Essays in Behavioral Public Economics

INAUGURAL-DISSERTATION

zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaft
doctor rerum politicarum (Dr. rer. pol.)

am Fachbereich Wirtschaftswissenschaft
der Freien Universität Berlin

vorgelegt von
Renke Schmacker, MSc

Berlin, Mai 2020

Gedruckt mit der Genehmigung des Fachbereichs Wirtschaftswissenschaft der Freien Universität Berlin.

Dekan:
Prof. Dr. Dieter Nautz

Erstgutachter:
Prof. Dr. Peter Haan, Freie Universität Berlin

Zweitgutachterin:
Prof. Dr. Dorothea Kübler, Technische Universität Berlin

Datum der Disputation:
26. Juni 2020

Erklärung über Zusammenarbeit mit KoautorInnen und Vorveröffentlichungen

Kapitel 1

- In Zusammenarbeit mit Jana Friedrichsen und Tobias König.
- Das in diesem Kapitel berichtete Experiment basiert auf einem Pilotexperiment, welches ich im Rahmen meiner Abschlussarbeit für den Master of Science in Economics an der Humboldt-Universität zu Berlin durchgeführt habe (Schmacker, 2015). Das Experiment wurde daraufhin substanziell erweitert und es wurden vollständig neue Daten erhoben.
- Eine frühere Version des Artikels wurde unter dem Titel "Welfare Stigma in the Lab: Evidence of Social Signaling" (2017, CESifo Working Paper No. 6519) zirkuliert.
- Veröffentlicht als Friedrichsen, J., König, T., and Schmacker, R. (2018), "Social Image Concerns and Welfare Take-Up", DIW Discussion Papers 1752, 2018.
- Veröffentlicht als Friedrichsen, J., König, T., and Schmacker, R. (2018), "Social Image Concerns and Welfare Take-Up", *Journal of Public Economics*, 168:174-192.

Kapitel 2

- In Zusammenarbeit mit Sinne Smed.
- Bislang unveröffentlicht.

Kapitel 3

- Bislang unveröffentlicht.

Kapitel 4

- In Zusammenarbeit mit Alessandro Castagnetti.
- Bislang unveröffentlicht.

Acknowledgments

I would like to express my deepest gratitude to all the people who helped me with this dissertation. I thank my first supervisor Peter Haan for his outstanding support throughout my PhD. He provided not only invaluable advice and guidance regarding my research, but has contributed to my professional and personal development in general with his patience and kindness. I would like to extend my gratitude to my second supervisor, Dorothea Kübler, who has inspired my interest in research when I worked in her department at WZB during my studies. I count myself lucky to have such caring and dedicated mentors. Next, I would like to show my greatest appreciation to Tobias König and Jana Friedrichsen, with whom I co-authored my first paper and whose advice has been tremendously helpful ever since. I would like to offer my special thanks to Sinne Smed for the great collaboration on our papers as well as her hospitality in Copenhagen. I also thank my co-author Alessandro Castagnetti for the excellent cooperation and his friendship.

I gratefully acknowledge financial support from the Leibniz Association and the German Research Foundation. Furthermore, I am grateful to Andy Schotter and Guillaume Frechette for hosting me during my research stay at NYU. Both of them dedicated a lot of their time to discussing parts of this dissertation. I would also like to thank Bertil Tungodden for welcoming me at NHH Bergen and Niels Johannesen for hosting me at CEBI in Copenhagen.

This dissertation was written while I was member of the Graduate Center and the Public Economics Department at DIW Berlin. At the Graduate Center, I would like to offer my special thanks to Georg Weizsäcker for his support and encouragement of my aspirations. Moreover, I would like to thank Helmut Lütkepohl and the entire GC team. I am indebted to Adam Lederer for proof-reading most of this dissertation. At the Public Economics department, I would like to thank all my colleagues for the friendly work environment. I am deeply grateful that I could start my PhD with an amazing cohort of people who have become dear friends. In particular, I thank Patrick Burauel, Sascha Drahs, Iuliia Grabova, Daniel Graeber, Melanie Koch, Max Schäfer, Julia Schmieder, Kevin Tran, and my office mate Aline Zucco.

Further, I thank my friends and my partner for their tremendous support over the years. Ultimately, I owe my deepest gratitude to my parents and my family for their enduring support and encouragement.

Table of Contents

List of Tables

List of Figures

Preface

When economists assess public policy, they traditionally employ the standard model of perfectly informed, fully rational, and self-interested decision-makers. Under this assumption, individuals always choose what is best for them and the scope for government intervention is mostly limited to correcting market failures. However, the assumption that individuals always act in their best interest is at odds with many empirical observations. For example, while both obesity rates and average calorie intake seem to be ever increasing, we observe at the same time the rise of a multi-billion-dollar weight loss industry (Ruhm, 2012). Many individuals show a high willingness to pay for bariatric surgery or Weight Watchers memberships in order to commit to eating less in the future. This suggests that observed consumption behavior does not necessarily comply with individual's own self-interest. In fact, many policy interventions are purposefully designed to protect individuals from making sub-optimal decisions. These include mandating warning labels on high-caloric food, setting deliberate defaults for saving plans, and establishing consumer protection agencies.

Since the emergence of behavioral economics, most economists acknowledge that individuals depart in various ways from the standard neoclassical model. Numerous empirical studies from the lab and the field identify systematic and economically relevant deviations from the benchmark model (see e.g. DellaVigna, 2009). There is ample evidence for non-standard *preferences* and non-standard *beliefs*. For example, individual's preferences are not just egocentric but many people care about what others think of them. Similarly, beliefs – i.e. the probabilities individuals assign to the possible states of the world – are not always always well-aligned with reality, as exemplified by many people's tendency to overestimate their own ability. A growing theoretical literature aims to incorporate more realistic assumptions about human behavior in a formal economic framework. For instance, people's self-control problems in food consumption can be modeled in the form of time-inconsistent discounting (Laibson, 1997), i.e. individuals put disproportionate value on the present. Importantly, changing the assumptions about what people want and believe has implications for optimal policy.

The question of how deviations from the neoclassical model affect the economic analysis of public policy is the subject of *behavioral public economics* (Bernheim and Taubinsky, 2018; Chetty, 2015). On the one hand, deviations from the standard model can necessitate policy interventions where the standard model would have prescribed to not interfere with individual choices. To give an example, if there is reason to believe that individuals behave time-inconsistently but would actually like to behave time-consistently, individuals can be made better off by giving them incentives to change their behavior (Bernheim and Taubinsky, 2018). On the other hand, behavioral economics can inform us how policies have to be implemented to reach their goals more effectively. While the standard model does not take contextual factors or framing effects into account, considering these behavioral factors can extend the policymaker's toolkit without having to restrict individuals' choices (e.g. Thaler and Sunstein, 2008). This dissertation analyzes different types of non-standard preferences and beliefs along with their implications for policy design. First, I consider social image concerns, i.e. the desire to be seen in a socially favorable way. I assess how such preferences can affect the take-up of a redistributive transfer. Second, I study self-control problems, i.e. the tendency to put excessive weight on immediate gratification that can, among others, lead to overconsumption of sugar. I analyze whether self-control problems can be corrected by taxing "sin goods" like sugar and fats. Third, I investigate motivated beliefs, i.e. the desire to hold certain beliefs because they provide utility to the individual. While miscalibrated beliefs are a major problem in many domains, I focus on how individuals can keep overconfident beliefs about their ability despite permanent feedback.

The first chapter investigates whether social image concerns lead potential transfer claimants to leave money on the table. In many welfare programs, eligible individuals do not take up the benefits they are entitled to. High non take-up can be problematic if it hinders policymakers from reaching political goals, such as redistribution of resources and amelioration of poverty. In the standard economic model, high non take-up can only be rationalized by lack of information or excessive transaction costs. However, the sociological literature about stigmatization of welfare claimants suggests that also contextual factors can influence welfare take-up (e.g. Goffman, 1963; Stuber and Schlesinger, 2006). When individuals have social image concerns, they may refrain from taking up transfers because they do not want to send negative signals about their ability or work motivation. While many papers discuss this idea (Moffitt, 1983; Besley and Coate, 1992), the literature does not provide causal evidence for welfare stigma.

In the first chapter, which is joint work with Jana Friedrichsen and Tobias König, we use a theory-based laboratory experiment to show that social image concerns can affect an individual's decision to claim a transfer. In the experiment, subjects have to decide whether to take up a redistributive transfer from other participants in the experiment. Eligibility to receive a transfer is based on performance in a general knowledge quiz: subjects who rank last within their group are entitled to receive the transfer. As a first treatment variation, we vary if the take-up decision is publicly observable or private. The results show that subjects are 30 percentage points less likely to take up the transfer if their take-up decision can be seen by others. We call this difference the "stigma effect."

In order to establish that social image concerns are responsible for the stigma effect, we vary the informational content of the take-up decision in further treatments. First, we make ranks and, hence, transfer eligibility random. Second, we let the transfer be paid by the experimenter and not by other subjects. Thereby, we can distinguish if the take-up decision conveys information about the claimant's poor performance in the knowledge quiz (ability signaling) or her willingness to live off others (free-rider signaling). We find that both ability and free-rider signaling have an independent and significant effect on take-up. These results support the predictions from our theoretical model of social image concerns and welfare take-up.

This chapter contributes to the literature that studies the reasons for incomplete program participation. The existing empirical literature finds that informational constraints, program complexity, and transaction costs contribute to low welfare take-up (Currie, 2006; Bhargava and Manoli, 2015; Finkelstein and Notowidigdo, 2019). In comparison to these factors, there is limited evidence for welfare stigma but the studies that aim to causally test its effect mostly focus on *self* image concerns (e.g. Bhargava and Manoli, 2015). Our findings suggest that there is a more important role for *social* image concerns in the take-up decision. Acknowledging the stylized setting of our study, we emphasize the need for future field studies in settings where social image concerns are likely to be relevant, e.g. in food stamps or social housing. If social image concerns turn out to be relevant in the field, there are many potential ways for policymakers to take these behavioral factors into account: amongst others, changing the way how transfers are paid out (cash transfers instead of foodstamps) or how application procedures are designed (online instead of in-person) may affect take-up.

The second and third chapters study the question whether sin taxes on unhealthy foods are effective measures against the "global obesity epidemic" (WHO, 2000). Obesity is one of the most important public health challenges of our time as it is a leading risk factor for, amongst others, type 2 diabetes and cardiovascular disease (WHO, 2015). The global increase in obesity rates is to a large degree attributed to the rising total calorie intake from processed food (Finkelstein et al., 2005). To counter this development, many countries, like France and the United Kingdom, are imposing "sin taxes" on sugar sweetened beverages (SSBs).

According to standard economic theory, taxes on sugar are only justified as long as they correct the *externality* that heavy sugar consumers impose on the public health sector (Diamond, 1973). Apart from that, the standard model would advise to not distort consumption since individuals themselves know what is best for them. However, there is a lot of evidence suggesting that individuals do not always behave in their long-term best interest. In particular, when it comes to food consumption, many individuals have self-control problems and find it hard to be on a diet for a sustained period of time. Hence, their behavior imposes an *internality* on their future self, for example in the form of health costs. Based on this observation, an influential theoretical literature argues that sin taxes can serve individuals with low self-control as a commitment device (O'Donoghue and Rabin, 2006; Allcott et al., 2019a). However, for such a tax to be welfare improving, the corrective gains for individuals with low self-control have to outweigh the distortionary costs for unbiased individuals. To ensure that this is the case, individuals with low self-control have to be at least as price responsive compared to those with high self-control (O'Donoghue and Rabin, 2006; Haavio and Kotakorpi, 2011). Whether this condition is fulfilled is an empirical question that we tackle in the next chapter.

The second chapter, which is based on joint work with Sinne Smed, investigates if sin taxes on soft drinks and fats successfully target individuals with low self-control. For identification of differential price responsiveness, we exploit tax hikes and cuts of the soft drink tax as well as the fat tax in Denmark. Our dataset is a unique household panel that comprises purchase records as well as a survey measure on self-control (Tangney et al., 2004). Using this measure, we can split the sample in consumers with high and low self-control, which allows us to estimate their differential response to the sin tax changes.

We find that consumers with low self-control decrease their purchases *less* when soft drink and fat taxes go up, compared to consumers with high self-control. This suggests that the tax does not help present-biased consumers to reduce their purchases. In contrast, when taxes are cut, we observe an increase in purchases of similar magnitude by both groups. Our identification strategy assumes that, without the tax change, there would be no differential change between the groups. The credibility of this assumption is supported by the parallel pre-tax trends. Further, the absence of differential responses by income, education, nutritional knowledge, and preferences for unhealthy foods suggests that the results are indeed driven by differences in self-control. We interpret our results in light of a model of self-control and rational habit formation based on O'Donoghue and Rabin (2002). In such a model, the asymmetric responses to tax hikes and cuts by self-control can be explained by differential patterns of habit formation.

This chapter contributes to the growing literature on using taxation to correct behavioral internalities, such as imperfect self-control (O'Donoghue and Rabin, 2006; Gruber and Köszegi, 2001; Allcott et al., 2019a; Dubois et al., 2019). We estimate a crucial statistic in theoretical models of sin taxes, that is, the correlation between the internality and the price responsiveness. Moreover, we contribute to the theoretical literature on habit formation and self-control. Here, we show that tax changes on habituating goods can lead to different responses by consumers with high and low self-control.

Our results indicate that policies using price incentives do not help individuals with low self-control to reduce their sugar and fat consumption. Instead, to help consumers with low self-control, it may be more effective to incentivize producers to make their products healthier. For example, instead of using volumetric taxes, the tax rate should depend on the sugar content of goods (Allcott et al., 2019b).

The second chapter suggests that habit formation plays a crucial role in shaping the response to soft drink taxes. This notion is supported by physiological research according to which tastes for sugar and sweetness are indeed habit-forming (e.g. Liem and de Graaf, 2004; Ahmed et al., 2013; Mennella et al., 2016). The third chapter builds on this finding and studies the impact of incorporating habit formation into a structural demand model.

In economics, habit formation is typically modeled in the form of intertemporal complementarities in consumption (Becker and Murphy, 1988). That means, consumption in the current period increases the utility from consumption in the next period. Ignoring this kind of positive state dependence can lead to underestimation of long-run price elasticities. The reason is that a price change does not just have a direct effect on consumption but also an indirect effect due to its impact on future consumption. Nonetheless, most demand models that simulate soft drink taxes do not model habit formation. The reason is that it is challenging to distinguish the impact of habit formation from negative state dependence due to stockpiling and unobserved heterogeneity (Wang, 2015).

In the third chapter, I investigate the impact of habit formation on long-run price elasticities of soft drink taxes. Therefore, I build a structural demand model that incorporates both habit formation and stockpiling. To identify these two forms of state dependence, I assume that consumers predominantly stockpile when a product is on sale. I estimate the model using household scanner data from the United States.

First, I provide descriptive evidence that both positive and negative state dependence are relevant. After controlling for proxy variables of stockpiling, there is a strong positive impact of past purchases

on current purchases. Second, I estimate a discrete choice model that includes lagged parameters indicating whether a purchase was made in the previous period and whether the purchased product was on sale. I estimate a nested logit model that links the decision to buy a soft drink with the decision which product to buy. The model includes random coefficients to allow for unobserved heterogeneity in preferences and takes into account that the initial state is not exogenous.
The estimation results show that there is evidence for true state dependence in soft drink purchases. The simulated long-run price elasticities are approximately 20 percent larger than short-run price elasticities. I simulate the effects of different taxes on purchases, such as a an excise tax and an *ad valorem* tax on sugary soft drinks, as well as an excise tax on all soft drinks. While the taxes increase prices by a similar extent on average, their impact on demand is different: The excise tax on sugary soft drinks is more effective in reducing sugar consumption than the *ad valorem* tax and the excise tax on all soft drinks. Due to habit formation, the long-run responses are 16 to 23 percent larger than the shor-run response. Moreover, the response to the tax is relatively uniform across the income distribution and across smaller and larger households.
These findings inform the literature that uses naturally occurring price variations to simulate the impact of taxes from demand models (e.g. Andreyeva et al., 2010; Allcott et al., 2019a; Dubois et al., 2019). I show that assuming time separability in utility can lead to an underestimation of the effectiveness of soft drink taxes. This complements the finding by Wang (2015) who argues that ignoring state dependence in the form of stockpiling leads to overestimation of price elasticities. I argue that both stockpiling *and* habit formation should be considered to get a more complete picture of the response to soft drink taxes.

While the first chapters are concerned with deviations from the standard assumption of self-concerned and time-consistent preferences, the final chapter turns to non-standard beliefs. Beliefs are essential to decision making in all kinds of domains, ranging from voting behavior to career decisions. Standard economic theory assumes that individuals' beliefs are, on average, correct and that agents cannot be systematically fooled. In contrast, the literature on "motivated beliefs" shows that this is not necessarily the case if individuals attach value to holding certain beliefs (see Bénabou and Tirole, 2016). For example, individuals are often shown to be overconfident about their ability (e.g. Benoît et al., 2015). The prevalence of motivated beliefs is puzzling considering that individuals receive continuous feedback from their surroundings. This chapter looks at a potentially important mechanism that has received little attention in the literature.
The fourth chapter, which is joint work with Alessandro Castagnetti, studies whether individuals select information structures in order to protect their motivated beliefs. Therefore, we conduct a lab experiment, in which we vary the ego-relevance of information: in one treatment, information concerns a subject's rank in the intelligence distribution and in another treatment it is based on a random number. Before subjects receive signals about these states of the world, they can choose the information structure from which they receive signals. Information structures vary in the informativeness and the framing of signals. Before and after receiving signals, we elicit subjects' beliefs about their intelligence rank or the random number in an incentive compatible way. Thereby, we can study how subjects update their beliefs in response to the signals.

Our results reveal that subjects choose less informative and positively framed signals when information is ego-relevant. While the framing of feedback should not matter for a Bayesian updater, we show that individuals update their beliefs less when negative feedback about own intelligence is made less explicit. However, this is not the case in the control treatment. This suggests that individuals select information that makes it easier for them to interpret signals in a self-serving way.

This chapter contributes to the literature on information avoidance (Golman et al., 2017) and asymmetric updating (e.g Eil and Rao, 2011; Möbius et al., 2014). We show that asymmetric updating is only possible if the information structure allows for misperceiving the feedback due to a more ambiguous framing. In the experiment, we show that only subjects in the positively framed information structure remain overconfident despite multiple signals. Hence, our findings contribute to a better understanding of the high prevalence of overconfident beliefs.

The findings have applicability to the provision of ego-relevant feedback in educational contexts and at the workplace. If individuals dislike adjusting their beliefs downward, they may make suboptimal career decisions due to differences in feedback culture (Sabot and Wakeman-Linn, 1991; Ahn et al., 2019). These findings open up a wide range of highly policy-relevant research questions.

Chapter 1

Social Image Concerns and Welfare Take-Up

Published as:

Friedrichsen, J., König, T., and Schmacker, R. (2018), "Social Image Concerns and Welfare Take-Up", *Journal of Public Economics*, 168:174-192.

Chapter 2

Sin Taxes and Self-Control

2.1 Introduction

The "global obesity epidemic" is a major public health challenge (WHO, 2000) and one of the leading risk factors for many non-communicable diseases like type 2 diabetes and coronary heart disease (Smith Jr., 2007). Poor diets that contain high levels of sugar and fat are among the main culprits of this phenomenon (Finkelstein et al., 2005). Hence, the World Health Organization (WHO, 2016) advises governments to consider the introduction of so-called "sin taxes" on unhealthy foods, e.g. taxes on sugar sweetened beverages (SSB). A number of countries have already implemented taxes on sugary beverages and other unhealthy foods, e.g. France, Mexico, the United Kingdom, and, until recently, Denmark.

There are two rationales for sin taxes: externalities and internalities. Externalities mean that high sugar consumers do not take the social costs of adverse health behavior into account and the tax is meant to internalize these costs. Internalities in the form of self-control problems imply that people underappreciate the long-term health costs that an unhealthy diet has on themselves. In this paper, we focus on the internality argument since it is very prominent in the public debate about sin taxes on foods.[19] The idea is that a sin tax could help consumers with low self-control to follow their long-run utility by increasing the instantaneous price. Such a tax can even be welfare-improving if the corrective gains for individuals with low self-control overweigh the distortionary costs for those without self-control problems. However, to ensure that this is the case individuals with low self-control have to reduce their purchases at least as much as those with high self-control (O'Donoghue and Rabin, 2006;

[19]First, soft drink taxes are often advocated based on the premise that particularly children, who belong to the heaviest soft drinks consumers, ignore the long-run consequences of high sugar intake (Dubois et al., 2019). Second, the effectiveness of these taxes is usually assessed by the reduction in consumption and not by tax revenue raised (for externality correction this distinction would not be as relevant). For example, on March 13, 2018, the former British finance minister and initiator of the British soft drink tax, George Osborne, tweeted: "In OBR *[Office for Budget Responsibility]* report today is news that our Sugar Tax is even more effective than hoped. Expected receipts halved [...]". (https://twitter.com/George_Osborne/status/973647500551827456, retrieved 09/23/19).

Haavio and Kotakorpi, 2011).

In this paper, we investigate empirically the causal effect of self-control on responsiveness to sin tax changes. For identification, we exploit exogenous variation in two Danish sin taxes: First, we consider the increase of the soft drink tax in 2012 and its complete repeal in 2014. Second, we investigate the fat tax on saturated fat, introduced in 2011 and repealed in 2013. We use a unique panel data set that comprises purchase records of 1,278 households who stay in the panel for the period of tax changes and who have also answered a well established survey on self-control (Tangney et al., 2004). Using the survey, we stratify the sample into high and low self-control consumers. The validity of the measure is supported by the observation that panelists with low self-control have larger Body Mass Index and report both the intention to reduce their weight and to improve their eating habits. In our empirical analysis, we estimate the differential effect of tax changes on consumers with low and high self-control. Finally, we propose a theoretical explanation for our results using a model of rational habit formation (Becker and Murphy, 1988).

In our empirical analysis, we find that consumers with low self-control reduce their purchases significantly less than those with high self-control when the soft drink tax is increased. In contrast, in response to the tax repeal all groups increase their purchases to a similar extent. We find the same pattern for the introduction and repeal of the fat tax. Here, we look at butter since it experienced substantial tax variation due to its high content of saturated fat. To provide credibility to our empirical strategy, we show that the trends in the years before the tax changes do not differ between self-control groups. Moreover, we show that the results are not driven by heterogeneous demand responses by education, income, taste for unhealthy food, nutritional knowledge, and proximity to the German border.

We provide a theoretical explanation for this asymmetry using a model of rational habit formation (Becker and Murphy, 1988; O'Donoghue and Rabin, 2002). Rational habit formation means that consumption today makes it more likely to also consume tomorrow and consumers take that into account. This explanation is supported in our dataset since panelists with low self-control report to be more strongly habituated with respect to sugar and fats. In the model, consumers are aware that a tax does not just change the instantaneous price but also all prices in the future. Since consumers with low self-control discount the future (and, hence, the future price changes) more, they react in general less strongly to tax changes. However, this effect is not symmetric if a tax hike is followed by a tax cut, as in our case. The reason is that there are more high self-control consumers who stop consuming due to the tax hike. Hence, they get used to not consuming and their habit stock decreases. Due to a lower habit stock, they are less likely to resume consumption when the tax is cut in the next period. Therefore, the difference in responsiveness between high and low self-control is smaller for the tax cut compared to the tax hike.

Our study is motivated by the theoretical literature on taxation of behavioral internalities like imperfect self-control. The idea is that a lack of self-control can lead consumers to over-consume goods that have long-run costs, which are not fully taken into account at the moment of consumption. A sin tax increases the instantaneous and future costs of consumption and reduces over-consumption. Gruber and Köszegi (2001) show that optimal taxes on cigarettes are substantially higher if addicted individuals are present-biased. O'Donoghue and Rabin (2006) and Haavio and Kotakorpi (2011) argue that an internality correcting tax can be welfare improving if individuals with low self-control are at least as responsive to a sin tax as those with high self-control. Further, the comprehensive model by Allcott et al. (2019a), in which they study the welfare effects and the distributional implications of sin taxes, takes the correction of internalities into account. However, these papers do not make predictions whether consumers with low self-control actually respond to sin taxes and leave this question to empirical research.

We contribute to the burgeoning empirical literature, which assesses targeting properties of sin taxes, by estimating heterogeneous tax responsiveness by self-control. Allcott et al. (2019a) estimate in their empirical section the share of soda consumption that is due to a self-reported lack of self-control.[20] They find that bias-induced consumption is decreasing in income, which means that poor consumers can benefit more from the corrective effects of the tax. However, due to their focus on the regressivity of sin taxes, they do not consider if the price elasticity varies with the level of self-control. In contrast, we use actual tax variation and investigate if the tax actually targets individuals with low self-control. The targeting properties of a soft drink tax are also investigated by Dubois et al. (2019) in a structural demand model. They estimate price elasticities of different consumer groups and hypothesize that the high soda preference of certain groups (e.g. young people and high sugar consumers) is more likely due to biases. They find that young people are more price responsive, but that high sugar consumers are less price responsive than the average consumer. We complement these findings by employing an established measure of self-control and by exploiting exogenous variation in prices to provide a causal test of the impact of self-control on price responsiveness.

Furthermore, we contribute to the empirical literature that uses quasi-experimental variation in sin taxes to estimate the impact of taxes on purchases. We are first to use tax variation to study heterogenous responses by self-control. There is a longstanding literature that uses tax variation in tobacco and alcohol taxes to estimate price elasticities (see the surveys in Chaloupka et al. (2012) for tobacco and in Wagenaar et al. (2009) for alcohol). With the increasing prevalence of sin taxes on food, there are also more and more evaluations of these kind of policies. Jensen and Smed (2013) analyze the short-run effects of the fat tax in Denmark in a pre-post design and document a significant drop in average purchases of saturated fat from butter and margarine. Cawley et al. (2019b) survey the empirical literature on soft drink taxes and conclude that average purchases decrease after tax introductions. This is documented for US cities like Berkeley and Philadelphia using geographical control groups (e.g. Cawley et al., 2019a; Rojas and Wang, 2017) and for the tax in Mexico using pre-post designs (Colchero et al., 2016, 2017).[21] In earlier work, we analyze the tax pass-through and average purchase response to the increase 2012 and repeal 2014 of the Danish tax on soft drinks using a pre-post design (Schmacker and Smed, 2020). Whereas the focus in almost all of these paper was on the average change in purchases, in this paper we use the exogenous variation in prices to test if different levels of self-control imply different degrees of price responsiveness.

[20] They use the Nielsen household panel and classify panelists as low self-control who respond “Definitely” to the statement “I drink soda pop or other sugar-sweetened beverages more often than I should”.

[21] Although a reduction in purchases is not necessarily equivalent with a reduction in (sugar) consumption. Seiler et al. (2019) show that many consumers avoid the tax in Philadelphia by shopping in neighboring jurisdictions and Aguilar et al. (2019) show that the reduction of calories from soft drinks due to the Mexican tax is offset by an increase of calories from untaxed sugary products.

Finally, we contribute to the literature on habit formation and responsiveness to taxes by providing empirical and theoretical evidence that tax hikes and cuts can have different effects. The seminal paper by Becker and Murphy (1988) already argues that a permanent change in prices of a habit-forming good may have an initially small effect on consumption that grows over time until a new steady state is reached. Zhen et al. (2011) provide empirical evidence for habit formation in demand for sugar sweetened beverages using a demand system model. Colchero et al. (2017) evaluate the long-run response to the sugar sweetened beverage tax in Mexico and find that the long-run response is in fact stronger than the short-run response. We add a new perspective to this literature and show, both empirically and theoretically, that tax increases have a smaller effect on purchases of habit-forming goods for people with low self-control. However, this effect seems to be not symmetric for tax increases and tax cuts, suggesting that individuals with low self-control find it hard to reduce consumption when prices go up but react to price incentives when prices go down.

The remainder of this chapter proceeds as follows. In Section 2.2, we present the conceptual framework that motivates our empirical analysis. Section 2.3 describes the institutional setting and the dataset that we are using. Section 2.4 specifies the empirical strategy. Section 2.5 presents the results and Section 2.6 provides a theoretical explanation for the results that we find. Section 2.7 concludes.

2.2 Conceptual Framework

In this section, we briefly summarize a key result of the sin tax literature that motivates our empirical investigation of heterogenous responses to sin taxes by self-control. O'Donoghue and Rabin (2006) and Haavio and Kotakorpi (2011) show that, in a simple two-good model, the optimal internality-correcting tax depends both on the average internality and on the covariance of the price responsiveness and the internality: The optimal tax is higher if individuals with low self-control respond stronger to price changes than individuals with high self-control and vice versa.

More formally, models in the literature typically assume that preferences can be characterized by a $\beta - \delta$ model of self-control (Laibson, 1997). That means, individuals maximize intertemporal utility:

$$U_t(u_1, ..., u_T) = u_t + \beta \sum_{\tau=t+1}^{T} \delta^{\tau-t} u_\tau. \tag{2.1}$$

Each period they receive instantaneous utility u_t and future utility is discounted by time-consistent discount factor δ and a hyperbolic discount factor β. If $\beta < 1$ the agents have a preference for immediate gratification (i.e. low self-control) and if $\beta = 1$ the agents behave time-consistently.

In a two-good model, consumer i decides whether to consume a sin good x_i that provides instantaneous utility $v(x_{it})$, but is associated with long-run costs $c(x_{i,t-1})$, and a numeraire good. Since consumers with low self-control ($\beta < 1$) underweigh the future costs of consumption, they overconsume the sin good. A social planner maximizing the long-run utility of all individuals (i.e. setting $\beta = 1$ for everyone), may decide to impose a tax t on the sin good to help consumers with low self-control to consume closer to their long-run utility. The idea is that the tax serves consumers with low self-control

as a commitment device by increasing the instantaneous costs of consumption.
Haavio and Kotakorpi (2011) show that, in this case, the optimal tax is given by

$$t = \frac{1}{N}\sum_i (1-\beta_i)c'(x_i) + \frac{\text{cov}((1-\beta)c'(x), \frac{\partial x}{\partial t})}{\partial \bar{x}/\partial t}. \tag{2.2}$$

We provide the derivation of the optimal tax formula in Appendix B.1. The first term in the optimal tax is the average internality in the population, i.e., the marginal costs that are not accounted for due to a lack of self-control. This first term is corrected by the targeting efficiency of the tax that is represented by the second term. The targeting of the tax is described by the covariance between the internality due to a lack of self-control and the responsiveness of consumption to tax changes (weighted by the average responsiveness of sin good consumption to tax changes). Intuitively, the optimal tax is larger if those with the largest internality reduce their consumption more than those without lack of self-control. In that case, the tax is relatively effective in correcting the internality. However, the tax is smaller if consumers with low self-control respond less to the tax. In that case, the distortionary effect on consumers without self-control problem overweighs the internality-correcting effect.
According to the previous literature, it is an empirical question whether the relationship between self-control and price responsiveness is positive or negative (O'Donoghue and Rabin, 2006). Hence, this is what we aim to investigate causally using the institutional setting described in the next section.

2.3 Data

2.3.1 Institutional Background

For identification, we exploit variation in two different sin taxes: the soft drink tax and the fat tax. Both were part of the Danish tax reform of 2010. The goal was to reduce income taxes and instead increase taxes on consumption goods that have detrimental effects on public health or the environment (The Danish Ministry of Taxation, 2009). Besides the taxes on soft drinks and fat, taxes on sweets, chocolate, ice cream, and tobacco were increased. Moreover, a tax on the content of sugar in all goods was planned but never realized.
The tax variation is illustrated in Figure 2.1. The first tax that we study is the tax on sugary soft drinks. The excise tax on soft drinks in Denmark has a longstanding tradition. Both its introduction and subsequent tax reforms were mainly motivated by the goal to raise tax revenues (Bergman and Hansen, 2019). However, the increase of the tax in January 2012 from 1.35 DKK to 1.98 DKK per liter (excise tax plus 25 percent value-added tax) aimed to improve public health. This is also illustrated by the fact that the tax on diet soft drinks remained constant. In previous work (Schmacker and Smed, 2020), we estimate the pass-through of the tax using a regression discontinuity approach and document a price increase by 1.01 DKK (10.7 percent) in reaction to the tax hike (see Figure B.1a in the Appendix). Hence, the tax hike is substantially overshifted, which is consistent with the study of Bergman and Hansen (2019) for earlier soft drink tax increases. In April 2013, the Danish government announced it would repeal the tax on soft drinks in order to secure jobs in the retail sector in the

Figure 2.1: Soft drink and fat tax variation in Denmark, incl. 25 percent VAT

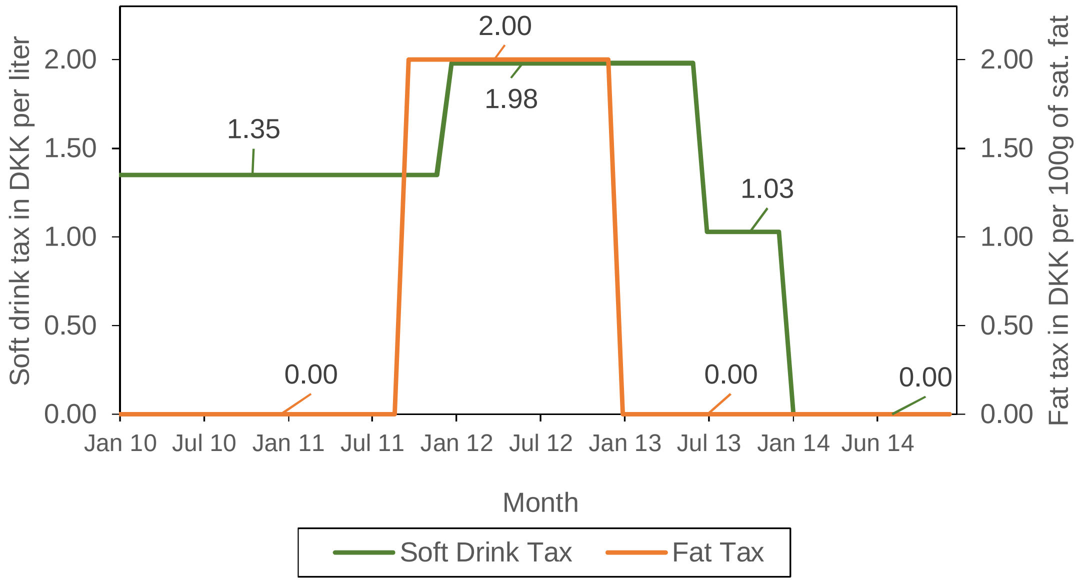

Notes: Graph shows soft drink and fat tax variation over time. The denoted taxes include 25 percent VAT.

Danish-German border region and to make up for tax revenue losses due to cross border trade. The tax was first decreased to 1.03 DKK (incl. VAT) in July 2013 and completely eliminated in January 2014. In Schmacker and Smed (2020), we estimate a price drop of 1.88 DKK (23.1 percent) in response to the tax repeal, i.e. approximately full pass-through (see Figure B.1b in the Appendix).

The second tax variation is the introduction (October 2011) and repeal (January 2013) of the fat tax. The fat tax was applied to all products that contain more than 2.3g saturated fats per 100g. It amounts to 1.60 DKK per 100g saturated fat plus 25 percent VAT, i.e. 2.00 DKK per 100g of saturated fats. Vallgårda et al. (2015) analyze the political debate around the introduction and repeal of the fat tax. They conclude that a change in the framing from public health arguments to economic arguments (cross-border shopping, administrative burden, and regressive effects on the poor) led from the introduction to the repeal. Since the tax was proportional to the amount of saturated fat, the tax affects product groups very differently. In the analysis, we consider butter since it contains a high amount of saturated fats (approximately 50 percent) and has, therefore, experienced substantial tax variation. In Appendix B.4.1, we show that the tax introduction is associated with an almost symmetric increase in butter prices by 0.76 DKK per 100g and the repeal with a decrease by 0.61 DKK per 100g.

2.3.2 Dataset

To investigate the response in purchases to the tax variation, we use household panel data from GfK Consumertracking Scandinavia for the years 2009 to 2014. Panelists are asked to track all their food purchases on a weekly basis. GfK aims for a representative panel with respect to geography, age, education, and family size. Panelists report quantities and prices paid for grocery purchases that they bring into the home. Moreover, once a year, households fill in a questionnaire on demographic and socioeconomic characteristics.

When looking at quantity purchased, we aggregate the purchases to monthly observations to account for potential stockpiling. We assign months a zero where purchases are observed but none of these purchases belongs to the product category in question (soft drinks or butter). Moreover, we use individual quantities defined as observed quantity divided by the members of the household where children aged younger than six years count as 0.5 household members.[22] Thus, we assume that the purchased products are shared equally within the household.

2.3.3 Measuring Self-Control

In 2013 and 2015, an additional survey containing a broad range of questions about self-control and dietary habits was sent to panelists. Self-control is measured using the scale developed by Tangney et al. (2004), which consists of 36 statements concerning different domains of self-control (see the items in Table B.1). The respondents indicate their approval to each of these statements on a 5-point Likert-scale. Whenever possible, we use the 2013 data and, if the panelist has not filled in the survey in 2013, we impute the missing data with data from 2015. Hence, we assume that self-control is a time-constant trait, which is supported by a high retest-reliability: among the 1,234 panelists, who have answered the self-control scale in both years, the scores from 2013 and 2015 correlate with r=0.783.

In order to reduce the large number of items and to find the latent dimension of self-control that matters for food choices, we perform a principal component factor analysis on the 36 items using all 2,387 panelists who filled in the self-control scale. Since we aim to identify the underlying factors of self-control, we do not restrict the sample to those panelists who report purchasing soft drinks or butter. Based on the original study by Tangney et al. (2004), we decide to extract five factors. In Appendix B.2, we describe the resulting factor structure.

In the analysis, we use the factor that is related to temptation tolerance and health-related habits, and, thus, is closest related to food choices. Based on the factor loadings and the responses given by the panelists, we compute a new variable that measures the level of self-control according to this factor. The 50 percent of individuals with the lowest parameter value are characterized as having low self-control and the individuals with the highest value as having high self-control.[23]

Table 2.1 illustrates that our measure of self-control correlates in meaningful ways with Body Mass Index and intentions to improve eating habits, even after controlling for socio-demographics like income, age, and education. Column (1) shows that people with low self-control have substantially higher BMI than those with high self-control and column (2) shows that they are more likely to be obese. Importantly, in column (3), we observe that respondents with low self-control indicate much more often that they would like to reduce their weight. Thus, individuals with different levels of self-control also differ in their intentions to achieve a healthy weight. Columns (4) and (5) show that the fraction

[22] This weighting is based on 2003-2006 survey data on dietary habits in Denmark (DTU Fødevareinstituttet, 2008).

[23] Since we do not want to make parametric assumptions about how self-control affects the responsiveness to tax changes, we use this non-parametric approach. In robustness checks, we show results for more granular sample splits and demonstrate that the results do not depend on the sample split.

Table 2.1: Correlations of self-control with characteristics and attitudes

	(1) Body Mass Index (BMI)	(2) Obesity (BMI>30)	(3) Intention to reduce weight	(4) "I should eat less sugar"	(5) "I should eat less animal fat"
Low self-control	2.124***	0.094***	0.202***	0.112***	0.115***
	(0.269)	(0.021)	(0.028)	(0.029)	(0.028)
Controls	Yes	Yes	Yes	Yes	Yes
Mean	26.021	0.175	0.620	0.483	0.354
Households	1237	1236	1197	1197	1197

Notes: The regressions control for income, age, education, labor market status, and number of children. Columns (1) and (2) are based on weight and height data from 2011. BMI is calculated as ([weight in kg]/[height in m]^2). Column (3) shows the fraction of respondents who indicate in the 2013 survey that they would like to weigh at least 1 kg less. Column (4) and (5) gives the fraction of respondents who indicate that they should eat "A lot less" or "A little less" sugar or animal fat to eat healthier. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

of respondents who agree that they should eat less sugar and animal fat is significantly larger among panelists with low self-control. To sum up, individuals with low self-control are more prone to risky health behavior and are aware of it, but apparently a lack of self-control prevents them from changing their eating habits. Consequently, individuals who are identified as having low self-control by this measure are those who should be targeted by an internality-correcting sin tax.[24]

2.3.4 Descriptive Statistics

In the analysis, we use data from 2009 through 2014, including only those households that report at least one purchase of the product in question per year and have responded to the self-control questionnaire. These restrictions leave us with 1,278 panelists. In Table 2.2, we show descriptive statistics of the overall sample used in the analysis, as well as descriptives of the sample split by self-control. Moreover, in the last column, we show descriptives for the unrestricted sample, which also includes panelists who report at least one purchase in every sample year but for whom we have no information on self-control.

The demographic characteristics appear quite similar across the different sample restrictions. However, there is an intuitive association between self-control and education, with high self-control respondents having higher education. In the robustness section, we address if the differential response by self-control is affected if we also control for heterogenous responses by education. Moreover, panelists who work full time are more likely to have low self-control than those working part-time and, further, the number of children is associated with lower self-control.

[24]In Table B.2, we show that these associations are substantially weaker or non-existent for the other extracted factors, highlighting that we have identified the most relevant factor of self-control in relation to food consumption.

Table 2.2: Descriptive statistics

	Overall	Low self-control	High self-control	Unrestricted sample
Equivalized household income in DKK				
<175K	18.7	17.8	19.6	19.3
175K-250K	26.1	28.0	24.4	26.0
250K-325K	18.2	17.4	18.9	17.8
325K-400K	19.6	20.1	19.1	19.5
≥400K	17.4	16.7	18.1	17.4
Age group				
<40	13.0	12.3	13.6	13.7
40-59	48.7	47.4	49.9	47.9
≥60	38.4	40.3	36.5	38.4
Labour market status				
Full time	38.7	42.3	35.2	38.5
Part time	27.2	24.4	30.0	27.7
Not employed	34.1	33.3	34.8	33.9
Education				
No tertiary education	59.4	62.8	56.3	59.4
1-3 years tertiary educ.	15.0	14.3	15.6	14.8
> 3 years tertiary educ.	25.6	22.9	28.1	25.8
Household size	1.921	1.942	1.901	1.909
	(0.985)	(1.041)	(0.928)	(0.988)
Number of child. age 0-6	0.066	0.090	0.044	0.068
	(0.321)	(0.382)	(0.249)	(0.326)
Number of child. age 7-14	0.130	0.147	0.113	0.128
	(0.458)	(0.501)	(0.413)	(0.454)
Number of child. age 15-20	0.100	0.102	0.099	0.099
	(0.367)	(0.365)	(0.370)	(0.365)
Households	1,278	623	655	1,412
Observations (Household-months)	78,086	37,940	40,146	85,349

Notes: Table shows descriptive statistics of the GfK Consumertracking Scandinavia data used in the analysis. Displayed are relative frequencies of values of categorical variables, as well as means and standard deviations (in parentheses) of continuous variables. Household income is equivalized using the OECD scale, i.e. dividing household income by the square root of the household size.

2.4 Empirical Strategy

In order to test if the demand response to tax changes differs by self-control, we estimate the within-household variation in soft drink purchases the year before and after the tax changes. Due to our bandwidth of one year, we can keep seasonal variation before and after the tax constant and also can capture changes that occur with a lag due to habit formation.
The empirical model for estimating purchase quantity in month t by consumer i is

$$\text{quantity}_{it} = X'_{it}\alpha = \alpha_0 + \alpha_1 \text{tax}_t + \alpha_2(\text{tax}_t \times \mathbb{1}(\beta_i = \beta^{high})) + \gamma_i + \eta_t + \alpha_4 Z_{it} + \epsilon_{it} \tag{2.3}$$

where the dependent variable is either the observed quantity, the purchase incidence in a given month (extensive margin), or the log-transformed quantity given a purchase (intensive margin). The variable tax_t is a dummy variable that is one after the tax change and zero before. We interact the tax dummy with indicator functions that specify if individual i is characterized by low or high levels of self-control as defined in the previous section. Hence, α_2 estimates the differential effect of the tax change on purchase quantity for consumers with high self-control compared to those with low self-control. γ_i denotes household fixed-effects, that are included to control for time-invariant unobserved heterogeneity, and η_t denote quarter fixed effects. Z_{it} is a set of household-specific controls that includes the number of kids at age groups 0-6, 7-14, and 15-20, the household size, income group, and labor market status of the main shopper.[25] In the analysis of the soft-drink tax, Z_{it} also includes the monthly average temperature in Denmark.
The main coefficient of interest is the interaction effect of the tax dummy and the self-control indicator, α_2. In order to identify if the differential responsiveness is due to self-control, we must make the following assumptions. First, we assume that consumers with low and high self-control exhibit parallel trends in consumption. We demonstrate the credibility of this assumptions and show that trends are parallel in the years absent the tax reforms. Second, we assume that differences in price responsiveness are due to self-control and not due to other correlated characteristics, like income and education. Therefore, we investigate if the differential response by self-control remains when we also interact the tax dummy with measured income and education.
As is often the case with household-level consumption data, the distribution of purchases is characterized by a mass at zero and a right-skewed distribution. Using ordinary least squares on the untransformed data can yield biased and inefficient estimates (Manning and Mullahy, 2001). For example, the presence of extreme outliers can have an undue influence on the parameter estimates. We employ multiple measures to alleviate this issue. First, we winsorize the reported quantities at the 99 percent level, i.e., the largest 1 percent of reported quantities are set to the quantity at the 99th percentile. Second, we use a two-part model that estimates, first, the likelihood to consume any soft-drinks (extensive margin) and, second, the amount of soft-drinks provided that a positive quantity is observed (intensive margin). Consequently, the expected value of the quantity is the product of the predicted purchase probability ($X'_{it}\alpha^{ext.}$) and the conditional (and re-transformed) purchase quantity ($exp(X'_{it}\alpha^{int.})$):

[25] We do not control for education since there is little within-household variation over time.

$$E(quantity_{it}|X_{it}) = (X'_{it}\alpha^{ext.}) \cdot exp(X'_{it}\alpha^{int.}) \cdot D \tag{2.4}$$

where $D = 1/N \sum \exp(ln(q_{it}) - X'_{it}\beta)$ is the Duan smearing factor that is needed for retransformation since $E(\exp(\epsilon_{it}))$ is not zero (Duan, 1983).[26] We compute the predicted purchase quantity separately for consumers with low and high self-control.

For each tax event, we consider one year before the tax change and one year after the tax change but we omit the months January and December of each year. Otherwise, we might e.g. overestimate the effect of the tax hike in January 2012 due to customers stockpiling soft drinks in December 2011 and living off stock in January 2012. In case of the tax repeal, we compare the year before the tax cut (July 2012 until June 2013, without January and December) to the year after the complete repeal (February 2014 until November 2014).

2.5 Results

In the empirical analysis, we investigate the differential responsiveness by self-control, first, for soft drink tax changes and, second, for fat tax changes. In both cases, we provide graphical evidence on the development of purchases surrounding the tax changes, before we present the regression results.

2.5.1 Differential Responsiveness to Soft Drink Tax Changes

Figure 2.2 shows predicted values after controlling for household fixed effects and the control variables specified in Section 2.4. First, the purchases of consumers with low and high self-control seem to follow parallel trends in the years before the first tax change, thus lending support to our identification strategy. When the tax is increased in 2012, soft drink purchases by consumers with low self-control did not change significantly, while we observe a significant drop for consumers with high self-control. In July 2013 the tax was cut in half and in January 2014 the tax is completely repealed. In response, we observe a marked increase in purchases by both consumer groups.

In order to quantify the purchase response to the tax variation, we show estimation results of the empirical model in Table 2.3 for the two parameters of interest (α_1 and α_2).[27] Panel A gives results for the tax hike and Panel B for the tax repeal. In the first column, we use the absolute quantity as dependent variable, in the second column the purchase incidence (extensive margin), and in the third column the log-transformed quantity given a purchase (intensive margin). In the fourth to sixth column, we add time-varying controls. The coefficients shown are the tax indicator variable, which gives the change in purchases by low self-control consumers (the reference category), and the interaction of the tax dummy with the high self-control indicator, which gives the differential change in purchases by high self-control consumers.

[26] The retransformation procedure assumes that $E[\exp(\epsilon)|q > 0, x] = 0$, i.e. that the error term is homoscedastic (Mullahy, 1998). As a robustness check, we estimate the second part using GLM, which does not require retransformation, and, thus, does not assume homoscedasticity with respect to the covariates. The results are very similar and are available upon request.

[27] The complete estimation tables are presented in Appendix B.5.

Figure 2.2: Predicted values of monthly soft drink purchase quantity by self-control

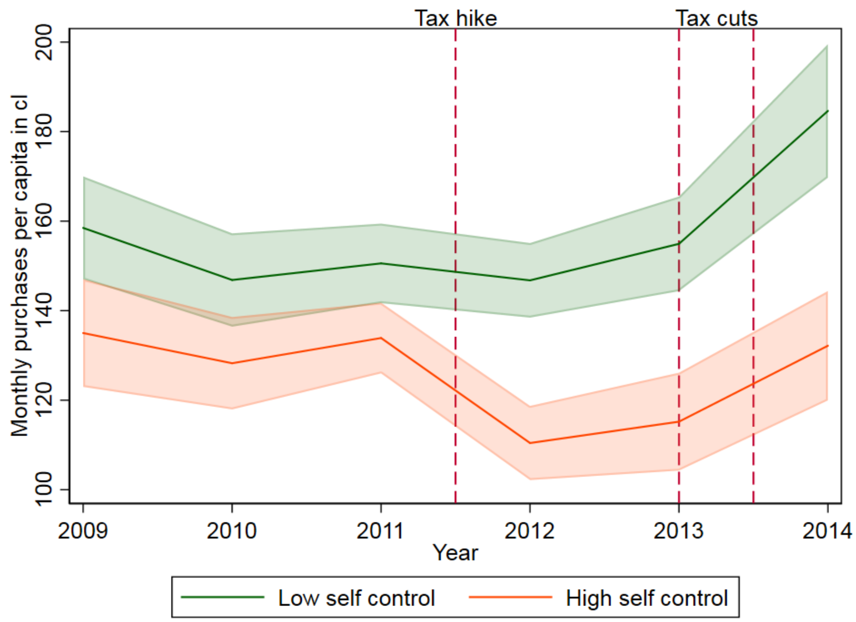

Notes: Graph shows annual predicted values after controlling for household fixed effects and the control variables specified in the Methods section, using GfK Consumertracking data. Household quantities are individualized by dividing the observed household quantity by the number of household members (children aged 0 to 6 enter as 0.5 household members). The shaded areas represent 95 percent confidence intervals. The vertical lines indicate the timing of tax changes.

Table 2.3: Soft drink purchases in response to soft drink tax changes by self-control

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Panel A: Tax Hike						
Tax Hike	0.319	-0.014	0.022	-1.912	-0.016*	-0.002
	(6.543)	(0.010)	(0.040)	(6.559)	(0.010)	(0.041)
High self-control × Tax Hike	-21.663***	-0.032**	-0.086	-21.097***	-0.030**	-0.075
	(8.198)	(0.013)	(0.054)	(8.134)	(0.013)	(0.052)
Panel B: Tax Repeal						
Tax Repeal	28.512***	0.041***	0.127***	31.827***	0.046***	0.151***
	(6.967)	(0.010)	(0.036)	(7.351)	(0.011)	(0.039)
High self-control × Tax Cut	-3.667	0.001	-0.013	-3.613	0.002	-0.015
	(8.817)	(0.014)	(0.050)	(8.814)	(0.014)	(0.050)
Controls	No	No	No	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In columns (1) and (4) the dependent variable is monthly quantity in centiliter per household member. In columns (2) and (5) it is purchase incidence in a given month. In columns (3) and (6) it is log-transformed quantity. In Panel A, we run the estimations with 1,278 households (22,197 household months) and for the intensive margin estimations with 1,104 households (7,466 household months). In Panel B, we use 1,278 households (22,747 households months) and for the intensive margin estimations with 1,122 households (7,782 household months). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.4: Change in soft-drink purchases based on predicted values from two-part model

	Low self-control	High self-control
Panel A: Tax Hike		
Relative change	-0.018	-0.190***
	(0.041)^{b}	(0.034)^{b}
Absolute change	-1.866	-20.247**
	(11.979)^{b}	(10.301)^{b}
Panel B: Tax Repeal		
Relative change	0.265***	0.274***
	(0.042)^{b}	(0.036)^{b}
Absolute change	24.804**	25.254**
	(11.790)^{b}	(9.962)^{b}

Notes: Table shows predicted values from the two-part model as described in Section 2.4. The predicted values are based on the extensive and intensive margin shown in Table 2.3. For the absolute change, the unit of measurement is in monthly centiliter per household member. Standard errors are bootstrapped with 400 replications and clustered on the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results in Panel A reveal that consumers with high self-control have decreased their purchases significantly stronger than consumers with low self-control in response to the tax hike. Consumers with low self-control have not reduced their purchases significantly from zero as the coefficient of the tax hike dummy tells us. In the second and third columns, we show the results for the extensive and the intensive margin only controlling for household fixed effects, and in the fifth and sixth column, we add time-variant control variables. On the extensive margin, consumers with high self-control have reduced the purchase probability in a given month by 3.0 percentage points more than consumers with low self-control. Moreover, in the last column it can be seen that consumers with high self-control reduce their purchase quantity by 7.5 percent more than those with low self-control, given that a purchase is observed. However, the differential response is only significant on the extensive margin. In Table 2.4, we use the estimates from the extensive and intensive margin to calculate predicted values of a two-part model, as described in Equation (2.4). The pattern resembles the results from the OLS using the untransformed quantity in the first column of Table 2.3. Consumers with low self-control exhibit a change in purchases that is not significantly different from zero. In contrast, consumers with high self-control reduce their purchases significantly, both in relative and absolute terms.

In Panel B of Table 2.3, we conduct the same exercise for the tax repeal. Here, we compare purchases one year after the tax repeal to one year before the first tax cut. The tax repeal dummy shows that consumers with low self-control have increased their purchases of soft-drinks in absolute terms, both on the extensive and intensive margins. Again, the results are not strongly affected by adding time-variant control variables. While the probability to purchase a soft-drink in a given month increases by 4.6 percentage points, the quantity conditional on a purchase increases by 15.1 percent. However, this time we do not observe a differential response by high self-control consumers. In Panel B of Table 2.4, the predicted values from the two-part model reiterate that the absolute and relative increases in purchases are, in fact, very similar across the consumer groups.

Robustness of Results

Our analysis assumes that, absent the tax changes, consumers with low and high levels of self-control would have exhibited the same trends. While we cannot directly test this assumption, we provide credibility for it by running the same estimation for placebo tax changes preceding the actual tax changes. In Table B.3, we complete this exercise for placebo tax changes on January 1, 2010, and January 1, 2011. We observe no differential change in purchases by high self-control consumers, thus lending support to the parallel trend assumption.

Another identifying assumption is that the change in purchases can really be attributed to differences in self-control and not to other correlated demographic variables like income and education. For example, one might suspect that self-control is positively correlated with income and that the reason for a differential response to the tax hike is liquidity constraints and not self-control.[28] Since we observe education and income in the data, we can address this concern by adding the triple interactions of tax and self-control with income and education to the main specification and check if the interaction with self-control becomes less pronounced.

In columns (1) to (3) of Table B.4, we re-run our main specification for the tax hike but add an interaction with a dummy indicating whether an individual is in the top half of the distribution of equivalized incomes. In Panel A, we observe that the coefficients for the interaction of tax hike and self-control are of a similar magnitude compared to our main specification. As is expected when including further interaction terms, the standard errors of the coefficients become larger and the coefficients are less significant compared to the main specification. In Panel B, we find that also the results for the tax repeal are reinforced. We do not find evidence that high and low income earners respond differently to the tax variation.

Moreover, in columns (4) to (6) of Table B.4, we add the triple interaction with an indicator for high education. High education means that the panelist has attended at least one year of tertiary education whereas low education indicates at most vocational education. In Panel A, the magnitude of the interaction coefficient of tax hike and self-control is almost unaffected compared to the main specification. Moreover, responsiveness to the tax hike does not differ significantly by education. In Panel B, we conduct the same exercise for the tax repeal. Again, there is not a clear pattern indicating that responsiveness to the tax differs by education. There is some evidence suggesting that for highly educated consumers there is a stronger differential response by self-control. However, as in our main specification we observe no significant interaction coefficient of the tax repeal and self-control.

It is conceivable that measured self-control is correlated with tastes for unhealthy food. To check if the differential response by self-control can be attributed to differences in taste, we add the triple interaction with a dummy variable that indicates if consumers approve to the statement "I believe I would make healthier food choices if unhealthy food was less tasty". In the first to third columns of Panel A in Table B.5, we observe that the interaction with high self-control remains of a similar magnitude irrespective of the taste for unhealthy food. While consumers who like unhealthy food seem to be less likely to reduce their purchases in response to the tax hike, this leaves the differential

[28] However, if that was the case, we would expect consumers with low self-control (and low income) to reduce purchases *more* than consumers with high self-control (and high income).

response by self-control almost unaffected. In Panel B, there is – as in the main specification – not much evidence for a differential effect by self-control.

Another alternative explanation is that our measure of self-control is associated with knowledge about the healthiness of food and that this drives the differential response. To account for that we add the triple interaction with a dummy indicating consumers who approve to the statement "I believe I would make healthier food choices if I had more information on how to eat healthy". In the fourth to sixth columns of Panel A in Table B.5, we show the results for the tax hike. Most importantly, the interaction with self-control remains of similar magnitude and significant. Further, in Panel B we observe that accounting for differences in nutritional knowledge does not alter the coefficient of self-control much.

As further robustness tests, we re-estimate our main specification on the subsample of single households. Thereby, we can be sure that measured self-control coincides with the self-control of the individual who is solely responsible for the purchasing decisions. Table B.6 presents the results and reiterates the previous finding: High self-control individuals reduce their purchases significantly more than low self-control consumers when the tax goes up, and the interaction coefficient is even larger than in the full sample. However, there is no differential change that is significantly different from zero when taxes go down. Moreover, in Table B.7 in Appendix B.3.2, we re-run the estimations using a sample split of four self-control quartiles. For the tax hike, we observe that the higher the level of self-control the stronger is the reduction in purchases. For the tax repeal, there is no significant difference between the groups.

As mentioned above the tax on soft drinks was mainly repealed to reduce cross-border shopping in Germany. In general, this should not be a concern for our analysis since in the GfK Consumertracking data, consumers also report purchases abroad. However, one may be concerned that cross-border purchases are underreported and consumers engage differently in border-shopping depending on self-control. Hence, in Table B.8 we add the triple interaction of the tax dummy and self-control with a dummy variable indicating whether the consumer has access to the German border without using a toll bridge or ferry.[29] Thus, the "No Toll" indicator is a proxy for how easy and economic it is to buy groceries in Germany. In Panel A, we observe that for consumers in "Toll" regions (i.e. where the border is not easily accessible) the difference between low and high self-control is even stronger than in the main specification. However, the difference between low and high self-control is slightly smaller in "No Toll" regions, although the triple difference estimator is not statistically significant. This seems to suggest that consumers with high self-control do not reduce their purchases as much when there are close-by opportunities to avoid the tax. Unfortunately, as there are relatively few panelists who live in the region close to the border, we lack the statistical power to investigate the relation between self-control and cross-border shopping more closely. In Panel B, we observe that in the "No Toll" region (i.e. close to the border) consumers with low self-control increased their purchases more after the tax is repealed. But also here a closer investigation is not feasible due to a lack of statistical power.

[29]While households in Jutland and Funen do not have to use a toll bridge or ferry to reach the German border, households in Sealand, Copenhagen or Bornholm must. The costs to use the ferry or bridge for a standard car start at 30 Euros each way. In Schmacker and Smed (2020) we provide descriptive evidence that this distinction captures differences in cross-border shopping reasonably well.

Figure 2.3: Predicted values of monthly butter purchase quantity by self-control

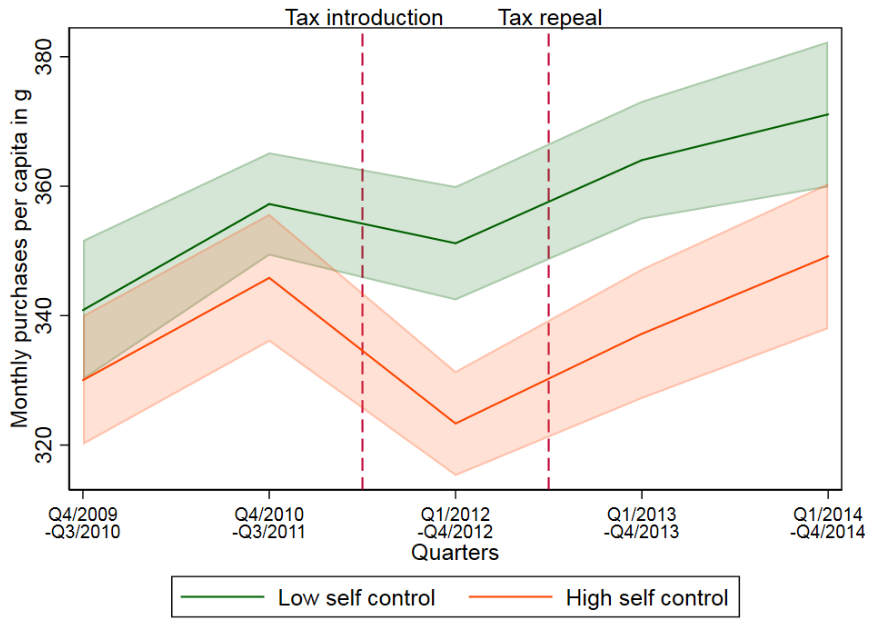

Notes: Graph shows predicted values after controlling for household fixed effects and the control variables specified in the methods section. Household quantities are individualized by dividing the observed household quantity by the number of household members (children aged 0 to 6 enter as 0.5 household members). The shaded areas represent 95 percent confidence intervals. The vertical lines indicate the timing of tax changes.

2.5.2 Differential Responsiveness to Fat Tax

In the previous sections, we show that consumers with low self-control respond less to increasing soft drink taxes than consumers with high self-control. In contrast, when soft drink taxes are cut, there is not a systematically different response. In this section, we check whether this pattern is particular to soft drink tax changes or whether it also emerges for the fat tax introduction and repeal.

In the following, we look at butter as it is one of the goods that contains the most saturated fat and is frequently purchased. The analysis of the fat tax complements the soft drink tax analysis in several dimensions. First, unlike the soft drink tax, the magnitude of the fat tax variation is very similar for tax hikes and cuts (see Appendix B.4.1). Hence, we can exclude that a difference in responsiveness is due to low and high self-control consumers responding differently to larger and smaller tax variation. Second, by looking at butter, we can exclude that the differential responsiveness is explained by low and high self-control consumers having different preferences for sugar. If we find a similar pattern for butter, there is further evidence that the reason for the differential responsiveness is due to self-control. We run the same estimations as described in Section 2.4 on the data for butter.[30] Figure 2.3 shows predicted values for butter purchases over the years. Since the tax was in place from the beginning

[30]The estimations mirror the estimations for soft drinks. The only notable differences are, first, that we restrict the sample to households who report a butter purchase in the years 2010 through 2013 since the tax variation occurs earlier in time. Second, we do not include the average temperature as a control variable since temperature is arguably less relevant for butter demand than it is for soft drink demand.

Table 2.5: Butter purchases in response to fat tax by self-control

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Panel A: Tax Introduction						
Tax Introduction	-11.926**	-0.018**	-0.023*	-11.224*	-0.018**	-0.026*
	(5.908)	(0.008)	(0.014)	(5.974)	(0.008)	(0.014)
High self-control × Tax	-15.692*	-0.022**	-0.007	-16.440**	-0.022**	-0.008
	(8.416)	(0.011)	(0.019)	(8.380)	(0.011)	(0.019)
Panel B: Tax Repeal						
Tax Repeal	12.913**	0.018**	0.027*	10.208	0.017**	0.027*
	(5.996)	(0.007)	(0.014)	(6.562)	(0.008)	(0.016)
High self-control × No Tax	-0.183	0.002	-0.008	-2.291	0.002	-0.013
	(8.467)	(0.010)	(0.021)	(8.516)	(0.010)	(0.021)
Controls	No	No	No	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In columns (1) and (4) the dependent variable is monthly quantity in gram per household member. In columns (2) and (5) it is purchase incidence in a given month. In columns (3) and (6) it is log-transformed quantity. In Panel A, we run the estimations with 1,324 households (27,192 household months) and for the intensive margin estimations with 1,291 households (17,056 household months). In Panel B, we use 1,323 households (27,507 households months) and for the intensive margin estimations with 1,298 households (17,460 household months). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

of the fourth quarter 2011 (starting October 2011) until the end of the fourth quarter 2012 (ending December 2012), we must exclude one of the five taxed quarters to compare entire years.[31] We observe that in the pre-tax years, consumers with low self-control purchase more butter than those with high self-control, but the overlapping confidence intervals suggest that the difference is not significant. When the tax is introduced, we find, once again, that consumers with high self-control reduce their purchases significantly more than those with low self-control. Furthermore, when the tax is repealed, both consumer groups increase their purchases to a similar extent, such that the difference in purchases between consumer groups remains until two years after the tax repeal.

We show estimation results of the coefficients of interest from the empirical model in Table 2.5.[32] Panel A illustrates that consumers with high self-control reduce their purchases significantly stronger than consumers with low self-control in response to the fat tax introduction. As seen in columns (2) and (4), the difference is mainly driven by a response on the extensive margin. The predicted values from the two-part model in Panel A in Table 2.6 illustrate that both in relative and absolute terms, the purchase response by high self-control consumers is stronger.

In Panel B of Table 2.5, we run the estimation for the tax repeal. However, now we do not find a differential response to the tax repeal, as illustrated by the small and insignificant interaction coefficients. Instead, we find an increase in purchases by all consumer groups. While the significance of the purchase response in Table 2.5 depends on the inclusion of covariates, it is strongly significant using the more robust two-part model, as seen in Panel B in Table 2.6. According to the two-part model, both the relative and absolute response is similar among low and high self-control consumers.

[31] In Figure 2.3 we exclude the fourth quarter of 2011. In Figure B.3 we show that we observe a similar pattern when we alternatively exclude the fourth quarter of 2012.

[32] The complete estimation tables are presented in Appendix B.5.

Table 2.6: Change in butter purchases based on predicted values from two-part model by self-control

	Low self-control	High self-control
Panel A: Tax Hike		
Relative change	-0.053***	-0.094***
	(0.014)^{b}	(0.013)^{b}
Absolute change	-16.071**	-28.487***
	(6.646)^{b}	(6.287)^{b}
Panel B: Tax Repeal		
Relative change	0.053***	0.057***
	(0.014)^{b}	(0.016)^{b}
Absolute change	15.270**	16.920**
	(6.712)^{b}	(7.415)^{b}

Notes: Table shows predicted values from the two-part model as described in Section 2.4. The predicted values are based on the extensive and intensive margin shown in Table 2.5. Standard errors are bootstrapped with 400 replications and clustered on the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Robustness of Results

In Table B.10, the results of placebo tax changes in January 2010 and October 2010 are shown. Most importantly, the interaction coefficients, which measure differential changes in response to the placebo tax changes, are insignificant and close to zero. There is a significantly positive common trend between 2009 and 2010, but the trend does not differ between consumer groups, as can be seen in Figure 2.3.
In Appendix B.4.3 we control for further interactions with education, income, tastes for unhealthy foods, and nutritional knowledge. As in the case of soft drinks, we find in Table B.11 that controlling for heterogenous responses by education does not explain the differential response by self-control to the tax introduction. Interestingly, we find in Table B.12 that the differential response by self-control to the tax introduction seems to be mainly driven by consumers with high income. In columns (1) to (3) of Table B.13, we additionally control for heterogeneous effects by taste for unhealthy food. The interaction of the tax dummy with high self-control has a similar magnitude and is still significant. In columns (4) to (6), we add the triple interaction with self-reported nutritional knowledge. Again, the magnitude of the interaction of the tax with self-control remains of a similar, but slightly smaller, magnitude. While the general pattern maintains, the results appear slightly more noisy than in the case of soft drinks. This could be explained by stronger controversies about the health effects of saturated fat (see the summary of the Danish public discourse in Vallgårda et al. (2015)).
In sum, we find evidence supporting the findings of the soft drink tax analysis. Also in response to the fat tax, consumers with low self-control respond less to increasing prices. This difference is primarily driven by consumers with high income. When the tax is repealed, consumers with high and low self-control increase their purchases just as much.

2.6 Price Responsiveness in a Model of Habit Formation

In the empirical analysis, we document an asymmetry in responses to tax increases and decreases depending on self-control. This asymmetry is hard to reconcile with standard models since price elasticities are typically symmetric to prices going up or down. In the following, we show that the

Table 2.7: Correlations of self-control with habit and addiction

	(1) "I am addicted to sugar"	(2) "I am addicted to fat"	(3) "Hard to establish healthy eating habits"
Low self-control	0.098***	0.066***	0.177***
	(0.026)	(0.020)	(0.026)
Controls	Yes	Yes	Yes
Mean	0.297	0.131	0.287
Households	1197	1197	1197

Notes: The dependent variable in each column is the fraction of panelists in the GfK ConsumerTracking panel who answer "Somewhat agree" or "Totally agree" to the respective statement. The regressions control for income, age, education, labor market status, and number of children. The complete statement in Column (3) is "I find it harder to establish healthy eating habits than it is to establish unhealthy eating habits". Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

asymmetry is to be expected when taking into account habit formation (Becker and Murphy, 1988). Previous analysis of price responsiveness by self-control have ignored habit formation although there is evidence that many sin goods, as for example sugary soft drinks, are habituating (Zhen et al., 2011) or have addictive properties (Ahmed et al., 2013). Also in our dataset there is descriptive evidence that habit formation could play an important role for the differential responsiveness: Table 2.7 shows that consumers with low self-control are more likely to agree to the statement that they are addicted to sugar or fat and that they find it harder to establish healthy eating habits than unhealthy eating habits.

In the following, we derive results about the price responsiveness by self-control when rational habit formation is taken into account. In such a model, habit formation and addiction can be used interchangeably since they rest on the same mechanism: Consumption today increases the utility from consumption in the future due to intertemporal complementarities. If an individual is aware of this property and takes it into account, we call it rational habit formation. The model of rational habit formation is based on the exposition in O'Donoghue and Rabin (2002). From here on, we adopt their approach to model the discrete choice of an agent to either engage in a habit-forming activity or not (i.e. an individual can consume or abstain). However, unlike in that paper, we introduce heterogeneity in self-control and focus on differential responses to tax variation by self-control.

Agents get utility in each period $t = (1, ..., T)$ with $T \to \infty$ from either consuming a sin good ($a_t = 1$) or abstaining ($a_t = 0$). By consuming the sin good they build up a habit stock k that evolves according to

$$k_t = \gamma k_{t-1} + a_{t-1}. \tag{2.5}$$

The habit stock in period t depends on the stock in the previous period, which decays with $\gamma \in [0, 1)$ and replenishes if the agents have consumed in the previous period (Becker and Murphy, 1988). The instantaneous utility of consumption is given by

$$u_t(a_t, k_t) = \begin{cases} v_t - p_t - c(k_t) & \text{if } a_t = 1 \\ -c(k_t) - g(k_t) & \text{if } a_t = 0 \end{cases} \tag{2.6}$$

and depends on an exogenous preference for the sin good v_t, the level of habituation k_t, and the price p_t. Consuming sin goods is associated with a negative internality ($c(k_t) > 0$), i.e. having consumed sin goods in the past has a negative effect on utility today. The internality costs of past consumption are incurred irrespective of today's consumption as, for example, the adverse health effects of being obese. For simplicity, we assume the internality costs to be linearly increasing in k with $c'(k_t) > 0$ and $c''(k_t) = 0$. Moreover, quitting consumption is associated with withdrawal costs ($g(k_t) > 0$), which are higher the more habituated an agent is, i.e. consumption is habit-forming. We assume that the withdrawal costs increase with the habit level ($g'(k_t) > 0$) and are weakly convex ($g''(k_t) \geq 0$).

In this model, v_t is exogenously given and assumed to be constant over time: $v_t = (\bar{v}, ...)$. The price p_t can be changed by the policy-maker by changing the tax rate but the individual takes p_t as given and assumes that it will not change in the future. In contrast, k_t depends on past decisions. Forward-looking agents anticipate that their current decisions will impact their future utility and will maximize for all periods s in $(t, t+1, ..., T)$ with $T = \infty$:

$$U_t(a, k_t) = \begin{cases} u_t(a_t, k_t) + \beta \sum_{\tau=t+1}^{T} \delta^{\tau-t} u_\tau(a_\tau, \overbrace{\gamma k_{\tau-1} + 1}^{k_\tau}) & \text{if } a_t = 1 \\ u_t(a_t, k_t) + \beta \sum_{\tau=t+1}^{T} \delta^{\tau-t} u_\tau(a_\tau, \gamma k_{\tau-1}) & \text{if } a_t = 0 \end{cases} \tag{2.7}$$

where in the first case the consumer decides to consume and in the second case to abstain in period t. We assume that consumers follow the strategy to either consume forever or to abstain forever. The reason is that if consumption is habit-forming (if $g(k_t) > 0$), it becomes harder to quit tomorrow compared to today. Hence, a consumer who decides to quit would rather quit today than at some point in the future.[33] An agent who consumes will eventually reach the steady-state habit stock $k^{max} \equiv \sum_{t=1}^{\infty} \gamma^{t-1} = \frac{1}{1-\gamma}$, while an individual who abstains approaches $k^{min} = 0$. In the following, we consider the case of naïve present-biased consumers, i.e. the consumers are not aware of their present-bias problem and believe they will behave as time-consistent individuals from the next period on.[34]

Assume consumers differ in their initial habit level k_{ti} and in their self-control β_i. Their initial habit level k_{ti} is independently drawn from a distribution that is characterized by a function K and their self-control β_i is independently drawn from a distribution F. Both K and F are continuous and have strictly positive density over their support $\beta \in (0, 1]$ and $k_t \in [0, k^{max}]$, respectively. Given their habit level k_{ti} and self-control β_i, a consumer would decide to consume if the utility from consumption

[33] O'Donoghue and Rabin (2002) show that for stationary preferences this is indeed the only perception-perfect strategy for time-consistent individuals with $\beta = 1$. For consumers with imperfect self-control ($\beta < 1$) there is another perception-perfect strategy where they plan to consume once and abstain thereafter (although they will not actually stop consuming). However, in this context, we do not consider the latter strategy.

[34] Naïve present-bias is a reasonable assumption for consuming soft-drinks as there is no effective commitment device that a sophisticated consumer could employ. See Gottlieb (2008) for a discussion.

starting today (i.e. $a_t = 1$ for all periods) exceeds the utility from abstaining starting today (i.e. $a_t = 0$ for all periods):

$$\bar{v} - p_t - c(k_{ti}) + \beta_i \sum_{\tau=t+1}^{\infty} \delta^{\tau-t}[\bar{v} - p_\tau - c(\sum_{n=1}^{\tau-t} \gamma^{n-1} + \gamma^{\tau-t} k_{ti})]$$
$$\geq -c(k_{ti}) - g(k_{ti}) + \beta_i \sum_{\tau=t+1}^{\infty} \delta^{\tau-t}[-c(\gamma^{\tau-t} k_{ti}) - g(\gamma^{\tau-t} k_{ti})] \quad (2.8)$$

Intuitively, an individual consumes the sin good if the utility from consumption, less the price and the internality costs in the current and all discounted future periods, are weakly larger than the internality and withdrawal costs incurred in this and the following periods due to the current level of k_{ti}.
We ensure a cut-off equilibrium in the sense that, for all k_{ti}, every individual (weakly) above a certain threshold ($\beta_i \geq \tilde{\beta}$) consumes the sin good ($a = 1$) and below the threshold ($\beta_i < \tilde{\beta}$) does not ($a = 0$). Formally, this threshold is defined by equation (2.8) with equality, or equivalently, by

$$\tilde{\beta} = -\frac{\bar{v} - p_t + g(k_t)}{\sum_{\tau=t+1}^{\infty} \delta^{\tau-t}[\bar{v} - p_\tau - c(\sum_{n=1}^{\tau-t} \gamma^{n-1}) + g(\gamma^{\tau-t} k_t)]}. \quad (2.9)$$

While the numerator in (2.9) describes the utility from consumption in the current period and is positive, the denominator describes the utility from consumption in all future periods and is negative.[35]
Define the utility from future consumption by Ψ.
To investigate how the cut-off type changes with a surprising price change, we differentiate (2.9) with respect to the price:

$$\frac{\partial \tilde{\beta}}{\partial p} = \frac{1 + \frac{\delta}{1-\delta} \tilde{\beta}}{\Psi} < 0 \quad (2.10)$$

For every k_{ti}, a price increase reduces the level of self-control, below which an individual finds it still worthwhile to consume. The reason is that an increasing price decreases utility from consumption today and in all future periods. Hence, we expect a tax hike to decrease consumption and a tax cut to increase consumption.
In the following, we focus on the question whether consumers with high and low levels of self-control are more likely to respond to price changes. Therefore, we differentiate (2.10) with respect to $\tilde{\beta}$:

[35] Define the cut-off implicitly by $J(k_t, \tilde{\beta}) = \bar{v} - p_t + g(k_t) + \tilde{\beta} \sum_{\tau=t+1}^{\infty} \delta^{\tau-t}[\bar{v} - p_\tau - c(\sum_{n=1}^{\tau-t} \gamma^{n-1}) + g(\gamma^{\tau-t} k_t)]$. To ensure existence and uniqueness, we assume that for every k_t, an individual with $\beta \to 0$ consumes the sin good ($J(k_t, \beta) > 0$) and an individual with $\beta = 1$ does not ($J(k_t, \beta) < 0$). If $J(k_t, \beta)$ is monotonically falling in β, the cutoff $\tilde{\beta}$ exists and is unique. We know that this is fulfilled since the denominator in (2.9) is negative. The proof is by contradiction: Suppose not. Since $\beta \in (0, 1]$, we know that either the numerator or denominator is positive while the other is negative. If the assumption was true, the numerator would be negative and the denominator positive. But since $\gamma \in [0, 1)$, every individual summand in the denominator is smaller than the numerator. However, then the numerator cannot be negative while the denominator is positive, which contradicts the assumption.

$$\frac{\partial^2 \tilde{\beta}}{\partial p \partial \tilde{\beta}} = \frac{\frac{\delta}{1-\delta}}{\Psi} < 0 \tag{2.11}$$

For every k_{ti}, a higher level of self-control implies a more negative price responsiveness. The reason is that individuals with high self-control take the future price change more into account. Hence, we predict that consumers with high self-control respond more to taxes than consumers with low self-control.

Result 2.1. *Consumers with high self-control are more likely to react to price changes than consumers with low self-control.*

Next, we are interested in the question if consumers react symmetrically to a tax hike and a subsequent tax cut. Here, we have to take into account that the habit stock k_t changes from one period to the next. Since we expect more consumers with a high level of self-control to respond to a tax hike (cf. Result 2.1), there are more high self-control consumers whose habit stock decreases. To make predictions regarding the question whether the response to a tax cut is symmetric, we have to evaluate how the price responsiveness depends on the habit stock.
Therefore, we differentiate (2.10) with respect to k_t:

$$\frac{\partial^2 \tilde{\beta}}{\partial p \partial k_t} = -\frac{(2\delta\beta + (1-\delta)) \sum_{\tau=t+1}^{\infty} (\delta\gamma)^{\tau-t} g'(\gamma^{\tau-t} k_t) + \delta g'(k_t)}{(1-\delta)\Psi^2} < 0 \tag{2.12}$$

The derivative is negative. Hence, the lower the habit stock, the less negative is the price responsiveness. The intuition is as follows: Suppose an individual with $(\beta_i, k_{ti}) = (\tilde{\beta}, k_{ti})$ consumes. A tax is introduced that increases the price, leading the individual to stop consuming. Hence, the habit stock k_{ti} goes down. In the next period, the tax is repealed, leading the price to return to its original level. However, since the individual now has a lower habit stock, she no longer finds it appealing to resume consumption again. The described effect is more pronounced for individuals with high self-control since, according to Result 2.1, we expect them to respond more strongly to the tax hike.

Result 2.2. *The difference in price responsiveness between low and high self-control is smaller when a tax cut follows a tax hike.*

2.7 Conclusion

Both in the policy debate and in the economic literature it is argued that sin taxes can help consumers with low self-control to act more in accordance with their own long-run interest. However, this requires that consumers with low self-control respond to tax changes by reducing consumption. This paper presents evidence that consumers with low self-control respond systematically less to increasing soft drink and fat taxes than high self-control consumers. However, we find no difference between the groups when the tax is reduced, indicating that it is not just a difference in price elasticity between the groups. We show in a theoretical model that this pattern can be explained by (rational) habit formation, an aspect so far largely neglected in the literature. The idea is that consumers who buy, e.g., soft drinks establish a habit for it and know that they will be more likely to also purchase in the future. If the taxed good is habituating (which is reasonable for many sin goods), sin taxes of modest magnitude may be less effective than previously thought in correcting internalities.

Our results suggest that other policy measures may be required to help consumers with low self-control to act according to their long-run interest. It is worth to consider, for example, time- and place-based restrictions regarding the sale of sugar sweetened beverages, as many jurisdiction have implemented for alcohol. Governments may also consider limiting the amount of sugar that beverages are allowed to contain or think about a ban on advertising sugary products.

It has to be noted that sin taxes can still correct externalities on public health, even if consumers with low self-control are not successfully targeted by these taxes. Sin taxes can make those consumers, who do not reduce their purchases, come up for the arising social costs of consumption. Furthermore, while consumers with low self-control may not respond to the price incentives themselves, smart sin tax design can still improve the diets of individuals with low self-control. If taxes are proportional to the harmful ingredient (e.g. sugar in soft drinks), producers are incentivized to make their product less unhealthy, as documented for the tiered soft drink tax in the UK (Public Health England, 2019). Since the Danish soft drink tax was volumetric, this incentive was not given. Moreover, taxes that increase the prices of the unhealthiest products the most, may invite consumers to substitute to less unhealthy alternatives (Grummon et al., 2019).

Chapter 3

Soft Drink Taxation and Habit Formation

3.1 Introduction

In the fight against obesity, the World Health Organization (WHO, 2016) recommends the introduction of a tax on sugary soft drinks. Although these taxes have been studied extensively (see the surveys in Cawley et al., 2019b; Allcott et al., 2019b), it is still under dispute whether they are effective in curbing consumption. Hence, there is a sustained need for demand models that allow to analyze the effects of soft drink taxes before they are introduced.

In order to simulate the total effect of soft drink taxes on purchases, it is important to take into account consumption dynamics. On the one hand, since consumers can stockpile soft drinks when prices are low, static demand models tend to overestimate price elasticities (Wang, 2015). On the other hand, if consumption of soft drinks is habit-forming, static demand models underestimate long-run price elasticities. The reason is that the tax does not just discourage instantaneous consumption, but due to habit formation, it also reduces the utility of consumption in the next periods. In fact, biological research shows evidence for addiction to sugar and sweetness (Ahmed et al., 2013; Mennella et al., 2016) and market-level data shows a high degree of habit persistence in demand for soft drinks (Zhen et al., 2011). However, until now the literature on soft drink taxation has largely neglected these two sources of state dependence and assumes consumption to be time-separable.

In this paper, I estimate a structural model of soft drink demand that accounts for both stockpiling and habit formation. I use weekly scanner data of household purchases in the US from 2003 to 2004 in order to analyze the effects of soft drink taxes in the presence of state dependence. The analysis proceeds in two steps. In the first step, I provide descriptive evidence for positive and negative state dependence. I follow the approach in Tuchman (2019) and find a very similar dynamic purchasing pattern compared to her analysis of the demand for cigarettes. After controlling for stockpiling, the proxy variables for habit formation have the expected sign and are highly significant. This highlights the importance of incorporating habit formation into the demand model. In the second step, I estimate a discrete choice model of soft drink demand. The model incorporates habit formation, as represented by the lagged purchase of the previous period, and stockpiling, as proxied by purchases made on sale. I estimate the model using nested logit to allow for correlated shocks to utility. Moreover, the panel data allows to control for unobserved heterogeneity in tastes, which is important in order to distinguish true from spurious state dependence (Heckman, 1981). The latter describes the situation when state dependence mainly picks up persistent taste differences.

The estimation results of the structural model provide evidence for habit formation in soft drink demand. As expected, the coefficients for state dependence decrease after controlling for persistent taste differences, but they remain sizeable. The effect of state dependence can be illustrated by the difference between short-run and long-run price elasticities. While the former only considers the price effect on instantaneous demand, the latter takes into account that a changed consumption pattern also affects future demand. I estimate a short-run price elasticity of 0.87 for sugary soft drinks and a long-run price elasticity of 1.07. Hence, the long-run price elasticity is approximately 20 percent larger than the short-run elasticity.

I use the estimated model to simulate the effect of different soft drink taxes on demand. The simulated taxes resemble actually implemented tax designs: (1) an excise tax on sugary beverages (as, e.g., in Mexico), (2) an *ad valorem* tax on sugary soft drinks (as in Chile), and (3) an excise tax on all soft drinks (as in France). The taxes are calibrated such that they generate the same relative price increase for the average product (22 percent). The simulation results show that the excise tax on sugary beverages is most effective and reduces the probability to buy a sugary soft drink by 24.4 percent. The *ad valorem* tax leads to a smaller reduction of 22.0 percent and consumers substitute to larger packaging sizes that experience a smaller price increase. The excise tax on all soft drinks leads to the smallest reduction by only 16.9 percent as it does not incentivize consumers to switch to diet soft drinks. The long-run responses to the simulated taxes are between 16 and 23 percent larger than the short-run responses. Finally, the taxes lead to a relatively uniform reduction in purchases across income groups and across household sizes.

This paper contributes to the literature that uses demand models to simulate the effect of soft drink taxes. This literature uses naturally occurring price variation and structural models to estimate demand parameters. While the literature used to rely mostly on market level data (see the surveys in Andreyeva et al., 2010; Powell et al., 2013), a rapidly growing number of studies uses disaggregated purchase data to estimate demand on the consumer level (Allcott et al., 2019a; Dubois et al., 2019; O'Connell and Smith, 2020; Wang, 2015; Bonnet and Réquillart, 2013). The use of disaggregated data has the advantage that it allows to study differentiated products and heterogeneity in consumer demand. However, the majority of these papers assumes consumption to be time-separable. Notable exceptions are Wang (2015) and Serse (2019).[36] Wang (2015) estimates a dynamic discrete choice model that explicitly models stockpiling, i.e., households stock up their inventory during price promotions and are less likely to buy when prices go up again. She finds that the price elasticity is substantially smaller compared to the results from static demand models. However, she does not consider positive state dependence arising from habit formation or addiction, which could lead to underestimation of

[36] O'Connell and Smith (2020) provide reduced-form evidence suggesting that habit formation and stockpiling are less relevant in their dataset from the United Kingdom compared to previous studies using US data.

the long-run price elasticity. Although she controls for persistent brand and sugary/diet preferences, I find that there is evidence for state dependence over and above unobserved heterogeneity. Serse (2019) estimates a mixed logit model for cola demand that incorporates habit formation. He circumvents the issue of stockpiling by modeling product choice conditional on making a purchase. He finds evidence for positive state dependence, which leads the long-run demand response to simulated taxes to be larger than the instantaneous response. However, he focuses only on the substitution between regular and diet cola beverages. In this paper, I also model the decision to make a purchase in order to provide a more comprehensive picture of the overall purchase response to taxes.

Moreover, I contribute to a literature that studies habit formation in soft drinks using other methodological approaches. Zhen et al. (2011) use market level data and a dynamic almost ideal demand system (AIDS). They show that long-run tax revenue is 15 to 20 percent lower than short-run revenue when habit formation is considered. Hence, they find a similar impact of habit formation compared to this study despite using different methods. Colchero et al. (2017) find a stronger long-term than short-term effect of the Mexican soft drink tax and conjecture that habit formation is the reason for it. Liem and de Graaf (2004) conduct randomized experiments with children and find that repeated exposure to sweetened beverages increases preference for sweetness already after a couple of days.

Finally, this paper relates to the literature that empirically estimates habit formation and state dependence. Here, habit formation is construed as intertemporal complementarities in consumption and is, following the convention in the economics literature, used synonymously with the concept of addiction (Becker and Murphy, 1988; Pollak, 1970). While the marketing literature typically understands habit formation as complementarities in brand choice probability (e.g. Keane, 1997; Dubé et al., 2010; Guadagni and Little, 1983), I consider category-level habit formation. This approach is also pursued in previous studies that use both static (Tuchman, 2019) and dynamic structural models (Gordon and Sun, 2015) of addiction to analyze demand for cigarettes. In this paper, I use a static demand model, in which the previous purchase enters utility as a lagged variable. Hence, I employ a myopic model of habit formation.

The rest of the article proceeds as follows. Section 3.2 describes the data and Section 3.3 provides reduced-form evidence indicating that stockpiling and habit formation are relevant. Section 3.4 introduces the model and Section 3.5 presents the results from estimating the model, including tax simulations using the estimated parameters.

3.2 Data

I use weekly scanner data collected by Information Resources, Inc. (IRI, Bronnenberg et al., 2008), which covers the years 2003 and 2004. The data consists of a household panel and a store panel. The household panel is comprised of households from Eau Claire, Wisconsin, and Pittsfield, Massachussetts. It tracks each household's purchases with respect to universal product code (UPC), quantity, store, and price paid. Moreover, some demographic characteristics of panelists such as income and household size are captured. The household panel is matched to the store panel based on the store, in which a household shopped in the respective week. The store panel contains store-level data on the prices

Table 3.1: Market shares of soft drink products

Brand	Brand share (%)	Type	Can: < 12 Cans	Can: ≥ 12 Cans	Bottle: < 2 Liter	Bottle: ≥ 2 Liter
Coca Cola	27.96	Sugary	0.87	6.68	0.49	4.47
		Diet	0.62	9.41	0.79	4.62
Pepsi	25.09	Sugary	0.41	6.44	0.61	5.11
		Diet	0.24	6.57	0.80	4.91
Dr Pepper	2.38	Sugary	0.00	0.66	0.12	0.49
		Diet	0.00	0.84	0.06	0.21
Mountain Dew	6.68	Sugary	0.04	2.90	0.46	1.16
		Diet	0.02	1.44	0.19	0.47
Sierra Mist	4.13	Sugary	0.05	0.97	0.17	0.80
		Diet	0.02	1.41	0.13	0.59
Sprite	5.05	Sugary	0.38	1.45	0.39	1.80
		Diet	0.07	0.51	0.01	0.45
Private Label	5.61	Sugary	0.18	1.74	0.16	2.18
		Diet	0.01	0.51	0.00	0.83
Other	23.11	Sugary	1.63	5.13	1.43	5.91
		Diet	0.68	3.79	0.90	3.63
Total	100.00		5.22	50.46	6.71	37.61

Notes: Table shows market shares in percent for products purchased. Products are differentiated by brand, packaging, and sugar content.

charged and potential promotional activities. This latter data allows to construct the choice set that the household faced during each shopping trip.

For the analysis, I only consider households who buy at least one soft drink product per year in the period from 2003 to 2004.[37] This leaves me with 3,062 households that make 64,222 soft drink purchases in 2013 and 59,761 purchases in 2014.[38] This dataset is merged with information on weekly shopping trips, during which no soft drink was purchased, adding a total of 159,484 store trips with no purchase. Table C.1 shows some demographic statistics about the households in the sample.

Table 3.1 shows the products considered and their respective market shares. A product is defined as the combination of brand, packaging type, and whether it is sugary or diet. While products in the dataset are observed on the UPC level, I follow the convention in the literature to aggregate UPCs to the brand/packaging level to construct tractable choice sets (e.g. Dubois et al., 2019; Gordon et al., 2013).[39] I consider the regular and diet products of eight major brands according to market share (Coca Cola, Pepsi, Dr. Pepper, Mountain Dew, Sierra Mist, Sprite) as well as those categorized as "Private Label". Together, these products have a cumulated market share of 76.9 percent and the remaining products constitute a composite brand named "Other". I further differentiate the products into cans or bottles and whether the packaging size is small (less than 12 cans and bottles of less than 2 liters) or large (at least 12 cans or bottles of 2 liters).

Using the store data, I can observe the marketing variables of all products in consumers' choice sets.

[37] Due to the computational burden the structural model is estimated using the data from 2003. The descriptive analysis uses the data from both years.

[38] If households buy more than one soft drink product per week, I keep one purchase decision per week at random.

[39] The store data provides price information for all products that were purchased during a given week on the UPC level. I calculate the price index for each brand/packaging combination by aggregating the share-weighted UPC prices per ounce in a given store and week. The share-weight of a UPC is calculated by dividing in each week the volume sold of that UPC by the total volume sold of the brand/packaging it belongs to.

First, the dataset provides information on prices of sold products in a given week. Table C.2 shows the average prices of each product. To make prices comparable, I denote prices in cents per ounce. The table shows that there is heterogeneity in prices across brands and across packaging types. Moreover, the dataset contains a variable that indicates whether products are on a temporary price reduction. The indicator equals one if the price reduction is at least 5 percent below its usual price. I use this variable as a proxy for stockpiling as discussed below.

3.3 Descriptive Evidence for Positive and Negative State Dependence

This section provides descriptive evidence for state dependence in soft drink demand. Two sources of state dependence in demand have been frequently identified: On the one hand, when addiction or habit formation matter, consumption in the current period increases the marginal utility of consumption in the following period (Becker and Murphy, 1988). On the other hand, when stockpiling matters, stockpiling purchases in the current period decrease the need to purchase in the next period (Hendel and Nevo, 2006). Thus, habit formation results in positive and stockpiling in negative state dependence of demand. In order to disentangle positive and negative state dependence, I follow the approach that Tuchman (2019) employs in her analysis of cigarette purchases.

I regress soft drink purchases in period t on lagged purchase variables which stand in as a proxy for habit formation and stockpiling:

$$X_{it} = \mu + \beta_1 \tilde{x}_{it-1} + \beta_2 x_{it-1} + \beta_3 \sum_{k=1}^{4} x_{it-k} + \beta_4 \mathrm{pr}_{it-1} \tilde{x}_{it-1} + \alpha_i + \alpha_t + \epsilon_{it}, \tag{3.1}$$

where the dependent variable X_{it} is either the incidence of a soft drink purchase, $\tilde{x}_{it}$, or the purchase quantity, x_{it}. α_i and α_t capture household and week fixed effects, $\tilde{x}_{it-1}$ is a binary variable indicating whether the consumer bought soft drinks at all in the previous period, x_{it-1} is the amount of soft drinks bought in the previous period and $\sum_{k=1}^{4} x_{it-k}$ is a stock variable that represents the total amount of soft drinks purchased in the previous four weeks. pr_{it-1} indicates whether the majority of purchased soft drinks in the previous period were on sale.

If habit formation matters, $\tilde{x}_{it-1}$ captures positive dependence between purchases of the previous period and the current period. However, if consumers stockpile, buying *more* in the previous period (x_{it-1}) decreases the probability that they buy in the current period. Conditional on the last period, soft drink purchases of the four preceding weeks (X_{it}) are arguably not so much affected by stockpiling and, thus, mainly capture the influence of habit formation. Finally, if a product was purchased on sale (pr_{it-1}), the likelihood is larger that it was stockpiled.

Table 3.2 shows the regression results. In the first column, I use purchase incidence as dependent variable. The probability to purchase soft drinks in the current week is larger if a consumer purchased in the previous week, but having purchased *more* soft drinks in the week before is associated with a lower propensity to purchase. Thus, demand shows positive persistence in general, but proxies for

Table 3.2: Reduced form evidence for demand persistence and stockpiling

	(1) Incidence	(2) Quantity	(3) Incidence	(4) Quantity
Soda purchase incidence previous week	0.019***	0.023***	0.022***	0.038***
	(0.003)	(0.008)	(0.003)	(0.009)
Soda purchase quantity previous week	-0.019***	-0.052***	-0.019***	-0.050***
	(0.001)	(0.004)	(0.001)	(0.004)
Soda purchase quantity previous 4 weeks	0.002***	0.014***	0.002***	0.014***
	(0.000)	(0.002)	(0.000)	(0.002)
Previous soda purchase was on sale			-0.009**	-0.033***
			(0.003)	(0.009)
Constant	0.438***	0.813***	0.438***	0.813***
	(0.002)	(0.007)	(0.002)	(0.007)
n	282,902	282,902	282,902	282,902
HH FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes

Notes: Table shows estimation results of regressing purchase incidence and quantity, respectively, on proxies for habit formation (incidence previous week, quantity previous 4 weeks) and stockpiling (quantity previous week, last purchase on sale). The estimations are performed on data from the years 2003 and 2004. Clustered standard errors on the household level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

stockpiling suggest a reduced need to purchase. Moreover, the stock of purchases in the previous four weeks is associated with a higher propensity to purchase. These estimates support the notion that both habit formation and stockpiling are relevant. The results are supported when using purchase quantity as dependent variable in the second column. Again, consistent with addiction and stockpiling, purchase incidence in the preceding week and purchase quantity in the preceding four weeks are associated with a larger quantity of purchased. In contrast, the larger the purchase quantity in the previous week, the smaller the quantity in the current week. All coefficients have the same sign as the ones Tuchman (2019) obtains for cigarette purchases.

In the third and fourth columns, I additionally control whether the household purchased a soft drink that was on sale during the previous purchase occasion. The third column shows that the proxies for positive and negative state dependence are similar as before and that a consumer is indeed less likely to make a soft drink purchase if she bought a soft drink during a sale the period before. The fourth column gives a similar picture when considering the total quantity purchased in a given week. While the coefficient for positive state dependence becomes more pronounced, having purchased during a price reduction the period before is associated with a lower quantity in the current period.

These results lead me to conclude that a model of soft drink demand should incorporate both positive and negative state dependence. Hence, in the next section, a model is introduced that incorporates both habit formation and stockpiling.

3.4 Model

Conditional on a store trip in week t, a consumer i, who belongs to income group $h(i)$, maximizes her utility from buying product j, which belongs to the sugary/diet segment $s(j)$:

$$
\begin{aligned}
U_{ijt} =& \theta_i x_j + \alpha_{h(i)} p_{jt} + Z_i \beta + \xi_t + \kappa_{s(j)} \text{purch}_{is(j),t-1} + \zeta \text{pr}_{ijt} \\
&+ \delta_{s(j)} \text{purch}_{is(j),t-1} \text{pr}_{is(j),t-1} + \gamma_{s(j)} \text{purch}_{is(j),0} + \epsilon_{ijt}.
\end{aligned} \tag{3.2}
$$

θ_i is a vector of consumer-specific, time-constant preferences for product characteristics x_j. The latter consist of a general preference for soft drinks, as well as brand, diet and packaging (large vs. small and can vs. bottle) fixed effects. For identification, the brand fixed effect of "Private Label" is normalized to zero. p_{jt} is the price of product j in week t and price sensitivity $\alpha_{h(i)}$ is allowed to vary by income group $h(i)$. Z_i contains income and household size as consumer-specific demographics. Thus, the general propensity to buy a soft drink depends on income and household size. ξ_t are quarterly fixed effects. $\kappa_{s(j)}$ captures the influence of state dependence which is allowed to differ between sugary and diet soft drinks. $\text{purch}_{is(j),t-1}$ indicates if a consumer purchased a sugary or diet soft drink in the previous period. pr_{ijt} indicates if a product is on sale and ζ measures the impact of the sale on purchase probability over and above the pure price effect. Moreover, the lagged purchase indicator is interacted with $\text{pr}_{is(j),t-1}$, which indicates whether the previous purchase of a sugary or diet product took place during a sale.[40] Hence, $\delta_{s(j)}$ is supposed to capture stockpiling. Finally, $\text{purch}_{is(j),0}$ indicates for sugary and diet products if a consumer purchased a product of the respective segment in the first observed period. By controlling for the first observed decision and exogenous covariates, Z_i, I aim to deal with the initial conditions problem (Wooldridge, 2005). I assume that the idiosynratic, unobserved shocks to utility, ϵ_{ijt}, are type 1 extreme value distributed. The outside option to not make a soft drink purchase is normalized to $U_{ijt} = \epsilon_{ijt}$.

I estimate the model using nested logit, as shown in Figure 3.1. Conditional on a store trip, a consumer first decides whether to purchase a soft drink or not, and, in case of a purchase, which product to buy. Thereby, I allow unobserved shocks to be correlated in case of a soft drink purchase but not with the outside option. This relaxes the independence of irrelevant alternatives (IIA) assumption.[41] In such a model, a price increase of a soft drink leads to a proportionate shift in the probabilities to buy soft drinks but not necessarily in the probability to choose the outside option. To allow more realistic substitution patterns between products, I allow for random coefficients in the product specific preferences θ_i. This allows products that are closer in characteristics space to be closer substitutes to each other. The unobserved heterogeneity is estimated non-parametrically using discrete mass points as described below.

Likelihood. Following Train (2003), I decompose U_{ijt} into utility components that are constant over products and the outside option, V_{it}, and components that depend on variables that vary between products, W_{ijt}. The individual likelihood contribution of consumer i over all weeks and choice alternatives can be written as

[40] The dummy variable "price reduction" is given in the IRI data and indicates a temporary price reduction of at least 5 percent. If a consumer bought more than one product, the dummy equals one if at least half the quantity purchased was on sale. Note that $pr_{is(j),t-1}$ refers to the previous calendar week as stockpiling should not be relevant for longer periods.

[41] The IIA assumption implies that, e.g., a price change of one particular soft drink leads to a proportional change in the purchase probability of all other soft drinks and the outside option.

Figure 3.1: Nesting structure of empirical model

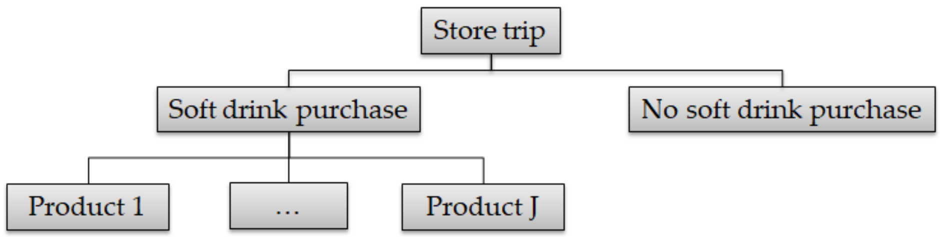

$$L_i = \prod_t \prod_j \left(\frac{exp(V_{it} + \lambda I_{it})}{1 + exp(V_{it} + \lambda I_{it})} \right)^{T_{it}} \left(\frac{1}{1 + exp(V_{it} + \lambda I_{it})} \right)^{1-T_{it}} \left(\frac{exp(W_{ijt}/\lambda)}{\sum_k exp(W_{ikt}\lambda)} \right)^{Y_{ijt}}, \tag{3.3}$$

where T_{it} indicates whether individual i makes a soft drink purchase in a week with a store visit t and Y_{ijt} indicates whether individual i buys product j in week t. $I_{it} = \ln(\sum_k exp(W_{ikt}/\lambda))$ is the inclusive value of making a soft drink purchase and links the purchase incidence decision and the product choice. It represents the expected utility that an individual receives from choosing a product. Its coefficient λ approximates the dissimilarity of alternatives within a nest (Train, 2003). If $\lambda = 1$, the unobserved utilities of alternatives are completely independent and the choice probabilities become logit. If λ lies within the unit interval, the model is consistent with random utility maximization (McFadden, 1978).
Consumer Heterogeneity. In order to identify true from spurious state dependence, it is crucial to allow for a rich and flexible structure of heterogeneity (Heckman, 1981; Keane, 1997). Here, I allow for random coefficients in the persistent preferences for each brand, as well as tastes for diet soft drinks (in contrast to sugary soft drinks), large packaging sizes, and cans (in contrast to bottles). The random coefficients are estimated non-parametrically with M discrete mass points (Heckman and Singer, 1984). The random coefficient vector is constant for each mass point. Hence, the mass points can be interpreted as different consumer types. The share of consumer type m is denoted by π^m:

$$\pi^m = \frac{exp(q^m)}{1 + \sum_{m'=2}^{M} exp(q^{m'})}, \quad \sum_{m=1}^{M} \pi^m = 1 \tag{3.4}$$

where the parameters q^m have to be estimated. For identification, q^1 is set to zero. The number of mass points is determined by successively increasing the number of types and evaluating the Bayesian Information Criterion (BIC) until there is no further increase.
The log-likelihood over all consumers thus becomes

$$LogL = \sum_{i=1}^{I} \log(\sum_{m=1}^{M} \pi_i^m L_i). \tag{3.5}$$

Empirical Bayes. To be able to perform simulations, I have to allocate values to the the random intercepts and to the random coefficients. Therefore, I assign each consumer to one of the discrete types m by applying Bayes rule (Bucklin and Gupta, 1992; Skrondal and Rabe-Hesketh, 2004). The average type probability in the sample π^m can be seen as the *prior* probability of belonging to a type. The posterior probability p_i^m is calculated using the likelihood of the purchase history conditional on each type:

$$p_i^m = \frac{\pi^m L_i^m}{\sum_{n=1}^{M} \pi^n L_i^n}. \tag{3.6}$$

where L_i^m is the likelihood of consumer i's observed purchase history given that she belongs to type m. I assign consumers the vector of random coefficients that belongs to the type m with the highest posterior, p_i^m.

Identification. Since the ultimate goal of the model is to perform tax simulations, a key parameter to be identified is the price sensitivity. I exploit price variation over time and across retailers (following Dubois et al., 2019). First, there is variation in the choice sets that consumers face since they make purchases at different retailers over time. This variation stems from different cost structures, varying product availability and occasional price reductions that I assume to be random. As I do not model the store choice, I have to assume that consumers do not select the store based on the prices of soft drinks. The second source of price variation is non-linear pricing within brands and across package sizes. As can be seen in Table C.2, cans are typically more expensive than plastic bottles and retailers offer discounts for larger sizes. The extent of non-linear pricing differs between retailers and over time. The identifying assumption is that, conditional on product-specific controls, the remaining price variation is exogenous. This assumption implies that demand shocks are not correlated with price changes. To make this assumption more plausible, I include (quarterly) time fixed effects in the decision to make a purchase. Thereby, I control for product-invariant, aggregate demand shocks (e.g. christmas) that are anticipated by the retailer in their pricing decision.

Moreover, the model aims to separately identify habit formation and stockpiling as sources for positive and negative state dependence. Therefore, I assume that consumers predominantly stockpile during sales periods. This is reasonable since they could otherwise wait until the next period allowing them, first, to not incur the storage costs, and, second, to hope for a price reduction in the next period.[42] Under this assumption, I can interact the lagged purchase indicator and the lagged price reduction indicator to take out the stockpiling component of state dependence. The remaining state dependence is interpreted as the impact of habit formation. This interpretation requires to assume that, conditional on the rich set of control variables, the lagged purchase coefficient is not correlated with any remaining unobserved preference heterogeneity. Under this assumption, habit formation can be identified from variation in choice sets and prices, which leads individuals to deviate from their purchasing pattern with longer lasting effects.

[42] This argument abstracts from transaction costs associated with making a shopping trip. However, in the dataset I observe most households making at least one shopping trip per week. If households make a shopping trip anyway, the additional costs of buying a soft drink are small.

3.5 Results

3.5.1 Parameter Estimates

Table 3.3 presents the estimated coefficients of the models with homogenous and heterogenous preferences. In the model with homogenous preferences, I do not account for unobserved consumer heterogeneity and assume all individual-specific coefficients to be fixed. In the model with heterogenous preferences, I introduce random coefficients for the taste parameters with three mass points, as described in Section 3.4. The first column of the heterogenous tastes model shows the baseline coefficients and the second and third column the interaction terms for the respective types. The average probability to belong to each type is 42.6, 38.4 and 19.1 percent, respectively. When introducing random coefficients, the Bayesian Information Criterion (BIC) substantially increases, hence, the model with heterogenous preferences is the preferred model.[43]

Panel A of Table 3.3 shows the variables that relate to product choice. In both models, the price coefficients are negative and the richest households are slightly less price sensitive. The indicator for habit formation (purch_{t-1}) is positive and significant in both specifications. However, the magnitude of state dependence is smaller in the model with heterogenous preferences. This suggests that the coefficient in the homogenous model picks up persistent taste heterogeneity. Such "spurious state dependence" (Heckman, 1981) reiterates that it is important to account for unobserved heterogeneity. Moreover, the coefficient indicating a purchase on sale in the previous period is negative, consistent with the idea that an individual who stockpiled in the previous period is less likely to purchase in this period. Regarding product preferences, there is substantial heterogeneity in brand and diet/sugary preferences when introducing consumer heterogeneity.

Panel B of Table 3.3 displays the variables that describe the decision to make a soft drink purchase. The coefficient of the inclusive value is close to one in the homogenous tastes model and smaller than one in the model with heterogenous tastes. This suggests that the model with heterogenous tastes is consistent with random utility maximization. The probability to make a purchase increases with the household size. Moreover, the constant, which measures a general preference for soft drinks, exhibits considerable heterogeneity between types.

3.5.2 Model Fit

In order to assess the in-sample fit of the estimated model, I simulate purchase decisions given the estimation sample. First, I assign values to the random intercepts and random coefficients using the "Empirical Bayes" approach described in Section 3.4. The resulting distribution of types closely resembles the average type probabilities in Table 3.3: The frequency of assigned types is 42.5 percent for type 1, 38.5 percent for type 2, and 19.0 percent for type 3. The average posterior probability of belonging to the assigned type is 95.9 percent.[44] Next, I take the first purchase decision as given

[43] I choose the model with three mass points since the BIC is larger compared to a model with two mass points, whereas a model with a fourth mass point could not be identified given the data at hand.

[44] Alternatively, the posterior probabilities could also be used as weights for the assignment of the estimated random coefficients (e.g. Haan, 2010). However, since the posterior is close to one for most consumers, the results would only change very marginally.

Table 3.3: Parameter estimates

	Homogenous preferences	Heterogenous preferences		
		Baseline	Type 2	Type 3
Panel A: Product choice				
Price (HH inc=1)	-0.376	-0.294		
	(0.010)	(0.009)		
Price (HH inc=2)	-0.387	-0.306		
	(0.008)	(0.008)		
Price (HH inc=3)	-0.354	-0.285		
	(0.008)	(0.008)		
Purch_{t-1} (sugary)	0.962	0.894		
	(0.022)	(0.021)		
Purch_{t-1} (diet)	1.202	0.696		
	(0.025)	(0.025)		
Price reduction (pr)	0.179	0.141		
	(0.011)	0.010		
Purch_{t-1} x pr_{t-1} (sugary)	-0.092	-0.094		
	(0.024)	(0.023)		
Purch_{t-1} x pr_{t-1} (diet)	-0.142	-0.182		
	(0.028)	(0.027)		
Diet Drinks	-0.180	-0.903	1.139	1.201
	(0.014)	(0.034)	(0.040)	(0.045)
Can Package	0.517	0.077	1.539	-0.853
	(0.013)	(0.025)	(0.048)	(0.045)
Large Size	0.412	-0.002	0.901	0.209
	(0.020)	(0.025)	(0.047)	(0.049)
Coca Cola	1.776	0.378	2.296	1.392
	(0.036)	(0.037)	(0.080)	(0.063)
Pepsi	0.026	-0.042	0.740	-0.080
	(0.029)	(0.036)	(0.079)	(0.068)
Dr Pepper	1.681	0.481	1.910	1.362
	(0.034)	(0.038)	(0.079)	(0.063)
Mountain Dew	0.378	-0.514	2.266	-0.223
	(0.028)	(0.049)	(0.090)	(0.091)
Sierra Mist	-0.166	-0.543	1.507	0.160
	(0.031)	(0.045)	(0.083)	(0.076)
Sprite	-0.616	-1.294	2.157	0.075
	(0.038)	(0.072)	(0.104)	(0.144)
Other	1.363	0.912	1.012	0.249
	(0.030)	(0.034)	(0.065)	(0.054)
Panel B: Purchase incidence				
Inclusive value	1.041	0.869		
	(0.016)	(0.018)		
HH size	0.177	0.182		
	(0.005)	(0.006)		
HH inc=2	0.050	0.030		
	(0.041)	(0.037)		
HH inc=3	-0.104	-0.079		
	(0.043)	(0.039)		
Constant	-5.010	-3.618	-3.035	0.196
	(0.080)	(0.066)	(0.112)	(0.065)
Type Share	1.000	0.426	0.384	0.191
BIC	-499,813	-472,696		

Notes: Table shows the estimated coefficients for a model with homogenous preferences (i.e. only fixed coefficients) and a model with heterogenous preferences (i.e. random coefficients with three discrete mass points). All models also control for quarter fixed effects and the observed purchase decision in the first period to account for the initial conditions problem. The model is estimated on the data from 2003. Standard errors are in parentheses.

Table 3.4: Sample fit of main model

	Observed (%)	Estimated (%)
Purchase incidence	46.88	46.02
Sugary products	55.24	59.98
Can packaging	55.68	54.55
Large sizes	88.07	87.77
Market shares (%)		
Coca Cola	25.55	27.96
Pepsi	5.77	5.05
Dr Pepper	23.35	25.09
Mountain Dew	6.34	6.68
Sierra Mist	4.09	4.13
Sprite	2.14	2.38
Private Label	6.97	5.61
Other	25.80	23.11

Notes: Table shows actually observed frequencies ("Observed") and simulated probabilities from the model with heterogenous preferences ("Estimated"). In the simulations, the first decision of each consumer is taken as given and the subsequent decisions are forward simulated conditional on the observed state space and the choice probabilities of the previous period. Choice probabilities are averaged over consumers.

and forward simulate the purchase decisions for all following periods. That means, in each period, I simulate the purchase decision conditional on the observed choice set (including product availability, prices, and price reductions) and the simulated choice probabilities in the previous period (that enter as lagged variables).

Table 3.4 shows the sample fit of the model. It can be seen that both the probability to make a purchase and the probability to buy a respective product are close to the empirically observed frequencies. Similarly, the estimated market shares of brands are close to the empirically observed market shares. This indicates that the model fits the data well.

3.5.3 Short- and Long-run Elasticities

To illustrate the impact of state dependence, I use the estimates from the model to simulate short- and long-run price elasticities of soft drink demand. In order to compute elasticities, I simulate the purchase decision a consumer makes with and without a 1 percent increase in prices. In the first simulated period, the price change affects the purchase probability only via the instantaneous utility. Hence, the arc elasticity in the first period is interpreted as *short-run elasticity*.[45] In the following periods, I forward simulate the purchase decision, taking into account the changed purchase probability in the preceding periods.[46] Since the model includes a lagged dependent variable, it can be described by a first-order Markov process, which is known to converge in the long run. I call the price elasticity that constitutes the limiting value after a number of periods *long-run elasticity*.

Table 3.5 shows the change in purchase probabilities when the prices of sugary soft drinks are increased by 1 percent. Since state dependence in this model is specific to sugary and diet beverages, there

[45] A different notion of short-run elasticity is the impact of a 1 percent price change that does not affect the lagged purchase probability in the next period (instead the actually observed purchase decisions enter the state space as lagged variable). This approach yields a very similar short-run elasticity.

[46] Since the reaction to a price increase partly depends on the market environment in the period of the price increase (e.g. due to the prices in the respective period and the availability of price reductions), I run ten simulations, in which I introduce the price increase at different points in time. The elasticities starting from the period of the price increase are averaged over simulations.

Table 3.5: Price elasticities over time (sugary soft drinks)

	Sugary soft drinks			All soft drinks		
Period	Price Elasticity	Difference to Period 1		Price Elasticity	Difference to Period 1	
1	0.87			0.33		
2	1.02	0.15	(0.03)	0.38	0.05	(0.01)
3	1.05	0.18	(0.04)	0.39	0.06	(0.02)
4	1.07	0.20	(0.04)	0.39	0.06	(0.02)
5	1.08	0.21	(0.04)	0.39	0.06	(0.02)
6	1.07	0.20	(0.04)	0.39	0.06	(0.02)
7	1.07	0.20	(0.04)	0.39	0.06	(0.03)
8	1.06	0.19	(0.04)	0.38	0.05	(0.03)
9	1.07	0.20	(0.04)	0.39	0.06	(0.02)
10	1.07	0.20	(0.04)	0.39	0.06	(0.02)

Notes: Table shows the respective elasticity over time starting with the period of the price increase. The elasticity is computed as the difference between simulated purchase probabilities for baseline prices and for a 1 percent increase in sugary soft drink prices. The simulations are performed for price increases at ten different points in time and elasticities are averaged over these simulations. Choice probabilities are averaged over consumers. Standard errors of the difference compared to the first period are bootstrapped with 200 replications and shown in parentheses.

are two ways in which state dependence can lead to differing short- and long-run elasticities. First, individuals can get used to not consuming soft drinks altogether, and, second, individuals can get used to consuming diet instead of sugary soft drinks. Table 3.5 shows that purchases of sugary soft drinks decrease by 0.87 percent in the short run and after ten periods the drop in purchases increases to 1.07 percent. Hence, the long-run elasticity is more than 20 percent larger than the short-run elasticity. The bootstrapped standard errors show that the difference between short- and long-run elasticity is already significant after the first period.[47] Table 3.5 also shows that a 1 percent price increase of sugary soft drinks leads to a drop in overall soft drink purchases by 0.33 percent, an effect that grows to 0.39 percent over time.

3.5.4 Soft Drink Tax Simulations

Using the estimates from the model, I simulate the impact of different taxes on demand. The considered taxes resemble taxes that are actually implemented in jurisdictions around the world. First, I consider an excise tax imposed on all sugary beverages as implemented in, among others, Mexico and in many cities in the United States. For example, Berkeley implemented a 1 cent per ounce tax in 2015 and most studies find that less than half of the tax was passed through to prices (Cawley et al., 2019b). Hence, I simulate the effect of a 0.5 cent per ounce price increase. Considering the mean price of all products weighted by their market shares, this amounts to an average price increase of 22 percent in my sample. Second, I study the impact of a 22 percent *ad valorem* tax on sugary beverages. The tax rate is chosen to raise the price of the mean product by the same magnitude as the excise tax. *Ad valorem* taxes are implemented, for example, in Thailand and Chile (see e.g. Nakamura et al., 2018). Third, I simulate the impact of a 0.5 cent per ounce excise tax on all soft drinks irrespective of whether they are sugary or diet. Soft drink taxes that do not distinguish between sugary and diet are implemented, for instance, in France and Philadelphia, Pennsylvania (Capacci et al., 2019; Cawley

[47] The bootstrap is performed by repeatedly drawing new parameter vectors from the multivariate normal distribution that is characterized by estimated parameters as distribution mean and by the estimated covariance matrix of parameters.

Table 3.6: Tax simulations, change in purchase probabilities (in percent)

	Excise tax (sugary, 0.5 ct/oz)		Ad valorem tax (sugary, 22%)		Excise tax (all drinks, 0.5 ct/oz)	
	Short-run	Long-run	Short-run	Long-run	Short-run	Long-run
Panel A: Sugary and diet						
Sugary soft drink	−20.13	−24.39	−17.84	−21.99	−14.56	−16.90
	(0.64)	(0.66)	(0.61)	(0.63)	(0.72)	(0.75)
Diet soft drink	7.68	10.76	6.79	9.65	−12.69	−13.34
	(0.90)	(0.95)	(0.88)	(0.93)	(0.86)	(0.89)
Any soft drink	−7.71	−8.74	−6.84	−7.90	−13.72	−15.31
	(0.49)	(0.50)	(0.47)	(0.48)	(0.57)	(0.60)
Panel B: Packaging and brand						
Sugary, large size	−20.05	−24.46	−16.16	−19.93	−14.41	−16.72
	(0.65)	(0.67)	(0.63)	(0.65)	(0.67)	(0.71)
Sugary, small size	−20.03	−24.14	−29.30	−32.39	−15.67	−18.33
	(1.31)	(1.30)	(1.16)	(1.14)	(1.36)	(1.35)
Sugary, private label	−19.89	−23.96	−8.31	−12.87	−16.13	−18.92
	(1.57)	(1.52)	(1.76)	(1.70)	(1.62)	(1.59)

Notes: Table shows the change in simulated purchase probabilities after imposing the respective tax compared to the baseline simulation. In the simulations, the first decision of each consumer is taken as given and the subsequent decisions are forward simulated conditional on the observed state space and the choice probabilities of the previous period. Choice probabilities are averaged over consumers. Long-run price elasticities are measured ten weeks after the tax introduction. The tax is implemented in ten different weeks and elasticities are averaged over these ten weeks. Bootstrapped standard errors with 200 replications are in parentheses.

et al., 2019a).

Table 3.6 shows the tax simulation results. While the excise tax on sugary beverages leads to an immediate drop in purchases of 20.1 percent, the effect increases up to 24.4 percent after ten weeks. Hence, the long-run effect is approximately 20 percent larger than the short-run effect. The impact of the tax can be decomposed into substituting to diet soft drinks and stopping to purchase soft drinks altogether. While there is an immediate increase in diet purchases by 7.7 percent, the effect grows substantially over time up to 10.8 percent. The probability to purchase any soft drink decreases by 7.7 percent and increases slightly up to 8.7 percent in the long run.

In Panel B of Table 3.6, I analyze if consumers substitute – within the sugary soft drink category – to cheaper soft drinks. As can be seen in Table C.2, larger package sizes and private label products are, for example, cheaper on average. However, we observe that the volumetric excise tax induces a uniform reduction in purchases across packaging types and brands.

The third and fourth column of Table 3.6 show the effect of the *ad valorem* tax on purchase probability. Although the *ad valorem* tax increases prices of the average product by the same extent as the excise tax, it is less effective in decreasing demand. The purchase probability decreases only by 17.8 percent in the short run and 22.0 percent in the long run. In contrast to the excise tax, there is less substitution to diet beverages and consumers are less likely to stop consuming soft drinks altogether. Panel B illustrates the reason for this pattern. Since the *ad valorem* tax leads to smaller price changes of cheaper products, consumers substitute to products that offer more value for money. While they strongly reduce their purchases of small packages (which are relatively expensive), they reduce their purchases of large sizes less strongly. This substitution decreases the effectiveness of the tax further

Table 3.7: Heterogenous changes in purchase probabilities of sugary soft drinks (in percent)

	Excise tax (sugary, 0.5 ct/oz)		Ad valorem tax (sugary, 22%)		Excise tax (all drinks, 0.5 ct/oz)	
	Short-run	Long-run	Short-run	Long-run	Short-run	Long-run
Panel A: Income						
Low Income	−19.96	−24.23	−17.63	−21.14	−14.81	−17.12
	(1.01)	(1.01)	(0.97)	(0.96)	(1.06)	(1.09)
Medium Income	−20.69	−25.16	−18.18	−21.87	−15.10	−17.51
	(0.74)	(0.76)	(0.70)	(0.71)	(0.78)	(0.80)
High Income	−19.14	−23.45	−17.08	−20.72	−13.64	−15.90
	(0.88)	(0.89)	(0.86)	(0.87)	(0.91)	(0.94)
Panel B: Household size						
Household size ≤ 2	−18.41	−23.20	−16.29	−20.28	−12.71	−15.15
	(0.75)	(0.77)	(0.72)	(0.74)	(0.77)	(0.81)
Household size > 2	−21.48	−25.49	−18.97	−22.29	−16.19	−18.42
	(0.69)	(0.69)	(0.66)	(0.66)	(0.72)	(0.75)

Notes: Table shows the change in simulated purchase probabilities of sugary soft drinks after imposing the respective tax compared to the baseline simulation. In the simulations, the first decision of each consumer is taken as given and the subsequent decisions are forward simulated conditional on the observed state space and the choice probabilities of the previous period. Choice probabilities are averaged over consumers. Long-run price elasticities are measured ten weeks after the tax introduction. The tax is implemented in ten different weeks and elasticities are averaged over these ten weeks. Bootstrapped standard errors with 200 replications are in parentheses.

as larger sizes contain more sugar.[48] Moreover, consumers substitute to the cheaper private label products.

Finally, in the fifth and sixth column, I simulate the impact of an excise tax of 0.5 cents per ounce on all soft drinks, irrespective of whether they are sugary or diet. As expected, the tax leads to a reduction in the purchase probability of both sugary and diet soft drinks. While the probability to buy a sugary soft drink decreases by 14.6 percent (16.9 percent) in the short run (long run), the probability to purchase a diet soft drink decreases by 12.7 percent (13.3 percent) in the short run (long run). Among all taxes, the reduction in the probability to buy a sugary soft drink is the smallest for the excise tax on all soft drinks. This is explained by the observation that the tax on all soft drinks discourages consumers from substituting to diet soft drinks, making this tax the least effective in reducing sugar consumption.

Table 3.7 addresses the question if taxes have heterogenous effects on different consumer groups. Since the discrete choice framework allows for observed and unobserved heterogeneity in consumer preferences, I can analyze heterogenous effects for different demographic groups.[49] I focus on subgroups that have received particular attention in the policy debate: poor consumers and households with children.

First, I focus on the regressivity of the tax by analyzing differential tax responsiveness by income. Panel A of Table 3.7 shows that the reduction in purchase probability is relatively uniform across the income distribution. There is some indication that the richest consumers reduce their purchases the

[48] While the small sizes of cans and bottles in the dataset contain on average 59.1 and 25.2 ounces, respectively, the large sizes contain 195.6 and 78.9 ounces. Thus, since the amount of sugar per ounce does not differ, large sizes contain substantially more sugar.

[49] On the one hand, the price coefficient in consumer's utility is allowed to differ by income group. On the other hand, I assign consumers to unobserved tastes based on their purchasing patterns. If consumers' preferences differ by demographic group, unobserved types will be distributed accordingly.

least but the differences are small and insignificant. These findings support the result from previous studies that price elasticities of poor consumers are at least as high as those of rich consumers (Dubois et al., 2019; Wang, 2015). Thus, the regressivity of soft drink taxes is alleviated when taking into account that poor individuals have larger health improvements from improving their diet (Allcott et al., 2019a).

Second, I differentiate the tax responsiveness into households that have more than two household members or not. I take that as a proxy for households with or without children (assuming that in most households with more than two members, there are children living in the household). Soft drink consumption by children is seen as particularly problematic as many consumption habits are formed early in life (Mennella et al., 2016). Hence, soft drink taxes can be seen as well targeted if they lead households with children to reduce their propensity to consume soft drinks (Dubois et al., 2019). Panel B of Table 3.7 shows that larger households respond to the taxes slightly more than smaller households. This suggests that the tax is well targeted in reducing the sugar purchases of households with children. However, the difference in the responses between households is relatively small.

3.6 Conclusion

There is evidence from biology and economics that sugar and soft drink consumption is habit forming. Nevertheless, most demand models of soft drink consumption do not take such inter-temporal complementarities into account. This paper incorporates habit formation into the analysis of soft drink taxes. Using scanner data from the US, I find that there is strong reduced-form evidence for both habit formation and stockpiling in soft drink purchases. I estimate a discrete choice model that includes habit formation and stockpiling, the latter being proxied by purchases of price reduced products. Using the estimates from the model, I simulate the short- and long-run responses to different soft drink taxes.

The results show that long-run price elasticities are approximately 20 percent larger than short-run elasticities. In the tax simulations, I find that an excise tax on sugary beverages is more effective in reducing purchases than an *ad valorem* tax. Although both increase the price of the average product by the same percentage, an excise tax reduces purchases more since it gives less incentives to substitute to cheaper products. An excise tax on all products is the least effective since it discourages substitution to diet soft drinks. The reduction in purchases are relatively uniform across demographic groups.

These results qualify the findings in Wang (2015) who has shown that not taking into account stockpiling leads to overestimation of price elasticities. I show that not taking into account habit formation can in turn lead to underestimation of long-run elasticities. Hence, future models of soft drink demand should ideally take both positive and negative state dependence into account. Moreover, policy evaluations of soft drink taxes should consider a sufficiently long post-treatment period to capture the entire treatment effect.

In this paper, I employ a myopic model of habit formation. If consumers are instead forward-looking, they will anticipate the intertemporal complementarities of consuming soft drinks. It would be interesting to assess the implications of a dynamic model of habit formation, in which consumers take into account that instantaneous consumption affects their future utility. This questions is left for further research.

Chapter 4

Protecting the Ego: Motivated Information Selection and Updating

4.1 Introduction

In many situations individuals can select what kind of feedback they receive. Standard models predict that individuals have a strict preference for choosing the most informative feedback source because this leads to more accurate beliefs and, thus, better-informed decisions. However, if individuals have motivated beliefs, i.e. if they derive utility from the belief that they have high ability (Bénabou and Tirole, 2002; Köszegi, 2006), they may prefer feedback structures that allow them to preserve this positive self-view. Consequently, motivated information selection can facilitate overconfident beliefs.[50]
Information structures in the real world differ in the way information is transmitted. In many instances, the desire to maintain a positive self-view may affect preferences over information structures. For example, individuals choose between mentors and supervisors with different feedback styles. Some supervisors always give explicit positive and negative feedback; others may give positive feedback but withhold negative feedback. In the latter example, the implicit negative feedback is less salient, so messages can be interpreted in a self-serving way. Similarly, students select into majors with different grading policies. They may prefer majors in which it is easier to receive a high absolute grade even though compressed grading distributions render grades less informative (Sabot and Wakeman-Linn, 1991; Ahn et al., 2019).
In this paper, we report results from a lab experiment showing that individuals select information structures in order to protect their belief that they have high ability. We provide evidence for two mechanisms of ego protection. First, individuals choose information structures, in which negative feed-

[50]Von Hippel and Trivers (2011) argue that self-deception and overconfidence evolved as an interpersonal strategy to gain a strategic advantage in persuading others. This conjecture is experimentally supported by Schwardmann and Van der Weele (2019), Solda et al. (2019), and Smith et al. (2017). However, holding inaccurate beliefs about ability is also costly in many domains, e.g., it leads to suboptimal management decisions (Malmendier and Tate, 2005) or over-entry into competition (Camerer and Lovallo, 1999).

back is less salient. Whereas the salience of feedback should not matter for a pure Bayesian updater, our results suggest that individuals choose less salient negative feedback because it facilitates misinterpreting negative feedback in a self-serving way. Second, when feedback is ego-relevant, individuals prefer less informative feedback structures in general. This is to avoid the risk of having to adjust their beliefs downwards.

In the experiment, subjects are asked to form beliefs about the probability of being in the top half of the distribution of subjects. The distribution is either based on (1) performance in an IQ test, or (2) a randomly drawn number. This creates exogenous variation in a between-subject design in whether the rank in the distribution is ego-relevant or not. First, we elicit participants' prior beliefs about whether they are in the top half of the distribution. We then give subjects three consecutive (noisy) signals informing them whether they are in the top half of the distribution. After each signal, we elicit the corresponding posterior belief in an incentive compatible way.

The key feature of our design is that we vary whether subjects receive signals from information structures, which they have selected themselves (endogenous treatment), or from information structures they are exogenously allocated to (exogenous treatment). While the endogenous treatment allows us to study subjects' information preferences, the exogenous treatment enables us to investigate updating behavior absent potential selection. Information structures are presented in the form of two urns with varying compositions of positive and negative signals. Depending on whether the subject is in the top or bottom half of the distribution, signals are drawn from one or the other urn. By varying the composition of signals in the urns, we vary the informativeness, salience, and skewness of feedback. *Informativeness* describes by how much beliefs are shifted by a given signal. *Salience* refers to the way in which positive and negative signals are framed, holding informativeness constant. We vary the salience of feedback by framing it as either green/red signals with an explicit description (high salience) or grey signals without description (low salience). An information structure is *skewed* if positive signals are more or less informative than negative signals. For example, an information structure is positively skewed if a potential positive signal shifts the posterior more than a negative signal.

In the endogenous treatment, subjects make five pairwise choices between information structures that vary in informativeness, salience, and skewness. Since beliefs are incentivized, we expect subjects to select the most informative feedback structure if beliefs are not ego-relevant. In contrast, subjects who derive utility from believing that they rank high in the IQ distribution may choose an information structure that is less informative, positively skewed, and makes negative feedback less salient, even if it means forming less accurate posterior beliefs.

We find stark differences in the way individuals seek information when the rank is ego-relevant and when it is not. When the ego is at stake, subjects are more likely to choose information structures that are less informative and that make negative feedback less salient. The results do not support the idea that individuals choose information structures that are positively skewed to protect their ego. Our findings are based on the analysis of information structure choices separately, and reinforced by looking at the within-individual choice patterns. Furthermore, we find that subjects who are classified as information avoiding according to the Information Preference Scale by Ho et al. (ming) are more likely to choose an information structure that is less informative and features less salient negative feedback.[51]

Moreover, we find that the subsequent belief updating process is heavily influenced by the information structure selected. Individuals in the IQ treatment update less to negative feedback, but only when negative feedback is less salient. We find a first indication for this in the endogenous treatment, where subjects receive signals from the information structure that they select into. The results are corroborated by our exogenous treatment in which subjects are placed into information structures to eliminate potential selection issues. Subjects update asymmetrically given ego-relevant signals, but only when the framing of the signals allows them to do so. When signals are not ego-relevant, subjects update to positive and negative feedback irrespective of its salience. Therefore we can reject the hypothesis that subjects not understanding less salient signals is the explanation for asymmetric updating in the ego-relevant treatment.

Our results shed light on the conditions under which individuals with a desire to protect their ego can distort their beliefs. Bénabou and Tirole (2002) show that it can be optimal for individuals to avoid or distort feedback when believing that one has high ability has a (concave) consumption value. However, since belief distortion in their two-selves model is costly, manipulation of beliefs is only beneficial within the realms of the "reality constraints." In our experiment, we show that reducing the salience of negative feedback can be one way to relax these reality constraints, leading individuals to hold overconfident beliefs about their intelligence. In fact, our results show that subjects who receive feedback that is less informative and in which negative feedback is less salient maintain overconfident beliefs about their intelligence. In contrast, subjects who receive balanced feedback are, on average, no longer overconfident about their rank at the end of the experiment.[52]

We assess explanations other than motivated reasoning for the treatment effects. The treatment variation in the ego-relevance of the state allows us to distinguish cognitive biases – general systematic errors regarding how people search and process new information (e.g. confirmation bias) – from motivated biases – biases that are driven by a desire to hold positive views of oneself. In the discussion section, we provide evidence that the treatment differences cannot be explained by cognitive biases like confirmation- or contradiction-seeking behavior, differences in cognitive ability, or confusion about the experimental design.

Our findings on motivated information selection contribute to the burgeoning literature on the production and maintenance of self-serving beliefs about oneself. Bénabou and Tirole (2016) claim that when self-relevant beliefs are involved, people tend to process information differently depending not just on its valence, but also in terms of attention, interpretation, and memory.[53] For instance, people tend to ignore or discount negative news, while more readily incorporating good news into their (posterior) beliefs. However, the resulting experimental evidence on this mechanism – asymmetric updating – is mixed.[54] On the one hand, Eil and Rao (2011), Möbius et al. (2014), and Charness

[51] We do not find evidence for heterogenous information preferences by gender, cognitive ability, or prior belief.

[52] Balanced feedback in our setting describes an information structure that is neither positively nor negatively skewed and in which both positive and negative signals are framed explicitly.

[53] A related literature on motivated reasoning pertains to the demand for and consumption of political news. This literature shows evidence that consumers prefer like-minded news (Garz et al., 2020; Gentzkow and Shapiro, 2010). Moreover, Chopra et al. (2020) show in a series of online experiments that people's demand for political news goes beyond the desire for acquiring more informative news.

[54] Selective recall of ego-relevant feedback is documented in the experiments of Chew et al. (2020) and Zimmermann

and Dave (2017) find positive asymmetry in updating. On the other hand, some studies either find no asymmetry (Grossman and Owens, 2012; Schwardmann and Van der Weele, 2019; Barron, ming; Gotthard-Real, 2017; Buser et al., 2018) or even the opposite asymmetry (Ertac, 2011; Kuhnen, 2015; Coutts, 2019). Moreover, in a paper related to how errors in updating can be driven by motivated beliefs, Exley and Kessler (2019) find that people update their beliefs based on completely uninformative signals, but only when the signals carry positive information and the updating state is ego-relevant. Our paper qualifies the existing results by showing that asymmetric updating is only observed when the information structure enables subjects to interpret the signals in a self-serving way.

Our results also contribute to the literature on information avoidance.[55] Eil and Rao (2011) and Möbius et al. (2014) present experimental evidence that a significant proportion of subjects who have received prior noisy information regarding their relative rank in an ego-relevant task (i.e., intelligence and attractiveness) have a negative willingness to pay to have their rank fully revealed. However, our study goes beyond pure information avoidance. First, in our study, subjects "choose" the signals that they would like to receive before any feedback is given. We identify preferences about information where information is open to interpretation and subjects can still remain in denial, while in the previous experiments information fully reveals the state. Second, we not only look at preferences for information avoidance, but we also seek to learn how individuals' preferences for information structures depend on the skewness and salience of feedback. Finally, we aim to understand whether there are systematic interactions between information source selection and updating behavior. In particular, our goal is to learn if and how motivated feedback selection interacts with belief formation leading to biased updating and overconfidence.

Our results also relate to an emerging literature on how complexity in the environment influences belief updating. While Epstein and Halevy (2019) and Fryer et al. (2019) find that ambiguous signals increase deviations from Bayes rule, Enke (ming) and Jin et al. (2018) illustrate that individuals find it difficult to make inferences from the absence of signals. In contrast to this literature, we look at how the informativeness and the framing of the signals affect updating. Moreover, by varying the ego-relevance of the state in a between-subject design, we can analyze the connection between cognitive and motivated biases in updating.

Finally, our research contributes to the literature on information structure selection. Within this literature, several papers study preferences about the timing and skewness of information disclosure in settings where information structures do not have any instrumental value (Falk and Zimmermann, 2016; Zimmermann, 2014; Nielsen, 2018). For example, Masatlioglu et al. (2017) find that individuals have a preference for positively skewed information sources; i.e., information structures that resolve more uncertainty regarding the desired outcome than the undesired one. Closer to our setting, some experimental papers study preferences towards information structures in settings where information has instrumental value. Charness et al. (2018) and Montanari and Nunnari (2019) study how people

(2020). Both papers find that negative feedback on IQ test performance is more likely to be forgotten, compared to positive feedback.

[55] For a review of the information avoidance literature, see the survey of Golman et al. (2017). In the health domain, Oster et al. (2013) and Ganguly and Tasoff (2016) provide empirical evidence that people avoid medical testing. In a financial context, Karlsson et al. (2009) and Sicherman et al. (2015) show that investors check their portfolios less often when the market is falling.

seek information from biased information structures. The findings of both papers show that a significant fraction of individuals make suboptimal choices. However, unlike these papers, our goal is to understand how the sub-optimality of information acquisition is driven by ego-relevant motives.

The remainder of this paper is organized as follows. In Section 4.2, we describe our experimental design, which comprises two treatment variations: ego-relevance of the rank and endogenous/exogenous information structure allocation. In Section 4.3, we present our experimental results. First, we study how participants select their preferred information structures depending on the ego-relevance of the rank. Second, we study subsequent belief updating. In Section 4.4, we discuss our findings and, in particular, we rule out cognitive biases as an alternative explanation for our main results. Finally, in Section 4.5, we conclude.

4.2 Experimental Design

To investigate whether individuals choose information structures that protect their ego, we design an experiment that contains (1) exogenous variation in ego-relevance of beliefs; (2) choices between different information structures; and (3) elicitation of updating behavior within different information structures. In a between-subject design we vary whether subjects receive feedback about their relative rank in IQ test performance (IQ treatment) or about a random number (random treatment). Separately, we vary whether subjects receive signals from the information structure they selected into (endogenous treatment) or from an information structure they are exogenously assigned to (exogenous treatment).

Figure 4.1 is an overview of the experiment.[56] It has two parts and only one is randomly selected for payout. In Part I, subjects are paid for their performance on an IQ test. They have 10 minutes to solve 20 matrices from the Raven Advanced Progressive Matrices (APM) test. They can earn £2.00 per correct answer out of three randomly chosen matrices. Although in Part II the IQ performance is only relevant for subjects in the IQ treatment, all subjects solve the IQ quiz in Part I. This ensures that there are no systematic differences in fatigue, timing, or average earnings between treatments.

In Part II, subjects express their belief about the state of the world, which is either related to their rank in the IQ test (IQ treatment) or to the draw of a random number (random treatment). First, we explain to the subjects the matching probabilities method (Karni, 2009), so that subjects maximize their chance to win a prize of £6.00 by stating their true belief (see Appendix D.2). Then, subjects give their prior belief about either their IQ performance rank or their random number depending on the treatment they are in. Afterwards, subjects in the endogenous treatment choose the information structure (feedback mode) they would like to receive signals from. Finally, there are three rounds of feedback and we elicit posterior beliefs. One out of the four belief elicitations is randomly selected for payout.

Subjects are incentivized to give their true belief, so a payoff-maximizing subject would always choose the most informative information structure and update according to Bayes rule. However, in the IQ treatment the motive to maximize payout can conflict with the desire to protect one's own beliefs

[56] In Appendix D.7, we provide screenshots of the instructions.

Figure 4.1: Timeline of the experiment

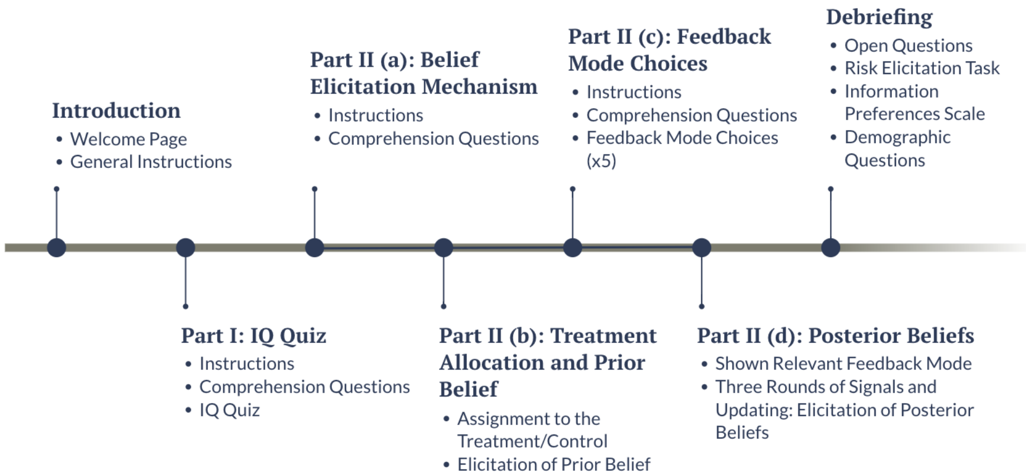

about one's (relative) ability. For instance, subjects may forego the expected payoff in order to not impair their belief that they have high intelligence. If that is the case, we expect a treatment difference in information structure selection and/or updating behavior depending on whether the beliefs are ego-relevant or not. In line with the experimental literature on motivated beliefs, we assume that further deviations from the rational benchmark are constant between the IQ and random treatments (this assumption is discussed in Section 4.4).

4.2.1 IQ and Random Treatment

We vary the ego-relevance of beliefs by randomizing subjects into an IQ treatment and a random treatment at the session level. Depending on the treatment, we tell subjects at the beginning of Part II whether we will ask for their beliefs about the IQ performance or the random number. Consistent with previous research, we argue that rank in the IQ treatment is ego-relevant (e.g. Eil and Rao, 2011). To increase the ego-relevance of the IQ treatment, we explicitly told subjects that the APM test is commonly used to measure fluid intelligence and that high scores in this test are regarded as a good predictor for academic and professional success, occupation, income, health, and longevity (Sternberg et al., 2001; Gottfredson and Deary, 2004).

4.2.1.1 IQ Treatment

In the IQ treatment, we informed subjects that the second part of the experiment was related to their relative performance on the IQ test they completed in Part I. We told them that the computer divided participants in their session into two groups: one group of subjects whose score was in the top half of the score distribution and the other with scores in the bottom half. The task was to assess whether their IQ performance was in the top or bottom half of the distribution, compared to all other participants in their session.

4.2.1.2 Random Treatment

In the random treatment, subjects were shown a randomly drawn number between 1 and 100. We told subjects that three other numbers between 1 and 100 (with replacement) had been drawn. They were not shown these numbers. Their task was to assess whether the number they saw was in the top or bottom half of the distribution among these four numbers.[57] The four numbers were randomized at the individual level. The task was deliberately designed to have variation in prior beliefs.[58] If subjects were not given the drawn number or if we compared their number against many numbers, we would expect a degenerate prior distribution.

4.2.2 Information Structure Selection

4.2.2.1 Feedback Modes

Table 4.1 shows the information structures in the experiment. Information structures consist of two urns with ten balls each. A ball drawn from an urn in the selected information structure constitutes a signal. If an individuals' IQ score or random number is in the top (bottom) half of the distribution, balls are drawn from the upper (lower) urn with replacement. Every subject receives three independently drawn signals from the urn.

Depending on the information structure, subjects can receive up to three different types of (noisy) signals. Figure D.6 displays how the signals are introduced in the instructions. Subjects in the IQ (random) treatment can either receive a green signal with a plus (+) sign and the description "You are in the top half" ("Your number is in the top half"), a red signal with a minus (-) sign and the description "You are in the bottom half" ("Your number is in the bottom half"), or a grey signal with the description "...". On the same page, we explain that the informational content of the respective signal depends on the feedback mode and the state (see Figure D.8 for a screenshot). For example, subjects are told that in feedback mode A, they are more likely to get the green (+) signal if they are in the top half of the distribution and that they are more likely to get the red (-) signal if they are in the bottom half of the distribution.[59] In all feedback modes, the green (+) signal increases the posterior that one is in the top half and the red (-) signal increases the posterior that one is in the bottom half. Depending on the feedback mode, the grey signal can be positive, negative, or non-informative feedback.

Information structures differ in their informativeness, skewness, and framing. The informativeness of signals in our experiment can be described, first, by their likelihood ratio (LR) and, second, by the probability of getting a non-informative signal. Both of these properties are given in the bottom panel of Table 4.1. The further away from one the likelihood ratio is, the more informative is the signal and the more it shifts the posterior belief of a Bayesian updater (e.g. the negative signal in Mode A is

[57] In both conditions, subjects were told that ties would be broken randomly.

[58] In the endogenous treatment, the standard deviation of prior beliefs in the control treatment turns out to be 25.065, compared to 19.993 in the IQ treatment. Hence, the control treatment generates similar variance compared to the IQ treatment.

[59] In the endogenous treatment, we explain these characteristics by always using one feedback mode choice as an example. We use three different examples, as illustrated in Figures D.8 to D.10, and check if the presented example matters for choices in Appendix D.4. In the exogenous treatment, we explain the signals using the feedback mode that the participant is assigned to: see the screenshots in Figure D.11 and D.12.

Table 4.1: Feedback modes

Mode A	Mode B	Mode C	Mode D	Mode E
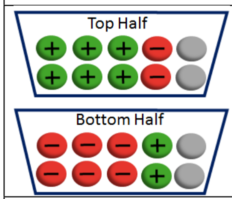	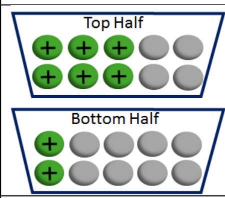	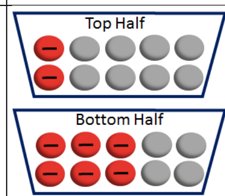	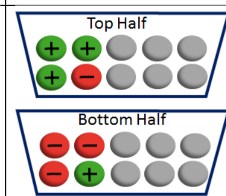	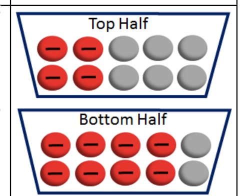
LR(Top|Green)=3	LR(Top|Green)=3	LR(Top|Grey)=2	LR(Top|Green)=3	LR(Top|Grey)=3
LR(Top|Red)=1/3	LR(Top|Grey)=1/2	LR(Top|Red)=1/3	LR(Top|Red)=1/3	LR(Top|Red)=1/2
Prob(No Info)=1/5	Prob(No Info)=0	Prob(No Info)=0	Prob(No Info)=3/5	Prob(No Info)=0

Notes: Table shows the feedback modes that can be selected in the experiment. Depending on the state (top or bottom half), a signal is drawn from the upper or lower urn. LR(State|Signal) describes the likelihood ratio of the signal concerning the state and is a measure for the informativeness of the signal. E.g., LR(Top|Green(+)) is the likelihood of receiving a green (+) signal when being in the top half divided by the likelihood of receiving a green (+) signal when being in the bottom half. Prob(No Info) describes the probability of receiving a non-informative signal.

more informative than the negative signal in Mode B).[60] The probability of receiving a non-informative signal only applies to Modes A and D, where grey signals are not informative (hence, Mode A is more informative than D).

We call an information structure positively skewed if the positive signals are more informative than the negative signals (as in Modes B and E) and negatively skewed if the negative signals are more informative (as in Mode C). An information structure is symmetric if positive and negative signals are equally informative (as in Modes A and D).

Finally, information structures differ in the salience of feedback. Since the color of the grey signal is not associated with positive or negative states and the description is not explicit, we call positive or negative feedback in the form of grey signals *less salient.* For example, in Mode B negative feedback is less salient, while in Modes C and E positive feedback is less salient. An information structure has a balanced framing if positive signals are green (+) balls, negative signals are red (-) balls, and non-informative signals are grey balls (as in Mode A and D).

4.2.2.2 Feedback Mode Choices

We let subjects make five pairwise choices (which we call "scenarios") between information structures. By carefully varying the information structures they can choose from, we are able to elicit whether subjects have preferences for informativeness, salience, or skewness of feedback depending on its ego-relevance. Every subject makes all five choices. To control for order effects, we vary the order of information structures (cf. Appendix D.4). Subjects are told that one of these five pairwise choices will be randomly selected and that the information structure chosen will be used to provide feedback.

Baseline Choice: Mode A vs Mode B In the baseline choice, we make subjects choose between two feedback modes that vary in informativeness, skewness, and salience. First, Mode A is more informative than Mode B. Second, while Mode A gives balanced positive and negative feedback depending

[60] In fact, a likelihood ratio of one implies that the signal is fully uninformative about the underlying state.

on the state, Mode B is positively skewed and negative feedback is less salient (i.e., the negative feedback is framed as grey signals). Hence, if more subjects choose Mode B in the IQ treatment compared to the random treatment, we can interpret this as evidence that subject protect their ego by choosing an information structure that gives less informative and positively skewed signals and where negative feedback is less salient. Using the remaining scenarios, we aim to disentangle the underlying preferences for informativeness, salience, and skewness.

Informativeness Choice: Mode A vs Mode D First, individuals might have a preference for less informative feedback structures if information is ego-relevant. The choice between Modes A and D isolates a preference for informativeness. In both modes, the skewness and salience of signals is held constant, only the probability of receiving a completely uninformative (grey) signal varies. Hence, subjects who have a preference for avoiding information prefer Mode D over Mode A.

Salience Choice: Mode B vs Mode E Second, individuals could have a preference for the salience of feedback, e.g., reducing the salience of negative feedback if information is ego-relevant. This could be simple aversion to explicit negative feedback, or anticipation of differential updating behavior (cf. results on updating in Section 4.3.2). To test for salience preferences, we let subjects choose between Modes B and E, which have the same informativeness and skewness but only differ in the salience of positive and negative feedback (in Mode B negative feedback is less salient and in Mode E positive feedback is less salient).

Skewness over Salience Choice: Mode A vs Mode E Third, individuals could prefer to receive positively skewed information, i.e. positive signals that are more informative than negative ones. We investigate the relative importance of preferences for positive skewness over preferences for less salient negative feedback by letting subjects choose between Mode A and Mode E. While the positive (grey) signals in Mode E have a higher likelihood ratio than the negative (red (-)) signals, positive signals in Mode E are less salient (positive signals are grey and negative signals red). Modes A and E do not just vary in skewness but also in salience and informativeness, so to get at a preference for positive skewness, in the analysis we look at this choice together with our baseline choice. Thus, if subjects have a stronger preference for positive skewness than aversion against salient negative signals, they would choose Mode E over Mode A and Mode B over A in our baseline choice.

Baseline Reversed Choice: Mode A vs Mode C Finally, we also check if individuals have a preference for or against a negatively skewed information structure with less salient positive feedback (Mode C). The signals in this information structure are equally informative about being in the bottom half of the distribution but less informative about being in the top half as compared to Mode A.

4.2.3 Updating Behavior: Endogenous and Exogenous Treatment

Besides information structure selection, we also analyze the updating behavior of subjects and how it interacts with the feedback mode. During the updating stage, subjects received three consecutive signals from one of the feedback modes. After each signal received, subjects were asked to report their

(posterior) beliefs.[61] Further, each time they received a signal and were asked about their beliefs, subjects could view a picture of the feedback mode urns from which they were receiving information by clicking a button (see Figure D.7 for a screenshot of the choice situation).
The feedback mode allocated to each subject depends on whether they are in the endogenous or the exogenous treatment.

4.2.3.1 Endogenous Treatment

In the endogenous treatment, one out of the five feedback mode choices explained above was randomly selected. The information structure that the subject selected when making their choice becomes relevant for updating. Before receiving signals, each subject was shown the choice they made in that scenario and the feedback mode from which they would be receiving the signals.

4.2.3.2 Exogenous Treatment

In the exogenous treatment, subjects are not asked to choose a feedback mode – assignments are exogenous. In particular, following the IQ test, subjects are randomly allocated to receive ego or non-ego relevant feedback from Mode A or Mode B.
The reason why we need the exogenous treatment in addition to the endogenous treatment to analyze updating behavior is twofold. First, subjects are randomly allocated into the feedback modes, so there are no systematic differences across groups, whereas in the endogenous treatment subjects self-select into feedback modes. Due to this self-selection, subjects in different feedback modes have, on average, different preferences over information structures. This could drive differences in updating behavior. Second, the exogenous treatment allows us to allocate subjects evenly into the feedback modes and across ego-relevance of the rank. Hence, in the exogenous treatment, we have more statistical power to analyze differences in updating between feedback modes.
Our main interest is in understanding deviations from Bayes rule across feedback modes and according to ego-relevance of the task. Our aim is to disentangle cognitive biases from motivated biases in updating. For this reason, we specifically focus on feedback modes A and B and their interaction with the ego-relevance of the task. On the one hand, a comparison in updating behavior across feedback modes A and B in the random treatment will allow us to understand if differences in the information structure drive cognitive biases. On the other hand, a comparison between updating across ego-relevant conditions will allow us to get at motivated biases in updating.
Importantly, as far as the updating stage is concerned, except for the random assignment into the feedback mode, there are neither differences in the experimental design nor in implementation between the endogenous and exogenous treatments. In this way, analyzing treatment differences in belief formation will permit us to also study whether and how selection affects updating.

[61] Depending on the condition they were assigned to, they were asked to report their beliefs regarding their IQ scores (IQ treatment) or their number (random treatment) being in the top half of the distribution.

4.2.4 Debriefing

In the last part of the experiment, we asked subjects a battery of questions. First, we asked subjects to complete the Information Preferences Scale by Ho et al. (ming), which is a 13-item questionnaire that measures an individual's desire to obtain or avoid information that has an instrumental value but is also unpleasant. The scale measures information preferences in three domains: consumer finance, personal characteristics, and health. Second, we asked subjects to complete the Gneezy and Potters (1997) risk elicitation task. Specifically, each subject received £1.00 and had to decide how much of this endowment to invest in a risky project with a known probability of success. The risky project returned 2.5 times the amount invested with a probability of one-half and nothing with the same probability. Third, we asked subjects to answer two questions in free-form text and they received £0.50 for their answers. We asked them to advice a hypothetical subject who would be performing the feedback mode choices and updating task. In the endogenous treatment, we additionally asked them to explain their motives for choosing the feedback modes across the five scenarios. Finally, we asked subjects a series of demographic questions including age, gender, and nationality. We also asked them a non-incentivized general willingness to take risks question (Dohmen et al., 2011).

4.2.5 Experimental Procedure

The experimental sessions were conducted from June to October 2019 in the Economics Laboratory of Warwick University, United Kingdom. Overall, we recruited 445 subjects through the Sona recruitment system to take part in the experiment. We conducted 14 sessions (216 subjects) for the endogenous treatment and 15 sessions (229 subjects) for the exogenous treatment. Sessions lasted an average of 60 minutes. Participants earned an average payment of £11.00, including the show-up fee of £5.00. We conducted the experiment using oTree (Chen et al., 2016). Descriptive statistics of the sample are provided in Table D.1.

In each session, subjects were randomly assigned a cubicle and general instructions were read aloud. The remaining instructions were provided onscreen. In both the endogenous and exogenous sessions, it was randomly determined whether the cubicle belonged to the IQ or random treatment. Moreover, in the exogenous treatment, it was randomly determined if the cubicle was allocated to feedback mode A or B.

4.3 Results

Our analysis proceeds in two steps: First, we investigate treatment differences between IQ and random in feedback mode choices. Second, we analyze how subjects update in response to signals from the corresponding feedback mode. We analyze updating in both the endogenous treatment, where subjects select into feedback modes, and in the exogenous treatment, where subjects are assigned to a feedback mode.

Figure 4.2: Share choosing feedback mode B over A in baseline choice

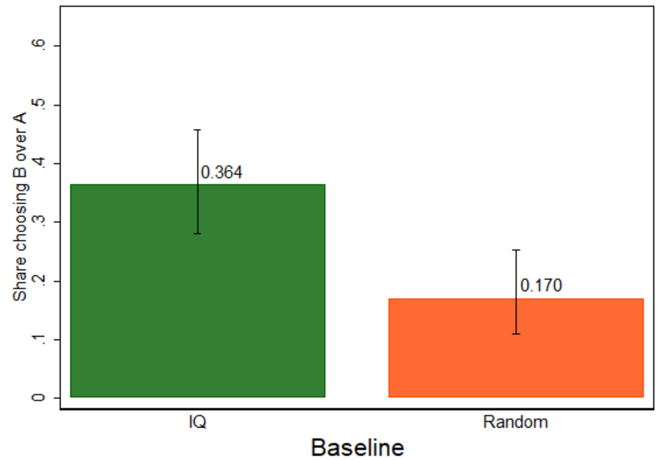

Notes: Plot shows the fraction of subjects who prefer Mode A over B in the baseline choice by treatment. In contrast to Mode A, Mode B is less informative, positively skewed and makes negative feedback less explicit. The 95 percent confidence intervals (Wilson) are shown by bar. In the IQ treatment there are N=110 subjects and in control N=106.

4.3.1 Information Selection

Information structures in our experiment differ in informativeness, salience, and skewness. The baseline choice is the choice between Mode A and Mode B, which varies in all three of these dimensions. While Mode A gives balanced feedback, Mode B produces less informative, positively skewed signals with less salient negative feedback. Hence, the choice of Mode B is costly because subjects are paid based on the accuracy of their posterior beliefs and Mode B provides less information.

Figure 4.2 illustrates the percentage of subjects who prefer to receive signals from Mode B over Mode A. While only 17.0 percent of subjects in the control treatment choose Mode B over Mode A, 36.4 percent in the IQ treatment prefer Mode B. The difference of 19.4 percentage points is statistically significant ($t(214) = 3.278, p = 0.001$).

In order to disentangle preferences for informativeness, salience, and skewness, we elicit subjects' preferences over information structures in four additional choices. Figure 4.3 plots the results. In the informativeness choice, subjects can choose between Mode A and Mode D, where both give balanced feedback but where Mode D is less informative than Mode A. The top left panel of Figure 4.3 shows that a higher proportion of subjects in IQ choose the less informative Mode D over Mode A. The difference of 13.5 percentage points is significant ($t(214) = 3.149, p = 0.002$). Hence, the results suggest that subjects in the IQ treatment have indeed a preference for less information compared to subjects in the random treatment.

In the salience choice, subjects choose between Mode B, in which negative feedback is less salient, and Mode E, in which positive feedback is less salient. We find that, in the IQ treatment, significantly more subjects exhibit a preference for less salient negative feedback, compared to the random treatment (difference of 26.5 percentage points, $t(214) = 4.116, p < 0.001$). Since the informativeness and skewness of the feedback modes are held constant, we expect subjects in the random treatment to be indifferent. Indeed, the share of 52.8 percent choosing Mode E in random is not significantly different from 50 percent ($t(105) = 0.581, p = 0.563$). In contrast, in IQ only 26.4 percent choose Mode E, which is significantly lower than the 50 percent predicted by indifference ($t(109) = 5.601, p < 0.001$). Hence, we infer that people care about the salience of signals when it concerns ego-relevant information.

In the skewness over salience choice, we give subjects the choice between Mode A and Mode E. Mode E gives – just like Mode B – positively skewed information but gives explicit negative feedback and less salient positive feedback. We find that in the IQ treatment, fewer subjects prefer Mode E over Mode A than in the random treatment. Taken together with our finding from the baseline choice, we conclude that subjects have a stronger preference against explicit negative feedback as they have a preference for positive skewness when information is ego-relevant. While in Mode B, positive feedback comes in the form of green (+) signals and negative feedback in the form of grey signals, in Mode E positive feedback is given in the form of grey signals and negative feedback as red (-) signals. This difference in framing is enough to overturn the treatment difference from the baseline choice, which suggests that the preference for positive skewness is not as strong as the preference against explicit negative signals.[62]

Finally, in the baseline reversed choice, we let subjects decide between feedback mode A and feedback mode C. Mode C gives less informative, negatively skewed signals with less salient positive feedback. In contrast to the baseline choice, we do not find a significant difference between treatments with fewer subjects in IQ choosing feedback mode C (difference of 5.2 percentage points, $t(214) = 1.041, p = 0.299$). This result suggests that subjects do not prefer less informative feedback modes in the IQ treatment if the feedback mode is negatively skewed and makes positive feedback less salient.

In Table D.2 in the appendix, we regress the five feedback mode choices on the IQ treatment dummy and different control variables. Controlling for demographics, prior beliefs, score on the IQ test and risk preferences does not alter the results in terms of treatment differences in feedback mode choices.

4.3.1.1 Information Selection - Within Individual

So far we have looked at each choice separately and analyzed what we can learn from these individual choices. We now address the within-subject choice patterns. To perform this within-analysis, we suppose that each subject has a fixed preference over information structures (conditional on the treatment) and chooses information structures according to this preference. In line with our experimental design, we focus on three preferences for information structures: maximize the informativeness, seek positive skewness, and reduce salience of negative feedback/increase salience of positive feedback. We estimate the fraction of subjects who consistently choose information structures that conform with these preferences.[63]

[62] Although the informativeness of Mode B and E is the same, we find different levels in the random treatment for the baseline choice and the skewness over salience choice. In Appendix D.4, we find evidence that the reason could be due to some subjects not understanding which feedback mode is more informative. When we explain that either Mode B or Mode E is less informative than Mode A, the levels in random are very similar.

[63] "Maximum information" predicts that subjects make choices according to $A \succ B, C, D, E$, "Positive skewness" predicts $B \succ A; E \succ A; A \succ C$, and "Salience of feedback" predicts $B \succ A; A \succ C, E; B \succ E$. Note that subjects can follow more than one preference but we assume that they have one dominant preference when these preferences conflict.

Figure 4.3: Selection of feedback mode under choice situation

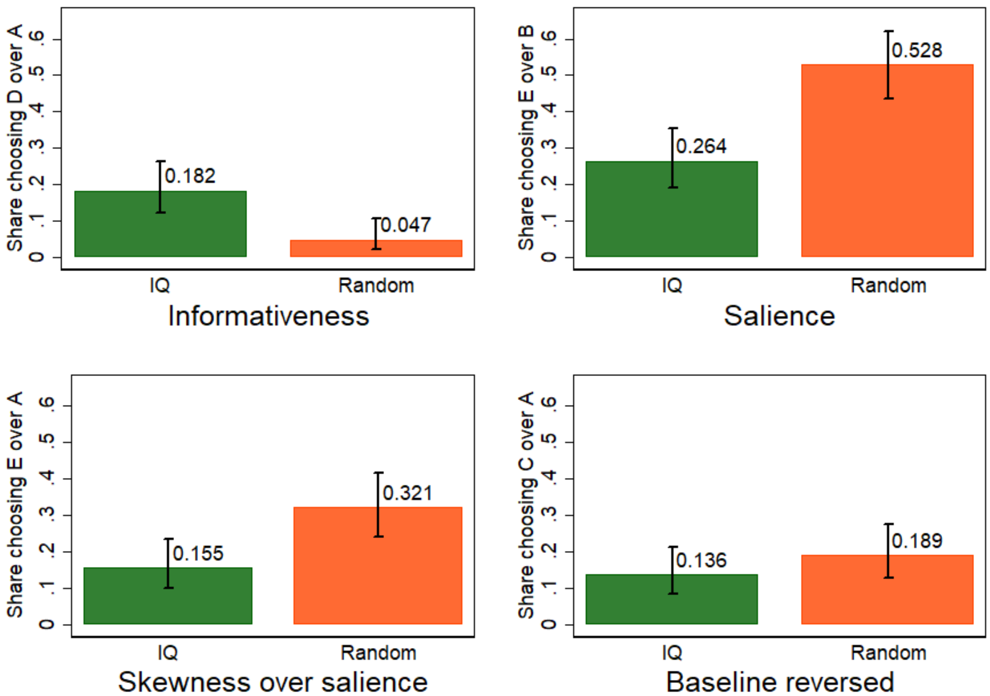

Notes: Plot shows the fraction of subjects who prefer one feedback mode over the other in the respective choice by treatment. Informativeness: Mode D is less informative than Mode A. Salience: Mode E makes negative feedback less salient than Mode B. Skewness over salience: Mode E is positively skewed and less informative than Mode A, but makes positive feedback less salient. Baseline reversed: Mode C is negatively skewed, less informative than Mode A, and makes positive feedback less salient. The 95 percent confidence intervals (Wilson) are shown by bar. In the IQ treatment there are N=110 subjects and in control N=106.

First, we calculate the fraction of subjects who make choices according to each of these preferences and the fraction of subjects who make choices that do not conform with one of these preferences. Second, we allow subjects to make mistakes and estimate which of the preferences can best explain subjects' choice patterns using a finite mixture model. Hence, we estimate the share of preference types and the amount of implementation noise (γ) necessary to classify subjects to one of the preferences. This estimation strategy is described in detail in Appendix D.3.

In Table 4.2, we compare the relative prevalence of the implied preferences by treatment. In the first two columns, we present the empirically observed fraction of subjects who adhere to a given preference when they are not allowed to make mistakes. In the third and fourth column, we show the estimated fractions using the finite mixture model. Note that 26.4 percent of subjects in the IQ and 36.8 percent in the random treatment are not classified if we do not allow for mistakes. In contrast, in the finite mixture model we use maximum likelihood to assign to every subject the preference that describes her choice pattern best.

First, consider the strategy to maximize the informativeness of information structures. There are fewer subjects who consistently maximize the informativeness in the IQ treatment than in the random treatment. When using the finite mixture model, the share increases from 40.9 to 65.5 percent in IQ and from 53.8 to 89.3 percent in random, suggesting that many subjects aim to maximize the informativeness of signals but make mistakes.[64] The treatment difference of 23.8 percent is significant ($p<0.01$).

In both treatments, there are only a few subjects who consistently choose feedback modes that are positively skewed. In particular, there are no subjects in IQ and 3.8 percent of subjects in random. In the finite mixture model, the share in random increases to 10.7 percent. However, note that it only requires three consistent choices to be attributed to this preference (in contrast to four for the other preferences). Hence, the results do not suggest that subjects choose positively skewed feedback to protect their ego.

Finally, there are significantly more subjects in IQ than in random who exhibit a preference for reducing the salience of negative feedback but not reducing the salience of positive feedback. While more than 30 percent of subjects follow such a preference in IQ, there are few to none subjects categorized as such in random. For the salience of feedback, we find the largest treatment difference, at 34.5 percentage points ($p<0.01$).

To sum up, the within-individual choices support our findings from the individual information structure choices. When looking at internally consistent choice patterns, subjects in IQ prefer less informative feedback modes and feedback modes in which positive feedback is explicit but negative feedback less salient. Moreover, we find few subjects who have a preference for positive skewness but, if anything, the share is higher in random than in IQ.

[64] In Appendix D.4, we exploit the order in which feedback modes are presented and find evidence that the difference is, to a large degree, driven by subjects who do not understand that Mode E reveals less information than Mode A.

Table 4.2: Share of subjects revealing a consistent preference by treatment

Preferences	No Mistakes		Maximum Likelihood		
	IQ	Random	IQ	Random	Difference
Maximum information	0.409	0.538	0.655^{***}	0.893^{***}	-0.238^{***}
			(0.053)	(0.042)	(0.069)
Positive skewness	0.000	0.038	0.000	0.107^{**}	-0.107^{**}
			(0.012)	(0.042)	(0.045)
Salience of feedback	0.327	0.057	0.345^{***}	0.000	0.345^{***}
			(0.056)	(0.000)	(0.056)
Implementation noise (γ)			0.492^{***}	0.585^{***}	
			(0.046)	(0.050)	
Not classified	0.264	0.368			
N	110	106	110	106	

Notes: Table shows the share of subjects who choose feedback modes consistent with the respective preference. The preference maximum information prescribes $A \succ B, C, D, E$. Positive skewness prescribes $B \succ A; A \succ C; E \succ A$. Salience of feedback means to seek explicit positive feedback but avoid explicit negative feedback and prescribes $B \succ A; A \succ C, E; B \succ E$. In "No Mistakes" we calculate the share choosing accordingly without allowing for implementation mistakes. In "Maximum Likelihood" we estimate the share allowing for implementation noise γ. Standard errors are bootstrapped with 300 replications. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.3.1.2 Heterogeneity in Information Structure Selection

We investigate heterogeneity in information structure selection based on self-reported information preferences, gender, prior beliefs, and performance on the IQ quiz. We focus on the baseline choice as it combines all three channels of ego protection: informativeness, skewness, and salience.

Information Preference Scale In the post-experimental questionnaire, subjects are asked to answer the Information Preference Scale (IPS) by Ho et al. (ming).[65] The scale consists of 13 scenarios from different domains (health, consumer finance, personal life), in which an individual can receive potentially unpleasant information (the items are shown in Appendix D.5). The respondent has to indicate her preference on a 4-point scale from "Definitely don't want to know" to "Definitely want to know."

In the first column of Table 4.3, we regress the choice of Mode B in the baseline choice on the IQ treatment indicator, an indicator if a subject scores above the median in the IPS scale, and an interaction of the two variables. The significant interaction term implies that being information seeking according to the IPS scale is associated with a lower probability of choosing a feedback mode that is less informative and in which negative feedback is less salient. Moreover, the IPS scale is not associated with information structure choice in the random treatment, as illustrated by the small and non-significant main coefficient of the IPS scale.

Gender In the second column of Table 4.3, we regress the choice of Mode B in the baseline choice on the IQ treatment indicator interacted with a dummy variable that indicates whether a subject is female. However, since the interaction term is small and far from significant, we conclude that there

[65] Ho et al. (ming) design and validate the Information Preference Scale in order to measure an individual's trait to obtain or avoid information. They show that it correlates strongly with related scales and that it even predicts information avoidance in the political domain, a domain not represented in the scale itself.

Table 4.3: Heterogeneity in baseline choice

	(1) IPS Scale	(2) Gender	(3) Prior Belief	(4) IQ Score
IQ treatment	0.323***	0.243**	0.196**	0.182**
	(0.083)	(0.107)	(0.073)	(0.091)
IPS (Info seeking)	0.064			
	(0.074)			
IPS (Info seeking) × IQ treatment	-0.259**			
	(0.117)			
Female		-0.121		
		(0.079)		
IQ treatment × Female		-0.064		
		(0.127)		
Low prior			-0.060	
			(0.073)	
IQ treatment × Low prior			-0.028	
			(0.125)	
Low IQ				0.018
				(0.074)
IQ treatment × Low IQ				0.020
				(0.120)
Constant	0.140***	0.244***	0.191***	0.159***
	(0.046)	(0.068)	(0.048)	(0.056)
R2	0.075	0.076	0.054	0.049
N	216	216	216	216

Notes: Table shows heterogenous treatment effects of the IQ treatment on the choice of mode B in the baseline choice by IPS Scale, Gender, Prior belief, and IQ score. "IPS (Info seeking)" is an indicator for subjects who score in the top half of the Information Preference Scale (Ho et al., ming). "Low prior" indicates if a subject reports a prior below 50 and "Low IQ" indicates if a subject has not scored more than the median in the IQ quiz. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 4.4: Beliefs before and after signals by feedback mode (endogenous/IQ treatment)

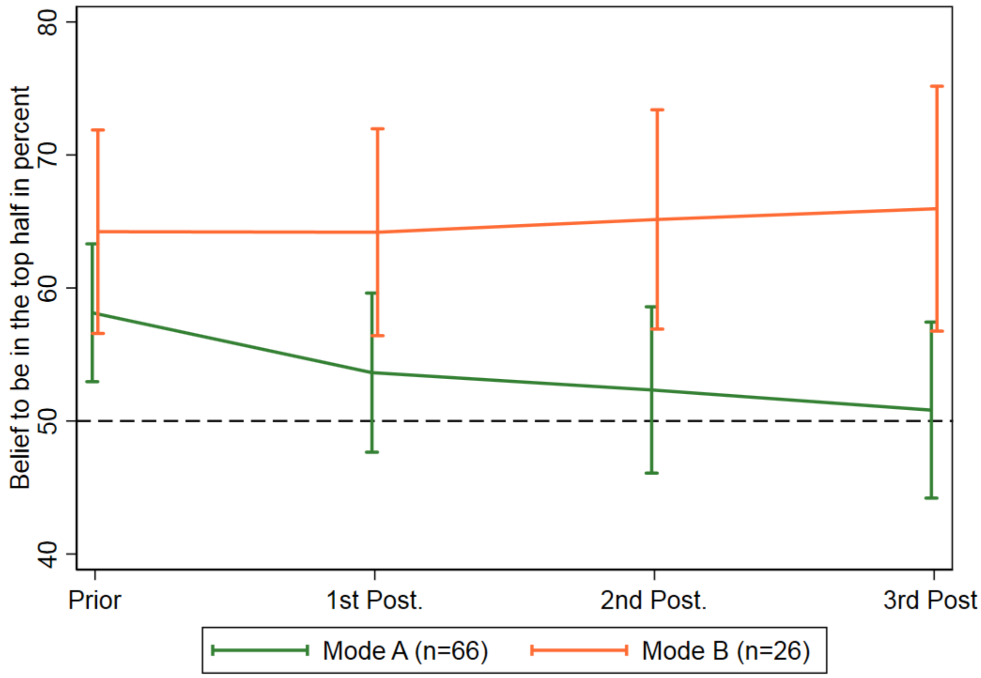

Notes: Plot shows the average prior and posterior beliefs from treatment endogenous/IQ for the selected feedback mode. The whiskers represent 95 percent confidence intervals.

is no evidence for heterogeneous treatment effects by gender in the experiment.

Prior Belief In the third column of Table 4.3, we investigate whether the treatment effect of the ego-relevant treatment is different depending on the reported prior belief. "Low Prior" indicates that an individual reports a prior that is lower than 50 percent, while the reference group reports a prior above or equal to 50 percent. We do not observe that subjects with priors below 50 percent are differently affected by the treatment compared to individuals with priors above 50 percent.

IQ Score Finally, in the last column of Table 4.3, we analyze whether there is a differential treatment effect for subjects who performed better or worse in the IQ quiz. "Low IQ" indicates that a subject has correctly solved the same number or fewer of the Raven matrices than the median (12). The non-significant and small interaction term suggests that there is no differential treatment effect depending on the performance in the IQ task (i.e., their measured cognitive ability).

4.3.1.3 Information Selection and Beliefs

After subjects made the five information structure choices, one of these choices was randomly selected and the corresponding decision of the subject was implemented.

In Figure 4.4, we plot how the average beliefs in the IQ treatment evolve, depending on the feedback mode from which subjects receive signals. First, we observe that in the IQ treatment, subjects are *overconfident* in their prior beliefs: on average, they report a likelihood higher than 50 percent of being in the top half of the IQ distribution, both in Mode A ($t(65) = 3.138, p = 0.003$) and in Mode B ($t(25) = 3.834, p = 0.001$).[66] Second, there is no significant difference in priors between individuals who end up in Mode A and Mode B ($t(90) = 1.284, p = 0.202$). Third, while the average beliefs of subjects in Mode A seem to converge toward 50 percent after receiving signals from Mode A, the beliefs in Mode B remain constant. In fact, after three rounds of feedback, we observe a significant difference in posterior beliefs between subjects in Modes A and B ($t(90) = 2.531, p = 0.013$).

These results suggest that selecting an information structure that is less informative and makes negative feedback less salient indeed leads to maintaining high beliefs in the IQ treatment. Moreover, in Appendix D.6, we show that this pattern is not observed in the control treatment. In the following section, we investigate more closely how individuals update their beliefs depending on the feedback mode from which they receive signals.

4.3.2 Belief Updating

We aim to investigate how subjects process the signals they receive from different feedback modes. First, we introduce the estimation framework for analyzing potential deviations from Bayesian updating. Then, we analyze updating in both the endogenous treatment and the exogenous treatment. While subjects in the endogenous treatment receive signals from the self-selected feedback mode, subjects in the exogenous treatment are allocated to a feedback mode.

4.3.2.1 Estimation Framework

We follow the approach developed by Grether (1980) and Möbius et al. (2014) to estimate updating behavior. The framework allows individuals to put different weights on the prior and the positive or negative signals they may receive, nesting the Bayesian benchmark as a special case. In the case of binary signals, Bayes rule can be written in the following form:

$$logit(\mu_t) = logit(\mu_{t-1}) + \mathbb{1}(s_t = pos)ln(LR_{pos}) + \mathbb{1}(s_t = neg)ln(LR_{neg}) \tag{4.1}$$

where μ_t is the belief at time t and LR_k is the likelihood ratio of the signal $s_t = k \in \{pos, neg\}$.

To estimate the model we add an error term and attach coefficients to the prior and to the positive and negative signals an individual receives:

$$logit(\mu_{it}) = \delta^{\text{prior}} logit(\mu_{i,t-1}) + \beta^{\text{pos}}\mathbb{1}(s_{it} = pos)ln(LR_{pos}) + \beta^{\text{neg}}\mathbb{1}(s_{it} = neg)ln(LR_{neg}) + \epsilon_{it} \tag{4.2}$$

[66] Benoît et al. (2015) show that true overconfidence is observed if the average stated probability to be in the top 50 percent of the distribution is significantly larger than 50 percent.

Table 4.4: Updating across feedback modes and treatments (endogenous treatment)

	(1) IQ Mode A	(2) Random Mode A	(3) IQ Mode B	(4) Random Mode B
δ^{prior}	0.776***	0.680***	0.828***	0.832***
	(0.086)	(0.075)	(0.107)	(0.163)
β^{Pos}	0.596***	0.762***	0.541**	0.408**
	(0.140)	(0.156)	(0.198)	(0.178)
β^{Neg}	0.527***	0.750***	0.191**	0.205
	(0.111)	(0.129)	(0.084)	(0.272)
p-value (δ^{Prior}=1)	0.011	0.000	0.119	0.316
p-value (β^{Pos}=1)	0.005	0.132	0.028	0.004
p-value (β^{Neg}=1)	0.001	0.057	0.000	0.009
p-value (β^{Pos}=β^{Neg})	0.658	0.952	0.114	0.372
R2	0.692	0.700	0.816	0.664
N	155	144	78	57

Notes: Table shows regression results of Equation (4.2) in the endogenous treatment, separately by IQ and random treatment and feedback mode A and B. We regress the posterior belief on the prior belief and the signal's likelihood ratio, interacted with an indicator if the signal is positive or negative. The model does not include a constant. Standard errors clustered on subject level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

where δ^{prior} captures the weight put on the prior while β^{pos} and β^{neg} measure the responsiveness to positive and negative signals, respectively. ϵ_{it} captures non-systematic errors in updating. A Bayesian updater would exhibit $\delta^{\text{prior}} = \beta^{\text{pos}} = \beta^{\text{neg}} = 1$. However, in this paper, we do not focus on the comparison with the Bayesian benchmark, instead we are mainly interested in differences in updating depending on the ego-relevance of the underlying state and across information structures. Thus, our analysis focuses on studying the β coefficients and their estimated differences across treatments and information structures.

More precisely, the estimated β coefficients in the control (i.e., where the state is not ego-relevant) will allow us to understand how updating behavior deviates from Bayes rule. Following the literature on belief updating, we interpret these deviations as being driven by "cognitive" biases, while the differential updating across the states allows us to identify "motivated" or "psychological" biases in processing information. In particular, we test whether there is asymmetric updating ($\beta^{pos} \neq \beta^{neg}$) in the ego-relevant treatment (e.g., subjects might have a desire to put more weight on positive rather than negative signals when forming their posteriors).

In doing so, we assume that cognitive biases in updating are held constant across the underlying valence of the state. Importantly, however, we do not assume that cognitive or motivated biases are constant across information structures. Indeed, the features of an information structure may have implications for both cognitive and motivated biases in updating.

4.3.2.2 Updating in the Endogenous Treatment

First, we focus on belief updating in the endogenous treatment. Here, subjects receive signals from a feedback mode, which they selected in (at least) one of the five scenarios. We restrict the analysis to feedback modes A and B because for these feedback modes we have the highest number of subjects:

Table 4.5: Updating across feedback modes and treatments (exogenous treatment)

	(1) IQ Mode A	(2) Random Mode A	(3) IQ Mode B	(4) Random Mode B
δ^{prior}	0.799***	0.726***	0.896***	0.779***
	(0.096)	(0.065)	(0.046)	(0.061)
β^{Pos}	0.631***	0.835***	0.482***	0.758***
	(0.122)	(0.124)	(0.077)	(0.174)
β^{Neg}	0.559***	0.761***	0.244***	0.938***
	(0.128)	(0.124)	(0.088)	(0.142)
p-value (δ^{Prior}=1)	0.042	0.000	0.026	0.001
p-value (β^{Pos}=1)	0.004	0.191	0.000	0.171
p-value (β^{Neg}=1)	0.001	0.059	0.000	0.662
p-value ($\beta^{\text{Pos}}=\beta^{\text{Neg}}$)	0.659	0.676	0.028	0.325
R2	0.633	0.729	0.677	0.735
N	132	139	165	186

Notes: Table shows regression results of Equation (4.2) in the exogenous treatment, separately by IQ and random treatment and feedback mode A and B. We regress the posterior belief on the prior belief and the signal's likelihood ratio, interacted with an indicator if the signal is positive or negative. The model does not include a constant. Standard errors clustered on subject level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In the IQ treatment, 66 (26) subjects update according to Mode A (B), and in the random treatment 61 (19) subjects update according to Mode A (B).

Table 4.4 shows estimation results of Equation (4.2) separately by treatment and feedback mode. We present results for Mode A in columns (1) and (2) and for Mode B in columns (3) and (4). We observe conservative updating ($\beta < 1$) in all feedback modes and treatments, which is consistent with previous evidence: subjects update less compared to the Bayesian benchmark (Möbius et al., 2014; Coutts, 2019).

In Mode A, we do not find evidence for asymmetric updating in either treatment since the coefficients β^{Pos} and β^{Neg} are of similar magnitude and we cannot reject the null hypothesis that they are equal (p-value = 0.658). In Mode B, in contrast, we observe that β^{Pos} is larger than β^{Neg}, both in IQ and random treatments. However, we do not have the statistical power to reject the null hypothesis that they are equal, since relatively few subjects end up in Mode B.

To have a sufficient number of subjects in Mode B and to exclude self-selection into feedback modes, in the next section we investigate updating in the exogenous treatment.

4.3.2.3 Updating in the Exogenous Treatment

In the exogenous treatment, subjects are randomly assigned into feedback mode A or B. In the IQ treatment, 55 (55) subjects update according to Mode A (B), and in the random treatment 57 (62) subjects update according to Mode A (B).

In Table 4.5, we display the estimation results for the exogenous treatment. As before, we do not find evidence for asymmetric updating in Mode A. Both in IQ and in random, the coefficients β^{Pos} and β^{Neg} are of similar magnitude and not significantly different from each other (p-value = 0.659 in IQ and p-value = 0.676 in random).

Table 4.6: Updating interacted with feedback modes by treatments

	Endogenous		Exogenous	
	(1)	(2)	(3)	(4)
	IQ	Random	IQ	Random
δ^{prior}	0.770***	0.671***	0.814***	0.716***
	(0.068)	(0.058)	(0.084)	(0.066)
β^{Pos}	0.594***	0.661***	0.623***	0.839***
	(0.108)	(0.126)	(0.119)	(0.127)
β^{Neg}	0.634***	0.751***	0.568***	0.767***
	(0.107)	(0.123)	(0.124)	(0.125)
δ^{prior} x Mode B	0.057	0.161	0.081	0.063
	(0.125)	(0.168)	(0.096)	(0.090)
β^{Pos} x Mode B	-0.053	-0.253	-0.142	-0.080
	(0.222)	(0.213)	(0.142)	(0.215)
β^{Neg} x Mode B	-0.443***	-0.546*	-0.324**	0.171
	(0.135)	(0.291)	(0.151)	(0.189)
R2	0.705	0.615	0.715	0.720
N	330	318	330	357

Notes: Table shows regression results of Equation (4.2), fully interacted with the feedback mode. The regressions are estimated separately by IQ and random treatment as well as by endogenous and exogenous treatment. We regress the posterior belief on the prior belief and the signal's likelihood ratio, interacted with an indicator if the signal is positive or negative as well as whether the signal comes from Mode A or B. The model does not include a constant. Standard errors clustered on subject level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

However, we find asymmetric updating in Mode B in the IQ treatment: when information is ego-relevant and negative signals are less salient (i.e., framed as grey signals), subjects update less to negative than to positive signals. The updating coefficient for positive signals is about twice as large as for negative signals and the coefficients differ significantly (p-value=0.028). However, this is not the case when feedback is not ego-relevant: in the random treatment, subjects update, if anything, more to negative feedback that is less salient but the difference in coefficients is not significant (p-value=0.735). In Table 4.6, we interact the feedback mode with the signal received, separately by treatment. In both endogenous and exogenous treatments, subjects update significantly less to negative feedback in Mode B when signals are ego-relevant (columns 1 and 3). However, when information is not ego-relevant, the interaction is only marginally significant in the endogenous treatment (column 2). The interaction is not significant and even slightly positive in the exogenous treatment (column 4), suggesting that the negative coefficient in the endogenous treatment is due to selection. Namely, a number of subjects who end up in Mode B in the endogenous/random treatment, chose Mode B in the baseline choice – in random this decision can neither be explained by payout maximization nor by motivated beliefs. Hence, this is a selected group of subjects, whose updating behavior should be interpreted with caution. Taken together these results suggest that subjects' belief formation is driven by motivated reasoning, as subjects update asymmetrically only in the IQ treatment. However, while individuals might have a preference for forming high beliefs about themselves, our results show that this does not seem to be always possible. In fact, our results show that asymmetric updating arises only in the information structure that features negative signals that are less salient and, thus, easier to misperceive. In the next section, we further discuss and interpret our experimental results.

4.4 Discussion

In our experimental findings there is a striking difference across treatments in the way subjects choose between different feedback modes. We interpret these results as evidence for differential preferences over information structures driven by motivated reasoning – biases that are driven by specific individuals' goals (e.g., of having high opinions of oneself and self-enhancement motives). We now discuss whether treatment differences could alternatively be explained by cognitive biases. In doing so, we follow the key features that distinguish motivated thinking from cognitive failures according to Bénabou and Tirole (2016).

4.4.1 Endogenous Directionality

A distinct feature of motivated reasoning is that it is directed toward some end (e.g., the belief that one has high intelligence). In contrast, general failures in cognitive reasoning that depend on one's prior beliefs, like confirmation-seeking and contradiction-seeking behavior, usually go in either direction. Here, we discuss whether these tendencies explain our results.

4.4.1.1 Confirmation-Seeking Behavior

Confirmation-seeking behavior, also known as confirmation bias, is the tendency to search for, interpret, favor, and recall information in a way that confirms one's preexisting beliefs.[67] In our experiment, this bias could explain treatment differences in information structure selection if participants with different priors have different beliefs over the informational content of our feedback modes and/or have a preference for receiving signals that confirm their priors. This could (partly) explain treatment differences as subjects in the treatment group have (slightly) higher prior beliefs than those in the control (60.1 vs. 54.2 percent).[68] For instance, it could be the case that participants with high priors believe that, in our baseline choice, Mode B is more informative than Mode A. However, even when controlling for prior beliefs, we see that subjects in the ego-relevant treatment are more likely to choose feedback mode B compared to those in the control. Similarly, as shown in Table D.2 in the appendix, our treatment differences in all feedback mode choices hold when controlling for prior beliefs. Finally, it is relevant to note that, in the informativeness choice, there is no role for confirmation-seeking behavior. Thus, this difference cannot be explained through confirmation bias. Taken together, these results show that confirmation bias falls short of explaining the treatment differences in the information selection stage of our experiment.

4.4.1.2 Contradiction-Seeking Behavior

Contrary to confirmation bias, contradiction-seeking behavior can be defined as the tendency to search for, interpret, favor, and recall information that goes against one's prior beliefs. As above, this bias could explain treatment differences in information structure selection if participants with different

[67] See Nickerson (1998) for a review of the psychological literature on confirmation bias.

[68] However, the difference is not statistically significant at conventional levels (p-value=0.101).

priors have different beliefs over the informational content of the feedback modes and/or have a preference for receiving signals that do not confirm their priors. However, by the same arguments as for confirmation-seeking behavior, we can rule out contradiction-seeking behavior as an alternative explanation for our treatment differences.

4.4.2 Bounded Rationality

4.4.2.1 Cognitive Ability

Cognitive errors in processing and interpreting information do vary by individuals' cognitive ability and analytical sophistication. That is, more able and more analytically sophisticated agents are less prone to cognitive biases. On the other hand, motivated reasoning does not necessarily imply a negative correlation.

By taking into account participants' abstract reasoning ability, measured by their scores in the IQ test, our results do not seem to be driven by cognitive ability. Two pieces of evidence support this conclusion. First, individuals across treatment groups do not vary in their cognitive ability.[69] Second, our treatment differences in feedback mode selection are robust to controlling for individuals' cognitive ability (see column 4 in Table D.2 in the appendix).

4.4.2.2 Confusion

Our experimental design features some fairly complex elements and, thus, might have affected participants' understanding. To tackle this possibility, we paid close attention to the way we presented the experimental instructions to our participants. We also ensured participants' understanding by letting them answer comprehension questions. Any participant confusion is unlikely to be a driving force of our results as it is held constant across our treatment conditions.

Furthermore, if we look at the informativeness choice in which the information structures only differ in the likelihood of receiving an uninformative signal and where we held constant their skewness and framing, we see that less than five percent of subjects in the control make the suboptimal choice. This finding is reassuring as it implies that subjects understood our experimental instructions and were sufficiently incentivized to choose the optimal decision.

4.4.3 Emotional Involvement: Heat vs. Light

Finally, Bénabou and Tirole (2016) argue that motivated beliefs evoke and trigger emotional reactions, whereas cognitively driven biases do not. While we do not measure participants' emotions in the experiment, there is suggestive evidence in favor of emotions arising in the treatment group. This supports our claim that motivated reasoning drives our results and not cognitive failures. First, in the open-text question where we asked subjects to describe how they chose between different information structures, participants in the IQ treatment were more likely to report answers that stated their willingness to avoid explicit negative (red) signals; this is not the case in the control treatment.

[69]Mean IQ scores in the IQ treatment is 11.34, while it is 11.41 in the control treatment. A Mann-Whitney test fails to reject the null that the difference in distributions is significantly different from zero (p-value=0.78).

This differential response by treatment suggests that participants in the IQ treatment chose specific feedback modes to avoid feeling negative emotions. Second, the informativeness choice result clearly shows that individuals in the IQ treatment are more inclined to protect their beliefs (and, presumably their emotions) since a Bayesian, or even boundedly rational thinker, without motivated beliefs would welcome more information.

In summary, in this section we argue that our results cannot be accounted for by cognitive biases and, in particular, by endogenous directionality and bounded rationality, including cognitive ability and confusion. On the other hand, there is evidence that participants' behavior could be driven by emotional reactions due to the ego-relevance of the underlying state.

4.5 Conclusion

We run an experiment to study individuals' preferences towards information structures and subsequent belief updating if information is ego-relevant or neutral. Our results from the information selection stage show that individuals in the ego-relevant treatment are more likely to choose feedback modes that are less informative and that make negative feedback less salient, compared to the control. These findings suggest that individuals selectively choose information structures that allow them to protect their ego. Moreover, the results from the belief updating stage indicate that individuals' belief formation is asymmetric (i.e., individuals respond more to positive news than to negative news), but only in the ego-relevant condition and when the negative feedback is less salient and therefore easier to misperceive. These results are also informative to the literature on asymmetric updating, and their mixed findings. Indeed, the different ways in which the signals are framed across studies could partly explain the differential results in the literature.

Our results suggest that while individuals might have a motivated tendency to process information differently depending on its valence, their ability to do so depends on the "reality constraints" in the environment. We show that the framing of feedback is one dimension of "reality constraints" that allows individuals to maintain and nurture motivated beliefs. Zimmermann (2020) shows that raising the incentives to recall negative feedback can constitute another "reality constraint." This raises a question for future research over which other dimensions in the environment may constrain individuals from holding motivated beliefs and how consciously people engage in motivated thinking.

Taken together, our findings suggest that motivated information acquisition might play a key role in producing overconfident beliefs. Indeed, it is often the case that in our everyday life, we can choose our information sources and exert some control over the type of signals that we receive. This choice can enable us to protect ourselves from receiving "bad" news about our abilities. Future research should pay more attention to the different ways in which individuals can and do self-select into environments to receive (avoid) more flattering (damaging) feedback to the self. Moreover, as found in our experimental data, information source selection also interacts with the subsequent belief formation process as some signals allow one to interpret the underlying information in a self-serving manner. These results have important implications. Namely, feedback procedures at the workplace that are intended to disclose unbiased information should not allow employees discretion over the information sources, to avoid biased information transmission. This could be implemented by having an objective third-party assessing workers' performance. Similarly, the anticipation of different feedback cultures (e.g., at the workplace, at the industry-level, or profession) may be an important factor that discourages people from undertaking certain career paths.

Appendix A

Appendix to Chapter 1

For copyright reasons, this appendix is not included in the online version of the dissertation. An electronic version of the appendix can be accessed at https://doi.org/10.1016/j.jpubeco.2018.10.008.

Appendix B

Appendix to Chapter 2

B.1 Derivation of optimal sin taxes

The following model closely follows O'Donoghue and Rabin (2003, 2006) and Haavio and Kotakorpi (2011) and derives the optimal tax formula in Section 2.2.
An individual i in period t has intertemporal utility from consumption that is given by

$$U_t(u_1, ..., u_T) = u_t + \beta \sum_{\tau=t+1}^{T} \delta^{\tau-t} u_\tau. \tag{B.1}$$

Each period she receives instantaneous utility u_t and future utility is discounted by time-consistent discount factor δ and by hyperbolic discounting factor β_i that differs between individuals. If $\beta_i < 1$, the agent has a preference for immediate gratification (low self-control) and if $\beta_i = 1$ the agent behaves time-consistent. For simplicity, we assume $\delta = 1$, i.e. there is no time-consistent discounting.
The instantaneous utility can be expressed as

$$u_t = v(x_t) - c(x_{t-1}) + z_t, \tag{B.2}$$

and consists of the utility $v(\cdot)$ from consuming a sin good, e.g. soft-drinks, in the current period x_t, the health costs $c(\cdot)$ with $c'(\cdot) > 0$ from having consumed soft-drinks in the past x_{t-1} and utility from a numeraire good z_t. The price of soft-drinks is p while the price of the numeraire is normalized to one. Thus, the per-period budget constraint is $px_t + z_t = y$, where y is income.
Since decisions are independent from other periods, each period the agent chooses x such as to maximize $u(x^*) = v(x^*) - \beta_i c(x^*) + z$, which yields the first order condition $v'(x^*) - \beta_i c'(x^*) = p$. However, if the agent had perfect self-control she would maximize $u(x^o) = v(x^o) - c(x^o) + z$ and consume according to the first order condition $v'(x^o) - c'(x^o) = p$. It can immediately be seen that a present-biased consumer with $\beta < 1$ overconsumes soft-drinks compared to their long-run optimal consumption x^o. Assuming that taste for soft-drinks $v(x)$ is independent of self-control β, we can expect that consumers with low self-control ($\underline{\beta}$) consume on average more soft-drinks than consumers with high self-control ($\bar{\beta}$) since they underweigh the costs.
A social planner may now decide to impose a tax t on soft-drinks in order to correct for the internality that is due to the low self-control. The social planner redistributes the tax revenues lump-sum back to consumers and the individual budget constraint becomes $(p + t)x_t + z_t = y + t\bar{x}$ where $\bar{x}$ is the average soft-drink consumption in the economy. The tax is chosen such as to maximize the social welfare function

$$\Omega(t) = \sum_i [v(x_i) - c(x_i) + (y + t\bar{x} - (p+t)x_i)], \tag{B.3}$$

which is the sum of individual long-run utility of all individuals. Solving for the first order condition yields

$$\frac{\partial \Omega(t)}{\partial t} = \sum_i [(v'(x_i) - c'(x_i) - (p+t))\frac{\partial x_i}{\partial t}] + Nt\frac{\partial \bar{x}}{\partial t} = 0, \tag{B.4}$$

where $\frac{\partial \bar{x}}{\partial t}$ is the average response in soft-drink consumption due to the tax change. Inserting the demand condition that allows for imperfect self-control $v'(x^*) - \beta c'(x^*) = p + t$ and rearranging gives (similar to Haavio and Kotakorpi, 2011):

$$t = \frac{1}{N}\sum_i (1-\beta_i)c'(x_i) + \frac{\text{cov}((1-\beta)c'(x), \frac{\partial x}{\partial t})}{\partial \bar{x}/\partial t}. \tag{B.5}$$

B.2 Factor structure of self-control scale

In order to extract the latent dimension of self-control that matters for food choices, we perform a principal component factor analysis. Following the original study by Tangney et al. (2004), we extract five factors. In Table B.1, we show the rotated factor loadings of the five factors. The first factor (13.4 percent of the variance) measures a general capacity for self-discipline and loads high on a variety of factors, e.g. "I blurt out whatever is on my mind" (0.647). The second factor (9.1 percent of the variance) is related to healthy habits and resistance against temptations. It has the highest loadings on "I eat healthy food" (0.712), "I have many healthy habits" (0.708), "I am resistant against temptations" (0.644), and "I have a hard time breaking bad habits" (0.608). The third factor (7.4 percent of the variance) is related to reliability, e.g. it has the highest loading on "I am always on time" (0.738). The fourth factor (6.6 percent of the variance) relates to self-restraint and has the highest loading on "I am self-indulgent at times" (0.620). The fifth factor (4.0 percent of the variance) describes being impulsive and loads highest on "People would describe me as impulsive" (0.552). Thus, the factor structure is very similar to that of Tangney et al. (2004).

Table B.1: Rotated factor loadings after principal component factor analysis (varimax)

	Factor1	Factor2	Factor3	Factor4	Factor5
I am good at resisting temptations	.213	.644	.109	.022	.051
(R) I have a hard time breaking bad habits	.298	.608	.004	.068	- .224
(R) I am lazy	.273	.439	.286	.135	- .299
(R) I often say inappropriate things	.551	.129	.130	.030	- .003
I never allow myself to lose control	- .150	.005	.111	-.152	.533
(R) I do certain things that are bad for me, if they are fun	.205	.231	.055	.539	.036
(R) Getting up in the morning is hard for me	.292	.173	.306	.084	- .405
(R) I have trouble saying no	.476	.234	.029	-.057	- .218
(R) I change my mind fairly often	.586	.104	.159	.008	- .154
(R) I blurt out whatever is on my mind	.647	.057	- .011	.063	.105
I refuse things that are bad for me	.114	.347	.152	-.284	.254
(R) I spend too much money	.340	.367	.177	.307	- .024
I keep everything neat	.082	.258	.512	.005	.088
(R) I am self-indulgent at times	.074	.029	- .024	.620	- .030
(R) I wish I had more self-discipline	.472	.459	.130	.054	- .142
I am reliable	.087	.058	.468	-.343	.306
(R) I get carried away by my feelings	.557	.134	-.062	.151	.043
(R) I do many things on the spur of the moment	.330	-.054	- .054	.450	.190
(R) I don't keep secrets very well	.470	-.041	.215	.045	- .040
(R) I have worked or studied all night at the last minute	.349	.097	.410	.300	-.208
I'm not easily discouraged	.258	.293	.245	-.514	.014
(R) I'd be better off if I stopped thinking before acting	.527	-.007	.128	.0367	.064
(R) Pleasure and fun sometimes keep me from getting work done	.338	.104	.314	.399	.003
(R) I have trouble concentrating	.550	.178	.229	-.076	-.253
I am able to work effectively toward long-term goals	.170	.305	.325	-.408	.122
(R) Sometimes I can't stop myself from doing something, even if I know it is wrong	.433	.316	.119	.407	.047
(R) I often act without thinking through all the alternatives	.575	.198	.106	.186	.220
(R) I lose my temper too easily	.537	.049	- .042	.001	.029
(R) I often interrupt people	.597	.062	.013	.071	- .027
I am always on time	.010	-.031	.737	-.011	-.043
People can count on me to keep the schedule	.048	.072	.719	-.014	- .042
(R) People would describe me as impulsive	.232	-.101	- .050	.307	.552
People would say that I have an iron self-discipline	.157	.397	.448	-.157	.083
I have many healthy habits	- .054	.708	.018	-.061	.021
I eat healthy foods	- .013	.712	.026	.007	- .015
(R) I sometimes drink too much alcohol	.085	.122	.139	.210	.188

Notes: Table shows factor loadings of 2,387 consumers in the sample (questionnaire from 2013 and, if missing, imputed with data from 2015). (R) indicates that the item is reverse coded.

Table B.2: Correlations of self-control factors with characteristics and attitudes

	(1) Body Mass Index (BMI)	(2) Obesity (BMI>30)	(3) Intention to reduce weight	(4) "I should eat less sugar"	(5) "I should eat less animal fat"
Low SC (Factor 1)	0.494*	0.032	0.089***	0.085***	0.057**
	(0.283)	(0.021)	(0.028)	(0.029)	(0.028)
Low SC (Factor 2)	2.124***	0.094***	0.202***	0.112***	0.115***
	(0.269)	(0.021)	(0.028)	(0.029)	(0.028)
Low SC (Factor 3)	0.453	0.026	0.021	-0.009	0.011
	(0.283)	(0.021)	(0.028)	(0.029)	(0.028)
Low SC (Factor 4)	0.727**	0.035	0.018	0.034	0.025
	(0.287)	(0.022)	(0.029)	(0.029)	(0.028)
Low SC (Factor 5)	0.175	-0.002	0.000	-0.063**	0.012
	(0.288)	(0.022)	(0.028)	(0.029)	(0.028)
Controls	Yes	Yes	Yes	Yes	Yes
Mean	26.021	0.175	0.620	0.483	0.354
Households	1237	1236	1197	1197	1197

Notes: Columns (1) and (2) are based on weight and height data from 2011. BMI is calculated as ([weight in kg]/[height in m]^2). Column (3) shows the fraction of respondents who indicate in the 2013 survey that they would like to weigh at least 1 kg less. Column (4) and (5) gives the fraction of respondents who indicate that they should eat "A lot less" or "A little less" sugar or animal fat to eat healthier. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3 Soft Drink Tax

B.3.1 Pass-through of soft drink tax to prices

Figure B.1: Average price residuals over time (Graph from Schmacker and Smed (2020))

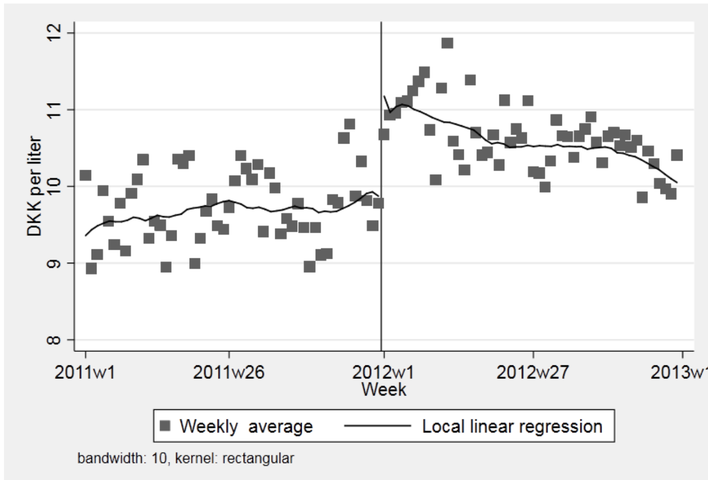

(a) Tax increase

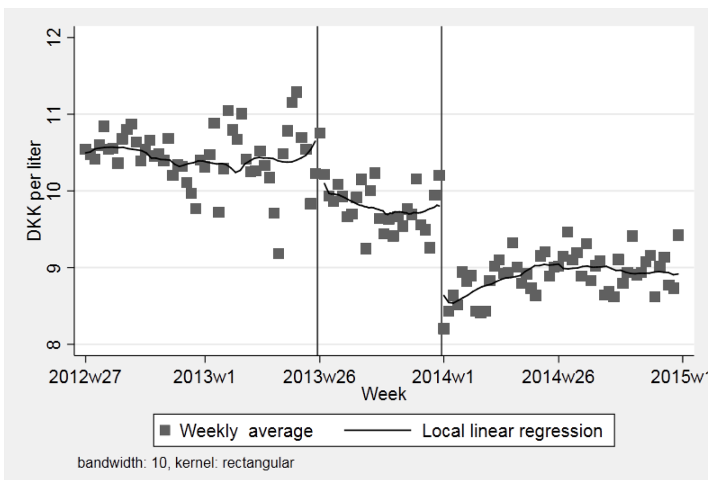

(b) Tax repeal

Notes: Graph shows soft drink prices around the tax increase in January 2012 and the tax cuts in July 2013 and January 2014. The graph plots residuals that are added to the sample mean after regressing prices on product fixed effects. Dots represent weekly averages and the lines local polynomials (triangular weights and 16 weeks bandwidth). The vertical lines indicate the timing of tax changes.

B.3.2 Robustness of soft drink tax estimations

Table B.3: Soft drink purchases in response to placebo tax changes by self-control

	(1) Quantity	(2) Quantity	(3) Quantity	(4) Quantity
Tax Placebo	-3.314	1.428	0.870	-2.683
	(6.069)	(6.006)	(6.015)	(5.774)
High self-control × Tax Placebo	-0.193	0.076	6.217	5.313
	(8.107)	(8.100)	(8.116)	(8.024)
Households	1171	1171	1260	1260
Household Months	20674	20674	21622	21622
Placebo	January 2010	January 2010	January 2011	January 2011
Controls	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. The dependent variable is monthly quantity in centiliter per household member. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.4: Soft drink purchases in response to soft drink tax changes controlling for income and education

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Panel A: Tax Hike						
Tax Hike	-2.137	-0.013	0.014	-0.713	-0.020	0.012
	(10.629)	(0.015)	(0.059)	(9.438)	(0.013)	(0.054)
Interactions with Tax hike						
High self-control	-25.081*	-0.048**	-0.137*	-25.918**	-0.030*	-0.105
	(13.058)	(0.020)	(0.077)	(12.392)	(0.017)	(0.074)
High Income	-0.600	-0.009	-0.031			
	(14.557)	(0.021)	(0.087)			
High self-control × High income	8.170	0.034	0.111			
	(18.305)	(0.028)	(0.116)			
High education				-3.465	0.009	-0.052
				(13.012)	(0.021)	(0.081)
High self-control × High education				11.929	-0.002	0.094
				(16.700)	(0.027)	(0.117)
Panel B: Tax Repeal						
Tax Repeal	30.456***	0.039**	0.110*	39.494***	0.059***	0.111**
	(11.268)	(0.015)	(0.056)	(9.641)	(0.013)	(0.046)
Interactions with Tax repeal						
High self-control	0.050	0.010	0.012	-16.417	-0.022	0.019
	(13.026)	(0.021)	(0.075)	(12.000)	(0.017)	(0.063)
High Income	2.826	0.013	0.081			
	(14.326)	(0.020)	(0.073)			
High self-control × High income	-6.955	-0.016	-0.058			
	(17.772)	(0.027)	(0.101)			
High education				-20.416	-0.036	0.116
				(14.251)	(0.022)	(0.078)
High self-control × High education				31.851*	0.059**	-0.101
				(17.923)	(0.028)	(0.105)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In columns (1) and (4) the dependent variable is monthly quantity in centiliter per household member. In columns (2) and (5) it is purchase incidence in a given month. In columns (3) and (6) it is log-transformed quantity. In Panel A, we run the estimations with 1,278 households (22,197 household months) and for the intensive margin estimations with 1,104 households (7,466 household months). In Panel B, we use 1,278 households (22,747 households months) and for the intensive margin estimations with 1,122 households (7,782 household months). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.5: Soft drink purchases in response to soft drink tax changes controlling for tastes and nutritional knowledge

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Panel A: Tax Hike						
Tax Hike	-7.406	-0.011	-0.078	-0.844	-0.015	-0.008
	(10.614)	(0.014)	(0.064)	(7.848)	(0.011)	(0.049)
Interactions with Tax hike						
High self-control	-20.837*	-0.049***	-0.039	-20.808**	-0.030**	-0.049
	(12.142)	(0.017)	(0.078)	(9.501)	(0.015)	(0.060)
Unhealthy taste	11.563	-0.004	0.140*			
	(13.407)	(0.020)	(0.081)			
High self-control × Unhealthy taste	3.628	0.046	-0.007			
	(17.467)	(0.028)	(0.106)			
Lacks knowledge				-2.934	0.008	0.027
				(14.955)	(0.024)	(0.088)
High self-control × Lacks knowledge				-11.094	-0.021	-0.129
				(21.548)	(0.037)	(0.158)
Panel B: Tax Repeal						
Tax Repeal	38.873***	0.046***	0.167***	31.241***	0.039***	0.147***
	(11.005)	(0.015)	(0.057)	(8.869)	(0.013)	(0.044)
Interactions with Tax repeal						
High self-control	-8.093	-0.001	0.008	-3.169	0.014	-0.014
	(12.895)	(0.018)	(0.070)	(10.433)	(0.015)	(0.056)
Unhealthy taste	-15.041	0.003	-0.054			
	(14.391)	(0.021)	(0.074)			
High self-control × Unhealthy taste	7.356	0.009	-0.044			
	(17.733)	(0.029)	(0.103)			
Lacks knowledge				-0.787	0.041	-0.039
				(14.789)	(0.025)	(0.092)
High self-control × Lacks knowledge				1.285	-0.071*	0.107
				(20.489)	(0.039)	(0.139)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In columns (1) and (4) the dependent variable is monthly quantity in centiliter per household member. In columns (2) and (5) it is purchase incidence in a given month. In columns (3) and (6) it is log-transformed quantity. "Unhealthy taste" identifies consumers that agree to the statement "I believe I would make healthier food choices if unhealthy food was less tasty", "Lacks knowledge" identifies consumers who agree to the statement "I believe I would make healthier food choices if I had more information on how to eat healthy". In Panel A, we run the estimations with 1,197 households (20,887 household months) and for the intensive margin estimations with 1,033 households (6,956 household months). In Panel B, we use 1,197 households (21,389 households months) and for the intensive margin estimations with 1,050 households (7,291 household months). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.6: Soft drink purchases in response to soft drink tax changes, only single households

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin
Panel A: Tax Hike			
Tax Hike	7.997	0.003	0.044
	(11.441)	(0.016)	(0.069)
High self-control × Tax Hike	-36.567**	-0.051**	-0.112
	(14.804)	(0.021)	(0.085)
Panel B: Tax Repeal			
Tax Repeal	39.102***	0.041**	0.104*
	(14.230)	(0.018)	(0.062)
High self-control × Tax Repeal	-10.256	0.003	0.027
	(15.831)	(0.023)	(0.084)
Sample	Single HH	Single HH	Single HH
Controls	Yes	Yes	Yes
Household FE	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In columns (1)the dependent variable is monthly quantity in centiliter per household member. In columns (2) it is purchase incidence in a given month. In columns (3) it is log-transformed quantity. In Panel A, we run the estimations with 467 households (7,893 household months) and for the intensive margin estimations with 391 households (2,466 household months). In Panel B, we use 467 households (8,064 households months) and for the intensive margin estimations with 394 households (2,523 household months). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.7: Soft drink purchases in response to soft drink tax by self-control (sample split by quartiles)

	(1) Overall	(2) Ext. Margin	(3) Int. Margin
Panel A: Tax Hike			
Tax Hike	6.060	-0.011	0.048
	(10.040)	(0.014)	(0.056)
Interactions with Tax Hike			
Medium low self-control	-15.812	-0.010	-0.109
	(12.898)	(0.019)	(0.077)
Medium high self-control	-21.595*	-0.027	-0.059
	(12.514)	(0.019)	(0.073)
High self-control	-36.275***	-0.043**	-0.195***
	(11.892)	(0.018)	(0.073)
Panel B: Tax Repeal			
Tax repeal	42.986***	0.048***	0.206***
	(11.261)	(0.015)	(0.053)
Interactions with Tax Repeal			
Medium low self-control	-22.349	-0.004	-0.123*
	(13.927)	(0.020)	(0.072)
Medium high self-control	-16.593	0.003	-0.113*
	(13.211)	(0.019)	(0.068)
High self-control	-13.082	-0.003	-0.028
	(13.435)	(0.019)	(0.074)
Controls	Yes	Yes	Yes
Household FE	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In Panel A, we run the estimations with 1,278 households (22,197 household months) and for the intensive margin estimations with 1,104 households (7,466 household months). In Panel B, we use 1,278 households (22,747 households months) and for the intensive margin estimations with 1,122 households (7,782 household months). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.8: Soft drink purchases in response to soft drink tax by access to German border

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin
Panel A: Tax Hike			
Tax Hike	1.422	-0.003	-0.002
	(10.445)	(0.015)	(0.065)
Interactions with Tax hike			
High self-control	-28.427**	-0.054***	-0.076
	(12.928)	(0.020)	(0.078)
No Toll	-5.651	-0.023	-0.000
	(13.239)	(0.020)	(0.078)
High self-control × No Toll	12.824	0.041	0.003
	(16.616)	(0.026)	(0.105)
Panel B: Tax Repeal			
Tax Repeal	28.108**	0.022	0.135**
	(12.482)	(0.017)	(0.056)
Interactions with Tax repeal			
High self-control	5.180	0.031	0.018
	(14.229)	(0.021)	(0.071)
No Toll	6.345	0.040*	0.028
	(14.592)	(0.021)	(0.073)
High self-control × No Toll	-15.402	-0.050*	-0.061
	(18.124)	(0.028)	(0.100)
Controls	Yes	Yes	Yes
Household FE	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In column (1) the dependent variable is monthly quantity in centiliter per household member. In column (2) it is purchase incidence in a given month. In column (3) it is log-transformed quantity. In Panel A, we run the estimations with 1,277 households (22,177 household months) and for the intensive margin estimations with 1,103 households (7,465 household months). In Panel B, we use 1,278 households (22,747 households months) and for the intensive margin estimations with 1,122 households (7,782 household months). While 730 households live in the "No Toll" region, 547 live in the "Toll" region. If a household has moved between "Toll" and "No toll", we use the region where they reported the most purchases. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.4 Fat tax

B.4.1 Pass-through of fat taxes to butter prices

In this section, we aim to show that the fat tax indeed had an effect on the price of butter. Figure B.2 illustrates the development of prices around the fat tax introduction and repeal. The graph plots residuals after controlling for product fixed effects in order to control for potentially changing purchasing patterns. It is clearly shown that during the time window when the fat tax was enacted, prices for butter were higher than before and after.

In Table B.9 we quantify the extent of the price changes by regressing absolute and log-transformed prices on a tax dummy while controlling for product fixed effects. Since we use a bandwidth of one year around the tax changes, the regression amounts to comparing the average prices one year before the tax change to one year after the tax change. We observe that prices per 100g of butter have increased by DKK 0.761 after the tax introduction and have decreased by DKK 0.611 after the tax repeal. Hence, the magnitude of price changes is indeed very similar for the tax introduction and the repeal.

Figure B.2: Average price residuals (controlling for product fixed effects) over time added to the sample mean

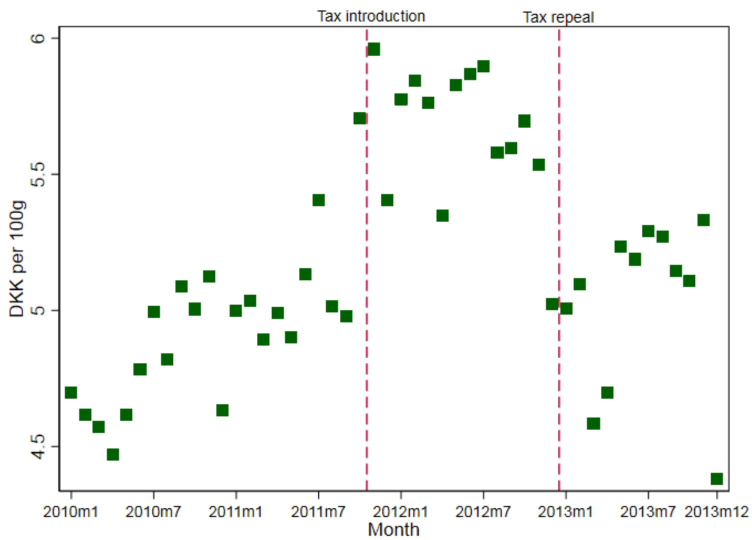

Notes: Graph shows butter prices around the tax increase in January 2012 and the tax cuts in July 2013 and January 2014. The graph plots residuals that are added to the sample mean after regressing prices on product fixed effects. Dots represent monthly averages. The vertical lines indicate the timing of tax changes.

Table B.9: Butter prices in response to tax changes

	Tax introduction		Tax repeal	
	(1)	(2)	(3)	(4)
	Absolute price	Log price	Absolute price	Log price
Tax change	0.761^{***}	0.151^{***}	-0.611^{***}	-0.124^{***}
	(0.042)	(0.010)	(0.051)	(0.009)
Constant	4.905^{***}	1.546^{***}	5.758^{***}	1.710^{***}
	(0.022)	(0.005)	(0.026)	(0.005)
EAN fixed effects	Yes	Yes	Yes	Yes
n	52198	52198	59123	59123

Standard errors clustered on EAN level in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.4.2 Robustness of fat tax responsiveness to sample window

Table B.10: Butter purchases in response to placebo tax changes by self-control

	(1) Quantity	(2) Quantity	(3) Quantity	(4) Quantity
Tax Placebo	16.743***	16.169***	-2.208	-3.170
	(5.969)	(5.908)	(6.040)	(6.031)
High self-control × Tax Placebo	-0.011	-0.255	1.846	2.813
	(8.250)	(8.202)	(8.000)	(7.951)
Households	1284	1284	1217	1217
Household Months	26139	26139	25355	25355
Placebo	January 2010	January 2010	October 2010	October 2010
Controls	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. The dependent variable is monthly quantity in gram per household member. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure B.3: Predicted values of monthly butter purchase quantity by self-control

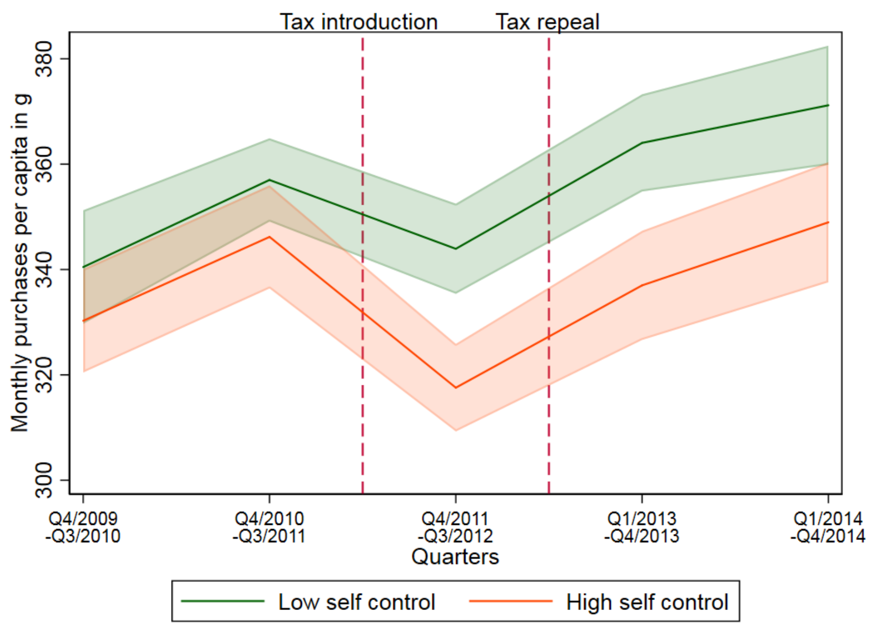

Notes: Graph shows predicted values after controlling for household fixed effects and the control variables specified in the methods section. Household quantities are individualized by dividing the observed household quantity by the number of household members (children aged 0 to 6 enter as 0.5 household members). The shaded areas represent 95 percent confidence intervals. The vertical lines indicate the timing of tax changes.

B.4.3 Robustness of fat tax estimations

Table B.11: Butter purchases in response to fat tax changes controlling for education

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin
Panel A: Tax introduction			
Tax introduction	-7.851	-0.015	-0.020
	(7.924)	(0.010)	(0.018)
Interactions with Tax introduction			
High self-control	-24.008**	-0.033**	-0.017
	(11.058)	(0.014)	(0.025)
High education	-9.217	-0.009	-0.014
	(12.376)	(0.016)	(0.029)
High self-control × High education	19.480	0.025	0.025
	(17.276)	(0.023)	(0.040)
Panel B: Tax Repeal			
Tax Repeal	15.202*	0.025**	0.031
	(8.571)	(0.010)	(0.020)
Interactions with Tax repeal			
High self-control	-8.504	-0.018	-0.009
	(11.790)	(0.014)	(0.028)
High education	-13.290	-0.022	-0.011
	(13.680)	(0.016)	(0.031)
High self-control × High education	14.863	0.047**	-0.007
	(18.408)	(0.022)	(0.044)
Controls	Yes	Yes	Yes
Household FE	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In column (1) the dependent variable is monthly quantity in gram per household member. In column (2) it is purchase incidence in a given month. In column (3) it is log-transformed quantity. In Panel A, we run the estimations with 1,324 households (27,192 household months) and for the intensive margin estimations with 1,291 households (17,056 household months). In Panel B, we use 1,323 households (27,507 households months) and for the intensive margin estimations with 1,298 households (17,460 household months). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.12: Butter purchases in response to fat tax changes using the subsample of low and high income households

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Panel A: Tax Introduction						
Tax Introduction	-12.755	-0.016	-0.032	8.020	-0.008	0.004
	(9.801)	(0.012)	(0.021)	(8.614)	(0.012)	(0.020)
High self-control × Tax	-5.068	-0.009	0.009	-28.953**	-0.034**	-0.022
	(13.957)	(0.016)	(0.030)	(11.707)	(0.016)	(0.027)
Households	739	739	706	815	815	794
n	12591	12591	7778	14601	14601	9278
Panel B: Tax Repeal						
Tax Repeal	18.219	0.041***	0.017	4.668	0.006	0.028
	(12.771)	(0.015)	(0.031)	(8.946)	(0.012)	(0.021)
High self-control × No Tax	-13.004	-0.019	-0.036	8.018	0.013	0.011
	(13.786)	(0.016)	(0.034)	(12.148)	(0.015)	(0.030)
Households	730	730	712	782	782	765
n	13087	13087	8217	14420	14420	9243
Income	Low	Low	Low	High	High	High
Controls	No	No	No	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. The estimations are performed separately on the sample of low and high income households. Low income indicates that the household is below the median of equivalized annual income (i.e. DKK 265,000) and high income that the household is above the median. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.13: Butter purchases in response to fat tax changes controlling for tastes and nutritional knowledge

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Panel A: Tax Hike						
Tax Hike	-2.005	-0.007	-0.003	-9.469	-0.022**	-0.004
	(8.517)	(0.011)	(0.019)	(6.784)	(0.009)	(0.015)
Interactions with Tax hike						
High self-control	-22.875**	-0.030**	-0.034	-17.139*	-0.017	-0.032
	(11.498)	(0.015)	(0.026)	(9.370)	(0.012)	(0.021)
Unhealthy taste	-22.545*	-0.021	-0.054*			
	(11.948)	(0.016)	(0.028)			
High self-control × Unhealthy taste	17.582	0.016	0.060			
	(17.461)	(0.023)	(0.041)			
Lacks knowledge				-18.580	0.022	-0.129***
				(14.680)	(0.021)	(0.037)
High self-control × Lacks knowledge				19.031	-0.017	0.143**
				(24.507)	(0.030)	(0.057)
Panel B: Tax Repeal						
Tax Repeal	7.767	0.026**	0.008	12.167*	0.022**	0.028
	(9.355)	(0.011)	(0.023)	(7.331)	(0.009)	(0.018)
Interactions with Tax repeal						
High self-control	2.074	-0.010	0.021	-3.269	-0.004	-0.009
	(11.819)	(0.014)	(0.028)	(9.543)	(0.011)	(0.023)
Unhealthy taste	6.160	-0.014	0.038			
	(12.378)	(0.015)	(0.029)			
High self-control × Unhealthy taste	-15.572	0.019	-0.096**			
	(17.218)	(0.022)	(0.043)			
Lacks knowledge				-6.742	-0.014	-0.011
				(15.963)	(0.020)	(0.038)
High self-control × Lacks knowledge				-13.550	0.014	-0.076
				(23.328)	(0.030)	(0.056)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table shows OLS regression results with standard errors clustered on household level. In columns (1) and (4) the dependent variable is monthly quantity in gram per household member. In columns (2) and (5) it is purchase incidence in a given month. In columns (3) and (6) it is log-transformed quantity. “Unhealthy taste” identifies consumers that agree to the statement “I believe I would make healthier food choices if uhealthy food was less tasty”, “Lacks knowledge” identifies consumers who agree to the statement “I believe I would make healthier food choices if I had more information on how to eat healthy”. In Panel A, we run the estimations with 1,241 households (25,602 household months) and for the intensive margin estimations with 1,211 households (16,019 household months). In Panel B, we use 1,241 households (25,914 households months) and for the intensive margin estimations with 1,220 households (16,418 household months).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5 Complete estimation tables

Table B.14: Soft drink purchases in response to soft drink tax hike by self-control

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Interaction of tax and self-control (Reference: Low self-control)						
Tax Hike	0.319	-0.014	0.022	-1.912	-0.016*	-0.002
	(6.543)	(0.010)	(0.040)	(6.559)	(0.010)	(0.041)
High SC × Tax Hike	-21.663***	-0.032**	-0.086	-21.097***	-0.030**	-0.075
	(8.198)	(0.013)	(0.054)	(8.134)	(0.013)	(0.052)
Equivalized HH income in DKK (Reference: <175K)						
175-250 DKK				16.556	0.028	0.023
				(15.061)	(0.021)	(0.078)
250-325 DKK				13.345	0.007	0.101
				(15.550)	(0.024)	(0.093)
325-400 DKK				13.420	0.001	0.053
				(17.779)	(0.029)	(0.116)
>500 DKK				12.521	-0.018	0.106
				(22.425)	(0.035)	(0.178)
Labor market status (Reference: Not employed)						
Full time				10.433	0.014	0.080
				(13.467)	(0.026)	(0.090)
Part time				31.778**	0.046**	0.152**
				(14.277)	(0.022)	(0.076)
Quarter Dummies (Reference: 1st Quarter)						
2nd Quarter				33.907***	0.049***	0.129**
				(7.880)	(0.014)	(0.053)
3rd Quarter				30.764***	0.043**	0.112
				(10.465)	(0.018)	(0.071)
4th Quarter				10.632*	-0.000	0.071
				(6.187)	(0.011)	(0.047)
Number of Kids (0-6)				-73.308***	-0.085	-0.276*
				(25.293)	(0.065)	(0.164)
Number of Kids (7-14)				-33.188***	-0.008	-0.284***
				(9.141)	(0.024)	(0.036)
Number of Kids (15-20)				-27.673	-0.025	-0.274
				(20.154)	(0.021)	(0.167)
Temperature				-0.771	-0.000	-0.005
				(0.717)	(0.001)	(0.005)
Constant	130.428***	0.351***	5.414***	102.136***	0.309***	5.359***
	(2.042)	(0.003)	(0.013)	(16.466)	(0.025)	(0.099)
Adj. R2	0.412	0.304	0.400	0.414	0.307	0.404
Households	1278	1278	1104	1278	1278	1104
n	22197	22197	7466	22197	22197	7466

Notes: Table shows OLS regression results with standard errors clustered on household level. $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.15: Soft drink purchases in response to soft drink tax repeal by self-control

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Interaction of tax and self-control (Reference: Low self-control)						
Tax Repeal	28.512***	0.041***	0.127***	31.827***	0.046***	0.151***
	(6.967)	(0.010)	(0.036)	(7.351)	(0.011)	(0.039)
High SC × Tax Repeal	-3.667	0.001	-0.013	-3.613	0.002	-0.015
	(8.817)	(0.014)	(0.050)	(8.814)	(0.014)	(0.050)
Equivalized HH income in DKK (Reference: <175K)						
175-250 DKK				13.744	0.041**	0.122
				(10.516)	(0.020)	(0.082)
250-325 DKK				0.248	0.022	0.101
				(14.571)	(0.029)	(0.105)
325-400 DKK				15.547	-0.007	0.204*
				(16.058)	(0.035)	(0.115)
>500 DKK				19.576	-0.010	0.126
				(25.757)	(0.045)	(0.156)
Labor market status (Reference: Not employed)						
Full time				-4.105	0.022	-0.082
				(13.473)	(0.030)	(0.110)
Part time				11.221*	0.023*	0.081
				(6.785)	(0.013)	(0.049)
Quarter Dummies (Reference: 1st Quarter)						
2nd Quarter				31.387***	0.039***	0.161***
				(7.885)	(0.013)	(0.050)
3rd Quarter				41.412***	0.061***	0.209***
				(10.485)	(0.017)	(0.071)
4th Quarter				9.923	0.006	0.093**
				(6.834)	(0.012)	(0.047)
Number of Kids (0-6)				-14.869	0.008	-0.184
				(21.006)	(0.051)	(0.162)
Number of Kids (7-14)				-28.951	0.021	-0.203
				(20.351)	(0.044)	(0.150)
Number of Kids (15-20)				-17.614	0.028	-0.175
				(18.735)	(0.041)	(0.146)
Temperature				-0.712	-0.001	-0.007
				(0.679)	(0.001)	(0.005)
Constant	119.302***	0.322***	5.370***	96.388***	0.266***	5.274***
	(2.179)	(0.003)	(0.013)	(12.387)	(0.025)	(0.106)
Adj. R2	0.412	0.300	0.424	0.414	0.302	0.426
Households	1278	1278	1122	1278	1278	1122
n	22747	22747	7782	22747	22747	7782

Notes: Table shows OLS regression results with standard errors clustered on household level. $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.16: Butter purchases in response to fat tax introduction by self-control

	(1) Quantity	(2) Extensive Margin	(3) Intensive Margin	(4) Quantity	(5) Extensive Margin	(6) Intensive Margin
Interaction of tax and self-control (Reference: Low self-control)						
Tax Introduction	-11.926**	-0.018**	-0.023*	-11.224*	-0.018**	-0.026*
	(5.908)	(0.008)	(0.014)	(5.974)	(0.008)	(0.014)
High SC × Tax Intro.	-15.692*	-0.022**	-0.007	-16.440**	-0.022**	-0.008
	(8.416)	(0.011)	(0.019)	(8.380)	(0.011)	(0.019)
Equivalized HH income in DKK (Reference: <175K)						
175-250 DKK				25.717**	0.018	0.061*
				(12.649)	(0.015)	(0.032)
250-325 DKK				29.178*	0.029	0.042
				(15.031)	(0.019)	(0.040)
325-400 DKK				55.229***	0.030	0.092**
				(17.658)	(0.023)	(0.047)
>500 DKK				39.027**	0.031	0.051
				(18.800)	(0.026)	(0.056)
Labor market status (Reference: Not employed)						
Full time				11.149	0.003	0.009
				(12.696)	(0.018)	(0.040)
Part time				-10.567	0.008	-0.045
				(12.260)	(0.016)	(0.032)
Quarter Dummies (Reference: 1st Quarter)						
2nd Quarter				20.119***	0.026***	0.010
				(5.412)	(0.007)	(0.013)
3rd Quarter				42.608***	0.046***	0.037***
				(5.267)	(0.007)	(0.013)
4th Quarter				61.383***	0.056***	0.097***
				(6.028)	(0.007)	(0.013)
Number of Kids (0-6)				-70.836***	-0.040	-0.329***
				(18.565)	(0.043)	(0.079)
Number of Kids (7-14)				-25.842**	-0.016	-0.181***
				(10.898)	(0.022)	(0.035)
Number of Kids (15-20)				-42.327*	-0.017	-0.229***
				(23.945)	(0.014)	(0.049)
Constant	349.137***	0.642***	6.002***	300.673***	0.591***	6.010***
	(2.136)	(0.003)	(0.005)	(14.135)	(0.019)	(0.040)
Adj. R2	0.384	0.275	0.481	0.387	0.277	0.485
Households	1324	1324	1291	1324	1324	1291
n	27192	27192	17056	27192	27192	17056

Notes: Table shows OLS regression results with standard errors clustered on household level. $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.17: Butter purchases in response to fat tax repeal by self-control

	(1)	(2)	(3)	(4)	(5)	(6)
		Extensive	Intensive		Extensive	Intensive
	Quantity	Margin	Margin	Quantity	Margin	Margin
Interaction of tax and self-control (Reference: Low self-control)						
Tax Repeal	12.913**	0.018**	0.027*	10.208	0.017**	0.027*
	(5.996)	(0.007)	(0.014)	(6.562)	(0.008)	(0.016)
High SC × Tax Repeal	-0.183	0.002	-0.008	-2.291	0.002	-0.013
	(8.467)	(0.010)	(0.021)	(8.516)	(0.010)	(0.021)
Equivalized HH income in DKK (Reference: <175K)						
175-250 DKK				17.008	-0.015	0.045
				(15.686)	(0.015)	(0.037)
250-325 DKK				19.140	-0.005	0.037
				(18.439)	(0.019)	(0.046)
325-400 DKK				34.319	0.009	0.058
				(22.088)	(0.025)	(0.058)
>500 DKK				23.555	0.011	0.072
				(23.764)	(0.031)	(0.062)
Labor market status (Reference: Not employed)						
Full time				-15.677	0.002	-0.026
				(12.565)	(0.018)	(0.040)
Part time				-4.661	-0.003	0.012
				(8.514)	(0.010)	(0.020)
Quarter Dummies (Reference: 1st Quarter)						
2nd Quarter				-12.448**	0.002	-0.054***
				(5.647)	(0.007)	(0.013)
3rd Quarter				5.548	0.019***	-0.034***
				(5.477)	(0.007)	(0.013)
4th Quarter				71.547***	0.074***	0.090***
				(6.264)	(0.007)	(0.013)
Number of Kids (0-6)				-67.879***	0.025	-0.299***
				(19.330)	(0.030)	(0.089)
Number of Kids (7-14)				-78.071***	0.029	-0.288***
				(20.746)	(0.029)	(0.080)
Number of Kids (15-20)				-76.105***	0.016	-0.297***
				(19.532)	(0.029)	(0.067)
Constant	335.547***	0.625***	5.970***	329.519***	0.595***	6.024***
	(2.115)	(0.003)	(0.005)	(15.764)	(0.020)	(0.047)
Adj. R2	0.382	0.274	0.470	0.388	0.278	0.476
Households	1323	1323	1298	1323	1323	1298
n	27507	27507	17460	27507	27507	17460

Notes: Table shows OLS regression results with standard errors clustered on household level. $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix C

Appendix to Chapter 3

C.1 Additional tables

Table C.1: Descriptive statistics

	2013	2014
Avg income per capita (K)	21.600	21.590
	(14.552)	(14.577)
Avg household size	2.528	2.535
	(1.248)	(1.242)
Avg weekly volume purchased (ounces)	214.922	197.177
	(194.611)	(188.237)
Avg weekly dollars spent	4.281	4.161
	(3.859)	(3.892)
Number of Households	3062	3062

Notes: Tables shows descriptive statistics of the sample for the years 2013 and 2014. The income is calculated by assigning households the midpoint of its discrete income bracket. Standard deviations are in parentheses.

Table C.2: Average prices of soft drink products (in cents per ounce)

		Can		Bottle	
Brand	Type	$<$ 12 Cans	$\geq$ 12 Cans	$<$ 2 Liter	$\geq$ 2 Liter
Coca Cola	Sugary	4.18	2.35	5.77	1.87
		(0.69)	(0.46)	(0.50)	(0.31)
	Diet	4.36	2.39	5.68	1.96
		(0.70)	(0.46)	(0.45)	(0.32)
Pepsi	Sugary	4.36	2.31	5.46	1.80
		(0.76)	(0.45)	(0.77)	(0.29)
	Diet	4.30	2.34	5.48	1.87
		(0.78)	(0.45)	(0.73)	(0.31)
Dr Pepper	Sugary	4.52	2.50	5.44	1.89
		(0.60)	(0.43)	(0.45)	(0.27)
	Diet	3.21	2.51	5.48	1.83
		(1.51)	(0.43)	(0.41)	(0.32)
Mountain Dew	Sugary	4.20	2.31	5.64	1.85
		(0.73)	(0.45)	(0.57)	(0.27)
	Diet	3.55	2.43	5.45	1.89
		(1.39)	(0.46)	(0.45)	(0.30)
Sierra Mist	Sugary	4.09	2.35	5.54	1.78
		(0.82)	(0.53)	(0.97)	(0.34)
	Diet	3.36	2.41	5.33	1.75
		(1.19)	(0.53)	(0.86)	(0.37)
Sprite	Sugary	3.83	2.39	5.59	1.88
		(1.17)	(0.45)	(0.74)	(0.34)
	Diet	3.67	2.54	5.47	1.73
		(1.41)	(0.44)	(0.81)	(0.33)
Private Label	Sugary	2.63	1.59	1.59	1.04
		(0.63)	(0.21)	(0.68)	(0.09)
	Diet	2.15	1.59	0.77	1.11
		(0.56)	(0.21)	(0.05)	(0.12)
Other	Sugary	2.90	2.32	4.21	1.69
		(1.00)	(0.44)	(0.63)	(0.29)
	Diet	2.84	2.43	3.32	1.69
		(0.75)	(0.42)	(0.98)	(0.25)
Mean		3.67	2.29	4.82	1.73

Notes: Table shows average prices of products in cents per ounces. Products are differentiated by brand, packaging and sugar content. Standard deviations are given in parentheses.

Appendix D

Appendix to Chapter 4

D.1 Additional Tables

Table D.1: Descriptive statistics

	Endogenous		Exogenous	
	IQ	Random	IQ	Random
Age (Mean)	20.973	21.830	20.236	20.168
	(2.960)	(4.095)	(2.933)	(2.304)
Female (Share)	0.664	0.613	0.618	0.529
	(0.475)	(0.489)	(0.488)	(0.501)
Native English speakers (Share)	0.427	0.377	0.364	0.370
	(0.497)	(0.487)	(0.483)	(0.485)
Studying (Share)	0.964	0.972	0.955	0.975
	(0.188)	(0.167)	(0.209)	(0.157)
First year students (Share)	0.455	0.481	0.545	0.504
	(0.500)	(0.502)	(0.500)	(0.502)
IQ puzzles solved (Mean)	11.336	11.406	10.964	10.924
	(3.425)	(3.397)	(3.038)	(3.051)
N	110	106	110	119

Notes: Table shows descriptive statistics of the experimental dataset. Standard deviations are in parentheses.

Table D.2: Information structure choices controlling for covariates

	(1)	(2)	(3)	(4)	(5)	(6)
Baseline	0.194***	0.192***	0.178***	0.193***	0.189***	0.161***
	(0.059)	(0.060)	(0.059)	(0.059)	(0.058)	(0.061)
Informativeness	0.135***	0.136***	0.132***	0.134***	0.132***	0.132***
	(0.042)	(0.044)	(0.043)	(0.042)	(0.042)	(0.042)
Framing	-0.264***	-0.266***	-0.243***	-0.265***	-0.263***	-0.263***
	(0.064)	(0.067)	(0.065)	(0.065)	(0.065)	(0.065)
Skewness over framing	-0.166***	-0.129**	-0.164***	-0.165***	-0.167***	-0.122**
	(0.057)	(0.059)	(0.057)	(0.057)	(0.057)	(0.058)
Baseline reversed	-0.052	-0.036	-0.049	-0.053	-0.052	-0.037
	(0.050)	(0.052)	(0.050)	(0.050)	(0.051)	(0.058)
Demographics		✓				✓
Prior			✓			✓
IQ score				✓		✓
Risk					✓	✓
N	216	216	216	216	216	216

Notes: Table shows the coefficient of the IQ treatment dummy in the regression of the feedback mode choice on the respective covariates. Demographics comprises controls for gender, age, years of study, and whether English is the native language. The risk measure is by Gneezy and Potters (1997). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

D.2 Belief Elicitation Mechanism

After Part I was completed but before we explained Part II, we told subjects that for the following part of the experiment we would ask them their beliefs regarding some events. In particular, they were told that they would be asked four belief questions and that one question would be chosen at random to count for payments (see Figure D.3 for a screenshot of the instructions).

Then, we explained our belief elicitation procedure to them. We used the belief elicitation mechanism proposed by Karni (2009) called the matching probabilities method.[71] Under this method, subjects are presented with two possible bets: the lottery and the event. Each bet either pays a prize p (£6.00 in our experiment) or nothing. More specifically:

- The Event: pays the prize p if the event occurs, 0 otherwise.
- The Lottery: pays the prize p with probability x for $x \in \{0, 1, 2, ..., 100\}$, and 0 otherwise;

Hence, subjects (through their answers to the belief question) indicate what probability x makes them indifferent between betting on the event or the lottery. After they indicate the indifference point, one probability $y \in \{0, 1, 2, ..., 100\}$ is drawn. If $x \geq y$, the subject bets on the event and earns the prize p if the event occurs. On the other hand, if $x < y$ the subject bets on the lottery, which has probability y of paying the prize p. Intuitively, by choosing x, the subject affects her chances of betting on the event or the lottery and the chances of earning the prize p in case she ends up betting on the lottery. Under this mechanism, reporting one's subjective probability of the event occurring maximizes the chances of earning the prize, regardless of risk preferences.

Given the complexity of this belief elicitation mechanism, we decided to make instructions intuitive for subjects by walking them through an example and explaining how their answer would affect the chances of them betting on the event or the lottery and their chances of winning the prize. We also emphasized that truthful reporting was the answer that maximized the chances of earning the prize. To ensure that subjects understood the main features of this elicitation procedure, we asked subjects to answer comprehension questions about the belief elicitation procedure.

[71] This method is also referred to as the "crossover mechanism," "reservation probabilities," and "lottery method." This belief elicitation mechanism is first introduced in an experiment by Möbius et al. (2014); since then it is extensively applied to other experiments in the asymmetric updating literature, including Coutts (2019), Buser et al. (2018), and Schwardmann and Van der Weele (2019).

D.3 Maximum Likelihood Estimation of Within-Subject Choice Patterns

We use a finite mixture model to estimate the share of subjects who exhibit consistent choice patterns that pertain to one of three preferences ("Maximum information," "Positive skewness," or "Salience of feedback"). We allow for a deviation between the observed choice and the choice prescribed by a subject's preference: $y_{ic} = I\left\{s_{ic}(s^p) + \gamma\epsilon_{ic} \geq 0\right\}$, where y_{ic} is the choice by subject i in choice situation c (0 for the first alternative and 1 for the second alternative). s_{ic} is the choice that is prescribed by the preference s^p (coded by -1 for the first alternative and 1 for the second alternative). $I\{.\}$ is 1 if the term in brackets is positive and 0 otherwise. ϵ_{ic} is an iid error term that is type 1 extreme value distributed. γ scales the variance of the error term and can be interpreted as the amount of implementation noise to be estimated. Thus, the more an individual's choices align with the choices prescribed by the respective preference, the smaller will be the estimated implementation noise.
The likelihood that subject i follows preference s^p over all choices c is

$$\pi_i(s^p) = \prod_C \Big(\frac{1}{1 + exp(-s_{ic}(s^p))/\gamma}\Big)^{y_{ic}} \Big(\frac{1}{1 + exp(s_{ic}(s^p))/\gamma}\Big)^{1-y_{ic}}. \tag{D.1}$$

The resulting log likelihood is $\sum_I \ln(\sum_P \pi(s^p)\pi_i(s^p))$, which is summed over all I subjects by treatment and where P represents the set of preferences we consider. $\pi(s^p)$ is the estimated fraction of the sample with preference p. For the estimation we adapt the code by (Dal Bó and Fréchette, 2011).

D.4 Investigation of Order Effects

The subjects in our experiment make five consecutive choices between information structures. Since only one of these five choices is randomly selected to be implemented, each individual choice can be treated as an independent choice. However, one may be concerned about potential order effects if subjects' subsequent choices are affected by their previous choices, e.g., due to a preference for consistency. Moreover, in the experimental instructions, we always use the first feedback mode choice as an example to explain the choice situation (see Figures D.8 to D.10). Hence, it is possible that subjects have a better understanding of the feedback mode choice that is presented first. However, note that potential order effects do not affect our results if they are constant between IQ and random treatments. To check if there are differential order effects between treatments, we vary the order in which subjects make pairwise choices. The different orders with the respective number of subjects are presented in Table D.3.

Table D.3: Order of feedback mode choices

Feedback choice	Order 1	Order 2	Order 3
Baseline (A vs B)	1^{st}	2^{nd}	3^{rd}
Informativeness (A vs D)	3^{rd}	3^{rd}	4^{th}
Framing (B vs E)	5^{th}	4^{th}	2^{nd}
Skewness over framing (A vs E)	4^{th}	5^{th}	1^{st}
Baseline reversed (A vs C)	2^{nd}	1^{st}	5^{th}
N	116	51	49

Notes: Table shows the order with which the respective feedback mode choice is presented.

In Table D.4 we interact the treatment dummy with a dummy indicating if a subject makes feedback mode choices according to the first, second, or third order. Most importantly, as indicated by the insignificant interaction effects, we do not find much support for differential order effects between treatments. This suggests that order effects are of no concern for our conclusions.

Interestingly, in column (4), we observe that in both treatments significantly fewer subjects select Mode E in the skewness over framing choice when this choice represents the first scenario (i.e., in the third order). This could be explained by the fact that in Order 3 this choice is presented first and is used as an example in the instructions (cf. Figure D.10). Hence, in Order 3 subjects may better understand that Mode E is, in fact, less informative than Mode A. This is supported by the observation that the fraction of subjects who follow a strict preference to maximize the informativeness in Table 4.2, increases substantially from 0.409 to 0.500 in the IQ treatment and from 0.538 to 0.689 in the control treatment when abstracting from the skewness over framing choice.

Table D.4: Feedback mode choice by presented order

	(1) Baseline	(2) Informativeness	(3) Framing	(4) Skewness over framing	(5) Baseline reversed
IQ treatment	0.199***	0.151***	-0.324***	-0.200**	-0.057
	(0.076)	(0.057)	(0.087)	(0.083)	(0.070)
Order 2	0.077	0.005	-0.081	-0.066	-0.073
	(0.092)	(0.047)	(0.121)	(0.115)	(0.085)
Order 3	0.127	0.048	-0.061	-0.219**	0.057
	(0.100)	(0.062)	(0.123)	(0.101)	(0.104)
IQ treatment x Order 2	0.024	0.001	0.075	-0.005	0.014
	(0.148)	(0.105)	(0.158)	(0.141)	(0.110)
IQ treatment x Order 3	-0.049	-0.075	0.184	0.153	0.007
	(0.154)	(0.110)	(0.167)	(0.131)	(0.140)
Constant	0.123***	0.035	0.561***	0.386***	0.193***
	(0.044)	(0.025)	(0.067)	(0.065)	(0.053)
R2	0.060	0.047	0.084	0.062	0.019
N	216	216	216	216	216

Notes: Table shows results from regressing the choice of the second alternative in the respective choice situation on the IQ treatment dummy and dummies indicating the order in which choices where presented. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

D.5 Information Preference Scale (Ho et al., ming)

- As part of a semi-annual medical checkup, your doctor asks you a series of questions. The answers to these questions can be used to estimate your life expectancy (the age you are predicted to live to). Do you want to know how long you can expect to live? [1: Definitely don't want to know; 4: Definitely want to know]

- You provide some genetic material to a testing service to learn more about your ancestors. You are then told that the same test can, at no additional cost, tell you whether you have an elevated risk of developing Alzheimer's. Do you want to know whether you have a high risk of developing Alzheimer's? [1: Definitely don't want to know; 4: Definitely want to know]

- At your annual checkup, you are given the option to see the results of a diagnostic test which can identify, among other things, the extent to which your body has suffered long-term effects from stress. Do you want to know how much lasting damage your body has suffered from stress? [1: Definitely don't want to know; 4: Definitely want to know]

- Ten years ago, you had the opportunity to invest in two retirement funds: Fund A and Fund B. For the past 10 years, you have invested all your retirement savings in Fund A. Do you want to know the balance you would have, if you had invested in Fund B instead? [1: Definitely don't want to know; 4: Definitely want to know]

- You decide to go to the theater for your birthday and give your close friend (or partner) your credit card so they can purchase tickets for the two of you, which they do. You aren't sure, but suspect that the tickets may have been expensive. Do you want to know how much the tickets cost? [1: Definitely don't want to know; 4: Definitely want to know]

- You bought an electronic appliance at a store at what seemed like a reasonable, though not particularly low, price. A month has passed, and the item is no longer returnable. You see the same appliance displayed in another store with a sign announcing 'SALE.' Do you want to know the price you could have bought it for? [1: Definitely don't want to know; 4: Definitely want to know]

- You gave a close friend one of your favorite books for her birthday. Visiting her apartment a couple of months later, you notice the book on her shelf. She never said anything about it; do you want to know if she liked the book? [1: Definitely don't want to know; 4: Definitely want to know]

- Someone has described you as quirky, which could be interpreted in a positive or negative sense. Do you want to know which interpretation he intended? [1: Definitely don't want to know; 4: Definitely want to know]

- You gave a toast at your best friend's wedding. Your best friend says you did a good job, but you aren't sure if he or she meant it. Later, you overhear people discussing the toasts. Do you want to know what people really thought of your toast? [1: Definitely don't want to know; 4: Definitely want to know]

- As part of a fund-raising event, you agree to post a picture of yourself and have people guess your age (the closer they get, the more they win). At the end of the event, you have the option to see people's guesses. Do you want to learn how old people guessed that you are? [1: Definitely don't want to know; 4: Definitely want to know]
- You have just participated in a psychological study in which all the participants rate one-anothers' attractiveness. The experimenter gives you an option to see the results for how people rated you. Do you want to know how attractive other people think you are? [1: Definitely don't want to know; 4: Definitely want to know]
- Some people seek out information even when it might be painful. Others avoid getting information that they suspect might be painful, even if it could be useful. How would you describe yourself? [1: If it could be painful, I don't want to know; 4: Even if it could be painful, I always want to know]
- If people know bad things about my life that I don't know, I would prefer not to be told. [1: Strongly agree; 4: Strongly disagree]

D.6 Information Selection and Beliefs in the Control Treatment

In Figure D.1, we plot the prior and posterior beliefs after three rounds of feedback in the endogenous/random treatment. First, as in the IQ treatment, we observe that prior beliefs between Modes A and B are not significantly different from each other ($t(78) = 0.853, p = 0.396$). However, unlike in the IQ treatment, we do not observe that beliefs in the feedback mode diverge with the arrival of signals and, in fact, also the posterior beliefs after three signals are not significantly different ($t(78) = 0.824, p = 0.413$).

Figure D.1: Beliefs before and after signals by feedback mode (endogenous/random treatment)

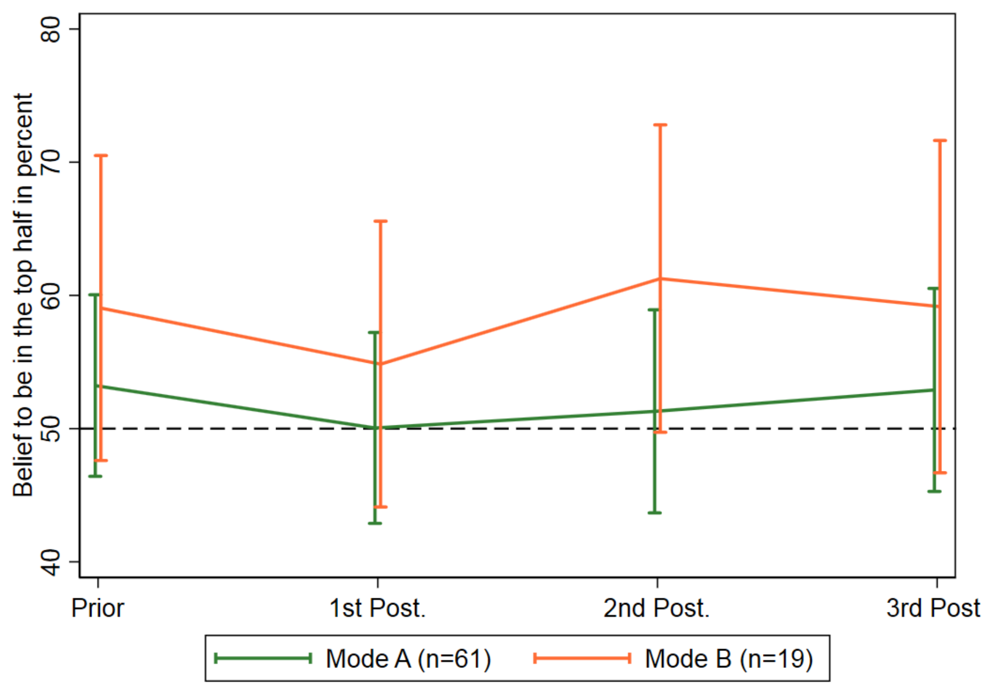

Notes: Plot shows the average prior and posterior beliefs (after each of the three signals) from treatment endogenous/random for the selected feedback mode. The whiskers represent 95 percent confidence intervals.

D.7 Screenshots of the Instructions

Figure D.2: Screenshot of the instructions' template about the IQ task

Instructions Task 1 - The Quiz

In this task you are asked to solve a quiz. The quiz is a non-verbal test that measures abstract reasoning and can be used to estimate fluid intelligence (IQ). High scores in this test are regarded as one of the best predictors for academic and professional success, occupation, income, health, and longevity.

More specifically, in the quiz you are asked to solve 20 puzzles in 10 minutes. This quiz consists of two sets of 10 puzzles. For each set, you will have 5 minutes (300 seconds) to solve them.

Each puzzle consists of a visual geometric design with a missing piece. You are asked to fill in the missing piece from eight possible choices. One example is below (here, the correct answer is piece number 1):

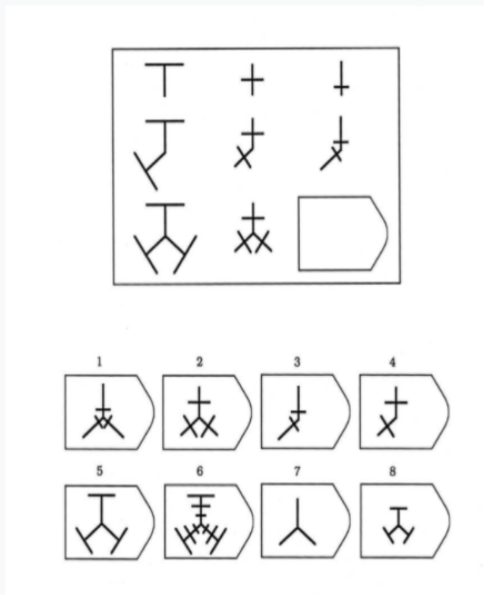

If this task is selected to count for payments, you will be paid for three randomly chosen puzzles. That is, for each of these randomly chosen puzzles, you will be paid £2.00 if your answer is correct.

On the next page, you will be asked comprehension questions about the instructions. You can only proceed with the experiment if you have solved them correctly.

To continue with the comprehension questions for this task, please type in the cell below the number "10".

Next

Notes: The figure displays a screenshot of the template in which we explained to participants the IQ task.

Figure D.3: Screenshot of the instructions' template about the belief elicitation mechanism

In this task, we will ask your beliefs regarding the probabilities of some events. The probabilities you state will affect your payment in this experiment. In particular, the payment method is such that you have the highest chance of earning (more) money by stating the true probability with which you think the event will occur.

We will now explain you how the payment method works. For ease of understanding, let us consider a specific event: The probability that it rained yesterday in New York. Note that this example is only for illustrative purposes, in the experiment it will be replaced by other events.

We give you £6.00 and you have to decide whether you want to place your bet on the lottery or the event:

- **The lottery**: you earn the £6.00 if a purple ball is drawn from an urn containing 100 purple and orange balls. The computer will randomly determine the composition of purple and orange balls (each possible composition is equally likely);
- **The event**: you earn the £6.00 if the event occurred (it rained yesterday in New York).

Should you place your bet on the lottery or the event (to have the highest chances of earning the £6.00)? This will depend on the number of purple balls in the urn and the probability you think it rained yesterday in New York.

For example, if there are 5 purple balls in the urn, the chance to win the £6.00 by picking the lottery is only 5%. Hence, most people would choose the event since the chance that it rained in New York is probably higher than 5%. However, if there are 90 purple balls in the urn, picking the event will only give you a higher chance to win the £6.00 if you believe the probability that it rained in New York is higher than 90%. Therefore in this case most people would pick the lottery.

How are you going to place your bet?

We will ask you to state your belief regarding the probability with which the event occurred. The mechanism ensures that it is optimal for you to state your true belief. In particular, this is because the number you give determines how many purple balls need to be there (in the lottery) for you to prefer to place your bet in the lottery instead of the event. The computer will then determine the composition of the urn. If there happen to be fewer (or equally many) purple balls than the minimum you chose, you will be betting on the event. Thus, you will earn £6.00 if it rained yesterday in New York and £0.00 otherwise. If there happen to be more purple balls than the minimum you chose, you will be betting on the lottery. Then the computer draws a ball. If the ball drawn is purple, you earn £6.00, if the ball is orange you earn £0.00. This method guarantees that you have the highest chances of earning the £6.00 if you report your true belief of the event occurring (if you wish, below the "next" button you can find a more detailed explanation of why this is the case).

While this method may look complicated, its implications are simple: you have the highest chance of earning (more) money if you honestly report your best guess of the probability of the event occurring. For example, if you believe that it rained yesterday in New York with 80%, you should state 80%.

Task 2: Payment

In this task, we will ask you four belief questions. Out of the four, the computer will randomly draw one. The randomly drawn question will count for your payments (if this task is selected to count for payments). In particular, you will be paid for your answer in that question following the method explained here.

Notes: The figure displays a screenshot of the template in which we explained to participants the belief elicitation mechanism.

Figure D.4: Screenshot of the instructions' template about the rank determination in the IQ treatment

Instructions Task 2 - Your IQ Rank

This task is related to your performance in the IQ quiz that you have just completed. In fact, it is about your assessment of your performance in the task compared to the performance of all other people in this session (and who have completed the same IQ quiz as you).

Your task is therefore to guess whether your performance (that is, the number of correctly solved questions) in the IQ quiz is in the top half of the distribution of these participants. Ties are broken randomly. In particular, we ask you to state the probability with which you think that your score is in the top half of the distribution.

To determine your payment, we will use the method we explained to you in the previous screen. Remember that you have the highest chance to win the £6.00 if you state your true belief regarding the probability with which you think that you are in the top half of the distribution.

To continue please type in the cell below the number "30".

Next

Notes: The figure displays a screenshot of the template in which we explain the participant how their rank in the IQ treatment is determined.

Figure D.5: Screenshot of the instructions' template about the rank determination in the random treatment

Instructions Task 2 - Your Number's Rank

From now on all the following tasks will **NOT** be related to the previous IQ quiz that you have just completed.

For this task the computer has drawn for you a random number between 1 and 100, with each number in this interval being equally likely to be drawn. We will call this draw your **number**.

In particular, the computer has drawn the following number:

10

Now, the computer will draw (with replacement) three other numbers (between 1 and 100 with each number in this interval being equally likely to be drawn). Unlike for your number, you will not know the realizations/draws of these three numbers.

Your task is to guess whether your number (10) is in the top half of the distribution of these four randomly drawn numbers (your number and the three other numbers drawn by the computer). Put differently, we ask you whether your number is among the two highest (top half) or among the two lowest numbers (bottom half). Ties are broken randomly. In particular, we ask you to state the probability with which you think that your number is in the top half of the distribution.

To determine your payment, we will use the method we explained to you in the previous screen. Remember that you have the highest chance to win the £6.00 if you state your true belief regarding the probability with which you think your number is in the top half of the distribution.

To continue please type in the cell below the number "30".

Next

Notes: The figure displays a screenshot of the template in which we explain the participant how their rank in the random treatment is determined.

Figure D.6: Screenshot of the prior belief elicitation template in the IQ treatment

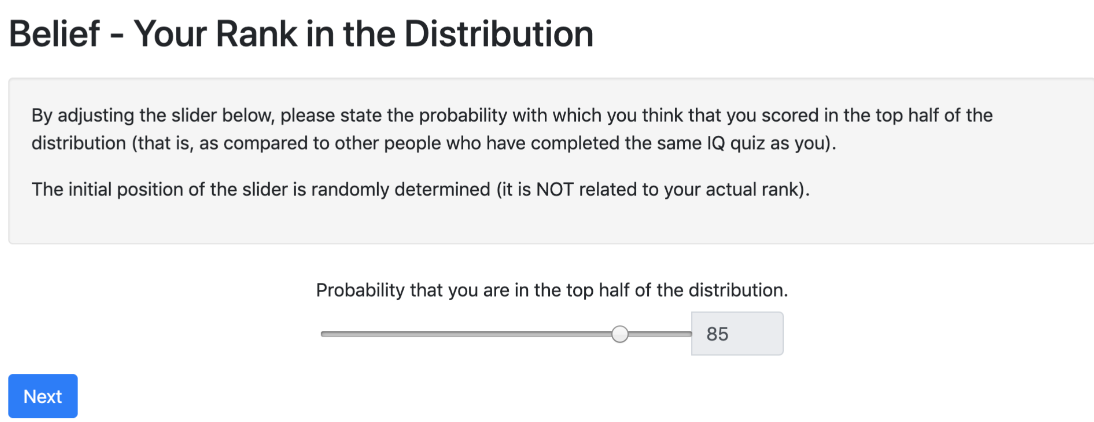
Belief - Your Rank in the Distribution

By adjusting the slider below, please state the probability with which you think that you scored in the top half of the distribution (that is, as compared to other people who have completed the same IQ quiz as you).

The initial position of the slider is randomly determined (it is NOT related to your actual rank).

Probability that you are in the top half of the distribution.

85

Next

Notes: The figure displays a screenshot of the template in which we asked the participant, in the IQ treatment, to state his/her prior belief about his/her relative rank.

Figure D.7: Screenshot of the instructions' template about the possible signals the participant can receive in the IQ/endogenous treatment

Instructions Task 2 - Feedback about your IQ Rank

Now, you will receive additional information (feedback) about your performance in the IQ quiz to help you assess whether or not you are in the top half of the distribution.

What is Feedback?
Depending on your rank in the distribution, you will receive feedback about your rank. You can receive three types of feedback in the form of evaluations:

- The green evaluation that tells you: "You are in the Top Half";
- The red evaluation that tells you: "You are in the Bottom Half";
- The grey evaluation that tells you: "...".

Figure 1 shows you the exact three possible evaluations that you can receive.

Figure 1: Feedback

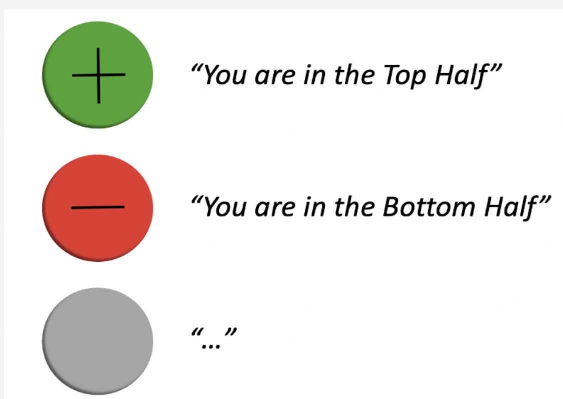

Notes: The figure displays a screenshot of the template in which we explain the participant, in the IQ/endogenous treatment, the possible signals he/she can receive about his/her performance.

Figure D.8: Screenshot of the instruction template regarding the feedback mode selection (Order 1) in the IQ/endogenous treatment

How is the feedback determined?

Which evaluation (feedback) you receive depends on your actual rank in the distribution in the IQ quiz and the "feedback mode" from which the feedback is generated. However, the feedback does not completely reveal your rank in the distribution.

Let's consider an example of two feedback modes to make things clearer:

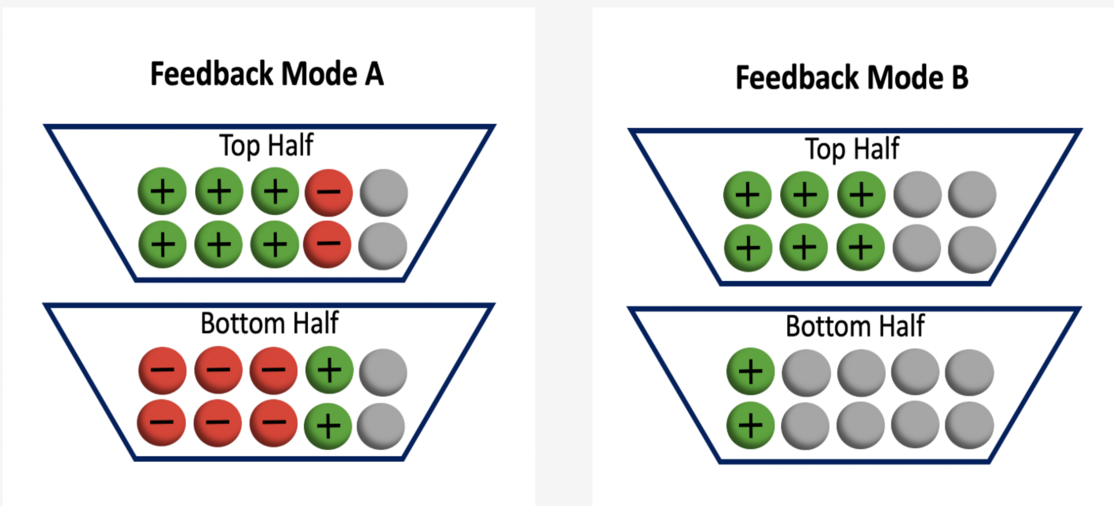

Notice that:

- Irrespective of the feedback mode you choose, if you are in the top half of the distribution your feedback will be determined by the urn at the top of each figure. If you are in the bottom half, your feedback will be determined by the urn at the bottom.
- Consider the example above. In both feedback modes you are more likely to get the green evaluation if you are in the top half of the distribution. However, if you choose to receive feedback from Feedback Mode A, you are more likely to get the red evaluation if you are in the bottom half of the distribution. In Feedback Mode B, in contrast to Feedback Mode A, you will never get the red evaluation but instead the grey evaluation. Thus, Feedback Mode A is more informative than Mode B in case that you are in the bottom half.

Notes: The figure displays a screenshot of the template in which we explain to the participant, in the IQ/endogenous treatment, the selection of feedback modes and how feedback modes differ. In Order 1, we use the baseline choice as an example.

Figure D.9: Screenshot of the instructions' template about the feedback mode selection (Order 2) in the IQ/endogenous treatment

How is the feedback determined?
Which evaluation (feedback) you receive depends on your actual rank in the distribution in the IQ quiz and the "feedback mode" from which the feedback is generated. However, the feedback does not completely reveal your rank in the distribution.

Let's consider an example of two feedback modes to make things clearer:

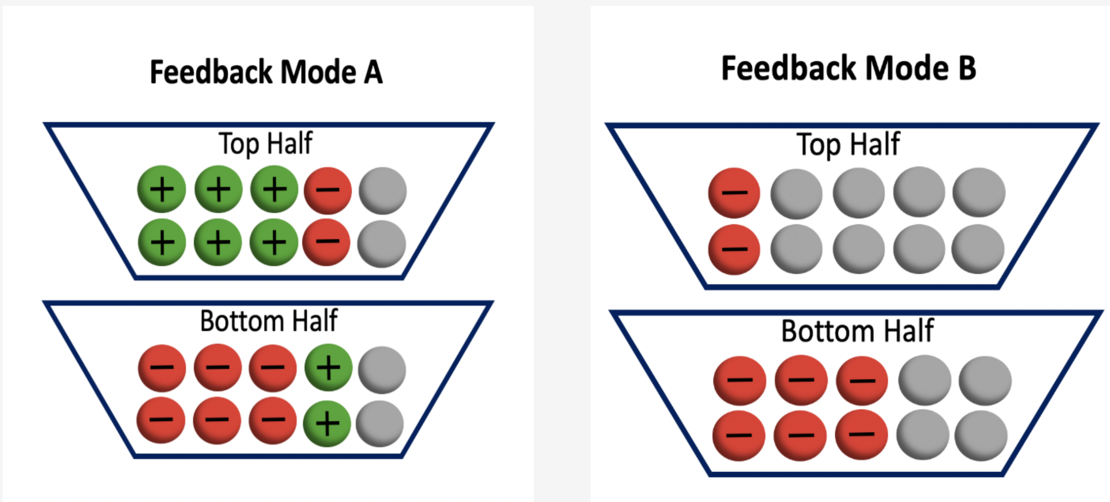

Notice that:

- Irrespective of the feedback mode you choose, if you are in the top half of the distribution your feedback will be determined by the urn at the top of each figure. If you are in the bottom half, your feedback will be determined by the urn at the bottom.
- Consider the example above. In both feedback modes you are more likely to get the red evaluation if you are in the bottom half of the distribution. However, if you choose to receive feedback from Feedback Mode A, you are more likely to get the green evaluation if you are in the top half of the distribution. In Feedback Mode B, in contrast to Feedback Mode A, you will never get the green evaluation but instead the grey evaluation. Thus, Feedback Mode A is more informative than Mode B in case that you are in the top half.

Notes: The figure displays a screenshot of the template in which we explain to the participant, in the IQ/endogenous treatment, the selection of feedback modes and how feedback modes differ. In Order 2, we use the baseline reversed choice as an example. Hence, Mode B in this screenshot is called Mode C in the remainder of the paper.

Figure D.10: Screenshot of the instructions' template about the feedback mode selection (Order 3) in the IQ/endogenous treatment

How is the feedback determined?
Which evaluation (feedback) you receive depends on your actual rank in the distribution in the IQ quiz and the "feedback mode" from which the feedback is generated. However, the feedback does not completely reveal your rank in the distribution.

Let's consider an example of two feedback modes to make things clearer:

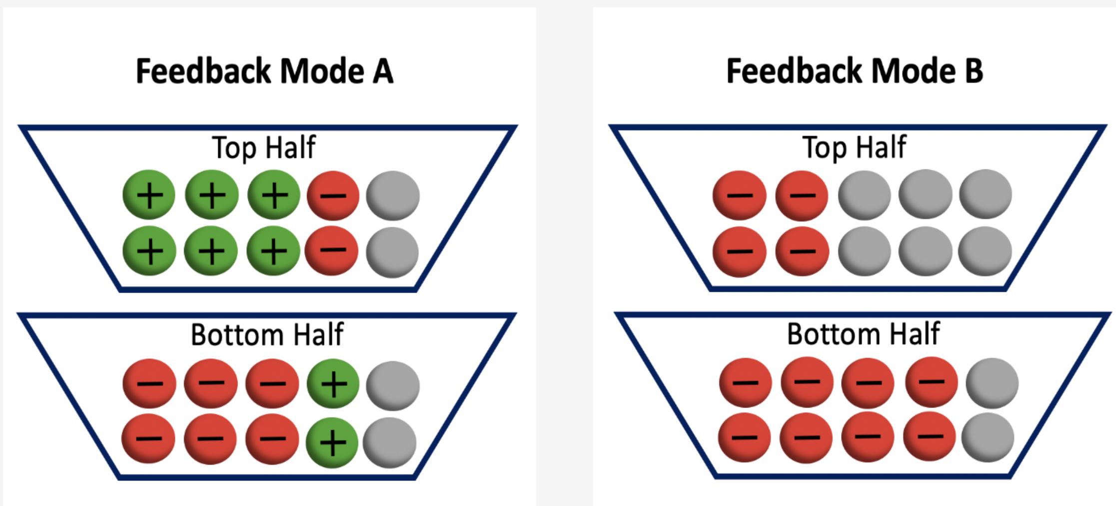

Notice that:

- Irrespective of the feedback mode you choose, if you are in the top half of the distribution your feedback will be determined by the urn at the top of each figure. If you are in the bottom half, your feedback will be determined by the urn at the bottom.
- Consider the example above. In both feedback modes you are more likely to get the red evaluation if you are in the bottom half of the distribution. However, if you choose to receive feedback from Feedback Mode A, you are more likely to get the green evaluation if you are in the top half of the distribution. In Feedback Mode B, in contrast to Feedback Mode A, you will never get the green evaluation. Note that Feedback Mode A is more informative than Mode B in case that you are in the bottom half.

Notes: The figure displays a screenshot of the template in which we explain to the participant, in the IQ/endogenous treatment, the selection of feedback modes and how feedback modes differ. In Order 3, we use the skewness over framing choice as an example. Hence, Mode B in this screenshot is called Mode E in the remainder of the paper.

Figure D.11: Screenshot of the instructions' template about feedback mode A (exogenous treatment)

How is the feedback determined?

Which evaluation you receive depends on your actual rank in the distribution in the IQ quiz. If you are in the top half of the distribution, your feedback will be determined by the urn at the top of the figure. If you are in the bottom half, your feedback will be determined by the urn at the bottom. However, the feedback does not completely reveal your rank in the distribution.

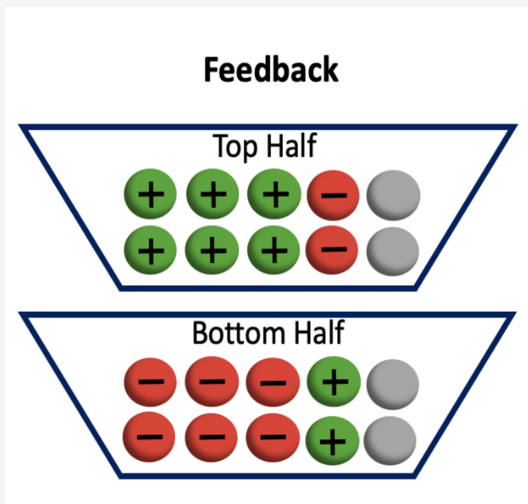

Notice that:

- You are more likely to get the green evaluation if you are in the top half of the distribution.
- You are more likely to get the red evaluation if you are in the bottom half of the distribution.

Notes: The figure displays a screenshot of the template in which we explain to the participant, in the IQ/exogenous treatment, how signals are drawn. In this example, the participant is exogenously assigned to Mode A.

Figure D.12: Screenshot of the instructions' template about feedback mode B (exogenous treatment)

How is the feedback determined?

Which evaluation you receive depends on your actual rank in the distribution in the IQ quiz. If you are in the top half of the distribution, your feedback will be determined by the urn at the top of the figure. If you are in the bottom half, your feedback will be determined by the urn at the bottom. However, the feedback does not completely reveal your rank in the distribution.

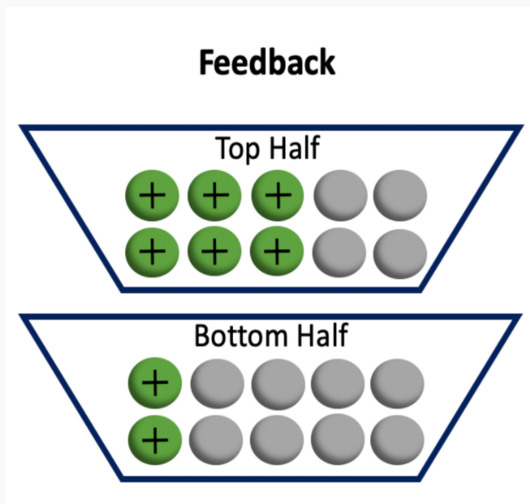

Notice that:

- You are more likely to get the green evaluation if you are in the top half of the distribution.
- You are more likely to get the grey evaluation if you are in the bottom half of the distribution.

Notes: The figure displays a screenshot of the template in which we explain to the participant, in the IQ/exogenous treatment, how signals are drawn. In this example, the participant is exogenously assigned to Mode B.

Figure D.13: Screenshot of the posterior belief elicitation template in the IQ treatment following a green (+) signal

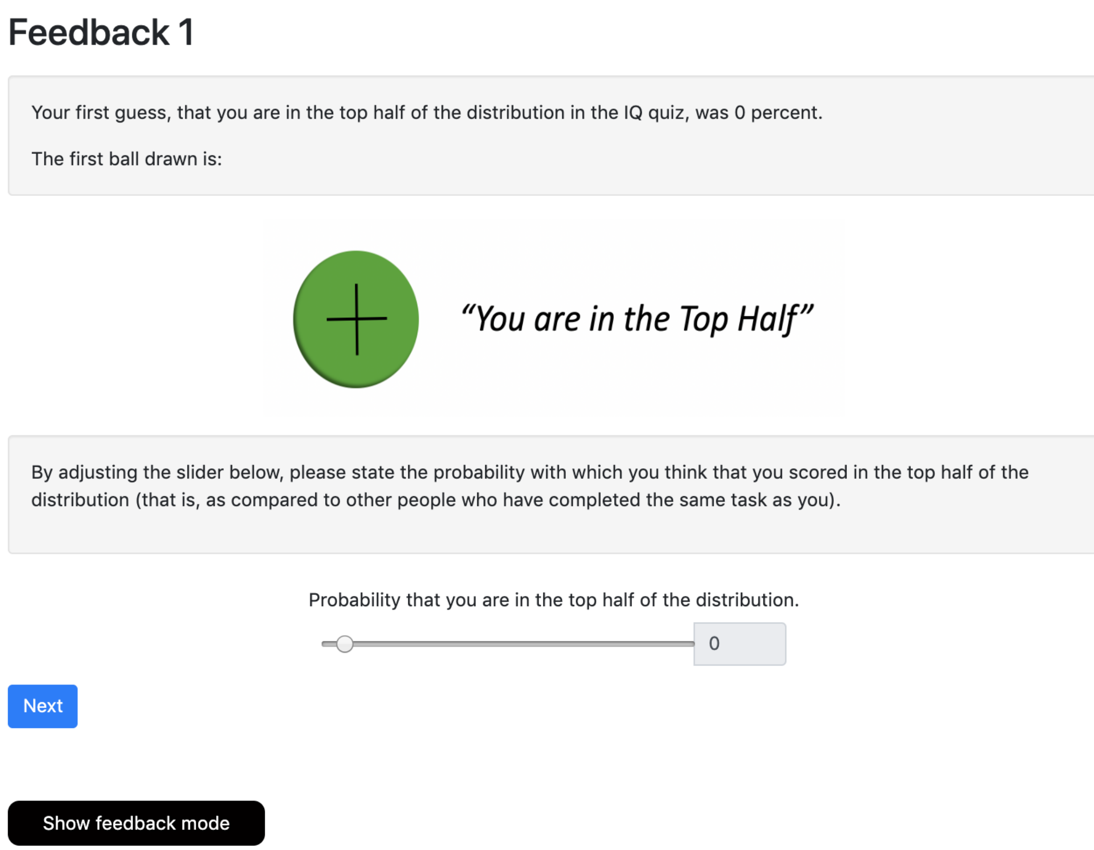

Notes: The figure displays a screenshot of the template in which we asked the participant, in the IQ treatment, to state his/her posterior belief about his/her relative rank following a green (+) signal.

Summary

This dissertation studies how insights from behavioral economics affect the economic analysis of public policy. The thesis consists of four chapters that make use of different methods: laboratory experiments, quasi-experimental and structural econometric methods, as well as theoretical analyses.

The first chapter investigates if social image concerns affect the take-up of a redistributive transfer. In a lab experiment, we vary the visibility of the take-up decision and find that subjects are substantially less likely to take up a public transfer. Moreover, we vary whether transfer eligibility is based on ability or luck, and how the transfer is financed. The results show that subjects avoid the inference both of being low-skilled (ability stigma) and of being willing to live off others (free-rider stigma). These results support our predictions from a theoretical model of social image concerns. Using a placebo treatment, in which the take-up is uninformative about the claimant's type, we exclude other explanations for the observed stigma effects. Although stigma reduces take-up, elicitation of political preferences reveals that only a minority of "taxpayers" vote for the public transfer.

The second chapter studies if sin taxes on soft drinks and fats are effectively targeting consumers with low self-control. For identification, we exploit upward and downward shifts of the soft drink tax and the fat tax in Denmark. We assess the response in purchases empirically using the GfK Consumerscan household panel. With this data, we can separate the sample in consumers with low and high levels of self-control using a survey measure. We find that consumers with low self-control reduce purchases less strongly than consumers with high self-control when taxes go up, but increase purchases to a similar extent when taxes go down. Hence, we document an asymmetry in the responsiveness to increasing and decreasing prices. We show theoretically that these observations are consistent with a model of self-control and rational habit formation. The results suggest that price instruments may not be an effective tool for targeting self-control problems.

The third chapter uses a structural demand model to analyze the impact of soft drink taxes in the presence of habit formation and stockpiling. The model is estimated using nested logit and incorporates unobserved heterogeneity in tastes. The estimated model is used to simulate short-run and long-run price elasticities, as well as the simulated impact of different soft drink taxes. The results show that long-run price elasticities are approximately 20 percent larger than short-run elasticities due to habit formation. Moreover, excise taxes on sugary soft drinks are more effective in reducing sugar consumption than *ad valorem* taxes and excise taxes that do not distinguish between sugary and diet beverages.

The fourth chapter investigates if individuals select information structures in order to protect their motivated beliefs. In a lab experiment, subjects can select the information structure that gives them feedback regarding their rank in the IQ distribution (ego-relevant treatment) or regarding a random number (control treatment). We find that individuals in the ego-relevant treatment select information structures, in which negative feedback is less salient. When receiving such negative feedback with lower salience they update their beliefs less, but only when feedback is ego-relevant. Hence, subjects select information structures that allow them to misinterpret negative feedback in a self-serving way. Moreover, individuals in the ego-relevant feedback choose less informative feedback.

German Summary

Diese Dissertation untersucht, wie Erkenntnisse aus der Verhaltensökonomik die volkswirtschaftliche Politikanalyse beeinflussen. Die Arbeit besteht aus vier Kapiteln und bedient sich verschiedener Methoden: laborexperimentelle, quasi-experimentelle und strukturell ökonometrische, sowie theoretische Analysen.

Das erste Kapitel befasst sich mit der Frage, ob sich der Wunsch nach einem positiven öffentlichen Ansehen ("social image concerns") auf die Inanspruchnahme eines sozialen Transfers auswirkt. In einem Laborexperiment wird die Sichtbarkeit der Inanspruchnahme variiert und gezeigt, dass Teilnehmer deutlich weniger geneigt sind einen öffentlichen Transfer in Anspruch zu nehmen. Darüber hinaus wird variiert, ob die Anspruchsberechtigung auf Können oder Glück basiert und wie der Transfer finanziert wird. Teilnehmer verzichten auf den Transfer, damit andere nicht folgern können, sie seien weniger fähig ("Fähigkeitsstigma") oder bereit auf Kosten anderer zu leben ("Trittbrettfahrerstigma"). Die Ergebnisse unterstützen die Vorhersagen des theoretischen Modells. Mit Hilfe weiterer Kontrollgruppen können alternative Erklärungen für den beobachteten Stigmaeffekt ausgeschlossen werden. Obwohl Stigma die Inanspruchnahme reduziert, stimmen nur wenige der experimentellen "Steurzahler" für einen öffentlichen Transfer, wenn sie nach ihren politischen Präferenzen gefragt werden.

Das zweite Kapitel widmet sich der Frage, ob Individuen mit niedriger Selbstkontrolle besonders auf Softdrinksteuern und Fettsteuern reagieren. Zur Identifikation dienen Erhöhungen und Senkungen dieser "Sündensteuern" in Dänemark. Die Auswirkungen der Steuervariation auf Softdrinkkäufe wird mit Hilfe des GfK Konsumentenpanels untersucht. Diese Daten beinhalten ein Maß für Selbstkontrolle, das es erlaubt die Stichprobe in Individuen mit niedriger und hoher Selbstkontrolle einzuteilen. Nach Steuererhöhungen reduzieren Konsumenten mit niedriger Selbstkontrolle ihre Käufe weniger stark als solche mit hoher Selbstkontrolle. Nach Steuersenkungen erhöhen jedoch beide Gruppen ihre Käufe in ähnlichem Ausmaße. Somit zeigt sich eine asymmetrische Reaktion auf steigende und fallende Preise. Die theoretische Analyse zeigt, dass die Ergebnisse konsistent sind mit einem Modell von Selbstkontrolle und Gewohnheitsbildung ("habit formation"). Die Ergebnisse legen nahe, dass Preisinstrumente nicht gut geeignet sind, Selbstkontrollprobleme zu lösen.

Das dritte Kapitel nutzt ein strukturelles Nachfragemodell, um Auswirkungen von Softdrinksteuern zu analysieren. Das Modell bezieht Gewohnheitsbildung und Vorratshaltung ("stockpiling") in die Betrachtung ein. Es wird als "nested logit" geschätzt und berücksichtigt unbeobachtete Präferenzheterogenität. Mit Hilfe der geschätzten Parameter werden kurzfristige und langfristige Elastizitäten simuliert, sowie die Auswirkungen verschiedener Steuermodelle. Die Ergebnisse zeigen, dass die langfristigen Preiselastizitäten circa 20 Prozent größer sind als die kurzfristigen Preiselastizitäten. Darüber hinaus erweist sich eine Mengensteuer auf zuckerhaltige Getränke als effektiver zur Zuckerredution verglichen mit einer Wertsteuer und einer Mengensteuer, die nicht zwischen zuckerhaltigen und zuckerfreien Getränken unterscheidet.

Das vierte Kapitel untersucht die Frage, ob Individuen Informationsstrukturen selektieren, in denen sie ihre gewünschten Überzeugungen ("motivated beliefs") eher aufrecht erhalten können. In einem Laborexperiment können Teilnehmer Informationsstrukturen wählen, aus denen sie anschließend Informationen über ihren IQ-Rang (ego-relevante Bedingung) oder eine Zufallszahl (Kontrollbedingung) erhalten. Individuen in der ego-relevanten Bedingung wählen Informationsstrukturen, in denen negative Signal weniger salient sind. Wenn die Teilnehmer negative Signale erhalten, die weniger salient sind, aktualisieren sie ihre Überzeugungen in geringerem Ausmaß – dies gilt jedoch nur in der ego-relevanten Bedingung. Individuen wählen also Informationsstrukturen, die es ihnen erlauben die Informationen auf eine eigennützige Weise zu interpretieren. Darüber hinaus wählen Teilnehmer in der ego-relevanten Bedingung weniger informative Signalstrukturen.

Bibliography

Agranov, M. and Palfrey, T. R. (2015). Equilibrium tax rates and income redistribution: A laboratory study. *Journal of Public Economics*, 130:45–58.

Aguilar, A., Gutierrez, E., and Seira, E. (2019). The effectiveness of sin food taxes: Evidence from Mexico. http://www.enriqueseira.com/uploads/3/1/5/9/31599787/obesidad_feb_13_2019_aejsub.pdf (retrieved 09/27/2019).

Ahmed, S. H., Guillem, K., and Vandaele, Y. (2013). Sugar addiction. *Current Opinion in Clinical Nutrition and Metabolic Care*, 16(4):434–439.

Ahn, T., Arcidiacono, P., Hopson, A., and Thomas, J. R. (2019). Equilibrium Grade Inflation with Implications for Female Interest in STEM Majors. NBER Working Paper Series 26556.

Alesina, A. and Angeletos, G.-M. (2005). Fairness and redistribution. *American Economic Review*, 95(4):960–980.

Allcott, H., Lockwood, B. B., and Taubinsky, D. (2019a). Regressive sin taxes, with an application to the optimal soda tax. *The Quarterly Journal of Economics*, 134(3):1557–1626.

Allcott, H., Lockwood, B. B., and Taubinsky, D. (2019b). Should we tax sugar-sweetened beverages? An overview of theory and evidence. *Journal of Economic Perspectives*, 33(3):202–227.

Andreoni, J. and Bernheim, B. D. (2009). Social image and the 50-50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, 77(5):1607–1636.

Andreyeva, T., Long, M. W., and Brownell, K. D. (2010). The impact of food prices on consumption: A systematic review of research on the price elasticity of demand for food. *American Journal of Public Health*, 100(2):216–222.

Ariely, D., Bracha, A., and Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1):544–55.

Bagwell, L. S. and Bernheim, B. D. (1996). Veblen effects in a theory of conspicuous consumption. *American Economic Review*, 86(3):349–373.

Barron, K. (forthcoming). Belief updating: Does the 'good-news, bad-news' asymmetry extend to purely financial domains? *Experimental Economics*.

Bartlett, S. and Burstein, N. (2004). Food stamp program access study: Eligible nonparticipants. *U.S. Department of Agriculture, E-FAN-03-013-2*.

Baumberg, B. (2016). The stigma of claiming benefits: A quantitative study. *Journal of Social Policy*, 45(2):181–199.

Becker, G. S. and Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy*, 96(4):675–700.

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, 117(3):871–915.

Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678.

Bénabou, R. and Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–64.

Benoît, J.-P., Dubra, J., and Moore, D. A. (2015). Does the better-than-average effect show that people are overconfident? Two experiments. *Journal of the European Economic Association*, 13(2):293–329.

Bergman, U. M. and Hansen, N. L. (2019). Are excise taxes on beverages fully passed through to prices? The Danish evidence. *FinanzArchiv*, 75(4):323–356.

Bernheim, B. D. and Taubinsky, D. (2018). Behavioral public economics. In Bernheim, B. D., DellaVigna, S., and Laibson, D., editors, *Foundations and Applications*, volume 1 of *Handbook of Behavioral Economics*, pages 381 – 516. North-Holland.

Besley, T. and Coate, S. (1992). Understanding welfare stigma: Taxpayer resentment and statistical discrimination. *Journal of Public Economics*, 48(2):165–183.

Bhargava, S. and Manoli, D. (2015). Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment. *American Economic Review*, 105(11):3489–3529.

Bhatia, R., Jones, P., and Reicker, Z. (2011). Competitive foods, discrimination, and participation in the national school lunch program. *American Journal of Public Health*, 101(8):1380–1386.

Blank, R. M. and Ruggles, P. (1996). When do women use AFDC and food stamps? The dynamics of eligibility vs. participation. *Journal of Human Resources*, 31(1):57–89.

Blumkin, T., Margalioth, Y., and Sadka, E. (2015). Welfare stigma re-examined. *Journal of Public Economic Theory*, 17(6):874–886.

Blundell, R., Fry, V., and Walker, I. (1988). Modelling the take-up of means-tested benefits: The case of housing benefits in the United Kingdom. *The Economic Journal*, 98(390):58–74.

Bonnet, C. and Réquillart, V. (2013). Tax incidence with strategic firms in the soft drink market. *Journal of Public Economics*, 106:77–88.

Bracha, A. and Vesterlund, L. (2017). Mixed signals: Charity reporting when donations signal generosity and income. *Games and Economic Behavior*, 104:24–42.

Bronnenberg, B. J., Kruger, M. W., and Mela, C. F. (2008). Database paper: The IRI marketing data set. *Marketing Science*, 27(4):745–748.

Bucklin, R. E. and Gupta, S. (1992). Brand choice, purchase incidence, and segmentation: An integrated modeling approach. *Journal of Marketing Research*, 29:201–215.

Bursztyn, L., Ederer, F., Ferman, B., and Yuchtman, N. (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, 82(4):1273–1301.

Bursztyn, L., Egorov, G., and Jensen, R. (2019). Cool to be smart or smart to be cool? Understanding peer pressure in education. *The Review of Economic Studies*, 86(4):1487–1526.

Bursztyn, L. and Jensen, R. (2017). Social image and economic behavior in the field: Identifying, understanding and shaping social pressure. *Annual Review of Economics*, 9:131–153.

Buser, T., Gerhards, L., and van der Weele, J. (2018). Responsiveness to feedback as a personal trait. *Journal of Risk and Uncertainty*, 56(2):165–192.

Camerer, C. and Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89(1):306–318.

Capacci, S., Allais, O., Bonnet, C., and Mazzocchi, M. (2019). The impact of the French soda tax on prices and purchases. an ex post evaluation. *PLOS One*, 14(10):e0223196.

Cappelen, A. W., Konow, J., Sørensen, E. Ø., and Tungodden, B. (2013). Just luck: An experimental study of risk-taking and fairness. *American Economic Review*, 103(4):1398–1413.

Cawley, J., Frisvold, D., Hill, A., and Jones, D. (2019a). The impact of the Philadelphia beverage tax on purchases and consumption by adults and children. *Journal of Health Economics*, 67:10225.

Cawley, J., Thow, A. M., Wen, K., and Frisvold, D. (2019b). The economics of taxes on sugar-sweetened beverages: A review of the effects on prices, sales, cross-border shopping, and consumption. *Annual Review of Nutrition*, 39:317–338.

Chaloupka, F. J., Yurekli, A., and Fong, G. T. (2012). Tobacco taxes as a tobacco control strategy. *Tobacco Control*, 21(2):172–180.

Charles, K. K., Hurst, E., and Roussanov, N. (2009). Conspicuous consumption and race. *The Quarterly Journal of Economics*, 124(2):425–467.

Charness, G. and Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, 104:1–23.

Charness, G., Oprea, R., and Yuksel, S. (2018). How Do People Choose Between Biased Information Sources? Evidence from a Laboratory Experiment. http://econ.ucsb.edu/~sevgi/CharnessOpreaYuksel_Dec2018.pdf (retrieved 07/26/2019).

Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree - an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Chetty, R. (2015). Behavioral economics and public policy: A pragmatic perspective. *American Economic Review: Papers & Proceedings*, 105(5):1–33.

Chew, S. H., Huang, W., and Zhao, X. (2020). Motivated false memory. https://ssrn.com/abstract=2127795 (retrieved 03/02/2020).

Chopra, F., Haaland, I., and Roth, C. (2020). Do people value more informative news? https://drive.google.com/file/d/1D4OAyAXM3gFWDjDwk9b8SycRKkKMzBwh (retrieved 03/26/2020).

Colchero, M. A., Popkin, B. M., Rivera, J. A., and Ng, S. W. (2016). Beverage purchases from stores in Mexico under the excise tax on sugar sweetened beverages: observational study. *BMJ*, 352:h6704.

Colchero, M. A., Rivera-Dommarco, J., Popkin, B. M., and Ng, S. W. (2017). In Mexico, evidence of sustained consumer response two years after implementing a sugar-sweetened beverage tax. *Health Affairs*, 36(3):564–571.

Corneo, G. and Jeanne, O. (1997). Conspicuous consumption, snobbism and conformism. *Journal of Public Economics*, 66(1):55 – 71.

Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22(2):369–395.

Currie, J. (2000). Do children of immigrants make differential use of public health insurance? In Borjas, G. J., editor, *Issues in the Economics of Immigration*, pages 271–308. University of Chicago Press, Chicago, US.

Currie, J. (2003). US food and nutrition programs. In Moffitt, R., editor, *Means-tested transfer programs in the United States*, pages 199–290. University of Chicago Press, Chicago, US.

Currie, J. (2006). The take-up of social benefits. In Auerbach, A., Card, D., and Quigley, J., editors, *Poverty, the Distribution of Income, and Public Policy*, pages 80–148. Russell Sage, New York.

Dal Bó, P. and Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–29.

Daponte, B. O., Sanders, S., and Taylor, L. (1999). Why do low-income households not use food stamps? Evidence from an experiment. *The Journal of Human Resources*, 34(3):612–628.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–372.

DellaVigna, S., List, J. A., Malmendier, U., and Rao, G. (2017). Voting to tell others. *Review of Economic Studies*, 84(1):143–181.

Department of Work and Pensions (2016). Income-related benefits: Estimates of take-up: Financial year 2014/15. https://www.gov.uk/government/statistics/income-related-benefits-estimates-of-take-up-financial-year-201415 (retrieved 06/14/2017).

Diamond, P. A. (1973). Consumption externalities and corrective imperfect pricing. *The Bell Journal of Economics*, 4:526–538.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3):522–550.

DTU Fødevareinstituttet (2008). Danskernes kostvaner 1995-2006. Status og udvikling med fokus på frugt og grønt samt sukker. http://www.food.dtu.dk/-/media/Institutter/Foedevareinstituttet/Publikationer/Pub-2008/Danskernes-kostvaner-1995-2006.ashx?la=da (retrieved 01/15/2019).

Duan, N. (1983). Smearing estimate: A nonparametric retransformation method. *Journal of the American Statistical Association*, 78(383):605–610.

Dubois, P., Griffith, R., and O'Connell, M. (2019). How well targeted are soda taxes? *TSE Working Papers*. 17-868.

Dubé, J.-P., Hitsch, G. J., and Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *RAND Journal of Economics*, 41(3):417–445.

Durante, R., Putterman, L., and van der Weele, J. (2014). Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4):1059–1086.

Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–38.

Enke, B. (forthcoming). What you see is all there is. *The Quarterly Journal of Economics*.

Epstein, L. and Halevy, Y. (2019). Hard-to-interpret signals. https://yoram-halevy.faculty.economics.utoronto.ca/wp-content/uploads/SignalAmbiguity.pdf (retrieved 07/26/2019).

Eriksson, T., Mao, L., and Villeval, M. C. (2017). Saving face and group identity. *Experimental Economics*, 20(3):622–647.

Ertac, S. (2011). Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback. *Journal of Economic Behavior & Organization*, 80(3):532–545.

Eurofound (2015). Access to social benefits: Reducing non-take-up. Publications Office of the European Union, Luxembourg.

Ewers, M. and Zimmermann, F. (2015). Image and misreporting. *Journal of the European Economic Association*, 13(2):363–380.

Exley, C. L. and Kessler, J. B. (2019). Motivated Errors. https://www.dropbox.com/s/fqb32hb1g7e6vrm/ExleyKessler_MotivatedErrors.pdf (retrieved 07/26/2019).

Falk, A. and Zimmermann, F. (2016). Beliefs and utility: Experimental evidence on preferences for information. IZA Discussion Paper No. 10172.

Finkelstein, A. and Notowidigdo, M. (2019). Take-up and targeting: Experimental evidence from SNAP. *The Quarterly Journal of Economics*, 134(3):1505–1556.

Finkelstein, E. A., Ruhm, C. J., and Kosa1, K. M. (2005). Implications of a sugar-sweetened beverage (SSB) tax when substitutions to non-beverage items are considered. *Annual Review of Public Health*, 26:239–257.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Friedrichsen, J. and Engelmann, D. (2018). Who cares about social image? *European Economic Review*, 110:61–77.

Fryer, R. G., Harms, P., and Jackson, M. O. (2019). Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization. *Journal of the European Economic Association*, 17(5):1470–1501.

Ganguly, A. and Tasoff, J. (2016). Fantasy and dread: the demand for information and the consumption utility of the future. *Management Science*, 63(12):4037–4060.

Garz, M., Sood, G., Stone, D. F., and Wallace, J. (2020). The supply of media slant across outlets and demand for slant within outlets: Evidence from us presidential campaign news. *European Journal of Political Economy*, 63(101877).

Gentzkow, M. and Shapiro, J. M. (2010). What drives media slant? Evidence from US daily newspapers. *Econometrica*, 78(1):35–71.

Glazer, A. and Konrad, K. A. (1996). A signaling explanation for charity. *American Economic Review*, 86(4):1019–1028.

Gneezy, U. and Potters, J. (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, 112(2):631–645.

Goffman, E. (1959). *The Presentation of Self in Everyday Life*. Doubleday, New York.

Goffman, E. (1963). *Stigma: Notes on the Management of Spoiled Identity.* Prentice-Hall Inc., New Jersey.

Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135.

Gordon, B. R., Goldfarb, A., and Li, Y. (2013). Does price elasticity vary with economic growth? A cross-category analysis. *Journal of Marketing Research*, 50(1):4–23.

Gordon, B. R. and Sun, B. (2015). A dynamic model of rational addiction: Evaluating cigarette taxes. *Marketing Science*, 34(3):452–470.

Gottfredson, L. S. and Deary, I. J. (2004). Intelligence predicts health and longevity, but why? *Current Directions in Psychological Science*, 13(1):1–4.

Gotthard-Real, A. (2017). Desirability and information processing: An experimental study. *Economics Letters*, 152:96–99.

Gottlieb, D. (2008). Competition over time-inconsistent consumers. *Journal of Public Economic Theory*, 10(4):673–684.

Gradwohl, R. and Smorodinsky, R. (2017). Perception games and privacy. *Games and Economic Behavior*, 104:293–308.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1):114–125.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, 95(3):537–557.

Grossman, Z. and Owens, D. (2012). An unlucky feeling: Overconfidence and noisy feedback. *Journal of Economic Behavior & Organization*, 84(2):510–524.

Gruber, J. and Köszegi, B. (2001). Is addiction "rational"? Theory and evidence. *The Quarterly Journal of Economics*, 116(4):1261–1303.

Grummon, A., Lockwood, B., Taubinsky, D., and Allcott, H. (2019). Designing better sugary drink taxes. *Science*, 365(6457):989–990.

Guadagni, P. M. and Little, J. D. C. (1983). A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2(3):203–238.

Haan, P. (2010). A multi-state model of state dependence in labor supply: Intertemporal labor supply effects of a shift from joint to individual taxation. *Labour Economics*, 17:323–335.

Haavio, M. and Kotakorpi, K. (2011). The political economy of sin taxes. *European Economic Review*, 55(4):575–594.

Haider, S. J., Jacknowitz, A., and Schoeni, R. F. (2003). Food stamps and the elderly: Why is participation so low? *Journal of Human Resources*, 38:1080–1111.

Hartley, R. P. and Lamarche, C. (2018). Behavioral responses and welfare reform: Evidence from a randomized experiment. *Labour Economics*, 54:135–151.

Heckman, J. and Singer, B. (1984). A method for minimizing the distributional assumptions in econometric models for duration data. *Econometrica*, 52:271–320.

Heckman, J. J. (1981). Heterogeneity and state dependence. In Rosen, S., editor, *Studies in Labor Markets*, pages 91–140. University of Chicago Press, Chicago.

Heffetz, O. (2011). A test of conspicuous consumption: Visibility and income elasticities. *The Review of Economics and Statistics*, 93(4):1101–1117.

Hendel, I. and Nevo, A. (2006). Measuring the implications of sales and consumer stockpiling behavior. *Econometrica*, 74(6):1637–1673.

Ho, E., Hagmann, D., and Loewenstein, G. (forthcoming). Measuring information preferences. *Management Science*.

Jensen, J. and Smed, S. (2013). The Danish tax on saturated fat: short run effects on consumption, substitution patterns and consumer prices of fats. *Food Policy*, 42:18–31.

Jin, G. Z., Luca, M., and Martin, D. (2018). Is no news (perceived as) bad news? An experimental investigation of information disclosure. NBER Working Paper Series 21099.

Karlsson, N., Loewenstein, G., and Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2):95–115.

Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2):603–606.

Keane, M. P. (1997). Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics*, 15(3):310–327.

Kleven, H. J. and Kopczuk, W. (2011). Transfer program complexity and the take-up of social benefits. *American Economic Journal: Economic Policy*, 3(1):54–90.

Kline, P. and Tartari, M. (2016). Bounding the labor supply responses to a randomized welfare experiment: A revealed preference approach. *American Economic Review*, 106(4):972–1014.

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4):673–707.

Kuhnen, C. M. (2015). Asymmetric learning from financial information. *The Journal of Finance*, 70(5):2029–2062.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Li, M. and Walker, R. (2017). On the origins of welfare stigma: Comparing two social assistance schemes in rural China. *Critical Social Policy*, 38(4):667–687.

Liebman, J. B. and Luttmer, E. F. P. (2015). Would people behave differently if they better understood social security? Evidence from a field experiment. *American Economic Journal: Economic Policy*, 7(1):275–299.

Liem, D. and de Graaf, C. (2004). Sweet and sour preferences in young children and adults: role of repeated exposure. *Physiology & Behavior*, 83(3):421–429.

Lindbeck, A., Nyberg, S., and Weibull, J. W. (1999). Social norms and economic incentives in the welfare state. *The Quarterly Journal of Economics*, 114(1):1–35.

Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics*, 116(4):1149–1187.

Malmendier, U. and Tate, G. (2005). CEO overconfidence and corporate investment. *The Journal of Finance*, 60(6):2661–2700.

Manning, W. G. and Mullahy, J. (2001). Estimating log models: to transform or not to transform? *Journal of Health Economics*, 20:461–494.

Mas, A. and Moretti, E. (2009). Peers at work. *American Economic Review*, 99(1):112–145.

Masatlioglu, Y., Orhun, A. Y., and Raymond, C. (2017). Intrinsic information preferences and skewness. http://econweb.umd.edu/~masatlioglu/Masatlioglu_Orhun_Raymond.pdf (retrieved 10/01/2019).

McFadden, D. (1978). Modelling the choice of residential location. In A. Karlqvist, F. S. and Weibull, J., editors, *Spatial Interaction Theory and Residential Location*, pages 75–96. North Holland Pub. Co., Amsterdam; New York.

McManus, T. C. and Rao, J. M. (2015). Signaling smarts? Revealed preferences for self and social perceptions of intelligence. *Journal of Economic Behavior & Organization*, 110:106–118.

Mennella, J. A., Bobowski, N. K., and Reed, D. R. (2016). The development of sweet taste: From biology to hedonics. *Reviews in Endocrine and Metabolic Disorders*, 17(2):171–178.

Milton, S., Buckner, S., Salway, S., Powell, K., Moffatt, S., and Green, J. (2015). Understanding welfare conditionality in the context of a generational habitus: A qualitative study of older citizens in England. *Journal of Aging Studies*, 34:113–122.

Möbius, M., Niederle, M., Niehaus, P., and Rosenblat, T. (2014). Managing self-confidence: Theory and experimental evidence. https://web.stanford.edu/~niederle/Mobius.Niederle.Niehaus.Rosenblat.paper.pdf (retrieved 07/26/2019).

Moffitt, R. (1983). An economic model of welfare stigma. *American Economic Review*, 73(5):1023–1035.

Montanari, G. and Nunnari, S. (2019). Audi alteram partem: An experiment on selective exposure to information. http://www.salvatorenunnari.eu/mn_selectexposure.pdf (retrieved 12/15/2019).

Mullahy, J. (1998). Much ado about two: reconsidering retransformation and the two-part model in health econometrics. *Journal of Health Economics*, 17:247–281.

Nakamura, R., Mirelman, A. J., Cuadrado, C., Silva-Illanes, N., Dunstan, J., and Suhrcke, M. (2018). Evaluating the 2014 sugar-sweetened beverage tax in chile: An observational study in urban areas. *PLOS Medicine*, 15(7):e1002596.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175–220.

Nielsen, K. (2018). Preferences for the resolution of uncertainty and the timing of information. https://kirbyknielsen.com/wp-content/uploads/kirby/RoU.pdf (retrieved 07/26/2019).

O'Connell, M. and Smith, K. (2020). Corrective tax design in oligopoly. https://martinoconnell85.github.io/Mywebsite/CorrectiveTaxDesigninOligopoly.pdf (retrieved 05/01/2020).

O'Donoghue, T. and Rabin, M. (2002). Addiction and present-biased preferences. https://blogs.cornell.edu/odonoghue/files/2015/10/addicts7_02-qh3epk.pdf (retrieved 12/10/2019).

O'Donoghue, T. and Rabin, M. (2003). Studying optimal paternalism, illustrated by a model of sin taxes. *American Economic Review*, 93(2):186–191.

O'Donoghue, T. and Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, 90(10-11):1825–1849.

Oster, E., Shoulson, I., and Dorsey, E. (2013). Limited life expectancy, human capital and health investments. *American Economic Review*, 103(5):1977–2002.

Perez-Truglia, R. and Cruces, G. (2017). Partisan interactions: Evidence from a field experiment in the United States. *Journal of Political Economy*, 125(4).

Pinard, C. A., Bertmann, F. M. W., Byker Shanks, C., Schober, D. J., Smith, T. M., Carpenter, L. C., and Yaroch, A. L. (2017). What factors influence SNAP participation? Literature reflecting enrollment in food assistance programs from a social and behavioral science perspective. *Journal of Hunger & Environmental Nutrition*, 12(2):151–168.

Pollak, R. A. (1970). Habit formation and dynamic demand functions. *Journal of Political Economy*, 78(4):745–763.

Powell, L. M., Chriqui, J. F., Khan, T., Wada, R., and Chaloupka, F. J. (2013). Assessing the potential effectiveness of food and beverage taxes and subsidies for improving public health: a systematic review of prices, demand and body weight outcomes. *Obesity Reviews*, 14(2):110–128.

Public Health England (2019). Sugar reduction: Report on progress between 2015 and 2018. https://www.gov.uk/government/publications/sugar-reduction-progress-between-2015-and-2018 (retrieved 10/13/2019).

Rainwater, L. (1982). Stigma in income-tested programs. In Garfinkel, I., editor, *Income-Tested Transfer Programs*, pages 19–46. Academic Press, London.

Rege, M. and Telle, K. (2004). The impact of social approval and framing on cooperation in public good situations. *Journal of Public Economics*, 88(7-8):1625–1644.

Rojas, C. and Wang, E. (2017). Do taxes for soda and sugary drinks work? Scanner data evidence from Berkeley, CA and Washington State. https://people.umass.edu/eywang/EYWang_BerkeleyTax.pdf (retrieved 09/27/2019).

Ruhm, C. J. (2012). Understanding overeating and obesity. *Journal of Health Economics*, 31:781–796.

Sabot, R. and Wakeman-Linn, J. (1991). Grade Inflation and Course Choice. *Journal of Economic Perspectives*, 5(1):159–170.

Schmacker, R. (2015). Welfare stigma, welfare take-up and voting: An experimental investigation. Master thesis, Humboldt-Universität zu Berlin.

Schmacker, R. and Smed, S. (2020). Do prices and purchases respond similarly to soft drink tax increases and cuts? *Economics & Human Biology*, 37:100864.

Schwardmann, P. and Van der Weele, J. (2019). Deception and self-deception. *Nature Human Behaviour*, 3(10):1055–1061.

Seiler, S., Tuchman, A., and Yao, S. (2019). The impact of soda taxes: Pass-through, tax avoidance, and nutritional effects. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3302335 (retrieved 09/27/2019).

Serse, V. (2019). Do sugar taxes affect the right consumers? CORE DP 2019/17.

Sexton, S. E. and Sexton, A. L. (2014). Conspicuous conservation: The Prius halo and willingness to pay for environmental bona fides. *Journal of Environmental Economics and Management*, 67(3):303–317.

Sicherman, N., Loewenstein, G., Seppi, D. J., and Utkus, S. P. (2015). Financial attention. *The Review of Financial Studies*, 29(4):863–897.

Skrondal, A. and Rabe-Hesketh, S. (2004). *Generalized Latent Variable Modeling*. Chapman and Hall, Boca Raton, Florida.

Smith, M. K., Trivers, R., and von Hippel, W. (2017). Self-deception facilitates interpersonal persuasion. *Journal of Economic Psychology*, 63:93–101.

Smith Jr., S. C. (2007). Multiple risk factors for cardiovascular disease and diabetes mellitus. *The American Journal of Medicine*, 120(3):S3–S11.

Soetevent, A. R. (2011). Payment choice, image motivation and contributions to charity: Evidence from a field experiment. *American Economic Journal: Economic Policy*, 3(1):180–205.

Solda, A., Ke, C., Page, L., and von Hippel, B. (2019). Strategically delusional. *Experimental Economics*.

Sommers, B. D. and Epstein, A. M. (2010). Medicaid expansion: The soft underbelly of health care reform? *New England Journal of Medicine*, 363(22):2085–2087.

Spence, A. M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–74.

Sternberg, R. J., Grigorenko, E. L., and Bundy, D. A. (2001). The predictive value of IQ. *Merrill-Palmer Quarterly*, 47(1):1–41.

Stuber, J. and Schlesinger, M. (2006). Sources of stigma for means-tested government programs. *Social Science & Medicine*, 63(4):933–945.

Tangney, J. P., Baumeister, R. F., and Boone, A. L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of Personality*, 72(2):271–322.

Thaler, R. H. and Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press, New Haven & London.

The Danish Ministry of Taxation (2009). Danish tax reform 2010. http://www.skm.dk/media/139042/danish-tax-reform_2010.pdf (retrieved 01/13/2019).

Train, K. (2003). *Discrete Choice Models using Simulation*. Cambridge University Press, Cambridge, UK.

Tuchman, A. E. (2019). Advertising and demand for addictive goods: The effects of e-cigarette advertising. *Marketing Science*, 38(6):913–1084.

Tyran, J.-R. and Sausgruber, R. (2006). A little fairness may induce a lot of redistribution in democracy. *European Economic Review*, 50(2):469–485.

US Department of Health and Human Services (2012). Estimates of child care eligibility and receipt for fiscal year 2009. https://aspe.hhs.gov/system/files/pdf/76541/ib.pdf (retrieved 06/14/2017).

US Department of Health and Human Services (2016). Welfare indicators and risk factors. Fifteenth report to Congress. https://aspe.hhs.gov/pdf-report/welfare-indicators-and-risk-factors-fifteenth-report-congress (retrieved 11/06/2016).

Vallgårda, S., Holm, L., and Jensen, J. D. (2015). The Danish tax on saturated fat: why it did not survive. *European Journal of Clinical Nutrition*, 69(2):223–226.

Von Hippel, W. and Trivers, R. (2011). The evolution and psychology of self-deception. *Behavioral and Brain Sciences*, 34(1):1.

Wagenaar, A. C., Salois, M. J., and Komro, K. A. (2009). Effects of beverage alcohol price and tax levels on drinking: A meta-analysis of 1003 estimates from 112 studies. *Addiction*, 104:179–190.

Wang, E. Y. (2015). The impact of soda taxes on consumer welfare: Implications of storability and taste heterogeneity. *RAND Journal of Economics*, 46(2):409–441.

Whelan, S. (2010). The take-up of means-tested income support. *Empirical Economics*, 39(3):847–875.

WHO (2000). Obesity: Preventing and managing the global epidemic. Technical report series 894, World Health Organization, Geneva.

WHO (2015). Guideline: Sugars intake for adults and children. Guideline, World Health Organization, Geneva.

WHO (2016). Fiscal policies for diet and prevention of noncommunicable diseases. Technical meeting report, World Health Organization, Geneva.

Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics*, 20(1):39–54.

Zhen, C., Wohlgenant, M. K., Karns, S., and Kaufman, P. (2011). Habit formation and demand for sugar-sweetened beverages. *American Journal of Agricultural Economics*, 93(1):175–193.

Zimmermann, F. (2014). Clumped or piecewise? Evidence on preferences for information. *Management Science*, 61(4):740–753.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, 110(2):337–361.

Declaration

Erklärung gem. §4 Abs. 2 der Promotionsordnung

Hiermit erkläre ich, dass ich mich noch keinem Promotionsverfahren unterzogen oder um Zulassung zu einem solchen beworben habe, und die Dissertation in der gleichen oder einer anderen Fassung bzw. Überarbeitung einer anderen Fakultät, einem Prüfungsausschuss oder einem Fachvertreter an einer anderen Hochschule nicht bereits zur Überprüfung vorgelegen hat.

Berlin, Mai 2020

Renke Schmacker

Erklärung gem. §10 Abs. 3 der Promotionsordnung

Hiermit erkläre ich, dass ich für die Dissertation folgende Hilfsmittel und Hilfen verwendet habe: Stata, Matlab und Microsoft Excel.

Berlin, Mai 2020

Renke Schmacker