

Responsible Consumer Behavior in the Digital Age

EMPIRICAL EVIDENCE ON THE IMPACT OF (MICRO-)ENVIRONMENTAL
FACTORS ON CONSUMERS' SOCIALLY RESPONSIBLE BEHAVIOR

Dissertation

A thesis submitted in fulfillment of the requirements for the degree of
Dr. rer. pol. from the Faculty of Economic Sciences University of
Göttingen

by

Nils Engelbrecht

born in Kassel

Göttingen, August 2022

First Supervisor: Prof. Dr. Kilian Bizer

Second Supervisor: Prof. Dr. Achim Spiller

Third Supervisor: Prof. Dr. Holger Rau

Contents

1	Introduction and Summary	1
2	The Interplay of Different Excuses in Determining Moral Wiggle Room <i>Mattheus Brenig, Nils Engelbrecht, and Daniel Hermann</i>	23
3	Consumer Behavior Under Benevolent Price Discrimination <i>Alexander Erlei, Nils Engelbrecht, and Mattheus Brenig</i>	65
4	The Virtual Online Supermarket <i>Nils Engelbrecht, Tim-Benjamin Lembcke, Alfred Benedikt Brendel, Kilian Bizer, and Lutz M. Kolbe</i>	115
5	Behavioral Design In Online Supermarkets: How Virtual Shopping Cart Functions Impact Sustainable Consumption <i>Tim-Benjamin Lembcke, Nils Engelbrecht, Mathias Willnat, and Sascha Lichtenberg</i>	153
6	Ethical Dimensions in Digital Nudging <i>Tim-Benjamin Lembcke, Nils Engelbrecht, Alfred B. Brendel, and Alan R. Dennis</i>	171

CHAPTER I

Introduction and Summary

*Problems cannot be solved at the same level of awareness that
created them.*

— Einstein (1905)

Sustainable consumption has become one of the major narratives of our time. Especially for developed countries, there are strong indications that predominant consumption and production patterns have a direct or indirect negative impact on the individual (and public) well-being, social equity, and the environment (SDG 2018; Willett et al. 2019). In particular, current *food systems*¹ are a leading cause of ecocide as well as for diseases caused by a poor or unhealthy diets (Food and United Nations; World Health Organization 2019). Hence, transforming food consumption is deemed to be essential for tackling major local and global societal challenges such as obesity or climate change (Vermeir et al. 2020). At present, it seems as if many players are ready to take up this challenge: An increasing number of governmental institutions, non-governmental organizations, and players from the industry have already launched own or joint efforts to promote "better" consumption patterns and eating habits (Schlaile et al. 2018).

However, free and liberal market economies are based on the paradigm of *consumer sovereignty*. Consequently, without *responsible consumer behavior (RCB)*, government regulations on the supply-side and companies' voluntary internal policies like *corporate social responsibility (CSR)* are doomed to become futile attempts (Lorek and Vergragt 2015; Vitell 2015; Parente 2020). In other words, sufficient demand for e.g., fair, environment-friendly, or ethically produced goods and services is crucial for sustainable shift in consumption.

In recent years, consumers have become more sensitive to the social, economic, and environmental consequences of their market choices. More and more consumers vocally oppose socially irresponsible production patterns (Shamir 2008; Fourcade and Healy 2007) and many studies have identified purchasing intention for e.g., fair, ethical, or "green" goods (Öberseder et al. 2011). However, as market shares for many of these goods remain small, an inconsistency between individuals' concerns and attitudes about ethical or environmental issues and their actual consumption behavior can be observed (Terlau and Hirsch 2015; Aschemann-Witzel and Zielke 2017). The literature offers manifold explanations for why this so-called *attitude-behavior gap* exists, including a broad spectrum of cognitive, social, and environmental influences (see

¹Food systems are defined as "compositions of interlinked elements and activities aimed for the production, processing, distribution, and consumption of food" (Smetana et al. 2019).

Vermeir and Verbeke 2006 for an overview). Some of these reasons are directly related to (i) some essential properties of market environments in a broad economic sense (e.g., not sufficiently transparent supply-chains in the agricultural market)² and (ii) micro-environmental factors in physical marketplaces like supermarkets.

RCB in Market Transactions

There is a growing strand of experimental research analyzing to what extent consumers exhibit social responsibility in market transactions and which market conditions inhibit or encourage this behavior (e.g., Falk and Szech 2013; Falk et al. 2020; Bartling et al. 2015; Bartling et al. 2017; Pigors and Rockenbach 2016; Irlenbusch and Saxler 2019). The starting point for these studies is the critical reflection of the following hypothesis: Compared to individual decision-making, markets erode social behavior because social responsibility in market transactions is diffused or shared.

Today, most markets are complex, multilateral, and highly competitive. Hence, each actor can easily believe that her individual behavior is not (or only little) pivotal for the aggregated market outcomes. This allows consumers to find an excuse for behaving selfishly and for ignoring (negative) effects on others (Falk and Szech 2013; Kirchler et al. 2016).

However, data on causal effects of market environments on RCB is still rare and evidence is inconsistent. While some empirical studies suggest that individuals in market environments tend to act less socially responsible and more selfishly compared to non-market transactions (e.g., Falk and Szech 2013), other studies do not (e.g., Bartling et al. 2015). Inconclusive findings could be due to variations in experimental designs that aim to mirror different market characteristics. For instance, previous studies investigate how certain aspects of the institutional environment like (i) competitive pressure (Pigors and Rockenbach 2016; Bartling et al. 2015), (ii) supply chain transparency (Kraft et al. 2016), (iii) social information about the behavior or intentions of others in markets (Irlenbusch and Saxler 2015; Friedrichsen and Engelmann 2018), (iv) different market regimes (e.g., price taking vs. bargaining; Kirchler et al. 2016), (v) the "scope" of a negative externality from market transactions (Bartling et al. 2019)

²The agricultural market describes the entirety of institutions that cover all production processes and services involved in moving groceries from the farm to consumers.

³, (vi) the diffusion to be pivotal for market outcomes (Falk et al. 2020), (vii) social norms (e.g., acceptance of outcome-maximizing behaviors Bartling and Özdemir 2017, or (viii) legal or freely imposed standards (Danz et al. 2012; Etilé and Teyssier 2016) affect CSR and RCB.

Consequently, it seems to depend on the specific market environment how easily consumers find and exploit *moral wiggle* rooms in order to feel less responsible for the negative impact of their consumption or market outcomes at large (Dana et al. 2007). For instance, as a bare minimum, for acting socially responsible, a consumer needs some degree of actual or perceived influence on the market. In more general terms, a single consumer's influence over market outcomes could be described as a function of market size and other consumers' behavior. For example, a socially irresponsible (responsible) choice in a small market with otherwise socially responsible (irresponsible) consumers is more pivotal, than it would be in a larger market with the same characteristics. How consumers perceive these market conditions and consequently their market power may influence their willingness to internalize the negative externalities (Falk et al. 2020).

Moreover, the availability of information is widely considered to be a key factor for RCB. In many markets, supply chains might be long and non-transparent. Information is distributed asymmetrically along the entire value chain and consumers deal with insufficient information about the environmental and social impact of goods and services. This results in uncertainty, mistrust, lower demand as well as less willingness to pay for e.g., fair or sustainable labeled products (Terlau and Hirsch 2015). For instance, it is difficult for consumers to assess whether price premiums for fairly produced goods are justified because they improve the situation of the workers or not (Pigors and Rockenbach 2016).

However, for consumers a lack of transparency may also be exploited as an excuse to justify harmful effects of their actions on others. Empirical results suggest that consumers remain willfully ignorant about the negative externalities of their consumption choices, even when they are aware that the information is freely available and can be retrieved with minimal effort (e.g. Ehrich and Irwin 2005; Thunström et al. 2016; Zane et al. 2016).

³It was investigated whether it makes a difference if market transactions had a strong and concentrated negative impact on a small group of workers outside the market or create relatively small harm for a large number of workers.

Therefore, there might be a strategic rationale for ignoring information on negative externalities (Golman et al. 2017). It allows individuals to behave selfishly while maintaining the illusion of not being inconsistent with respect to one's own or others' social norms (Bénabou and Tirole 2011; Reczek et al. 2018). In other words, acting selfishly might entail "moral costs". These costs might be smaller if individuals avoid information about the consequences of their decisions (Serra-Garcia and Szech 2019).

In addition, not only information on the impacts of consumption choices might affect RCB, but also social information about the observed, implied, or imagined behavior or presence of other social entities. For a variety of reasons, such information can have a significant impact on an individual's behavior. People may have ambitions for conformity or – on the contrary – strive for individuality. They could seek social esteem or try to avoid social disapproval. They could use the observed behavior of others as a reference point for efficient transactions or as a guide to prevent risks (Cialdini 2007; Farrow et al. 2017). It seems that consumers care more about not appearing to buy a good that is unethical, unsustainable, or unsocially produced than they do about actually not buying it.

Among other factors, the degree and manner of social influence depends on the specific market context. Traditionally, market transactions are social activities. They are characterized by direct or unintended consumer-to-consumer interactions and are therefore rich and saturated with numerous social cues. Hence, it is relatively easy for individuals to observe the behavior of other consumers and get an idea of prevailing social norms. For instance, to reduce uncertainty and purchase risks, consumers' may refer to the observed behavior of these unacquainted consumers, e.g., by imitating their product choices (Uhrich and Luck 2012). Such a mimicking or herding behavior has been well documented by a large body of research and in different settings (e.g., Tanner et al. 2008).

In contrast, transactions in other markets may be much more functional, impersonal, and anonymous. Hence, consumers can not refer to (observable) social information like the behavior of other social entities around them. In other words, there is no option to "outsource" difficult decisions by simply doing what others do. Equally, this uncertainty could give consumers *moral wiggle room* not to behave in a socially responsible manner. This particularly concerns situations where their impact on market outcomes is uncertain as well (Garcia et al. 2002; Fischer et al. 2011; Shleifer 2004). Consequently, behavior may vary depending on how strongly consumers' perceptions of being pivotal for a market outcome reflect social information.

Hence, for a deeper understanding of the intricate relationship between specific market settings and RCB, there is a need for comparable data on how various market conditions affect individual RCB and RCB in interaction. This is especially relevant since consumers make transactions in a variety of (sub)markets, characterized by different institutional rules, competitive structures, or levels of transparency. Moreover, due to technological and social change, markets are undergoing massive change. In this vein, alternative market structures (e.g., platform economies), types of distributions (e.g., sharing economies), and manufacturing processes (e.g., community-supported production, mass customization, or other made-to-order concepts)⁴ have emerged. In addition, modern information and communication tools have opened up opportunities to increase transparency in the supply chain and to disclose information to consumers (e.g., via websites, smart labels, or QR codes on packaging; Kraft et al. 2016). Equally, customer-level data availability has exploded, creating unprecedented opportunities for suppliers to implement segmentation strategies or differentiating pricing models (e.g., dynamic or personalized pricing). Not lastly due to these trends, evidence-based policy-making aiming to support RCB has an ongoing need for empirical insights to identify effective interventions for the respective market setting.

The impact of micro-environmental factors on RCB

While the composition of markets seems to be important for RCB, purchase decisions still take place at the Point-of-Sale (POS). Therefore, there is also a growing body of literature that discusses the impact of the micro-environment (e.g., of a store) on consumers' purchase decisions (e.g., Lunn 2014; Hollands et al. 2013; Hoek et al. 2017).

Purchases for many consumer goods like groceries are frequently repeated decisions characterized by a habitual nature and without a high level of involvement (Kalnikaite et al. 2013). In these cases, consumers do not spend much time and (cognitive) effort searching for information and comparing different product features and prices. Instead, choices tend to be "fast and frugal".

Modern marketplaces like supermarkets usually offer thousands of different items and use a variety of in-store communication and marketing instruments. Hence, they are full of physical, verbal, emotional, and/or social stimuli. Consequently, consumers'

⁴While Made-to-order was predominantly used for complex and highly customized products such as industrial machines or cars, it has recently also be applied to mass products such as clothes (Bernard et al. 2012)

decision-making at the POS takes place in a cognitively stressful environment. Such environments trigger unconscious and automatic responses to micro-environmental factors like product placement (Baker et al. 1994; Cohen and Babey 2012). Consumers' reaction to these "seemingly irrelevant" (Thaler 2016) factors remain unaffected even when they are aware of them (Turley and Milliman 2000). Therefore, it does not seem to be very surprising that, up to 80 percent of all purchase decisions for groceries are impulsive and unplanned (Nordfalt 2009; Pornpitakpan and Han 2013; Amos et al. 2014).

At the same time, conscious decision-making plays an important role in allowing individuals to resist their impulses, desires, or bad habits. Hence, it is argued that thoughtless designed micro-environments might hinder consumers to behave according to their intentions or attitudes, even if this was not intended (Terlau and Hirsch 2015). Therefore, analyzing the POS represents a promising approach to gain further insights into the determinants driving RCB. Especially policy-makers are increasingly interested in empirical evidence on (i) how a store's micro-environment can influence consumer behavior and (ii) how this knowledge can be used to design behaviorally informed policy interventions aimed at promoting more RCB (Hartmann-Boyce et al. 2018; Just and Gabrielyan 2018; Just and Byrne 2020).

In particular, one concept has gained attention for its claim to take up empirical insights about systematic behavior patterns in human decision-making (e.g., biases; Tversky and Kahneman 1974; Camerer et al. 2004; Fehr and Gächter 2002) and transfer them into a set of less intrusive behavioral change techniques termed as *nudges*. A nudge is a micro-targeted interventions that: "alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates. Putting fruit at eye level counts as a nudge. Banning junk food does not" (Thaler and Sunstein 2008).

However, for several reasons, there is still a controversial debate about whether such intervention techniques are appropriate to be used in consumer policy or not. First, there is an ethical debate on whether and under which conditions nudges should be used to alter an individual's behavior in the desired direction (e.g., Schubert 2017; Lembcke, Engelbrecht, Brendel, Herrenkind, et al. 2019). Second, due to a lack of empirical evidence, the effectiveness of such an approach is still questioned (Hollands et al. 2013; Hummel and Maedche 2019). Third, only little is known about the extent

to which findings from physical stores can be transferred to nudging in digital or hybrid retail channels (see Lembcke, Engelbrecht, Brendel, and Kolbe 2019).

The last point is particularly relevant since retail in the 21st century is undergoing a digital transition. Most importantly, this manifests in increasingly faster growing e-commerce. There is initial evidence that digital choice environments such as online shops should not be considered mirrors of their real world-equivalents (e.g., Pozzi 2012). For instance, differences exist in product presentation (e.g., symbolically vs. physically; Huyghe et al. 2017, in integrating AI-assisted decision support tools (Kozyreva et al. *in press*), or in the amount of interpersonal interactions and social elements (Ogonowski et al. 2014; Lu et al. 2016). Taking this into account, it is quite unclear (i) how those contextual differences between both retail channels can affect RCB and (ii) how (digital) nudges should be designed in order to be effective yet ethical in e-commerce (Pitts et al. 2018).

However, digital information does also influence shopping in physical stores and has contributed to emerging hybrid formats like *click-and collect*. For instance, many consumers already use their smartphones in brick-and-mortar stores to get further product information or for price comparisons. Additionally, retailers have already started to integrate digital services like in-store terminals or even humanoid robots (e.g., for customer services) into physical stores (Karar et al. 2019; Vannucci and Pantano 2019). Hence, former separated and sometimes competing retail channels are increasingly being integrated into an *omni-channel* experience. This results in a “user journey” through several analog, digital, and blended (e.g., in the case of augmented reality) micro-environments before making a purchase (Lembcke, Engelbrecht, Brendel, Herrenkind, et al. 2019; Schaer and Stanoevska-Slabeva 2019). So far, only little is known about how *omni-channel* retailing might affect the effectiveness and design requirements of (digital) nudges.

Overall, previous research highlights the importance of the decision context for individuals’ sustainable consumption behavior. However, even though research on RCB has grown rapidly in the last two decades, this trans-disciplinary field of research is still in its infancy (Reisch and Sandrini 2015). Hence, there is still need for more research that systematically analyzes the impact of specific compositions of market-environments and micro-environmental factors on RCB (Eckhardt et al. 2010).

Especially when it comes to not losing sight of current trends in a rapidly changing and increasingly digital economy and society. This is emphasized by the need to not

lose sight of current trends in a rapidly changing and increasingly digital economy and society.

Chapter Overview

This book contributes to both above-mentioned research strands. In particular, a selection of five research papers is presented that examine how different aspects of modern and digital market- and micro-environments affect RCB. This overview provides a brief summary of all five research papers, their key results, and (where applicable) implications for researchers and policy-makers.

The first two papers focus on the interplay between RCB and certain key characteristics of markets in general, like uncertainty about external effects of one's market decisions or pricing strategies. Both studies use abstract experimental set-ups where a subject's choice has an impact on the distribution of outcomes between himself and other players. At the same time, the processes that leads up to this distribution and other non-monetary factors are also examined. Hence, both studies rather focus on the macro and meso perspective of RCB.

In contrast, the remaining papers turn their attention to far more subliminal and seemingly irrelevant influencing factors that are to be located at the micro level of consumer behavior, like nudges. Here, we are primarily concerned with pointing out new tools and approaches that can promote research and evaluation of (digital) nudges and similar intervention techniques. These are envisioned to enable an evidence-based discussion on how they can be applied for political incentives or regulations.

Chapter Two: The Interplay of Different Excuses in Determining Moral Wiggle Room

The first paper focuses on one key component inherent to most market-decisions: uncertainty about the (positive or negative) impact of one's decision on others or the environment. Several experimental studies propose that many individuals aim to be fair (e.g., in behaving socially responsible) or at least aim to appear to be (for a review see e.g., Johnson and Mislin 2011). However, as soon as the relationship between their actions and outcomes is blurred, individuals defer from this good intention because it is easier for them to excuse their more selfish and less socially responsible behavior (Dana et al. 2007; Bartling and Fischbacher 2012; Coffman 2011). Such an excuse-driven behavior might be favored or restricted by the scope of uncertainty.

Depending on the specific decision-context, consumers have more or less information about the impact of their own decisions as well as on the behavior of other consumers. For instance, on the aggregated level, one's negative impact might be compensated by socially responsible behavior of other market participants and vice versa. Hence, the question arises under which circumstances and to what extent disclosure policies that reduce uncertainty and increase transparency about the consequences of consumption decisions can successfully promote RCB.

In this chapter, we present the results of a large-scale online experiment that examined the interplay of various excuses and different degrees of being pivotal for negative externalities. Further, we analyzed RCB in a context where subjects could always prevent at least some of their negative externality from occurring. Thus, they could not argue that market outcomes would be the same regardless of their own decision.

For our study, we utilized a two-round ultimatum bargaining situation where transactions between two proposers (A_1 and A_2) and one responder (B) might negatively affect the payoffs of at least one of two inactive third players (C_1 and C_2). Group composition was fixed above both rounds. In the first round, a responder faced a decision between offers from two proposers: One offer was more beneficial to the responder but entailed a loss for one of the inactive third players ("selfish offer"). The other offer was less beneficial for the responder but did not (negatively) affect the inactive third party's endowment ("unselfish offer"). We measured consumer social responsibility as accepting the proposer offer that did not entail a loss for an inactive third player. Round two followed the same procedure, but offers were not chosen by a proposer but determined by a computerized urn draw ("move by nature"). Each move of nature led to one of two scenarios in round 2: (i) a selfish market where only two "selfish offers" existed or (ii) an unselfish market where only two "unselfish" offers existed. Specific parameters of the move by nature were set by the experimental condition.

Within this set-up, we analyzed the interplay of three experimental conditions: First, to simulate different degrees of consumer agency in markets, we varied how responders' choices in the first round affected the move by nature and therefore the market structure in round two. This allowed us to consider three specific excuses for selfish behavior: the bystander excuse ("if I don't make the transactions, someone else will"), the replacement excuse ("someone else will avert harmful effects on others"), and diffusion of being pivotal. For instance, we varied whether a responder's socially responsible decision in round one was either (i) a necessary and sufficient condition

(the responder is fully pivotal); (ii) neither necessary nor a sufficient condition (the responder is not pivotal at all); (iii) a necessary but not sufficient condition (replacement excuse), or (iv) not necessary but sufficient condition (bystander excuse) for a socially responsible future market outcome. Second, in accordance with two different kind of manufacturing processes in markets, we varied the timing of when the negative effect occurred for the third party. In our Made-to-Stock condition (MTS), all the consequences associated with the production process had already occurred as soon as an offer was made. In our Made-to-Order condition (MTO), the negative externalities only occurred when the offer was accepted by the responder. Third, to test whether responders take advantage of the opportunity to remain ignorant about the impact of their decision on third parties (*strategic ignorance*), we implemented two information conditions. In one condition, subjects were informed about the payoff consequences for third parties in the experimental instructions. In the other condition, information about third-party payoff consequences were not disclosed in the instructions, but subjects knew that they would be given the opportunity to access the information before deciding. Introducing strategic ignorance allowed us to closely mirror how consumers retrieve information at the POS (e.g., via product labels).

Our data revealed that as soon as subjects had to deal with uncertainty about their impact on the future market, selfish behavior increased by approximately one-fourth. This result held both for different degrees of pivotality and for contexts where a proposer could apply the bystander or replacement excuse. However, behavior seemed to be robust to changes in uncertainty: RCB significantly decreased when consumers no longer had any impact on the market. Here, one implication is that RCB can be promoted by policies that increase consumers' perceived market power.

Further, we found no significant effect of strategic ignorance on RCB. About two-thirds of subjects retrieved information about the consequences of their actions on third parties. This result implies that for situation where consumers have neither any prior information nor expectations about the impact of their market transactions on others, (smart) disclosure policies might be eagerly embraced.

Chapter Three: Consumer Behavior under Benevolent Price Discrimination

In chapter three, we present results from a research paper aiming to assess consumer reactions towards a socially responsible differentiating pricing model, termed as *benevolent price discrimination* (BPD). BPD is defined as a pricing strategy that aims at

more equitably distributing economic gains. For this purpose, it benefits financially disadvantaged groups or individuals whilst not putting anyone else at a disadvantage in absolute terms. Meaning, it will decrease prices for certain individuals or groups, but not increase prices for the others. Thus, nobody is worse off as a consequence of BPD and it can be termed socially responsible. Analyzing consumers' reactions to this very moderate form of price discrimination provides indications on the potential of price segmentation strategies and poaching prices. Moreover, our approach allows to gain insight on meaningful degrees of strategic price obfuscation or the targeted deployment of certain price strategies as part of a CSR communication strategy.

Across a series of five incentivized, context-neutral, and controlled online-experiments, we observed consumers' purchase decisions and (costly) switching behavior to or away from a price-discriminating store. Thus, we were able to quantify both preferences for and against BPD. Overall, we found clear and robust evidence that a large share of consumers is averse to BPD. Across all studies, between 30 and 40 percent of consumers were willing to costly switch to a second seller who did not offer a discount for low-income consumers. These results held, even when it was highly salient that each consumer is priced according to the same rule and that income differences are based on a random draw. The same applied when consumers were aware that they have control over future prices and thereby other consumer outcomes. Further, our results indicate that these behavioral constraints were not driven by (i) consumers' perceptions of sellers' intentions for introducing BPD, (ii) reciprocity, or (iii) a lack of transparency about the process or effects of BPD. In contrast, costly support for BPD was very rare and needed a high level of transparency as well as consumers' agency over future market transactions. Hence, we found no evidence that price discriminating stores may experience positive consumer migration net-effects.

In sum, our results indicate that distributive fairness concerns play only a minor role for consumer reactions towards differential pricing. BPD does not meet consumers' expectations about the firm's CSR. Rather, sellers should expect consumer rejection and a decline in demand if they choose to establish such a pricing model. Therefore, it seems to be reasonable for sellers to hesitate to deviate from uniform pricing.

Chapter Four: The Virtual Online Supermarket

In chapter four, we turn our focus to in-store interventions in (digital) retail channels that aim to promote RCB. We introduce VOS – a novel, modular, and open-source web-

application for designing and conducting computer-simulated shopping experiments in a state-of-the art e-commerce environment. Its front-end was designed to emulate the user interface and functions (e.g., navigation tools) of a modern online shop as realistically as possible. At the same time, nearly every element of this digital shopping environment can be easily manipulated and redesigned. This is possible due to a unique feature of our tool: the visual administration interface (VAI). Using the VAI, a large number of configurations can be implemented with just a few mouse clicks and without programming knowledge. For instance, it can be used to manage the product database as well as to create, configure, and test experimental treatments in a demo environment (sandbox testing). Here, researchers can built on a set of pre-defined modification options (use cases). These use cases cover different types of in-store interventions including economic incentives (e.g., taxes and/or subsidies), changes to the store's micro-environment (e.g., product arrangement), knowledge-based interventions (e.g., labeling and scores), and decision-support tools (e.g., swap options).

In this manner, the impact of several in-store interventions on consumers' shopping behavior can be analyzed under controlled experimental conditions. Since it is not always feasible to conduct field studies in (online) stores, VOS offers a complementary or alternative way to gain further insights into the determinants of consumers' sustainable consumption patterns. Evidence-based consumer policy and CSR strategies have a constant need for those insights which can inform new policy interventions or be used to (re)design specific elements of a store environment (e.g., product packaging, marketing communication).

In a pilot study with 29 students and university staff members, VOS has proven to be ready to use: The application's user-friendliness and usability was evaluated predominately positive. In addition, subjects stated that their shopping experience and behavior corresponded largely to their shopping reality. Technical functionality was also demonstrated. All data was recorded as expected and we were able to generate and analyze a broad set of variables about subjects purchases and in-store behavior dynamics.

In sum, VOS provided all means to run computer-simulated shopping experiments in a realistic e-commerce environment. Using VOS, valid purchase and in-store behavior data can be collected at a relatively low cost and effort. Given these benefits, we hope that our tool will prove useful for many researchers to learn more about the

impact of micro-environmental factors on consumers in-store behavior and to examine key issues for promoting sustainable consumption

Chapter Five: The impact of Virtual Shopping Cart Functions on Sustainable Consumption

In chapter five, we present the results of a computer-simulated shopping experiment. Using an early version of VOS, the study aimed to analyze to what extent real-time spending feedback (RSF) could be part of an in-store intervention to promote sustainable food choices.

In particular, we built upon the budgeting and spending literature (e.g., van Ittersum et al. 2010; van Ittersum et al. 2013) to design an incentive compatible shopping task: Subjects were asked to buy groceries from a predetermined shopping list with eight broad and common categories (e.g., “one item from fruits and vegetables”). For this task, each subject received a budget of 30 EUR. Each subject who did not overspend her budget had a 20 percent chance to win her final shopping cart. Moreover, subjects were informed that they did not need to spend their entire budget. Instead, in case of winning, they would additionally receive their remaining budget in cash. Consequently, there were clear incentives for subjects to stay within their budget and to make authentic purchases.

Within this set-up, we ran two experimental treatments: In our baseline, subjects did not receive any RSF during shopping (No RSF). In contrast, in our intervention treatment subjects were able to retrieve information about their spending at any time of their shopping trip (RSF).

Our results indicate that RSF significantly reduced underspending and prevented overspending. Consequently, subjects in RSF were able to utilize their entire budget more effectively. Further in RSF, subjects did neither significantly shop more items in total nor more different items but spent significantly more on organic food items. Surprisingly, this effect was driven by subjects with weak purchasing intentions for organic food whilst effects for the expected “target group” of subjects with strong organic purchasing intention remain small and not significant.

Chapter Six: Ethical Dimensions in Digital Nudging

The final chapter six reflects ethical considerations in (digital) nudging. In a research opinion paper we discuss whether and to what extent ethical dimensions that have

been identified and well-established for analog nudges can be transfer into the digital sphere.

We argue that individuals have adopted specific behavioral patterns in their use of digital devices. These particularly include (i) attention spans and consequently also interactions tend to become shorter, (ii) multi-screen content consumption becomes more common, and (iii) more shallow forms information processing (e.g., browsing or skimming) are spreading. Consequently, human attention has become a scare resource.

Consequently, users might be even more susceptible to modifications in the choice architecture that could intentionally or unintentionally alter their behavior. Digital choice architects can purposefully control which elements are presented and how these may steer, enable, or restrict users' behavior. In contrast to (re-)design choices in the analog world, this can be implemented with less physical or monetary efforts. Further, AI-assisted tools offer innovative ways to support or alter user behavior. It is obvious that these possibilities are a doubled-edged sword: On the one hand, well designed digital artifacts may steer users towards their own interests or for benevolent reasons. On the other hand, the same choice architecture techniques might be used for "dark patterns" (e.g., tricking users into signing up for something they do not even like or need; Gray et al. 2018).

We conclude that ethical considerations derived from analog nudging must be adapted to and reconsidered for digital choice environments. Three ethical dimensions proposed by Thaler and Sunstein 2008 structure our discussion on ethical issues in digital nudging: (i) transparent disclosure of nudges, (ii) preserving an individual's freedom of choice, and (iii) user and social goal-oriented justification of nudging. For each dimension, we are highlighting differences and similarities between analog and digital environments, underline the respective necessities for adapting the dimensions, present avenues for future research, and propose implications for practice.

References

- Amos, C., G. R. Holmes, and W. C. Keneson (2014). "A meta-analysis of consumer impulse buying". In: *Journal of Retailing and Consumer Services* 21.2, pp. 86–97.
- Aschemann-Witzel, J. and S. Zielke (2017). "Can't buy me green? A review of consumer perceptions of and behavior toward the price of organic food". In: *Journal of Consumer Affairs* 51.1, pp. 211–251.
- Baker, J., D. Grewal, and A. Parasuraman (1994). "The Influence of Store Environment on Quality Inferences and Store Image". In: *Journal of the Academy of Marketing Science* 22.4, pp. 328–339.
- Bartling, B. and U. Fischbacher (2012). "Shifting the Blame: On Delegation and Responsibility". In: *The Review of Economic Studies* 79.1, pp. 67–87.
- Bartling, B. and Y. Özdemir (2017). "The limits to moral erosion in markets: Social norms and the replacement excuse". In:
- Bartling, B., V. Valero, and R. Weber (2017). "On the scope of externalities in experimental markets". In: *Experimental Economics*, pp. 1–15.
- (2019). "On the scope of externalities in experimental markets". In: *Experimental Economics* 22.3, pp. 610–624.
- Bartling, B., R. A. Weber, and L. Yao (2015). "Do Markets Erode Social Responsibility? *". In: *The Quarterly Journal of Economics* 130.1, pp. 219–266.
- Bénabou, R. and J. Tirole (2011). "Identity, morals, and taboos: Beliefs as assets". In: *The Quarterly Journal of Economics* 126.2, pp. 805–855.
- Bernard, A., J. Daaboul, F. Laroche, and C. Da Cunha (2012). "Mass customisation as a competitive factor for sustainability". In: *Enabling Manufacturing Competitiveness and Economic Sustainability*. Springer, pp. 18–25.
- Camerer, C. F., G. Loewenstein, and M. Rabin (2004). *Advances in behavioral economics*. Princeton university press.
- Cialdini, R. B. (2007). "Descriptive Social Norms as Underappreciated Sources of Social Control". In: *Psychometrika* 72.2, p. 263.
- Coffman, L. C. (2011). "Intermediation Reduces Punishment (and Reward)". In: *American Economic Journal: Microeconomics* 3.4, pp. 77–106.
- Cohen, D. A. and S. H. Babey (2012). "Contextual Influences on Eating Behaviors: Heuristic Processing and Dietary Choices". In: *Obesity reviews : an official journal of the International Association for the Study of Obesity* 13.9, pp. 766–779.

- Dana, J., R. A. Weber, and J. X. Kuang (2007). "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness". In: *Economic Theory* 33.1, pp. 67–80.
- Danz, D., D. Engelmann, and D. Kübler (2012). "Do Legal Standards Affect Ethical Concerns of Consumers? An Experiment on Minimum Wages". Working Paper Series. Mannheim: University of Mannheim.
- Eckhardt, G. M., R. Belk, and T. M. Devinney (2010). "Why don't consumers consume ethically?" In: *Journal of Consumer Behaviour* 9.6, pp. 426–436.
- Ehrich, K. R. and J. R. Irwin (2005). "Willful Ignorance in the Request for Product Attribute Information". In: *Journal of Marketing Research* 42.3, pp. 266–277.
- Einstein, A. (1905). "Zur Elektrodynamik bewegter Körper. (German) [On the electrodynamics of moving bodies]". In: *Annalen der Physik* 322.10, pp. 891–921.
- Etilé, F. and S. Teyssier (2016). "Signaling corporate social responsibility: Third-party certification versus brands". In: *The Scandinavian Journal of Economics* 118.3, pp. 397–432.
- Falk, A., T. Neuber, and N. Szech (2020). "Diffusion of being pivotal and immoral outcomes". In: *The Review of Economic Studies* 87.5, pp. 2205–2229.
- Falk, A. and N. Szech (2013). "Morals and Markets". In: *Science* 340.6133, pp. 707–711.
- Farrow, K., G. Grolleau, and L. Ibanez (2017). "Social Norms and Pro-environmental Behavior: A Review of the Evidence". In: *Ecological Economics* 140, pp. 1–13.
- Fehr, E. and S. Gächter (2002). "Altruistic Punishment in Humans". In: *Nature* 415.6868, pp. 137–140.
- Fischer, P., J. I. Krueger, T. Greitemeyer, C. Vogrincic, A. Kastenmüller, D. Frey, et al. (2011). "The bystander-effect: a meta-analytic review on bystander intervention in dangerous and non-dangerous emergencies." In: *Psychological bulletin* 137.4, p. 517.
- Food and A. O. of the United Nations; World Health Organization (2019). *Sustainable healthy diets: guiding principles*.
- Fourcade, M. and K. Healy (2007). "Moral views of market society". In: *Annual review of Sociology* 33.
- Friedrichsen, J. and D. Engelmann (2018). "Who cares about social image?" In: *European Economic Review* 110, pp. 61–77.
- Garcia, S. M., K. Weaver, G. B. Moskowitz, and J. M. Darley (2002). "Crowded minds: The implicit bystander effect". In: *Journal of Personality and Social Psychology* 83.4, pp. 843–853.

- Golman, R., D. Hagmann, and G. Loewenstein (2017). "Information Avoidance". In: *Journal of Economic Literature* 55.1, pp. 96–135.
- Gray, C. M., Y. Kou, B. Battles, J. Hoggat, and A. L. Toombs (2018). "The Dark (Patterns) Side of UX Design". In: *CHI 2018*. New York, NY: ACM, pp. 1–14.
- Hartmann-Boyce, J., F. Bianchi, C. Piernas, S. P. Riches, K. Frie, R. Nourse, et al. (2018). "Grocery store interventions to change food purchasing behaviors: a systematic review of randomized controlled trials". In: *The American journal of clinical nutrition* 107.6, pp. 1004–1016.
- Hoek, A. C., D. Pearson, S. W. James, M. A. Lawrence, and S. Friel (2017). "Healthy and environmentally sustainable food choices: Consumer responses to point-of-purchase actions". In: *Food Quality and Preference* 58, pp. 94–106.
- Hollands, G. J., I. Shemilt, T. M. Marteau, S. A. Jebb, M. P. Kelly, R. Nakamura, et al. (2013). "Altering micro-environments to change population health behaviour: towards an evidence base for choice architecture interventions". In: *BMC public health* 13, p. 1218.
- Hummel, D. and A. Maedche (2019). "How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies". In: *Journal of Behavioral and Experimental Economics* 80, pp. 47–58.
- Huyghe, E., J. Verstraeten, M. Geuens, and A. van Kerckhove (2017). "Clicks as a Healthy Alternative to Bricks: How Online Grocery Shopping Reduces Vice Purchases". In: *Journal of Marketing Research* 54.1, pp. 61–74.
- Irlenbusch, B. and D. Saxler (2015). "Social responsibility in market interaction". In: Irlenbusch, B. and D. J. Saxler (2019). "The role of social information, market framing, and diffusion of responsibility as determinants of socially responsible behavior". In: *Journal of Behavioral and Experimental Economics* 80, pp. 141–161.
- Johnson, N. D. and A. A. Mislin (2011). "Trust games: A meta-analysis". In: *Journal of economic psychology* 32.5, pp. 865–889.
- Just, D. R. and A. T. Byrne (2020). "Evidence-based policy and food consumer behaviour: how empirical challenges shape the evidence". In: *European Review of Agricultural Economics* 47.1, pp. 348–370.
- Just, D. R. and G. Gabrielyan (2018). "Influencing the food choices of SNAP consumers: Lessons from economics, psychology and marketing". In: *Food Policy* 79, pp. 309–317.
- Kalnikaitė, V., J. Bird, and Y. Rogers (2013). "Decision-making in the aisles: informing, overwhelming or nudging supermarket shoppers?" In: *Personal and Ubiquitous Computing* 17.6, pp. 1247–1259.

- Karar, A., S. Said, T. Beyrouthy, et al. (2019). "Pepper humanoid robot as a service robot: a customer approach". In: *2019 3rd International Conference on Bio-engineering for Smart Technologies (BioSMART)*. IEEE, pp. 1–4.
- Kirchler, M., J. Huber, M. Stefan, and M. Sutter (2016). "Market Design and Moral Behavior". In: *Management Science* 62.9, pp. 2615–2625.
- Kozyreva, A., S. Lewandowsky, and R. Hertwig (in press). "Citizens versus the internet: Confronting digital challenges with cognitive tools". In: *Psychological Science in the Public Interest*.
- Kraft, T., L. Valdés, and Y. Zheng (2016). "Transparency in social responsibility: Investigating consumers' behaviors and motives". In: *Social Science Research Network, Rochester, NY*.
- Lembcke, T.-B., N. Engelbrecht, A. B. Brendel, B. Herrenkind, and L. M. Kolbe (2019). "Towards a Unified Understanding of Digital Nudging by Addressing its Analog Roots". In: *Twenty-Third Pacific Asia Conference on Information Systems*, p. 123.
- Lembcke, T.-B., N. Engelbrecht, A. B. Brendel, and L. Kolbe (2019). "To Nudge or not to Nudge: Ethical Considerations of Digital Nudging Based on Its Behavioral Economics Roots". In: *In Proceedings of the 27th European Conference on Information Systems (ECIS), Stockholm & Uppsala, Sweden, June 8-14, 2019*.
- Lorek, S. and P. J. Vergragt (2015). "Sustainable consumption as a systemic challenge: inter-and transdisciplinary research and research questions". In: *Handbook of research on sustainable consumption*. Edward Elgar Publishing.
- Lu, B., W. Fan, and M. Zhou (2016). "Social presence, trust, and social commerce purchase intention: An empirical research". In: *Computers in Human Behavior* 56, pp. 225–237.
- Lunn, P. (2014). *Regulatory policy and behavioural economics*. Paris: OECD Publishing.
- Nordfalt, J. (2009). "Unplanned grocery purchases: the influence of the shopping-trip type revisited". In: *Journal of Consumer Behaviour* 8.1, pp. 1–13.
- Öberseder, M., B. B. Schlegelmilch, and V. Gruber (2011). "'Why don't consumers care about CSR?': A qualitative study exploring the role of CSR in consumption decisions". In: *Journal of business ethics* 104.4, pp. 449–460.
- Ogonowski, A., A. Montandon, E. Botha, and M. Reyneke (2014). "Should new online stores invest in social presence elements? The effect of social presence on initial trust formation". In: *Journal of Retailing and Consumer Services* 21.4, pp. 482–491.

- Parente, R. (2020). "Digitalization, Consumer Social Responsibility, and Humane Entrepreneurship: Good news from the future?" In: *Journal of the International Council for Small Business* 1.1, pp. 56–63.
- Figors, M. and B. Rockenbach (2016). "Consumer Social Responsibility". In: *Management Science* 62.11, pp. 3123–3137.
- Pitts J., S. B., S. W. Ng, J. L. Blitstein, A. Gustafson, and M. Niculescu (2018). "Online grocery shopping: promise and pitfalls for healthier food and beverage purchases". In: *Public health nutrition*, pp. 1–17.
- Pornpitakpan, C. and J. H. Han (2013). "The effect of culture and salespersons' retail service quality on impulse buying". In: *Australasian Marketing Journal (AMJ)* 21.2, pp. 85–93.
- Pozzi, A. (2012). "Shopping Cost and Brand Exploration in Online Grocery". In: *American Economic Journal: Microeconomics* 4.3, pp. 96–120.
- Reczek, R. W., J. R. Irwin, D. M. Zane, and K. R. Ehrich (2018). "That's Not How I Remember It: Willfully Ignorant Memory for Ethical Product Attribute Information". In: *Journal of Consumer Research* 45.1, pp. 185–207.
- Reisch, L. A. and J. Sandrini (2015). *Nudging in der Verbraucherpolitik: Ansätze verhaltensbasierter Regulierung*. 1. Auflage. Vol. 36. Schriftenreihe des Instituts für Europäisches Wirtschafts- und Verbraucherrecht e.V. Baden-Baden: Nomos.
- Schaer, A. and K. Stanoevska-Slabeva (2019). "Application of Digital Nudging in Customer Journeys – A Systematic Literature Review". In: *Twenty-fifth Americas Conference on Information Systems*. Cancun, Mexiko.
- Schlaile, M. P., K. Klein, and W. Böck (2018). "From bounded morality to consumer social responsibility: A transdisciplinary approach to socially responsible consumption and its obstacles". In: *Journal of Business Ethics* 149.3, pp. 561–588.
- Schubert, C. (2017). "Exploring the (behavioural) political economy of nudging". In: *Journal of Institutional Economics* 13.03, pp. 499–522.
- SDG, U. (2018). "Sustainable development goals". In: *United Nations*.
- Serra-Garcia, M. and N. Szech (2019). "The (in) elasticity of moral ignorance". In:
- Shamir, R. (2008). "The age of responsabilization: On market-embedded morality". In: *Economy and society* 37.1, pp. 1–19.
- Shleifer, A. (2004). "Does competition destroy ethical behavior?" In: *American Economic Review* 94.2, pp. 414–418.

- Smetana, S. M., S. Bornkessel, and V. Heinz (2019). "A path from sustainable nutrition to nutritional sustainability of complex food systems". In: *Frontiers in nutrition* 6, p. 39.
- Tanner, R. J., R. Ferraro, T. L. Chartrand, J. R. Bettman, and R. van Baaren (2008). "Of Chameleons and Consumption: The Impact of Mimicry on Choice and Preferences". In: *Journal of Consumer Research* 34.6, pp. 754–766.
- Terlau, W. and D. Hirsch (2015). "Sustainable Consumption and the Attitude-Behaviour-Gap Phenomenon - Causes and Measurements towards a Sustainable Development: 159 - 174 Pages / International Journal on Food System Dynamics, Vol 6, No 3 (2015)". In:
- Thaler, R. H. (2016). "Behavioral Economics: Past, Present, and Future". In: *American Economic Review* 106.7, pp. 1577–1600.
- Thaler, R. H. and C. R. Sunstein (2008). *Nudge: Improving Decisions About Health, Wealth and Happiness*. Illustrated ed. New Haven: Yale University Press.
- Thunström, L., J. Nordström, J. F. Shogren, M. Ehmke, and K. van't Veld (2016). "Strategic self-ignorance". In: *Journal of Risk and Uncertainty* 52.2, pp. 117–136.
- Turley, L. and R. E. Milliman (2000). "Atmospheric Effects on Shopping Behavior". In: *Journal of Business Research* 49.2, pp. 193–211.
- Tversky, A. and D. Kahneman (1974). "Judgment under Uncertainty: Heuristics and Biases". In: *Science* 185.4157, pp. 1124–1131.
- Uhrich, S. and M. Luck (2012). "Not too many but also not too few". In: *Qualitative Market Research: An International Journal* 15.3, pp. 290–308.
- van Ittersum, K., J. M. Pennings, and B. Wansink (2010). "Trying Harder and Doing Worse: How Grocery Shoppers Track In-Store Spending". In: *Journal of Marketing* 74.2, pp. 90–104.
- van Ittersum, K., B. Wansink, J. M. Pennings, and D. Sheehan (2013). "Smart Shopping Carts: How Real-Time Feedback Influences Spending". In: *Journal of Marketing* 77.6, pp. 21–36.
- Vannucci, V. and E. Pantano (2019). "Digital or human touchpoints? Insights from consumer-facing in-store services". In: *Information Technology & People*.
- Vermeir, I. and W. Verbeke (2006). "Sustainable Food Consumption: Exploring the Consumer "Attitude – Behavioral Intention" Gap". In: *Journal of Agricultural and Environmental Ethics* 19.2, pp. 169–194.

- Vermeir, I., B. Weijters, J. De Houwer, M. Geuens, H. Slabbinck, A. Spruyt, et al. (2020). "Environmentally sustainable food consumption: a review and research agenda from a goal-directed perspective". In: *Frontiers in psychology* 11, p. 1603.
- Vitell, S. J. (2015). "A case for consumer social responsibility (CnSR): Including a selected review of consumer ethics/social responsibility research". In: *Journal of Business Ethics* 130.4, pp. 767–774.
- Willett, W., J. Rockström, B. Loken, M. Springmann, T. Lang, S. Vermeulen, et al. (2019). "Food in the Anthropocene: the EAT–Lancet Commission on healthy diets from sustainable food systems". In: *The Lancet*.
- Zane, D. M., J. R. Irwin, and R. W. Reczek (2016). "Do less ethical consumers denigrate more ethical consumers? The effect of willful ignorance on judgments of others". In: *Journal of Consumer Psychology* 26.3, pp. 337–349.

CHAPTER II

The Interplay of Different Excuses in Determining Moral Wiggle Room

Mattheus Brenig, Nils Engelbrecht, and Daniel Hermann

In most market decisions, the relationship between one's actions and the consequences for others is subject to uncertainty. Previous research suggests that this creates moral wiggle room and allows individuals to excuse selfish behavior, leading to the claim that markets erode morality. We extend the literature by examining how sensitive selfish behavior, and thus moral wiggle room, is to changes in uncertainty that affect individuals' ability to use different excuses simultaneously, as well as the plausibility with which an excuse can be used. We vary the extent to which the probability of the externality is affected by individuals' choices, namely how pivotal they are. We also vary whether an individual can apply either the bystander excuse (externality may be prevented despite a selfish decision), the replacement excuse (externality may occur despite an unselfish decision), or both. In addition, we limit the scope of uncertainty to a portion of the externality and introduce the possibility of intentional uncertainty through ignorance. We find that introducing uncertainty increases selfish behavior by approximately one-fourth. This result is robust to our experimental variations in pivotality, the scope of uncertainty, and the possibility of invoking the bystander and/or replacement excuses. Once pivotality is reduced to zero, selfish behavior further increases. Surprisingly, two-thirds of subjects disclose information about their influence on the externality, leading to no significant effect when allowing for willful ignorance.

Keywords Diffusion of responsibility · Replacement logic · Bystander intervention · Moral wiggle room

1 Introduction

In his seminal work on moral disengagement, Bandura (1999) describes several psychosocial mechanisms that decision-makers use to justify harmful effects of their actions on others. These mechanisms include denial of personal agency through diffusion of responsibility, as well as consciously ignoring effects on others. Recent evidence on moral disengagement comes from an emerging strand of experimental economic research suggesting that decision-makers are more selfish once the relationship between their actions and others' consequences is blurred through e.g. (i) delegating tasks to agents (Hamman et al. 2010; Bartling and Fischbacher 2012), (ii) responsibility diffusion when outcomes depend on the interaction of multiple persons (Brütt et al. 2020; Falk et al. 2020; Ziegler et al. 2021; Falk and Szech 2013), and (iii) willful ignorance (Bartling et al. 2014; Grossman and Weele 2016; Reczek et al. 2018). All of these examples introduce ex-ante uncertainty about the consequences of one's actions on others. This allows decision-makers to excuse selfish behavior and avoid attribution of blame and responsibility not only by themselves but also by others (e.g. Bartling and Fischbacher 2012; Coffman 2011). Dana et al. (2007) have fittingly described such behavior as exploiting "moral wiggle room".

While the existing literature has examined how the ability to use single excuses influences selfish behavior, the question of how individuals behave when they have the option to resort to multiple excuses remains unanswered. In other words, does adding or subtracting excuses lead to a corresponding change in the moral wiggle room? This is particularly relevant in market transactions, which can be as diverse as they are complex. Markets are characterized by different sizes, institutional rules, competitive structures or levels of transparency. Hence, compared to other more stylized decisions, individuals might be able to use different excuses to justify harmful effects of their actions on others. For example, some externalities may be directly attributable to a single act of consumption, while responsibility for other externalities may be diffused among both consumers and producers. This diffusion allows participants to justify selfish behavior with an uncertain impact on the outcome by using the replacement excuse ("if I don't create the harmful effect on others, someone else will") and/or bystander excuse ("someone else will avert harmful effects on others"). Moreover, responsibility diffusion may make it easier for market participants to justify self-induced uncertainty by ignoring relevant information (Stoll-Kleemann et al. 2001; Norgaard 2006). In increasingly fragmented and often quite obscure supply chains,

market participants have ample opportunity for willful ignorance. Thus, how various excuses interact is not trivial and might add a new dimension to the discussion about morals and markets. Depending on how excuse-driven behavior works, adding more excuses may or may not result in increasing selfishness. Such findings would have important policy implications. For example, the disclosure of information about harmful effects may reduce selfish behavior when responsibility can be attributed to a single person, but may be ignored when responsibility is diffused.

To address this, the purpose of this study is to experimentally examine how different excuses emanating from ex-ante uncertainty about the consequences of one's actions on third parties affect behavior. We focus on excuses that are relevant in market transactions, and we are particularly interested in how sensitive individuals' behavior is to the possibility of using multiple excuses for selfish behavior. We investigate the interplay of four factors:

First, we consider the interplay of the bystander and replacement excuse. The latter can be applied if it is uncertain whether an unselfish decision prevents the externality. For example, a market transaction may cause some externality with certainty, but can be justified by using the excuse "if I don't make the transactions, someone else will" (Falk et al. 2020; Ziegler et al. 2021). On the other hand, the bystander excuse can be applied if it is uncertain whether a selfish decision results in the externality. For example, a bystander can help a victim and prevent harm with certainty, but may rely on other bystanders to respond to the emergency (Darley and Latane 1968; Fischer et al. 2011). In the examples above, only one of the two choice options leads to uncertainty. However, in most decisions, both the selfish and unselfish choice have uncertain consequences (e.g. regarding a consumer's influence on future market outcomes). As a result, decision-makers can combine both the replacement and bystander excuse.

Second, the extent to which individuals believe that they influence the likelihood of the externality probably also matters. The more bystanders that are present at an emergency or the more participants in a market, the less pivotal the individual may feel that they are to the outcome. This can also be interpreted as the plausibility of the excuse. For example, the excuse "someone else will help the victim" is certainly more plausible or convincing if there are many bystanders than if there is only one other bystander. More technically, the extent to which an individual is pivotal (ex-ante) can be described by how much the probability of externality is affected by the individual's decision (Engl 2018).

Third, in market transactions, some externalities may be directly attributable to an individual consumer, while responsibility for other externalities may be diffused among market participants. Accordingly, the scope of uncertainty may be limited to some of the externalities, while consumers are fully pivotal for the remaining externalities, limiting their ability to use excuses and thus reducing selfish behavior. This can similarly be interpreted as the plausibility of the excuse.

Fourth, we add the possibility of willfully-induced uncertainty through ignorance. In introducing strategic ignorance, we closely mirror how consumers in markets retrieve information, e.g., via product labels, just before making a purchase.

Our study is closely related to economic experiments in which unanimous group decisions are required to either cause or prevent an externality. In Falk et al. (2020) and Brütt et al. (2020), groups need a unanimous vote to prevent a negative externality. In Hauser et al. (2014), a single group member can exploit a common-pool resource to the point that it is not replenished for future generations. Since a single selfish vote or choice is sufficient to cause the negative externality, the ex-ante probability of each individual being pivotal is reduced. Similarly, in the multi-unit market treatment in Ziegler et al. (2021), one buyer-seller pair can trade all available units in a market, leading to the largest possible negative externality. In all of the above-mentioned studies, subjects can use the replacement excuse to justify selfish behavior. Other studies have investigated unanimity voting in favor of a negative externality (Dana et al. 2007; Irlenbusch and Saxler 2019; Behnk et al. 2017; Brütt et al. 2020). Here, subjects can apply the bystander excuse.

Our study extends and contributes to this literature in various ways. First and foremost, our study examines the interplay of several excuses that are relevant in market transactions and have previously only been studied in isolation. We compare and combine both the bystander and replacement excuse, different degrees of being pivotal, investigate the influence of the scope of uncertainty, and add the possibility of strategic ignorance. This allows us to draw conclusions about whether selfish behavior depends on the possibility of using different excuses, combinations thereof and their plausibility. Second, we replace uncertainty through the behavior of other decision-makers with moves of nature. This allows us to isolate the pure effect of uncertainty, while holding constant other dimensions of responsibility diffusion, like shared guilt (Rothenhäusler et al. 2018) or information on social norms. In group decisions, perceptions of uncertainty depend on subjective beliefs about the morality of other group members and are thus imbued with social information. Third, and related

to the second point, the use of moves of nature allows us to manipulate uncertainty in a more controlled manner. By contrast, in group decisions recorded for previous studies, perceptions of uncertainty depend on individuals' endogenous beliefs and are thus ambiguous. The formation of beliefs is difficult to control and cannot be precisely manipulated by experimental treatments. While interesting areas of research in their own right, the formation of beliefs about pivotality and the role of ambiguity are not central to perceptions of pivotality or the bystander and replacement excuse.

2 Experimental Design

Setting

Subjects participate in groups of five in an ultimatum bargaining situation in which there are three roles: proposers (A), responders (B), and inactive third parties (C). In each group, there are two proposers A_1 and A_2 , two inactive third parties C_1 and C_2 , and one responder B . Each proposer is assigned one inactive third party. The game is played for two rounds, both of which are relevant for payoffs. The group composition is fixed over both rounds. In each round, all subjects first receive an endowment of $a = b = c = 100$ coins.¹ As explained in more detail below, proposers and responders have varying degrees of influence over the final distribution of coins in each round.

In round one, each proposer is first given a binary choice between offering the responder a distribution of either $(a_i = 110, b = 120, c_i = 70)$ or $(a_i = 110, b = 110, c_i = 100)$ coins. For both possible offers, the sum of the payments that the proposer and responder receive is 30 or 20 coins higher than the sum of their initial endowments of 100 coins each. However, one distribution does not affect the inactive third party's endowment, while the other entails a loss of 30 coins. Throughout this paper, we will refer to the first and second offer as selfish and unselfish, respectively. In the instructions, we used neutral terms (A and B) and the colors blue and red to differentiate the offers.² Matching of proposers and responders only takes place after the first stage of round one is implemented, and it is not random. Instead, groups are matched in such a way that each responder can choose between the unselfish and

¹Throughout the study, we used the term "coins" for our experimental currency.

²We randomized both colors to control for possible color preferences.

selfish offer.³ To control for possible positional preferences on the computer screen (i.e. left and right), we randomized the position of the two offers. The responder has to accept one of the two offers. Thereby, the final allocation of coins to all group members is determined. The responder and the “winning” proposer receive the coins associated with the accepted offer. The “losing” proposer keeps the initial endowment of 100 coins. The determination of the payoffs for the inactive third parties depends on the experimental condition, as explained below.

The process of round two is identical to round one, except that offers are no longer chosen by proposers. Instead, offers depend on a move by nature in which the probabilities of future offers depend on the choice of the responder. The parameters of the move by nature depend on the experimental condition, as explained below. Independent of the specific parameter values, the move by nature leads to either two selfish or unselfish offers.

Experimental Conditions

Move by Nature

While offers in round one are chosen by proposers, the offers in round two are determined by a move by nature. The move leads to one of two possible “market structures” in round two: a market where both proposers make the selfish or unselfish offer. We call the former the selfish and the latter the unselfish market. By varying how the responder’s decision affects the probability of the selfish and unselfish market occurring in round two, we determine the responder’s degree of being pivotal and whether the replacement and/or bystander excuse can be applied. To explain the logic of our experimental conditions, we speak of two properties that determine the move by nature. The first property α (with $0 \leq \alpha \leq 1$) is the pivotality of the responder, i.e. the extent to which the probability of a (un)selfish market increases with a (un)selfish choice. The second property β (with $0 \leq \beta \leq 1$) is the inherent probability of the selfish market, which is independent of the subject’s choice. Of course, $\alpha + \beta \leq 1$.

If the responder is fully pivotal ($\alpha = 1, \beta = 0$), a selfish choice leads to a selfish market with certainty and vice versa. Consequently, there is no moral wiggle room and no excuse can be applied. By contrast, if the responder is not pivotal at all ($\alpha = 0, 0 \leq \beta \leq 1$), their decision has no influence on the move of nature. In this

³Responders are not informed about this detail of the matching process. However, it is important to emphasize that no deception is involved.

case, only responders following a deontological moral reasoning would still refrain from choosing the selfish option. For any $0 < \alpha < 1$, varying β affects whether the responder can apply the replacement and/or bystander excuse. First, if $\beta = 0$, an unselfish choice leads to an unselfish market with certainty, but there is uncertainty concerning whether a selfish choice results in the selfish market (bystander excuse). Second, if $\alpha + \beta = 1$, a selfish choice leads to a selfish market with certainty, but there is uncertainty concerning whether an unselfish decision prevents the unselfish market (replacement excuse). Third, if $\alpha + \beta < 1$, both choice options are subject to uncertainty, which is why both excuses can be used.

Rounds one and two in extensive form are represented by Figure 1.⁴ In reading the game tree, remember that groups are matched so that both the unselfish and selfish offer are available to the responder in round one. Hence, if a responder accepts a selfish offer by e.g. A_1 , an unselfish offer by A_2 is rejected and vice versa. Due to this symmetry, only half of the game tree is shown in detail.

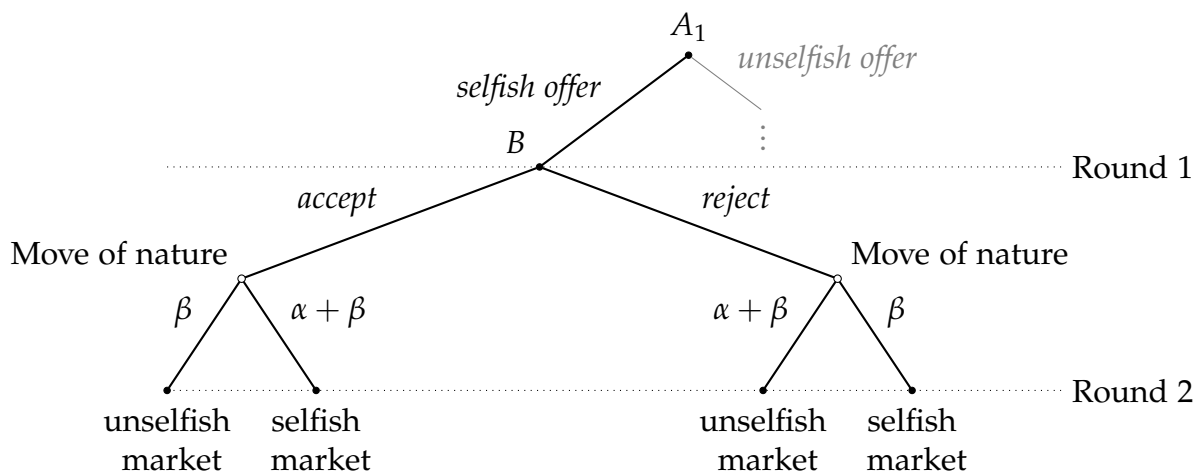


Figure 1 — Game tree over both rounds

Scope of uncertainty

Proposers' and responders' payoffs are always determined once the responder accepts an offer. On the other hand, the timing of the payoff consequences for third parties depends on the experimental condition. This applies to both rounds regardless of

⁴Payoffs are not shown and can be deduced from Figure 2. Instead, the game tree is intended to show which path of the decision tree leads to which market structure.

whether the proposer has actively selected the offer (round one) or whether the offer for the proposer was determined by the market rule (round two). Analogous to the Made-to-Stock (*MTS*) and Made-to-Order (*MTO*) manufacturing processes, pivotality for externalities of offers that have already been made is implemented as a binary variable, i.e., responders are either pivotal or not.⁵ In *MTO*, a negative externality arises by consumption. In *MTS*, all of the injustices associated with the production process have already occurred once a product is offered to consumers. Experimental studies on socially responsible behavior in markets have applied either *MTO* or *MTS* (e.g., *MTO* in Bartling et al. 2015 and *MTS* in Pigors and Rockenbach 2016). In *MTS*, a third-party payoff is determined as soon as the assigned proposer chooses the offer. In *MTO*, all payoffs are determined once the responder accepts an offer. Hence, while responders cannot prevent the externality by rejecting the selfish offer in *MTS*, they can do so in *MTO*. Round one in extensive form is represented by Figure 2 for *MTO* and *MTS*. The payoffs associated with the “winning” offer are highlighted by a box.

Importantly, *MTO* splits the avoidable externality over both rounds. If a responder accepts the selfish offer in round one, 30 coins of externality are created. When the selfish market occurs in round two, two selfish offers are made, but only 30 coins of externality are created by consumption. By contrast, all avoidable externalities in *MTS* are realized in round two. When round two is the selfish market, two selfish offers are made, resulting in 60 coins of externality regardless of the responder choice in round two. Because the move by nature only applies to offers in round two, the scope of uncertainty is reduced by half in *MTO*. Because this experimental condition depends on subjects understanding the consequences of their choices on payoffs, we included three comprehension questions that asked about all relevant payoffs for a selfish choice in round one and the payoffs for both round two scenarios (see Online Appendix A.5 for all comprehension questions).

Strategic Ignorance

In the *Full information* condition, subjects are informed about the payoff consequences for all roles in the instructions and their understanding of these consequences is tested

⁵Traditionally, goods are produced ahead, i.e. made-to-stock (*MTS*). However, due to new production technologies and more differentiated consumer demand, goods are increasingly made-to-order (*MTO*). While *MTO* was predominantly used for complex and highly customized products such as industrial machines or cars, it is now applied to mass products such as clothes or even groceries (Gilmore and Pine 1997; Holweg and Pil 2001).

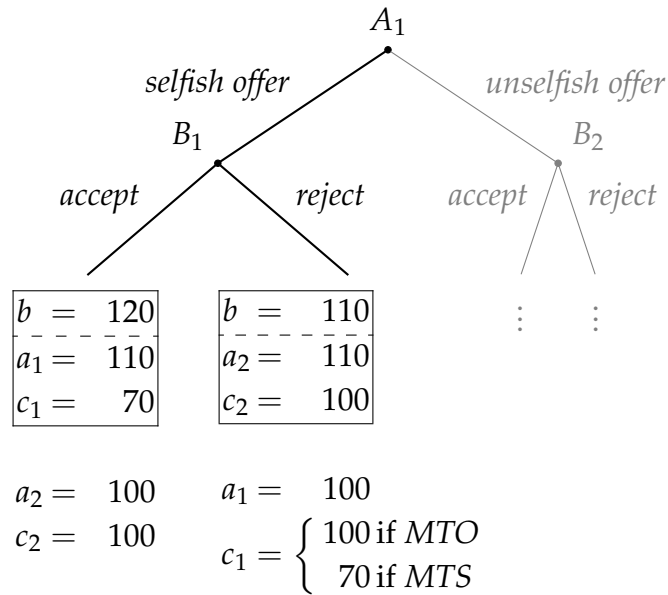


Figure 2 — Game tree for round 1

with several comprehension questions. In the *Strategic ignorance* condition, payoff consequences for third parties are hidden from the subjects in the instructions. Accordingly, subjects do not know the direction or magnitude of the payoff consequences. However, subjects are informed that they have the opportunity to reveal the payoff consequences before making a decision. To reveal the information, subjects only have to click a button on the screen, i.e. there are no transaction costs.

Treatments and Behavioral Hypotheses

Table 1 provides an overview of the treatments. In the first step of our analysis, we investigate behavior in *MTS*, i.e. when the externality of 60 coins is realized in round two and is therefore subject to uncertainty in its entirety.

We hypothesize that selfishness increases in the order of treatments as shown in the first section of Table 1, i.e. rows one to five. In our *Baseline* condition, subjects are maximally pivotal, i.e. a selfish choice leads to the externality with certainty and vice versa. In this case, there is no moral wiggle room. In our *Bystander* condition, the intrinsic probability corresponds to zero and pivotality to fifty percentage points. In this case, subjects can prevent the externality. However, they can excuse selfish behavior on the grounds that there is a chance that the externality will be prevented

Table 1 — Treatments

Name	Experimental setup			
	Pivotality α	Inherent probability β	MTO/MTS	Strategic ignorance
<i>Baseline</i>	1	0	MTS	No
<i>Bystander</i>	0.5	0	MTS	No
<i>Replacement</i>	0.5	0.5	MTS	No
<i>Compound</i>	0.1	0.45	MTS	No
<i>Deontological</i>	0	0.5	MTS	No
<i>CompoundMTO</i>	0.1	0.45	MTO	No
<i>BaselineMTO</i>	1	0	MTO	No
<i>Compound*</i>	0.1	0.45	MTS	Yes
<i>CompoundMTO*</i>	0.1	0.45	MTO	Yes

either way. In our *Replacement* condition, we set the inherent probability of the externality and pivotality to fifty percentage points. A selfish choice leads to the externality with certainty, while there is a fifty percent probability of the externality even if the subject is not selfish. Although pivotality is still relatively high, subjects may use the uncertainty regarding their ability to prevent the externality as an excuse for selfishness. Pivotality is equally strong in *Bystander* and *Replacement*, with the only difference being the inherent probability of the externality. This is analogous to unanimity voting in favor (*Bystander*) or against (*Replacement*) the externality. In our *Compound* condition, we set the inherent probability to forty-five and pivotality to ten percentage points. In *Compound*, pivotality is substantially reduced and subjects can apply both the bystander and replacement excuse. Accordingly, we expect increased selfishness compared to *Bystander* and *Replacement*. Finally, in *Deontological*, the inherent probability is fifty percent and subjects have no influence on the market at all. In this scenario, only individuals following a deontological moral reasoning would still refrain from choosing the selfish option (Falk et al. 2020).

Hypothesis 1 *Selfishness is lowest in Baseline and largest in Deontological. The treatments Bystander, Replacement, and Compound lie in-between, with the latter exhibiting relatively strong selfish behavior.*

In step two (see rows 6 and 7 in Table 1), we extend our main analysis by investigating behavior in *MTO*. In this choice setting, subjects always generate half of the externality,

i.e. 30 coins, by choosing the selfish offer in round one. If round two is a selfish market, another 30 coins of externality arise by consumption. Essentially, the move by nature now only applies to half of the externality, while subjects can prevent the other half with certainty. Accordingly, subjects can never justify that the ex-post outcome will be the same, independent of their behavior. We chose the *Compound* condition for the *MTO* setting because it generates the largest difference between *MTO* and *MTS* in terms of expected externality, while allowing for the largest number of excuses due to low pivotality as well as the bystander and replacement excuse. We call this treatment *CompoundMTO*.

Hypothesis 2 *Selfish behavior is lower in CompoundMTO compared to Compound.*

Furthermore, as a control, we also run a *BaselineMTO* treatment with full pivotality. For the purpose of our study, we are only interested in the differences between *MTO* and *MTS* in terms of the scope of uncertainty. However, there may be a concern that the two conditions introduce other differences that are relevant to behavior. For example, simply by framing the timing of the externality differently, subjects may feel like second movers in *MTS*, and as a result they may feel less responsible for the externality. At full pivotality, there is no uncertainty and thus no difference between *MTO* and *MTS* on the dimension of interest for our study. In both treatments, subjects have full pivotality for 60 coins of externality. Comparing *Baseline* and *BaselineMTO* thus serves as a control for potential confounds due to unintended changes.

In a third and final step (see rows 8 and 9 in Table 1), we allow for strategic ignorance in both *CompoundMTO* and *Compound*. *Ceteris paribus*, giving subjects the opportunity to remain strategically ignorant about the consequences of their decisions on third parties should increase selfishness.

Hypothesis 3 *Selfishness is higher in CompoundMTO* and Compound*, compared to CompoundMTO and Compound, respectively.*

Experimental Procedure

The experiment was programmed with oTree (Chen et al. 2016) and administered to subjects via MTurk between December 2020 and February 2021. We prevented retakes and restricted participation to US residents.⁶ Since we were only interested in the

⁶In addition, due to readability issues, we excluded participants who reported accessing the study via smartphone or tablet.

responder decisions and needed to ensure that both offers were made in each group, we collected the proposer and third-party observations beforehand. We applied a many-to-one matching, collecting at least one selfish and unselfish proposer and two third-party observations per treatment. These participants received a fixed payment immediately, and the rest of their payment after all responder observations had been collected.

In total, 1,130 subjects participated in the role of responder. After reading a plain language statement and giving consent to participate, subjects read the instructions (see Online Appendix A.4) and had a maximum of two attempts to answer a few comprehension questions correctly. Subjects who failed twice were excluded from participation. We performed two additional tests to check the validity of our experimental design and subjects' understanding of the experiment and instructions (see Online Appendix A.2). After completion of the main task, subjects completed a post-experimental questionnaire. Following our pre-registration and based on arrival time, we excluded observations in excess of 125 observations per treatment. Our final sample thus comprises 1,125 observations. About 49% of subjects were female, and participants were on average 36 years old. For more information on participants' socio-demographics, see Online Appendix A.1. Participants took an average of 21 minutes, with a mean payment of \$3.05, which equals \$13.30 per hour.⁷

3 Results

Main Results

Figure 3 shows the share of selfish choices per treatment in *MTS* and *Full information*. As expected, subjects are least selfish in *Baseline* (48.8%), with deterministic influence on second-round offers. Once uncertainty is introduced, selfish behavior increases by about ten percentage points to 58.4%, 59.2%, and 60.0%. This increase is on the borderline of significance in a pairwise comparison (chi²-test, $p = 0.13, 0.10, \text{ and } 0.08$). There is no significant difference (chi²-test, $p = 0.967$) among the three treatments with uncertainty. Accordingly, it does not seem to matter for selfish behavior whether

⁷We pre-registered the experiment at <https://aspredicted.org/blind.php?x=be7c8c>. The complete data, do-files containing all of the commands we execute for the analysis, oTree app as well as additional information about the data structure can be accessed via the following repository: https://osf.io/md7wu/?view_only=02fe724834be413d8ef84370adb20f97.

uncertainty allows subjects to use the replacement and/or bystander excuse. Even more surprisingly, reducing pivotality from 50 % to 10 % does not significantly increase selfishness either. However, subjects differentiate between uncertain causal responsibility and no causality at all. Once responders have no influence on the market in *Deontological*, selfishness increases to 72 %. This increase is significant with respect to all other treatments in a pairwise comparison (χ^2 -test, $p = 0.05, 0.03, 0.02$, and 0.00). The treatment *Deontological* also reveals that 28 % of responders still reject the selfish offer, which may be due to deontological reasons or to punish the proposer.




Figure 3 — Share of selfish choices in treatments

Result 1 *Introducing uncertainty about a responder’s influence on future market outcomes increases selfish behavior.*

Result 2 *Subjects behavior seems to be inelastic to changes in pivotality and different combinations of the replacement and bystander excuse.*

Result 3 *Selfish behavior is highest in Deontological. Nonetheless, 28 % of responders still reject the selfish offer.*

Similar to the different excuses relating to uncertainty, reducing its scope from *Compound* to *CompoundMTO* does not influence behavior significantly. Despite the possibility to prevent half of the externality with certainty and therefore considerable differences in aggregate pivotality, selfish behavior only slightly decreases to 58.4 %, i.e. by 1.6 percentage points. Concordantly, a Pearson’s χ^2 test is not significant ($p = 0.8$). A comparison of *Baseline* (48.8 %) and *BaselineMTO* (47.2 %) shows that there are no other confounding differences between *MTS* and *MTO*.

Result 4 *Reducing the scope of uncertainty by half does not significantly affect selfish behavior.*

Allowing responders to remain ignorant about the consequences of their decisions on third parties does not significantly affect selfish behavior in our experiment. While selfishness increases in *MTS* as expected (from 60.0 % to 63.2 %), there is a comparable decrease in *MTO* (from 58.4 % to 54.4 %). About two-third of subjects chose to reveal the payoff consequences for third parties in both *MTO* (68.8 %) and *MTS* (61.6 %). This difference is not statistically significant ($\chi^2 = 1.43, p = 0.23$). Subjects who actively disclose the payoff information are less likely to choose the selfish offer than subjects who are (passively) informed in the instructions. In *MTO*, the difference is 16.54 percentage points (41.86 % vs. 58.40 %; $\chi^2 = 5.58, p = 0.02$), while in *MTS* the difference is 10.65 percentage points (49.35 % vs. 60.00 %; $\chi^2 = 2.19, p = 0.14$).

Result 5 *Allowing responders to remain ignorant about the consequences of their decision on third parties does not significantly affect selfishness.*

Result 6 *About two-thirds of subjects reveal information on the consequences of their actions on third parties.*

Individual Heterogeneity

Table 2 shows the results of a logistic regression including all observations. Model 1 includes the treatments as explanatory variables and a dummy variable that equals one for subjects revealing the payoff information in the *Strategic ignorance* condition. Controlling for those who actively revealed the information, *Strategic ignorance* leads to a significant increase in selfish behavior. Unsurprisingly, subjects who remain ignorant predominantly choose the offer that maximizes their payoff.

Model 2 adds several socio-demographic variables from the post-experimental questionnaire. Overall, the odds ratios and standard errors of the treatments remain fairly stable. Regarding employment status, those who are self-employed and not working but looking for a job are significantly less selfish than those who are employed (reference group). Regarding education, those holding an associate degree are significantly less selfish than those holding a bachelor's degree (reference group). Ethnicity, age, and gender all have no significant effect on selfishness. We also asked subjects about their position on the political spectrum on a nine-item scale ranging from 1 (left-wing) to 9 (right-wing). The average marginal effect of a one-point increase on the political spectrum scale is 1.24 percentage points. Furthermore, we asked subjects a few questions regarding their income situation (see verbatim questions and response scale

below Table 2). Subjects who reported that the COVID-19 pandemic has improved their financial situation are significantly less likely to choose selfishly.

As mentioned in the description of our experimental design, we randomized both the color and screen position of the two offers. While the latter had no effect, subjects were significantly less likely to choose the selfish offer when it was visually highlighted by the color blue. The average marginal effect of using the color blue is -5.71 percentage points.

Table 2 — Logistic regression

	Model 1		Model 2	
	Odds ratio	se	Odds ratio	se
<i>Experimental conditions</i>				
Baseline	Reference group		Reference group	
Bystander	1.473	0.375	1.529	0.403
Replacement	1.522*	0.389	1.534	0.404
Compound	1.574*	0.402	1.580*	0.416
Deontological	2.698***	0.722	2.429***	0.669
CompoundMTO	1.473	0.375	1.412	0.371
BaselineMTO	0.938	0.237	0.866	0.227
CompoundMTO*	4.695***	1.737	4.800***	1.834
Compound*	6.275***	2.318	6.370***	2.433
<i>Decisions</i>				
Revealed Payments	0.162***	0.054	0.149***	0.051
<i>Employment status</i>				
Working (paid employee)	Reference group			
Working (self-employed)	0.577***		0.111	
Not working (temporary layoff)	1.174		0.530	
Not working (looking for work)	0.628**		0.147	
Not working (retired)	0.684		0.330	
Not working (disabled)	2.079		1.363	
Not working (other)	0.925		0.223	
Prefer not to answer	0.933		0.522	
<i>Education</i>				
Less than high school degree	0.406		0.509	
High school graduate	0.805		0.225	
Some college but no degree	0.896		0.162	
Associate degree in college	0.669*		0.156	

continued

Table 2 — *continued*

	Model 1		Model 2	
	Odds ratio	se	Odds ratio	se
Bachelor's degree in college			Reference group	
Master's degree			0.793	0.142
Doctoral degree			0.711	0.388
Professional degree (JD, MD)			0.884	0.387
<i>Ethnicity</i>				
African american			1.216	0.408
American indian			0.741	0.397
Asian			1.741	0.604
Hispanic			0.778	0.266
White caucasian			0.937	0.290
<i>Gender</i>				
Male			Reference group	
Female			0.835	0.110
Prefer not to answer			0.271*	0.203
<i>Other socio-demographics</i>				
Age			0.990	0.007
Political spectrum			1.057*	0.032
<i>Income</i>				
Relative Income			1.021	0.066
Effect COVID-19 (current)			0.847*	0.083
Effect COVID-19 (until end 2021)			1.104	0.117
<i>Other controls</i>				
Deal B is blue			0.774**	0.099
Deal B is left			0.917	0.119
<hr/>				
N	1125		1125	
LR $\chi^2(8)$	59.61		105.66	
p-value	0.00		0.00	
Pseudo-R ²	0.04		0.07	

Note. Table reports results of a logistic regression. Dependent variable is a binary variable that equals 1 if a responder chose the fair offer, 0 otherwise. Independent variables: Education ("What is the highest level of school you have completed or the highest degree you have received?"; Employment status ("Which statement best describes your current employment status?"); Ethnicity ("Choose one or more ethnicities that you consider yourself to be."); Political spectrum ("Where would you classify yourself on the left/right political spectrum?"), response scale ranges from 1 (left-wing) to 9 (right-wing); Relative Income ("How do you think your income and financial situation currently compare to those of others in the U.S. who are of similar age?"), response options are Don't know/No answer, Much below average, Somewhat below average, About the average, Somewhat above average, Much above average; Effect COVID-19 (current, "How did the COVID-19 pandemic affect your financial situation?", response options are Worsened a lot, Worsened, Remained unchanged, Improved, Improved a lot; Effect COVID-19 (until end 2021, "How do you expect the COVID-19 pandemic will affect your financial situation until the end of 2021?"), response options are Worsen a lot, Worsen, Remain unchanged, Improve, Improve a lot — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Behavioral Intentions

In the post-experiment questionnaire, we asked subjects to indicate on a scale from 1 (“not at all important”) to 7 (“very important”) how important a number of given reasons were for their decision.

Table 4 shows the mean values of the seven reasons for the full sample and the different combinations of α and β . Looking at the full sample, punishment of the unfair proposer was the least important reason (1.96). Maximizing one’s own payoff (4.95) and the impact on the second round (4.41) were the most important. Table 4 also reports the results of Kruskal-Wallis equality-of-populations rank tests. For all items except punishment of the unfair proposer and deontological reasoning, the test reveals a significant difference across experimental conditions (i.e. combinations of α and β). Overall, the comparison of means across experimental conditions is consistent with expectations. For example, subjects in $\alpha = 0$ report the lowest importance of the reasons “Reduce inequality”, “Reduce externality” and “Impact on second round”. By contrast, subjects in $\alpha = 1$ report the strongest importance of the reasons “Reduce inequality” and “Reduce externality”.

Besides differences between experimental conditions, the items can also be used to compare motives between subjects who chose the selfish versus the unselfish offer. One way to make this comparison is to identify for each subject the reason that was attributed the most (least) importance and determine the relative frequencies depending on the choice.⁸ Reducing payout inequality (30.23 %) and the externality (28.75 %) were most frequently valued as the most important reason by unselfish subjects, whereas punishing (33.50 %) or rewarding (28.96 %) the unfair proposer were most frequently regarded as least important. The most and least important reason most frequently chosen by selfish subjects were the maximization of one’s own payoff (61.04 %) and the impact on the second round (17.33 %) compared to punishment of the proposer of the selfish offer (33.28 %) and deontological reasoning (24.69 %). See Tables 10–12 in Appendix A.3 for more information.

Table 3 shows the means for the items depending on the choice in round one. A Mann-Whitney test is significant for all reasons aside from “Impact on second round”. Of course, this reason should be equally important regardless of whether someone is driven by material self-interest or concern for the third party. The difference for all other items is in the expected direction. Subjects choosing the unselfish offer are

⁸If subjects gave equal importance to more than one reason, ties were randomly broken.

relatively strongly driven by the intention to (i) reduce payoff inequality (5.79 vs. 2.64), (ii) reduce the externality (5.74 vs. 2.43), (iii) “not get their hands dirty” (4.43 vs. 2.09), and (iv) punish the unfair proposer (2.39 vs. 1.65). By contrast, they are relatively weakly driven by the intention to maximize their own payoff (3.18 vs. 6.23) and reward the unfair proposer (2.47 vs. 4.11).

Table 3 — Behavioral intentions by choice

Behavioral intention	Choice		Mann-Whitney test	
	Selfish Mean	Unselfish Mean	<i>z</i>	<i>p</i> -value
Reduce inequality	2.64	5.79	−22.18	0.00
Punish unfair proposer	1.65	2.39	−6.96	0.00
Reduce externality	2.43	5.74	−22.69	0.00
Impact on second round	4.37	4.45	−0.75	0.45
Deontological reasoning	2.09	4.43	−17.99	0.00
Maximize own payoff	6.23	3.18	22.83	0.00
Reward unfair proposer	4.11	2.47	12.37	0.00

Note. Wording of question: For each of the following reasons, please indicate how important it was for your decision to accept Deal [A/B] by Proposer [1/2] in Round 1; Item 1: “Reduce the inequality of payments between participants.”; Item 2: “Punish Proposer [A/B] for choosing a deal associated with a loss for Third Party [1/2]”; Item 3: “Reduce or eliminate the loss for the third parties.”; Item 4: “Increase the probability of an ‘Only [A/B]’ scenario in Round 2.”; Item 5: “Not ‘get my hands dirty’ by accepting an offer that reduced the payment of a third party.”; Item 6: “Maximize my own payoff.”; Item 7: “Reward Proposer [A/B] for an offer that was financially advantageous for me.”; All questions were answered on a 7-point Likert-type scale, where 1 is “not important at all” and 7 is “very important.”

4 Discussion and Conclusion

This paper examines how selfish behavior is affected by uncertainty about ex-ante causal responsibility for a negative externality suffered by third parties. Using moves of nature, we investigate the role of pivotality, the bystander and replacement excuse, and the scope of uncertainty. In addition, we introduce the possibility of intentional uncertainty through ignorance.

Table 4 — Behavioral intentions by combinations of α and β

Behavioral intention	Experimental condition (α, β)										Kruskal-Wallis test ^a	$\chi^2(4)$	p -value		
	Full sample		0,0.5		0.1,0.5		0.5,1		0.5,0					1,-	
	Mean	Mean	Mean	Mean	Mean	Mean	Mean	Mean	Mean	Mean				Mean	Mean
Reduce inequality	3.97	3.54	3.94	3.70	4.01	4.35	13.46	0.01							
Punish unfair proposer	1.96	1.89	1.88	2.02	2.22	2.00	3.14	0.53							
Reduce externality	3.82	3.25	3.79	3.53	3.94	4.27	19.94	0.00							
Impact on second round	4.41	2.52	4.32	5.20	4.98	4.83	120.61	0.00							
Deontological reasoning	3.07	2.78	3.09	2.82	3.26	3.22	7.48	0.11							
Maximize own payoff	4.95	5.17	5.02	5.18	4.94	4.59	9.30	0.05							
Reward unfair proposer	3.42	3.82	3.46	3.34	3.58	3.12	10.16	0.04							

Note. ^a χ^2 adjusted for ties; Wording of question: For each of the following reasons, please indicate how important it was for your decision to accept Deal [A/B] by Proposer [1/2] in Round 1; Item 1: "Reduce the inequality of payments between participants."; Item 2: "Punish Proposer [A/B] for choosing a deal associated with a loss for Third Party [1/2]"; Item 3: "Reduce or eliminate the loss for the third parties."; Item 4: "Increase the probability of an 'Only [A/B]' scenario in Round 2."; Item 5: "Not 'get my hands dirty' by accepting an offer that reduced the payment of a third party."; Item 6: "Maximize my own payoff."; Item 7: "Reward Proposer [A/B] for an offer that was financially advantageous for me."; All questions were answered on a 7-point Likert-type scale, where 1 is "not important at all" and 7 is "very important."

We find that selfish behavior increases by about 25 % once uncertainty is introduced. Interestingly, neither the degree of being pivotal, which specific excuse can be applied, nor the scope of uncertainty affect this result. However, once subjects have no influence on the externality, selfishness significantly increases. Thus, our results suggest an inverse S-shaped relationship between uncertainty and selfish behavior, which can be taken as evidence for self-interested interpretations of uncertainty and excuse-driven behavior (Haisley and Weber 2010; Exley 2015). Our results suggest that only two types of measures can significantly reduce selfish behavior: first, measures that (almost) completely eliminate uncertainty; and second, and starting from a situation where consumers have no (perceived) influence on externalities, measures or factors that make consumers feel that they indeed have influence.

We find that 28 % of subjects forgo payoff even when they have no influence on the externality, and thus appear to engage in deontological moral reasoning. The results of other studies also suggest deontological reasoning (Casal et al. 2019; Falk et al. 2020), but the precise numbers are likely context-specific and thus difficult to compare. For example, Falk et al. (2020) find that about 18 % of subjects who believe that they are not pivotal at all in a simultaneous group decision forgo payoff by voting against the externality. Regarding pivotality, our results contrast with Falk et al. (2020), who find that perceived diffusion of being pivotal in group decisions is significantly correlated with selfish behavior. The differences in decision context between their study and ours may explain these results. Ex-ante pivotality in groups depends on subjective beliefs about the behavior of other group members. Accordingly, high pivotality corresponds with the belief that a relatively large proportion of group members vote against the externality. Accordingly, perceptions of ex-ante pivotality reflect the perceived social norm and level of morality in the group and thus “carry” social information. By contrast, in our experiment, uncertainty and thus pivotality is the result of moves of nature. Thus, our results indicate that behavior may vary depending on how strongly perceptions of pivotality reflect social information. Applying the results to the market context, one could argue that social information plays a lesser role in global markets because consumers have little shared identity. In such decisions, behavior might follow the pattern observed in our study rather than that observed in studies with small groups in laboratory experiments.

In our *Strategic ignorance* condition, two-thirds of subjects do not use the opportunity to remain ignorant about the consequences of their actions on third parties. Overall, we find no significant effect of *Strategic ignorance* on selfish behavior. This is in contrast

to findings in previous literature. One explanation could be that subjects in our experiment had no information about the possible direction of influence on third parties. The high percentage of information disclosure may thus be explained by ambiguity. In studies that find widespread strategic ignorance, such as Dana et al. (2007), Bartling et al. (2014), and Grossman and Weele (2016), subjects know the potential consequences for third parties, whereas only which action leads to which outcome is hidden information. Thus, for decisions where individuals have no prior information or expectations about the potential impact of their actions on others, strategic ignorance may not be that prevalent. From a policy perspective, our finding implies that a large proportion of consumers are interested in the consequences of their actions on others and that information disclosure policies that allow easy access to information at the time of purchase (Weil et al. 2006), such as standards and labels, may have a significant impact on consumer behavior.

We see several interesting avenues for future research. First, in this study, we examined comparatively large changes in pivotality from one hundred percent down to fifty, ten, and zero percent. Our results suggest an inverse S-shaped relationship between pivotality and selfish behavior, which is consistent with excuse-driven behavior. A more fine-grained examination of this relationship may further support this finding and ultimately reveal differences between the replacement and bystander excuse, as well as a combination of the two. Second, while changing the scope of uncertainty had no influence in our experimental setup, we only considered one variation in scope. This finding could also be explored in a more fine-grained study. Third, the role of social information for the perception of ex-ante causal responsibility represents an interesting area of research. How decision-makers perceive social cues may depend on the decision environment, such as the (shared) identity with others or the importance of moves by nature compared to human decisions. Moreover, belief formation and susceptibility to social information may be correlated with personality traits. Fourth, the high proportion of information disclosure in our experiment highlights the dependence of strategic ignorance on the experimental setup and in particular the information that decision-makers have about the possible outcomes for others. Future research could examine how strategic ignorance changes with different prior information about others' outcomes, including uncertainty or vague cues about the direction of influence, given that people rarely know the exact outcomes for others.

Funding: Daniel Hermann acknowledges financial support from Deutsche Forschungsgemeinschaft (DFG) [grant number HE 8478/1-1]. DFG had no involvement in the study design, the collection, analysis and interpretation of data, the writing of the report, and the decision to submit the article for publication.

Declarations of interest: none.

Author Contributions

Conceptualization: M.B., N.E., and D.H.; Data curation: M.B.; Formal analysis: M.B.; Investigation: M.B., N.E. and D.H.; Methodology: M.B., N.E. and D.H.; Project administration: M.B., N.E., and D.H.; Supervision: D.H.; Validation: M.B., N.E., and D.H.; Visualization: M.B., N.E.; Writing—original draft: M.B.; Writing—review editing: M.B., N.E., and D.H.

References

- Bandura, A. (1999). "Moral Disengagement in the Perpetration of Inhumanities". In: *Personality and Social Psychology Review* 3.3, pp. 193–209.
- Bartling, B., F. Engl, and R. A. Weber (2014). "Does willful ignorance deflect punishment? – An experimental study". In: *European Economic Review* 70, pp. 512–524.
- Bartling, B. and U. Fischbacher (2012). "Shifting the Blame: On Delegation and Responsibility". In: *The Review of Economic Studies* 79.1, pp. 67–87.
- Bartling, B., R. A. Weber, and L. Yao (2015). "Do Markets Erode Social Responsibility?" In: *The Quarterly Journal of Economics* 130.1, pp. 219–266.
- Behnk, S., L. Hao, and E. Reuben (2017). *Partners in Crime: Diffusion of Responsibility in Antisocial Behaviors*. IZA Discussion Papers 11031. Institute of Labor Economics (IZA).
- Brütt, K., A. Schram, and J. Sonnemans (2020). "Endogenous group formation and responsibility diffusion: An experimental study". In: *Games and Economic Behavior* 121, pp. 1–31.
- Casal, S., F. Fallucchi, and S. Quercia (2019). "The role of morals in three-player ultimatum games". In: *Journal of Economic Psychology* 70, pp. 67–79.
- Chen, D. L., M. Schonger, and C. Wickens (2016). "oTree—An open-source platform for laboratory, online, and field experiments". In: *Journal of Behavioral and Experimental Finance* 9, pp. 88–97.
- Coffman, L. C. (2011). "Intermediation Reduces Punishment (and Reward)". In: *American Economic Journal: Microeconomics* 3.4, pp. 77–106.
- Dana, J., R. A. Weber, and J. X. Kuang (2007). "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness". In: *Economic Theory* 33.1, pp. 67–80.
- Darley, J. M. and B. Latane (1968). "Bystander intervention in emergencies: Diffusion of responsibility". In: *Journal of Personality and Social Psychology* 8.4, Pt.1, pp. 377–383.
- Engl, F. (2018). "A theory of causal responsibility attribution". In: *Available at SSRN* 2932769.
- Exley, C. L. (2015). "Excusing Selfishness in Charitable Giving: The Role of Risk". In: *The Review of Economic Studies* 83.2, pp. 587–628.
- Falk, A., T. Neuber, and N. Szech (2020). "Diffusion of Being Pivotal and Immoral Outcomes." In: *Review of Economic Studies* 87.5, pp. 2205–2229.
- Falk, A. and N. Szech (2013). "Morals and Markets". In: *Science* 340.6133, pp. 707–711.

- Fischer, P., J. I. Krueger, T. Greitemeyer, C. Vogrincic, A. Kastenmüller, D. Frey, et al. (2011). "The bystander-effect: A meta-analytic review on bystander intervention in dangerous and non-dangerous emergencies." In: *Psychological Bulletin* 137.4, pp. 517–537.
- Gilmore James H., I. and B. J. Pine (1997). "The four faces of mass customization". In: *Harvard Business Review* 75, pp. 91+.
- Grossman, Z. and J. J. van der Weele (2016). "Self-Image and Willful Ignorance in Social Decisions". In: *Journal of the European Economic Association* 15.1, pp. 173–217.
- Haisley, E. C. and R. A. Weber (2010). "Self-serving interpretations of ambiguity in other-regarding behavior". In: *Games and Economic Behavior* 68.2, pp. 614–625.
- Hamman, J. R., G. Loewenstein, and R. A. Weber (2010). "Self-Interest through Delegation: An Additional Rationale for the Principal-Agent Relationship". In: *American Economic Review* 100.4, pp. 1826–46.
- Hauser, O. P., D. G. Rand, A. Peysakhovich, and M. A. Nowak (2014). "Cooperating with the future". In: *Nature* 511.7508, pp. 220–223.
- Holweg, M. and F. K. Pil (2001). "Successful build-to-order strategies start with the customer". In: *MIT Sloan Management Review* 43.1, p. 74.
- Irlenbusch, B. and D. J. Saxler (2019). "The role of social information, market framing, and diffusion of responsibility as determinants of socially responsible behavior". In: *Journal of Behavioral and Experimental Economics* 80, pp. 141–161.
- Norgaard, K. M. (2006). "'People Want to Protect Themselves a Little Bit': Emotions, Denial, and Social Movement Nonparticipation*". In: *Sociological Inquiry* 76.3, pp. 372–396.
- Pigors, M. and B. Rockenbach (2016). "Consumer Social Responsibility". In: *Management Science* 62.11, pp. 3123–3137.
- Reczek, R. W., J. R. Irwin, D. M. Zane, and K. R. Ehrich (2018). "That's Not How I Remember It: Willfully Ignorant Memory for Ethical Product Attribute Information". In: *Journal of Consumer Research* 45.1, pp. 185–207.
- Rothenhäusler, D., N. Schweizer, and N. Szech (2018). "Guilt in voting and public good games". In: *European Economic Review* 101, pp. 664–681.
- Stoll-Kleemann, S., T. O'Riordan, and C. C. Jaeger (2001). "The psychology of denial concerning climate mitigation measures: evidence from Swiss focus groups". In: *Global Environmental Change* 11.2, pp. 107–117.
- Weil, D., A. Fung, M. Graham, and E. Fagotto (2006). "The effectiveness of regulatory disclosure policies". In: *Journal of Policy Analysis and Management* 25, pp. 155–181.

Ziegler, A., G. Romagnoli, and T. Offerman (2021). *Morals in multi-unit markets*. Tinbergen Institute Discussion Papers 20-072/I. Tinbergen Institute.

Appendix

A.1 Socio-demographics

Table 5 — Education

What is the highest level of school you have completed or the highest degree you have received?	N	%
Less than high school degree	3	0.27
High school graduate	70	6.22
Some college but no degree	221	19.64
Associate degree in college (2-year)	104	9.24
Bachelor's degree in college(4-year)	474	42.13
Master's degree	212	18.84
Doctoral degree	16	1.42
Professional degree (JD, MD)	25	2.22

Table 6 — Employment status

Which statement best describes your current employment status?	N	%
Working (paid employee)	668	59.38
Working (self-employed)	177	15.73
Not working (temporary layoff from a job)	26	2.31
Not working (looking for work)	104	9.24
Not working (retired)	23	2.04
Not working (disabled)	12	1.07
Not working (other)	97	8.62
Prefer not to answer	18	1.60

Table 7 — Ethnicity

Choose one or more ethnicities that you consider yourself to be:	N	%
African American	86	7.64
American Indian	16	1.42
Asian	109	9.69
Hispanic	58	5.16
White Caucasian	918	81.60
Other	13	1.16

Table 8 — Political spectrum

Where would you classify yourself on the left/right political spectrum?	N	%
1 - left-wing	152	13.51
2	199	17.69
3	172	15.29
4	110	9.78
5	220	19.56
6	104	9.24
7	83	7.38
8	58	5.16
9 - right-wing	27	2.40

Table 9 — Relative income

How do you think your income and financial situation currently compare to those of others in the U.S. who are of similar age?	N	%
Don't know/No answer	25	2.22
Much below average	181	16.09
Somewhat below average	266	23.64
About the average	355	31.56
Somewhat above average	267	23.73
Much above average	31	2.76




Figure 4 — Age of participants

Note. Subjects were asked about their year of birth. To approximate age, we assumed that subjects who participated in December 2020 already had a birthday. For subjects who participated between January and February 2021, we assumed the opposite.

A.2 Data Quality Checks

Quality of answers to an open-ended question

As a test of the quality of our data, we asked subjects to explain their decision using an open-ended question in the post-experimental questionnaire. Overall, the data quality for this question is quite good, both linguistically and in terms of scope. The vast majority of participants gave comprehensible reasons for their decisions, which demonstrated their understanding of the experiment and payoff scheme. There were no copied texts or meaningless combinations of strings. Less than 5 percent gave answers that were inconsistent with their choices in the experiment or that showed a misunderstanding of the experiment. On average, subjects wrote 18 words (with a maximum of 151) or 94 characters (with a maximum of 820). Subject that chose the unselfish offer wrote significantly more (104.81 vs. 86.13 characters; t -stat = 3.71, p -value < 0.00).

Sensitivity to changes in incentives

To test the validity of our experimental set-up and subjects' understanding of the instructions and experiment, we tested subjects' sensitivity to changes in incentives under full pivotality in MTS. We ran an additional study with three variations of payoffs for the third party associated with the selfish offer: 70 Coins (as in the main study), 40 Coins, and 10 Coins. We collected 50 observations per treatment. We preregistered this study at <https://aspredicted.org/blind.php?x=xr7e79>. The proportion of selfish decisions is 56, 38, 28 percent for the third-party payoffs of 70, 40, 10. Pearson $\chi^2 = 8.34$ with $p = 0.015$.

A.3 Behavioral intentions

Table 10 — Percentage of *selfish* subjects reporting reason as most important

Behavioral intention	%
Reduce inequality ¹	30.23
Reduce externality ³	28.75
Impact on second round ⁴	15.43
Deontological reasoning ⁵	12.47
Maximize own payoff ⁶	8.03
Punish unfair proposer ²	2.96
Reward unfair proposer ⁷	2.11

Note. Ties were randomly broken; Wording of question: For each of the following reasons, please indicate how important it was for your decision to accept Deal [A/B] by Proposer [1/2] in Round 1; ¹: “Reduce the inequality of payments between participants.”; ²: “Punish Proposer [A/B] for choosing a deal associated with a loss for Third Party [1/2]”; ³: “Reduce or eliminate the loss for the third parties.”; ⁴: “Increase the probability of an ‘Only [A/B]’ scenario in Round 2.”; ⁵: “Not ‘get my hands dirty’ by accepting an offer that reduced the payment of a third party.”; ⁶: “Maximize my own payoff.”; ⁷: “Reward Proposer [A/B] for an offer that was financially advantageous for me.”; All questions were answered on a 7-point Likert-type scale, where 1 is “not important at all” and 7 is “very important.”

Table 11 — Percentage of *selfish* subjects reporting reason as least important

Behavioral intention	%
Punish unfair proposer ²	33.40
Reward unfair proposer ⁷	28.96
Maximize own payoff ⁶	12.68
Impact on second round ⁴	10.36
Deontological reasoning ⁵	9.51
Reduce externality ³	4.02
Reduce inequality ¹	1.06

Note. Ties were randomly broken; Wording of question: For each of the following reasons, please indicate how important it was for your decision to accept Deal [A/B] by Proposer [1/2] in Round 1; ¹: “Reduce the inequality of payments between participants.”; ²: “Punish Proposer [A/B] for choosing a deal associated with a loss for Third Party [1/2]”; ³: “Reduce or eliminate the loss for the third parties.”; ⁴: “Increase the probability of an ‘Only [A/B]’ scenario in Round 2.”; ⁵: “Not ‘get my hands dirty’ by accepting an offer that reduced the payment of a third party.”; ⁶: “Maximize my own payoff.”; ⁷: “Reward Proposer [A/B] for an offer that was financially advantageous for me.”; All questions were answered on a 7-point Likert-type scale, where 1 is “not important at all” and 7 is “very important.”

Table 12 — Percentage of *unselfish* subjects reporting reason as most important

Behavioral intention	%
Maximize own payoff ⁶	61.04
Impact on second round ⁴	17.33
Reward unfair proposer ⁷	12.27
Reduce inequality ¹	5.37
Reduce externality ³	2.61
Deontological reasoning ⁵	1.23
Punish unfair proposer ²	0.15

Note. Ties were randomly broken; Wording of question: For each of the following reasons, please indicate how important it was for your decision to accept Deal [A/B] by Proposer [1/2] in Round 1; ¹: “Reduce the inequality of payments between participants.”; ²: “Punish Proposer [A/B] for choosing a deal associated with a loss for Third Party [1/2]”; ³: “Reduce or eliminate the loss for the third parties.”; ⁴: “Increase the probability of an ‘Only [A/B]’ scenario in Round 2.”; ⁵: “Not ‘get my hands dirty’ by accepting an offer that reduced the payment of a third party.”; ⁶: “Maximize my own payoff.”; ⁷: “Reward Proposer [A/B] for an offer that was financially advantageous for me.”; All questions were answered on a 7-point Likert-type scale, where 1 is “not important at all” and 7 is “very important.”

Table 13 — Percentage of *unselfish* subjects reporting reason as least important

Behavioral intention	%
Punish unfair proposer ²	33.28
Deontological reasoning ⁵	24.69
Reduce inequality ¹	13.65
Reduce externality ³	12.42
Reward unfair proposer ⁷	7.67
Impact on second round ⁴	7.21
Maximize own payoff ⁶	1.07

Note. Ties were randomly broken; Wording of question: For each of the following reasons, please indicate how important it was for your decision to accept Deal [A/B] by Proposer [1/2] in Round 1; ¹: “Reduce the inequality of payments between participants.”; ²: “Punish Proposer [A/B] for choosing a deal associated with a loss for Third Party [1/2]”; ³: “Reduce or eliminate the loss for the third parties.”; ⁴: “Increase the probability of an ‘Only [A/B]’ scenario in Round 2.”; ⁵: “Not ‘get my hands dirty’ by accepting an offer that reduced the payment of a third party.”; ⁶: “Maximize my own payoff.”; ⁷: “Reward Proposer [A/B] for an offer that was financially advantageous for me.”; All questions were answered on a 7-point Likert-type scale, where 1 is “not important at all” and 7 is “very important.”

A.4 Experimental instructions

Text

Group composition and role assignment

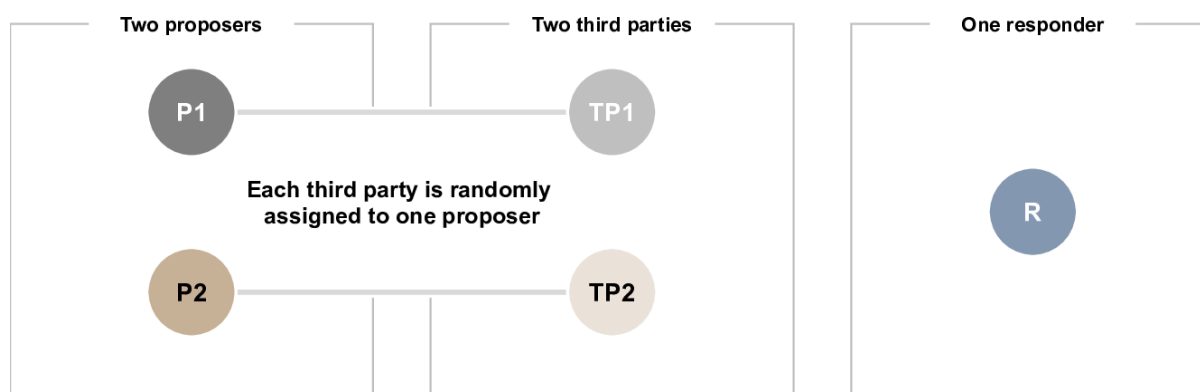
After you have read the instructions and answered the comprehension questions correctly, you will be matched into groups of five participants.

All participants in your group will be assigned to one of three different roles. In each group, there are

- two proposers (P1 and P2),
- two third parties (TP1 and TP2) and
- one responder (R).

Each third party is randomly assigned to one proposer. To indicate this assignment, we will denote the two pairs as “Third Party 1, Proposer 1” and “Third Party 2, Proposer 2”. The figure below illustrates the group composition.

The group composition and role assignment remain unchanged throughout the experiment.



- 1 *The task*
- 2 This task lasts for two rounds. In each round, each participant initially receives 100
- 3 coins.

4 The final earnings of all participants per round may differ from this initial endowment
 5 and depend on the decisions of the two proposers and the responder. Third parties
 6 cannot make decisions.

7 Each round comprises two consecutive steps:

- 8 • **Step 1:** Both proposers offer a deal to the responder.
- 9 • **Step 2:** The responder decides which of the two offered deals he/she wants to
 10 accept.

11 Deals are always beneficial for proposers and the responder, but they can also affect
 12 the earnings of third parties. The earnings of a third party are [*MTO*: only affected
 13 by a deal if it is accepted by the responder; *MTS*: affected by a deal as soon as it is
 14 offered by a proposer. Therefore, irrespective of the responder's choice, the earnings
 15 of third parties are determined by the deals offered by the proposers.]

16 The following illustration shows the two consecutive steps of the deal-making process
 17 and at which step payment consequences occur for each participant. As an example,
 18 the illustration assumes that the responder accepts the deal offered by Proposer 1.

19 [*MTO*: Figure 5; *MTS*: Figure 6]

20 *The two deals*

21 Proposers and responder can make two types of deals: [Deal A](#) and [Deal B](#). The table
 22 below shows the payment consequences of the deals and the resulting earnings per
 23 round, taking into account the initial endowment of 100 coins. [*Strategic ignorance*: As
 24 you can see from the table, the third parties' values are not visible.

25 After the comprehension questions but before any decisions are made, each participant
 26 has the opportunity to reveal the payment consequences of the two deals for the third
 27 party. A participant's decision whether to reveal or not is not shared with other
 28 participants.]

29 [*Full information*: Table 14; *Strategic ignorance*: Table 15]

30 *Earnings per round*

31 **Proposers:** The proposer of the accepted deal receives 110 coins (+10 coins). The
 32 proposer of the rejected deal, keeps the initial endowment of 100 coins (± 0 coins).

33 **Third parties:** Third party earnings are [MTO: only affected by a deal if it is accepted
 34 by the responder. If a responder accepts Deal A, the assigned third party receives
 35 [Full information: 70 coins (-30 coins); Strategic ignorance: ??? coins (??? coins)]. If a
 36 responder accepts Deal B, the assigned third party receives [Full information: 100 coins
 37 (± 0 coins); Strategic ignorance: ??? coins (??? coins)]. The third party assigned to the
 38 rejected deal remains unaffected and keeps the initial endowment of 100 coins (± 0
 39 coins); MTS: affected by a deal as soon as it is offered by a proposer and therefore they
 40 do not depend on the responder's choice. If a proposer offers Deal A, the assigned
 41 third party receives [Full information: 70 coins (-30 coins); Strategic ignorance: ???
 42 coins (??? coins)]. If a proposer offers Deal B, the assigned third party receives [Full
 43 information: 100 coins (± 0 coins); Strategic ignorance: ??? coins (??? coins)].

44 **Responder:** If the responder accepts Deal A, he/she receives 120 coins (+20 coins). If
 45 the responder accepts Deal B, he/she receives 110 coins (+10 coins).

46 *Offers in Round 2*

47 While deals offered in Round 1 are deliberately chosen by proposers, the deals offered
 48 in Round 2 are determined for the proposers by a computerized urn draw. The
 49 computerized urn draw can lead to one of two possible scenarios:

- 50 • **Only A:** Both proposers offer Deal A
- 51 • **Only B:** Both proposers offer Deal B

52 *Composition of the urn and the responder's influence on the offers in Round 2*

53 The urn is filled with 20 balls. Each ball is labeled with either A or B and has the same
 54 chance of being drawn. If an A ball is randomly drawn, Round 2 will be the **Only A**
 55 scenario. If a B ball is randomly drawn, Round 2 will be the **Only B** scenario.

56 The following illustration shows the composition of the urn depending on the re-
 57 sponder's choice in Round 1. In addition, the probabilities for both scenarios are
 58 shown.

59 [Figure 7 for $\alpha = 0.1$ and $\beta = 0.5$ as an example]

60 If the responder chooses Deal A in Round 1, Round 2 will be either an **Only A** scenario
 61 with a probability of 55 percent or an **Only B** scenario with a probability of 45 percent.

If the responder chooses Deal B in Round 1, Round 2 will be either an Only A scenario with a probability of 45 percent or an Only B scenario with a probability of 55 percent.

Tables and figures

Table 14 — Tabular representation of the payoffs associated with the deals in *Full information*

	Deal A		Deal B	
	Payment consequences	Earnings	Payment consequences	Earnings
Proposer	+10 coins	110 coins	+10 coins	110 coins
Third Party	-30 coins	70 coins	±0 coins	100 coins
Responder	+20 coins	120 coins	+10 coins	110 coins

Table 15 — Tabular representation of the payoffs associated with the deals in *Strategic ignorance*

	Deal A		Deal B	
	Payment consequences	Earnings	Payment consequences	Earnings
Proposer	+10 coins	110 coins	+10 coins	110 coins
Third Party	??? coins	??? coins	??? coins	??? coins
Responder	+20 coins	120 coins	+10 coins	110 coins

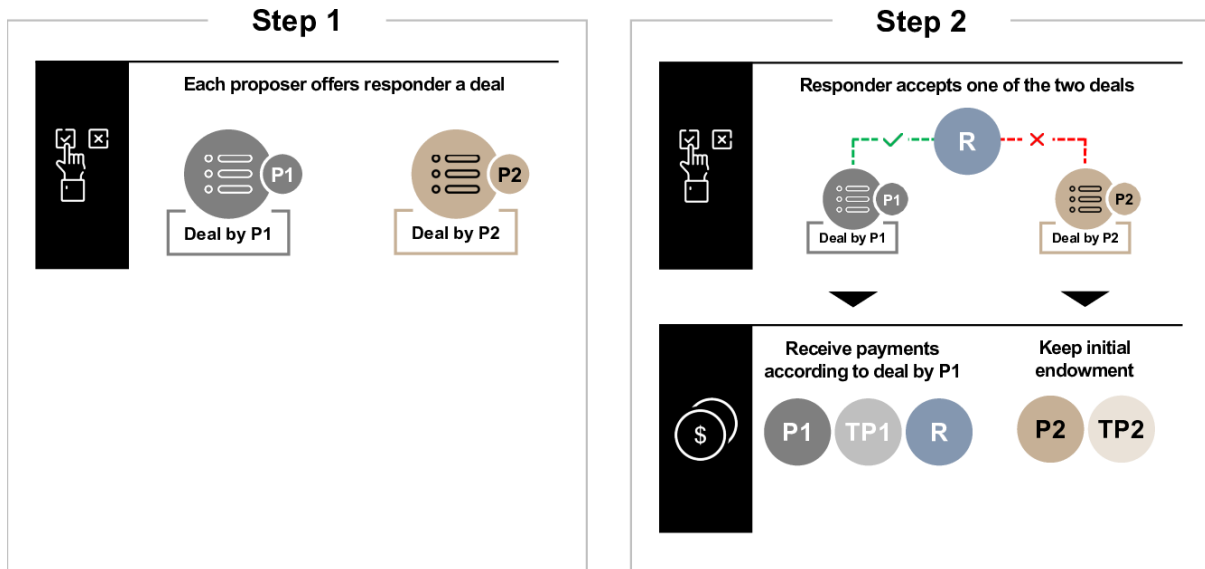


Figure 5 — Graphical representation of the steps in *MTO*

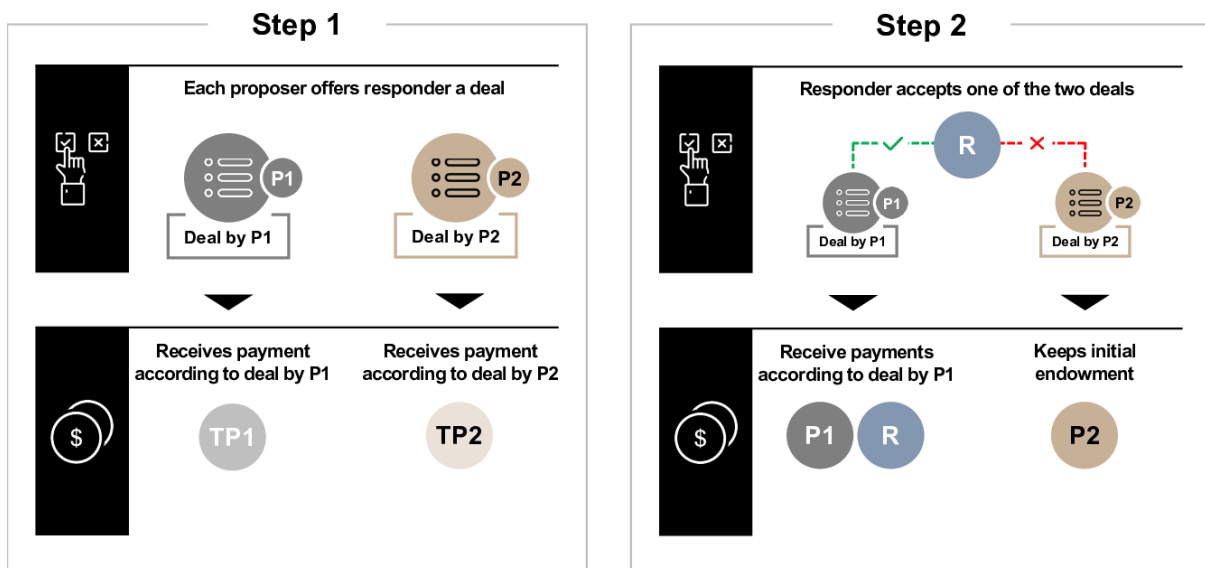


Figure 6 — Graphical representation of the steps in *MTS*

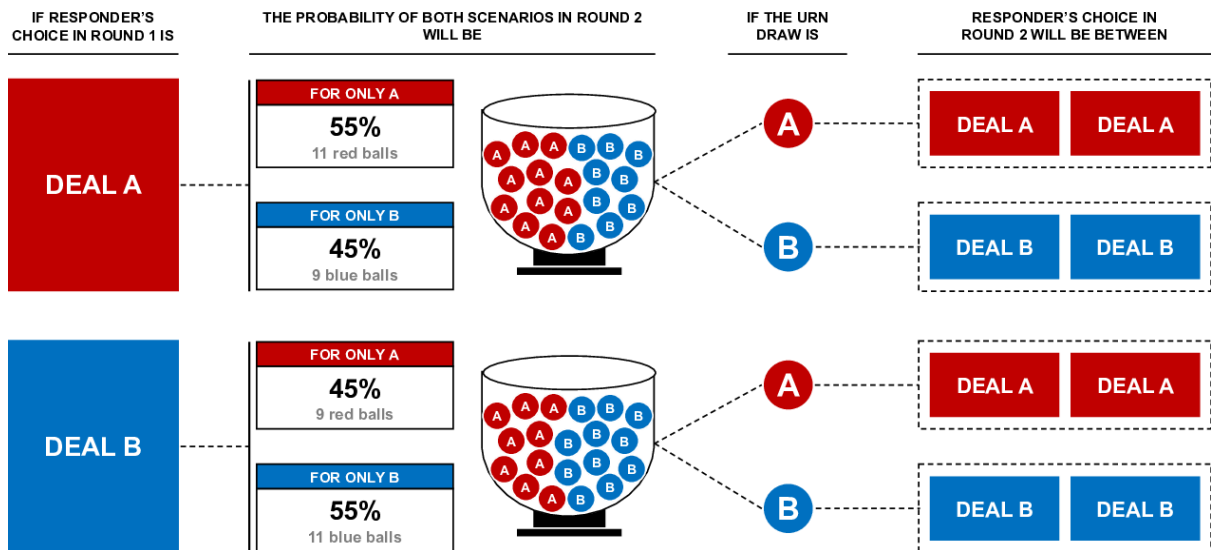


Figure 7 — Graphical representation of the urn drawing for $\alpha = 0.1$ and $\beta = 0.1$

A.5 Comprehension questions

Table 16 — Comprehension questions

Question	Type
<p>At which step of the deal-making process do third parties bear the payment consequences of a deal?</p> <ul style="list-style-type: none"> • Third parties bear no payment consequences at all. • At step 1, as soon as a proposer offers a deal. • At step 2, only if the responder accepts a deal. 	Single-answer multiple choice question
<p>The deals offered in Round 1 are ...</p> <ul style="list-style-type: none"> • deliberately chosen by proposers • determined by a computerized urn drawing • negotiated between proposers and third parties 	Single-answer multiple choice question
<p>Imagine that in Round 1 Proposer 1 offers Deal A and Proposer 2 offers Deal B. The responder accepts Deal B by Proposer 2. How many Coins does each participant earn in Round 1?^a</p> <ul style="list-style-type: none"> • Responder • Proposer 1 • Third Party 1 • Proposer 2 • Third Party 2 	Dropdown question; List of answers ranges from 0 to 150 Coins in 10 Coin increments and “???”

continued

Table 16 — *continued*

Question	Type
<p>Imagine an Only A scenario, where a responder accepts the deal offered by Proposer 1. How many Coins do the third parties earn in that round?^a</p> <ul style="list-style-type: none"> • Third Party 1 • Third Party 2 	<p>Dropdown question; List of answers ranges from 0 to 150 Coins in 10 Coin increments and “???”</p>
<p>Imagine an Only B scenario, where a responder accepts the deal offered by Proposer 1. How many Coins do the third parties earn in that round?^a</p> <ul style="list-style-type: none"> • Third Party 1 • Third Party 2 	<p>Dropdown question; List of answers ranges from 0 to 150 Coins in 10 Coin increments and “???”</p>
<p>Which scenario, i.e. which combination of deals is possible in Round 2?^b</p> <ul style="list-style-type: none"> • Deal A and Deal B: One proposer offers Deal A and the other proposer offers Deal B • Only A: Both proposers offer Deal A • Only B: Both proposers offer Deal B 	<p>Multiple choice question</p>
<p>If the responder chooses Deal A in Round 1, what is the probability of an Only A scenario in Round 2?^c</p>	<p>Open-ended question, answers were limited to integers between 0 and 100</p>
<p>If the responder chooses Deal B in Round 1, what is the probability of an Only B scenario in Round 2?^c</p>	<p>Open-ended question, answers were limited to integers between 0 and 100</p>

Note. For some questions, we provided an additional hint, which was shown directly below the question in italics: ^a“Please take into account the initial endowment of 100 Coins, i.e. if a participant gains 10 Coins as a result of a deal, choose ‘110’. Remember that Third Party 1 is assigned to Proposer 1 and Third Party 2 is assigned to Proposer 2. [*Strategic ignorance*: If a payment is unknown, please choose the option ‘???.’.]”, ^b“Remember that the deals offered by the proposers in Round 2 are determined by a computerized urn drawing.”, ^c“Please enter the probability in percent. Remember that the deals offered by the proposers in Round 2 are determined by a computerized urn drawing.”

CHAPTER III

Consumer Behavior Under Benevolent Price Discrimination

Alexander Erlei, Nils Engelbrecht, and Mattheus Brenig

Extensive research shows that consumers are generally averse to price discrimination. However, instruments of differential pricing can benefit consumer surplus and alleviate inequity through targeted price discounts. This paper examines whether and how these outcome considerations influence consumer reactions to price discrimination. The authors introduce the concept of benevolent price discrimination (BPD) as a downward-bound instrument that always benefits financially disadvantaged groups or individuals, thereby leading to a convergence of outcomes. Five experiments with 3415 participants show that a large share of consumers is willing to costly switch away from a store that offers a discount to low-income consumers. This happens irrespective of whether income differences are due to luck or merit. While the price-discriminating store does attract some new high-income consumers, it cannot compensate the loss of existing consumers. Simulating market transactions by endowing consumers with agency over future prices increases costly support for BPD. These results contrast previous findings on social preferences and inequity aversion, highlighting the importance of context and procedure for economic bargaining. Strong behavioral constraints persist even when price discrimination reduces unmerited income differences and no consumer experiences price increases.

Keywords Price Discrimination · Fairness and Inequity Aversion · Behavioral Constraints · Decision-Making · Pricing

1 Introduction

Technological progress, especially the advent of big data, has reduced the technological and informational constraints on sellers to differentiate prices among consumers. Despite these developments, sellers still hesitate to deviate from uniform pricing strategies (Fudenberg and Villas-Boas 2012). One commonly evoked explanation is that sellers anticipate the risk of antagonizing consumers (Fabiani et al. 2006), as exemplified by research showing that differential price strategies are often at odds with consumers' fairness perceptions (see e.g., Xia et al. 2004; Haws and Bearden 2006), which translates into shopping intentions (Bolton et al. 2010), purchase satisfaction (Shor and Oliver 2006), intentions to spread negative word-of-mouth (Ferguson et al. 2014) and consumption choices (Leibbrandt 2020). Thus, in addition to technological constraints, profit-seeking through demand-based price discrimination may also be inhibited by negative consumer reactions, i.e., "behavioral constraints" (Kahneman et al. 1986).

However, consumers do not appear to be universally averse to demand-based price discrimination. Most prominently, price discounts for specific groups of the population, like students and the elderly, are common in many countries. One explanation for this ostensible contradiction is that consumers accept price discrimination if it benefits lower income groups, thereby leading to an equalization of economic outcomes (Rotemberg 2011; Wu et al. 2012). Whilst theoretically appealing, this assertion has thus far not been empirically verified. Existing behavioral research primarily focuses on processes, without direct control or transparency over the realized distribution of economic outcomes between consumers. Alternative explanations for the prevalence of certain price discounts can therefore not be ruled out. Regarding e.g. the elderly, the acceptance of lower prices may instead be driven by habituation, learned societal norms, or compensation for other kinds of inequality like quality of living. Therefore, isolating the effect of economic consequences on consumer behavior requires a more abstract examination of the problem. If consumer reactions to price discrimination do depend on the concurrent equalization of consumer outcomes, this insight could open up new possibilities for sellers to implement and market tools of differential price setting. This includes a broader range of potential price discount recipients, poaching prices, improved judgments regarding useful degrees of strategic price obfuscation and transparency, or promoting certain pricing strategies as part of corporate social responsibility communications.

This paper experimentally addresses whether and how consumer reactions depend on the effect of interpersonal price discrimination on the overall distribution of consumer welfare. In particular, we introduce the concept of *benevolent price discrimination (BPD)*. We define BPD as a policy of differential pricing that is downward-bound and always benefits financially disadvantaged groups or individuals, thereby leading to a convergence of outcomes. Whilst a seller can decrease prices for certain individuals or groups, they cannot increase prices for anyone. Thus, nobody is worse off as a consequence of price-discrimination. Consumers without access to the lower price experience *relatively disadvantageous*, rather than *disadvantageous* price discrimination. Most importantly, BPD always leads to a more equal outcome distribution. We thus quantify behavioral constraints towards price discrimination when implemented as a mechanism that increases joint overall welfare while creating a more equitable distribution of economic gains. By abstracting from various cultural and context-dependent factors, our analysis avoids behavioral distortions induced by residual variables and derives generalizable results that isolate the effects of economic outcome equalization on consumer behavior.

Procedural and distributive price fairness

The price fairness literature has long recognized the importance of outcome-related reasoning, identifying two distinct main criteria that determine consumer price judgments: procedural and distributive price fairness (Maxwell 2007). Procedural price fairness refers to the procedures or processes behind a pricing strategy, including intentions (Campbell 1999) and the specific variables determining differential pricing (Variable of Discrimination, VOD; Bayer 2010; Kuo et al. 2016). Distributive price fairness refers to comparisons with respect to the outcome of other actors, which has thus far often been interpreted as the offered price (see e.g. Ferguson et al. 2014; Xia et al. 2004). This conceptualization, while informative, does not address distributive fairness with regard to the distribution of consumer welfare. To integrate outcome equalization into the analysis of consumer behavior under price discrimination, research needs to relate price differences to differences in purchasing power and the subsequent distribution of e.g. goods or income, rather than equating offered prices with outcomes. For example, prices that are tailored towards different income groups might be seen as distributively unfair because people have to pay different amounts for the same good or service, but might simultaneously lead to a convergence of

consumer welfare, which could increase distributive fairness. This article therefore refers to "outcomes" as the final distribution of consumer welfare, and interprets the unequal treatment of consumers through price differences as a procedural variable.

Despite these standing conceptual differences, prior research clearly establishes the importance of between-consumer comparisons for adverse reactions to price discrimination. One central tenet is that a consumer explicitly or implicitly refers to an internal (e.g. Thaler 1985; Janiszewski and Lichtenstein 1999; Koszegi and Rabin 2006) or external (e.g. Mayhew and Winer 1992; Dholakia and Simonson 2005) reference point to assess whether an offered price is fair. Internal, memory-based reference points are the prices that the consumer themselves paid in similar previous transactions (Herz and Taubinsky 2018). External, interpersonal reference points are the prices paid by one's peers (Ho and Su 2009), i.e. other consumers (Haws and Bearden 2006; Jin et al. 2014), for an identical good. Although both factors are important, there is evidence to suggest that interpersonal comparison has a larger impact on consumers' fairness perception than self-comparison (see e.g. Xia et al. 2004), which is particularly relevant in modern online markets that exhibit high transparency and extensive information sharing (Anderson and Simester 2008).

However, moving the analysis to the level of realized economic outcome distributions requires a conceptual shift whereby potential preferences for equal outcomes are tested within a price setting framework that replicates the important prerequisites of consumer price comparisons.

Preferences for Equal Outcomes

There is ample evidence that people generally care about the equal distribution of monetary outcomes (Fehr et al. 1993; Harrington Jr et al. 2016), particularly if existing income inequalities are seen as arbitrary (Rutström and Williams 2000; Alesina and Angeletos 2005; Durante et al. 2014).

Two very influential economic models propose that a person's utility depends on the outcome of other people, either because people experience disutility when their outcome is different from other people's outcomes (Fehr and Schmidt 1999), or because their utility function depends on their share of total payoffs (Bolton and Ockenfels 2000). Thus, as long as differential pricing leads to a convergence of outcomes without

any (or many) consumer losses, it may not only be accepted by consumers, but even be desirable.¹

However, supporting results from the literature on inequity aversion and studies showing human preferences for redistribution under unequal economic outcomes cannot be readily applied to the context of price discrimination in consumer markets. Foremost, this literature purposely abstracts from many process-related variables in order to focus exclusively on outcome distributions. It thus factors out many elements of the price setting mechanism, which could well interact with or mediate distributive preferences as well as their translation into behavior. For example, under price discrimination, redistribution is not actively chosen by consumers, but imposed by sellers. This evokes considerations regarding reciprocity (Fehr and Gächter 2000) and consumers act as second, rather than first movers, strongly inhibiting their agency in determining prices and thereby outcomes (Choshen-Hillel and Yaniv 2011; Alexopoulos et al. 2013). Moreover, moral responsibility for causing fairer market outcomes is shifted away from the consumer. At best, consumers may ascribe themselves moral responsibility through their influence on future prices and thereby future outcomes (Pigors and Rockenbach 2016).

Outcomes vs. Procedures

Overall, it remains unclear whether outcome equalization can alleviate behavioral constraints to price discrimination. This will crucially depend on the extent to which consumers weigh unequal price treatment (e.g., not being eligible for a price discount other consumers receive; Ho and Su 2009; Ho et al. 2014) against the equalization of consumer outcomes. This relationship is not obvious, since there is evidence from outside the price literature that both outcomes (Lerner and Whitehead 1980; Rutte and Messick 1995) and procedures (Folger and Konovsky 1989; Van den Bos, Vermunt, et al. 1997) could be more important in forming overall fairness judgments. In particular, a large body of work highlights the role of procedural elements e.g., in bargaining contexts (Trautmann 2009) or for contextual variables like norms (Garbarino and Maxwell 2010) and consumer loyalty (Anderson and Simester 2010). Other research finds equivalent effects on allocation acceptability (Bolton et al. 2005). According to Van

¹See Krawczyk and Le Lec (2010) for a model of procedural and distributive fairness that predicts higher support of redistribution in such cases. See also Guo (2015) and Li and Jain (2016) for theoretical analyses on how consumer inequity aversion and social preferences affect optimal buyer strategies.

den Bos, Lind, et al. (1997), the availability of outcome information could mediate the influence of procedural fairness. While procedures were significant in shaping fairness judgments without outcome information, subjects who knew about the outcomes of other participants did not appear to be significantly influenced by procedures. In Bolton et al. (2005), fair outcomes were accepted irrespective of procedural bias.

Thus, in so far as these results translate into purchasing environments, consumers who know that BPD only benefits those with lower incomes might accept or even support it. However, for consumers who do not receive information on other consumers' outcomes and experience relatively disadvantageous price discrimination, the literature on price fairness and reference points suggests comparatively strong aversions.

Study Overview

To assess consumer reactions towards BPD, we conducted a series of incentivized, context-neutral, and controlled experimental studies on Amazon Mechanical Turk (MTurk). Sellers could offer a price discount for low-income consumers, leading to a convergence in consumer outcomes. Price discrimination was transparent and high-income consumers could either costly switch to a second seller who did (not) offer the discount, choose not to buy at all, or maximize their income by staying with the same seller across all rounds. By quantifying both, behavioral shifts away from and towards the price discriminating store, we were able to capture bidirectional consumer migration patterns which have so far been largely neglected by the experimental literature. This allowed us to make qualitative judgments about the net-effect of BPD, rather than focusing solely on the behavioral constraints elicited by existing consumers.

In the first study, high-income consumers received no explanation as to who benefited from the price discount or how it was implemented. Whilst this masked an important element of our framework, abstracting from differences in outcomes allowed us to establish a first benchmark of consumer reactions towards differential downward pricing.

The second study addressed the role of reciprocity by keeping sellers out of the price-setting process. By shifting responsibility for the pricing decision from sellers to an algorithm, we were also able to examine the role of human intentionality in consumer reactions to price discrimination. Studies 3 to 5 then progressively introduced new

process-related variables, endowing consumers with information and agency that mirror important contextual variables of real-world consumer markets and enable concrete between-consumer comparisons. In the third study, high-income consumers learned that price discounts were available for low-income consumers only. This change revealed the benevolent nature of our price discounts, and thus introduced the interaction between unequal price treatment and an equalization of consumer outcomes. Moreover, we used two mechanisms to allocate initial incomes. In one condition, income was randomly assigned to consumers, in the other, consumer income was assigned based on their performance in a real-effort task. This allowed us to test the prediction of equity theory (i.e. distributive fairness), whereby an outcome distribution is perceived as fair if outcomes are proportional to an individuals' input. In Study 4, we increased consumers' moral responsibility for their purchasing behavior by introducing a market mechanism which gave consumers agency over future prices and thereby outcomes. Finally, we conducted three robustness checks addressing transparency, the costs of switching and consumer motivations behind their purchasing decisions.

Our results suggest widespread consumer aversions towards BPD. Across all studies, between 30 %–40 % of consumers exhibited strictly BPD-averse purchasing patterns. Reactions were the strongest in Study 1, where consumers did not receive any contextual explanations for price discrimination. As shown in Study 2, this behavior can not be explained by consumers' assumptions about seller intentions. That is, holding seller intentions constant does not substantially affect consumer aversion towards BPD. Even in Study 3, with a clear understanding that price discrimination only benefited consumers who were arbitrarily assigned a lower income, every third consumer chose to costly switch away from the price discriminating store. Endowing consumers with agency through a market mechanism in Study 4 did slightly decrease costly switching, but only if income differences were generated randomly. Contrary to studies 1–3, a sizeable number of consumers was willing to *support* BPD by costly switching *towards* the price-discriminating seller when consumer purchases influenced the store's future pricing strategy. Finally, Study 5 subsumes three robustness checks that confirm the validity of our behavioral interpretations and show that results hold when consumers decide under full transparency as well as for varying switching costs.

Overall, we show that outcome-fairness related reasoning mediates behavioral consumer constraints only under very specific conditions. In particular, consumers need to participate in determining future prices. Highlighting the re-distributive nature of

BPD has little effect on behavior, and contrary to much of the existing literature on social preferences, consumers do not differentiate between merited and unmerited income inequality. Even in the most benevolent case, a substantial share of consumers is willing to give up money in order to switch to a competitor without BPD despite no economic gains.

The complete data set, additional information and this project's code can be accessed under the following repository:

https://osf.io/sztnh/?view_only=fc170dcbfed4431e994b23ed095d8b4b

2 Study 1

Study 1 first used a simple one-shot design to elicit general consumer preferences for price discrimination. Second, we examined whether these preferences generalize to a more comprehensive multiple round within-subject scenario with monetary “switching costs”. The latter reflects that in reality, punishing sellers for their pricing strategy by purchasing elsewhere usually entails transaction costs – even in the case of homogeneous goods (Nilssen 1992). Throughout both experiments, subjects did not receive any explanation about the introduction of price discrimination. The beneficiaries of BPD were kept private and consumers could not infer that price discounts led to more equal outcomes. In reality, it is likely that BPD would be used in situations where consumers cannot know how different prices are generated and who benefits. As we will show in later studies, the main effects of this study replicate to purchase environments where outcome equalization is transparent to consumers.

2.1 Experimental Design

Because the basic experimental setup was the same across all studies, we will describe it here in more detail and refer to this description in the remainder of the paper.

We ran purchasing experiments where consumers had the option to buy a homogeneous good at one of two stores, each run by another participant acting as store manager. We measure preferences for or against price discrimination by a consumer's store choice. We included a post-experimental questionnaire measuring participants' fairness perceptions, attitudes towards price discrimination as well as a questionnaire on social comparison (Gibbons and Buunk 1999). All experiments in this paper were

conducted online using MTurk and oTree (Chen et al. 2016). Participants enrolled on their own accord and were randomly assigned to one treatment and one role. After the instructions, participants had two attempts to answer five comprehension questions correctly. Those who failed both attempts were excluded from participating. We used “Coins” as our experimental currency. Coins were later converted into dollars, where 10 coins equaled 3 cents. Additionally, subjects playing the role of “consumer” received a fixed payment of \$1.20 for completing the survey. There were two main experimental paradigms to analyze consumer behavior: *One-Shot* and *Repeated*.

One-Shot

Upon arrival, participants were randomly assigned the role of either a *high-income* or a *low-income* consumer. High-income (low-income) consumers learned that they would receive an endowment of 100 (50) Coins. Coins could be used to buy a good at one of two stores (A, B). In both stores, the good returned a value of 150 Coins to a consumer’s final payoff. Thus, consumers were incentivized to always buy the good. The two stores sold the same good and were identical except for one thing: one store offered all consumers the good for a price of 100 Coins. The other store offered low-income consumers the good for a discounted price of 50 Coins and charged high-income consumers the regular price of 100 Coins. Consumers learned that each store was run by a manager, who had decided on the pricing strategy of the store beforehand, and that managers earned money for each good sold in their store. This ensured that consumer purchasing choices were meaningful.

Participants decided at which of the two stores they wanted to purchase the good. High-income consumers could observe that one store offered the good at a price of 50 Coins and learned with a click that they were not eligible for that discount. However, in Studies 1 and 2, no explanation for the different prices was given.

We randomized the position (left, right) and name (A, B) of the store introducing price discrimination. Thus, if price discrimination had no impact, we would expect consumers to be equally distributed across both stores.

Repeated

Compared to the one-shot design, consumers completed four purchasing rounds. We did not change the basic parameters of the experiment. High-income consumers received 100 Coins each round and could use them to purchase the good in either

store A or store B. Low-income consumers were not able to purchase a good in the first two rounds and had to rely on an outside option that was a simple multiplier of one. Endowments could not be transferred from one round to another. Thus, we can rule out that consumers may respond to a switch in sellers pricing strategies by delaying their purchases in anticipation of future discounts (Coase 1972).

To make switching meaningful, we introduced monetary costs for switching between stores in two consecutive rounds. Consumers always started a round in the store they chose the previous round. For example, if a consumer purchased the good in store A in the first round, they started the second round in store A. If the consumer then decided to switch and purchase the good in store B, they had to pay a fee of 10 Coins, i.e. 10% of their endowment per round. Thus, consumers were monetarily disincentivized to switch and maximized their payoff by always purchasing from the same store.

In the first two purchasing rounds, both stores offered the good for the same price. After the second round, consumers learned that one of the two store managers changed their pricing strategy and would offer low-income consumers the good for a price of 50 Coins in the remaining two rounds. Depending on treatment, this was either (i) the manager of the store consumers had purchased in during the second round (*Avoid*) or (ii) the manager of the store consumers had *not* purchased in during the second round (*Approach*). Price discrimination was transparent to all consumers.

Compared to the one-shot design, this setup imposes stricter conditions on observing consumer preferences against price discrimination. Subjects who switch away from (or towards) the price discriminating store are willing to substantially reduce their own payoff. Second, it allows us to analyze benevolent price discrimination, where consumers who already participate in the market do not experience a price increase and only low-income consumers who priorly could not purchase any goods benefit from lower prices. Whereas in the one-shot design, consumers might perceive the lower price to be at their own costs, the within-subject design establishes the reference price of 100 Coins over two rounds and no consumer loses money as a result of price discrimination. Importantly, low-income consumers miss out on substantial payoff over the first two rounds, and the price discount in rounds three and four serves to mitigate, but not fully eliminate unequal outcomes.

Manager

In each treatment, two participants were assigned the role of a store manager. Before consumers made their choices, managers were free to decide whether they wanted to implement a price discrimination strategy (a discount for low-income consumers) in their store. In the *Repeated* treatments, price discounts could only be introduced after the second round. Managers had full information about the experimental setup. We gathered manager observations until we had one manager who decided for and one manager who decided against price discrimination for each experimental treatment. Hence, no deception was involved. Managers were rewarded with a base reward of \$0.50 and earned 1 Cent for each good sold in their store.

Study 1

In Study 1, information on consumer endowment was private. High-income consumers did neither know about the existence of low-income consumers, nor the endowment of anybody else. They also did not receive any explanation about the price discount. Participants were randomly assigned to *One-Shot*, *Repeated-Avoid* or *Repeated-Approach*. We gathered data until we had 100 high-income consumers in each treatment and set a fixed probability of 10% for a participant to become a low-income consumer.

2.2 Results

We excluded all observations where a subject chose the outside option in the second round. This left us with 96 independent observations in *Avoid*, 97 independent observations in *Approach* and 100 in *One-Shot* (47% female). Our results confirm widespread consumer aversion to price discrimination. Even more so, we find that a large share of consumers is willing to incur costs in order to switch away from a benevolently price-discriminating store without any possibility of subsequent monetary gains.

One-Shot

From now on, we will refer to high-income consumers as “consumers”. Low-income consumers were irrelevant for our analysis and only served to avoid deception. As hypothesized, a large majority of consumers preferred the non-price-discriminating store (*NoPD Store*). Without any explanation for the observed price discrimination, 86% of consumers purchased the good in the store that charged all consumers the

same price, 9% purchased in store with price discounts and 5% chose the outside option. A one sample t -test also confirmed consumers to significantly prefer the non-price-discriminating store ($t = 14.25, p = 0.000$). The name ($\tilde{\chi}^2 = 1.33, p = 0.25$) and position ($\tilde{\chi}^2 = 1.66, p = 0.20$) of the store introducing price discrimination did not have a significant effect on store choice.

Result 1 *Under scarce information, consumers exhibit strong preferences against a price discriminating store.*

Repeated-Avoid

In the first two rounds, both stores offered the good for the same price. After the second round, consumers learned that the store they chose in the second round offered some consumers lower prices. Figure 1 illustrates the two most common behavioral patterns in the *Avoid* treatment.




Figure 1 — The two most common behavioral patterns in *Avoid*

Note. The two most common purchase patterns were “self-interested” and “anti-pd”. Consumers classified as self-interested never switched to the other store. Consumers classified as anti-pd stayed within the same store for the first two rounds and then switched after their store introduced price discrimination.

We find that 47% of consumers never switched and maximized their payoff. We classify those consumer as “self-interested”. Surprisingly, almost as many consumers decided to costly switch away from a benevolently price-discriminating store. We classify those consumers who stayed in the same store for the first two rounds,

switched after the second round, and stayed in the other store for the third and fourth round as “anti-pd”.² Using this conservative measure, 40% of consumers exhibited the corresponding pattern. To support this conjunction, we ran a paired *t*-test to test whether consumers were less likely to purchase in the price discriminating store after it had introduced price discrimination, i.e. in round three and four. On average, consumers purchased 1.91 out of 2 goods in round one and two in the price-discriminating store, and only 1.04 out of 2 in rounds three and four ($t = -8.64, p = 0.000$).

Throughout this paper, we will use a random effects panel logistic regression model with clustered standard errors to confirm and expand on our findings. In accordance with the results above, we report a significant and large effect of the BPD-dummy on a subject’s probability to purchase in the price-discriminating store (see Tables 8 and 9 in the appendix).

Result 2 *Under scarce information, roughly 40% of consumers costly punish a benevolently price-discriminating store.*

Repeated-Approach

In *Approach*, price discrimination was introduced by the store not chosen by a consumer in the second round. This setup was designed to capture positive reactions towards BPD. 80% of consumers followed the self-interested pattern, whereas only 4% switched to the price-discriminating store. A paired *t*-test showed no significant differences in the likelihood that a consumer purchased the good in the price-discriminating store between rounds 1 and 2 versus 3 and 4. On average, consumers purchased 0.07 out of 2 goods in round one and two in the price-discriminating store, and 0.12 out of 2 after the introduction of price discrimination in rounds three and four ($t = 0.96, p = 0.339$).

Result 3 *Consumers are generally not willing to costly support a benevolently price-discriminating store under scarce information.*

²We can rule out egalitarian motives towards the managers’ income distribution as a reason for consumer switching, because we did not tell participants how many rounds they would play. Thus, consumers who wanted both managers to earn the same would have already switched in the second round. A few consumers seemed to exhibit these egalitarian preferences. Text data from an open-ended question in the questionnaire of later studies also confirms the validity of our categorization.

Consumer Attitudes

Table 1 shows subjects' evaluations for a selection of post-experimental questions from *Avoid*.³ There are only few differences between self-interested and anti-pd consumers. Descriptively, those who costly switched away from the price-discriminating store found price discrimination on average fairer (but still very unfair) than those who decided to stay in the same store throughout all four purchasing rounds. Since both groups also did not differ on the Social Comparison Orientation scale, the recorded attitudes suggest that in the no-information context, peer-induced fairness concerns cannot explain adverse consumer reactions towards BPD. While price-discrimination was seen as unfair by almost everybody, perceptions of unfairness do not appear to be a sufficient condition for costly switching.

Instead, consumer behavior might be partially driven by perceptions of intention and exploitation. Subjects categorized as *anti-pd* were less likely to ascribe good intentions to the introduction of price discrimination and thought more often that their manager wanted to take advantage of them ($t = 1.80, p = 0.075$). They also found costly punishment to reduce the manager's income fairer. This interpretation would be in line with attribution theory (Fiske and Taylor 2013), whereby unfair acts by a causal agent can cause the attribution of blame and motivational beliefs.

2.3 Discussion

We show that a large share of consumers costly punishes store managers even when they are not negatively affected by the introduction of price discrimination. Forty percent of consumers exhibit strictly averse behavioral patterns by engaging in inefficient switching and punishing the differential price decrease. In total, the number of goods sold by the price-discriminating store decreases by roughly 45%.

These results cannot be explained by fairness concerns or social comparison. Rather, subjects appear to assign bad intentions to the introduction of price discrimination. One reason might be that subjects in this study received no information about either the goal, the reasons or the beneficiaries of price discrimination. Without any salient point of comparison or justification mechanism, subjects were free in their interpretation of the manager's actions. These were largely judged to be unfair, regardless of the subject's behavior. This interpretation is coherent with a number of studies on

³Because participants in *Approach* did not significantly react towards price discrimination, we refrain from reporting the results here.

Table 1 — Consumer attitudes in *Avoid* for full sample and by behavioral pattern

Short description of questionnaire item	By behavioral pattern				Diff.	t-stat.	
	Full sample		<i>anti-pd</i>				<i>self-interested</i>
	Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)			
Fairness of PD ¹	2.21 (1.55)	2.47 (1.67)	2.11 (1.54)	1.03		1.03	
Fairness of punishment by switching ²	5.81 (1.60)	6.32 (0.93)	5.58 (1.86)	2.22**		2.22**	
Good intentions of PD manager ³	2.68 (1.52)	2.42 (1.43)	2.96 (1.62)	-1.58		-1.58	
Exploitation by PD manager ⁴	4.79 (1.84)	5.18 (1.67)	4.44 (2.01)	1.80*		1.80*	
SCO ⁵	4.50 (1.24)	4.42 (1.17)	4.71 (1.25)	-1.07		-1.07	

Note. All questions were answered on a 7-point scale. Superscripts 3, 4: scale ranged from 1 (unfair) to 7 (fair). Superscripts 1, 2, 5: scale ranged or from 1 (strongly disagree) to 7 (strongly agree); 1stPlease rate how fair you think it is that manager A/B decided to offer some consumers lower prices for the same good.; 2ndThe more goods a store sells, the higher the income of the participant acting as its manager. Do you consider it fair to switch from Store A/B to Store B/A in order to reduce manager A/B's profit because of their pricing strategy?"; 3rdManager A/B's intention for introducing interpersonal price differences was good."; 4thManager A/B intended to take advantage of me (the consumer)."; 5thMean of Social comparison orientation scale by Gibbons and Buunk (1999); The reported *t*-statistic results from a two-sample mean-comparison *t*-tests by behavioral pattern; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

attribution theory, showing that people have a strong tendency to infer causes and assign responsibility, especially when outcomes are perceived to be unfair (Heider 1982; Blount 1995; Falk et al. 2008).

In order to better understand the motives behind consumer aversions towards BPD, we conducted a follow-up study designed to control for manager intentions. If rather than fairness, the ascription of intentions and bad motives drive consumer punishment, we would expect less inefficient switching as a response to price discrimination. Study 2 also controls for negative reciprocity towards the store manager as a reason for switching between stores.

3 Study 2

We eliminated any intentionality behind price discrimination by keeping managers out of the price-setting process. Consumers learned that both managers decided to let a pricing algorithm determine store prices, instead of setting prices themselves. They were told that managers had no information about how the algorithm determined prices and could therefore not predict changes caused by the algorithm. This setup ensures that we isolate the effect of differing intentions: subjects can still punish the manager of the price-discriminating store for bearing responsibility for choosing the algorithm in the first place.⁴ To the extent that subjects are willing to attribute motives, e.g. carelessness, these motives should not differ between both managers since they made the same decision under the same information.

3.1 Experimental Design

To test for the effect of intentionality, we ran two treatments (*Intention* and *No Intention*) of a slightly adjusted version of our (repeated) *Avoid* setup from Study 1.⁵ Whereas *Intention* simply replicated the first study, *No Intention* introduced the pricing algorithms mentioned in the previous paragraph.⁶ One algorithm was randomly selected to

⁴Importantly, this setup ensures meaningful punishment, since switching decreases the income of a related actor, rather than an innocent bystander. The popular choice of substituting the human decision-maker with e.g. a computer or an algorithm would not only render punishment arbitrary, but eliminate a whole host of other variables we cannot account for.

⁵Because consumers appeared to have no preferences for price discrimination in Study 1, we omit a follow-up analysis for *Approach*.

⁶Before launching the main experiment, we gathered two manager observations per treatment. In *No Intention*, both managers decided to delegate price-setting to their respective algorithm.

price discriminate, the other one did not change prices throughout the task. We gathered 200 consumer observations per treatment and extended the post-experimental questionnaire to better evaluate the effectiveness of our manipulation (see the online appendix for an overview).

3.2 Results

We excluded all observations where a subject chose to rely on the outside option in the second round. This left us with 198 independent observations in *Intention* and 193 independent observations in *No_Intention* (52% female). Our results indicate that intentions are not the main driver of consumer aversions against BPD. Across both treatments, a significant share of consumers is willing to costly punish the price discriminating store manager.

Manipulation Check

Our manipulation of intentionality should have two effects: First, when asked about the intentions behind the manager action triggering price discrimination⁷, we expected subjects in *No Intention* to ascribe more “neutral” intentions. Second, since both managers made the exact same decision, we expected no differences in the ascribed intentions between both managers in *No Intention*. Table 2 shows average consumer sentiments towards the price-discriminating manager over the whole sample by treatment.

Subjects in *Intention* evaluated the intention of the price-discriminating store manager to be both significantly less good as well as significantly more bad. For *No Intention*, the averages are closer to the “neutral” point of 4 on the response scale. Furthermore, subjects in *Intention* were more likely to think that the price-discriminating manager had the intention of exploiting them. Thus, our manipulation seems to have worked in the desired direction.

The differences in ascribed intentions between the two managers in *No Intention* are more ambivalent. For all three questions capturing a manager’s intentions, subjects ascribed more positive intentions to the manager whose algorithm did not introduce the lower price for some consumers. That was true when asked whether (i) the

⁷i.e. the intentions behind offering some consumers lower prices in *Intention*, and the intentions behind choosing the pricing algorithm in *No Intention*

Table 2 — Consumer attitudes by experimental condition

Short description of questionnaire item	By experimental condition		
	<i>Intention</i>	<i>No intention</i>	Diff.
	Mean (SD)	Mean (SD)	<i>t</i> -stat.
Fairness of PD ¹	2.27 (1.44)	2.41 (1.55)	-0.90
Good intentions of PD manager ²	2.56 (1.51)	4.47 (1.66)	-11.91***
Bad intentions of PD manager ³	5.31 (1.60)	3.39 (1.82)	11.11***
Exploitation by PD manager ⁴	5.02 (1.75)	3.57 (2.01)	7.63***
Accountability of PD manager ⁵	5.69 (1.70)	3.81 (2.08)	9.78***

Note. All questions were answered on a 7-point scale. Superscript 1: 1 (unfair) to 7 (fair). Superscripts 2–4: 1 (strongly disagree) to 7 (strongly agree). Superscript 5: 1 (not at all) to 7 (very much); ¹“Please rate how fair you think it is that [*Intention*: manager A/B; *No intention*: Store A/B’s algorithm] decided to offer some consumers lower prices for the same good”, ²“Manager A/B’s intention for [*Intention*: introducing interpersonal price differences; *No intention*: choosing the pricing algorithm] was good.”, ³Wording is equal to row 2, except last word is “bad”, ⁴“Manager A/B intended to take advantage of me (the consumer).”, ⁵“How much do you hold Manager A/B accountable for not being offered the good at the lower price?”; The reported *t*-statistic results from a two-sample mean-comparison *t*-tests by behavioral pattern; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

manager’s intention to choose the pricing algorithm was bad (PD: 3.39, NoPD: 2.76; $t = 4.96$, $p = 0.000$), (ii) the manager’s intention to choose the pricing algorithm was good (PD: 4.47, NoPD: 5.11; $t = -5.58$, $p = 0.000$) or (iii) the manager tried to take advantage of the consumer (PD: 3.57, NoPD: 2.34; $t = 8.49$, $p = 0.000$).

However, when compared to *Intention*, the differences between the evaluation of both managers are much smaller. For differences of 0.63, 0.65 and 1.23 in *No Intention*, we record differences of 2.64, 2.47 and 2.64 in *Intention*. Hence we still affirm that our manipulation worked in the desired direction.

Main Results

First, we replicate our results from Study 1. In *Intention* (*No Intention*), 37% (30%) of consumers were willing to costly switch away from a benevolently price-discriminating store, while 45% (49%) acted monetarily self-interested. We further identify a third group of subjects who switched within the first two rounds and settled for the non-price discriminating store afterwards (10%). Since subjects did not know about the number of rounds beforehand, this might be interpreted as egalitarian preferences, i.e. a preference for both managers to earn the same, until price discrimination swayed

them towards one store. The average number of goods sold by the price-discriminating store decreased significantly with the introduction of price discrimination in round three and fourth in both *Intention* ($t = -13.23, p = 0.000$) and *No Intention* ($t = -11.02, p = 0.000$).

Second, we find moderate evidence for small treatment differences. While the share of subjects classified as *anti-pd* drops by seven percentage points in *No Intention*, a *t*-test comparing the average number of goods sold by the price-discriminating store during rounds three and four in *Intention* (0.98) and *No Intention* (1.09) shows no significant difference ($t = 1.1, p = 0.270$). The regression model reveals a significant interaction between the introduction of BPD and *Intention* (see Table 8), suggesting that subjects were less likely to purchase in the price-discriminating store after the introduction of BPD when the decision was made by a store manager instead of a pricing algorithm. Overall, the main results hold irrespective of whether price-discrimination is introduced by the manager themselves or an autonomous pricing algorithm. Our manipulation succeeded in partially neutralizing consumer perceptions of manager intentions and significantly reduced the perceived differences in intentions between the two managers. We take these results as evidence that consumer aversion towards benevolent price-discrimination is largely independent of manager intentions. Despite controlling for intentions, all three behavioral patterns are largely consistent across the two treatments.

Result 4 *Consumer aversion to benevolent price discrimination under information scarcity is not driven by perceived manager intentions.*

Consumer Attitudes

Table 3 shows subjects' attitudes for a selection of post-experimental questionnaire items.

Notably, subjects categorized as *anti-pd* were significantly more neutral in their assessment of the price-discriminating manager's intentions when they learned that the manager had delegated price-setting to a "black-box" algorithm. This again suggests that manager intentions were not the primary driver of costly switching. Subjects in *No Intention* also tended to hold the manager less accountable. However, there were no differences in the perception of fairness regarding the introduction of price discrimination as well as perceived levels of exploitation. Irrespective of

Table 3 — Consumer attitudes by behavioral patterns and experimental condition

<i>Behavioral pattern</i>	By experimental condition		
	<i>Intention</i>	<i>No intention</i>	Diff.
Short description of questionnaire item	Mean (SD)	Mean (SD)	<i>t</i> -stat.
<i>self-interested</i>			
Fairness of PD ¹	2.62 (1.57)	2.51 (1.58)	0.48
Fairness of punishment by switching ²	5.30 (1.64)	4.29 (1.85)	3.92***
Good intentions of PD manager ³	2.89 (1.64)	4.72 (1.57)	-7.76***
Bad intentions of PD manager ⁴	4.99 (1.79)	3.18 (1.69)	7.05***
Exploitation by PD manager ⁵	4.69 (1.91)	3.23 (1.95)	5.12***
Feel exploited by PD manager ⁶	5.08 (1.64)	4.90 (1.78)	0.69
Accountability of PD manager ⁷	5.12 (1.90)	3.00 (1.93)	7.52***
SCO ⁸	4.26 (1.23)	4.27 (1.08)	-0.04
<i>anti-pd</i>			
Fairness of PD ¹	1.82 (1.13)	2.05 (1.47)	-1.01
Fairness of punishment by switching ²	6.57 (0.97)	5.89 (1.63)	2.94***
Good intentions of PD manager ³	2.19 (1.28)	4.11 (1.69)	-7.40***
Bad intentions of PD manager ⁴	5.62 (1.32)	3.88 (1.91)	6.17***
Exploitation by PD manager ⁵	5.54 (1.42)	4.30 (2.04)	4.12***
Feel exploited by PD manager ⁶	6.00 (1.27)	6.02 (1.32)	-0.08
Accountability of PD manager ⁷	6.49 (0.97)	4.95 (1.85)	6.16***
SCO ⁸	4.57 (1.12)	4.31 (1.06)	1.30

Note. All questions were answered on a 7-point scale. Superscript 1: 1 (unfair) to 7 (fair). Superscripts 2–4: 1 (strongly disagree) to 7 (strongly agree). Superscript 5: 1 (not at all) to 7 (very much); ¹“Please rate how fair you think it is that [*Intention*: manager A/B; *No intention*: Store A/B’s algorithm] decided to offer some consumers lower prices for the same good”, ²“The more goods a store sells, the higher the income of the participant acting as its manager. Do you consider it fair to switch from Store A/B to Store B/A in order to reduce manager A/B’s profit because of their pricing strategy?”, ³“Manager A/B’s intention for [*Intention*: introducing interpersonal price differences; *No intention*: choosing the pricing algorithm] was good.”, ⁴Wording is equal to row 2, except last word is “bad”, ⁵“Manager A/B intended to take advantage of me (the consumer).”, ⁶“The fact that [*Intention*: Manager A/B, *No intention*: Store A/B’s algorithm] decided to offer some consumers lower prices makes me feel taken advantage of.”, ⁷“How much do you hold Manager A/B accountable for not being offered the good at the lower price?”, ⁸Mean of Social comparison orientation scale by Gibbons and Buunk (1999); The reported *t*-statistic results from a two-sample mean-comparison *t*-tests by behavioral pattern; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

treatment, subjects judged price discrimination to be very unfair and felt exploited. However, *anti-pd* consumers found price discrimination significantly more unfair than *self-interested* consumers (*Intention*: $t = -3.66, p = 0.000$).⁸ The fact that both groups found price discrimination unfair does suggest, however, that fairness perceptions are neither the sole, nor a sufficient driver of costly punishment. For all other items, the differences between treatments and categories are as expected and in line with Study 1.

3.3 Discussion

We show that consumer aversions to benevolent price discrimination under scarce information persist even when we control for manager intentions. The share of consumers willing to costly switch away from a store that introduces downward price discrimination remains high, and the store sells between 48 % (*Intention*) and 40 % (*No Intention*) less as a result. In sum, the two studies suggest consumer aversion to benevolent price discrimination to be a widespread phenomenon that does not depend on the ascription of bad motives or intentions. Rather, a large share of consumers seems to object to the general concept, at least in a situation of information scarcity. While the prevalence of costly switching may appear high, the finding itself is in line with the existing literature on consumer attitudes towards and perceptions of price discrimination. What happens, however, when we allow for some amount of information on the consumer side? It is likely that consumers often have some information about the rules by which interpersonal price differences are determined. Intuitively, one probable reason for the wide acceptance of student discounts is that students tend to have less money. Consumers might thus be more forgiving to the practice of downward price adjustments, because it decreases inequality. More generally, one might hypothesize that benevolent price discrimination becomes more accepted once consumers are aware who benefits, i.e. consumers with a lower income who otherwise would not be able to participate in the market. If true, transparency could be an efficient way of dealing with negative consumer reactions in the context of price discrimination. To test these assertions, the following studies extend our analysis of consumer behavior under BPD to situations where consumers know how the seller discriminates and who the beneficiaries are.

⁸For *No Intention*, the results are not significant ($t = -1.77, p = 0.079$).

4 Study 3

We extended the experimental design by two crucial features. First, we made the variable of discrimination, i.e. income, transparent to all consumers. Second, we introduced two different causes of income inequality. Endowments were either allocated randomly (*Random*) or based on the performance (*Effort*) in a real-effort slider task (Gill and Prowse 2012). Consumers only knew by which mechanism endowments were allocated, but not the exact endowment levels. That is, consumers in *Effort* knew that their performance would influence their endowment, but were not informed about the exact functional relationship. Consumers received either the high or low endowment based on a fixed performance threshold.⁹

High-income consumers in Study 3 knew that (i) low-income consumers existed, (ii) low-income consumers could not afford to buy a good for the initial uniform price, (iii) only low-income consumers benefited from price discrimination and (iv) high- and low-income consumers could afford the same amount of goods under price discrimination.

Because price discounts do not negatively affect high-income consumer welfare, increase the number of consumers able to participate in the market and increase overall welfare, we expected relatively little switching away from the price discriminating store (*Avoid*). Likewise, we expected some high-income consumers to support the price discriminating store by switching towards it (*Approach*). We also expected high-income consumers in *Effort* to exhibit stronger aversions towards BPD than those in *Random*. That is because high-income consumers might feel entitled to higher payoffs when they are earned by greater effort in the slider task. BPD, in that sense, undermines a merit-based advantage.

4.1 Experimental Design

The experiment consisted of six treatments.

We gathered observations until we had 100 high-income consumers in each treatment with a fixed probability of 10 % to become a low-income consumer in *Random*.

The basic procedure of the experiment mirrored Study 1. In *Random*, we added one page where participants learned about the random endowment mechanism. In

⁹The threshold was 31 sliders over three rounds of 60 seconds each. We followed a pretest with 100 participants and selected the cut-off for the 10th percentile.

Effort, participants first completed three rounds of an effort task based on sliders. After being informed about their endowment, consumers indicated their agreement to two statements on a seven point Likert-scale to validate our manipulation. In both treatments, we asked whether it would be fair if every consumer received the same endowment, and whether it was fair that (1) endowments were allocated randomly or (2) better performances in the slider task were rewarded with a higher endowment. Consumers then proceeded with the original instructions. Participants who completed the effort task but failed to answer the comprehension questions correctly were paid the base reward of \$1.20.


4.2 Results

After excluding all observations where a subject chose the outside option in the second round, we were left with 597 participants (47 % female). Our results show that consumer aversions towards BPD persist under full transparency and are not affected by the cause of income inequality. Even with the information that price discrimination only benefits consumers who otherwise cannot participate in the market, a large share of consumers is willing to engage in costly, inefficient switching to punish the price-discriminating store. Similar to Study 1, almost no consumer costly supports BPD.

One-Shot

In *Effort*, consumers found a hypothetically equal distribution of endowments significantly more unfair than in *Random* (*Effort*: 2.52, *Random*: 6.42; $t = 18.31$, $p = 0.000$). They also found the mechanism by which differences in endowments were achieved significantly fairer (*Effort*: 6.37, *Random*: 4.22; $t = -9.34$, $p = 0.000$). This affirms that subjects in *Effort* perceived differences in income and thereby different purchasing abilities as more fair, presumably because they were rooted in merit.

Regarding store choices, consumers in the one-shot experiment still exhibited strong aversions against price discrimination, albeit weaker than in Study 1. In *Random*, almost 67 % of consumers chose to purchase in the non-price-discriminating store. In *Effort*, the share was around 72 %. The difference was not significant ($t = 1.46$, $p = 0.146$).

Figure 2 — Behavioral patterns across the four conditions from Study 3.*Repeated-Avoid*

Results for the manipulation check are in line with the ones from *One-Shot* and confirm the success of our intervention. In *Avoid*, we largely replicate the results from Study 1 for both *Effort* and *Random* (see Figure 2).

Surprisingly, we find no differences between price discrimination over merit-based compared to arbitrary income inequality. In both treatments, one third of consumers was willing to costly switch to a non-price-discriminating store. Note that this was with the information that only low-income consumers who otherwise were not able to participate in the market benefited from price discrimination, while nobody incurred any income loss. Consumers purchased significantly less goods in the price discriminating store after the introduction of price discrimination (*Effort*: 1.11 vs. 1.86, $t = -7.61, p = 0.000$; *Random*: 1.11 vs. 1.88, $t = -8.22, p = 0.000$). There was no significant treatment difference ($t = -0.10, p = 0.923$) and a random effects logistic regression revealed no significant interaction between *Effort* and the introduction of BPD over the whole sample (see Table 8).

Result 5 *Consumer aversions to BPD persist under full information about the VOD and the beneficiaries.*

Result 6 *Consumers do not differentiate between BPD over merit-based inequality and BPD over arbitrary inequality.*

Repeated-Approach

In accordance with Study 1, almost no consumer was willing to costly support a benevolently price discriminating store. Instead, the majority of consumers exhibited self-interested behavior patterns (*Effort*: 79%, *Random*: 81%). Only four and five percent respectively followed the pro-pd pattern. Paired *t*-tests suggest that consumers did not significantly change their behavior once the other store introduced price discrimination (*Effort*: 0.17 vs. 0.11, $t = 1.10$, $p = 0.275$; *Random*: 0.15 vs. 0.11, $t = 0.68$, $p = 0.495$). There are no significant treatment differences.

Consumer Attitudes

Table 4 shows consumer attitudes of the whole sample as well as segmented by treatment. Differences in *Avoid* point in the expected direction, but are not significant. For *Approach*, despite equal switching behavior, subjects in *Random* perceived price discrimination as significantly fairer and ascribed more positive intentions to the manager of the price discriminating store. These results are in accordance with Studies 1 and 2, whereby feelings of unfairness and the ascription of intentions are not sufficient to explain consumer behavior under BPD.

4.3 Discussion

Consumer aversions towards BPD are not restricted to decision-making under scarce information. Our results indicate that a large share of consumers is willing to costly switch away from a benevolently price-discriminating store, even when price discrimination only benefits low-income consumers. Given that it was, by design, impossible for low-income consumers to earn a higher total payment than high-income consumers, this behavior is largely inconsistent with inequality aversion and the concept of distributive fairness. Although we document a qualitative drop in *anti-pd* subjects compared to the first two studies, equalizing outcomes does not seem to substantially alleviate behavioral constraints towards BPD. As before, net consumer switching is clearly negative for the store that introduces price discrimination. What is more, consumers appear to disregard the cause of income inequality. In contrast to the existing literature on economic inequality and re-distributive preferences, consumer choice under BPD does not depend on whether differences in initial endowments are

Table 4 — Consumer attitudes by behavioral patterns and experimental condition

<i>Experimental condition</i>	By experimental condition		
	<i>Effort</i>	<i>Random</i>	Diff.
Short description of questionnaire item	Mean (SD)	Mean (SD)	<i>t</i> -stat.
<i>Avoid</i>			
Fairness of PD ¹	3.71 (2.09)	4.07 (2.12)	-1.20
Fairness of punishment by switching ²	5.05 (1.88)	4.79 (1.97)	0.95
Good intentions of PD manager ³	4.33 (1.84)	4.38 (1.93)	-0.19
Exploitation by PD manager ⁴	3.25 (1.99)	3.55 (2.10)	-1.04
SCO ⁵	4.38 (1.29)	4.33 (1.42)	0.26
<i>Approach</i>			
Fairness of PD ¹	3.68 (1.93)	4.80 (2.08)	-3.95***
Fairness of switching to increase profit ²	4.47 (1.80)	4.80 (1.85)	-1.28
Good intentions of PD manager ³	4.17 (1.77)	5.18 (1.78)	-4.02***
Exploitation by PD manager ⁴	3.27 (1.76)	2.79 (1.89)	1.86*
SCO ⁵	4.54 (1.21)	4.31 (1.26)	1.33

Note. All questions were answered on a 7-point scale. Superscripts 3, 4: scale ranged from 1 (un-fair) to 7 (fair). Superscripts 1, 2, 5: scale ranged or from 1 (strongly disagree) to 7 (strongly agree); ¹"Please rate how fair you think it is that manager A/B decided to offer consumers who [*Effort*: earned a lower endowment in the slider task, *Random*: randomly received a lower endowment] lower prices for the same good.", ²"The more goods a store sells, the higher the income of the participant acting as its manager. Do you consider it fair to switch from Store A/B to Store B/A in order to [*Avoid*: reduce, *Approach*: Increase] manager [*Avoid*: A/B's, *Approach*: B/A's] profit because of their pricing strategy?", ³"Manager A/B's intention for introducing interpersonal price differences was good.", ⁴"Manager A/B intended to take advantage of me (the consumer).", ⁵Mean of Social comparison orientation scale by Gibbons and Buunk (1999); The reported *t*-statistic results from a two-sample mean-comparison *t*-tests by behavioral pattern; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

earned or arbitrary. One possible explanation might simply be that procedural fairness judgments, i.e. charging different prices for the same good, are much more important in determining consumer behavior than outcome-related reasoning. If consumers do not think about the distribution of economic outcomes but are instead purely focused on the process, it follows that the source of income inequality is a negligible factor. This would also explain why so many consumers are willing to give up a substantial portion of their endowment to decrease the income of a seller who did nothing to reduce their outcome.

An alternative explanation is that the top-down approach by which sellers impose differential pricing reduces both the moral responsibility and agency of consumers in causing more equal outcomes. Because in our setup the “future market” is unaffected by consumer decisions, and consumers do not decide whether low-income consumers receive discounts, they might exhibit patterns that are different from e.g. voting for re-distributive policies. To test this explanation, Study 4 extends our design by a second period to simulate agency through a market environment.

5 Study 4

We examine whether consumer aversions towards BPD as well as consumers’ disregard of arbitrary vs. merit-based inequality generalize to a market-analogous situation where consumers’ purchasing decisions affect a store’s future pricing strategy. By endowing current consumers with agency over the price-setting process for a later period, we induce externalities on their purchasing decisions and thus agency regarding future consumer outcomes. Consumers are responsible for future overall welfare as well as the potential equalization of outcomes induced by BPD. Those who still switch away from the price discriminating store essentially pay to decrease seller and low-income consumer profits as well as to prevent a more equal outcome distribution.

5.1 Experimental Design

We implemented two (*Avoid* vs. *Approach*) 2x2 mixed factorial designs manipulating the source of income inequality (*Random* vs. *Effort*) as well as the agency of consumers regarding the price-discriminating store’s pricing strategy in a future task (*Agency* vs. *No agency*). The *No Agency* treatments were equivalent to the corresponding *Avoid* conditions in Study 3. For the agency treatments, we added one screen after the

store's introduction of price discrimination. On this screen, subjects were informed that the discriminating store's sales in the remaining two rounds would determine its price-setting in a future HIT on MTurk. If sales exceeded a certain threshold, a rule would automatically implement the lower prices for low-income consumers in the last two rounds of a future setup-equivalent task. Otherwise, future low-income consumers would not be able to purchase the good at the reduced price. We thereby foreclosed any considerations regarding differing manager income levels or reverse outcome inequalities.¹⁰

For both *Avoid* and *Approach*, we gathered 150 high-income consumer observations per treatment and set a fixed probability of two percent to become a low-income consumer in *Random*.

5.2 Results Avoid

After dropping subjects who chose the outside option in round two, the final sample consisted of 586 observations (49 % female). Our results largely replicate Study 3, once again establishing the robustness of consumer aversions towards BPD, but also hint that under agency, consumers are more likely to reject BPD when income differences are rooted in merit.


Manipulation Check Avoid

To check that subjects understood our manipulation, we extended the post-experimental questionnaire by a control question about the future impact of their purchase decisions, asking subjects to choose the correct statement out of four options. In total, 76 % of subjects answered the questions correctly. Since there could be multiple reasons why participants might answer the question incorrectly, we do not drop observations based on the manipulation check. Instead, we will use the results from the sub-sample of subjects who correctly answered the question as a robustness check.

Main Results Avoid

In line with Study 3, 58 % (56 %) of subjects in the *No Agency* treatments of *Effort (Random)* decided to stay in the same store for the whole experiment. Similarly, 31 %

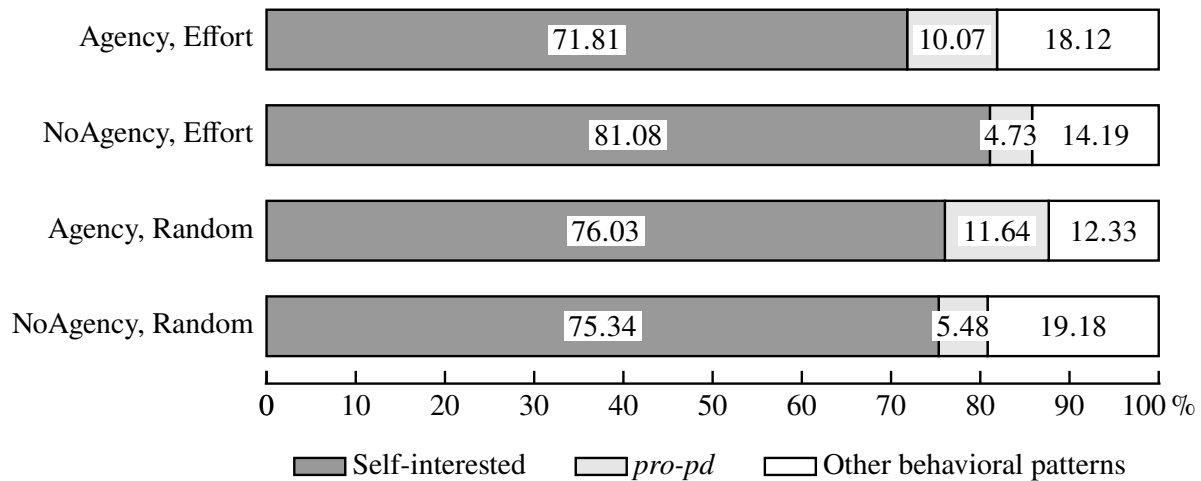
¹⁰Specifying that low-income consumers would receive the discount only for two rounds ensures that it is impossible for them to earn more money in total than regular consumers.

Figure 3 — Behavioral patterns across the four conditions from Study 4, *Avoid*

and 29% respectively exhibited strictly anti price-discrimination behavioral patterns by switching after the introduction of discounted prices. Endowing subjects with agency regarding future consumer welfare does not appear to fundamentally shift behavior (see Figure 3). In *Agency-Effort*, 32% are categorized as *anti-pd*. In *Random*, that share drops to 23%. Contrary to the *No Agency* treatments, we document a significant difference in the drop of sales conditional on the source of endowment differences. In *Random*, the store introducing BPD experienced a significantly smaller decline in demand (*Effort*: -0.72, *Random*: -0.54; $t = -1.56$, $p = 0.06$), which is accentuated when restricting the sample to subjects who correctly answered the control question ($t = -1.89$, $p = 0.03$). The regression model replicates this effect across the whole sample through a significant interaction between the *Effort*-Dummy and the BPD-Dummy. Further and in line with the descriptive results, agency and benevolent price discrimination are positively – but not significantly – correlated in predicting purchases at the price-discriminating store (see Table 8).

Overall, there is moderate evidence that under agency, the source of endowment differences becomes more meaningful and consumers tend to switch less. However, differences are small and only marginally significant, depending on the sample.

Result 7 *Consumer aversions to BPD persist when consumer choices determine future prices.*

Figure 4 — Behavioral patterns across the four conditions from Study 4, *Approach*

As expected given the procedural nature of the intervention, subject answers in the post-experimental questionnaire did not differ substantially (see Table 6 in the appendix).

5.3 Results Approach

Our final sample consisted of 589 (50 % female) observations. 76 % of subjects answered the control question correctly. In *No Agency*, consumers generally do not support BPD by switching towards the store. Once consumers are endowed with some agency, the price-discriminating store sells significantly more goods in rounds three and four. There are no differences between *Effort* and *Random*.

Main Results Approach

The *No Agency* treatments closely replicate our results from Study 3 (see Figure 4). The majority of consumers maximized their income by staying in the same store for all four rounds and the price-discriminating store sold roughly the same amount before and after the introduction of discounts for low-income consumers (*Effort*: $t = 0.29, p = 0.769$; *Random*: $t = 0.27, p = 0.787$).

Once consumers were endowed with some agency regarding future store prices, their purchasing behavior changed significantly. The share of subjects classified as *pro-pd* roughly doubled with 10 % in *Effort-Agency* and 12 % in *Random-Agency*. In

Effort-Agency, the price discriminating store sold 0.09 goods on average over the first two rounds, and 0.30 goods over rounds three and four – a significant increase ($t = 3.89, p = 0.000$). The results for *Random-Agency* are similar (0.08 vs. 0.32; $t = 4.04, p = 0.000$) and there are no significant treatment differences between the two ($t = 0.25, p = 0.802$). The seller introducing BPD after the second round sold significantly more in *Effort-Agency* than in *Effort-No Agency* ($t = -2.94, p = 0.013$) and more in *Random-Agency* than in *Random-No Agency* ($t = -2.21, p = 0.028$). Concurrently, the regression model reveals a significant interaction between BPD and the *Agency* treatment dummy on store choice. The source of endowment does not predict switching.

Result 8 *Endowing consumers with agency over future prices increases costly support for BPD.*

Differences in self-reported consumer attitudes between agency treatments are again minimal (see Table 7 in the appendix).

5.4 Discussion

The results show that a salient market mechanism, which endows consumers with agency over future prices for low-income consumers significantly increases costly support of BPD, while aversions persist. A large share of consumers continues to switch away from the price discriminating store, and there is only little evidence that costly switching might decrease as a function of consumer influence. However, we do find a relatively large effect of agency on consumer switching *towards* the price discriminating store. This effect holds irrespective of the cause for income differences.

We think that there are two main reasons for the increased acceptance and support of BPD. First, some consumers might feel morally obligated to accept or support lower prices for poor consumers once their actions determine future outcomes. Second, introducing a second period allows (or forces) consumers to ascribe themselves some responsibility for changing the outcome of low income consumers. Thus, introducing agency increases the significance of outcome-related reasoning, as consumers are forced to take the future impact of their purchase decisions into account. We deduce that greater consumer agency increases the importance of outcome equalization and distributive fairness for judging BPD, consequently leading to more support and potentially less aversion. The more differences in income are based on merit, the higher the behavioral constraints on BPD.

6 Study 5

Study 5 combines three robustness checks to our main results. First, we increased the salience of low-income consumers and disclosed their endowment of 50 Coins per round to all consumers. Before, consumers only knew that low-income consumers received a lower endowment, they did not know their exact endowment. This information does not change any behavioral predictions regarding distributive fairness and outcome equalization. However, it made it harder for high-income consumers to e.g. strategically ignore that low-income consumers could not purchase anything before the introduction of BPD. It also highlighted that price discrimination merely led to a convergence in payoff – low-income consumers remained considerably poorer than high-income consumers. Moreover, while consumers were treated unequally with respect to absolute prices, they were treated equally in that everybody was priced relative to their endowment.

Second, we varied switching costs to assess how sensitive consumer choices were towards the ease of switching. In reality, this could for instance be determined by the purchase environment (online vs. analog), the level of market competition or the availability of substitutes.

Third, we summarize and analyse text data from an open-ended question on consumer behavior to demonstrate data quality, confirm the validity of our behavioral labels and highlight some alternative consumer reasoning. We show that consumer animosity towards BPD exceeds quantifiable switching rates and that economic self-interest disciplines inefficient consumer reactions.

6.1 Robustness Check: Salience

We implemented a 2 (*Approach vs. Avoid*) × 2 (*Effort vs. Random*) mixed-factorial design that replicated the agency-treatments from Study 4 with additional information on the income of low-income consumers as well as their ability to purchase exactly one good per round after the introduction of price discounts. In such a framework, we would expect maximal impact of outcome-related fairness concerns. We collected 150 independent observations per treatment. Dropping all subjects who chose the outside option in the second round left us with 592 (51 % female) total subjects.

Results

Results closely follow the findings from Study 4 (see Table 5).

Consumer still exhibited substantial aversions towards BPD while the majority followed a self-interested behavioral pattern. There was both significant consumer emigration from (*Avoid*), as well as migration to (*Approach*) the price discriminating store (see also Tables 5 and 8 in the appendix). However, differences for average number of goods sold in *Avoid* (*Effort*: -0.67 ; *Random*: -0.53) are much larger than in *Approach* (*Effort*: 0.31 ; *Random*: 0.25). Hence, the price discriminating store still experienced a net loss of consumers. Looking only at the 78% who correctly answered the control question does not change the results. We thus affirm that our main results are not dependent on subjects' failure to understand the implications for low-income consumers or a wrong attribution of differential procedural treatment between consumers regarding the variable of discrimination.

6.2 Robustness Check: Switching Costs

To assess the influence of switching costs on purchasing behavior, we opted for the *Avoid* paradigm where subjects earned their endowment through the effort task and had agency over the store's future prices. We implemented three treatments with switching costs of either 5, 15 or 30 Coins. From 125 independent observations per treatment, we dropped 8, leaving us with 367 independent observations in total (54% female).

Results

Results reveal moderate evidence for downward elasticity of consumer BPD aversions. Increasing switching costs does not affect behavior (see Tables 5 and 8). For switching costs of 5 Coins (5% of a participant's endowment), the drop in average goods sold after the introduction of price discounts was significantly higher than for switching costs of 15 (-0.91 vs. -0.70 ; $t = -1.71$, $p = 0.045$) and weakly significantly higher than for switching costs of 30 Coins (-0.91 vs. -0.74 ; $t = -1.40$, $p = 0.082$). There was no difference between 15 and 30 Coins ($t = 0.30$, $p = 0.619$).¹¹ Finally, the regression

¹¹Restricting the sample to subjects who correctly answered the control question about consumer purchases affecting store prices in a future HIT (75%) amplifies these results for switching costs of 15 (5: -0.82 vs. 15: -0.51 ; $t = -2.30$, $p = 0.011$), but not for switching costs of 30 (5: -0.82 vs. 30: -0.65 ; $t = -1.23$, $p = 0.110$).

analysis reveals no significant treatment dummy (Table 8 in the appendix), confirming that consumer aversions towards BPD are relatively inelastic to changes in switching costs.

Table 5 — Behavioral patterns from the robustness checks

<i>Robustness check</i> Treatment	Behavioral pattern (%)	
	<i>anti-pd</i> in <i>Avoid</i> <i>pro-pd</i> in <i>Approach</i>	<i>self-interested</i>
<i>Salience</i>		
Avoid, Random	19.59	62.16
Approach, Random	10.00	65.33
Avoid, Effort	28.57	58.50
Approach, Effort	14.97	76.19
<i>Switching costs</i>		
5 Coins	38.21	43.09
15 Coins	31.40	53.72
30 Coins	33.33	51.22

Note. Results are consistent with the main studies. Aversions towards BPD are substantially stronger than respective support, and behavior is relatively inelastic to changes in switching costs.

6.3 Robustness Check: Consumer Reasoning

The post-experimental questionnaire of Studies 2, 4, and 5 included an open-ended question that allowed participants to explain why they “did (not) decide to switch stores after [PD-Store] offered some consumers lower prices for the same good” in their own words.¹²

First, the data largely confirmed the authenticity of subject feedback and thus validates the quality of our data. A large majority of participants offered straightforward reasoning for their choices which demonstrated their understanding of the task, the payoff scheme and the actors involved. Each participant wrote an average of 114

¹²To draw conclusions from these data, we opted for an inductive and manual coding scheme to identify different response labels (coding framework). Further, we used the independent coder method to heighten the reliability of our coding. Results from three coders were compared for consistency and amended by the authors if needed.

characters, and across all considered studies, only 136 out of 3229 stated reasons (4%) were independently judged to be nonsensical.¹³

Second, subject answers confirmed that our interpretation of the two predominant behavioral patterns as *self-interested* and *anti-pd/pro-pd* was overwhelmingly in line with subject reasoning. From 590 identified reasons that were stated by subjects categorized as *anti-pd* in *Avoid*, the majority (334) explicitly stated that they switched because they were snubbed by the price-discriminating store. The second most cited reason was that consumers wanted to punish the manager who introduced price discounts (73). For the 1160 reasons from subjects labeled as *self-interested* in *Avoid*, the two main recurring explanations were switching costs (444) and a lack of economic benefits to switching (389). Thus, it appears that economic considerations disciplined many consumers to abstain from inefficient switching. The analysis further highlighted that a significant portion of consumers “silently” supported BPD (165), which could not be captured by the *Avoid* paradigm. All these 165 subjects participated in studies where consumers were able to influence the price setting in a future market and knew about low-income consumers as well as income as the variable of discrimination (Study 4 and 5).

For *Approach*, the most prominent reason by subjects classified as *pro-pd* (104 identified reasons) was support for the price-discriminating store (51). Beyond that, 14 answers mentioned their agency over future prices as well as the option to equally split manager profits over the four rounds. Consumers who stayed in the same store over all rounds (880 reasons) predominantly mentioned that there were either no benefits to switching (459), that switching was costly (225) or that they felt loyalty towards the other store (66).

Overall, the data supports our prior interpretations while underlining the importance of some crucial design choices like switching costs and the consideration of both switching directions.

7 General Discussion

In this paper, we introduced the concept of *benevolent price discrimination* as a downward-bound implementation of differential pricing that always benefits financially disadvan-

¹³We extracted up to two reason from each subject’s answer. In case of two reasons, each received equal weight in the analysis. For the full data set, please refer to the web appendix.

tagged groups or individuals and thus leads to more equal economic outcomes while no consumer incurs any losses. Across five studies, we show that a large share of consumers rejects BPD and pays to switch to a competitor that treats all consumers equally. Effects are the strongest when the benevolent nature of price discounts is obfuscated. However, even when consumers know that everybody is priced according to the same rule, and that differences in endowment are purely due to random chance, the main results hold. This behavior is not driven by reciprocity, intentions, a lack of transparency or a perception of merited outcome inequality. Costly support for BPD is rare, and largely contingent on high levels of transparency as well as salient agency over a store's future prices. While consumers do move bidirectionally, net migration for the BPD-store is consistently negative. Thus, distributive fairness of realized economic outcomes appears to play a relatively small role in determining consumer reactions towards differential pricing.

Our results partially contrast well-researched phenomena from the literature on social preferences and inequity aversion. First, there is good evidence to suggest that consumers should be expected to support BPD as a less invasive analog to other re-distributive policies. Second, many consumers do not differentiate whether pricing equalizes outcome differences based on luck or based on merit, despite being similarly salient as in other studies on social preferences relating to e.g. economic redistribution (Alesina and Giuliano 2011), inequity aversion (Engelmann and Strobel 2004), the fair process effect (Van den Bos, Vermunt, et al. 1997) or ultimatum bargaining (Camerer 2011). Under price discrimination, effects on the overall distribution of economic welfare appear to be uniquely minuscule. Our results support the interpretation that this is partially due to consumers being second movers and thus unable to ascribe themselves agency. The introduction of a second period increases costly support for BPD, while decreasing consumer aversion if income is allocated by a random mechanism as opposed to an effort task. Hence, caring about other consumer's payoff appears to be partially dependent on one's own direct influence. Put differently, second-mover social preferences differ from first-mover social preferences.

One reason why price discrimination might elicit unique counter reaction lies in people's general aversion towards demand-driven pricing changes. We think that many people neither judge prices, nor price-setting as something endogenously determined by supply and demand, but rather something that should reflect the costs of a good. This has also been proposed by Kahneman et al. (1986) in their seminal work on price fairness perceptions. Setting different prices for the same good might be viewed as

uniquely undesirable by relatively adversely affected consumers, since (i) there is no objective reason why prices should differ in any way and (ii) price discounts signal that the “true” value of a product lies below what has been charged before.

Implications

While some studies have found that obfuscating interpersonal price differences reduces consumer fairness concerns and thereby increases sellers’ pricing power (see e.g. Allender et al. 2016), other researchers have suggested overt transparency to signal benevolence if price discrimination is based on consumer income (Rotemberg 2011). Our results indicate that the latter is, by and large, not necessarily a good strategy for sellers. At the very least, firms should expect a drop in demand by current consumers, which could become quite large and almost certainly exceeds the attraction of new consumers who would not benefit from price discounts. Given that price discrimination in our study is downward-bound, further segmentation by increasing prices would presumably culminate in even larger counter reactions. Instead, our findings are more in line with results from Li and Jain (2016). Firms that are expecting consumer fairness concerns reduce differential pricing, e.g. low poaching prices, which reduces inefficient switching and consumer disutility from perceived price unfairness, but hurts overall consumer surplus. Thus, consumer reactions induce constraints for price interventions that increase market participation as well as joint welfare. We conclude that seller hesitation to deviate from uniform pricing appears to be a well-adapted strategic reaction towards substantial behavioral constraints that cannot be alleviated by equalizing unequal outcome distributions and transparently targeting low-income consumers.

Limitations and Future Research

Our results support the conjunction that contextual and procedural elements inherent to the price setting framework play a considerable role for consumer behavior and, in our setting, clearly outweigh outcome-related reasoning. Nevertheless, contextual and procedural elements may naturally differ between markets or even be strategically influenced by sellers, possibly tipping the scales in favor of outcome-related reasoning. For instance, our results cannot explain the widely observed acceptance of price discounts for specific low-income groups. Since the experiments were explicitly designed to be context-neutral, effects elicited by e.g. social norms, cultural norms,

inertia or habituation are beyond the scope of this paper. They might, however, moderate the influence of outcome reasoning on consumer behavior. It is also possible that people are more accepting of different prices for population groups they (i) have been part of in the past (e.g. people below 25) or (ii) will be part of in the future (e.g. people above 65).

One explanation for our results may be that consumers expect prices to correlate with costs, which is violated by demand-based price discrimination. If that is the case, the introduction of price *increases* for high-income consumers might counter-intuitively *reduce* behavioral constraints to price discrimination, because they signal a potential necessity to compensate price discounts, rather than relatively large profit margins at the original price. One hidden implication is that consumers could react differently to price discrimination over physical products than price discrimination over services. Presumably, it is much harder to define the “true” value of a service than that of a product, as the former’s cost function more saliently captures intangibles like opportunity costs, rather than costs of production. This question, however, requires a less abstract approach to the quantification of consumer behavior.

Finally, we were able to identify consumers’ agency over prices and thereby outcomes, as one relevant procedural element that interacts with outcome-driven behavioral changes. Although consumers’ agency was already quite salient in our study, there may be other interventions that are more effective in increasing perceptions of agency. First, a logical extension of our setup is to implement a multi-round dynamic market environment in which consumers can “experience” their agency over market outcomes. Other interventions may increase agency by involving consumers more directly in the price setting process, e.g. by some form of voluntary extra payment or “Pay-as-You-Wish” pricing (Chen et al. 2017). Here, we hope that our results help future researchers in exploring innovative pricing mechanisms that better fit the social and economic preferences of consumers while avoiding the pitfalls of antagonizing consumers against one another. Consumer social and fairness preferences are a demonstrably complicated phenomenon, influenced by a multitude of procedural and distributive variables, that often translate into non-obvious behavioral patterns. Without careful considerations, instruments of differential pricing might provoke consumer reactions with at least ambiguous, if not outright negative welfare effects. An increasing public recognition of price discrimination could be detrimental to consumer welfare, which should caution policy calls for more transparency.

Acknowledgments

The authors gratefully acknowledge support from the “Ministry of Science and Arts” of Lower Saxony (Germany), the Volkswagen Foundation (grant number ZN3492, "Vorab"), and the Center for Digital Innovations (ZDIN).

Author Contributions

Conceptualization: A.E., N.E., and M.B.; Data curation: A.E., M.B.; Formal analysis: A.E., M.B.; Investigation: A.E., N.E., and M.B. Methodology: A.E., N.E., and M.B.; Project administration: A.E.; Supervision: none; Validation: A.E., N.E., and M.B.; Visualization: M.B.; Writing—original draft: A.E.; Writing—review editing: A.E., N.E., and M.B.

References

- Alesina, A. and G.-M. Angeletos (2005). "Fairness and Redistribution". In: *American Economic Review* 95.4, pp. 960–980.
- Alesina, A. and P. Giuliano (2011). "Preferences for redistribution". In: *Handbook of social economics*. Vol. 1. Elsevier, pp. 93–131.
- Alexopoulos, J., D. M. Pfabigan, F. Göschl, H. Bauer, and F. P. S. Fischmeister (2013). "Agency matters! Social preferences in the three-person ultimatum game". In: *Frontiers in human neuroscience* 7, p. 312.
- Allender, W. J., J. Liaukonyte, and T. J. Richards (2016). "Strategic Obfuscation and Price Fairness". In: *SSRN Electronic Journal*.
- Anderson, E. T. and D. I. Simester (2008). "Research note—does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes". In: *Marketing Science* 27.3, pp. 492–500.
- (2010). "Price stickiness and customer antagonism". In: *The quarterly journal of economics* 125.2, pp. 729–765.
- Bayer, R.-C. (2010). "Intertemporal price discrimination and competition". In: *Journal of Economic Behavior & Organization* 73.2, pp. 273–293.
- Blount, S. (1995). "When social outcomes aren't fair: The effect of causal attributions on preferences". In: *Organizational behavior and human decision processes* 63.2, pp. 131–144.
- Bolton, G. E., J. Brandts, and A. Ockenfels (2005). "Fair procedures: Evidence from games involving lotteries". In: *The Economic Journal* 115.506, pp. 1054–1076.
- Bolton, G. E. and A. Ockenfels (2000). "ERC: A theory of equity, reciprocity, and competition". In: *American economic review* 90.1, pp. 166–193.
- Bolton, L. E., H. T. Keh, and J. W. Alba (2010). "How Do Price Fairness Perceptions Differ across Culture?" In: *Journal of Marketing Research* 47.3, pp. 564–576.
- Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- Campbell, M. C. (1999). "Perceptions of Price Unfairness: Antecedents and Consequences". In: *Journal of Marketing Research* 36.2, pp. 187–199.
- Chen, D. L., M. Schonger, and C. Wickens (2016). "oTree—An open-source platform for laboratory, online, and field experiments". In: *Journal of Behavioral and Experimental Finance* 9, pp. 88–97.

- Chen, Y., O. Koenigsberg, and Z. J. Zhang (2017). "Pay-as-You-Wish Pricing". In: *Marketing Science* 36.5, pp. 780–791.
- Choshen-Hillel, S. and I. Yaniv (2011). "Agency and the construction of social preference: Between inequality aversion and prosocial behavior." In: *Journal of personality and social psychology* 101.6, p. 1253.
- Coase, R. H. (1972). "Durability and monopoly". In: *The Journal of Law and Economics* 15.1, pp. 143–149.
- Dholakia, U. M. and I. Simonson (2005). "The effect of explicit reference points on consumer choice and online bidding behavior". In: *Marketing Science* 24.2, pp. 206–217.
- Durante, R., L. Putterman, and J. van der Weele (2014). "Preferences for Redistribution and Perception of Fairness: An Experimental Study". In: *Journal of the European Economic Association* 12.4, pp. 1059–1086.
- Engelmann, D. and M. Strobel (2004). "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments". In: *American economic review* 94.4, pp. 857–869.
- Fabiani, S., M. Druant, I. Hernando, C. Kwapil, B. Landau, C. Loupias, et al. (2006). "What Firms' Surveys Tell Us about Price-Setting Behavior in the Euro Area". In: *International Journal of Central Banking*.
- Falk, A., E. Fehr, and U. Fischbacher (2008). "Testing theories of fairness—Intentions matter". In: *Games and Economic Behavior* 62.1, pp. 287–303.
- Fehr, E. and K. M. Schmidt (1999). "A Theory of Fairness, Competition, and Cooperation". In: *The Quarterly Journal of Economics* 114.3, pp. 817–868.
- Fehr, E. and S. Gächter (2000). "Fairness and retaliation: The economics of reciprocity". In: *Journal of economic perspectives* 14.3, pp. 159–181.
- Fehr, E., G. Kirchsteiger, and A. Riedl (1993). "Does fairness prevent market clearing? An experimental investigation". In: *The quarterly journal of economics* 108.2, pp. 437–459.
- Ferguson, J. L., P. S. Ellen, and W. O. Bearden (2014). "Procedural and Distributive Fairness: Determinants of Overall Price Fairness". In: *Journal of Business Ethics* 121.2, pp. 217–231.
- Fiske, S. T. and S. E. Taylor (2013). *Social cognition: From brains to culture*. Sage.
- Folger, R. and M. A. Konovsky (1989). "Effects of procedural and distributive justice on reactions to pay raise decisions". In: *Academy of Management journal* 32.1, pp. 115–130.

- Fudenberg, D. and J. M. Villas-Boas (2012). "Price discrimination in the digital economy". In: *The Oxford handbook of the digital economy*, p. 254.
- Garbarino, E. and S. Maxwell (2010). "Consumer response to norm-breaking pricing events in e-commerce". In: *Journal of Business Research* 63.9-10, pp. 1066–1072.
- Gibbons, F. X. and B. P. Buunk (1999). "Individual differences in social comparison: Development of a scale of social comparison orientation." In: *Journal of Personality and Social Psychology* 76.1, pp. 129–142.
- Gill, D. and V. Prowse (2012). "A Structural Analysis of Disappointment Aversion in a Real Effort Competition". In: *American Economic Review* 102.1, pp. 469–503.
- Guo, L. (2015). "Inequity aversion and fair selling". In: *Journal of Marketing Research* 52.1, pp. 77–89.
- Harrington Jr, J. E., R. H. Gonzalez, and P. Kujal (2016). "The relative efficacy of price announcements and express communication for collusion: Experimental findings". In: *Journal of Economic Behavior & Organization* 128, pp. 251–264.
- Haws, K. L. and W. O. Bearden (2006). "Dynamic Pricing and Consumer Fairness Perceptions". In: *Journal of Consumer Research* 33.3, pp. 304–311.
- Heider, F. (1982). *The psychology of interpersonal relations*. Psychology Press.
- Herz, H. and D. Taubinsky (2018). "What makes a price fair? An experimental study of transaction experience and endogenous fairness views". In: *Journal of the European Economic Association* 16.2, pp. 316–352.
- Ho, T.-H. and X. Su (2009). "Peer-Induced Fairness in Games". In: *American Economic Review* 99.5, pp. 2022–2049.
- Ho, T.-H., X. Su, and Y. Wu (2014). "Distributional and Peer-Induced Fairness in Supply Chain Contract Design". In: *Production and Operations Management* 23.2, pp. 161–175.
- Janiszewski, C. and D. R. Lichtenstein (1999). "A range theory account of price perception". In: *Journal of Consumer Research* 25.4, pp. 353–368.
- Jin, L., Y. He, and Y. Zhang (2014). "How Power States Influence Consumers' Perceptions of Price Unfairness". In: *Journal of Consumer Research* 40.5, pp. 818–833.
- Kahneman, D., J. L. Knetsch, and R. Thaler (1986). "Fairness as a constraint on profit seeking: Entitlements in the market". In: *The American economic review*, pp. 728–741.
- Koszegi, B. and M. Rabin (2006). "A Model of Reference-Dependent Preferences". In: *The Quarterly Journal of Economics* 121.4, pp. 1133–1165.
- Krawczyk, M. and F. Le Lec (2010). "'Give me a chance!' An experiment in social decision under risk". In: *Experimental economics* 13.4, pp. 500–511.

- Kuo, A., D. H. Rice, and P. Fennell (2016). "How fitting! The influence of fence-context fit on price discrimination fairness". In: *Journal of Business Research* 69.8, pp. 2634–2640.
- Leibbrandt, A. (2020). "Behavioral constraints on price discrimination: Experimental evidence on pricing and customer antagonism". In: *European Economic Review* 121, p. 103303.
- Lerner, M. J. and L. A. Whitehead (1980). "Procedural justice viewed in the context of justice motive theory". In: *Justice and social interaction: Experimental and theoretical contributions from psychological research*, pp. 219–256.
- Li, K. J. and S. Jain (2016). "Behavior-based pricing: An analysis of the impact of peer-induced fairness". In: *Management Science* 62.9, pp. 2705–2721.
- Maxwell, S. (2007). *The price is wrong: Understanding what makes a price seem fair and the true cost of unfair pricing*. John Wiley & Sons.
- Mayhew, G. E. and R. S. Winer (1992). "An empirical analysis of internal and external reference prices using scanner data". In: *Journal of consumer Research* 19.1, pp. 62–70.
- Nilssen, T. (1992). "Two kinds of consumer switching costs". In: *The RAND Journal of Economics*, pp. 579–589.
- Pigors, M. and B. Rockenbach (2016). "Consumer social responsibility". In: *Management Science* 62.11, pp. 3123–3137.
- Rotemberg, J. J. (2011). "Fair Pricing". In: *Journal of the European Economic Association* 9.5, pp. 952–981.
- Rutström, E. and M. B. Williams (2000). "Entitlements and fairness: an experimental study of distributive preferences". In: *Journal of Economic Behavior & Organization* 43.1, pp. 75–89.
- Rutte, C. G. and D. M. Messick (1995). "An integrated model of perceived unfairness in organizations". In: *Social Justice Research* 8.3, pp. 239–261.
- Shor, M. and R. L. Oliver (2006). "Price discrimination through online couponing: Impact on likelihood of purchase and profitability". In: *Journal of Economic Psychology* 27.3, pp. 423–440.
- Thaler, R. (1985). "Mental accounting and consumer choice". In: *Marketing science* 4.3, pp. 199–214.
- Trautmann, S. T. (2009). "A tractable model of process fairness under risk". In: *Journal of Economic Psychology* 30.5, pp. 803–813.

- Van den Bos, K., E. A. Lind, R. Vermunt, and H. A. Wilke (1997). "How do I judge my outcome when I do not know the outcome of others? The psychology of the fair process effect." In: *Journal of personality and social psychology* 72.5, p. 1034.
- Van den Bos, K., R. Vermunt, and H. A. Wilke (1997). "Procedural and distributive justice: What is fair depends more on what comes first than on what comes next." In: *Journal of Personality and Social Psychology* 72.1, p. 95.
- Wu, C.-C., Y.-F. Liu, Y.-J. Chen, and C.-J. Wang (2012). "Consumer responses to price discrimination: Discriminating bases, inequality status, and information disclosure timing influences". In: *Journal of Business Research* 65.1, pp. 106–116.
- Xia, L., K. B. Monroe, and J. L. Cox (2004). "The Price is Unfair! A Conceptual Framework of Price Fairness Perceptions". In: *Journal of Marketing* 68.4, pp. 1–15.

Appendix

<i>Experimental condition</i>	By experimental condition		
	<i>Agency</i>	<i>No agency</i>	<i>Diff.</i>
Short description of questionnaire item	Mean (SD)	Mean (SD)	<i>t</i> -stat.
<i>Effort</i>			
Fairness of PD ¹	3.79 (1.96)	3.41 (1.90)	1.68*
Fairness of punishment by switching ²	5.21 (1.68)	5.03 (1.87)	0.89
Good intentions of PD manager ³	4.24 (1.78)	4.19 (1.97)	0.23
Bad intentions of PD manager ⁴	3.65 (1.84)	3.75 (2.04)	-0.44
Exploitation by PD manager ⁵	3.66 (2.00)	3.47 (2.08)	0.81
Feel exploited by PD manager ⁶	3.82 (2.07)	3.90 (2.07)	-0.35
SCO ⁷	4.52 (1.24)	4.47 (1.11)	0.34
<i>Random</i>			
Fairness of PD ¹	4.26 (1.88)	4.26 (2.07)	0.00
Fairness of punishment by switching ²	4.78 (1.81)	4.91 (1.97)	-0.57
Good intentions of PD manager ³	4.62 (1.75)	4.75 (1.80)	-0.63
Bad intentions of PD manager ⁴	3.30 (1.83)	3.09 (1.91)	0.99
Exploitation by PD manager ⁵	3.55 (1.95)	3.26 (2.09)	1.26
Feel exploited by PD manager ⁶	3.72 (1.99)	3.68 (2.18)	0.16
SCO ⁷	4.70 (1.05)	4.33 (1.19)	2.81***

Table 6 — Consumer attitudes by experimental condition (Study 4, *Avoid*)

Note. All questions were answered on a 7-point scale. Superscript 1: 1 (unfair) to 7 (fair). Superscripts 2–4: 1 (strongly disagree) to 7 (strongly agree). Superscript 5: 1 (not at all) to 7 (very much); ¹“Please rate how fair you think it is that manager A/B decided to offer consumers who [*Effort*: earned a lower endowment in the slider task, *Random*: randomly received a lower endowment] lower prices for the same good.”, ²“The more goods a store sells, the higher the income of the participant acting as its manager. Do you consider it fair to switch from Store A/B to Store B/A in order to reduce manager A/B’s profit because of their pricing strategy?”, ³“Manager A/B’s intention for [*Intention*: introducing interpersonal price differences; *No intention*: choosing the pricing algorithm] was good.”, ⁴Wording is equal to row 2, except last word is “bad”, ⁵“Manager A/B intended to take advantage of me (the consumer).”, ⁶“The fact that [*Intention*: Manager A/B, *No intention*: Store A/B’s algorithm] decided to offer some consumers lower prices makes me feel taken advantage of.”, ⁷Mean of Social comparison orientation scale by Gibbons and Buunk (1999); The reported *t*-statistic results from a two-sample mean-comparison *t*-tests by behavioral pattern; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

<i>Experimental condition</i>	By experimental condition		
	<i>Agency</i>	<i>No agency</i>	<i>Diff.</i>
Short description of questionnaire item	Mean (SD)	Mean (SD)	<i>t</i> -stat.
<i>Effort</i>			
Fairness of PD ¹	3.94 (2.02)	4.03 (1.88)	-0.42
Fairness of switching to increase profit ²	4.71 (1.65)	4.55 (1.79)	0.82
Good intentions of PD manager ³	4.50 (1.66)	4.57 (1.66)	-0.40
Bad intentions of PD manager ⁴	3.52 (1.84)	3.28 (1.71)	1.16
Exploitation by PD manager ⁵	3.32 (1.82)	2.93 (1.79)	1.86*
Feel exploited by PD manager ⁶	3.91 (2.02)	3.25 (1.91)	2.90***
SCO ⁷	4.67 (1.08)	4.63 (1.11)	0.29
<i>Random</i>			
Fairness of PD ¹	4.42 (1.84)	4.66 (1.85)	-1.11
Fairness of switching to increase profit ²	4.66 (1.77)	4.50 (1.69)	0.78
Good intentions of PD manager ³	4.79 (1.55)	4.93 (1.60)	-0.78
Bad intentions of PD manager ⁴	3.10 (1.66)	2.97 (1.70)	0.70
Exploitation by PD manager ⁵	3.25 (1.81)	2.81 (1.80)	2.11**
Feel exploited by PD manager ⁶	3.66 (1.87)	3.30 (1.90)	1.64
SCO ⁷	4.66 (1.09)	4.66 (1.17)	-0.02

Table 7 — Consumer attitudes by experimental condition (Study 4, *Approach*)

Note. All questions were answered on a 7-point scale. Superscript 1: 1 (unfair) to 7 (fair). Superscripts 2–4: 1 (strongly disagree) to 7 (strongly agree). Superscript 5: 1 (not at all) to 7 (very much); ¹“Please rate how fair you think it is that manager A/B decided to offer consumers who [*Effort*: earned a lower endowment in the slider task, *Random*: randomly received a lower endowment] lower prices for the same good.”, ²“The more goods a store sells, the higher the income of the participant acting as its manager. Do you consider it fair to switch from Store A/B to Store B/A in order to reduce manager A/B’s profit because of their pricing strategy?”, ³“Manager A/B’s intention for [*Intention*: introducing interpersonal price differences; *No intention*: choosing the pricing algorithm] was good.”, ⁴Wording is equal to row 2, except last word is “bad”, ⁵“Manager A/B intended to take advantage of me (the consumer).”, ⁶“The fact that [*Intention*: Manager A/B, *No intention*: Store A/B’s algorithm] decided to offer some consumers lower prices makes me feel taken advantage of.”, ⁷Mean of Social comparison orientation scale by Gibbons and Buunk (1999); The reported *t*-statistic results from a two-sample mean-comparison *t*-tests by behavioral pattern; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	Study 1	Study 2	Study 3	Study 4a	Study 5a	Study 5b
	Coef. (se)	Coef. (se)	Coef. (se)	Coef. (se)	Coef. (se)	Coef. (se)
PD	-5.353*** (1.206)	-3.935*** (0.459)	-4.700*** (0.806)	-3.994*** (0.587)	-4.136*** (0.674)	-5.273*** (0.726)
PD × <i>Intention</i>		-1.543** (0.617)				
PD × <i>Effort</i>			0.507 (0.852)	-1.300* (0.754)	-1.074 (0.852)	
PD × <i>Agency</i>				0.380 (0.739)		
PD × <i>Effort</i> × <i>Agency</i>				0.351 (1.075)		
PD × 15 Coins						0.598 (0.849)
PD × 30 Coins						0.732 (0.812)
<i>Intention</i>		0.969* (0.520)				
<i>Effort</i>			-0.412 (0.731)	1.119 (0.683)	0.506 (0.772)	
<i>Agency</i>				0.270 (0.646)		
<i>Effort</i> × <i>Agency</i>				-1.050 (0.951)		
15 Coins						0.504 (0.731)
30 Coins						0.096 (0.683)
SCO	0.513* (0.304)	-0.237 (0.172)	-0.086 (0.199)	-0.375*** (0.144)	0.316 (0.216)	-0.117 (0.165)
Constant	3.383** (1.584)	5.582*** (0.965)	5.647*** (1.263)	6.907*** (0.958)	4.754*** (1.279)	5.758*** (0.992)
N	384	1564	788	2344	1180	1468
AIC	292.194	1267.133	634.181	1734.051	819.050	1129.807
BIC	307.996	1299.263	662.198	1791.647	849.489	1172.140

Table 8 — Panel logistic regression using random effects for *Avoid*

Note. Table reports results of panel logistic regressions using random effects and a cluster-robust VCE estimator. Dependent variable is a dummy variable that equals 1 if participant bought a good in the price discriminating store, 0 otherwise — Independent variables: “PD” equals 1 in Round 2 and 3, 0 otherwise; *Intention*, *Effort*, *Agency*, 15 ECU and 30 ECU are dummy variables for the experimental conditions; Reference groups are omitted from the table; “SCO” is mean of Social comparison orientation scale by Gibbons and Buunk (1999) — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	Study 1	Study 3	Study 4b	Study 5a
	Coef. (se)	Coef. (se)	Coef. (se)	Coef. (se)
PD	0.636 (0.666)	0.409 (0.586)	0.120 (0.437)	1.689*** (0.411)
PD × <i>Effort</i>		0.184 (0.774)	0.026 (0.654)	1.619** (0.770)
PD × <i>Agency</i>			1.923*** (0.650)	
PD × <i>Effort</i> × <i>Agency</i>			-0.258 (0.913)	
<i>Effort</i>		-0.012 (0.568)	-0.384 (0.448)	-2.031*** (0.705)
<i>Agency</i>			-0.992** (0.503)	
<i>Effort</i> × <i>Agency</i>			0.527 (0.709)	
SCO	-0.308 (0.239)	-0.076 (0.193)	-0.064 (0.121)	0.410** (0.193)
Constant	-2.731** (1.229)	-3.709*** (1.037)	-3.420*** (0.675)	-6.070*** (1.140)
N	388	800	2356	1188
AIC	153.150	381.708	1246.513	728.275
BIC	168.994	409.816	1304.160	758.755

Table 9 — Panel logistic regression using random effects for *Approach*

Note. Table reports results of panel logistic regressions using random effects and a cluster-robust VCE estimator. Dependent variable is a dummy variable that equals 1 if participant bought a good in the price discriminating store, 0 otherwise — Independent variables: “PD” equals 1 in Round 2 and 3, 0 otherwise; *Effort* and *Agency* are dummy variables for the experimental conditions; Reference groups are omitted from the table; “SCO” is mean of Social comparison orientation scale by Gibbons and Buunk (1999) — * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

CHAPTER IV

The Virtual Online Supermarket

Nils Engelbrecht, Tim-Benjamin Lembcke, Alfred Benedikt Brendel, Kilian Bizer, and
Lutz M. Kolbe

It is controversially discussed if and which interventions policymakers should implement to promote healthier, more sustainable, and more ethical food choices. Often, policy measures suffer from a lack of data. This is especially true for the growing field of online grocery shopping. Yet, it is not always feasible to test the impact of each possible policy intervention in the field. Here, computer-simulated shopping experiments offer a complementary approach. Recent evidence suggests that they heighten the realism of consumer experiments and collect valid data at a relatively low cost. In this paper, we introduce an open-source tool-set that offers multiple avenues to develop and run experiments in the context of online grocery shopping. Hence, it supports researchers and policymakers in evaluating instore-intervention aiming to support more sustainable food choices.

Keywords Online Supermarket · E-Grocer · Experimental Consumer Research · Research Platform

Published as Engelbrecht, N.; Lembcke, T.-B.; Brendel, A.B.; Bizer, K.; Kolbe, L.M.: *The Virtual Online Supermarket - An Open-Source Research Platform for Experimental Consumer Research*. *Sustainability* 2021, 13, 4375. <https://doi.org/10.3390/su13084375>.

1 Introduction

Today, more and more governments, NGOs, and players from the food industry have started to follow the United Nations' call to transition to a more healthy and sustainable food system (Hertwich 2010). In Western countries particularly, empirical evidence emphasizes the negative impact of the predominant dietary patterns on individuals' health, the environment, and society (Willett et al. 2019). Hence, policy measures limited to the supply side seem insufficient to trigger a sustainable transition of the food system (Hoek et al. 2017). Consequently, there is a controversial debate about what public and private institutions should or should not promote "better food choices" (Guthrie et al. 2015; Reisch et al. 2013).

These days, most decisions about food still take place in traditional brick-and-mortar supermarkets. However, driven by the development and diffusion of new communication technologies, grocery shopping is undergoing a change in the 21st century (Pitts et al. 2018). As one major element of this change, online grocery shopping is becoming an increasingly more important retail channel, especially in urban centers (Anesbury et al. 2016). Thus, in-store interventions in traditional brick-and-mortar stores and online supermarkets are crucial instruments for policymakers aiming to alter consumers' sustainable food choices (Hartmann-Boyce et al. 2018). Here, they can draw on a variety of different intervention types ranging from economic interventions (e.g., taxes) to changes in a store's micro-environment (e.g., choice architecture techniques; see Münscher et al. 2016 for an overview).

Due to this variety, evidence-based policy-making has a constant need for data to identify the right intervention type or a mix of interventions for the respective case. For instance, due to a lack of empirical evidence, the effectiveness of changes in a store's micro-environment to promote sustainable food choices is still questioned (Hollands et al. 2013; Hummel and Maedche 2019). Furthermore, little is known about the extent to which findings from traditional brick-and-mortar stores can be transferred to online supermarkets (Anesbury et al. 2016; van Herpen et al. 2016). For instance, there is initial evidence that online supermarkets should not be regarded as perfect mirrors of their real-world equivalents. While some elements of the shopping environment like shelf placement strategies seem to be a relevant factor for both online and "offline" supermarkets (van Herpen 2016; van Herpen and Bosmans 2018), both channels differ in aspects like (i) in product presentation, e.g., physical vs. virtual (Huyghe et al. 2017), (ii) navigation pathways (Santos et al. 2012), or interpersonal interactions (Ogonowski

et al. 2014; Goldfarb et al. 2015), Moreover, compared to physical contexts, online environments allow an easier, faster, and more flexible integration of different design choices and interventions and provide enhanced functionalities like decision support systems (Schneider et al. 2018).

Hence, as part of a transition to a sustainable food system, it is mandatory to gain further insights into the determinants of consumers' in-store behavior patterns and food choices in online supermarkets and traditional brick-and-mortar stores. As one of the first, a survey study in Poland analyzed determinants and barriers of organic online shopping (Bryła 2018). Only in this way, it will be possible to evaluate the current legal requirements (e.g., packing information) on their effectiveness in analog and digital choice environments and modify them if necessary. Nevertheless, empirical insights can support policymakers in developing, testing, and adopting new policies to govern sustainable food choices (OECD 2016; OECD 2017).

As it is not always feasible to run studies in actual (online) supermarkets, researchers have recently started to conduct studies in simulated virtual supermarkets (Epstein et al. 2015; Forwood et al. 2015; Demarque et al. 2015). Such an approach has the potential to heighten the realism level of consumer experiments and allow researchers to collect valid purchase data at a relatively low cost.

To make computer-simulated shopping experiments as accessible as possible to interested researchers, we developed an open-source, modular, and highly customizable virtual online supermarket application, called VOS. The application allows researchers to easily implement and perform experiments in the context of (online) grocery shopping. Thus, it can help to develop and evaluate policy measures aiming to support more sustainable food choices. All it takes is a server computer (e.g., a cloud server) to host the experiment and participants with access to a device using a modern web browser.

The tool's front-end was designed to emulate the store design and functions (e.g., navigation tools) of a realistic online grocer environment. In addition, the back-end of the tool allows researchers to modify the research conditions and to configure and implement different experimental treatments. For universal access, the project's source code, python scripts for automated treatment administration, and configuration snippets for local and server hosting are available on GitHub [27]. Moreover, user documentation, set-up instructions, and sample data are available in conjunction with the repository. Everyone should feel invited to use a VOS for academic purposes

and improve our application and submit it via GitHub (<https://github.com/Kuiter/vegs-repo>). However, we ask to cite this paper if the VOS is used for your publication.

The technical barriers for using our application were lowered so the product database and a couple of research conditions can be edited without any programming knowledge by using a visual administration interface (VAI). For instance, researchers can use predefined modification options (“Use Cases”; UC) to adjust food prices, implement different labeling strategies, or change the arrangement of food items. The aim was to make a preselection of modification options covering the broadest possible spectrum of varying research interests.

Hence, use cases range from traditional economic instruments (e.g., taxes) over knowledge-based interventions (e.g., labeling strategies) to choice architecture and decision support tools. Researchers familiar with programming in “*Angular*” are not limited to basic features. Instead, they can extend or change any aspect of the tool’s visual or functional implementation by altering the program’s code. In total, our application allows researchers to test a broad spectrum of interventions on consumers’ food purchases (i) at relatively low cost, (ii) without a complex implementation process, and (iii) without having to collaborate with a specific retailer. In addition to evaluating such interventions based on outcome variables like purchase data, the tool offers extensive possibilities for recording subjects’ in-store behavior dynamics. This way, a VOS generates both general information about users’ in-store behavior (e.g., the average use of filters) and detailed information about a single user’s “journey” (e.g., single user’s navigation pathway) during the shopping trip. In this paper, we will first introduce the development, features, and implementation of our application. With this, we pay particular attention to describing opportunities for practical implementation. Secondly, we will present the results of an initial evaluation study on the tool’s functionality and realism level.

2 VOS in a Nutshell: Features and Functions

The VOS application was designed to enable researchers to conduct experiments in a realistic online shopping environment. It contains two core elements: The shop view and a visual administration interface (VAI). For more information about the technology and programming of the VOS application, see Appendix A.

2.1 Shop View

The shop view mimics the appearance and functionality of a realistic online supermarket. The design is not prohibitive and provides users with the services expected of an online store. Users can browse for products without any (visible) restrictions, e.g., by clicking on categories and subcategories or using a search bar for word searches (Figure 1). Participants see a list of all food items belonging to the selected category (or subcategory) on category pages. In this list, product names, product images, price, the price per unit (e.g., EUR 1.00/kilogram), and the package sizes (e.g., 100 g) are displayed. Researchers are free to decide which currency and quantity units (e.g., gram vs. ounces) are used for the shop. Clicking on a food item takes consumers to a product page, where more detailed information on the product is available (e.g., nutritional information). Further, our tool's base version provides category, attribute, and open text filter options for filtering items. In addition, the items can be sorted via ascending or descending prices, and a virtual shopping cart (VSC) lets users add items to and delete them from their cart.

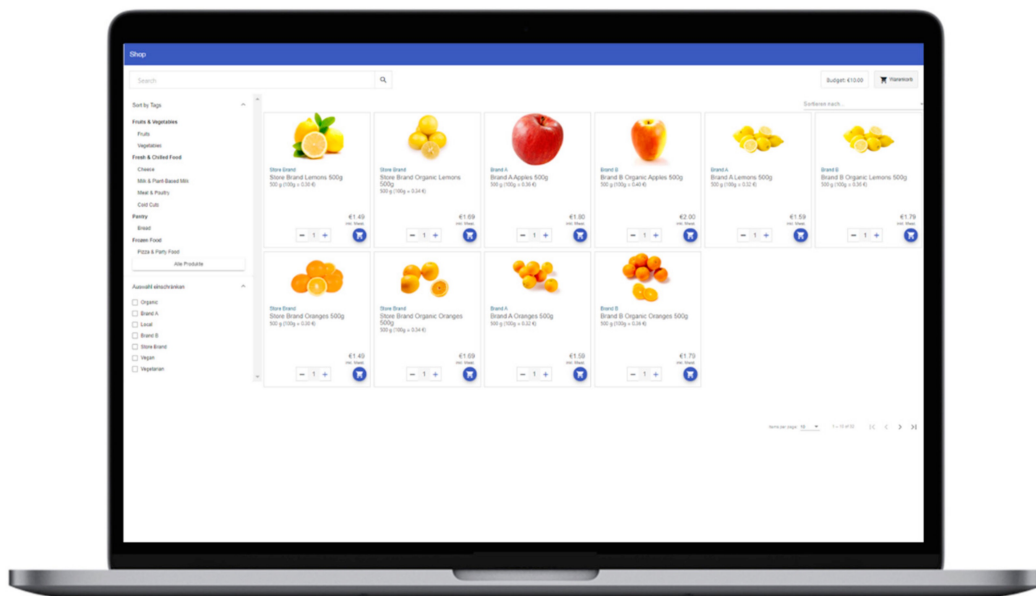


Figure 1 — Screenshot of the shop view.

In the tool's base version, the front-end design is very neutral to prevent any existing customer relationship bias towards the store's design elements. However, as the application code is open-source, the design can be customized to suit researchers'

individual needs. The same applies to most of the elements, functions, and information displayed in the shop. As explained in the next section, researchers can use the visual administration interface's support for some of these modifications.

2.2 Visual Administration Interface (VAI)

The VAI is a visual subdomain where each authenticated user (administrator) has their own workspace (Figure 2). It provides the means for managing the product database and various modification options that allow researchers to determine which information or functions the shop view presents. In the following, we refer to these predefined modification options as "Use Cases" because researchers can use and modify them to create their own experimental treatments. For instance, without any mandatory knowledge in programming, it is possible to create or edit the items, filter mechanisms, taxes, labels, scores, configure VCS's functions, and implement swap interventions. Swap options offer consumers the opportunity to replace a selected food item, for instance, with a healthier or more sustainable one (Forwood et al. 2015). Furthermore, researchers can limit participants' shopping budget for each treatment. This might be particularly relevant for spending tracking experiments (van Ittersum et al. 2010). In this manner, it can be safeguarded that subjects invest reasonable effort into the shopping task and do not merely click through it. Here, just one click is necessary for immediately testing all created treatments in a demo mode.

Via the VAI, it is also possible to administrate participants and to conduct experiments. Here, one can easily create (personal) links for all participants. These links refer participants to a unique URL, which presents the SV modified to the assigned treatment conditions. On this webpage, the VOS records participants' behavior automatically and saves it on the employed server. Hence, a VOS solely requires the link, a connection to the hosting server, and a modern web browser for participating in the experiment.

2.3 Preconfigured Modification Options ("Use Cases")

The application's base version includes five use cases (UC.1 to UC.5). These cover four intervention types: (i) economic interventions (UC.1 taxes and subsidies), (ii) changes to the store's microenvironment (UC.2 product arrangement), (iii) knowledge-based interventions (UC.3 product labeling and scores), and (iv) two decision-support tools

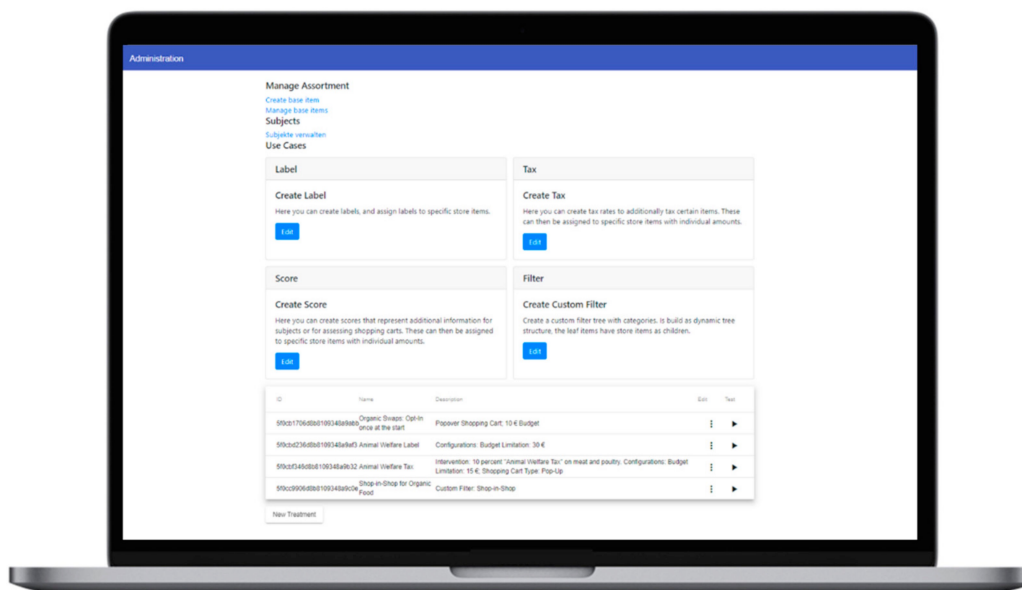


Figure 2 — Screenshot of the visual administration interface.

(UC.4 different types of VSCs and UC.5 swap options). All use cases were selected and designed based on research findings on in-store interventions in (online) supermarkets (e.g., Hartmann-Boyce et al. 2018; Liberato et al. 2014; Cameron et al. 2016; Bianchi et al. 2018; Cadario and Chandon 2020). In addition, we conducted expert interviews with researchers from different disciplines to identify and prioritize their expectations and requirements for a research tool.

Use Case One: Taxes and Subsidies

In many cases, the price of a food product is the central criterion underlying the purchase decision. Hence, it does not seem very surprising that several studies analyze how economic interventions alter consumers' food choices, for instance, by applying taxes, subsidies, or other monetary incentives (e.g., discounts and food coupons) on selected food items (see Thow et al. 2014 for a review).

Overall, there is empirical evidence that monetary incentives effectively alter consumers' food purchases and consumption (e.g., Hartmann-Boyce et al. 2018; An 2012). With this, salience seems to be an important determinant for the effectiveness of a tax or a subsidy (Chetty et al. 2009). This might be particularly relevant for online supermarkets as individuals tend to shift towards more shallow forms of information

processing in digital environments such as browsing, scanning, or skimming (Loh and Kanai 2016; Mangen et al. 2013; Liu 2005). Hence, an individual's actual food choice in an online supermarket could be based more on a superficial first visual impression like the final price than consciously considering different internal and external aspects like health, price, or convenience (Scheibehenne et al. 2007). Accordingly, the effect of taxes in an online supermarket could also depend on their visual presentation and salience.

Consequently, the need for empirical evidence for evidence-based policy-making prevails for at least two reasons: first, food prices remain a politically controversial topic, especially against the background of debates on sustainability and animal welfare (e.g., Bonnet et al. 2020). Second, little particular research on the effects of taxes in online supermarkets exists.

To this end, a VOS offers researchers an easy way to create their own or implement existing taxation models (e.g., meat tax) and evaluate their effect on consumers in a realistic online shopping environment. When defining an experimental treatment, researchers can configure whether and for which products taxes should be displayed and charged in the shop. Additionally, further information on the tax can be provided. Consumers will find this information on the product page of each item affected by the tax (Figure 3).

Here, the VOS automatically tracks whether a consumer has retrieved this information or not. Despite being labeled "taxes", the function is not limited to taxation. Instead, it can also be used for other pricing mechanisms and strategies like subsidies, discount models, or dynamic pricing. In particular, the latter approach is coming into focus as recent technological advances have opened up unprecedented opportunities for retailers to implement interpersonal pricing strategies in a more flexible and prevalent manner (e.g., based on consumers' individual characteristics).

Use Case Two: Product Arrangement

Previous research has demonstrated that even marginal and seemingly irrelevant changes in a store's micro-environment can alter individuals' food choices in a predictable way, e.g., a more "prominent" positioning of a food item [42]. This resonates in the well-known phrase "Eye-Level Is Buy-Level". Indeed, the impact of these and other shelf placement strategies on consumers' purchases are extensively documented in marketing literature (e.g., van Herpen and Bosmans 2018). However, whether these

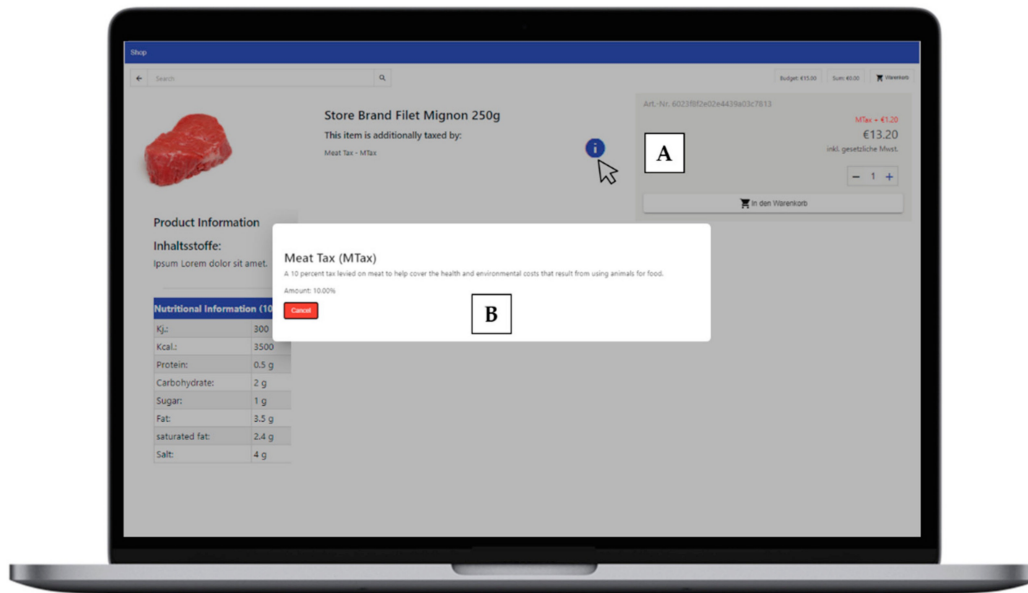


Figure 3 — Shop view - (A) tax display in food details component and (B) tax information dialog, for showing the description of the additional tax.

interventions can also be successfully used to promote better food choices has not been completely evaluated (Keller et al. 2015).

Even if product presentation methods differ between traditional brick-and-mortar stores and online supermarkets, (digital) shelf placement seems to remain a relevant factor for online stores. Overall, initial findings suggest that small differences in assortment organization affect consumers' behavior similarly when tested in an online store (van Herpen et al. 2016; van Herpen and Bosmans 2018). However, there are still a couple of unresolved research questions that concern both physical and online supermarkets. For instance, little is known about how product bundling strategies (e.g., Organic Box Schemes) or store-in-store concepts affect consumers' sustainable food choices. Even if some shelf management strategies show similar effects on consumers' behavior for both retail channels, it is much easier to change or even customize shelf or product arrangements within an online supermarket (Kummer and Milestad 2020; Cheung et al. 2016). Hence, future research may benefit from a systematic evaluation of how customized digital shelf placement strategies affect consumers' food choices.

Accordingly, findings obtained in computer-simulated shopping experiments can provide generalizable insights. Thus, the VOS is suitable for both: studies aiming to gather general findings on how organizing an assortment affects food choices

and analyzing particularities of absolute and relative display locations in online supermarkets. For product sorting and bundling, the application offers two kinds of filter types: Filters based on tags (e.g., name of a product category) and filters based on attributes (e.g., organic). For each item, tags and attributes can be defined directly in the product database (Figure 4).

	A	B	C	D
1	external_ID	Name	Tags	Niceness
2	100	Store Brand Lemons 500g	['Fruits & Vegetables', 'Fruits']	1
3	200	Store Brand Organic Lemons 500g	['Fruits & Vegetables', 'Fruits']	1
4	300	Brand A Lemons 500g	['Fruits & Vegetables', 'Fruits']	0.5
5	400	Brand B Organic 500g	['Fruits & Vegetables', 'Fruits']	0

Figure 4 — Filter definition via the data model.

In this case, we created a filter tree for tags with one parent category and one child category, displayed automatically on the left-hand side of the shop view. In the same manner, a filter based on items' base attributes will be automatically implemented. This can be used to limit the selection, e.g., only including items with the trait “organic”. Figure 5 shows a possible implementation.

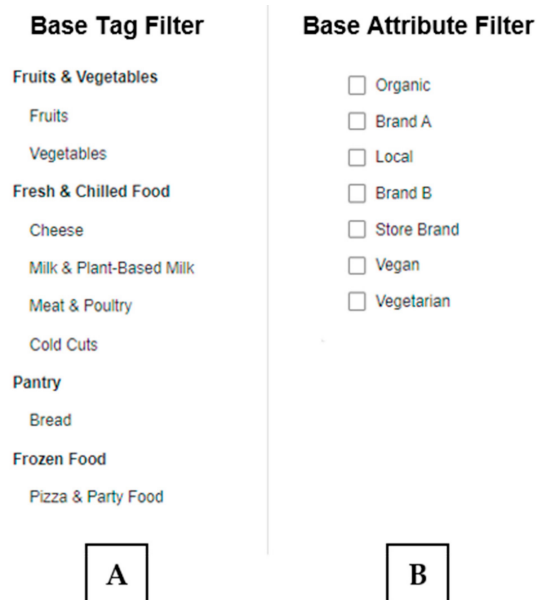


Figure 5 — Shop view—(A) base tag filter and (B) base attribute filter.

The base attribute filter is a multiple-selection whereas the tag filter requires a single choice. Alternatively, researchers can create filter trees in the VAI. This way, there are no limits to the filter tree's customizability and depth (see the VOS User Guide for further instructions on custom filtering). Further, researchers can use an additional sorting mechanism: "niceness.". Niceness is an attribute on the item level, represented by a numeric value between zero and one. Based on this attribute, items from the database are sorted, beginning with the least nice items. This ordering mechanism allows individual researchers to define which items always are displayed first (see Figure 4).

Using these functions, researchers can edit (i) how products are arranged into categories and subcategories, (ii) how these categories are titled in the store, and (iii) which products are displayed in which order. Furthermore, they can use it to implement shop-in-shop concepts (e.g., sustainable food substore).

Use Case Three: Product Labels and Scores

Besides interventions geared towards human affection, cognitively-oriented interventions aiming to transfer knowledge and information remain widely used in food policies (Cadario and Chandon 2020). Prior research has illustrated that grocery shoppers do not pay much attention to detailed product information like nutritional information (Cheung et al. 2016; Benn et al. 2015) but base their decisions on key information that is readily available and processable (Schwarz 2004; Shah and Oppenheimer 2007). Consequently, simplified and salient front-of-package label formats may facilitate consumers' processing of nutritional information and product attributes (e.g., organic) at the point of sale (Weinrich and Spiller 2016).

Based on these empirical insights, many different food labels have been developed, empirically tested, and in some cases, launched (Newman et al. 2016; Campos et al. 2011; Grunert et al. 2014; van Kleef and Dagevos 2015). However, the empirical evidence on the effectiveness of labeling strategies to promote more sustainable and healthier food choices is overall inconclusive and inconsistent on the question of which label format works best (Cadario and Chandon 2020; Sanjari et al. 2017). Consequently, there are numerous starting points for further research, such as (i) evaluating new labeling strategies (e.g., for animal welfare-friendly standards), (ii) analyzing the effectiveness of existing labels in visual online shops, or (iii) exploring the impact of the above-mentioned new technological possibilities in online supermarkets.

A VOS can support this type of research by allowing researchers to create a new binary or multi-level food labeling schemes and scores. These can be assigned to specific store items directly in the interface. Further, legal definitions or other explanations can be provided for each claim on the food label or score (Figure 6). Moreover, all product information and the nutrition facts labels (nutrition information panel) for each item in the database can be edited via the VAI.

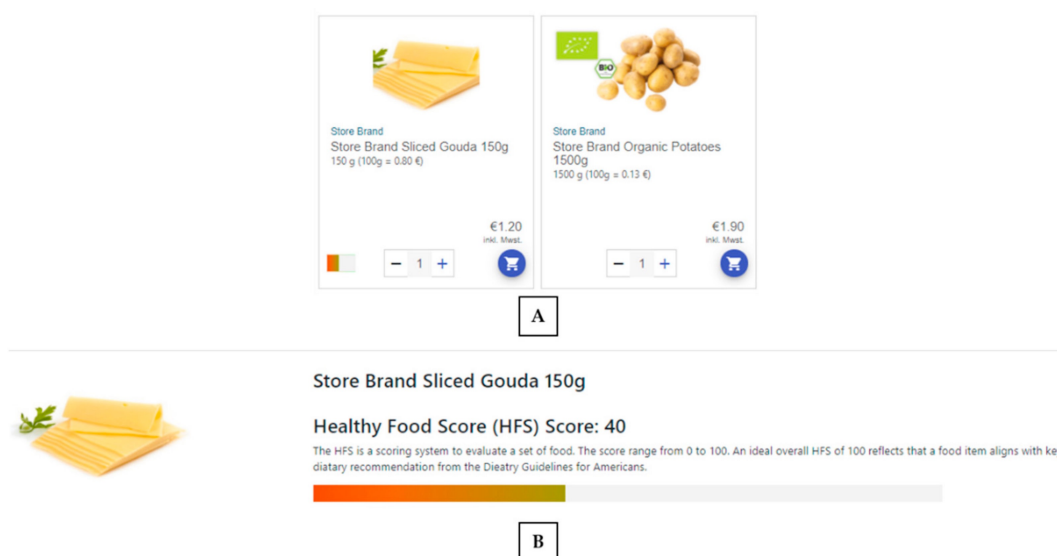


Figure 6 — .Shop view — (A) exemplary score display (left) and food labeling scheme (right) in the food-card component and (B) score display in the food-details component.

Shoppers can access this information with just one click in shop view, and it is automatically recorded whether they have done so or not. This might be interesting because empirical evidence points out that some consumers willfully ignore information if it helps them avoid an inner conflict (e.g., animal welfare concerns vs. meat consumption; Onwezen and van der Weele, Cor N 2016). Moreover, those labels that guarantee only low sustainability standards or animal welfare might bias consumers to overrate the actual product quality in favor of these two aspects (Etilé and Teyssier 2016). This has been referred to as the label halo effect (Küst 2019).

Use Case Four: Virtual Shopping Carts

In contrast to traditional brick-and-mortar stores, consumers cannot physically interact with a salesperson in an online store. To compensate for this, vendors have already started implementing a wide variety of decision-support tools within their online shops. These tools aim to assist shoppers with their purchase; for instance, by providing information on demand (e.g., via a search-bar or a chatbot), (ii) providing real-time feedback (e.g., about spending), or (iii) present personalized product recommendations. According to Häubl and Trifts 2000, the way in which consumers search for product information and make purchase decisions is always a result of the sum of all the single interactions with different decision-support tools available in an online shopping environment.

One decision-support tool and at the same time an integral part of every online store is a VSC. Even the most simplistic VSC enables consumers to accumulate all want-to-buy products in a list and provide real-time feedback about the price of their goods. In contrast to a shopping trip without decision support (e.g., in a physical store), this can increase consumers' total spending and the spending for higher-priced, hedonic, or organic products (van Ittersum et al. 2013; Lembcke et al. 2020). Moreover, VSC's lower transaction costs make it relatively convenient for consumers to adjust their current shopping cart by adding (or removing) single items or making changes in product quantity at any time and without much effort (Chintagunta et al. 2012). However, to the best of our knowledge, no research exists that analyzes the impact of different VSC designs on consumers' food choices and in-store behavior.

To capture this, VOS allows researchers to choose between three different types of VSCs: (i) "Icon Only", (ii) "Icon Plus", and (iii) "Pop-Up" (Figure 7). They vary in how many "clicks" a user has to invest in reviewing their present spending and adding, removing, or replacing items.

For the first type (Icon Only), only a small dynamic cart icon is displayed in the upper right corner of the header on category pages and product pages. This icon provides three basic functionalities: (i) notifies them when a new item is added to the shopping cart, (ii) shows the number of total items in the cart, and (iii) provides a link to a separate fullpage cart. Only on the full-page cart can consumers find all the details of their transactions and edit their cart before continuing to shop or proceeding to the checkout page. The second type (Icon Plus) is identical to the previous version, except that the dynamic cart icon additionally shows the total spending. The third type (Pop-

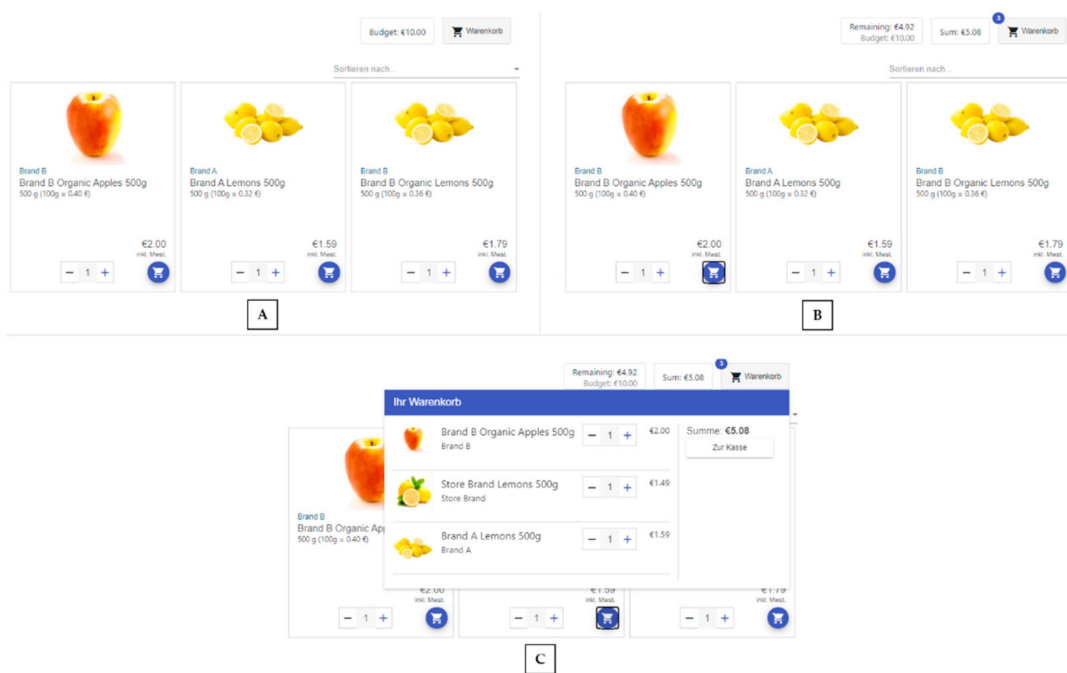


Figure 7 — Shop view — overview about different types of virtual shopping carts (VSCs) available in the application’s base version. (A) Icon Only, (B) Icon Plus, and (C) Pop-Up.

Up) includes a mini shopping cart, which will be displayed if users hover their mouse above the shopping cart icon. It is designed to provide users with a compact version of the main shopping cart page with all information and functionalities while keeping them on the product pages to continue shopping. Via the treatment configuration functions in the admin view, researchers can choose which VSC type they want to implement for which experimental treatment.

Use Case Five: Swap Options

We decided to include swaps as a second use case for decision-support tools because they represent a highly transparent and, at the same time, effective recommendation agent. On the one hand, users can easily identify when, where (type-transparency), how, and for what purpose (token-transparency) swap interventions were used (Hansen et al. 2016; Barton and Grüne-Yanoff 2015). On the other hand, various empirical studies document a statically significant effect of swap interventions on consumers' food purchases in real and virtual supermarket settings (Eyles et al. 2017; Riches et al. 2019; Koutoukidis et al. 2019).

However, there are still many open questions on the impact of swap interventions (Koutoukidis et al. 2019). For instance, it is unknown which product categories and attributes consumers are more likely to accept swaps. In particular, there is no study to examine how effective this type of intervention is in promoting the sale of products with credence attributes like animal welfare or sustainability. Moreover, more evidence is needed to determine if swap interventions are more effective when presented immediately (e.g., after clicking on the "Add to Cart" button) or when bundled at checkout (Forwood et al. 2015). The same applies to the degree of freedom of choice that consumers should be given. For instance, should swaps options be automatically displayed per default, or should consumers be obliged to determine if and when swaps are displayed (forced-choice)?

To address these and similar research questions, VOS provides swap options for food items at different points of the shopping trip: (i) when adding items to the shopping cart or (ii) when finishing the shopping trip by checking out (see Figure 8).

In each case, a pop-up notification informs participants that a swap option is available for the chosen product(s) (swap dialog). In addition, it can also be left for participants to decide whether they want to receive swap options or not (forced-choice). In this case, the swap dialog does not automatically provide a swap option. Instead,

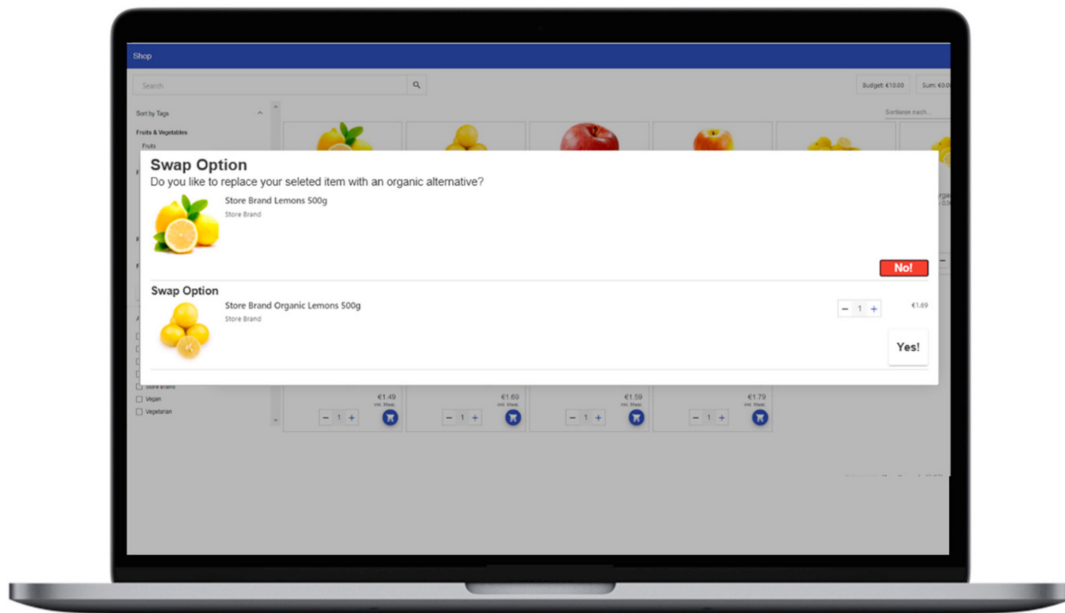


Figure 8 — Shop view — example of a swap dialog.

participants can determine whether they want to (i) see swap options if available or (ii) not to receive future notifications about swap options. Again, researchers can configure these aspects directly in the VAI. Unfortunately, so far, the VAI cannot define which swap options should be shown for which item. Instead, swaps must be assigned in the original product database on which the application is based. You can find more information about this procedure in the VOS User Guide.

Data Recording: Behavioral Outputs and Instore Behavioral Dynamics

The data's scope automatically recorded by the application includes many outcome variables that cover participants' purchases by default. However, by focusing on these outputs, in-store decision-making happens largely within a black box (Hui et al. 2009; Sheehan and van Ittersum 2018). Hence, researchers interested in more in-depth insights about subjects' shopping also need data on in-store behavioral dynamics. However, this kind of data is often not available or has to be reconstructed laboriously, e.g., by using complex procedures like video screen-capturing (Hui et al. 2013).

The VOS records several actions taken by a subject during an experimental shopping trip automatically in code to counteract this. This includes the navigation path (routing), filtering or sorting options used, all actions taken on a page (pagination

events), and all additions, changes, and removals to and from the shopping cart. This data can later be converted into variables and thus used for statistical analysis. In this manner, we receive (among others) information about a shopper's (i) shopping duration, (ii) first orientation time on the website, (iii) the add-purchase-ratio for items in the VSC, or (iv) the number of detailed product views. Furthermore, some use case-specific data is recorded. For example, for "Sustainable Swap Interventions", data is recorded on (i) which item triggered a swap dialogue, (ii) the time when a swap dialogue starts/ends and (iii) whether a swap option offered was accepted or rejected. As each action has a unique timestamp, instore-behavior data can be used for aggregate-level analyses and in-depth analysis like tracking a single user's behavior on the website. All data can be downloaded from the server in JavaScript Object Notation format and reformatted in any number of ways. Example scripts for extracting and reformatting the data can be found in the VOS User Guide.

In addition to behavioral data, it is possible to obtain further self-reported data by including configurable questionnaires. These can be positioned before and/or after the actual shopping task. Thus, they can function as screening questionnaires, comprehension checks, manipulation checks, or post-experimental questionnaires.

3 Using a VOS: Results from a Pilot Study

This section describes the results gained from a pilot study aiming to evaluate the app's technical functionality and gain feedback about users' shopping experience. For this study, we used the base version of our application, which was populated with a representative stock of food-items. Subjects in our study were free to choose between 6619 food items from seven general categories. However, this product database comes with a few drawbacks, especially compared to the offline shopping experience. This includes an over-representation of certain food categories. For instance, convenience foods and goods with long shelf lives are more commonly available in online supermarkets.

Since it is the application's primary objective to enable researchers to conduct meaningful experiments, it needs to be able to generate genuine user data. Hence, this study should also verify if users have the sensation of interacting with a real online shopping environment. Additionally, it aimed to generate feedback for further refinement and development of the online shopping experience. Finally, yet notably,

the pilot study served as a showcase to illustrate which data can be obtained from using the VOS and how it can be analyzed. The design and results of this pilot study are described below.

3.1 Procedure

The study was conducted online over ten days in December 2019. Subjects were recruited from a pool of students and university staff members. In total, the VOS was intensively tested among 29 people. All subjects had to complete the same shopping task. In particular, they were asked to shop for groceries to cover their household's needs for one week. Subjects were prompted to select items and amounts similar to their actual grocery shopping behavior. There were few restrictions except that subjects were obliged to shop at least 12 unique items to complete the task. This was imposed for two reasons: (i) 12 items is the average number of products bought per supermarket trip (e.g., Sorensen et al. 2017) and (ii) we wanted to ensure that participants are interacting with the application seriously. Yet, subjects in our study did not actually purchase the selected items. Nonetheless, previous studies have illustrated that even hypothetical purchasing scenarios can provide pertinent data, in particular when focusing on in-store behavior rather than on analyzing the final VSC (Yang and Lynn 2014; Janiszewski and Cunha Jr 2004).

After subjects completed this task, they were asked to evaluate their shopping experience with a questionnaire. Besides, basic demographical facts, individuals' grocery shopping habits, and basic economic data were queried.

3.2 Results and Discussion

None of our subjects reported any major technical problems with the VOS, and all data was accurately sent to our server. The subjects covered a wide age range from 18 up to 54 years ($M = 28.76$). Most subjects had a high level of education (this criterion was fulfilled if subjects indicated to hold at least a high school diploma) (75.9 percent), were single (51.7 percent), and lived in their own apartment (62.1 percent). This suggests that most subjects were responsible for grocery shopping in their households. Indeed, on average, they went grocery shopping 2.55 times per week and spent €30.10 per shopping trip.

User Experience

Further, we used four constructs to measure users' experience with the VOS regarding their (i) general satisfaction (GS) with the application, (ii) its information quality (IQ), (iii) its system quality (SQ), and (iv) its realism level (RL). Table 1 provides an overview of all subconstructs belonging to these constructs and the associated descriptive statistics (for a complete overview of all questionnaire items used, see Appendix B). All items were measured using a seven-point Likert scale.

Overall, subjects were satisfied with their shopping experience in our online supermarket (GS.01_{mean}: 4.89), and their expectations of the application were mainly met (GS.02_{mean}: 4.96). In addition, the subjects agreed rather than disagreed on average with the statement "I would buy groceries in a real online supermarket, similar to the VOS" (GS.03_{mean}: 4.93). Concerning the information quality, subjects predominantly rated the understandability (USS_{mean}: 5.68), reliability (RS_{mean}: 5.73), and usefulness (UFS_{mean}: 5.02) of the available product information positively.

The subjects were satisfied with information quality and system quality (SQ_{mean}: 5.43). In particular, subjects rate usability of our tool positively (USA_{mean}: 5.80). For instance, they perceived good responsiveness (SQ.01_{mean}: 6.00) and fast loading times (SQ.02_{mean}: 6.04). Additionally, the tool's user-friendliness was evaluated positively (UFS_{mean}: 5.22). For instance, subjects stated that the shop had a straightforward design (SQ.10_{mean}: 5.61), a simple layout (SQ.04_{mean}: 5.7), where all functions could be found as expected (SQ.06_{mean}: 5.71) and was easy to use (SQ.05_{mean}: 5.96). In sum, the feedback indicates that the design and structure of the shopping user interface (UI) successfully guided our users. In this article, we refer to the term user interface (UI) to describe the graphical environment (shop view) that allows users to interact with an online shop.

However, comments in a free text field yielded valuable hints at what users perceived to impede their shopping experience. The main criticism referred to the free-text search bar at the top of the page because only specific terms yielded the desired results (e.g., baked beans vs. beans). This was mainly due to the information quality of the item pool. Here, the free-text search's current implementation was based on pattern matching of the search term in the item name and brand attributes. For a more refined product search experience, well-annotated item data would be necessary. Secondly, several participants had issues with the absence of components usually featured in

Variable	Description	Mean	Median	SD
A. General Satisfaction Score (GS)	To measure general satisfaction, we asked users (GS.01) about their satisfaction with their shopping experience, (GS.02) if their expectations of shopping in VOS have been met, and (GS.03) if they would buy from a real shop similar to VOS.	4.92	5.66	1.15
B. Information Quality Score (IQ)	Information Quality comprises three second-order constructs: understandability, reliability, and usefulness.	5.52	5.59	0.71
B.1. Understandability Score (USS)	To measure understandability, we asked subjects to rate whether or not product information was (IQ.01) unambiguous, (IQ.02) easy to read, (IQ.03) comprehensible, and (IQ.04) generally understandable.	5.68	5.75	0.91
B.2. Reliability Score (RS)	To measure reliability, we asked subjects to rate whether or not product information was (IQ.05) trustworthy, (IQ.06) accurate, (IQ.07) credible, and (IQ.08) reliable for making a purchase decision.	5.73	5.75	0.69
B.3. Usefulness Score (UFS)	To measure usefulness, we asked subjects to rate whether or not product information was (IQ.09) informative for the purchase decision, (IQ.10) valuable for making a purchase decision, and (IQ.11) useful for making a purchase decision.	5.02	5.33	1.22
C. System quality Score (SQ)	System Quality comprises three second-order constructs: usability, user-friendliness, and navigation.	5.43	5.81	0.98
C.1. Usability Score (USA)	To measure usability, we asked subjects to evaluate VOS's (SQ.01) responsiveness (e.g., quick reactions to queries), (SQ.02) performance (e.g., loading times of graphics), and (SQ.03) technological modernness.	5.80	6.00	1.02
C.2. User-friendliness Score (UFS)	To measure user-friendliness, we asked subjects to rate whether or not (SQ.04) the store's layout is simple, (SQ.05) it is easy to use, (SQ.06) it is well organized, (SQ.07) it is possible to see as many products as possible at a glance, (SQ.08) it is possible to easily compare different products, (SQ.09) multiple product images and display formats are available (e.g., zoomed images), (SQ.10) its design is straightforward, and (SQ.11) it is user-friendly in general.	5.22	5.63	1.09
C.3. Navigation Score (NS)	To evaluate the site's navigation, we asked subjects to rate whether or not (SQ.12) it has made it possible to (SQ.12) easily go back and forth between pages, (SQ.13) locate the information they need with just a few clicks, (SQ.14) locate the products they prefer as quickly as possible, (SQ.15) edit their shopping cart as quickly and easily as possible (e.g., to add products), and (SQ.16) easily navigate.	5.56	6.00	1.21
D. Realism Level Score (RL)	To measure the realism level of our online shopping simulation, we asked subjects whether or not (RL.01) VOS has given them the feeling of using a real online store, (RL.02) their purchases correspond to their regular shopping behavior, (RL.03) their decisions reflect their regular in-store behavior (e.g., product comparisons), and (RL.04) their gathered information while shopping reflects their behavior on a regular shopping trip.	5.28	5.50	0.94

Note. Each score was calculated as an unweighted sum index from the items underlying the respective construct.

Table 1 — List of selected questionnaire constructs and items.

online shops. For example, they mentioned missing a landing page, remarked an unusual color scheme for a grocery store, and the absence of advertising, up-selling, and product suggestions. These are valid points of criticism, but they were deliberately not included in the project's scope. These features are broad and can be implemented in any fashion. We decided against having these in the base version of the application because their functionality is specific and not easily configurable.

Despite these limitations, subjects reported that VOS gave them the feeling of using a real online store (RL0.1_{mean}: 5.32). Further, they stated that their purchases (RL.02_{mean}: 5.61) and in-store behavior (e.g., use of information) were mostly in line with their usual shopping habits (RL.03_{mean}: 4.93; RL.04_{mean}: 5.29). Hence, the average realism level score of 5.28 points indicates that a VOS can convey an experience similar to a real online supermarket and can generate meaningful data about customer behavior exhibited. The next section shows exemplarily how this data might be analyzed.

Purchase Data

Besides user experience, we were also interested in finding out whether the application is able to provide a usable data structure for analyzing participants' purchases. For this purpose, participants' final shopping cart was recorded. Table 2 provides an overview of the variables we used exemplary for our study. Presented variables are only a selection of possible outcome-variables that might be calculated from the recorded data. For instance, as evident from the table, we focused on organically produced foods and store brands. However, these variables can easily be adapted for other product attributes like vegan, animal welfare, local origin, and much more.

It must be kept in mind that we primarily wanted to evaluate the functionalities of our application. Due to the small number of participants, no generalizable statements about shopping behavior in an online supermarket can be derived from this pilot study. Moreover, for the same reason, our data is biased by outliers. For this reason, we opted to report the median instead of the mean in the following.

However, in total, 788 items have been purchased, which is an indication that the VOS was tested extensively. On median average, subjects purchased 13 different food items. Moreover, total spending (median: 33.03 €) was relatively close to most subjects' self-reported expenses for groceries per shopping trip (median: 25.00 €). Together with the median average basket size (18 items), subjects showed realistic

Variable	Description	Min.	Max	Mean	Median	SD
Total Spending (in €)	The sum of all items purchased.	20.66	371.55	46.87	33.03	63.69
Total Spending on Organic Items (in €)	The sum of all organic items purchased.	0.00	33.13	9.27	6.87	8.34
Relative Spending on Organic Items	The share of spending for organic items in relation to total purchases.	0.00	0.77	0.27	0.19	0.22
Basket Size	The total number of all items purchased.	12.00	234.00	27.17	18.00	40.47
Number of Unique Items	The sum of all unique items purchased.	0.00	0.61	0.22	0.18	0.17
Share of Organic Items	The share of organic items in total purchases.	1.00	15.00	6.72	6.00	3.61
Number of Store Brands	The sum of items from store brands purchased.	0.01	0.75	0.34	0.35	0.18
Number of Organic Items by a Store Brand	The share of items from store brands in total purchases.	0.00	11.00	3.10	3.00	2.76
Share of Store Brands on Organic Items	The sum of purchased organic items offered by store brands.	0.00	1.00	0.74	0.85	0.31

Table 2 — Purchase Data.

shopping behavior. This impression coincides with subjects' self-reported realism level presented in the previous section. The share of organic items in total purchases was relatively high (median: 26.67 percent) compared to the market share of organic food in Germany (11.97 percent in 2019; see **BOELW**). Further, the majority of purchased organic food items were store brands (median: 85.00 percent). This is not surprising, as price premiums for organic food are high, and store brands can offer cheaper alternatives to other organic food brands.

In-Store Behavior

We used data about participants' interactions with the website to derive several variables that helped us analyze participants' in-store behavior dynamics while shopping online for groceries. Table 3 provides an overview of all considered variables and how these were conceptualized. The selection of variables presented here is based on our own research interests. Of course, researchers are not limited to this set of variables but can adapt or extend it to their own research question.

Variable	Description	Min.	Max	Mean	Median	SD
Shopping Duration (min:sec)	The duration between when a subject had successfully logged onto our virtual online supermarket website and subjects' last visit to the VSC full-screen site *	01:22	501:16	37:05	06:19	109:14
First Orientation Time (min:sec)	Time taken between the start of the shopping task and adding the first item to the cart.	00:59	00:04	02:38	00:50	00:44
Product Adds (to cart)	Number of unique items added to VSC.	12.00	29.00	14.38	13.00	4.35
Product Removes (from cart)	Number of unique items removed from the VSC.	0.00	3.00	0.55	0.00	0.87
Product amount changes (in cart)	Number of amount changes in the VSC during shopping.	0.00	13.00	1.00	0.00	2.74
Total Cart Adjustments	Sum of all adjustments (PA, PR, and PC) to the VSC.	12.00	35.00	15.93	14.00	6.12
Detailed product views	Sum of all product pages a consumer has visited during shopping.	0.00	16.00	2.93	2.00	3.85
Cart-to-Detail Ratio	Ratio of the number of items added to the VSC to the number of product-detail views.	0.00	1.14	0.22	0.14	0.28
Add-to-Purchase Ratio	Ratio of the number of items bought to the number of items added to VSC.	0.80	1.00	0.96	1.00	0.53
Total user interactions	Number of all user interactions that link a subject to another (sub)page of the VSC or changes the view of the currently seen (sub)page.	3.00	460.00	83.00	47.00	105.28
Interactions-to-Purchase Ratio	Number of user interaction in relation to the basket size.	0.23	23.21	5.71	3.62	6.09

Table 3 — In-store Behavior Variables.

As evident in the table, subjects spent on median an average of 6 min and 19 s shopping in our virtual online supermarket. Hence, shopping was faster than previous studies, indicating where estimates for average (online) supermarket shopping durations ranged between 13 and 40 min (Anesbury et al. 2016; Hui et al. 2009; Sorensen et al. 2017). The average number of total user interactions was relatively low (median: 47.00), and it took subjects on median an average of just 3.62 interactions on the website to purchase an item. Consequently, many selections were made without having viewed many different pages beforehand.

Moreover, most subjects did not or only slightly adjusted their VSC during their shopping trip. In particular, items were rarely removed or edited once they had been added to the shopping cart (e.g., change in amount). Consequently, the add-to-purchase ratio was almost one on average (mean: 0.96), respectively, the median average. Further, the cart-to-detail ratio (median: 0.14) is an indicator that subjects have not viewed the detailed product pages for the majority of items selected. Hence, it can be stated that most subjects in this sample showed a straightforward shopping behavior. They were not interested in extensively browsing and comparing different products in our online supermarket. Instead, purchasing decisions were made relatively quickly in most cases, without looking at individual items in more detail. This behavior is consistent with shopping patterns found in brick-and-mortar supermarkets (e.g., Dickson and Sawyer 1990; Cohen and Babey 2012). On median average, after logging onto our virtual online supermarket, it took subjects less than one minute to add their first item to the cart. In addition to subjects' self-reported user experience, we considered this another indicator for good usability and user-friendliness of our tool.

Overall, data from this pilot study suggests that subjects exhibited a reasonable purchasing behavior and the application gave them the impression of shopping at a real online supermarket. This indicates that the VOS can convey an experience similar to a real online grocery store, thereby generating meaningful data about its customers' behavior. Moreover, it recorded all actions taken by the individual subjects and provided this data as expected. Hence, VOS supplied the necessary means for conducting our study.

While users evaluated their experience predominantly positively, feedback also suggests areas for improvement: Regarding system quality, future research projects may improve query and data retrieval between front- and back-end. The free-text search bar, which currently only operates on searching the queried term in the product's name or brand, should be extended to provide higher versatility. Furthermore,

researchers may direct their attention towards augmenting the shopping experience, e.g., by implementing a dedicated landing page where subjects start their shopping task. Similar considerations apply to product pages. In addition, the user experience may benefit from viewing multiple product images, choosing between different display formats (e.g., zoom), or having complementary media such as product videos available.

4 Conclusions

This paper introduced the VOS, an open-source web-based application designed to run computer-simulated shopping experiments. While its front-end emulates a modern online supermarket's design and functions, its back-end provides a visual administration interface that enables researchers to create and modify different experimental conditions easily. This feature makes VOS a modular, highly customizable, and usable research tool for analyzing consumer behavior in (online) supermarkets. Researchers can build on several preconfigured use cases to implement and test different in-store interventions, including economic incentives and choice architecture techniques. The base version already includes possibilities to configure (i) taxes and subsidies, (ii) product arrangement and placement, (iii) product labeling, and to choose between different (iv) types of VSCs and (v) swap options. Further, we present the results from a pilot study, which showed that the functionalities of our tool have been working and subjects evaluated their shopping experience to be realistic and user-friendly.

We considered the VOS to be a useful tool for testing a broad range of policy interventions in a realistic online shopping environment. This can be done at relatively low cost, without a complex implementation process, and without collaborating with a specific retailer. Hence, the VOS offers researchers the opportunity to heighten the realism level of their experimental designs. Moreover, vast options for automatic data recording allow analyzing purchases and bringing light into the black box of consumers' in-store decision-making. In further analyzing the particularities of digital interactions in online grocery, scholars could also shed further light on ethical strings attached to influencing users and customers in digital environments (Lembcke et al. 2019).

In summary, we hope that our tool will support researchers in exploring key issues that might be conducive to understanding and promoting sustainable consumption.

Policy makers and industrial stakeholders can benefit from knowledge that informs the design of future policy interventions, (in-store) marketing communication, product packaging, or the shop design of (online) supermarkets in general.

Supplementary Materials

The following materials are available online at <https://www.mdpi.com/article/10.3390/su13084375/s1>, VOS User Guide.

Acknowledgments

We acknowledge support by the Open Access Publication Funds of the University of Goettingen. Further, we would like to thank Sebastian Kuitert for the extensive programming and technical support. Moreover, we thank our following interview partners for their time and consideration: (1) Paolo Crosetto, Senior Researcher at the French National Research Institute for Agriculture, Food and Environment (INRAE); (2) Achim Spiller, Professor of “Marketing for Food and Agricultural Products” at the Georg-August University of Göttingen; (3) Stefan Morana, Senior Researcher with the group Information Systems and Service Design at the Karlsruhe Institute of Technology (KIT); and (4) Alfred Benedikt Brendel, Professor for Business Informatics, especially Intelligent Systems and Services at the Technische Universität Dresden.

Author Contributions

Conceptualization: N.E., T.-B.L., and A.B.B.; Data curation: N.E.; Formal analysis: N.E.; Investigation: N.E., T.-B.L.; Methodology: N.E., T.-B.L., and A.B.B.; Project administration: N.E., T.-B.L. and A.B.B.; Supervision: A.B.B., K.B., and L.M.K.; Validation: N.E., T.-B.L., and A.B.B.; Visualization: N.E. and T.-B.L.; Writing—original draft: N.E.; Writing—review editing: T.-B.L., A.B.B.

References

- An, R. (2012). "Effectiveness of Subsidies in Promoting Healthy Food Purchases and Consumption: A Review of Field Experiments". In: *Public health nutrition* 16.7, pp. 1215–1228.
- Anesbury, Z., M. Nenycz-Thiel, J. Dawes, and R. Kennedy (2016). "How do shoppers behave online? An observational study of online grocery shopping". In: *Journal of Consumer Behaviour* 15.3, pp. 261–270.
- Barton, A. and T. Grüne-Yanoff (2015). "From libertarian paternalism to nudging—and beyond". In: *Review of Philosophy and psychology* 6.3, pp. 341–359.
- Benn, Y., T. L. Webb, B. P. I. Chang, and J. Reidy (2015). "What information do consumers consider, and how do they look for it, when shopping for groceries online?" In: *Appetite* 89, pp. 265–273.
- Bianchi, F., E. Garnett, C. Dorsel, P. Aveyard, and S. A. Jebb (2018). "Restructuring physical micro-environments to reduce the demand for meat: a systematic review and qualitative comparative analysis". In: *The Lancet Planetary Health* 2.9, e384–e397.
- Bonnet, C., Z. Bouamra-Mechemache, V. Réquillart, and N. Treich (2020). "Regulating meat consumption to improve health, the environment and animal welfare". In: *Food Policy*, p. 101847.
- Bryła, P. (2018). "Organic food online shopping in Poland". In: *British Food Journal* 120.5, pp. 1015–1027.
- Cadario, R. and P. Chandon (2020). "Which healthy eating nudges work best? A meta-analysis of field experiments". In: *Marketing Science* 39.3, pp. 465–486.
- Cameron, A. J., E. Charlton, W. W. Ngan, and G. Sacks (2016). "A Systematic Review of the Effectiveness of Supermarket-Based Interventions Involving Product, Promotion, or Place on the Healthiness of Consumer Purchases". In: *Current Nutrition Reports* 5.3, pp. 129–138.
- Campos, S., J. Doxey, and D. Hammond (2011). "Nutrition labels on pre-packaged foods: a systematic review". In: *Public health nutrition* 14.8, pp. 1496–1506.
- Chetty, R., A. Looney, and K. Kroft (2009). "Salience and taxation: Theory and evidence". In: *American Economic Review* 99.4, pp. 1145–1177.
- Cheung, T. T. L., A. F. Junghans, G. B. Dijksterhuis, F. Kroese, P. Johansson, L. Hall, et al. (2016). "Consumers' choice-blindness to ingredient information". In: *Appetite* 106, pp. 2–12.

- Chintagunta, P. K., J. Chu, and J. Cebollada (2012). "Quantifying Transaction Costs in Online/Off-line Grocery Channel Choice". In: *Marketing Science* 31.1, pp. 96–114.
- Cohen, D. A. and S. H. Babey (2012). "Contextual Influences on Eating Behaviors: Heuristic Processing and Dietary Choices". In: *Obesity reviews : an official journal of the International Association for the Study of Obesity* 13.9, pp. 766–779.
- Demarque, C., L. Charalambides, D. J. Hilton, and L. Waroquier (2015). "Nudging Sustainable Consumption: The Use of Descriptive Norms to Promote a Minority Behavior in a Realistic Online Shopping Environment". In: *Journal of Environmental Psychology* 43, pp. 166–174.
- Dickson, P. R. and A. G. Sawyer (1990). "The price knowledge and search of supermarket shoppers". In: *Journal of marketing* 54.3, pp. 42–53.
- Epstein, L. H., E. Finkelstein, H. Raynor, C. Nederkoorn, K. D. Fletcher, N. Jankowiak, et al. (2015). "Experimental analysis of the effect of taxes and subsidies on calories purchased in an on-line supermarket". In: *Appetite* 95, pp. 245–251.
- Etilé, F. and S. Teyssier (2016). "Signaling corporate social responsibility: Third-party certification versus brands". In: *The Scandinavian Journal of Economics* 118.3, pp. 397–432.
- Eyles, H., R. McLean, B. Neal, Y. Jiang, R. N. Doughty, R. McLean, et al. (2017). "A salt-reduction smartphone app supports lower-salt food purchases for people with cardiovascular disease: findings from the SaltSwitch randomised controlled trial". In: *European journal of preventive cardiology* 24.13, pp. 1435–1444.
- Forwood, S. E., A. L. Ahern, T. M. Marteau, and S. A. Jebb (2015). "Offering within-category food swaps to reduce energy density of food purchases: a study using an experimental online supermarket". In: *The international journal of behavioral nutrition and physical activity* 12, p. 85.
- Goldfarb, A., R. C. McDevitt, S. Samila, and B. S. Silverman (2015). "The Effect of Social Interaction on Economic Transactions: Evidence from Changes in Two Retail Formats". In: *Management Science* 61.12, pp. 2963–2981.
- Grunert, K. G., S. Hieke, and J. Wills (2014). "Sustainability labels on food products: Consumer motivation, understanding and use". In: *Food Policy* 44, pp. 177–189.
- Guthrie, J., L. Mancino, and C.-T. J. Lin (2015). "Nudging Consumers toward Better Food Choices: Policy Approaches to Changing Food Consumption Behaviors". In: *Psychology & Marketing* 32.5, pp. 501–511.
- Hansen, P. G., L. R. Skov, and K. L. Skov (2016). "Making Healthy Choices Easier: Regulation versus Nudging". In: *Annual review of public health* 37, pp. 237–251.

- Hartmann-Boyce, J., F. Bianchi, C. Piernas, S. P. Riches, K. Frie, R. Nourse, et al. (2018). "Grocery store interventions to change food purchasing behaviors: a systematic review of randomized controlled trials". In: *The American journal of clinical nutrition* 107.6, pp. 1004–1016.
- Häubl, G. and V. Trifts (2000). "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids". In: *Marketing Science* 19.1, pp. 4–21.
- Hertwich, E. (2010). *Assessing the environmental impacts of consumption and production: priority products and materials*. UNEP/Earthprint.
- Hoek, A. C., D. Pearson, S. W. James, M. A. Lawrence, and S. Friel (2017). "Healthy and environmentally sustainable food choices: Consumer responses to point-of-purchase actions". In: *Food Quality and Preference* 58, pp. 94–106.
- Hollands, G. J., I. Shemilt, T. M. Marteau, S. A. Jebb, M. P. Kelly, R. Nakamura, et al. (2013). "Altering micro-environments to change population health behaviour: towards an evidence base for choice architecture interventions". In: *BMC public health* 13, p. 1218.
- Hui, S. K., P. S. Fader, and E. T. Bradlow (2009). "Research Note —The Traveling Salesman Goes Shopping: The Systematic Deviations of Grocery Paths from TSP Optimality". In: *Marketing Science* 28.3, pp. 566–572.
- Hui, S. K., J. J. Inman, Y. Huang, and J. Suher (2013). "The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies". In: *Journal of Marketing* 77.2, pp. 1–16.
- Hummel, D. and A. Maedche (2019). "How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies". In: *Journal of Behavioral and Experimental Economics* 80, pp. 47–58.
- Huyghe, E., J. Verstraeten, M. Geuens, and A. van Kerckhove (2017). "Clicks as a Healthy Alternative to Bricks: How Online Grocery Shopping Reduces Vice Purchases". In: *Journal of Marketing Research* 54.1, pp. 61–74.
- Janiszewski, C. and M. Cunha Jr (2004). "The influence of price discount framing on the evaluation of a product bundle". In: *Journal of Consumer Research* 30.4, pp. 534–546.
- Keller, C., F. Markert, and T. Bucher (2015). "Nudging product choices: The effect of position change on snack bar choice". In: *Food Quality and Preference* 41, pp. 41–43.
- Koutoukidis, D. A., S. A. Jebb, J. M. Ordóñez-Mena, M. Noreik, M. Tsiountsioura, S. Kennedy, et al. (2019). "Prominent positioning and food swaps are effective

- interventions to reduce the saturated fat content of the shopping basket in an experimental online supermarket: a randomized controlled trial". In: *International Journal of Behavioral Nutrition and Physical Activity* 16.1, p. 50.
- Kummer, S. and R. Milestad (2020). "The Diversity of Organic Box Schemes in Europe—An Exploratory Study in Four Countries". In: *Sustainability* 12.7, p. 2734.
- Küst, P. (2019). "The Impact of the organic label halo effect on consumers' quality perceptions, value-in-use and well-being". In: *Junior Management Science* 4.2, pp. 241–264.
- Lembcke, T.-B., N. Engelbrecht, A. B. Brendel, and L. Kolbe (2019). "To Nudge or not to Nudge: Ethical Considerations of Digital Nudging Based on Its Behavioral Economics Roots". In: *In Proceedings of the 27th European Conference on Information Systems (ECIS), Stockholm & Uppsala, Sweden, June 8-14, 2019*.
- Lembcke, T.-B., N. Engelbrecht, M. Willnat, and S. Lichtenberg (2020). "Behavioral Design in Online Supermarkets: How Virtual Shopping Cart Functions Impact Sustainable Consumption". In:
- Liberato, S. C., R. Bailie, and J. Brimblecombe (2014). "Nutrition interventions at point-of-sale to encourage healthier food purchasing: a systematic review". In: *BMC public health* 14, p. 919.
- Liu, Z. (2005). "Reading behavior in the digital environment". In: *Journal of Documentation* 61.6, pp. 700–712.
- Loh, K. K. and R. Kanai (2016). "How Has the Internet Reshaped Human Cognition?" In: *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry* 22.5, pp. 506–520.
- Mangen, A., B. R. Walgermo, and K. Brønnick (2013). "Reading linear texts on paper versus computer screen: Effects on reading comprehension". In: *International Journal of Educational Research* 58, pp. 61–68.
- Münscher, R., M. Vetter, and T. Scheuerle (2016). "A Review and Taxonomy of Choice Architecture Techniques". In: *Journal of Behavioral Decision Making* 29.5, pp. 511–524.
- Newman, C. L., E. Howlett, and S. Burton (2016). "Effects of Objective and Evaluative Front-of-Package Cues on Food Evaluation and Choice: The Moderating Influence of Comparative and Noncomparative Processing Contexts". In: *Journal of Consumer Research* 42.5, pp. 749–766.
- OECD (2016). *Behavioural Insights for Environmentally Relevant Policies: Review of Experiences from OECD Countries and Beyond*. Paris: OECD Publishing.

- (2017). *Behavioural Insights and Public Policy: Lessons from Around the World*. Paris: OECD Publishing.
- Ogonowski, A., A. Montandon, E. Botha, and M. Reyneke (2014). “Should new online stores invest in social presence elements? The effect of social presence on initial trust formation”. In: *Journal of Retailing and Consumer Services* 21.4, pp. 482–491.
- Onwezen, M. C. and van der Weele, Cor N (2016). “When indifference is ambivalence: Strategic ignorance about meat consumption”. In: *Food Quality and Preference* 52, pp. 96–105.
- Pitts J., S. B., S. W. Ng, J. L. Blitstein, A. Gustafson, and M. Niculescu (2018). “Online grocery shopping: promise and pitfalls for healthier food and beverage purchases”. In: *Public health nutrition*, pp. 1–17.
- Reisch, L., U. Eberle, and S. Lorek (2013). “Sustainable food consumption: an overview of contemporary issues and policies”. In: *Sustainability: Science, Practice and Policy* 9.2, pp. 7–25.
- Riches, S. P., P. Aveyard, C. Piernas, M. Rayner, and S. A. Jebb (2019). “Optimising swaps to reduce the salt content of food purchases in a virtual online supermarket: A randomised controlled trial”. In: *Appetite* 133, pp. 378–386.
- Sanjari, S. S., S. Jahn, and Y. Boztug (2017). “Dual-process theory and consumer response to front-of-package nutrition label formats”. In: *Nutrition reviews* 75.11, pp. 871–882.
- Santos, B. D. I., A. Hortaçsu, and M. R. Wildenbeest (2012). “Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior”. In: *American Economic Review* 102.6, pp. 2955–2980.
- Scheibehenne, B., L. Miesler, and P. M. Todd (2007). “Fast and frugal food choices: Uncovering individual decision heuristics”. In: *Appetite* 49.3, pp. 578–589.
- Schneider, C., M. Weinmann, and J. Vom Brocke (2018). “Digital Nudging-Influencing Choices by Using Interface Design”. In: *Communications of the ACM* January, pp. 1–11.
- Schwarz, N. (2004). “Metacognitive experiences in consumer judgment and decision making”. In: *Journal of Consumer Psychology* 14.4, pp. 332–348.
- Shah, A. K. and D. M. Oppenheimer (2007). “Easy does it: The role of fluency in cue weighting”. In:
- Sheehan, D. and K. van Ittersum (2018). “In-Store Spending Dynamics: How Budgets Invert Relative-Spending Patterns”. In: *Journal of Consumer Research* 45.1, pp. 49–67.

- Sorensen, H., S. Bogomolova, K. Anderson, G. Trinh, A. Sharp, R. Kennedy, et al. (2017). "Fundamental patterns of in-store shopper behavior". In: *Journal of Retailing and Consumer Services* 37, pp. 182–194.
- Thow, A. M., S. Downs, and S. Jan (2014). "A systematic review of the effectiveness of food taxes and subsidies to improve diets: understanding the recent evidence". In: *Nutrition reviews* 72.9, pp. 551–565.
- van Herpen, E. (2016). "Product Category Layout and Organization: Retail Placement of Food Products". In: *Reference module in food science*. Ed. by G. W. Smithers. [Amsterdam]: Elsevier.
- van Herpen, E. and A. Bosmans (2018). "Arranging the assortment to arouse choice: Effects of goal-relevant assortment organization on food choice and variety perceptions". In: *Food Quality and Preference* 64, pp. 192–204.
- van Herpen, E., E. van den Broek, H. C. M. van Trijp, and T. Yu (2016). "Can a virtual supermarket bring realism into the lab? Comparing shopping behavior using virtual and pictorial store representations to behavior in a physical store". In: *Appetite* 107, pp. 196–207.
- van Ittersum, K., J. M. Pennings, and B. Wansink (2010). "Trying Harder and Doing Worse: How Grocery Shoppers Track In-Store Spending". In: *Journal of Marketing* 74.2, pp. 90–104.
- van Ittersum, K., B. Wansink, J. M. Pennings, and D. Sheehan (2013). "Smart Shopping Carts: How Real-Time Feedback Influences Spending". In: *Journal of Marketing* 77.6, pp. 21–36.
- van Kleef, E. and H. Dagevos (2015). "The growing role of front-of-pack nutrition profile labeling: A consumer perspective on key issues and controversies". In: *Critical reviews in food science and nutrition* 55.3, pp. 291–303.
- Weinrich, R. and A. Spiller (2016). "Developing food labelling strategies: Multi-level labelling". In: *Journal of Cleaner Production* 137, pp. 1138–1148.
- Willett, W., J. Rockström, B. Loken, M. Springmann, T. Lang, S. Vermeulen, et al. (2019). "Food in the Anthropocene: the EAT–Lancet Commission on healthy diets from sustainable food systems". In: *The Lancet*.
- Yang, S. and M. Lynn (2014). "More evidence challenging the robustness and usefulness of the attraction effect". In: *Journal of Marketing Research* 51.4, pp. 508–513.

Appendix

A.1 Brief Information about Technology and Programming

Application Development Frameworks

Our application is built upon two development frameworks to decrease development-time, reduce maintenance cost, and quality assurance reasons. For the back-end, Node.js is used to make a web-based service, provide an Application Programming Interface (API) for handling requests, and run Create Read Update Delete (CRUD) operations on the database. Second, for implementing the web, the front-end Angular version 8 is used. Both frameworks are based on JavaScript.

Model View Controller (MVC)

The implementation style of our application follows the model view controller (MVC) concept. Working with MVC in web application development is different from conventional application development. The architecture has to be partitioned between the client and the server-side. The client-side always handles a web application's view, but the model and controller can be partitioned in various ways between the client and server. Hence, a compelling architecture would rely exclusively on the server to refresh the client's screen. In this case, the model and the view-generating logic for the client's browser would reside entirely on the server. Moreover, the controller would partially reside on the client (detecting user interaction) but mostly reside on the server (code that updates the state of the model's business objects based on a HTTP request). This describes a thin-client approach with the advantages of decreasing the client machine's performance demand and providing greater security, performance, and data consistency for the application. Web application frameworks that reflect this paradigm are Django and ASP.NET.

The other extreme is maintaining the bulk of the application on the client-side (fatclient approach). This means that the model mostly resides on the client-side, but the database remains on the server-side. In particular, the view is exclusively implemented on the client-side, and the controller mostly resides there. This provides a more seamless and interactive experience through fewer load times, minimizing the need to make server calls. Frameworks that support this style of partitioning are AngularJS, EmberJS, and JavaScriptMVC.

Front-End Programming

The views are built directly with HTML5 in conjunction with style sheets written in SCSS (Sassy CSS) to provide a visually appealing and user-friendly experience. This combination is supported by most browsers natively, which means a wide range of devices can be supported. In this way, a responsive front-end web design is provided, which is consistent between devices and browsers. This combination allows for the separation of presentation and content, the reduction of repetitive code, flexibility, control of the presentation, and sharing of formats between views. The web-view also utilizes Bootstrap and Angular Material, which are CSS libraries that offer standardized web-content styling and component options. Therefore, developers benefit from the ease of use and accessibility of these frameworks for building visually engaging views. Moreover, they are still able to determine custom styles and layouts centrally.

The front-end codebase is designed and implemented to deliver an accessible experience both to users and developers. Features and design elements are designed to encapsulate specific functionality or business logic. Hence, the code is partitioned into feature modules (which house the model's logic), view, and controller, represented by the components contained. Utilizing this implementation logic makes it easy to tease apart singular design elements, use case implementations, and feature sets for subsequent expansion and development. Singular design elements of the application are thus contained in one subfolder easily recognizable (see 9).

Changing the appearance or behavior of such a design feature would mean locating the associated component and altering the source files. These component definitions are divided into (i) the template (*.html file), which handles the display elements, (ii) the (*.scss file), which defines the styles for this component, and (iii) the (*.ts file), which handles the data CRUD, data binding, and event handling of the specific component. In addition to this, the implementation of our application follows a strict separation of control. Functionality like data handling, CRUD operations, and event recording is centralized and separated from component view logic. Injectable services offer reusable access to functions, which are generally consumed by multiple components. Furthermore, the services are descriptively modeled to offer all functionality connected with a specific data model. For example, CRUD operations connected to the shopping cart component are combined into a shopping-cart.services.ts file. If any view component needs access to the specific functions and data of that topic, it has to inject the shared instance of the service into its constructor, thereby gaining access.

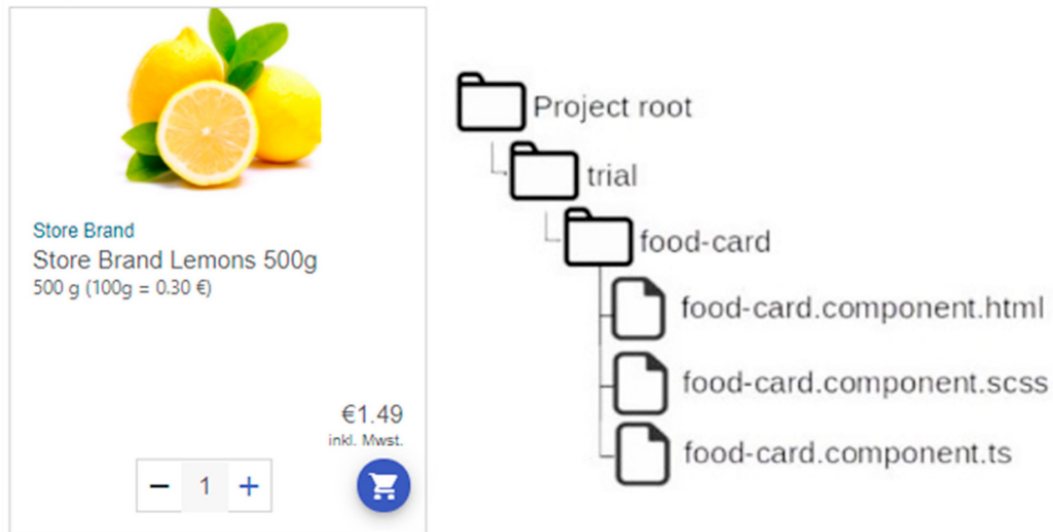


Figure 9 — Example of component structure with design element “food card” on the left and associated folder structure representing the code on the right.

This offers organized, reusable, and easy access to all operations needed throughout view components. Even if any implementation details change, these changes need only be applied in one file.

Back-End Programming

The back-end is a representational state transfer (REST) API. This exposes the data, functions, and facilitates the interaction between the database and the front-end application. Moreover, it exposes endpoints that respond to client requests in a predictable manner. The web services are stateless, as they do not maintain the state of each client application accessing the web-service, instead offering predefined sets of stateless operations. This allows it to remain independent of the front-end application, meaning that these web services may serve different client applications and can be interacted with or without using the front-end application. This independence offers the advantage that researchers are not limited to using the visual treatment edit interface to interact with and change treatment aspects or items. Scripts can be written that automate treatment creation, modification, and data analysis tasks (examples can be found in the VOS User Guide).

As mentioned above, the back-end is based on the application development frameworks Node.js, an event-driven JavaScript run-time environment that works outside of

the browser. This allows for continuous utilization of JavaScript in both application areas and reduces entry barriers for developers, as only knowledge of one programming language is required.

However, using JavaScript for the REST API does not incur performance decreases, as could be generally expected. Node.js is built on the libraries V8 and libuv; these are responsible for partly converting the JavaScript code to C++ code, thereby combining the ease of use attributed to JavaScript and its high performance attributed to C++. It is also highly scalable without threading, instead of utilizing a simplified model of event-driven programming with callbacks to signal the task's completion. However, this essentially single-threaded approach means that the application cannot scale vertically, which means that merely adding computing power to a given system will not directly translate to an increase in application performance. Despite this, it is still capable of scaling by running several concurrent instances of the same application within one cluster manager (cluster mode). This distributes the workload among the available application instances. A production-ready and open-source load balancing software for Node.js applications is already available free of charge; see, for example, PM2.

Moreover, the repository structure of the back-end application is modeled to promote easy access and understandability. Thereby, the subfolder structure represents the data structure utilized throughout the project. For instance, an item and all its associated CRUD operations are contained in one subfolder (see Figure 9): models (*.model.js files) enforce the document structure, routes (*.route.js files) define the request endpoints, and with that access to all the operations that can be performed on the data objects. Additionally, functions collect reusable logic used throughout the route definitions. Middleware functions (*.middleware.js files) provide necessary state information to the otherwise stateless endpoints. This makes the project accessible for further development, as the application structure can be deduced from the repository structure. Effects from changing aspects of the data structure are contained in this subfolder and do not affect the overall application. This makes it easy to add or change and customize the functions, data models, and route specifications implemented in the base application.

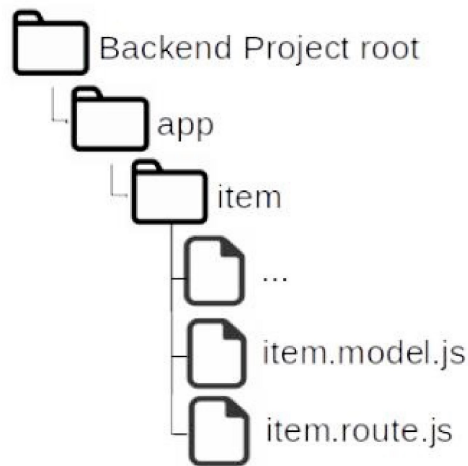


Figure 10 — Sample repository structure as used in the back-end.

Back-End Programming

The code is developed using Git, a source-code versioning system. This encourages good backup and versioning practices and allows developers to synchronize files across computers, develop collaboratively, manage separate branches, and merge synchronization conflicts. The project is open for other researchers to join, collaborate, or just download the source code on GitHub. Developers will find more detailed information on the VOS programming in the VOS User Guide.

A.2 List of questionnaire constructs and items.

ID	Description	Mean	Median	SD
Please rate the following statements:				
GS.01	I am satisfied with my shopping experience in this online supermarket.	4.89	6.00	1.55
GS.02	Overall, my expectations of shopping in this online supermarket have been met.	4.86	5.00	1.26
GS.03	I would buy food from a real online supermarket similar to this VOS.	4.93	5.00	1.46
Please rate the understandability of the product information provided in the store. Product information provided in the store was generally...				
IQ.01	... unambiguous.	5.50	6.00	0.92
IQ.02	... easy to read.	5.64	6.00	1.25
IQ.03	... comprehensible.	5.71	6.00	1.05
IQ.04	... in general, understandable for me.	5.86	6.00	0.89
With regard to reliability, the product information was...				
IQ.05	... trustworthy.	5.68	6.00	1.02
IQ.06	... accurate.	5.71	6.00	0.85
IQ.07	... credible.	5.68	6.00	0.86
IQ.08	... in general, reliable for making my purchase decision.	5.86	6.00	0.71
With regard to usefulness, the product information was...				
IQ.09	... informative for my purchase decision.	4.86	5.00	1.35
IQ.10	... valuable for making my purchase decision.	5.18	6.00	1.25
IQ.11	... in general, useful for making my purchase decision.	5.04	5.00	1.35
Please rate the usability of the online store:				
SQ.01	Responsiveness (e.g., the online shop reacts quickly to queries).	6.00	6.00	1.25
SQ.02	Performance (e.g., loading times of text and graphics)	6.04	6.00	1.20
SQ.03	Technological modernness in general	5.36	6.00	1.51
Please rate the user-friendliness of the online store with regard to whether...				
SQ.04	... its layout is simple.	5.71	6.00	1.54
SQ.05	... it is easy to use	5.96	6.00	1.11
SQ.06	... it is well organized (e.g., functions are where you would expect them)	5.36	6.00	1.70
SQ.07	... it is possible to see as many products as possible at a glance	5.04	6.00	1.69
SQ.08	... it is possible to compare different products easily	4.43	5.00	1.20
SQ.09	... multiple product images and display formats are available (e.g., zoomed images)	4.61	5.00	1.51
SQ.10	... its design is clear	5.61	6.00	1.10
SQ.11	... it is user-friendly in general	5.07	6.00	1.63
Navigation: The navigation on the site has made it possible...				
SQ.12	... to easily go back and forth between pages.	5.50	6.00	1.37
SQ.13	... to locate the information I need with just a few clicks.	5.35	6.00	1.04
SQ.14	... to locate the products I prefer as quickly as possible.	5.07	6.00	1.61
SQ.15	... to edit my shopping cart as quickly and easily as possible (e.g., to add products).	5.79	6.00	1.32
SQ.16	... to easily navigate in general.	5.68	6.00	1.19
Please tell us to what extent you agree with the following statements:				
RL.01	The design of the website has given me the feeling of using a real online store.	5.32	6.00	1.59
RL.02	Product selections and the price of my total shopping cart corresponds to my regular shopping behavior.	5.61	6.00	0.92
RL.03	The decisions I have made reflect exactly how I would behave on a regular shopping trip.	4.93	5.00	1.65
RL.04	The information which I gathered while shopping reflects my behavior on a regular shopping trip.	5.29	6.00	1.27

Notes: Items have been translated from German. All items have been measured on a 7-point Likert scale (1 = "I do not agree at all" and 7 = "I completely agree").

CHAPTER V

Behavioral Design In Online Supermarkets: How Virtual Shopping Cart Functions Impact Sustainable Consumption.

Tim-Benjamin Lembcke, Nils Engelbrecht, Mathias Willnat, and Sascha Lichtenberg

In recent years, the negative environmental impact of consumers' dietary habits has become more visible. Accordingly, in-store interventions to promote more sustainable (e.g., organic) food choices have received increased scholarly attention. Thereby, online grocery shopping is gaining momentum as web-applications provide decision support tools such as real-time spending feedback (RSF). Building on budgeting and spending literature, this study provides initial insights on the impact of RSF on consumers' organic food choices in online supermarkets. Using a free simulation experimental approach, we were able to track participants' real grocery shopping behavior within a realistic online shopping environment. Within a baseline (no RSF) and an intervention (RSF) condition (between subject design), we show that RSF facilitated participants to stay within their budget and significantly reduced underspending. Some-what surprisingly in response to the RSF, participants who usually buy fewer organic products purchased significantly more organic food items, both in absolute and relative term.

Keywords Real-Time Spending Feedback · Online Supermarkets · Organic and Sustainable Food Choices · Behavioral Design

Published as *Lembcke, T. -B.; Engelbrecht, N.; Willnat, M.; Lichtenberg, Sascha: Behavioral Design in Online Supermarkets - How Virtual Shopping Cart Functions Impact Sustainable Consumption. AMCIS 2020 Proceedings.*

1 Introduction

Due to the negative environmental impact of the predominant dietary patterns in western countries, there is a controversial debate about the role of policies in promoting “better” food choices for the consumer, the environment, and the society at large (e.g., Willett et al. 2019). As most purchasing decisions for groceries take place in supermarkets (i.e., in-store), public and private institutions aiming to alter consumers’ food choices have become interested in in-store interventions (e.g., Hartmann-Boyce et al. 2018). Typically, food-related purchases in a supermarket are frequently repeated decisions with a habitual nature and without a high level of involvement. Moreover, as consumers are surrounded by numerous environmental cues, decisions are made “fast and frugal” and consumers sometimes unconsciously react to even “seemingly irrelevant” (Thaler 2016) factors like product arrangements—even if they are aware of these (Baker et al. 1994; Cohen and Babey 2012; Turley and Milliman 2000). Therefore, a store’s micro-environment provides grounds for the development, evaluation, and implementation of policy interventions aiming to promote more sustainable food choices (e.g., Escaron et al. 2013). Thereby, online grocery shopping is gaining momentum for multiple reasons: it (i) is becoming an increasingly more important retail channel (Anesbury et al. 2016), (ii) offers numerous customization and design options for “choice architects” (Thaler et al. 2014), and (iii) offers innovative decision support tools.

However, despite the growing interest in e-grocery research (e.g., Anesbury et al. 2016; Benn et al. 2015; Bhatnagar et al. 2000; Martín et al. 2019; Shin et al. 2019), only little is known about how different “micro-environmental design choices” impact consumers’ online purchase behavior for groceries (Anesbury et al. 2016; Huyghe et al. 2017)). For instance, in traditional brick-and-mortar stores, consumers do not receive any feedback about the price of their shopping cart until checkout. Simultaneously, a stimuli-rich supermarket necessitates a lot of cognitive effort to keep track of actual in-store spending. Research on in-store spending thus proposes that consumers predominantly adapt mental computing strategies to approximate the price instead of trying to calculate the price of their total shopping cart to the cent (van Ittersum et al. 2010). As these mental computing strategies are less accurate and prone to errors, consumers are less confident about their total spending (van Ittersum et al. 2013).

In contrast, virtual shopping carts (VSC) offer some advantageous functions including real-time spending feedback (RSF). Even the most simplistic shopping cart

software enables consumers to accumulate all want-to-buy products in a list and automatically provides key information like the price and the quantity of each chosen item as well as the total price of the shopping cart (i.e., RSF). Initial empirical results indicate that under these conditions consumers should struggle less with estimation biases. In turn, this can have an impact on consumers' purchasing behavior, e.g. by increasing total spending as well as spending for higher-priced and hedonic products (van Ittersum et al. 2013). Therefore, the question arises to what extent a RSF could be part of a digital nudging intervention (e.g., Lembcke et al. 2019) to enable online grocery shoppers to make more sustainable food choices. However, to the best of our knowledge, no research exists that investigates the impact of VSC on consumer's pro-environmental consumption patterns. This study makes an initial contribution to fill this gap by investigating how a real-time spending feedback mechanism affects food choices in terms of total spending and basket share of organic products bought by budget shoppers.

To capture these variables, we are using a free simulation experimental approach (e.g., Demarque et al. 2015). For this purpose, an experimental web-based online shopping platform was developed (see Engelbrecht et al. 2020). Using the websites of the leading supermarket providers in Germany for orientation, the tool's front-end was designed to emulate the store design and functions (e.g., navigation tools) of a fictitious online supermarket.

2 Theoretical Background

2.1 Estimation Bias and Spending Uncertainty

According to Bénabou and Tirole (2004), budgeting describes a behavior whereby individuals earmark a certain amount of their income into mental accounts for specific uses; for instance expenditures for food, housing, or leisure activities. Applied to the context of our study, we have defined budget buyers as consumers who have an explicit budget in mind while shopping in a supermarket, which they do not want to exceed or fall short of (van Ittersum et al. 2013). Mental accounting research has shown that such a behavioral pattern alters budget shoppers' in-store decision-making in multiple ways (Heath and Soll 1996; Thaler 1985). One factor is that individuals do not have knowledge about their actual in-store spending and are not able to calculate the total price of their shopping cart at any time (van Ittersum et al. 2010).

Consequently, it is very difficult for budget shoppers to utilize their entire budget. Instead of calculating the exact value of their shopping cart, consumers tend to use simple mental computation strategies to approximate total spending. This results in less confident, less accurate, and more biased calculations, i.e., estimation biases (Chandon and Wansink 2012).

Against this background, we are following van Ittersum et al. (2013) and hypothesize that implementing a real-time spending feedback mechanism in our experimental online supermarket will reduce the perceived spending uncertainty among budget shoppers:

Hypothesis 1 *RSF decreases participants' perceived difficulty of keeping track of their spending.*

2.2 The Pain of Paying and its Impact on Spending

However, as illustrated by van Ittersum et al. (2010), mental computing strategies to track in-store spending do not only result in a higher spending uncertainty, but also have a lowering impact on total spending. Mental accounting research suggests that consumers have to deal with so-called “pain of paying” during a shopping trip (Loewenstein and Lerner 2003). This refers to perceived emotional distress when thinking about spending money, which is rooted in consumers' perceptions of opportunity costs. To what extent consumers experience pain of paying depends on their internal reference points. Whereas this reference point is always zero (i.e., no spendings) for non-budget shoppers, budget shoppers assess the opportunity costs of their in-store spendings in reference to their self-set budget. Spending an amount of money for one product means there is less scope for other purchases during the same shopping trip (Sheehan and van Ittersum 2018; Soster et al. 2014).

Moreover, van Ittersum et al. (2013) suggest that budget shoppers consider products they do not consider affordable as forgone gains, whereas overspending is associated with losses. Due to individuals' tendency to prefer avoiding losses over acquiring equivalent gains (loss aversion; e.g., Kahneman and Tversky 1979), budget shoppers who are uncertain about their total spending should try to avoid overspending rather than underspending. Therefore, one common strategy is to build a “safety margin” when using mental computing strategies to track in-store spending. This margin represents an amount of money within an explicit budget that a consumer is not willing

to spend in order to minimize the risk of overspending. In other words, the individual subjectively evaluates the risk of exceeding the allocated budget with a monetary value (van Ittersum et al. 2013). In that regard, we expect that a RSF decreases consumers' perceived spending uncertainty and makes usage of a "safety margin" obsolete (Sheehan and van Ittersum 2018; van Ittersum et al. 2013). In consumers' perception, this is equivalent to an increase in their financial leeway. Hence, we hypothesize that a RSF mechanism will also have a significant positive impact on participants' in-store spending generosity:

Hypothesis 2 *RSF a) decreases underspending and b) increases total spending among budget shoppers.*

2.3 Impact on Product Choices

In order to avoid negative consequences associated with pain of paying and overspending, budget shoppers tend to be very price salient and price sensitive (Raghubir and Srivastava 2008; van Ittersum et al. 2013). Consequently, both factors alter budget shoppers' relative spending-patterns: Budget shoppers are less likely to buy the relatively expensive products within a product category. Accordingly, they often find themselves in situations where they have to find lower-priced substitutes for the higher-priced products that they could afford without budget constraints (Sheehan and van Ittersum 2018); for instance, replacing a preferred organic cheese with a similar lower-priced but non-organic product of the store brand.

However, initial empirical results suggest that this behavioral pattern can be influenced by using RSF to draw consumers' attention to their in-store spending (Sheehan and van Ittersum 2018; van Ittersum et al. 2007; van Ittersum et al. 2013). For instance, within two lab and one field study van Ittersum et al. (2013) find evidence that budget shoppers who received a RSF increased their spending by reducing the share of lower-priced products from store brands and by increasing the share of higher-priced products from well-known national brands in their shopping cart in return. This decline in consumers' price sensitivity can be explained by the fact that budget shoppers interpret their newly won financial leeway as an unexpected gain. Yet, besides these initial empirical insides (e.g., Sheehan and van Ittersum 2018; van Ittersum et al. 2010), it is fairly unknown what kind of products profit from this "budget boost".

Due to the widely cited intention-behavior gap for organic food purchases, it is worth investigating whether this effect can also be harnessed to promote sustainable consumption. One expression of this phenomenon is that many consumers express highly positive intentions with regard to consuming products of organic origin. In contrast, these purchasing intentions are not mirrored in the sales figures for organic products. Even though sales of organic products have increased fivefold in the last two decades, the global market share of organic products totals on average less than five percent (Frank and Brock 2018). Because of contradictory empirical results and a lack of behavioral data, much of this gap remains unexplained. The existing literature on consumer research largely deals with reported or hypothetical purchases, while empirical evidence on incentive-compatible purchase behavior is lacking (Frank and Brock 2018).

However, a frequently reported reason hindering consumers from buying organic food is its price (Aschemann-Witzel et al. 2013; Buder et al. 2014). Indeed, considering conventional food as a benchmark, there is a price premium for organic food on average. Based on the theoretical and empirical foundations from spending tracking literature, we expect that without receiving RSF, budget shoppers – and especially those with positive purchasing intentions for organic products – are more price-sensitive than they need to be. Thus, a real-time spending feedback could support consumers to follow their preferences and increase organic product sales. We will analyze this effect based on the following hypothesis:

Hypothesis 3 *RSF increases a) the number organic food items purchased, b) the share of organic items in the shopping cart, as well as c) the absolute and d) relative spending on organic food items among budget shoppers.*

3 Design and Procedure

The experiment took place in a university laboratory of behavioral economics in September 2019. Using a pool of university employees and students, participants were recruited via the web-based online recruitment system ORSEE (Greiner 2004). Participants did not have any prior knowledge about the experiment and were allowed to participate in only one session. Each session lasted approximately 30 minutes. In total, 97 participants took part in this experiment. In total, 100 subjects participated.

We excluded three participants' data with overspending of more than 50 percent (above 45 Euro) due to the assumption that they had either not understood their budget or made unintentional errors that resulted in major overspending. For participation, each subject received a guaranteed payment of 5 € (show-up fee).

3.1 Shopping Task

The task was to shop for products on a predetermined shopping list in our experimental online grocery store. The shopping list contained eight broad and common product categories (e.g., bread, meat, cheese). Participants were asked to shop for at least two products per category and to choose products they would actually like to purchase. While shopping, subjects could view their shopping list at any time by clicking a special button labeled "shopping list" (see Figure ??). The shopping list was randomized to avoid having the order of the shopping list affect their purchasing patterns. If subjects tried to end their shopping trip before all required products had been added to shopping cart, a notification message appeared. This notification indicated the categories from which products were still missing and asked participants to continue their shopping. Moreover, participants were informed that they receive a budget of 30 € for their shopping. Other than that, they received no further directions on what to shop for.

In order to increase the external validity of our experimental design and encourage the expression of true preferences, the experiment was incentive compatible: Subjects were informed that if they do not overspend, they would be part of a lottery at the end of the experiment. In this case, there was a 20 percent chance of winning their purchase in form of an actual shopping cart. Moreover, participants were told that if they do not spend their entire budget, in case of winning they would receive their purchased groceries and the remaining budget in cash. In doing so, we created a set-up in which subjects had to seriously consider how to spend their money. Because all participants had a 1-in-5 chance to win all product selected during the experiment, there was an incentive to reveal their preferences truthfully. Choosing a product that one does not like is therefore linked to the risk of having to take it home at a price that is higher than the value for the consumer and vice versa.

After all participants finished the task, they were given a post-experimental questionnaire. Next, a computer program randomly determined which participant won his/her shopping cars. Last, all participants received a notification about whether or

not they had won the lottery. All winners received a verification code that allowed them to pick up their purchases at a downtown supermarket. Within this basic set-up, participants were randomly assigned to either a baseline or a treatment condition described below.

No RSF (Baseline): Participants in this treatment did not receive any RSF (see left side of Figure 1). The implemented shopping cart application displayed a small dynamic cart icon in the upper right corner of the product page header. This dynamic icon provided three basic functionalities for consumers: (i) it notified him/her when a new item was added to the shopping cart, (ii) it showed the number of total items in the cart, and (iii) it provided a link to the fullpage cart. On this separate page, subjects had the opportunity to edit their cart before continuing to shop or proceeding to the checkout page. Again, only the names and the amount of their chosen items were displayed on the full-page cart. Consequently, participants' total spending was displayed for the first time at checkout, where editing was no longer possible.

RSF via VSC (Intervention): The intervention treatment was identical to the baseline, except for the difference that subjects found all details of their potential transactions (including prices) in the full-page cart (see right side of Figure1). There, they also had the opportunity to edit their cart before continuing to shop or proceeding to the checkout page.

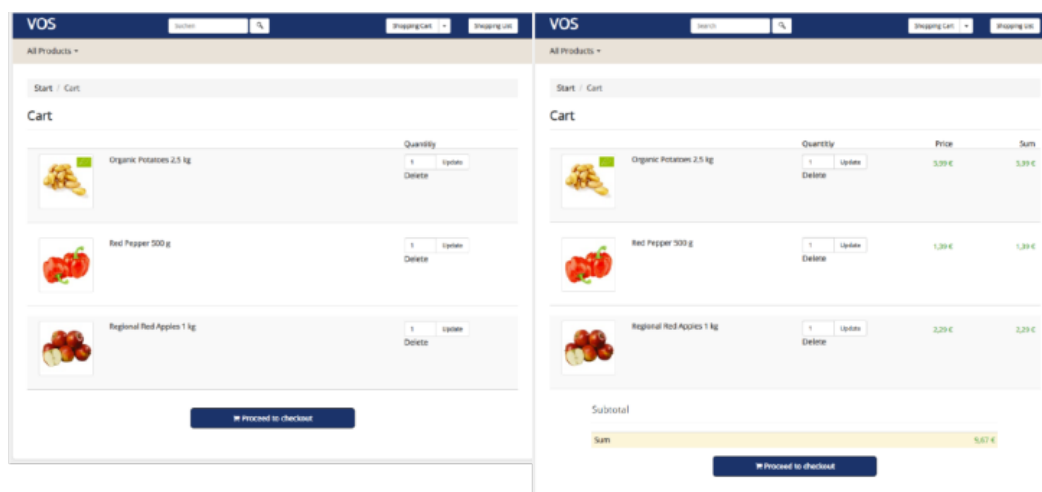


Figure 1 — No RSF Baseline (Left) and RSF via VSC Intervention (Right) Screenshots.

Note. Screenshots exhibit the virtual shopping cart and have been translated from German.

3.2 Interventions

3.3 Measures

As primary outcomes, participants purchasing data – namely the total spending and the number of products in each participant’s shopping cart, the share of organic products bought as well as the percentage of revenue spent for organic products – was recorded. Moreover, we used the post-experimental questionnaire (see Table 1) to ask participants about the effort they had put in for the task (TE). Additionally, three items asked participants about their perception on the task’s difficulty (TD1-3). For our analysis, we conducted a confirmatory factor analysis (CFA) for TD1-3 and – with respect to the factor loading values – calculated a weighted sum index score out of these three items (TDI).

Code	Item	Factor loading Bold values indicate loadings on OPI	Source Item based on
TE	I tried extremely hard to keep track of the total price of my shopping cart.	.	van Ittersum et al. (2013)
Task Difficulty (TDI): CR = 0.820; AVE = 0.604			
TD1	The online supermarket supported me to keep track of my spending.	0.838	Donthu and Gilliland (1996); Davis (1989)
TD2	It was very cumbersome for me to keep track of my total spending.	0.695	
TD3	Overall, I found it easy to keep track of my spending. ^R	0.797	
Organic Purchasing Intention (OPI): CR = 0.852; AVE = 0.659			
<i>Attitude (ATI): CR = 0.942; AVE = 0.845</i>		0.908	
AT1	I like the idea of purchasing organic food.	0.892	Paul et al. (2016)
AT2	Purchasing organic food is a good idea.	0.908	
AT3	I have favorable attitude towards purchasing organic versions of foods.	0.957	
<i>Subjective Norm (SDI): CR = 0.912; AVE = 0.777</i>		0.696	
SN1	Most people who are important to me think I should purchase organic food when doing grocery shopping.	0.901	Paul et al. (2016)
SN2	Most people who are important to me would want me to purchase organic food when doing grocery shopping.	0.930	
SN3	People whose opinion I value would prefer that I purchase organic products.	0.813	
<i>Perceived Behavioral Control (PBCI): CR = 0.782; AVE 0.554</i>		0.824	
PBC1	If it were up to me, I would surely buy organic food.	0.882	Paul et al. (2016)
PBC2	I see myself as capable of purchasing organic food in the future.	0.685	
PBC3	I have sufficient resources, time, and willingness to purchase organic food.	0.616	

Table 1 — Items, measures, and factor loadings.

Note. All items have been measured on a 7-point Likert scale (1 = “I do not agree at all” and 7 = “I totally agree”). R = Reverse-Coded Item.

Furthermore, based on the Theory of Planned Behavior, three sub-constructs were measured: attitudes towards organic products (ATI), the subjective norm (i.e., social influences of the interviewees' social and peer group; SNI) and the perceived behavioral control to actually shop organic products (PBCI). For these sub-constructs, a CFA was conducted as well, revealing a sufficient internal validity of our three respective items per sub-construct (DiStefano et al. 2009). The respective items of a sub-construct were sum-weighted according to their factor loadings as well. The three sub-constructs, in turn, served as the basis for our overall organic purchasing intention (OPI) index. Therefore, we conducted a second CFA, this time between the three sub-constructs and one overall OPI index. The three sub-constructs demonstrated sufficient composite reliability to form an overall OPI and the OPI was, again, calculated by a weighted sum score of the three sub-constructs, with the factor loadings serving as weights (DiStefano et al. 2009). Table 1 summarizes the constructs including their composite reliability (CR) and average variance (AVE) extracted, as well as the items and measured factor loadings. All constructs exhibited sufficient CR ($> .80$) and AVE ($> .50$) with respect to the levels proposed by Urbach and Ahlemann (2010).

4 Results

All participants stated that they had made a serious effort to keep track of their spending (TE: T0 = 5.19 vs. T1 = 5.27). However, in line with our first hypothesis, participants who received no real-time spending feedback perceived this task on average as much more difficult (TDI: T0 = 5.45 vs. T1 = 4.01). A Mann-Whitney-U-Test determined that the observed difference in perceived task difficulty to be statistically significant ($U = 506.000$, $p < 0.001$) for a positive and large effect ($d = 1.127$; Cohen 1992). Thus, the findings support our first hypothesis (H1). At first glance, this difference in perceived task difficulty does not manifest in spending differences between shoppers who received real time-spending feedback and those who did not. Under both conditions, participants spent on average about 86 percent of their budget (T0 = 25.8 € vs. T1 = 25.9 €). Nevertheless, there are large differences in the level of under- and overspending between our baseline and our intervention condition (see Figure 2). While all but one subject who received RSF (T1) remained within budget, 25 percent of subjects in T0 overspent. Moreover, the average underspending was significantly lower if subjects received RSF (T0 = 6.49 € vs. T1 = 4.15 €; $U = 572.000$, p

< 0.05). According to Cohen (1992) the effect size $d = 0.601$ corresponds to a medium effect. Moreover, an ordinary least squares regression confirmed that participants' underspending increases the more difficult they consider the task ($B = 0.61$, $p < 0.05$). These results support H2a but not H2 b.

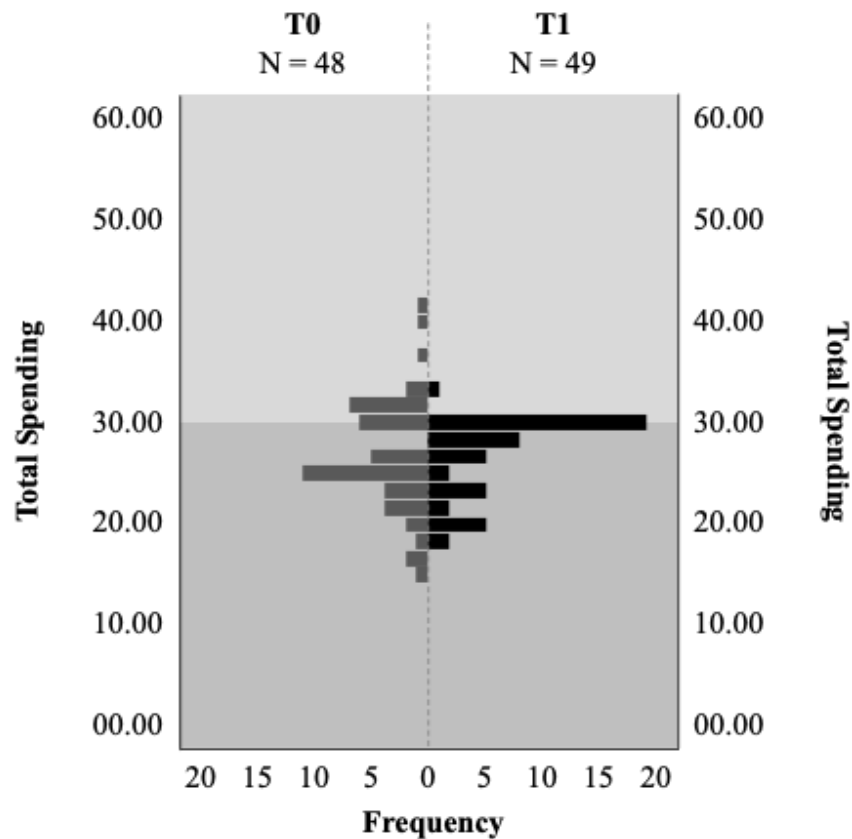


Figure 2 — Subjects Total Spending per Treatment.

Next, we focus on the effect of a RSF on consumers' product choices: First, it is noticeable that participants have – on average – purchased neither significantly more items in total ($T0 = 19.8$ vs. $T1 = 20.0$; $U = 1163.000$; $p > 0.1$) nor more different items ($T0 = 18.9$ vs. $T1 = 18.7$; $U = 1090.000$; $p > 0.1$) in response to the RSF. However, with regard to organic food, it can be observed that all participants purchased on average only a few organic items (2.3), representing on average 12.1 percent of their purchase. However, as shown in Table 2, the data indicate significant differences between our control and our interventions condition for all four observed dependent variables. In response to the RSF, participants purchased on average significantly more organic food items in absolute (number of organic items purchased: $T0 = 1.83$ vs. $T1 = 2.77$; U

= 867.500, $p < 0.05$) and relative terms (share of organic products in total shopping cart: T0 = 9.6% vs. T1 = 14.6%; $U = 1483.500$, $p < 0.05$). Furthermore, they spent on average significantly more money on organic food items in comparison to their total spending (T0 = 10.7% vs. T1 = 17.9%; $U = 1511.000$; $p < 0.05$) as well as in absolute terms (T0 = 2.88 € vs. T1 = 4.53 €; $U = 1480.500$, $p < 0.05$).

However, these differences could also result from an unequal distribution of subjects with strong (weak) purchasing intentions for organic food items between both observed groups. Hereby, we considered individuals to have strong organic purchasing intentions if they scored an OPI value of 5 and higher on a scale from 1 to 7. Indeed, there are more subjects with strong purchase intentions for organic food items in the intervention group (T0 = 37.5% vs. T1 = 53.1%). To control for this imbalance, we conducted two separate Mann-Whitney pairwise comparisons for our treatment conditions within the group of subjects with weak and strong purchasing intentions for organics (see Table 2).

Independent Variable	Overall	Weak OPI	Strong OPI
ON: Number of organic food items	$M_{T0} = 1.83$ $M_{T1} = 2.77$ $U = 572.000$ $p = 0.028^{**}$ $d = 0.45$	$M_{T0} = 1.10$ $M_{T1} = 2.09$ $U = 254.000$ $p = 0.0038^{**}$ $d = 0.575$	$M_{T0} = 3.06$ $M_{T1} = 3.38$ $U = 212.000$ $p = 0.596$ $d = 0.159$
OS: Share of organic food items	$M_{T0} = 9.6\%$ $M_{T1} = 14.6\%$ $U = 1483.500$ $p = 0.025^{**}$ $d = 0.462$	$M_{T0} = 5.9\%$ $M_{T1} = 10.8\%$ $U = 456.500$ $p = 0.042^{**}$ $d = 0.572$	$M_{T0} = 15.9\%$ $M_{T1} = 17.9\%$ $U = 266.000$ $p = 0.444$ $d = 0.232$
TS: Total spending on organic food items (€)	$M_{T0} = 2.88$ $M_{T1} = 4.53$ $U = 1480.500$ $p = 0.027^{**}$ $d = 0.458$	$M_{T0} = 1.87$ $M_{T1} = 3.09$ $U = 455.500$ $p = 0.044^{**}$ $d = 0.566$	$M_{T0} = 4.57$ $M_{T1} = 5.80$ $U = 275.000$ $p = 0.321$ $d = 0.302$
RS: Relative spending on organic food items	$M_{T0} = 10.7\%$ $M_{T1} = 17.9\%$ $U = 1511.500$ $p = 0.015^{**}$ $d = 0.507$	$M_{T0} = 7.1\%$ $M_{T1} = 12.5\%$ $U = 453.000$ $p = 0.049^{**}$ $d = 0.552$	$M_{T0} = 16.6\%$ $M_{T1} = 22.8\%$ $U = 290.000$ $p = 0.180$ $d = 0.411$

Table 2 — Results of the pairwise comparisons between treatments.

Note: * significant at 0.1; ** significant at 0.05; *** significant at 0.01

For the first group (weak organic purchasing intentions), RSF had a medium-sized and significant effect for all four independent variables ($p < 0.05$). However, in the

second group (strong organic purchasing intentions) the Mann-Whitney-U-Test showed small but not statistically significant effects of RSF for all of the observed independent variables. Overall, the findings do not fully support H3a-H3d. Nevertheless, the results indicate that treatment effects are particularly apparent for subjects with low OPI.

5 Discussion And Conclusion

The aim of our study was to explore the impact of real time spending feedback (RSF) on consumers' grocery spending and (organic) food choices. To this end, we developed an online shopping simulation that allowed us to analyze the effects of modifications in the user interface. Within the framework of an incentive compatible experiment, we were able to record the real shopping behavior of participants within a realistic shopping environment.

Result 1 *In line with Sheehan and van Ittersum (2018) and van Ittersum et al. (2013), we found that RSF prevented overspending and significantly reduced underspending. RSF seems to allow budget shoppers to remain within and fully utilize their budget more effectively. The positive effects of the intervention in reducing their perceived task difficulty (i.e., staying within budget) were clearly tangible for the participants.*

Result 2 *Confirming previous research, this study demonstrated that having a favorable attitude towards organic food and intending to buy it does not necessarily translate into fully organic shopping carts. Despite the fact that shoppers preferring organic foods did buy significantly more products of organic origin, the overall shopping cart share remained remarking low. Hence, the much-cited intention-behavior gap seems to be persistent.*

Result 3 *Theoretically, we had assumed that the gap might be mitigated by this digital nudging intervention. In this light, it seems surprising that effects become particularly apparent for individuals with weak purchasing intentions for organic goods whilst effects for those favorable of organics remain small and not significant.*

In light of these results, the question arises why the effects do not translate to the “target group” of individuals with already strong OPI. One possible explanation may be rooted in the design of the virtual online supermarket: To render the shopping experience as realistic as possible, we mirrored the product range of Germany’s leading online supermarket (approx. 5,200 items). Similar to reality, the share of organic products was much smaller than that of conventional items (approx. five percent). As a result, not every item had a perfect substitute of organic origin. Within certain sub-categories, that might have led to a trade-off situation where shoppers had to select between their general attitude towards organic products and their product-specific preferences (e.g., yogurt not being available in their favorite flavor and organic at the same time). This explanation is in accordance with the results provided by Buder et al. (2014): Using a qualitative approach, the authors point out insufficient availability of organic products to be one reason for the existence of so-called “regular” organic buyers (i.e., consumers who spent on average less than half of their budget for organic products). However, within a free simulation experimental approach it is analytically difficult to capture which factor prevails in every single purchasing decision. In one of our follow-up experiments, we plan to tackle this analytical difficulty by focusing on the decision-making process when it comes to substituting a food item for another. By adapting the product selection to offer less variety but more suitable organic alternatives in each (sub) category, we will be able to reduce this situational trade-off.

An alternative interpretation of our key results is that subjects hardly made use of the technical possibilities online shopping provides, such as adjusting the shopping cart without much effort and at any time during the shopping process. Instead, participants still adopted a “straightforward shopping behavior” similar to that in a brick-and-mortar supermarket. In particular, this could be true for customers who have not yet had much experience with shopping for groceries online. According to a representative PWC survey, this applies to a majority of German consumers (PWC n.d.). Consequently, RSF shoppers who realize the extent of their left-over budget may not replace products that are already in their shopping cart with other items but add further products towards the end of the shopping trip. Moreover, previous studies suggest this leftover budget is considered a monetary “windfall” and is spent more frivolously for hedonic and relatively expensive items consumers would not have bought otherwise (van Ittersum et al. 2013). Organic food items could meet these criteria as purchasing them becomes more and more the expression of a respected sustainable lifestyle and they are generally regarded to be expensive yet

qualitatively superior (Bravo et al. 2013). By analyzing the outputs (i.e., the final shopping cart), decision-making happens largely within a black box. Future research should move towards a detailed analysis of in-store behavior to scrutinize which products are substituted at which stage of the shopping trip and for which alternatives. Finally, one must critically reflect on the composition and size of our pure student sample and the respective reliability of the results. However, our experimental virtual online supermarket application is easily scalable to more representative consumer panels. Future studies should use this opportunity to increase the external validity of the empirical results. In sum, this study demonstrates and further underlines the fruitful potential of digital nudging to improve individuals' decisions about grocery purchasing decisions. Utilizing the additional freedom that digital environments provide for choice architects, our results show how seemingly small interventions can already demonstrate a significant effect on (organic) grocery purchasing choices.

Author Contributions

Conceptualization: T.B.L., N.E.; Data curation: T.B.L., N.E., S.L.; Formal analysis: N.E.; Investigation: N.E., T.-B.L.; Methodology: T.B.L., N.E.; Project administration: T.B.L., N.E.; Supervision: none; Validation: T.B.L., N.E.; Visualization: N.E.; Writing—original draft: N.E.; Writing—review editing: T.B.L., M.E., and S.L.

References

- Anesbury, Z., M. Nenycz-Thiel, J. Dawes, and R. Kennedy (2016). "How do shoppers behave online? An observational study of online grocery shopping". In: *Journal of Consumer Behaviour* 15.3, pp. 261–270.
- Aschemann-Witzel, J., K. G. Grunert, H. C. M. van Trijp, S. Bialkova, M. M. Raats, C. Hodgkins, et al. (2013). "Effects of nutrition label format and product assortment on the healthfulness of food choice". In: *Appetite* 71, pp. 63–74.
- Baker, J., D. Grewal, and A. Parasuraman (1994). "The Influence of Store Environment on Quality Inferences and Store Image". In: *Journal of the Academy of Marketing Science* 22.4, pp. 328–339.
- Bénabou, R. and J. Tirole (2004). "Willpower and personal rules". In: *Journal of Political Economy* 112.4, pp. 848–886.
- Benn, Y., T. L. Webb, B. P. I. Chang, and J. Reidy (2015). "What information do consumers consider, and how do they look for it, when shopping for groceries online?" In: *Appetite* 89, pp. 265–273.
- Bhatnagar, A., S. Misra, and H. R. Rao (2000). "On risk, convenience, and Internet shopping behavior". In: *Communications of the ACM* 43.11, pp. 98–105.
- Bravo, C. P., A. Cordts, B. Schulze, and A. Spiller (2013). "Assessing determinants of organic food consumption using data from the German National Nutrition Survey II". In: *Food Quality and Preference* 28.1, pp. 60–70.
- Buder, F., C. Feldmann, and U. Hamm (2014). "Why regular buyers of organic food still buy many conventional products". In: *British Food Journal*.
- Chandon, P. and B. Wansink (2012). "Does food marketing need to make us fat? A review and solutions". In: *Nutrition reviews* 70.10, pp. 571–593.
- Cohen, D. A. and S. H. Babey (2012). "Contextual Influences on Eating Behaviors: Heuristic Processing and Dietary Choices". In: *Obesity reviews : an official journal of the International Association for the Study of Obesity* 13.9, pp. 766–779.
- Demarque, C., L. Charalambides, D. J. Hilton, and L. Waroquier (2015). "Nudging Sustainable Consumption: The Use of Descriptive Norms to Promote a Minority Behavior in a Realistic Online Shopping Environment". In: *Journal of Environmental Psychology* 43, pp. 166–174.
- DiStefano, C., M. Zhu, and D. Mindrila (2009). "Understanding and using factor scores: Considerations for the applied researcher". In: *Practical Assessment, Research, and Evaluation* 14.1, p. 20.

- Escaron, A. L., A. M. Meinen, S. A. Nitzke, and A. P. Martinez-Donate (2013). "Supermarket and grocery store-based interventions to promote healthful food choices and eating practices: a systematic review". In: *Preventing chronic disease* 10.
- Frank, P. and C. Brock (2018). "Bridging the intention-behavior gap among organic grocery customers: The crucial role of point-of-sale information". In: *Psychology & Marketing* 35.8, pp. 586-602.
- Greiner, B. (2004). "An online recruiting system for economic experiment". In: *Forschung und wissenschaftliches Rechnen 2003*. Ed. by K. Kremer and V. Macho. GWDG Bericht. Göttingen, pp. 79-93.
- Hartmann-Boyce, J., F. Bianchi, C. Piernas, S. P. Riches, K. Frie, R. Nourse, et al. (2018). "Grocery store interventions to change food purchasing behaviors: a systematic review of randomized controlled trials". In: *The American journal of clinical nutrition* 107.6, pp. 1004-1016.
- Heath, C. and J. B. Soll (1996). "Mental Budgeting and Consumer Decisions". In: *Journal of Consumer Research* 23.1, pp. 40-52.
- Huyghe, E., J. Verstraeten, M. Geuens, and A. van Kerckhove (2017). "Clicks as a Healthy Alternative to Bricks: How Online Grocery Shopping Reduces Vice Purchases". In: *Journal of Marketing Research* 54.1, pp. 61-74.
- Kahneman, D. and A. Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk". In: *Econometrica* 47.2, p. 263.
- Lembcke, T.-B., N. Engelbrecht, A. B. Brendel, and L. Kolbe (2019). "To Nudge or not to Nudge: Ethical Considerations of Digital Nudging Based on Its Behavioral Economics Roots". In: *In Proceedings of the 27th European Conference on Information Systems (ECIS), Stockholm & Uppsala, Sweden, June 8-14, 2019*.
- Loewenstein, G. and J. S. Lerner (2003). "The role of affect in decision making". In: *Handbook of affective science* 619.642, p. 3.
- Martín, J., F. Pagliara, and C. Román (2019). "The Research Topics on E-Grocery: Trends and Existing Gaps". In: *Sustainability* 11.2, p. 321.
- PWC (n.d.). *Online-Lebensmittelhandel vor dem Durchbruch in Deutschland (The Advent of Online Grocery in Germany)*. URL: <https://www.pwc.de/de/handel-und-konsumguter/pwc-studie-online-lebensmittelhandel-2018.pdf>.
- Raghubir, P. and J. Srivastava (2008). "Monopoly money: The effect of payment coupling and form on spending behavior". In: *Journal of experimental psychology: Applied* 14.3, p. 213.

- Sheehan, D. and K. van Ittersum (2018). "In-Store Spending Dynamics: How Budgets Invert Relative-Spending Patterns". In: *Journal of Consumer Research* 45.1, pp. 49–67.
- Shin, J., H. Qh, and A. Kankanhalli (2019). "Effects of Expert- and User-Generated Evaluations on Food Product Choices via a Food Literacy App". In: *Fortieth International Conference on Information Systems*.
- Soster, R. L., A. D. Gershoff, and W. O. Bearden (2014). "The bottom dollar effect: the influence of spending to zero on pain of payment and satisfaction". In: *Journal of Consumer Research* 41.3, pp. 656–677.
- Thaler, R. (1985). "Mental accounting and consumer choice". In: *Marketing Science* 4.3, pp. 199–214.
- Thaler, R. H. (2016). "Behavioral Economics: Past, Present, and Future". In: *American Economic Review* 106.7, pp. 1577–1600.
- Thaler, R. H., C. R. Sunstein, and J. P. Balz (2014). "Choice Architecture". In: *SSRN Electronic Journal*.
- Turley, L. and R. E. Milliman (2000). "Atmospheric Effects on Shopping Behavior". In: *Journal of Business Research* 49.2, pp. 193–211.
- Urbach, N. and F. Ahlemann (2010). "Structural equation modeling in information systems research using partial least squares". In: *Journal of Information technology theory and application* 11.2, pp. 5–40.
- van Ittersum, K., J. M. E. Pennings, B. Wansink, and H. C. M. van Trijp (2007). "The validity of attribute-importance measurement: A review". In: *Journal of Business Research* 60.11, pp. 1177–1190.
- van Ittersum, K., J. M. Pennings, and B. Wansink (2010). "Trying Harder and Doing Worse: How Grocery Shoppers Track In-Store Spending". In: *Journal of Marketing* 74.2, pp. 90–104.
- van Ittersum, K., B. Wansink, J. M. Pennings, and D. Sheehan (2013). "Smart Shopping Carts: How Real-Time Feedback Influences Spending". In: *Journal of Marketing* 77.6, pp. 21–36.
- Willett, W., J. Rockström, B. Loken, M. Springmann, T. Lang, S. Vermeulen, et al. (2019). "Food in the Anthropocene: the EAT–Lancet Commission on healthy diets from sustainable food systems". In: *The Lancet*.

CHAPTER VI

Ethical Dimensions in Digital Nudging: Reflections on Freedom of Choice, Transparency and Goal-Oriented Justification

Tim-Benjamin Lembcke, Nils Engelbrecht, Alfred B. Brendel, and Alan R. Dennis

The design and implementation of digital nudging raise several important ethical considerations. In this paper, we discuss three ethical considerations of nudging in non-digital environments and transfer this discourse into the digital world: (i) transparent disclosure of digital nudges, (ii) preserving an individual's freedom of choice and (iii) user and social goal-oriented justification of digital nudging. We show that these ethical considerations are also a concern in digital environments, and may be even more prominent given the higher degree of freedom in digital environments. Contrary to their analogue counterparts, it is easy to design and craft highly persuasive digital interventions in digital environments. Here, choice architects can draw on a growing toolkit of artificial intelligence-enabled choice architectures that can make nudging decisions autonomously and independently. This may lead to digital nudging being used in biased, comprised, or other unethical ways. We present a framework for understanding ethical considerations in digital nudging and discuss several avenues for future research. With this framework, we support choice architects from politics and business in being more aware of potential ethical considerations surrounding their interventions. This helps (i) to analyze, reflect on, and potentially revise existing interventions as well as (ii) to consciously design and implement ethically sound new interventions.

Keywords Digital Nudging · Digital Choice Architecture · Information Systems Design · Manipulation · Autonomy

Submitted to *Journal of the Association for Information Systems*.

1 Introduction

As digitization permeates our everyday life, more decisions are being made on digital screens such as computers, smartphones or smartwatches. Driven by this digital transformation, scholars have increasingly put forth research on how user interfaces can be designed to steer users' behavior in the desired direction (e.g., Fogg 2003; Oinas-Kukkonen 2013). Here, a growing stream of IS research is focusing on how different behavior change techniques identified in other fields of research such as psychology or behavioral economics can be transferred into the digital sphere (Abraham and Michie 2008; Michie et al. 2013).

One approach called digital nudging has gained attention for taking empirical insights from judgment and decision making research (e.g., Camerer et al. 2004; Fehr and Gächter 2002; Tversky and Kahneman 1974) and integrating them into the digital environment (e.g., Dennis et al. 2020; Hummel et al. 2018; Hummel and Maedche 2019; Weinmann et al. 2016). Proponents of the "Nudge Theory" (Thaler and Sunstein 2008) argue that nudges can address individuals' systematic and predictable deviations from rational behavior to alter individuals' behavior in a manner that is better for the individual or to achieve a socially desirable goal. They view these nudges as very "soft" behavioral interventions as they should be easy to avoid, do not restrict people's choices (e.g., by forbidding any options), and do not significantly change the economic incentives of different alternatives. Nonetheless, nudges are intended to induce individuals to make different economic choices (Sunstein 2015b).

Digital nudging has been the subject of controversy for two very different reasons. First, it is unclear how effective digital nudges are (Hollands et al. 2013; Hummel and Maedche 2019). Second, there is an ethical debate on whether, for what reasons, and under which conditions digital nudges should be used to alter individuals' behavior (e.g., Lembcke, Engelbrecht, Brendel, and Kolbe 2019). While the lack of empirical evidence can easily benefit from further data, ethical aspects are more difficult to address.

Scholars have pointed out that there are several ethical dimensions inherent in nudging (Hagman et al. 2015): most nudges are non-conscious or subliminal manipulations that do not respect an individual's autonomy or desire for informed consent (e.g., Bovens 2009; Hausman and Welch 2010; Nys and Engelen 2017). A large body of conceptual and empirical work has examined these ethical considerations, for example, by formulating ethical guidelines (Meske and Amojo 2020; Renaud and Zimmermann

2018) or by asking how nudged individuals think about the ethical justification of them being nudged (Abhyankar et al. 2014; Bruns 2019; Michaelsen et al. 2021). However, few research papers address the importance of understanding which ethical considerations matter when nudging is transferred into the digital sphere (e.g., Gregor and Lee-Archer 2016; Weinmann et al. 2016). There is a lack of research considering which ethical considerations need to be taken into account given the systematic differences between analogue/offline and digital/online environments. ¹

This research opinion paper aims to address this lack in research and to pick up the call of Weinmann et al. (2016) for more research on the ethics of digital nudging. The following questions guide our article and envision to facilitate our academic discussion:

1. What ethical considerations arise when applying nudging in digital environments (RQ1)?
2. How should researchers and practitioners act on these ethical considerations (RQ2)?

We begin by examining the particularities of digital choice environments and briefly introduce the main concepts of the ethical considerations surrounding digital nudging (Section 2). Based on these theoretical foundations, Section 3 presents three ethical considerations surrounding (analogue) nudging and transfers these considerations into the digital sphere. Finally, in Section 4 we summarize our key findings, offer practical implications and suggest future research directions.

2 Digital Nudging

According to the fundamental definition put forth by Thaler and Sunstein (2008), a “nudge [...] is any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives.” ² The potential for nudging has gained increased attention in the IS domain, leading to the coinciding adaption of “digital nudging” (Dennis et al.

¹See Appendix A.1 for further details.

²See Appendix A.3 for an overview on exemplary definitions of (digital) nudging and Appendix A.6 for an overview of choice architecture techniques.

2020; Weinmann et al. 2016). Initially, digital nudging was defined as “the use of user-interface design elements to guide people’s behavior in digital choice environments” (Weinmann et al. 2016, p. 433). One year later, a slightly more concise definition was introduced, defining digital nudging to be “a subtle form of using design, information, and interaction elements to guide user behavior in digital environments, without restricting the individual’s freedom of choice” (Meske and Potthoff 2017, p. 2589). More recently, Lembcke, Engelbrecht, Brendel, Herrenkind, et al. (2019) argued that — to count as a digital nudge — an intervention needs to work by making use of “cognitive boundaries, biases, routines, and habits in individual and social decision-making” while not limiting the set of alternatives (Lembcke, Engelbrecht, Brendel, Herrenkind, et al. 2019, p. 10).

Long before Thaler and Sunstein (2008) popularised the term “nudge”, retailers were using nudges to influence consumers to spend more. For instance, the grocery cart was invented in 1937 as a nudge to get consumers to buy more food, larger carts were introduced to nudge consumers to buy more, and music nudges consumers to buy different items (Schwartz 2020). Therefore, the innovative novelty of the work of Thaler and Sunstein (2008) is not that they developed completely new intervention techniques but rather that they transferred findings from cognitive and behavioral psychology into a more general idea of an action-guiding decision design. In doing so, they used empirical evidence that even simple modifications of decision contexts can be sufficient to steer behavior in a different direction. The concept of nudging makes it possible to collect a variety of interventions (e.g., from marketing) under the common umbrella of nudging. For instance, there are countless examples of retailers using nudges to change consumer behavior in physical contexts (Johnson et al. 2012).

More recently, digital nudges have begun to receive more attention because digital devices increasingly permeate most areas of daily life (Schneider et al. 2018). However, the strict demarcation between digital and analogue environments has blurred as (i) analogue and digital retail channels are increasingly integrated into one “omnichannel” experience (Shi et al. 2020) and (ii) new technologies like augmented reality have increasingly intermingled these two dimensions (Floridi 2014). Research finds a deeper integration of digital devices into our daily lives and increasing inter-dependencies between individuals’ use of digital devices and their “real world” behavior and vice versa (Leidner et al. 2018).

To conceptualize this, we follow Lembcke, Engelbrecht, Brendel, Herrenkind, et al. (2019) and adapt the term Blended Environments, which is rooted in education

(e.g., Garrison and Kanuka 2004). In its original domain, the connotation “blended” describes an approach that combines online education with traditional face-to-face learning (Boelens et al. 2015). Figure 1 depicts a high-level framework for decisions in blended environments, which can be applied for digital nudging. This framework integrates a process and environment perspective, differentiating between three process steps of (i) an initial situation in which (ii) some intervention is taking place that leads to (iii) some kind of outcome.³ With an increasing augmentation of our analogue lives with digital devices, these three steps can occur in an analogue, digital, or blended decision-making environment.

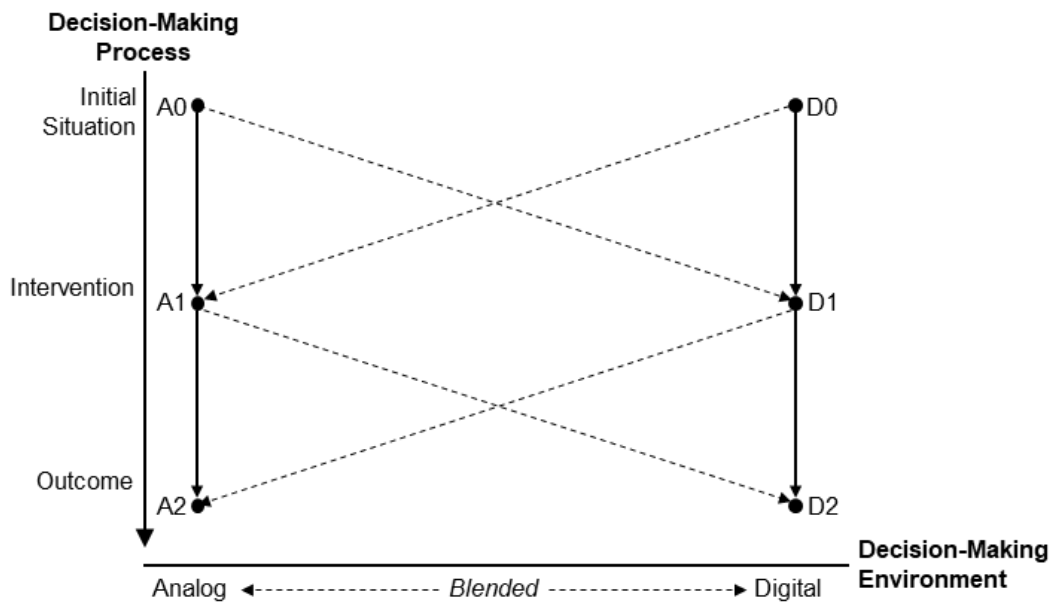


Figure 1 — Blended-Environment Framework for Interventions. Source:Lembcke, Engelbrecht, Brendel, Herrenkind, et al. (2019).Arrows indicate a sequential relationship.

However, this development comes not without consequences for consumers’ behavior. Digital and blended environments differ significantly from their pure analogue counterparts concerning the scope, availability, visual presentation, and information. There is initial evidence that these contextual factors have, in turn, influenced how individuals process information and what decisions they make (see Subsection 2.1).

³See Appendix A.2 for an exemplary instantiation of the Blended Environment Framework for digital nudging.

2.1 IS as Attention-Guiding Environments

The rise of information technology has led to broad overlapping between the digital and analogue world but also rapidly increased the quantity, accessibility and availability of information and choices. However, as theoretical (e.g., Simon 1969) and empirical work (e.g., Qiu et al. 2017) suggest, this abundance of information and choices is not always helpful for decision-making. Instead, phenomena such as information overload (Eppler and Mengis 2004) and choice overload (Scheibehenne et al. 2010) may occur, hindering effective decision making. Moreover, there is evidence that individuals' processing of digital information on a computer screen differs from analogue information such as on a sheet of paper. Empirical evidence illustrates that screen-based reading behavior tends to be non-linear, more selective, and less in-depth and concentrated compared to reading on paper. Consequently, understanding and memorizing such information is more difficult e.g., Liu 2005; Mangen 2008; Oppenheimer et al. 2009).

It is also argued that the advent of digital and connective technologies has accelerated information and communication patterns, favoring shrunken attention spans and contributing to the prevalence of multitasking. For instance, neuro-science suggests that "Digital Natives" (those who grew up in a digital world) generally tend to have more shallow forms of information processing such as browsing, scanning, or skimming compared to "Digital Immigrants" who grew up analogue but have moved into the digital world (Loh and Kanai 2016). Furthermore, as digital devices have become omnipresent in our everyday lives, fluently jumping from one task or conversation to another becomes more common. Individuals are triggered to multitask or consume different contents in parallel, with the borders between the analogue and digital world becoming increasingly blurred. Of course, multitasking can both improve and impair cognition, performance and enjoyment, depending upon how it is done (Dennis et al. 2010; Lee and Choi 2018; Smith et al. 2019; Uncapher and Wagner 2018). However, there is evidence that the mere presence of a digital device like a smartphone can adversely affect individuals' cognitive capacity in terms of working memory capacity (Ward et al. 2017). Further, smartphone notifications can reduce individuals' concentration and task performance

Some of these effects can be balanced by IS, for example, through an appealing user interface design ⁴, alternative ways to present information (e.g., video clips), or by

⁴See Appendix A.3 for an example of display biases which may distract individuals from their rational and conscious decision-making.

providing personalized content and decision support tools. Hence, it is one role of an IS to influence, direct and focus the user's attention (e.g., Banker and Kauffman 2004; Fogg 2003). We consider IS, thus, as a form of attention-guiding environment. In other words, IS are suitable platforms for behavioral interventions aiming to change or guide individuals' behavior like digital nudges.

2.2 IS as Platforms for Digital Nudges

One key difference between the analogue and the digital worlds is how choice architects⁵ can control what users see. In essence, the entire digital world is an artificial artifact that designers can fully influence and control. User interface designers — in their role as digital choice architects — can control so much more than analogue choice architects. It is much easier and cheaper for digital choice architects to design both the digital nudge and the entire information environment where the user finds themselves. In the analogue world, large environmental changes are costly or even impossible (e.g., physical structures like walls, buildings, or plants require significant effort to be adapted or removed). For instance, a store owner can usually only operate within the floor plan and given space. In contrast, a digital choice architect can control everything on the user screen, from the layout of the web shop over the display of the shop items to the variety of products offered. It is also much easier in the digital world to change the information environment for each user so that different users (or groups of users) receive personalized versions of an electronic environment that are customized to influence their behavior better (cf., Benartzi et al. 2017; Simanowski 2016; Mills 2022).

In other words, digital choice architects have a high degree of design freedom and can build on this knowledge to modify elements of the digital decision environment (e.g., by varying the arrangement of objects or changing default settings) in a manner that steers users' behavior in the desired direction. For example, how information is presented and where it is located on a screen (Milosavljevic et al. 2012) may alter users' choices (Chang and Liu 2008; Kahn and Wansink 2004). Phenomena like middle bias, top-left bias, or the visual salience bias are only three examples of so-called “display

⁵Here, we use the term choice architect as a synonym for the person or group of people responsible for the design of a digital artifact and the purposeful use of digital nudging interventions. Depending on the hierarchy in a company, this can include managers (e.g., CIOs) as well as programmer or interface designer.

biases” (Benartzi et al. 2017).⁶ When considering this increased degree of design freedom, the question arises whose interests are served by digital nudging—those of the choice architect, the user, society or some mutual interest.

Digital applications also enable AI-assisted choice architectures, including tools like search algorithms and ranking systems, filtering algorithms, recommender systems, smart virtual assistants and bots or customized content moderation (Kozyreva et al. *in press*). These tools can increase the efficiency of the respective choice architecture and may support users in handling the abundance of online information. In other words, “automated algorithmic systems act as buffers between the abundance of information and the scarcity of human attention” (Kozyreva et al. *in press*, p. 30). However, as we argue later, the opacity and lack of interpretability of complex machine-learning algorithms—i.e., the AI black box problem (Rahwan et al. 2019) — also enables manipulation and may negatively affect individuals’ autonomy.⁷

In sum, there are manifold ways to alter user behavior in digital choice environments. With its potential to manipulate individuals’ behaviors non-consciously, digital nudges may be implemented to cater to the benevolence of users and for a choice architect’s or organizational agenda (e.g., increasing sales and profits). If they had not been already prevalent for analogue nudging (Sunstein 2015b), ethical considerations would be so much more important in digital and blended environments (Weinmann et al. 2016). In digital spheres, users are way more susceptible to being subconsciously influenced due to the general “control of everything” that digital choice architects own (ib.). Likewise, Benartzi et al. 2017 argue that digital choice environments should not simply be regarded as mirrors of their real-world equivalents, leading to a necessary re-examination of ethical considerations in digital nudging.

3 Ethical Considerations in Digital Nudging

The possibilities offered by digital choice architectures are a double-edged sword. The same smart, easy and highly customizable digital choice architecture can be used for both benevolent reasons and ethically questionable “dark patterns” (cf., Gray et al. 2018)⁸ that aim to manipulate user behavior for commercial gain (Mirsch et al. 2017; Schneider et al. 2018; Weinmann et al. 2016).

⁶See Appendix A.3 for a more detailed account on these three biases.

⁷See Appendix A.7 for an overview of AI-assisted digital choice architectures.

⁸See Appendix A.4 for an overview about different types of dark patterns.

Different interpretations exist about what kind of interventions can be subsumed under term “nudging”. For instance, Michalek et al. (2016) differentiate between heuristic-triggering, heuristic-blocking and informational nudges. Heuristic triggering nudges work by addressing individuals’ intuitive and automatic perceptions and exploiting existing biases (biasing) or setting new ones (re-biasing). In contrast, the two other types aim to push away individuals from non-conscious and reactive decision-making by initiating and stimulating a more reflective mode of thinking (de-biasing). However, it is indisputable that heuristic triggering nudges are part of the nudge approach. Hence, in this paper, we follow the suggestion by Hansen and Jespersen (2013) and focus on heuristic-triggering nudges.

There has been criticism that especially this type of nudges infringes on individuals’ autonomy (Hansen and Jespersen 2013; Nys and Engelen 2017; Wilkinson 2013). In response to such criticism, Thaler and Sunstein (2008) argued for creating “rules of engagement that reduce fraud and other abuses [...], and that create incentives to make it more likely that architects will serve the public interest” (p. 240). For this purpose, they proposed three ethical requirements in using nudges: (i) transparency (Sunstein 2015a; Sunstein 2015b), (ii) preserving freedom of choice (Sunstein 2011), and (iii) influencing choices “in a way that will make choosers better off, as judged by themselves” (Thaler and Sunstein 2008, p. 8). However, it is controversial whether such criteria are necessary or sufficient for an ethical application of nudging (Bovens 2009).

This Section discusses three key proposals for ethical nudging in digital environments. Following Thaler and Sunstein (2008), we adopt the user perspective in discussing these ethical considerations, although we do consider other perspectives. For instance, choice architects in a company might arrive at a different view based on a different goal-oriented justification (e.g., a shareholder value perspective). Digital choice architects, thus, may or may not incorporate user-oriented ethical considerations behind nudging.

We will consider each of these three elements in turn: First, for each element, we present a key consideration for the fulfillment of the element. Second, we adapt and discuss these challenges in light of digital contexts to address RQ1 (Which ethical challenges arise when applying nudging mechanisms in digital environments?). If necessary, adaptations or extensions of the ethical concerns are suggested. Furthermore, potential approaches and resolutions to address these ethical concerns are

introduced, answering RQ2 (How can researchers and practitioners act on these ethical challenges?).

3.1 Transparency

Criteria to Evaluate the Degree of Transparency

To avoid a nudge, one must know that they are being nudged. Hence, a sufficient level of transparency is a precondition to make nudges are easy to avoid. First of all, individuals should be informed about the general existence of nudges (type-transparency). Two further criteria are more crucial at the individual level: For each specific intervention, individuals should be effortlessly able to recognize when and where they are subjected to a nudge (token-transparency; Barton and Grüne-Yanoff 2015) and understand how and why the nudge is intended to work (we term it mechanism-transparency).

There has been controversy regarding the circumstances under which these criteria can be considered satisfying (Lepenes and Małecka 2015). Hansen and Jespersen (2013) developed an initial evaluation framework based on these transparency criteria and psychological dual-process models of cognition (e.g., Evans and Stanovich 2013; Kahneman 2012; Stanovich and West 2000). In simple terms, dual-process models divide cognitive processes into separate yet interacting types: An intuitive, automatic, emotion-driven, and quickly operating (e.g., labelled as “Type 1” because it occurs first), and a deliberate, logically calculating, but slower processing (e.g., labelled as “Type 2” because it occurs second; Lades 2014).⁹

Thus, Hansen and Jespersen (2013) identified four types of nudges in a 2x2 matrix differentiating between Type 1 versus 2 as one dimension and transparent versus non-transparent on the other axis (see Figure 2). This approach provides a good starting point understanding a nudge’s transparency level with two caveats. First, hybrid forms may occur: “Some [interventions] may fall into grey zones or seem to qualify for being several types due to their multi-layered structure of mechanisms and long-term dynamics that allow them to wander between categories” (Hansen and Jespersen 2013, p. 20). Second, the perception of a nudge can vary depending on an individual’s

⁹Although the terms “System 1” and “System 2” were popularized by Kahneman (2012), this does not suggest that “the two types of processes are located in just two specific cognitive or neurological systems” (Evans and Stanovich 2013, p. 225). In order to avoid such confusion, we follow the suggestion of Evans and Stanovich (2013) in using the terms “Type 1” and “Type 2” instead.

perspective. The existence, the purpose and the underlying mechanism of a nudge intervention may be fully transparent to one individual while being non-transparent to another due to limited cognitive capacities or other social disadvantages (Gregor and Lee-Archer 2016). The latter brings another important factor into play: Avoiding even transparent external cues created by a nudge intervention might entail cognitive efforts (see Subsection 3.2).

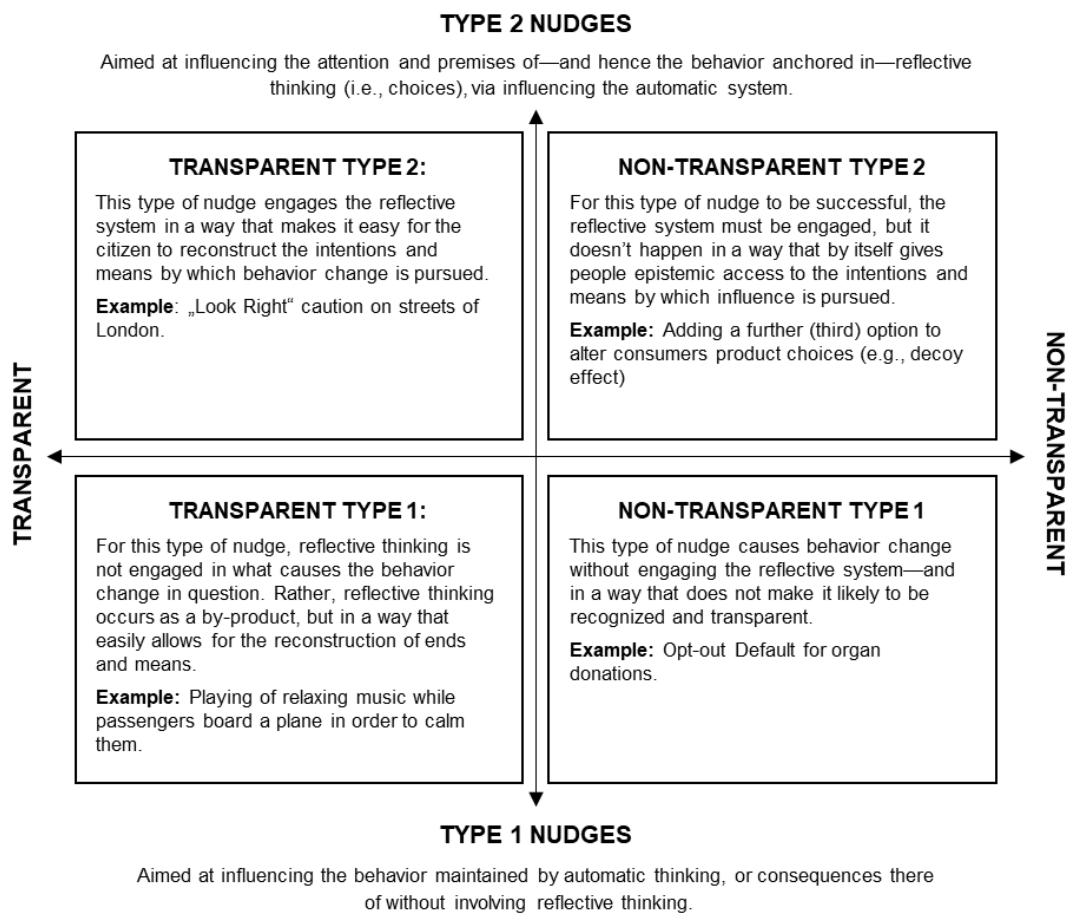


Figure 2 — A Framework of Four Types of Nudge Interventions. Source: Based on Hansen and Jespersen (2013).

Token-Transparency and Mechanism-Transparency in Digital Choice Environments

Token-transparency (when and where users are subjected to a nudge) is relatively straightforward to implement in digital choice environments because IS allows for various forms of notification, alert, consent, dialogues, or the like. Digital choice

architects can design an IS to require consent before proceedings (e.g., users may only continue to use an IS after providing consent). Nonetheless, several practical issues arise: When, where, how and to what extent should users be made aware of an application aiming to influence their behavior via nudging?

Which form and frequency of information are sufficient to make a user aware of a nudge intervention is ultimately an empirical question. Choice architects should evaluate digital nudge's notice-ability in pilot tests before rolling them out to a broader audience. Users should be able to identify nudge interventions using a basic level of attention. However, as transparency also refers to a nudge's visual salience, display biases and hot-spots should be considered. For example, a "transparent" hint may be positioned or displayed in the bottom-right of the screen. From a rational point, this hint would be transparent; however, users have been shown to pay lower attention to the information displayed on the margins of a screen (Drew et al. 2013; see Appendix 4 for further details). Thus, the transparency of a nudge needs to be considered from the user's perspective. If the IS provided notification of a nudge, but users do not notice the notification, then we cannot consider it to be transparent.

However, it can also be annoying for users to be repeatedly reminded of nudges in the name of transparency, especially if nudges may only marginally influence their user behavior. We argue that digital nudges need to exceed a certain threshold of goal-setting and human consideration before they need to be transparent. Hence, to be subject to the requirement of mechanism-transparency, a digital nudge needs to be directed at something more than simply switching to the following webpage or interacting with a website.

The criteria of mechanism-transparency (how and why a nudge is working) can be a more significant challenge given the proliferation of AI-assisted and highly personalized nudges (e.g., Mills 2022; Sætra 2019). Combining insights into the mechanisms behind human decision-making with highly detailed data about individuals' personalities, habits or preferences may help choice architects design highly effective interventions that respect heterogeneity within the target group (Sunstein 2013). However, besides privacy concerns (see Subsection 3.2), it remains an ethical challenge to disclose and explain why, for instance, a machine learning algorithm has chosen a particular nudge intervention (e.g., a default option) or a specific form (e.g., a specific wording) over another (André et al. 2018; Lipton 2018). Even to the experts who built a highly sophisticated algorithm (e.g., based on machine learning and big data), it is sometimes difficult or even impossible to make its decisions transparent, i.e., the

black box problem of AI algorithms (Castelvecchi 2016). In other words, an algorithm being in charge as the choice architect can be “opaque in the sense that if one is a recipient of the output of the algorithm (the classification decision), rarely does one have any concrete sense of how or why a particular classification has been arrived at from inputs [sic]” (Burrell 2016, p. 1).

3.2 Freedom of Choice

Freedom of Choice as an Opportunity Concept

One extensively discussed ethical premise is that each nudge intervention needs to preserve an individual’s freedom of choice (Hansen and Jespersen 2013; Wilkinson 2013). Freedom of choice is mostly interpreted as the absence of obstacles, barriers, or constraints (Thaler and Sunstein 2003) that may hinder individuals from choosing among alternative “non-nudged” options (Marchiori et al. 2017). According to such an “opportunity concept” (Carter 2004), real freedom of choice can only exist if (i) individuals are capable of making “reasonable” decisions and (ii) their options are not restricted.

Hence, to qualify as an ethical nudge, the nudge must be relatively easy to avoid. This presents some ambiguity (Marchiori et al. 2017). Hausman and Welch (2010) argue that this means that choosing a non-nudged option is not significantly more costly in terms of money, time, effort, or social sanctions. Irrespective of the fact that these criteria are hard to operationalize, it is controversial whether these criteria remain adequate to preserve an individual’s autonomy or not (Veetil 2011). Schnellenbach (2012) suggests that if the process of dismissing a transparent digital nudge necessitates extra cognitive effort, it is impossible to design a transparent digital nudge that can be avoided with no cognitive effort.

This presents an important — and potentially unresolvable - consideration. Either the digital nudge is effective (i.e., it changes users’ behavior non-consciously), or it is made transparent and preserves the user’s freedom of choice (i.e., it requires cognitive effort to understand and possibly avoid). Hence, if users want to sustain their transparent freedom of choice, they will need to expend some cognitive effort (Barton and Grüne-Yanoff 2015).

Therefore, ethical nudges include a “cognitive tax”, a cognitive effort that must be invested to understand the nudge. Thus, there is a delicate balance in understanding

how much cognitive tax users need to pay to preserve their freedom of choice and how to balance this cognitive tax against the value derived from nudging (and, as we discuss below, for which stakeholder this value materializes).

Freedom of Choice in Digital Choice Environments

As mentioned above (see Section 2.2), digital nudges can be implemented with little cost (Schneider et al. 2018). Thus, there might be a risk of “over-nudging” individuals (e.g., sending hourly notifications). Over-nudging might increase psychological costs for ignoring the intervention (e.g., a single notification) and elicit counter-reactions. For instance, users might be annoyed by too many reminders and intentionally reject a nudge, even though they approve of its goal in principle (Damgaard and Gravert 2018).

Besides its frequency, the timing and the form of a nudge can influence individuals’ freedom of choice by increasing the cognitive tax. For example, vibration or audio notifications can disrupt other activities and prevent individuals from fully concentrating (Duke and Montag 2017). Thus, inappropriate timing and frequency of notifications may mean individuals need to devote more cognitive capacities to deal with these interruptions, which increases psychological costs (Barton and Grüne-Yanoff 2015). Such notifications may even contribute to smartphone addiction (Duke and Montag 2017). Consequently, architects of ethical digital nudges need to be particularly sensitive towards their responsibility of designing nudge interventions in a non-exploitative, non-prohibiting and choice-preserving way that does not require users to pay too high a cognitive tax.

Simultaneously, digital choice environments could be designed to enhance and preserve individuals’ freedom of choice. Thereby, individuals could be empowered to influence the type, form, and frequency of nudge interventions they face. For example, users could be allowed to configure their IS experience by choosing which kind of mechanisms to which they want to be subjected. In a web shop, this would mean that users can deactivate or personalize several functions like (1) recommender systems (e.g., what other customers have bought, additional items suggested on their shopping cart view), (2) reminders (e.g., push notifications or website layovers), (3) conversational agents (e.g., product specification chatbots), (4) simplification mechanisms (e.g., one-click-buying options that make ordering easier but also more tempting) or (5) website personalization means (e.g., personally welcoming a user).

These personalization scenarios to preserve and freedom of choice also present new challenges. Such an active choice and consent would be required every time users visit new websites, which increases effort. For instance, from an ethical point of view, it would be desirable to inform users how even marginal changes in an online store's product placement (e.g., Breugelmans et al. 2007) or arrangement (e.g., van Herpen and Bosmans 2018) may alter their choices (see also Appendix 4).

However, forcing a consumer to manually choose how to sort the displayed items (e.g., by price, retailer's recommendation) could easily be overwhelming, resulting in less cognitive capacity for other important decisions (Sunstein 2015b). This is particularly true if user preferences are inquired after again and again for each situation and at each visit. Therefore, in an ethically ideal scenario, individuals should be asked once, for instance, upon registration, about their chosen default options that remain active but can be changed at any time. Users should be informed actively about updates and subsequent changes in the IS (e.g., whether new default options were introduced or existing ones changed or removed). The European Union's General Data Protection Regulation (GDPR) can serve as an example of good practice: GDPR requires all website operators to disclose information on privacy practices through cookie notices. In addition, each user can individually determine which form of cookie use they wish to agree to and select these preferences for each website as a future default (Bornschein et al. 2020).

In sum, a certain amount of cognitive effort to understand and manage nudging mechanisms cannot be avoided. This cognitive tax needs to be kept as insignificant as possible. Ideally, it manifests as a one-time effort to ignore a nudge ("do not show again"). If the nudge is presented repeatedly, there must be an argument to justify it as ethical. Nonetheless, for freedom of choice to be met and individuals' autonomy to be respected, nudges have to be clearly and easily identifiable by individuals, i.e., they need to be transparent (Lepenies and Małecka 2015; Sunstein 2015b).

3.3 Goal-Oriented Justification

Nudging for whose Benefit?

Nudges can be used for three goals, not necessarily mutually exclusive (Clavien 2018; Hagman et al. 2015): (1) goals of the choice architect (e.g., nudge users to buy more products), (2) user goals which are in line with users' long-term preferences

(e.g., retirement savings) and (3) social goals (e.g., gender equality). Thaler and Sunstein (2008) argued that ethical justification exists only for user-goal-oriented nudges, nudges that “help individuals steer away from irrational behavior (or bounded rationality), which decreases their long-term well-being” (Hagman et al. 2015, p. 441). Metaphorically speaking, this form of nudge assumes a kind of supportive GPS function by making it easier for people to make the subjectively best decisions for themselves (Sunstein 2015b).

According to Sunstein (2015b), the choice architect’s goals take second place behind users’ and social goals. In contrast, from a corporate perspective, companies have an obligation to serve their shareholders’ interests (shareholder value approach). There is an ongoing debate on how to balance shareholder value (usually the maximization of profits) with concepts like stakeholder management or corporate social responsibility (Jones and Nisbet 2011; Manchiraju and Rajgopal 2017; Martin et al. 2009). Competing interests may arise in commercial transactions, and there may be no singular, “ethically” right way.

We consider the goals of choice architects to be legitimate and argue in favor of introducing them from a different perspective. In addition to all justified criticism of corporate misconduct, we are convinced that we—as a society—are better off with corporate and entrepreneurial actions than without them (Smith 1776). Hence, we propose that choice architects’ welfare compliments the “ethical dyad” (Gray et al. 2012) of user and social welfare, forming an ethical triad that needs to be balanced. In some situations, digital nudges align with all three goals, thus balancing all stakeholders’ interests. Nevertheless, conflicting perspectives and goals can arise, requiring architects to make trade-offs.

Various authors have pointed out that the practical implementation of a user-goal strategy is difficult because choice architects lack knowledge about individual users’ goals (e.g., Barton and Grüne-Yanoff 2015; Rebonato 2014). From an idealized ethical standpoint, tailor-made nudges would be preferred, catering to each user’s individual goals. In analogue environments, the personalization of nudges is quite limited (e.g., the order of products on a particular physical supermarket shelf cannot be displayed to different users in “multiple” fundamentally different arrangements at the same time). In digital environments, such personalization is possible. There is a wealth of social media and search data on users so that companies can design different nudges for different people. Digital choice architects could personalize information displays to best match each user; however, it can be challenging to find options that meet the

heterogeneous preferences of all users. As a trade-off, one solution may be to mass personalize digital nudges, e.g., to suit different user segments.

Because of this heterogeneity in preferences, some nudges may steer some individuals away from their true preferences (Barton and Grüne-Yanoff 2015). Thus, to be ethical, a user-goal nudge has to aim at justifiable goals (Ismaili M'hamdi et al. 2017), such as enabling individuals to avoid “decisions they would not have made if they had paid full attention and possessed complete information, unlimited cognitive abilities, and complete self-control” (Thaler and Sunstein 2008, p. 5).

Most choice architects work for companies, and from a corporate perspective, companies have an obligation to serve their shareholders' interests. Is it ethical to spend shareholders' resources to create a nudge that benefits consumers but harms shareholders' interests? Balancing shareholder value (e.g., profits) and corporate social responsibility is not simple (Jones and Nisbet 2011; Manchiraju and Rajgopal 2017; Martin et al. 2009). Competing interests often arise in commercial transactions, and there may be no singular, “ethical” right way to balance the interests of the company (and its shareholders) against the interests of consumers.

Ethical nudges can also focus on public welfare goals, with proponents for this including Clavien (2018), Hagman et al. (2015), and Thaler and Sunstein (2008). This line of reasoning is based on the utilitarian harm principle (limiting individuals' actions to prevent harm to others) and concerns associated with market failures like negative externalities or public goods (Barton and Grüne-Yanoff 2015). Consequently, in many cases, there can be multiple lines of reasoning. For instance, nudging individuals to reduce their energy consumption could be justified as a user-goal nudge (saving money) and a social welfare nudge (avoiding negative externalities).

Social welfare nudges can likewise be justified if individuals have not actively consented in some way or form (for example, through government programs to increase individuals' health or retirement savings). However, “nudges will leave a trace of moral violation as long as there are reasons to think that some choice architects would find it unbearable to be nudged” (Clavien 2018, p. 6). Even if the socially desirable behavior is in the user's own best interest, some individuals will still refuse (e.g., the resistance to mandatory seatbelt laws). In a similar vein, some users may disagree with the goals of choice architects or disagree with being nudged at all.

Understanding Users' Goals in Digital Choice Environments

Choice architects of digital environments are faced with the same challenges as their analogue peers: how to know what individuals want and how to create interventions suitable for a heterogeneous user group. Thaler and Sunstein (2008) use the phrase “as judged by themselves”, which indicates that there could be an empirical answer to this question. In recent years several studies were conducted to capture the public opinion on nudging in general (e.g., Arad and Rubinstein 2018; Reisch and Sunstein 2016; Sunstein et al. 2018) as well as to investigate ethical judgment from individuals subjected to such interventions (e.g., Abhyankar et al. 2014; Bruns 2019; Michaelsen et al. 2021). Overall, these empirical results are somewhat ambivalent: On the one hand, the data suggests that most nudges receive individuals' support. On the other hand, people still express ethical concerns regarding a possible threat to their freedom of choice and autonomy (Michaelsen et al. 2021).

The particularities of digital choice environments offer architects at least two approaches to address this ambivalence. First, online surveys can be integrated into the digital application to gather feedback from users before, during or after being digitally nudged. This could also allow users to participate in a nudge intervention's designing process (participation). Following offline approaches, where the nudge is jointly designed and proliferated by civil society groups (e.g., Nielsen et al. 2017), it would be conceivable for (online) communities to vote on which behavioral changes they would like to achieve and which nudge intervention should support them.

Second, search engines, social media applications and corporate databases provide user data that enables choice architects to understand users' preferences better. This understanding facilitates the ethical considerations, as the user-goal or social welfare dimension can be more easily justified if the target group is better understood. For instance, this enables choice architects to argue that their intervention is “shared by, or in line with the preferences of individuals impacted by the nudge” (Clavien 2018), p. 6).

AI-assisted choice architectures can also analyze users' data to provide them with customized nudges (Kannan 2017; Taken Smith 2012). For instance, personalized or targeted advertising techniques could be used to show users customized reminders based on their prior behavior (Altmann and Traxler 2014). An interesting area for future research is the adaptation of such digital marketing mechanisms to inform the design and implementation of digital nudges.

These AI-assisted information architectures do not come without risks because their inherent opacity and complexity can be problematic. Users may agree with the output of a tool (e.g., an app that nudges users to exercise more) but may have reservations about the processing of their data (e.g., due to the transferal of personal data to a service provider). Trends like big data and analytics have inspired corporations to become increasingly interested in collecting as much data as possible to reach some competitive advantage (Kubina et al. 2015).

Individuals, however, are often unaware of the data and data sources that a service provider uses (Pasquale 2015). This is particularly relevant as individuals claim deep concerns about their personal data in digital contexts but at the same time tend to disclose more sensitive private information in virtual than in analogue, face-to-face interactions (Acquisti et al. 2015; Brandimarte et al. 2013). This “privacy paradox” shows a disconnect between what individuals claim to value and what their actions show they actually value (Adjerid et al. 2018). It may be that individuals overgeneralize or respond in socially desirable ways when asked about the importance of privacy but make more accurate decisions when asked to disclose a specific piece of information (Goel et al. 2020). Thus, digital choice architects should be cautious in designing AI-assisted information architectures to respect individuals’ autonomy by preserving their ability to keep some elements private (i.e., to assume authorship over your own life; Dworkin 1988).

4 Discussion

Digital nudges are valuable tools for encouraging behavior changes, yet there may be repercussions worthy of discussion as well (for instance, Sunstein 2015b). This study set out to examine the important ethical considerations for nudging in digital and blended environments. We introduced IS as attention-guiding environments in which users are particularly susceptible to being influenced. We highlighted important aspects of digital nudges and provided an introductory background about digital nudging.

To address our first research question (What ethical considerations arise when applying nudging in digital environments?), we built on three considerations widely acknowledged and discussed in psychology or behavioral economics—namely, transparency, freedom of choice and goal-oriented justification. We do not presume these

aspects to be mutually exclusive or collectively exhaustive. Instead, these aspects have to be seen as a complementary and interrelated triad.

As a synthesis of our ethical deliberations, we propose the concept of an ethical digital nudge, integrating the perspectives of Sunstein (2015b) and Lembcke, Engelbrecht, Brendel, Herrenkind, et al. (2019). We define an ethical digital nudge as a digital nudge that:

1. Enables users to easily recognize when they are subject to being nudged (token-transparency) and how and why the nudge is intended to work (mechanism-transparency);
2. Preserves users' freedom of choice and does not make alternatives appreciably more costly in terms of money, time, effort, social sanctions and so forth; and
3. Balances the welfare of choice architects, the welfare of users, and the social welfare of society as a whole.

To examine our second research question (How should researchers and practitioners act on these ethical considerations?), we elaborated on our three considerations in digital choice environments, highlighting similarities and differences to analogue choice environments. To each dimension, we provide theoretical and practical references. As a result, all three considerations need to be orchestrated both in research and practice rather holistically and collectively, as changes in one dimension may entail changes in other dimensions.

Our paper is not free of limitations. First, this paper introduces digital nudging but cannot be exhaustive on the complex ethical discourse surrounding digital nudging. Second, we aimed to facilitate future research with our three key considerations for ethical digital nudging. We recognize that other points of view to digital nudging may stem from a different ethical worldview, goal orientation, or understanding of what constitutes a nudge. Hence, we encourage additional viewpoints and discussions to advance or contrast our opinions. Third, as this paper is conceptual, we can only provide an initial framework on how to operationalize our ethical considerations in research and practice. In sum, we argue that these ethical considerations are becoming more critical as digitization gradually permeates more aspects of our daily lives and influences us both in fully digital and blended environments. Table 1 summarizes our main conclusions. The following subsections build on these to present essential avenues for future research and implications for practice.

Key Considerations	Potential Resolution Approaches	Key Literature
Transparency		
<ul style="list-style-type: none"> ▪ Users shall recognize when and where they might be nudged (<i>token-transparency</i>) ▪ Users shall understand the how and why of a nudge intervention (<i>mechanism-transparency</i>) ▪ The complexity of modern IS rendering full token- and mechanism-transparency a major challenge (e.g., due to AI-assisted IS) 	<ul style="list-style-type: none"> ▪ Increase algorithmic transparency in AI-assisted information architectures ▪ Empirical investigations on the noticeability of a digital nudge ▪ Ensure that a digital nudge is perceivable with a basic level of attention 	<ul style="list-style-type: none"> ▪ Hansen and Jespersen (2013) ▪ Barton and Grüne-Yanoff (2015) ▪ Rahwan et al. (2019)
Freedom of Choice		
<ul style="list-style-type: none"> ▪ Choice options may be restricted (i.e., no <i>freedom of choice</i>) ▪ IS allow for a high degree of freedom for choice architects ▪ Users may become overwhelmed by opting into every digital nudge intervention 	<ul style="list-style-type: none"> ▪ Users can choose among nudge interventions they want to be subject to (e.g., as part of a signup process) ▪ Users are proactively informed and asked to agree with new and updated nudge interventions ▪ Keeping cognitive effort to dismiss a digital nudge as insignificant as possible (e.g., by “never” or “do not ask again” options) 	<ul style="list-style-type: none"> ▪ Hansen and Jespersen (2013) ▪ Thaler and Sunstein (2003) ▪ Carter (2004)
Goal-Oriented Justification		
<ul style="list-style-type: none"> ▪ Digital nudges can pursue the goals of a choice architect (e.g., own benefit or profit), a user (e.g., pro-self), or society (e.g., pro-social) ▪ Challenge for choice architects to know what users want in their individual and best interests ▪ Challenge to design one-size-fits-all digital nudge interventions (e.g., for a publicly available webshop) 	<ul style="list-style-type: none"> ▪ If an individual legitimization may not be feasible, digital nudges could be legitimized via common interests (e.g., public welfare) ▪ Gather feedback from users about their goals before, during, or after being digitally nudged (e.g., online surveys) ▪ Involve target-group users into the design process ▪ Use existing data and surveys to understand the individual preferences of a target group (e.g., similar to market/customer research) 	<ul style="list-style-type: none"> ▪ Clavien (2018) ▪ Hagman et al. (2015) ▪ Ismaili M’hamdi et al. (2017) ▪ Michaelsen et al. (2021)

Table 1 — Summary of Ethical Considerations of Digital Nudging

4.1 Implications for Research

As a consequence of our deliberations on ethical considerations in digital nudging, we propose the following avenues for future research. First and foremost, we see a necessary trade-off in different goals of choice architects, different users and society as a whole (Clavien 2018). Various goals and conflicting perspectives can arise, necessitating further research on reaching an ethically justified equilibrium between different stakeholders. For example, the debate on shareholder value and the ethical obligations of corporations can be applied to digital nudging as well: Is the ethical responsibility for managers to maximize profits in favor of their shareholders less or more important than ethical responsibilities to recognize users’ goals if both cannot

be met simultaneously? We need more research to understand how to balance these conflicting ethical considerations better.

Second, there is a trade-off in the transparency of digital nudging (Hansen and Jespersen 2013). Here, the question is, how we could design easy to use and transparent digital nudges that do not impose a high cognitive tax on users. Every user interaction is not free. Users need to pay a cognitive tax to circumvent digital nudges, and even just to understand the information provided about nudges (Barton and Grüne-Yanoff 2015; Loewenstein and O'Donoghue 2006). For example, suppose a user has come to our site to buy a specific book. Is it ethical to require the user to expend cognitive effort to understand the website's nudging mechanisms if they just want to buy the book and leave? In other words, the provision of information about digital nudges imposes a cognitive tax on the user in the form of cognitive effort that gets in the way of the actual shopping task. Is it ethical to increase the transparency of such digital nudges because this also increases the cognitive tax on the user? Once again, we need more research to understand how to balance these conflicting ethical considerations better.

Third, users may be overwhelmed if their entire freedom of choice is always maintained (cf., Besedeš et al. 2015). For routine shopping, it might not always be helpful to offer every option to each user and make each nudge transparent. The question arises, how such limited choice environments can—from a user perspective — still be designed as ethically sound as possible, or with the least possible restrictions.

Fourth, supposing that we have created a digital choice environment that balances all three considerations, does this ethical implementation change the effectiveness of digital nudges? In making digital nudges more transparent, we likely shift formerly non-conscious and automatic decisions (Type 1) to more effort-full deliberate cognition (Type 2). Does this deliberation change individuals' likelihood to elaborate similar decisions more consciously (Type 2) in the future? Or do users fall back to Type 1, forget about the nudge after learning about it and behave again in the same way as if they would not have been informed? From a meta-perspective, one question is whether digital nudges can be a valuable part of stimulating long-term behavior changes contrary to one-time changes.

Finally, the mechanism-transparency of digital nudges creates a tremendous challenge concerning algorithmic complexity as AI-assisted choice architectures become more and more capable (Burrell 2016). Hence, digital nudging driven by AI and big data and its impact on users' autonomy is an emerging field for further research

(Puaschunder 2017). For instance, future developments in AI may allow choice architects to automatically generate highly individualized digital nudges, potentially moving the decision of when and how to nudge to some kind of black box. This raises the question of digital nudges that may be unethical by design (i.e., through conscious deliberation of the choice architect) or by accident (i.e., through a self-learning AI algorithm making “wrong” decisions).

These issues are challenging and are likely to be present in many digital nudges, so scholars should examine ways to make ethical considerations manageable and practical. Since all issues are hard to address simultaneously, digital choice architects need to pay great attention to design and implementation decisions.

4.2 Implications for Practice

From a practical perspective, how can choice architects design digital nudges that are ethically sound with the necessary tools and information in a form that is ready to use? We propose the following four implications for practice:

First, be deliberate about the design of technologies that nudge individuals. Digital choice architects need to understand what they are doing and why. They also need to understand how they are influencing or nudging users. What are the goals of a digital nudge and whom does it serve? These deliberations become complicated and complex very quickly, but choice architects would be fully aware of their actions and associated consequences in an ideal world.

Second, we cannot all become ethicists or philosophers. The first implication suggests that choice architects should become amateur ethicists and philosophers to understand the ethical considerations of digital nudging truly. This is impractical. From a managerial perspective, we cannot leave the ethical decisions to every individual designer or programmer who finally implements solutions because different choice architects might make different decisions, leading to uneven designs, where one application might be highly ethical and another not. This partly resembles the ethical considerations of user privacy. Initially, companies provided uneven responses to the need for user privacy. Ultimately, companies realized that centralized privacy guidelines, e.g., due to legislation, were necessary. Thus, we believe that organizations should develop and provide coordinated and consistent nudging policies.

These policies can build on existing ethical guidelines (for instance, Meske and Amojo 2020) but should include both management recommendations and practical

design proposals. Otherwise, management risks that digital nudges will be inconsistent regarding their ethical considerations: One does not want to have a portfolio of digital nudges which are, by chance, ethical or not, depending on who was assigned to develop them. Many choice architects come from different cultural backgrounds, both those who are internal to the organization, as well as contractors, which exacerbates the situation.

Third, do not assume you know what your customers want best. It is better to formally evaluate your customers' long-term preferences than over-estimate your assumptions. You should not be too paternalistic as an organization, implying that your goals are your customers' goals. Conduct surveys, studies, workshops or the like, and shift from a product-centred perspective to a user-centred perspective.

Fourth, be prepared to bear the costs of designing ethically sound digital nudges. Incorporating the ethical considerations around digital nudging will incur significant extra deliberation, time and cost for choice architects. Large firms may more easily bear this extra cost, but do costs of the ethical considerations for digital nudging make digital nudging less attractive for small and medium enterprises with limited budgets? Individual firms will need to make their own decisions, but until there are widely accepted guidelines (or legal regulations) for managing the ethical considerations of digital nudging, each firm may go its own way, and there may be unintended consequences: "Caveat emptor - Let the buyer beware."

5 Conclusion

In sum, this paper adds to the ethical discourse on digital nudging. With the increasing digitization and proliferation of IS, there is abusive potential if digital nudging is misused. This paper hopes to spark a greater interest in important ethical considerations within the entire IS community about digital nudging ethics. With multiple conflicting perspectives on the interplay of individuals' behavior, information processing, IS design and ethics, the context of digital choice environments and choice architects is still an emerging area. We have an opportunity to lead IS practice and push IS research to more critical yet constructive ethical reflections harnessing the interdisciplinary strength of the IS community (cf., Rowe 2018).

Author Contributions

Conceptualization: T.B.L., N.E., A.B.B., and A.R.D.; Data curation: none; Formal analysis: none; Investigation: T.B.L., N.E. Methodology: T.B.L., N.E.; Project administration: T.B.L., N.E.; Supervision: A.B.B., A.R.D.; Validation: T.B.L., N.E., A.B.B., and A.R.D.; Visualization: T.B.L., N.E.; Writing—original draft: T.B.L., N.E.; Writing—review editing: T.B.L., N.E., A.B.B., and A.R.D.

References

- Abhyankar, P., B. A. Summers, G. Velikova, and H. L. Bekker (2014). "Framing Options as Choice or Opportunity: Does the Frame Influence Decisions?" In: *Medical decision making : an international journal of the Society for Medical Decision Making* 34.5, pp. 567–582.
- Abraham, C. and S. Michie (2008). "A taxonomy of behavior change techniques used in interventions". In: *Health psychology : official journal of the Division of Health Psychology, American Psychological Association* 27.3, pp. 379–387.
- Acquisti, A., L. Brandimarte, and G. Loewenstein (2015). "Privacy and human behavior in the age of information". In: *Science (New York, N.Y.)* 347.6221, pp. 509–514.
- Adjerid, I., E. Peer, and A. Acquisti (2018). "Beyond the privacy paradox: objective versus relative risk in privacy decision making". In: *MIS Quarterly* 42.2, pp. 465–488.
- Albrechtslund, A. (2007). "Ethics and technology design". In: *Ethics and Information Technology* 9.1, pp. 63–72.
- Altmann, S. and C. Traxler (2014). "Nudges at the dentist". In: *European Economic Review* 72, pp. 19–38.
- André, Q., Z. Carmon, K. Wertenbroch, A. Crum, D. Frank, W. Goldstein, et al. (2018). "Consumer choice and autonomy in the age of artificial intelligence and big data". In: *Customer Needs and Solutions* 5.1-2, pp. 28–37.
- Arad, A. and A. Rubinstein (2018). "The People's Perspective on Libertarian-Paternalistic Policies". In: *The Journal of Law and Economics* 61.2, pp. 311–333.
- Aragón, O. R., M. S. Clark, R. L. Dyer, and J. A. Bargh (2015). "Dimorphous expressions of positive emotion: Displays of both care and aggression in response to cute stimuli". In: *Psychological Science* 26.3, pp. 259–273.
- Atalay, A. S., H. O. Bodur, and D. Rasolofoarison (2012). "Shining in the center: Central gaze cascade effect on product choice". In: *Journal of Consumer Research* 39.4, pp. 848–866.
- Banker, R. D. and R. J. Kauffman (2004). "50th anniversary article: The evolution of research on information systems: A fiftieth-year survey of the literature in management science". In: *Management Science* 50.3, pp. 281–298.
- Barton, A. and T. Grüne-Yanoff (2015). "From libertarian paternalism to nudging—and beyond". In: *Review of Philosophy and psychology* 6.3, pp. 341–359.

- Benartzi, S., J. Beshears, K. L. Milkman, C. R. Sunstein, R. H. Thaler, M. Shankar, et al. (2017). "Should Governments Invest More in Nudging?" In: *Psychological science* 28.8, pp. 1041–1055.
- Besedeš, T., C. Deck, S. Sarangi, and M. Shor (2015). "Reducing Choice Overload without Reducing Choices". In: *Review of Economics and Statistics* 97.4, pp. 793–802.
- Boelens, R., S. van Laer, B. de Wever, and J. Elen (2015). "Blended learning in adult education: towards a definition of blended learning". In:
- Bornschein, R., L. Schmidt, and E. Maier (2020). "The Effect of Consumers' Perceived Power and Risk in Digital Information Privacy: The Example of Cookie Notices". In: *Journal of Public Policy & Marketing* 39.2, pp. 135–154.
- Bovens, L. (2009). "The Ethics of Nudge". In: *Preference Change*. Ed. by T. Grüne-Yanoff and S. O. Hansson. Theory and Decision Library. Dordrecht et al.: Springer, pp. 207–219.
- Brandimarte, L., A. Acquisti, and G. Loewenstein (2013). "Misplaced Confidences". In: *Social Psychological and Personality Science* 4.3, pp. 340–347.
- Breugelmans, E., K. Campo, and E. Gijsbrechts (2007). "Shelf sequence and proximity effects on online grocery choices". In: *Marketing Letters* 18.1, pp. 117–133.
- Brignull, H. and A. Darlington (2020). *Types of Dark Pattern*. URL: <https://darkpatterns.org/types-of-dark-pattern.html>.
- Bruns, H. (2019). "Are Distracted People Easier to Nudge? An Experiment on the Interaction of Cognitive Scarcity and Defaults in a Public Goods Game". In: *An Experiment on the Interaction of Cognitive Scarcity and Defaults in a Public Goods Game* (July 9, 2019).
- Burrell, J. (2016). "How the machine 'thinks': Understanding opacity in machine learning algorithms". In: *Big Data & Society* 3.1, p. 205395171562251.
- Camerer, C. F., G. Loewenstein, and M. Rabin (2004). *Advances in behavioral economics*. Princeton university press.
- Caraban, A., E. Karapanos, D. Gonçalves, and P. Campos (2019). "23 ways to nudge: A review of technology-mediated nudging in human-computer interaction". In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–15.
- Carter, I. (2004). "Choice, freedom, and freedom of choice". In: *Social Choice and Welfare* 22.1, pp. 61–81.
- Castelvecchi, D. (2016). "Can we open the black box of AI?" In: *Nature News* 538.7623, p. 20.

- Chang, C.-C. and H.-H. Liu (2008). "Which is the compromise option? Information format and task format as determinants". In: *Journal of Behavioral Decision Making* 21.1, pp. 59–75.
- Clavien, C. (2018). "Ethics of nudges: A general framework with a focus on shared preference justifications". In: *Journal of Moral Education* 47.3, pp. 1–17.
- Clement, J., J. Aastrup, and S. Charlotte Forsberg (2015). "Decisive visual saliency and consumers' in-store decisions". In: *Journal of Retailing and Consumer Services* 22, pp. 187–194.
- Damgaard, M. T. and C. Gravert (2018). "The hidden costs of nudging: Experimental evidence from reminders in fundraising". In: *Journal of Public Economics* 157, pp. 15–26.
- Dennis, A. R., J. A. Rennecker, and S. Hansen (2010). "Invisible whispering: Restructuring collaborative decision making with instant messaging". In: *Decision Sciences* 41.4, pp. 845–886.
- Dennis, A. R., L. Yuan, X. Feng, E. Webb, and C. J. Hsieh (2020). "Digital Nudging: Numeric and Semantic Priming in E-Commerce". In: *Journal of Management Information Systems* 37.1, pp. 39–65.
- Drew, T., M. L.-H. Võ, and J. M. Wolfe (2013). "The invisible gorilla strikes again: Sustained inattention blindness in expert observers". In: *Psychological Science* 24.9, pp. 1848–1853.
- Duke, É. and C. Montag (2017). "Smartphone addiction, daily interruptions and self-reported productivity". In: *Addictive behaviors reports* 6, pp. 90–95.
- Dworkin, G. (1988). *The theory and practice of autonomy*. Cambridge University Press.
- Eppler, M. J. and J. Mengis (2004). "The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines". In: *The Information Society* 20.5, pp. 325–344.
- Evans, J. S. B. T. and K. E. Stanovich (2013). "Dual-Process Theories of Higher Cognition: Advancing the Debate". In: *Perspectives on psychological science : a journal of the Association for Psychological Science* 8.3, pp. 223–241.
- Fehr, E. and S. Gächter (2002). "Altruistic Punishment in Humans". In: *Nature* 415.6868, pp. 137–140.
- Floridi, L. (2014). *The fourth revolution: How the infosphere is reshaping human reality*. OUP Oxford.
- Fogg, B. J. (2003). *Persuasive Technology: Using Computers to Change What We Think and Do*. Interactive Technologies. San Francisco: Morgan Kaufmann.

- Garrison, D. and H. Kanuka (2004). "Blended Learning: Uncovering Its Transformative Potential in Higher Education". In: *The Internet and Higher Education* 7.2, pp. 95–105.
- Goel, S., J. Huang, A. R. Dennis, and K. J. Williams (2020). "Is Privacy Really a Paradox? Understanding the Gap between Privacy Intentions and Behavior". In: *AMCIS 2020 Proceedings*. Salt Lake City, US.
- Gray, C. M., Y. Kou, B. Battles, J. Hoggat, and A. L. Toombs (2018). "The Dark (Patterns) Side of UX Design". In: *CHI 2018*. New York, NY: ACM, pp. 1–14.
- Gray, K., A. Waytz, and L. Young (2012). "The moral dyad: A fundamental template unifying moral judgment". In: *Psychological Inquiry* 23.2, pp. 206–215.
- Greenberg, S., S. Boring, J. Vermeulen, and J. Dostal (2014). "Dark patterns in proxemic interactions: A critical perspective". In: *Proceedings of the 2014 ACM SIGCHI Conference on Designing Interactive Systems, June 21 - 25, 2014, Vancouver, BC, Canada*. Ed. by R. Wakkary. New York, NY: ACM, pp. 523–532.
- Gregor, S. and B. Lee-Archer (2016). "The digital nudge in social security administration". In: *International Social Security Review* 69.3-4, pp. 63–83.
- Hagman, W., D. Andersson, D. Västfjäll, and G. Tinghög (2015). "Public Views on Policies Involving Nudges". In: *Review of Philosophy and psychology* 6.3, pp. 439–453.
- Hansen, P. G. and A. M. Jespersen (2013). "Nudge and the Manipulation of Choice". In: *European Journal of Risk Regulation* 4.01, pp. 3–28.
- Hausman, D. M. and B. Welch (2010). "Debate: To Nudge or Not to Nudge". In: *Journal of Political Philosophy* 18.1, pp. 123–136.
- Hollands, G. J., I. Shemilt, T. M. Marteau, S. A. Jebb, M. P. Kelly, R. Nakamura, et al. (2013). "Altering micro-environments to change population health behaviour: towards an evidence base for choice architecture interventions". In: *BMC public health* 13, p. 1218.
- Hummel, D. and A. Maedche (2019). "How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies". In: *Journal of Behavioral and Experimental Economics* 80, pp. 47–58.
- Hummel, D., P. Toreini, and A. Maedche (2018). "Improving Digital Nudging Using Attentive User Interfaces: Theory Development and Experiment Design". In: *Designing the Digital Transformation: DESRIST 2018 Research in Progress Proceedings of the 13th International Conference on Design Science Research in Information Systems and Technology*. Chennai, India. 3 Jun-6 Jun.

- Ismaili M'hamdi, H., M. Hilhorst, E. A. P. Steegers, and I. de Beaufort (2017). "Nudge me, help my baby: on other-regarding nudges". In: *Journal of medical ethics* 43.10, pp. 702–706.
- Johnson, E. J., S. B. Shu, B. G. C. Dellaert, C. Fox, D. G. Goldstein, G. Häubl, et al. (2012). "Beyond Nudges: Tools of a Choice Architecture". In: *Marketing Letters* 23.2, pp. 487–504.
- Jones, B. and P. Nisbet (2011). "Shareholder value versus stakeholder values: CSR and financialization in global food firms". In: *Socio-Economic Review* 9.2, pp. 287–314.
- Kahn, B. E. and B. Wansink (2004). "The Influence of Assortment Structure on Perceived Variety and Consumption Quantities". In: *Journal of Consumer Research* 30.4, pp. 519–533.
- Kahneman, D. (2012). *Thinking, fast and slow*. London: Penguin Books.
- Kannan, P. K. (2017). "Digital marketing: A framework, review and research agenda". In: *International Journal of Research in Marketing* 34.1, pp. 22–45.
- Kozyreva, A., S. Lewandowsky, and R. Hertwig (in press). "Citizens versus the internet: Confronting digital challenges with cognitive tools". In: *Psychological Science in the Public Interest*.
- Kubina, M., M. Varmus, and I. Kubinova (2015). "Use of big data for competitive advantage of company". In: *Procedia Economics and Finance* 26, pp. 561–565.
- Lacey, C. and C. Caudwell (2019). "Cuteness as a 'Dark Pattern' in Home Robots". In: *14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*.
- Lades, L. K. (2014). "Impulsive consumption and reflexive thought: Nudging ethical consumer behavior". In: *Journal of Economic Psychology* 41, pp. 114–128.
- Lee, J. and Y. Choi (2018). "Expanding affective intelligence theory through social viewing: Focusing on the South Korea's 2017 presidential election". In: *Computers in Human Behavior* 83, pp. 119–128.
- Leidner, D. E., E. Gonzalez, and H. Koch (2018). "An affordance perspective of enterprise social media and organizational socialization". In: *The Journal of Strategic Information Systems* 27.2, pp. 117–138.
- Lembcke, T.-B., N. Engelbrecht, A. B. Brendel, B. Herrenkind, and L. M. Kolbe (2019). "Towards a Unified Understanding of Digital Nudging by Addressing its Analog Roots". In: *Twenty-Third Pacific Asia Conference on Information Systems*, p. 123.
- Lembcke, T.-B., N. Engelbrecht, A. B. Brendel, and L. Kolbe (2019). "To Nudge or not to Nudge: Ethical Considerations of Digital Nudging Based on Its Behavioral

- Economics Roots". In: *In Proceedings of the 27th European Conference on Information Systems (ECIS), Stockholm & Uppsala, Sweden, June 8-14, 2019*.
- Lepenes, R. and M. Małecka (2015). "The Institutional Consequences of Nudging – Nudges, Politics, and the Law". In: *Review of Philosophy and Psychology* 6.3, pp. 427–437.
- Lipton, Z. C. (2018). "The Mythos of Model Interpretability: In machine learning, the concept of interpretability is both important and slippery." In: *Queue* 16.3, pp. 31–57.
- Liu, Z. (2005). "Reading behavior in the digital environment". In: *Journal of Documentation* 61.6, pp. 700–712.
- Loewenstein, G. and T. O'Donoghue (2006). "'We can do this the easy way or the hard way': Negative emotions, self-regulation, and the law". In: *The University of Chicago Law Review* 73.1, pp. 183–206.
- Loh, K. K. and R. Kanai (2016). "How Has the Internet Reshaped Human Cognition?" In: *The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry* 22.5, pp. 506–520.
- Manchiraju, H. and S. Rajgopal (2017). "Does corporate social responsibility (CSR) create shareholder value? Evidence from the Indian Companies Act 2013". In: *Journal of Accounting Research* 55.5, pp. 1257–1300.
- Mangen, A. (2008). "Hypertext fiction reading: haptics and immersion". In: *Journal of research in reading* 31.4, pp. 404–419.
- Marchiori, D. R., M. A. Adriaanse, and D. T. de Ridder (2017). "Unresolved questions in nudging research: Putting the psychology back in nudging". In: *Social and Personality Psychology Compass* 11.1, e12297.
- Martin, W. C., N. Ponder, and J. E. Lueg (2009). "Price fairness perceptions and customer loyalty in a retail context". In: *Journal of Business Research* 62.6, pp. 588–593.
- Meske, C. and I. Amojó (2020). "Ethical Guidelines for the Construction of Digital Nudges". In: *Proceedings of the 53rd Hawaii International Conference on System Sciences*.
- Meske, C. and T. Potthoff (2017). "The DINU-model—a process model for the design of nudges". In: *n Proceedings of the 25th European Conference on Information Systems (ECIS) 25*, pp. 2587–2597.
- Michaelsen, P., L.-O. Johansson, and M. Hedesström (2021). "Experiencing default nudges: autonomy, manipulation, and choice-satisfaction as judged by people themselves". In: *Behavioural Public Policy*, pp. 1–22.

- Michalek, G., G. Meran, R. Schwarze, and Y. Özgür (2016). "Nudging as a New "Soft" Tool in Environmental Policy: An Analysis Based on Insights from Cognitive and Social Psychology". In: *Zeitschrift für Umweltpolitik & Umweltrecht* 8.4, pp. 169–207.
- Michie, S., M. Richardson, M. Johnston, C. Abraham, J. Francis, W. Hardeman, et al. (2013). "The behavior change technique taxonomy (v1) of 93 hierarchically clustered techniques: building an international consensus for the reporting of behavior change interventions". In: *Annals of behavioral medicine : a publication of the Society of Behavioral Medicine* 46.1, pp. 81–95.
- Mills, S. (2022). "Personalized nudging". In: *Behavioural Public Policy* 6.1, pp. 150–159.
- Milosavljevic, M., V. Navalpakkam, C. Koch, and A. Rangel (2012). "Relative visual saliency differences induce sizable bias in consumer choice". In: *Journal of Consumer Psychology* 22.1, pp. 67–74.
- Mirsch, T., C. Lehrer, and R. Jung (2017). "Digital Nudging: Altering User Behavior in DIgital Enviroments". In: *Tagungsband: ‚qTowards Thought Leadership in Digital Transformation‘q*. Ed. by J. M. Leimeister and W. Brenner, pp. 634–648.
- Missbach, B. and J. S. König (2016). "Middle choice preference and snack choice: The role of self-regulatory resources to nudge healthier food choice". In: *Food Quality and Preference* 53, pp. 127–131.
- Münscher, R., M. Vetter, and T. Scheuerle (2016). "A Review and Taxonomy of Choice Architecture Techniques". In: *Journal of Behavioral Decision Making* 29.5, pp. 511–524.
- Nielsen, A. S. E., H. Sand, P. Sørensen, M. Knutsson, P. Martinsson, E. Persson, et al. (2017). *Nudging and pro-environmental behaviour*. Nordisk Ministerråd.
- Nys, T. R. V. and B. Engelen (2017). "Judging Nudging: Answering the Manipulation Objection". In: *Political Studies* 65.1, pp. 199–214.
- Oinas-Kukkonen, H. (2013). "A Foundation for the Study of Behavior Change Support Systems". In: *Personal and ubiquitous computing* 17.6, pp. 1223–1235.
- Oppenheimer, D. M., T. Meyvis, and N. Davidenko (2009). "Instructional manipulation checks: Detecting satisficing to increase statistical power". In: *Journal of Experimental Social Psychology* 45.4, pp. 867–872.
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Puaschunder, J. M. (2017). "Nudging in the digital big data era". In: *European Journal of Economics, Law and Politics* 4.4, pp. 18–23.

- Qiu, X., D. FM Oliveira, A. Sahami Shirazi, A. Flammini, and F. Menczer (2017). "Limited individual attention and online virality of low-quality information". In: *Nature Human Behaviour* 1.7, pp. 1–7.
- Rahwan, I., M. Cebrian, N. Obradovich, J. Bongard, J.-F. Bonnefon, C. Breazeal, et al. (2019). "Machine behaviour". In: *Nature* 568.7753, pp. 477–486.
- Rebonato, R. (2014). "A Critical Assessment of Libertarian Paternalism". In: *Journal of Consumer Policy* 37.3, pp. 357–396.
- Reisch, L. and C. R. Sunstein (2016). "Better than a Whip? European Attitudes Toward Health Nudges". In: *SSRN Electronic Journal*.
- Renaud, K. and V. Zimmermann (2018). "Ethical guidelines for nudging in information security & privacy". In: *International Journal of Human-Computer Studies* 120, pp. 22–35.
- Reutskaja, E., A. Lindner, R. Nagel, R. A. Andersen, and C. F. Camerer (2018). "Choice overload reduces neural signatures of choice set value in dorsal striatum and anterior cingulate cortex". In: *Nature Human Behaviour*.
- Reutskaja, E., R. Nagel, C. F. Camerer, and A. Rangel (2011). "Search Dynamics in Consumer Choice Under Time Pressure: An Eye-Tracking Study". In: *American Economic Review* 101.2, pp. 900–926.
- Rowe, F. (2018). "Being critical is good, but better with philosophy! From digital transformation and values to the future of IS research". In: *European Journal of Information Systems* 27.3, pp. 380–393.
- Sætra, H. S. (2019). "When nudge comes to shove: Liberty and nudging in the era of big data". In: *Technology in Society* 59, p. 101130.
- Schaer, A. and K. Stanoevska-Slabeva (2019). "Application of Digital Nudging in Customer Journeys – A Systematic Literature Review". In: *Twenty-fifth Americas Conference on Information Systems*. Cancun, Mexiko.
- Scheibehenne, B., R. Greifeneder, and P. M. Todd (2010). "Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload". In: *Journal of Consumer Research* 37.3, pp. 409–425.
- Schneider, C., M. Weinmann, and J. Vom Brocke (2018). "Digital Nudging-Influencing Choices by Using Interface Design". In: *Communications of the ACM* January, pp. 1–11.
- Schnellenbach, J. (2012). "Nudges and Norms: On the Political Economy of Soft Paternalism". In: *European Journal of Political Economy* 28.2, pp. 266–277.

- Schwartz, E. (2020). *Six Facts About How Supermarkets Influence What We Buy*. URL: <https://econlife.com/2018/12/six-facts-about-supermarket-shopping/>.
- Shi, S., Y. Wang, X. Chen, and Q. Zhang (2020). "Conceptualization of omnichannel customer experience and its impact on shopping intention: A mixed-method approach". In: *International Journal of Information Management* 50, pp. 325–336.
- Simanowski, R. (2016). *Data love: The seduction and betrayal of digital technologies*. Columbia University Press.
- Simon, H. A. (1969). "The sciences of the artificial". In: *Cambridge, MA*.
- Smith, A. (1776). *The Wealth of Nations: An inquiry into the nature and causes of the Wealth of Nations*. Year 2010. Harriman House Limited.
- Smith, L. R., A. Pegoraro, and S. A. Cruikshank (2019). "Tweet, retweet, favorite: The impact of Twitter use on enjoyment and sports viewing". In: *Journal of Broadcasting & Electronic Media* 63.1, pp. 94–110.
- Stanovich, K. E. and R. F. West (2000). "Individual differences in reasoning: Implications for the rationality debate?" In: *Behavioral and Brain Sciences* 23.5, pp. 645–665.
- Sunstein, C. R. (2013). "Impersonal default rules vs. active choices vs. personalized default rules: A triptych". In: *Active Choices vs. Personalized Default Rules: A Triptych (May 19, 2013)*.
- (2011). "Empirically Informed Regulation". In: *University of Chicago Law Review* 78.4, pp. 1349–1430.
- (2014). "Nudging: A Very Short Guide". In: *Journal of Consumer Policy* 37.4, pp. 583–588.
- (2015a). "The Ethics of Nudging". In: *Yale Journal on Regulation* 32.
- (2015b). "The ethics of nudging". In: *Yale J. on Reg.* 32, p. 413.
- Sunstein, C. R., L. A. Reisch, and J. Rauber (2018). "A Worldwide Consensus on Nudging? Not Quite, but Almost". In: *Regulation & Governance* 12.1, pp. 3–22.
- Taken Smith, K. (2012). "Longitudinal study of digital marketing strategies targeting Millennials". In: *Journal of Consumer Marketing* 29.2, pp. 86–92.
- Thaler, R. H. and C. R. Sunstein (2003). "Libertarian Paternalism". In: *American Economic Review* 93.2, pp. 175–179.
- (2008). *Nudge: Improving Decisions About Health, Wealth and Happiness*. Illustrated ed. New Haven: Yale University Press.
- Tversky, A. and D. Kahneman (1974). "Judgment under Uncertainty: Heuristics and Biases". In: *Science* 185.4157, pp. 1124–1131.

- Uncapher, M. R. and A. D. Wagner (2018). "Minds and brains of media multitaskers: Current findings and future directions". In: *Proceedings of the National Academy of Sciences* 115.40, pp. 9889–9896.
- Valenzuela, A. and P. Raghurir (2009). "Position-based beliefs: The center-stage effect". In: *Journal of Consumer Psychology* 19.2, pp. 185–196.
- van Herpen, E. and A. Bosmans (2018). "Arranging the assortment to arouse choice: Effects of goal-relevant assortment organization on food choice and variety perceptions". In: *Food Quality and Preference* 64, pp. 192–204.
- Veetil, V. P. (2011). "Libertarian Paternalism is an Oxymoron: An Essay in Defence of Liberty". In: *European Journal of Law and Economics* 31.3, pp. 321–334.
- Venema, T. A. G., F. M. Kroese, E. de Vet, and D. T. D. de Ridder (2019). "The One that I Want: Strong Personal Preferences render the Center-stage Nudge Redundant". In: *Food Quality and Preference*, p. 103744.
- Ward, A. F., K. Duke, A. Gneezy, and M. W. Bos (2017). "Brain drain: The mere presence of one's own smartphone reduces available cognitive capacity". In: *Journal of the Association for Consumer Research* 2.2, pp. 140–154.
- Weinmann, M., C. Schneider, and J. Vom Brocke (2016). "Digital Nudging". In: *Business & Information Systems Engineering* 58.6, pp. 433–436.
- Wilkinson, T. M. (2013). "Nudging and Manipulation". In: *Political Studies* 61.2, pp. 341–355.

Appendix

A.1 Importance of Ethical Considerations on (Digital) Nudging

The importance of ethical considerations in IS research has been generally acknowledged, e.g., by March and Niederman (2012): “[...] we encourage research that contributes to making a better world, focuses on ethical and critical goals, increases the diversity of methodological approaches and welcomes other disciplines to contribute to our understanding of IS core issues” (p. 98). As highlighted in Section 3, ethical considerations seem particularly pertinent to ethical (digital) nudging. Accordingly, IS scholars have noted the importance and provided initial research mainly focusing on digital nudging (see Table 2 below).

Suggesting Further Research on Ethical Considerations of Digital Nudging	
Reference	Key Quotations
Weinmann et al. (2016, p. 434)	“A key consideration when making such design decisions is the ethical implications of using nudges. [...] While these unethical nudges may lead to short-term gains for the company, they may have long-term repercussions in terms of loss of goodwill, negative publicity, or even legal action. Therefore, designers must be aware of the ethical implications of nudges [...]”
Gregor and Lee-Archer (2016, p. 67)	“When applying a digital nudge, additional consideration of public concerns regarding ethics and privacy are crucial, as implementation occurs in a dynamic manner at an individual level rather than a cohort level.”
Meske and Potthoff (2017, p. 8)	“In this paper, we have not yet discussed ethical and moral concerns about digital nudging in detail, which is of high importance but should be studied separately. For now, we refer to the discussion paper of Hansen and Jespersen (2013) and call for future research about the ethical implications of digital nudging.”
Meske and Amojó (2020a, p. 415)	“We believe that contributing more ethic-related research on digital nudging, especially research that uses and evaluates ethical guidelines for designing digital nudges, will further help digital nudging research to strongly demarcate aspects in the concept.”
Investigating and Addressing Ethical Considerations	
Reference	Findings
Meske and Amojó (2020b)	Provision of initial ethical guidelines for the construction of digital nudges, highlighting the overall importance of ethical considerations in digital nudging. The paper, however, focuses primarily on the aspect of transparency, thus reflecting on freedom of choice and goal-oriented justification rather briefly.

Table 2 — Summary of Pivotal Literature Suggesting or Addressing Ethical Considerations of Digital Nudging.

A.2 Exemplary Blended Environment Framework Instantiation for (Digital) Nudging

The Blended Environment Framework is instantiated with a prototypical everyday situation of using a fitness app (see 3). For instance, a user has just come home from work and is sitting in her living room reading a book (an analogue situation, i.e., 3, node A0). Now the fitness app could send its user a notification, which motivates her to go on a run (3, node D1). Here, the primary impulse for action occurs in a digital environment (interacting with her smartphone and reading the notification). However, the actual act remains a physical act of exercise (3, node A2), taking place in an analogue environment. In addition, the app may provide our user with feedback on her performance both during and after the workout (3, node D2). Furthermore, the app may also encourage her to set training goals or remind her of further training sessions, and so on. In sum, this results in a “user journey” through several analogue and digital contact points, leading to strong interdependencies between both environments: A digital application aims to induce analogue behaviour and at the same time tries to bind users to the app (Lembcke, Engelbrecht, Brendel, Herrenkind, et al. 2019; Schaer and Stanoevska-Slabeva 2019).




Figure 3 — Blended Environment Framework Instantiation for a Mobile Fitness App. Source: Adapted from Lembcke, Engelbrecht, Brendel, Herrenkind, et al. (2019).

A.3 Middle-, Top-Left- and Visual-Saliency Biases as Exemplary Display Biases

Individuals generally tend to be attracted to things displayed in the middle of a screen, thereby ignoring information on the margins (Drew et al. 2013). This middle bias does not only impact users' attention to specific information or options but also affects their choices. For instance, when presented with the choice between horizontally arranged items, users tend to select the middle option in the array (Atalay et al. 2012; Valenzuela and Raghurir 2009). This effect has been demonstrated for a wide range of choices (Missbach and König 2016). However, according to Venema et al. (2019), this effect is not strong enough to override particularly strong preferences for another option. Moreover, Reutskaja et al. (2018), illustrate that users tend to concentrate their attention on alternative hot spots in situations where no clear middle option exists. In a 2x2 matrix, this is usually the upper left corner (top-left bias; see Figure 4).




Figure 4 — Hotspots for three different arrangements. Measured as a percentage of participants' first fixations on a field. Source: adapted from Reutskaja et al. (2011).

In addition to the arrangement of information on a screen, an option's visual saliency can also affect users' decisions (Clement et al. 2015; Milosavljevic et al. 2012). For instance, Milosavljevic et al. (2012) show that when forcing rapid decision-making, visual saliency determinates consumers' product choices. In particular, consumers are significantly more likely to buy a product that is displayed brighter or longer than other alternatives. The effect of these visual saliency biases increases under cognitive load like multitasking and is particularly strong when individuals do not have strong preferences for or against any available option. Although these three

phenomena also occur in the “analogue world”, they are particularly pronounced in digital environments due to a higher visual impact of screens (Reutskaja et al. 2011).

A.4 Overview of Dark Patterns

With IS becoming more prevalent in our everyday lives, so does the impact of IS on societies (Albrechtslund 2007; Benartzi et al. 2017). Accordingly, Human-Computer-Interaction (HCI) and UX research have become interested in critical aspects that may accompany IS design and implementation (Gray et al. 2018). Besides clear social and economic merits, IS can also be used for deceptive and ethically unsound endeavours. Besides unintended negative consequences that happen by chance, modern IS can also include deliberate deception measures. One way to conceptualise these has been the introduction of so-called “dark patterns”, i.e., “instances where designers use their knowledge of human behaviour (e.g., psychology) and the desires of end-users to implement deceptive functionality that is not in the user’s best interest” (Gray et al. 2018). From a technical perspective, such dark patterns can closely resemble (unethical) nudge interventions as they may utilise the same principles and mechanisms as nudges do.

Although there is a broad consensus about the dangerousness of such manipulative design patterns, this topic has received limited attention besides general calls for devoting more attention (see, e.g., Greenberg et al. 2014, for a rare exception). Based on their practical experience as UX designers, Brignull and Darlington (2020) has begun to collect prominent examples for different types of dark patterns on their website (see Table 3). This collection served as a starting point for a systematisation of dark pattern strategies in which Gray et al. (2018) derive five different strands reflecting potential design motivations:

1. *Nagging*: Redirection of expected functionality that persists beyond one or more interaction(s)
2. *Obstruction*: making a process more difficult or complicated than it needs to be, with the intent of dissuading certain action(s).
3. *Sneaking*: attempting to hide, disguise or delay the divulging of information that is relevant to the user.

4. *Interface interference*: manipulating the user interface that privileges specific actions over others.
5. *Forced action*: requiring the user to perform a particular action to access (or continue to access) certain functionality.

Patterns do not necessarily need to be very sophisticated or directly harmful to be considered as dark (Gray et al. 2018): For instance, a social media app like Instagram asking for the allowance to notify a user may already be seen as dark if it only presents an “OK” and “Not now” option to choose from, but no “Never” or “Do not ask again” alternative (as with “not now”, the app legitimises itself to ask for this consent, again and again, all at the choice architect’s discretion). It shall be noted that, recently, scholars have begun to argue that the discourse on dark patterns shall not only cover traditional screen-based digital interactions but rather include smart devices as well—e.g., voice-based devices like Alexa and Siri or domestic robots like Pepper (Lacey and Caudwell 2019). With these seemingly “cute” devices, users may be even more prone to form closer relationships, opening the door for even more subtle manipulations (Aragón et al. 2015).

Type of Dark Pattern	Description
Bait and Switch	You set out to do one thing, but a different, undesirable thing happens instead.
Confirm shaming	This is the act of guiltting the user into opting into something. The option to decline is worded in such a way as to shame the user into compliance.
Disguised Ads	Advertisements are disguised as other kinds of content or navigation to get you to click on them.
Forced Continuity	When your free trial with a service comes to an end, and your credit card silently starts getting charged without any warning. In some cases, this is made even worse by making it difficult to cancel the membership.
Friend Spam	The product asks for your email or social media permissions under the pretence it will be used for a desirable outcome (e.g., finding friends), but then spams all your contacts in a message that claims to be from you.
Hidden Costs	You get to the last step of the checkout process, only to discover some unexpected charges have appeared, e.g., delivery charges, tax.
Misdirection	The design purposefully focuses your attention on one thing to distract your attention from another.
Price Comparison Prevention	The retailer makes it hard for you to compare the price of an item with another item, so you cannot make an informed decision.
Privacy Zuckering	You are tricked into publicly sharing more information about yourself than you really intended to and named after Facebook CEO Mark Zuckerberg.
Roach Motel	You get into a situation very easily, but you find it hard to get out of it (e.g., a premium subscription).
Sneak into Basket	You attempt to purchase something, but somewhere in the purchasing journey, the site sneaks an additional item into your basket, often through the use of an opt-out radio button or checkbox on a previous page.
Trick questions	While filling in a form, you respond to a question that tricks you into giving an answer you did not intend. When glanced upon quickly, the question appears to ask one thing, but it asks another thing entirely when read carefully.

Table 3 — Exemplary Types of Dark Patterns. Source: Quoted and Adapted from Brignull and Darlington (2020).

A.5 Overview on Exemplary Definitions of (Digital) Nudging

Table 4 depicts pivotal definitions of nudging in behavioural economics, while Table 4 highlights selected definitions for digital nudging stemming from IS research. It should be noted that only IS-related definitions for (digital) nudging have been considered that consider the concept of nudging as it is rooted in behavioural economics.

Reference	Exemplary Nudging Definition (all quotations)
Thaler and Sunstein (2009, p. 6)	A nudge [...] is any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates.
Thaler and Sunstein (2009, p. 5)	[A nudge, the authors] tries to influence choices in a way that will make choosers better off, as judged by themselves.
Hausman and Welch (2010, p. 126)	Nudges are ways of influencing choice without limiting the choice set or making alternatives appreciably more costly in terms of time, trouble, social sanctions, and so forth.
Barton and Grüne-Yanoff (2015, p. 343)	A nudge is defined here as an intervention on the choice architecture that is predictably behaviour-steering, but preserves the choice set and is (at least) substantially non-controlling, and does not significantly change the economic incentives.
Barton and Grüne-Yanoff (2015, p. 347)	[...] to be avoidable, nudges should not only be type-transparent (the general existence of such nudges is made transparent to the nudgee), but also token-transparent (each specific intervention is made transparent to the nudgee).
Sunstein (2015a, p. 511)	Nudges are interventions that steer people in particular directions but that also allow them to go their own way. [...] To qualify as a nudge, an intervention must not impose significant material incentives.
Sunstein (2015b, p. 416)	A nudge might preserve freedom of choice, but it might manipulate people and be objectionable for that reason. [...] transparency and accountability are indispensable safeguards, and both nudges and choice architecture should be transparent.
Hansen (2016, p. 174)	<p>A nudge is a function of (I) [sic] any attempt at influencing people's judgment, choice or behaviour in a predictable way, that is</p> <ol style="list-style-type: none"> 1. made possible because of cognitive boundaries, biases, routines, and habits in individual and social decision-making posing barriers for people to perform rationally in their own self-declared interests, and which 2. works by making use of those boundaries, biases, routines, and habits as integral parts of such attempts. <p>Thus a nudge amongst other things works independently of:</p> <ol style="list-style-type: none"> i. forbidding or adding any rationally relevant choice options, ii. changing incentives, whether regarded in terms of time, trouble, social sanctions, economic and so forth, or iii. the provision of factual information and rational argumentation.

Table 4 — Exemplary Definitions of Nudging in BE

Reference	Exemplary Digital Nudging Definition (all quotations)
Weinmann et al. (2016, p. 433)	Digital nudging is the use of user-interface design elements to guide people’s behavior in digital choice environments. Digital choice environments are user interfaces – such as web-based forms and ERP screens – that require people to make judgments or decisions.
Meske and Potthoff (2017, p. 2589)	Digital nudging is a subtle form of using design, information, and interaction elements to guide user behavior in digital environments, without restricting the individual’s freedom of choice.
Lembcke, Engelbrecht, Brendel, Herrenkind, and Kolbe (2019, p. 10)	<p>A digital nudge is any intended and goal-oriented intervention element (e.g. design, information or interaction elements) in digital or blended environments attempting to influence people’s judgment, choice, or behavior in a predictable way, that</p> <ol style="list-style-type: none"> 1. is made possible because of cognitive boundaries, biases, routines, and habits in individual and social decision-making, 2. works by making use of those cognitive boundaries, biases, routines, and habits as integral parts of such attempts, 3. preserves the full freedom of choice without forbidding or adding any rationally relevant choice options, 4. does not limit the choice set or making alternatives appreciably costlier in terms of time, trouble, social sanctions, and so forth, 5. nudgees must be able to easily recognize when and where they are subject to being nudged (type-transparency), as well as what the nudger’s goals of this intervention are, in addition to how and why the nudge is working (token-transparency) and 6. increases the private welfare of the nudged individual (pro-self) or the social welfare in general (pro social).

Table 5 — Exemplary Definitions of Nudging in IS

A.6 Overview of Choice Architecture Categories and Techniques

This appendix provides an overview of choice architecture techniques: Firstly, we present a taxonomy of choice architecture techniques with examples (see Table 6). In addition, there have been other approaches to structure the broad field of choice architecture techniques (Caraban et al. 2019). Secondly, we outline some frequently used analogue nudges and provide examples for potential digital counterparts (see Table 7).

Category	Technique	Examples
Decision information	Translate information Includes: reframe, simplify	Reframing call for blood donations as death-preventing rather than life-saving
	Make information visible Includes: make own behaviour visible (feedback), make external information visible	Feedback about one's behaviour (e.g., digital wellbeing), information in the form of graphics rather than text
	Provide a social reference point Includes: refer to the descriptive norm, refer to opinion leader	Information about people's behaviour from one's own peer group or people who are valued for special purposes, experts or role models
Decision structure	Change choice defaults Includes: set no-action as default, use prompted choice	Pre-selected options that leave the freedom to select a different option (or not), such as seen on cookie consent banners on websites
	Change option-related effort Includes: increase/decrease physical/financial effort	(Re-)Arrangement of food items in e-grocery stores (e.g., making more expensive food items easier to reach), present a no button significantly smaller than a yes button
	Change range or composition of options Includes: change categories, change the grouping of options	Combine different product options (e.g., the main product and accessories)
	Change option consequences Includes: connect decision to benefit/cost, change social consequences of the decision	Information access requires additional action (e.g., taking place in a survey), restricting articles through a paywall
Decision assistance	Provide reminders Includes: notifications	Push notifications on a smartphone (e.g., of WhatsApp, Facebook, or Instagram)
	Facilitate commitment Includes: support self-commitment/public commitment	Shaping a community to combine efforts (e.g., stickK.com), blocking Internet access (e.g., through focus apps after a specific period)

Table 6 — Choice Architecture Categories and Techniques. Source: Adapted from Münscher et al. (2016).

Technique	Analogue-Focused Examples	Digitally-Focused Examples
Default rules	Automatic enrolment in programs, including education, health, savings	Pre-set default options (e.g., preselect a marketing consent in cookie consent banners)
Simplification	In part to promote the take-up of existing programs (e.g., ease legislation or sign-up procedures for public health programs)	Make mobile devices and apps easier to use (e.g., lower app complexity, provide standard and expert interfaces, design interfaces as self-explanatory as possible)
Use of social norms	Emphasising what most people do, e.g., “most people plan to vote” or “most people pay their taxes on time” or “nine out of ten hotel guests reuse their towels”	Present “customers who browsed this article also browsed” or “customers who bought this article also bought” items in an e-commerce shop (e.g., Amazon)
Increases in ease and convenience	Making low-cost options or healthy foods visible	Offer one-click ordering (e.g., Amazon)
Disclosure	The economic or environmental costs associated with energy use or the total cost of certain credit cards	Releasing large amounts of data, as in the case of data.gov and the Open Government Partnership
Warnings, graphic or otherwise	As for cigarettes, private or public warnings in large fonts, bold letters and bright colours	Notifications on daily smartphone usage time, show concrete steps that can be taken to reduce a risk
Precommitment strategies	By which people commit to a particular course of action, e.g., sign up for a regular sport exercise course	Facilitate frequent digital engagement at precise future moments in time, e.g., a regular online language class or a regular time to train in an e-gym together with friends
Reminders	By telephone or mail to be reminded of important events or appointments.	By push-notification, e-mail or text message for overdue bills and coming obligations or appointments
Eliciting implementation intentions	Asking people implementation-related questions (e.g., “do you plan to vote” or “do you plan to vaccinate your child”) emphasised with individuals’ identity (e.g., “you are a voter, as your past practices suggest”)	Accompanying user implementation intentions (e.g., through digital Pomodoro timer apps), suggesting certain options based on previous usage patterns
Informing people of the nature and consequences of their own past choices	Smart disclosure (e.g., informing users about previous health care or electricity expenditures), comparing spending in a period with the previous spending in a similar period	Comparing previous with current behaviour (e.g., “yesterday you picked up your phone 30 times less than today and had 2 hours more for individual leisure”)

Table 7 — Ten Frequent Nudges. Source: Adapted from Sunstein (2014).

A.7 Overview of AI-Assisted Digital Choice Architectures

Table 8) highlights key choice architectures that are increasingly fuelled by AI and may be adapted to influence or even disinform, misguide or manipulate users, such as some kind of digital nudge intervention. However, as outlined in Sections 3.3.1 and 3.3.2, such AI-assisted choice architectures have a higher risk of becoming ethically unsound, both by design and accidentally.

Category / Technique	Description / Examples
Algorithmic Curation and Personalisation	
Recommender systems	Recommending media items and products (e.g., Netflix, YouTube, Spotify, Amazon) Suggestions to follow friends or accounts (e.g., Instagram, Facebook, Twitter, LinkedIn) Potential match recommendations in online dating (e.g., <i>Tinder, OkCupid</i>)
Search algorithms and ranking systems	Presentation of search results (e.g., <i>Google, Bing, Baidu</i>) Predictive search suggestions/auto-complete functions (e.g., <i>Google's personalised search suggestions</i>) Newsfeed and timeline customisation (e.g., <i>Facebook, Instagram, Twitter</i>)
Advertising algorithms	Customising and targeting advertising to specific audiences (e.g., <i>Facebook, Twitter</i>) Algorithmic pricing (e.g., competitive pricing suggestions in e-commerce like Amazon Marketplace)
Bots and Smart Assistants	
Virtual assistants	Software agents reacting or interacting based on commands by human users (e.g., <i>Siri, Amazon Alexa, Google Assistant</i>)
Social media bots	Software agents designed to behave like human users (e.g., fake news bots that are used to comment or share posts on social media)
Chatbots	Software agents designed to converse with human users (e.g., <i>in customer support</i>)
Algorithmic Tools	
Translation and speech recognition	Machine-learning-based translation services (e.g., <i>Google Translate, DeepL</i>)
Filtering algorithms	Automated e-mail filters for separating e-mails into categories (e.g., <i>spam, promotions, social, primary, important</i>)
Maps and navigation	Providing optimised or customised directions and orientation on maps (e.g., <i>Google Maps, Waze, Here</i>)
Content moderation	Detection and automated removal of harmful content (e.g., fake accounts, hate speech, offensive graphic content, disinformation)

Table 8 — Exemplary AI-Assisted Digital Choice Architectures Applicable to Digital Nudging. Source: Excerpted from Kozyreva et al. (in press).

Declaration of contribution to each essays of the cumulative dissertation (based on the CRediT taxonomy by Brand et al. 2015)¹⁰

To the five essays of the cumulative dissertation I contributed as follows:

(1) The Interplay of Different Excuses in Determining Moral Wiggle Room

co-authored with Mattheus Brenig and Daniel Hermann:

Conceptualization, Investigation, Methodology, Project administration, Validation, Visualization, Writing—review editing.

(2) Consumer Behavior under Benevolent Price Discrimination

co-authored with Alexander Erlei and Mattheus Brenig:

Conceptualization, Investigation, Methodology, Writing—review editing.

(3) The Virtual Online Supermarket

co-authored with Tim-Benjamin Lembcke, Alfred Benedikt Brendel, Kilian Bizer, and Lutz M. Kolbe:

Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Validation, Visualization, Writing—original draft.

(4) Behavioral Design In Online Supermarkets

co-authored with Tim-Benjamin Lembcke, Mathias Willnat, and Sascha Lichtenberg:

Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Validation, Visualization, Writing—original draft.

¹⁰Brand et al. (2015). “Beyond authorship: attribution, contribution, collaboration, and credit”. In: *Learned Publishing* 28.2, pp. 151–155

(5) Ethical Dimensions in Digital Nudging:

co-authored with Tim-Benjamin Lembcke, Alfred B. Brendel, and Alan R. Dennis:

Conceptualization, Investigation, Methodology, Project administration, Visualization, Writing—original draft, Writing—review editing.

Date, Signature

Ph.D. program in Economics Declaration for admission to the doctoral examination

I confirm

1. that the dissertation that I submitted “Determinants of Other-Regarding Behavior and the Private Provision of Public Goods” was produced independently without assistance from external parties, and not contrary to high scientific standards and integrity,
2. that I have adhered to the examination regulations, including upholding a high degree of scientific integrity, which includes the strict and proper use of citations so that the inclusion of other ideas in the dissertation are clearly distinguished,
3. that in the process of completing this doctoral thesis, no intermediaries were compensated to assist me neither with the admissions or preparation processes, and in this process, no remuneration or equivalent compensation were provided and no services were engaged that may contradict the purpose of producing a doctoral thesis,
4. that I have not submitted this dissertation or parts of this dissertation elsewhere.

I am aware that false claims (and the discovery of those false claims now, and in the future) with regards to the declaration for admission to the doctoral examination can lead to the invalidation or revoking of the doctoral degree.

Date, Signature