. (2022). Using a Personality-Profiling Algorithm to Investigate Political Microtargeting: Assessing the Persuasion Effects of Personality-Tailored Ads on Social Media. *Communication Research*, *49*(8), 1066–1091. https://doi.org/10.1177/0093650220961965

Zhang, J., & Tao, D. (2021). Empowering Things With Intelligence: A Survey of the Progress, Challenges, and Opportunities in Artificial Intelligence of Things. *IEEE Internet of Things Journal*, *8*(10), 7789–7817. IEEE Internet of Things Journal. https://doi.org/10.1109/JIOT.2020.3039359

Zhang, X., Wu, Y., & Liu, S. (2019). Exploring short-form video application addiction: Socio-technical and attachment perspectives. *Telematics and Informatics*, *42*, 101243. https://doi.org/10.1016/j.tele.2019.101243

Zuboff, S. (2019). *The age of surveillance capitalism*. profile books.