Buy Baits and Consumer Sophistication: Theory and Field Evidence from Large-Scale Rebate Promotions

Matthias Rodemeier*

This version: January 3, 2021
Link to latest version

Abstract

Can firms exploit behavioral biases to increase profits? Does consumer sophistication about these biases limit the scope of exploitation? To answer these questions, I run a series of natural field experiments with over 600,000 consumers and estimate novel sufficient statistics of consumer sophistication. The empirical application is a ubiquitous and widely regulated form of price discrimination: rebates that need to be actively claimed by consumers. These promotions are suspected of boosting sales even though many consumers eventually fail to claim the rebate—a phenomenon marketers refer to as "slippage." I show theoretically that consumers' *subjective* redemption probabilities can be inferred from how demand responds to rebates as opposed to simple price reductions. I identify these elasticities in three natural field experiments with a major online retailer, in which I randomize prices, redemption requirements, and reminders. Results reveal that claimable rebates in fact increase sales substantially even though 47% of consumers do not redeem the rebate. However, consumers exhibit a remarkable degree of sophistication: the demand response to a rebate is only 76% of the demand response to an equivalent price reduction. Structural estimates imply that consumers almost perfectly anticipate their inattention but vastly underestimate the hassle of redemption by 20 EUR per consumer. Exploiting this misperception increases the profitability of rebates by up to 260%.

JEL Codes: D18, D61, D83, D49, D90

Keywords: consumer sophistication, price discrimination, field experiments, behavioral public economics

---

Acknowledgments: I am deeply grateful to my advisors, Leonardo Bursztyn, John List, and Andreas Löschel, for their support and guidance. I also thank Alec Brandon, Luigi Butera, Felix Chopra, Markus Dertwinkel-Kalt, Andreas Gerster, Sam Hirshman, Justin Holz, Kilian Huber, Chris Knittel, Botond Kőszegi, Thibaut Lamadon, Rob Metcalfe, Maximilian Mueller, Nadine Riedel, Carlo Schwarz, Frederik Schwerter, Dmitry Taubinsky, Jurre Thiel and seminar participants at Aachen, Chicago, Copenhagen Business School and ZEW Mannheim for helpful discussions that improved this paper. The pilot studies have been pre-registered at the AEA RCT registry under trial numbers AEARCTR-0003150 and AEARCTR-0004475. The main field experiment including the theoretical model has been pre-registered under trial number AEARCTR-0005830. All remaining errors are mine.
*University of Münster and Becker Friedman Institute at the University of Chicago, rodemeier@uni-muenster.de.

"These companies are only offering the illusion of a rebate to the many people like me who never get around to claiming it. Because of such thick sludge, redemption rates for rebates tend to be low, yet the lure of the rebate still can stimulate sales—call it 'buy bait.'"

- Richard H. Thaler, 2018

1 Introduction

A growing literature in psychology and economics suggests that behavioral anomalies affect consumption choices. These findings raise the question of whether firms can design contracts that profitably exploit consumers' optimization errors. The scope of exploitation should be governed by consumers' sophistication about their own behavior: consumers who are aware of their behavioral biases will avoid contracts designed to exploit them. Consumer protection regulations are, therefore, frequently motivated by the notion that consumers select into suboptimal contracts because they do not correctly anticipate their behavioral tendencies. Despite the importance for consumer welfare and regulatory policy, empirical evidence on the prevalence of consumer sophistication and its role for contract design is scant.

This paper derives a novel test of consumer sophistication that relies on a small set of sufficient statistics and takes the test to a series of large-scale "natural field experiments" (Harrison and List 2004) in cooperation with an online retailer that frequently offers claimable rebates. Consumer protection agencies suspect that rebate promotions cause excessive demand responses because consumers do not anticipate that they might fail to claim the rebate. I distinguish between two types of redemption frictions that may explain why consumers make mistakes when responding to rebates. First, cognitive *inattention* may cause consumers to forget to redeem the rebate. Second, the *hassle* required to redeem the rebate may be too large to be worth the redemption. Based on a simple choice-theoretic model, I argue that these redemption barriers only cause consumers to make systematic mistakes if consumers do not anticipate them when deciding to buy. Low rebate redemption rates alone are not sufficient to conclude that consumers make mistakes, and should not be the sole basis for consumer protection regulations. A theoretically founded test of sophistication relies on a small set of sufficient statistics: the buying and redemption *elasticities* with and without redemption frictions. Intuitively, a 2 USD rebate that is only redeemed with a probability of 0.5 should increase the buying probability by approximately the same amount as a 1 USD reduction in price. This relationship can be inverted to obtain an estimate of consumers' perceived redemption probability, which then can be compared with the true redemption rates.

To test these predictions empirically, I partner with a large online retailer for furniture and homeware that offers rebates to customers on a regular basis. I design three large-scale natural field experiments that are tightly linked to the theoretical model and that allow me to directly quantify the degree of consumer sophistication about inattention and hassle costs. In particular, website visitors are offered a type of rebate that is ubiquitous in online shopping: the discount only applies to the purchase value if consumers enter a rebate code during the checkout. I cross-randomize the rebate value with a reminder that reduces inattention during checkout and with a feature that automatically redeems the rebate, thereby eliminating hassle costs and inattention simultaneously. In addition, I vary consumer beliefs through an informational treatment that explicitly informs them about the features of the rebate at the start of their visit. A unique feature of the data is that it involves an empirical moment typically unobserved in most observational studies and field trials: the probability of making a purchase at the store for the universe of store visitors. This margin is important both theoretically and empirically as it allows for a direct test of consumer sophistication.

The main experiment produces four main insights. First, making the rebate claimable introduces economically large redemption barriers. Moving from an automatically applied discount to a rebate that needs to be actively claimed is associated with a decrease in redemption rates of 35 percentage points. Second, these barriers appear to be driven by both inattention and hassle costs. Redemption rates are 10 percentage points larger when the firm offers a reminder. Eliminating hassle costs is associated with an increase in redemption rates of 25 percentage points. Third, consumers make large adjustments on the extensive margin in response to redemption barriers. The increase in the buying probability due to a claimable rebate is only 76% of the increase due to an automatically applied discount of equal value. However, the difference in elasticities vanishes when the firm offers a reminder, even though hassle costs remain a redemption friction. Put differently, consumers' buying probability responds to inattention but not to hassle costs. Fourth, firm profits are a monotonically increasing function of redemption frictions on the observed interval. The introduction of hassle costs and inattention increases the rebate's profitability by up to 263% relative to an automatically applied discount. A large share of this effect is driven by consumers not fully anticipating that the rebate needs to be actively claimed. Redemption frictions increase the profitability by only 13% if consumers also receive an information treatment at the start of their visit, informing them that the rebate needs to be actively claimed. The information treatment provides additional evidence that profits depend on consumer expectations about redemption frictions.

Overall, results show that consumers are partially sophisticated. Full naivety can be ruled out because it would imply that consumers' buying elasticity is the same for rebates as it is for automatically applied discounts. Instead, I observe that even though many consumers do not redeem the rebate, a substantial share of them incorporate the low redemption probability in their buying decision. In terms of mechanisms, consumers exhibit sophistication regarding inattention but are almost fully naive with respect to the hassle of rebate redemption. For the supply side, exploiting these frictions is an impressive lever of profit.

Although differences in the buying probability have a causal interpretation by random treatment assignment, differences in the redemption probability with and without the reminder may not be causal, because the reminder may also affect sorting of consumers into the pool of buyers. I address this potential confound by estimating a sample selection model with an arguably credible exclusion restriction: regional and temporal variation in sudden internet outages. Estimation results indicate that the causal effect of the reminder on redemption frictions is almost the same to the previously reported correlations, and that the degree of systematic selection is small, given the underlying distributional assumptions.

To quantify the degree of consumer sophistication, I estimate the sufficient statistics derived from the theoretical model and identify the perceived and true values of redemption frictions. Structural estimates indicate that consumers slightly overestimate their probability of forgetting to redeem the rebate but largely underestimate the effort required for successful redemption. The perceived probability of remembering rebate redemption is 73%, whereas the true probability is 78%. Perceived hassle costs only equal 1 EUR even though true costs are as large as 20.16 EUR. The large underestimation of hassle costs outweighs the slight overestimation of inattention, such that the demand response to rebates is excessively high relative to the fully sophisticated benchmark.

These estimates also directly inform consumer protection policies in the US and across the globe. Policymakers have substantially regulated the features of rebates due to the suspicion that consumers confuse rebates with simple price reductions. Despite the ubiquity of rebate laws, there is little evidence supporting this claim. This paper provides the empirical foundation for these widely used regulatory policies and substantiates the underlying motivation: the demand response to rebates is excessively high because choices do not fully incorporate redemption frictions.

Contributions to the Literature. The paper makes a number of contributions to the existing literature at the intersection of behavioral economics, industrial organization and public finance. First, it provides empirical evidence for the theoretical literature on firm practices of exploiting consumer naiveté (DellaVigna and Malmendier 2004, Eliaz and Spiegler 2006, Gabaix and Laibson 2006, Eliaz and Spiegler 2008, Grubb 2009, Heidhues and Kőszegi 2010, Kőszegi 2014, Heidhues and Kőszegi 2017). Rebate promotions can be seen as price schedules that effectively condition the price a consumer has to pay on her behavioral bias. For instance, consumers that remember rebate redemption receive a lowered price, whereas inattentive consumers pay the full price. The degree of sophistication governs whether this self-selection harms consumers: sophisticated consumers buy and receive the price reduction only if it is optimal for them, whereas naive consumers might buy and pay the full price even though they should not have made the purchase in the first place. If all consumers are sophisticated, rebates are just a simple form of price discrimination that cannot cause suboptimal consumption choices.

Most theory papers on naiveté-based exploitation are motivated by descriptive facts of real-world contracts or common anecdotes, but hard empirical evidence is scarce. The first empirical study comes from DellaVigna and Malmendier (2006) and uses observational data on gym membership contracts. They show that many gym members lose money by choosing flat-rate contracts, likely because they overestimated their gym attendance ex ante. To the best of my knowledge, my study provides the first evidence from a natural field experiment on the economic implications of behaviorally-motivated price discrimination for consumer welfare and firm profits.

The price schedules analyzed in this paper are distinct from well-studied firm practices of shrouding fixed product surcharges, as empirically analyzed in the context of shipping surcharges (Hossain and Morgan 2006, Brown, Hossain, and Morgan 2010), sales taxes (Chetty, Looney, and Kroft 2009, Taubinsky and Rees-Jones 2017), and hidden fees of cinema tickets (Dertwinkel-Kalt, Köster, and Sutter 2020). Shrouding surcharges simply makes a part of the total price less salient, but the actual price for a given product remains the same for all consumer types. By contrast, the contract design studied in this paper varies the price of the goods as a function of the behavioral bias. Consumers do not simply make mistakes because costs are not salient, but because they do not correctly predict their own behavior (i.e., rebate redemption).

Second, the current paper joins a rapidly growing literature that attempts to identify deep primitives of economic models motivated by insights from psychology (see DellaVigna 2018 for an overview). Methodologically, I provide a novel test of consumer sophistication about behavioral biases that relies only on aggregate elasticities. The test builds on the sufficient-statistics approach in Chetty, Looney, and Kroft (2009) and extends it by modeling a new extensive margin on which consumers respond to frictions before they are even observed in many transaction-level datasets. This extension has important implications for the welfare effects of behavioral frictions and the design of policies aimed at correcting these frictions.[1] Because sophisticated consumers do not make mistakes in expectation, corrective policies may introduce substantial inefficiencies if policymakers ignore the degree of sophistication, thereby implicitly assuming full naiveté.[2]

A small number of existing papers identify structural parameters of agents' sophistication about their behavioral biases. Almost all of the papers study sophistication in the context of present bias. One elicitation method is to let subjects predict their future behavior and to incentivize accurate predictions. Structural parameters of perceived and true present bias are then inferred from a comparison of predicted and true future behavior (Augenblick and Rabin 2019, Allcott et al. 2020). An alternative approach estimates willingness to pay for commitment devices and compares the estimates with the true value of these devices (Bai et al. 2017, Carrera et al. 2019, Chaloupka IV, Levy, and White 2019).[3] A closely related paper by Bronchetti et al. (2020) uses this approach in the context of inattention and shows that willingness to pay for a reminder is below the true returns of being attentive.

[1] Bernheim and Taubinsky (2019) provide an overview of policies aimed at correcting behavioral biases.

[2] Also recall that if all consumers are sophisticated, rebates are just a traditional form of price discrimination. Consumers only obtain a discount if its value exceeds the individual transaction costs related to redemption. Paying attention to rebate redemption is simply a costly activity that is fully incorporated into choices, such as in models of rational inattention (see, e.g., Caplin, Dean, and Leahy 2017).

[3] Löschel, Rodemeier, and Werthschulte (2020) estimate willingness to pay for "soft" commitments such as self-set goals in the context of time inconsistency.

The test developed in my study follows an alternative but related strategy to the latter approach. I implicitly infer consumers' willingness to pay to avoid their behavioral biases, by observing whether they decide to buy at a store that exploits the biases. Choosing not to buy at the store is a quasi-commitment against falling for the traps set by the store. Willingness to pay to avoid these traps can then be quantified by a comparison of buying elasticities with respect to the rebate value and with respect to exogenous variation in the behavioral frictions.

Third, this paper adds novel insights to an interdisciplinary literature that studies whether rebates can exploit overconfidence. Prior work in marketing has analyzed "slippage" or "breakage," which is the terminology for the phenomenon whereby consumers are enticed by a rebate even though they fail to redeem it (Jolson, Wiener, and Rosecky 1987, Silk 2004).[4] Subjects are either directly asked how often they forget to redeem rebates or are asked to predict their redemption probability. Even though results indicate that subjects overestimate their redemption probability, the elicitation methods rely entirely on non-incentivized statements. These measures are a helpful starting point in understanding redemption rates. However, they are unlikely to provide sufficient evidence to draw conclusions about consumer sophistication, because self-reported and non-incentivized predictions suffer from several reporting biases that are documented in the literature (e.g., Cummings, Harrison, and Rutström 1995, List and Gallet 2001).

Two laboratory experiments in behavioral economics on overconfidence use incentivized choices to study whether subjects overestimate their probability of redeeming a rebate. Ericson (2011) and Tasoff and Letzler (2014) let university students choose between a rebate that needs to be claimed in the future and an automatically applied discount. In line with the self-reported measures, the authors find subjective redemption probabilities are above true redemption rates.

I make several complementary contributions to these two laboratory studies. First, the two experiments only implicitly test whether demand responses to rebates are excessively high, by eliciting perceived and true redemption probabilities. By contrast, I directly observe demand responses to rebates, which allows me to explicitly test whether the buying probability is too large. Second, elicitation methods used in both laboratory studies only identify lower bounds of the perceived redemption probabilities, whereas the sufficient-statistics approach I present in this paper derives point estimates. Third, my experimental design isolates the underlying mechanisms of overconfidence and rules out naiveté about time discounting because the rebate redemption choice happens almost immediately. From the previous studies, it is unclear how naiveté about time preferences affects the results because there is a long gap between the elicitation of willingness-to-pay and the rebate redemption choice. Fourth, subjects in my study are actual market participants that are not aware of being part of an experiment. The difference in the subject pool may be important because prior work documents that market experience can eliminate certain behavioral

[4] Gilpatric (2009) develops a theoretical model to study the implications of time inconsistency for rebate redemption rates and firm profits.

biases (List 2003, List 2011, Allcott et al. 2020). The field experiments presented in this paper document behavior of real market participants in their natural environment, not knowing they are being observed by a researcher.

I find that subjects in my experiments are, in fact, more sophisticated than in the laboratory studies but that overconfidence remains a serious anomaly that firms can exploit. Related to this finding, a fifth novel contribution of this paper is that it sheds light on the implications of rebates for the supply side by estimating firm profits as a function of behavioral frictions.

The rest of this paper is structured as follows. Section 2 introduces the institutional background of rebate regulations and the type of rebates that firms use. In Section 3, I present a simple model of consumer responses to rebates and derive empirically testable predictions of sophisticated behavior. The model guides the field experimental design discussed in Section 4. Reduced-form results are presented in Section 5. In Section 6, I estimate the structural parameters of sophistication and discuss policy implications. Section 7 discusses additional mechanisms that can explain the data. Section 8 concludes.

2 Institutional Background

Firms offer a number of different rebates to consumers with various degrees of redemption requirements. The most traditional type are mail-in rebates that need to be redeemed by sending in a form to the firm. These rebates can often be redeemed only several weeks after the purchase, and the time window for redeeming the rebate may be short. Another rebate type are coupons that can be found in newspapers and advertisements. Coupons can typically be redeemed at the point of sale by handing them to the cashier. With the rise of digital technologies, most rebates can be more conveniently redeemed using a mobile phone or a web browser. For instance, customers may show a number or bar code on their mobile phone to a cashier in order to validate the discount at the point of sale. These so-called "instant rebates" are also the most common form in online shopping where a rebate code must be entered in a field on the checkout page of the shop.

The fact that consumers may forget to redeem the rebate is well acknowledged in marketing and has been coined "slippage." Instant rebates may seem to involve less slippage and fewer redemption efforts than mail-in rebates because no large temporal lag occurs between the act of buying and redemption. However, as I show in the empirical application, instant rebates involve substantial redemption barriers despite the fact that they are redeemed immediately.[5]

Many governments have decided to regulate rebates out of fear that consumers mistake them for simple price discounts. In the US, the Federal Trade Commission warns consumers to not be

[5] Although redemption barriers may be lower for instant rebates than for traditional mail-in rebates, a less obvious point is whether consumer sophistication varies between these rebate types. For instance, consumers may anticipate that rebates with a larger temporal redemption lag have a lower redemption probability, such that the level of sophistication is similar between mail-in and instant rebates.

"'baited' by rebates that never arrive" Federal Trade Commission (2020). The regulation of these promotions lies in the responsibility of each state. The Connecticut "Unfair Trade Practices Act" allows firms to advertise a price net of the rebate only if the rebate is immediately applied. The "Unfair Trade Practice and Consumer Protection Act" in Rhodes Island explicitly states that the burden of rebate redemption should be on the side of the retailer if prices are advertised net of the rebate.[6] California, Maryland, and New Jersey have similar regulations that illustrate the regulators' concerns that consumers may confuse rebates with simple price reductions. Other countries around the world have also heavily regulated the use of rebates in order to protect consumers, such as Australia, Canada, and Germany.

Marketers defend rebates as a way for consumers to receive attractive offers. The Promotion Marketing Association (PMA) argues that low redemption rates are misleading because they are calculated as a percentage of the total number of transactions, thereby including inframarginal consumers who would have purchased in the absence of the rebate (PMA 2005). According to the authors, the more appropriate denominator for calculating redemption rates would be incremental sales, because it captures only consumers attracted by the rebate.

In fact, a benefit of this ratio is that it is informative about sophistication in the special case in which the number of redemptions is smaller than the number of incremental sales. If the rebate causes 100 additional transactions but only 70 consumers redeem it, this result would be sufficient to conclude that some consumers were not fully aware of the redemption barriers when making the decision to buy. However, in any other case, the ratio does not sufficiently answer the question of whether consumers who are attracted by the rebate also claim it. For example, imagine an extreme case in which none of the consumers attracted by the rebate manage to claim it but some of the inframarginal consumers do. In this case, the share of redemptions as a fraction of incremental sales can be equal to or larger than 100% even though *all* consumers attracted by the rebate fail to redeem it.

More formally, the empirical challenge is that the joint distribution of redemption choice and treatment effects of rebates on sales remains unobserved, which prevents us from knowing whether marginal or inframarginal consumers take up the rebate. The approach presented in this paper solves this issue by deriving empirically testable predictions of how *aggregate* moments of sales and redemption must behave if consumers fully take into account redemption barriers. It is the first consistent empirical test to inform public policy regarding whether rebates harm consumers, and helps resolve the dispute between regulators and marketers through empirical evidence founded in economic choice theory.

As a preview to the upcoming analysis, applying the ratio suggested by PMA to my empirical results would likely be misinterpreted as optimistic news by regulators. The number of redemptions as a fraction of incremental sales is greater than 300% in my empirical application even though the

[6] See Connecticut Unfair Trade Practices Act Regulations (2004) and Rhode Island General Laws (2016).

sufficient statistics indicate that the aggregate demand response to rebates is substantially above the response of sophisticated consumers.

3 Theoretical Framework

3.1 A Simple Model of Consumer Responses to Rebates

In this section, I develop a simple model of consumer behavior and derive an empirical test of consumer sophistication. The model has been pre-registered at the AEA RCT Registry under trial ID AEARCTR-0005830 and directly produced the experimental design.

In the most concise version of the model, the consumer faces two choices.[7] On the extensive margin, she chooses where to buy, $b \in \{0, 1\}$, where $b = 1$ means she buys at a location I call "the store," and $b = 0$ means she chooses the outside option. The outside option may represent one or multiple competitors of the store, but it may also mean the consumer does not buy anywhere. The store offers a rebate of value s that is only redeemed if the consumer actively claims it. Without loss of generality, I assume the outside option does not offer a rebate.

On the intensive margin, she decides whether to claim the rebate conditional on buying at the store. Let her rebate redemption choice be represented by $r \in \{0, 1\}$, where $r = 1$ means she redeems and $r = 0$ means she does not redeem. Claiming the rebate causes hassle costs c. For instance, the consumer might have to search for a discount code and show it to the cashier during the check out. Claiming it may also require effort to understand the details of redeeming the rebate correctly.

Another redemption barrier is inattention: the consumer may eventually forget to claim the rebate even if redemption would have been worth the hassle. Her probability of remembering rebate redemption is given by $\theta \in [0, 1]$.

I allow for the possibility that the consumer is not perfectly aware about these frictions when deciding whether to buy at the store. Let $\hat{c}$ and $\hat{\theta}$ be the consumer's *perceived* hassle costs and probability of remembering to redeem the rebate, respectively. The difference between the perceived and true values of inattention and of hassle costs measures the degree of sophistication. A consumer is said to be *sophisticated* about her redemption frictions if $\hat{\theta} = \theta$ and $\hat{c} = c$. She is *naive* if $\theta \neq \hat{\theta} = 1$ and $c \neq \hat{c} = 0$. The consumer is said to be *partially naive* if neither of the two preceding definitions applies. A sophisticated consumer correctly anticipates redemption barriers, whereas a naive consumer falsely ignores redemption barriers, thereby confusing a rebate with a simple reduction in prices. Partially naive consumers realize a difference exists, but will either over- or underestimate redemption friction to some degree.

[7] In Appendix A.3, I extend the model by a third stage in which the consumer also decides which goods to buy. This extension does not affect the predictions of the model. In Appendix A.2, I also extend the model by introducing heterogeneity in inattention and hassle costs, and discuss implications for the empirical identification.

We can characterize consumer behavior as a two-stage decision process and solve it backwards. Given that the consumer buys at the store and remembers rebate redemption, she chooses $r = 1$, if and only if

$$s - c \geq \kappa, \tag{1}$$

where κ is an idiosyncratic taste parameter. This parameter may represent the utility the consumer receives from choosing not to redeem (i.e., the outside option utility) or it may also represent other forms of noise affecting the redemption decision. In my empirical application, κ may, for example, represent random technical issues that affect rebate redemption in the online shop.

On the extensive margin, the consumer chooses whether to buy at the store, taking into account rebate-redemption frictions. She chooses $b = 1$ if and only if the perceived expected utility from buying at the store exceeds the utility she gets from buying at the outside option:

$$\hat{\theta}\left[\hat{r}(s - \hat{c}) + (1 - \hat{r})\kappa\right] + (1 - \hat{\theta})\kappa \geq \epsilon, \tag{2}$$

where ϵ is the outside option utility and $\hat{r} = r(s, \hat{c}, \kappa)$ the perceived redemption decision. Note that the expectation is taken over the distribution of perceived inattention. Similarly, perceived rather than true hassle costs matter for her extensive-margin choice. The model does not make any assumption regarding the formation of consumers' beliefs about their redemption frictions. Consumers may be Bayesian agents that form their beliefs about inattention and hassle costs rationally. They may also have systematically distorted beliefs due to some of the biases in belief formation as documented in the literature.[8]

Expected utility is a function of the true redemption frictions and given by

$$U(b) = b\left[\theta\left(r(s - c) + (1 - r)\kappa\right) + (1 - \theta)\kappa\right] + (1 - b)\epsilon, \tag{3}$$

whereas the actual buying decision, denoted $\hat{b}$, is determined by the perceived redemption barriers:

$$\hat{b} = \arg\max_b\{\hat{U}(b)\} = \arg\max_b\{b\left[\hat{\theta}\left(\hat{r}(s - \hat{c}) + (1 - \hat{r})\kappa\right) + (1 - \hat{\theta})\kappa\right] + (1 - b)\epsilon\}. \tag{4}$$

Let $\Psi = U(\hat{b}) - \hat{U}(\hat{b})$ be the difference between the utility the consumer obtains from the buying

[8] See, e.g., Falk and Zimmermann (2018) and Enke and Zimmermann (2019).

decision and the utility she believes she obtains from that decision. Then,

$$\Psi = b(\hat{c}, \hat{\theta})\Big[\theta r^* \left(s - c - \kappa\right) - \hat{\theta}\hat{r}(s - \hat{c} - \kappa)\Big], \tag{5}$$

where $r^* = r(s, c, \kappa)$ is the realized redemption choice. The difference Ψ is commonly referred to as an "internality" or "behavioral wedge" in the behavioral economics literature. It is a price metric that quantifies the welfare loss the consumer bears due to making a mistake. Importantly, this difference is only nonzero in my model if there is a lack of sophistication. The existence of redemption barriers or behavioral biases, that is, $\theta \neq 1$ and $c \neq 0$, is not sufficient to conclude that consumers make systematic mistakes. The degree of sophistication is what governs whether redemption barriers create an internality, not the behavioral friction itself. This insight has an important implication for consumer protection regulations: the existence of behavioral frictions alone should not be the sole motivation for regulatory policy. Moreover, low rebate redemption rates are not informative regarding the question of whether consumers make mistakes.

Equation 5 also has important implications for the literature on inattention more generally because it changes the interpretation of inattention as a behavioral bias. Most models of inattention do not explicitly allow for sophistication, thereby implicitly assuming full naiveté. A typical research design would randomize a reminder treatment and conclude that consumers make mistakes due to inattention whenever the reminder increases the desired behavior (e.g., redemption rates).[9] This approach ignores that consumers may anticipate their inattention and make appropriate adjustments on an extensive margin. If consumers are fully aware of their probability of being attentive, inattention does not create a behavioral wedge but is just a simple form of transaction cost. In this case, internality taxes on rebates, for example, cannot improve consumer welfare, but instead would create additional inefficiencies.[10] Optimal policy, therefore, crucially depends on the degree of sophistication.

3.2 Aggregating Consumer Choices

To aggregate individual behavior, let the idiosyncratic taste parameters on the intensive and extensive margin follow an atomless joint distribution $G(\kappa, \epsilon)$. The marginal distributions are denoted by $H(\kappa)$ and $F(\epsilon)$. The probability of redeeming the rebate (unconditional on buying), denoted $R(s, \theta, c)$, is given by

$$R(s, \theta, c) = \theta \int^{s-c} dH(\kappa). \tag{6}$$

[9]An exception is Bronchetti et al. (2020), who explicitly allow for ex-ante beliefs about inattention.

[10]By contrast, informational interventions that directly eliminate inattention may yield large welfare gains to consumers by reducing the *transaction costs* of paying attention. See Farhi and Gabaix (2020) and Rodemeier and Löschel (2020) for theoretical frameworks and empirical evidence comparing the efficiency effects of informational interventions with internality taxes.

The probability of buying at the store, denoted $B(s, \hat{\theta}, \hat{c})$, can be written as

$$B(s, \hat{\theta}, \hat{c}) = \underbrace{\int^{s-\hat{c}} \int^{\hat{\theta}(s-\hat{c})+(1-\hat{\theta})\kappa} dF(\epsilon|\kappa)dH(\kappa)}_{\text{share of buyers intending to redeem}} + \underbrace{\int_{s-\hat{c}} \int^{\kappa} dF(\epsilon|\kappa)dH(\kappa)}_{\text{share of buyers not intending to redeem}}, \quad (7)$$

which consists of the share of buyers who intend to redeem and the share of buyers who do not intend to redeem.

Both the buying and the redemption probability depend on the size of the rebate. However, the redemption probability is determined by the true redemption frictions, whereas the buying probability is a function of the consumers' beliefs about these frictions. This formulation has intuitive implications for how claimable rebates may act as buy baits. The demand response to the rebate may be excessively high because consumers do not fully take into account redemption barriers when deciding whether to buy at the store.

3.3 Testable Predictions of Consumer Sophistication

A simple test of full naiveté is to observe whether the buying probability responds *at all* to exogenous variation in redemption barriers. If the buying elasticity with respect to redemption barriers is fully inelastic, consumers are fully naive about redemption barriers. Conversely, any reduction in the buying probability as a response to an increase in redemption barriers is evidence of some degree of consumer sophistication.

The model also allows us to go one step further and quantify the degree of consumer sophistication empirically by testing whether the buying probability decreases *sufficiently* with more difficult rebates. Proposition 1 establishes that we can empirically identify consumer beliefs about inattention and hassle costs through a small set of reduced-form elasticities: the effects on the buying and redemption probabilities of treatment variation in the rebate size, the hassle costs, and the probability of paying attention. I let $\Delta_c R$ denote the effect of a change in the hassle costs by Δc on the redemption probability. Analogously, $\Delta_{\hat{c}} B$ denotes the change in the buying probability in response to a change in the perceived hassle costs of $\Delta \hat{c}$.

Proposition 1. *Perceived inattention and perceived hassle costs can be approximated by reduced-form responses on the extensive margin:*

$$\hat{\theta} \approx \frac{\frac{\partial}{\partial s}B(s,\hat{\theta},\hat{c})}{\frac{\partial}{\partial s}B(s,1,\hat{c})} \tag{8}$$

$$\hat{c} \approx \frac{\Delta_{\hat{c}}B(s,1,\hat{c})}{\frac{\partial}{\partial s}B(s,1,\hat{c})} \tag{9}$$

True inattention and true hassle costs can be identified through reduced-form responses on the intensive margin:

$$\theta = \frac{R(s,\theta,c)}{R(s,1,c)} \tag{10}$$

$$c = \frac{\Delta_{c}R(s,1,c)}{\frac{\partial}{\partial s}R(s,1,c)}. \tag{11}$$

Consumers are sophisticated if and only if

$$\frac{\frac{\partial}{\partial s}B(s,\hat{\theta},\hat{c})}{\frac{\partial}{\partial s}B(s,1,\hat{c})} \approx \frac{R(s,\theta,c)}{R(s,1,c)} \tag{12}$$

and

$$\frac{\Delta_{\hat{c}}B(s,1,\hat{c})}{\frac{\partial}{\partial s}B(s,1,\hat{c})} \approx \frac{\Delta_{c}R(s,1,c)}{\frac{\partial}{\partial s}R(s,1,c)}. \tag{13}$$

The approximations in equations 8 and 9 each require that $f(\epsilon|\kappa)$ is roughly constant on the interval $[\hat{\theta}(s-\hat{c})+(1-\hat{\theta})\kappa, s-\hat{c}]$ for all κ, which will be more accurate for small levels of inattention. Perceived inattention is approximated by the ratio of two treatment effects on the buying probability of a small change in the rebate size. The numerator is the treatment effect under perceived inattention, $\hat{\theta}$, whereas the denominator is the effect under perceived full attention. The treatment effect in the numerator can be empirically identified by observing how the buying probability responds to a small change in the value of a claimable rebate. For the treatment effect in the denominator, we need to observe the same margin but to a rebate that consumers remember with certainty. For instance, the store could offer a reminder at the checkout telling consumers not to forget rebate redemption.

To understand the intuition behind equation 8, imagine the share of consumers who buy at the store increases by two percentage points as a response to a 1 USD rebate. It increases by four percentage points in response to the same rebate if the store additionally offers a reminder during checkout that eliminates inattention. Then, this implies that consumers believe to forget rebate redemption half of the time in the absence of a reminder, that is, $\hat{\theta} = 0.5$. It also implies that consumers willingness to pay for the reminder is approximately 1 USD.

Perceived hassle costs are identified in a similar way. The numerator in equation 9 is the treatment effect on the buying probability of fully eliminating the hassle of rebate redemption, given that the store already offers a reminder. For instance, the store could move to a new policy in which the rebate is automatically redeemed and therefore requires no effort. Since $\hat{c}$ has a money metric, the treatment effect on the buying probability in the numerator needs to be scaled by the demand derivative with respect to the rebate value. Money-metric hassle costs therefore equal the necessary change in the rebate value that would cause the same increase in the buying probability as a complete elimination of hassle costs would generate. If the buying probability increases by one percentage point to a 2 USD increase in rebate value, but by four percentage points in response to the elimination of hassle costs, then perceived hassle costs equal approximately 8 USD.

The true redemption frictions are identified through behavioral responses on the intensive margin. The true probability of remembering rebate redemption is simply the redemption probability under inattention as a share of the redemption probability without inattention. True hassle costs are identified analogous to perceived hassle costs: they equal change in the rebate value necessary to generate the same effect on the redemption probability as a treatment that fully eliminates hassle costs. Both identification strategies parallel Chetty, Looney, and Kroft (2009), who identify inattention to sales taxes by comparing demand elasticities with and without a treatment increasing the salience of taxes. However, their work does not allow for an identification of consumer sophistication, because extensive-margin responses are not modeled theoretically and remain unobserved empirically.

If consumers are sophisticated, we can equate perceived and true values for each redemption friction, which yields the empirically testable predictions in equations 12 and 13. These predictions have an intuitive interpretation. Equation 12 tells us that a rebate that is only redeemed with probability 0.5 should increase the buying probability half as much as a rebate that is remembered with certainty.[11] Equation 13 implies that if the elimination of hassle costs increases the expected value of the rebate by, say, 1 USD, the buying probability should increase by the same amount as if the firm had directly decreased the rebate value by 1 USD.

These general predictions only rely on reduced-form elasticities that can be empirically identified. Proposition 1 therefore provides us with a recipe for an experimental design. The ideal experiment observes how both the buying and the redemption probability respond to exogenous variation in rebate value, hassle costs, and inattention.

[11] With risk aversion, the buying probability of sophisticated consumers increases less in response to a rebate than under risk neutrality. However, for a small change in rebate value and prices, any reasonable degree of risk aversion should not substantially affect the predictions of Proposition 1.

4 Field Experiments

4.1 Cooperation with Online Retailer

I test the theoretical predictions in a series of natural field experiments in cooperation with a major online retailer for furniture and homeware in the European Union. The retailer operates online stores in the majority of European countries and sells a large variety of products. The main experiment presented in this paper is implemented among consumers in Germany in 2020. Two pilot studies with smaller samples and fewer treatments precede this main experiment.

4.2 Pilot Studies

The first pilot was implemented in the United Kingdom in 2018 with a sample size of 19,811 consumers. The second pilot was implemented with 52,302 consumers in Germany in 2019. Even though the pilot studies involve fewer treatments and more specific target groups, both studies replicate the findings from the main experiment, qualitatively. A description of the pilot studies and the results are presented in Appendix B.

All pilot studies have been pre-registered at the AEA RCT Registry.

4.3 Ethics

All interventions designed for the purpose of this study intend to make consumers better off relative to the standard rebate policy of the firm. Treatments either increase attention toward rebate redemption or intend to eliminate all redemption barriers simultaneously. The firm typically offers rebates that need to be actively redeemed during the checkout in the web shop. They also run regular pricing experiments in which a control group that does not receive a discount typically exists.

It is important to highlight that the cooperating firm does not use any uncommon promotion practices, but rather follows the standard pricing policies in the industry. Claimable rebates are the predominant form of price promotions in online shopping and can be found in virtually every major web shop.

4.4 Design

Figure 1 illustrates the experimental design. Upon visiting the website of the shop, the subject is randomized into one of eight experimental cells with equal probability. Subjects are individually identified and remain in the same experimental condition on follow-up visits.[12] The experimental cells can be categorized into four main groups: A, B, C, and D. Subjects in group A receive an automatically applied discount of 10% on all products. In group B, subjects receive the same rebate

[12]The firm uses the HTTP cookie to individually identify subjects.

Figure 1: Experimental Design

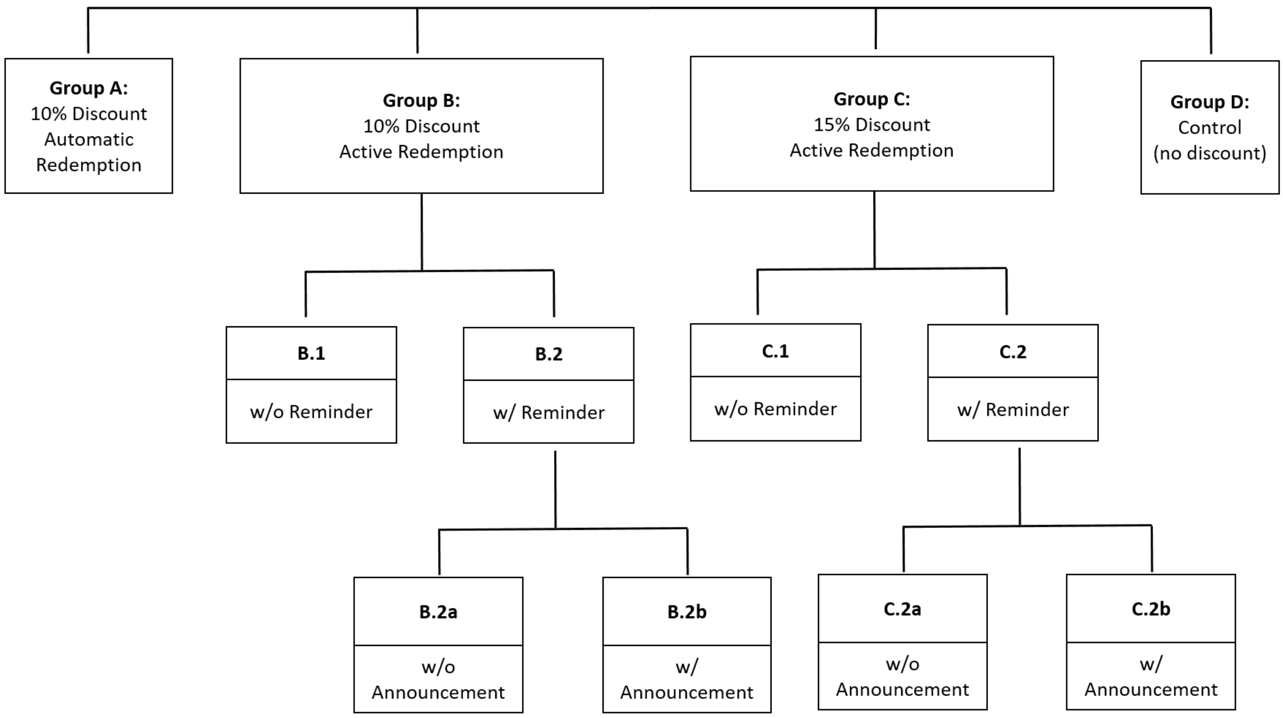

Note: This figure illustrates the experimental design. Subjects are randomized into one of eight groups with equal probability upon visiting the website.

value but only if they actively claim the rebate during checkout. Group C receives a larger discount value of 15% but is otherwise identical to group B. Groups B and C involve several subgroups in which I randomize a reminder and an information treatment as described below. Group D serves as the control group. Subjects are not offered a rebate.

Group A (automatic rebate redemption, 10% Discount): In group A, the banner in Figure 2a is displayed at the top of the browser, saying *"Only for a short period of time: 10% off everything."* If subjects click on the information icon next to this text, they are informed that the rebate will be automatically applied at the checkout (see Figure 2b). Subjects can then browse through the online store and add products to their shopping basket. Once they click on the shopping basket, they see the checkout page presented in Figure 4. On the checkout page, another banner is displayed telling subjects the discount has been automatically applied. At the center left of the page is a field into which the rebate code has already been entered. An additional small pop-up box above the rebate-code field tells subjects that the code has already been applied. The actual rebate code and some product descriptions have been censored with black bars to protect the anonymity of the company.

Note that even though the discount is automatically applied, the discount can still be invalidated by actively deleting the rebate code from the field. This may happen by accident or due to technical issues related to the device or browser of the subject. In addition, some tax deduction motives may encourage business customers to refrain from receiving the rebate. The redemption rate will therefore not exactly equal 100%. However, because subjects are randomly assigned to the groups, idiosyncrasies reducing the rebate redemption rate below 100% are also present in the other experimental conditions, and therefore neither threaten the internal validity of the experiment nor the identification of the structural parameters.

Group B (active rebate redemption, 10% discount): In this group, subjects are also offered 10% off all products but need to actively claim the rebate. The banner displayed at the top of the browser is shown in Figure 2c and includes the same text as the banner in group A. However, instead of the information icon, a hyperlink says, *"Go to rebate"*. If subjects click on the hyperlink, they are forwarded to another subpage of the website shown in Figure 3. The text on the subpage informs consumers that they need to copy the rebate code shown on the page and enter it into the respective field during the checkout. Subjects only get 10% off their purchase value if they read the text to find the rebate code, copy it onto their clipboard, and remember to paste it into the rebate redemption field on the checkout page.

Subjects in group B can be divided into the following three experimental subgroups in which I vary a reminder and an information treatment during checkout.

Group B.1 (without reminder): In this group, subjects are not reminded to redeem the rebate during the checkout process. An example of the checkout page is shown in Figure 5a and does not involve any reference to the rebate. Subjects see the products they added to the shopping card, the quantity of each product, and the total purchase price. To redeem the rebate, subjects must paste the rebate code into the respective field that is located below the list of products and on the left side of the page. Once they paste in the code, they need to click on the button saying "apply" right next to the field in order to redeem it. The total purchase price is then reduced by the size of the rebate.

Group B.2a (with reminder): Subjects receive the same treatment as in group B.1 but are offered an additional reminder during the checkout process. Figure 5b shows an additional banner that tells subjects to not forget rebate redemption. In addition, a pop-up box just above the rebate redemption field highlights the rebate redemption field and tells subjects to enter the rebate code. These two reminders are designed to be salient in order to best capture subjects' attention.

Group B.2b (with reminder and announcement of reminder): This group features an additional information treatment that explicitly announces the reminder at the start of the website visit.

Figure 2: Examples of Rebate Banners

(a) Banner in Group A: 10% Rebate, Automatic Redemption

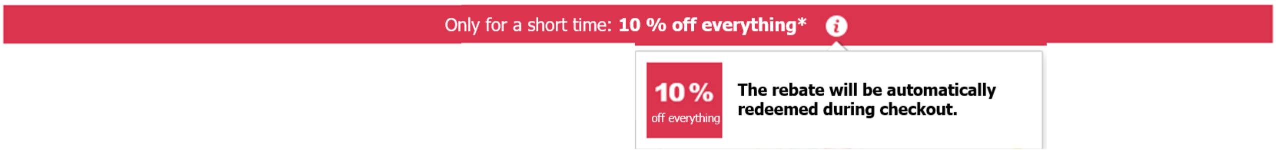

(b) Banner in Group A after click on information icon

(c) Banner in Groups B.1 and B.2a: 10% Rebate, Active Redemption

(d) Banner in Groups B.2b: 10% Rebate, Active Redemption with Announcement of Reminder

Note: This figure shows an English translation of the banners displayed in the experimental groups A, B.1, B2.a, and B2.b.

Figure 3: Subpage with Rebate Code in Group B.1, B.2a and B.2b

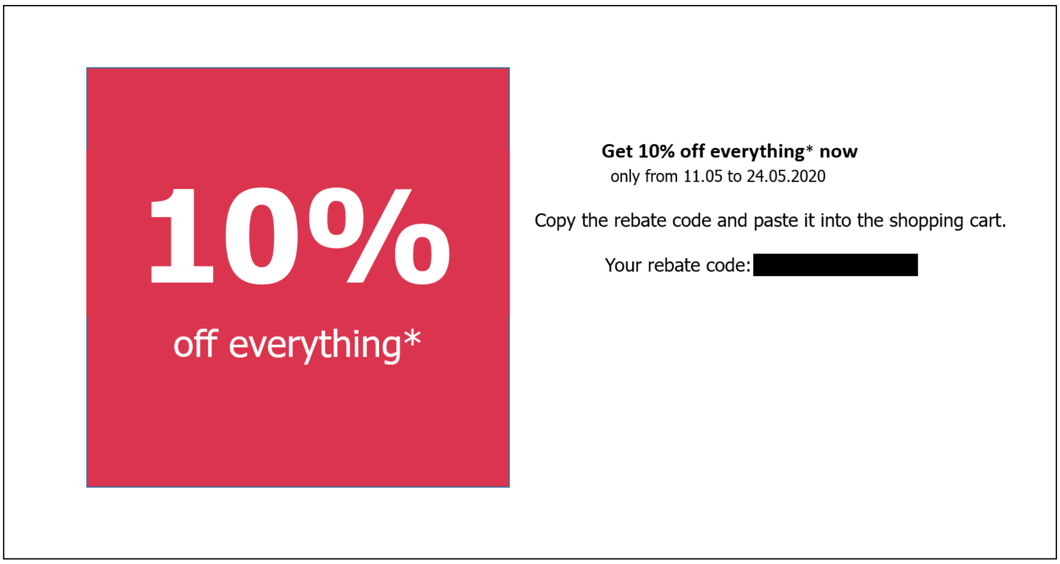

Note: This figure shows an English translation of the subpage showing the rebate code in experimental groups B.1, B2.a, and B2.b.

Note that in group B.2a, subjects receive a reminder during the checkout but the initial banner is the same as in group B.1. Therefore, although subjects are fully attentive during the checkout, they might not anticipate the reminder, simply because the firm has not announced it. To shed more light on the role of subjects' beliefs at the start of the website visit, I randomize the additional information treatment shown in Figure 2d. The banner is accompanied by a pop-up box telling subjects that they will be reminded to redeem the rebate during the checkout process.

The information treatment should have two effects on subjects' beliefs. First, it ensures subjects correctly anticipate that they will be attentive during checkout. This point is important because only then does the buying probability reflect behavior under fully anticipated attention. Second, a more nuanced effect is that the treatment also implicitly informs subjects that the rebate is not automatically redeemed but needs to be actively claimed.

Therefore, whether the information increases or decreases the buying probability relative to group B.1a is ambiguous. If subjects' initial beliefs in group B.1a are that the rebate needs to be actively claimed and that they might forget about the redemption, informing them that they will eventually receive a reminder should increase the buying probability. On the other hand, if subjects initially believe the rebate is automatically redeemed, the announcement of the reminder implicitly informs them that the rebate needs to be actively claimed and that rebate redemption involves hassle costs. In this case, the announcement is expected to reduce the buying probability relative to group B1.a. Thus, the effect on the buying probability of the announcement crucially depends on consumers' belief about the form of rebate redemption in group B1.a.

Group C (active rebate redemption, 15% discount): This group is structured in the same way as group B, but subjects receive a 15% instead of a 10% discount. Groups C.1, C.2a and C.2b are analogous to groups B.1, B.2a, and B.2b, respectively. The figures showing the treatments can be found in Appendix E.

Group D (no discount): No banner is displayed at the top of the browser and subjects are not offered a rebate. They simply see the status quo of the website without any price promotions.

Figure 4: Checkout Page in Group A: Automatic Redemption

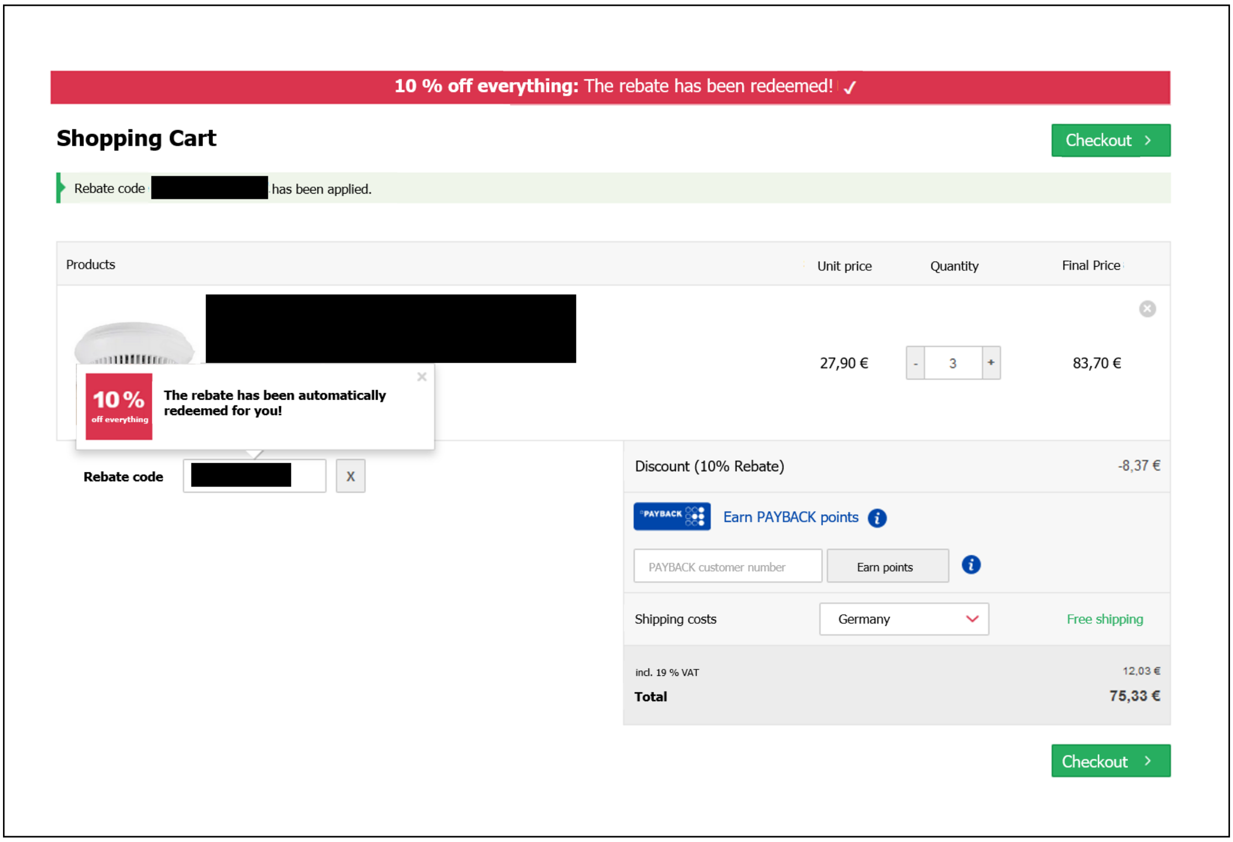

Note: This figure shows an English translation of the checkout page in experimental group A.

Figure 5: Checkout Pages in B-Groups

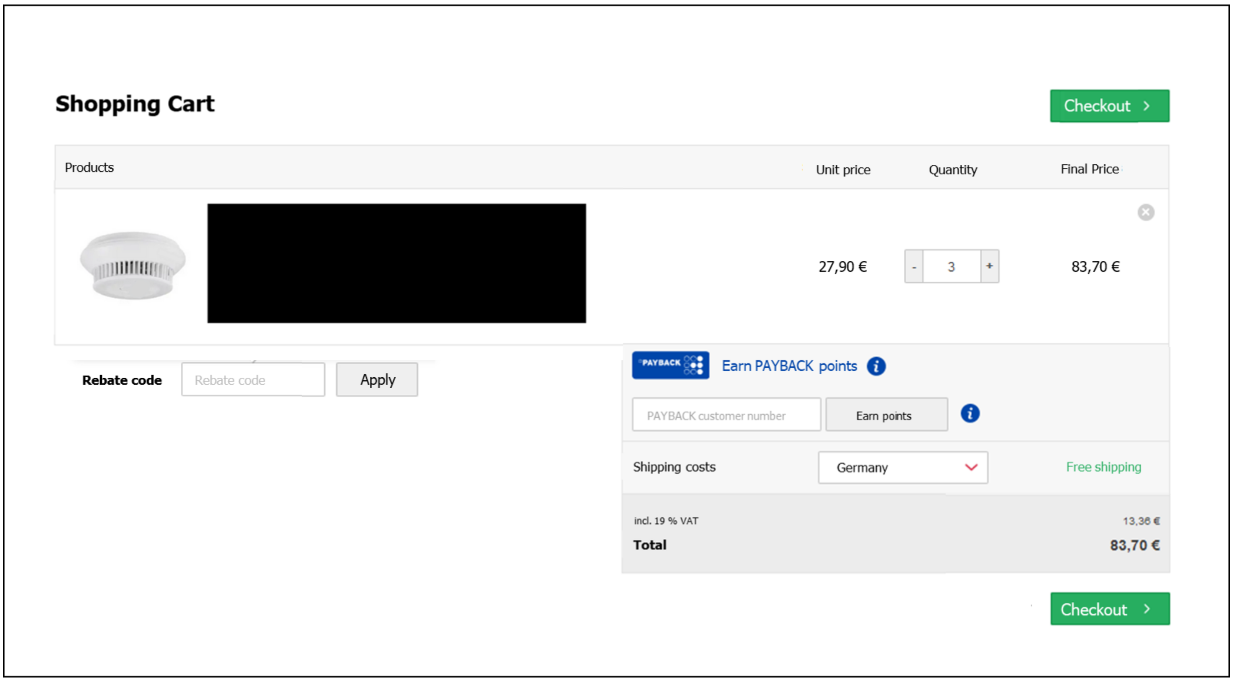

(a) Checkout Page in Group B.1: No Reminder

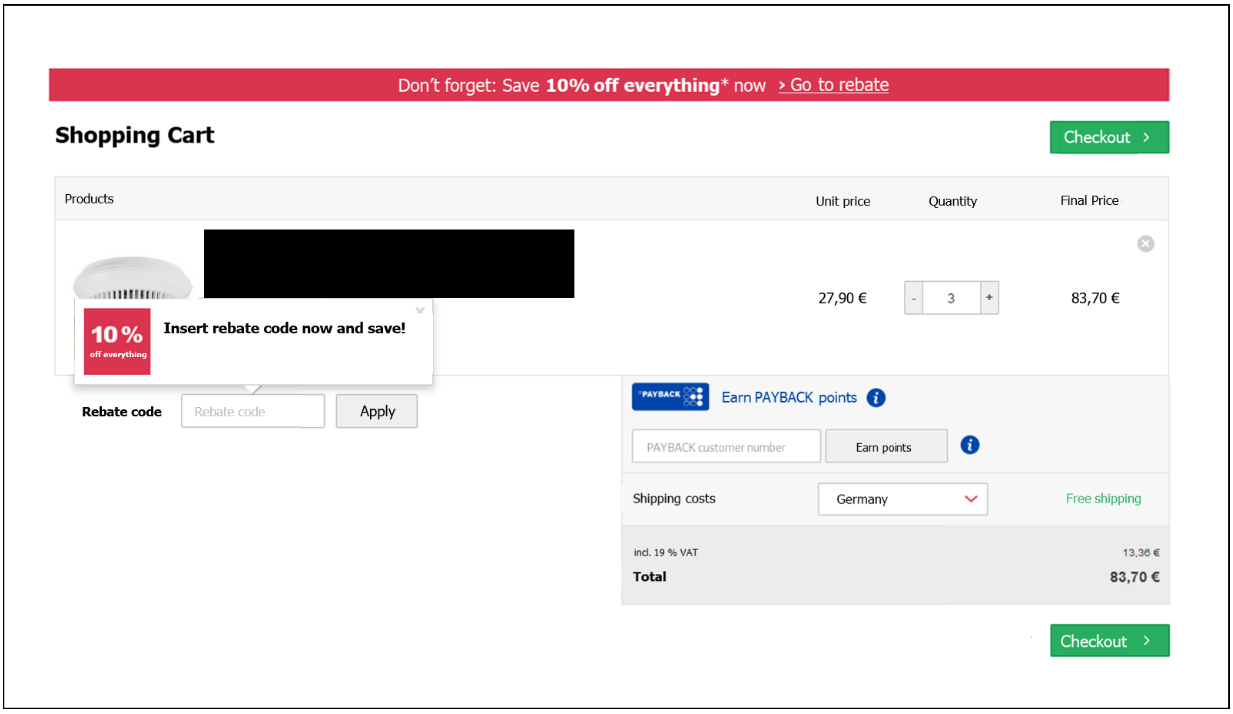

(b) Checkout Page in Group B.2a and B.2b: With Reminder

Note: This figure shows an English translation of the checkout pages in experimental groups B.1, B.2a and B.2b.

4.5 Identification of Consumer Sophistication

The experimental design allows for both a qualitative test of consumer sophistication and a structural quantification of the degree of sophistication. The qualitative test has the advantage that it does not rely on the structure imposed by the theoretical model in Section 3 and is more general. The downside is that it only allows us to conclude whether consumers exhibit *some* degree of sophistication, but not how pronounced it is. To quantify welfare implications and guide policy-making, we need to identify the deep primitives underlying consumer choices. The experiment is, therefore, designed to identify the theoretically derived sufficient statistics. In the following, I connect the experimental design to the qualitative and structural tests of sophistication.

4.5.1 A Reduced-Form Test of Sophistication

A simple comparison of differences in the buying probability between rebates with and without redemption barriers is sufficient to conclude whether consumers are fully or (at least) partially naive. *Any* decrease in the buying probability as a response to the introduction of redemption barriers is evidence of some degree of consumer awareness about the lower redemption probability. Again, the intuition is that the consumer's valuation of a rebate depends not only on its absolute monetary value but also on the likelihood of being redeemed. Thus, as we exogenously reduce this likelihood, we should also observe fewer consumers be attracted by the rebate. By contrast, full naiveté implies consumers' buying probability is the same for an automatically applied discount as for a claimable rebate with redemption barriers.

If consumers anticipate that they might forget about rebate redemption, the buying probability must be higher in group B.2b (C.2b) than in group B.1 (C.1). Similarly, if they perceive that hassle costs might affect their rebate redemption choice, the buying probability must be higher in group A than in group B.1. We can even test whether sophistication increases with stakes by comparing whether the buying probability responds more strongly to the reminder treatment with a 15% than to the treatment with a 10% rebate.

Moreover, we are able to quantify which redemption friction consumers perceive themselves as being more affected by. Imagine, for instance, that the buying probability in group B.2b is only slightly smaller than the buying probability in group A, but it is much smaller in group B.1 than in group B.2b. This observation implies consumers perceive their inattention to be a more severe redemption friction than their effort required for redemption.

4.5.2 Identification of Structural Parameters

Before discussing how the experiment identifies the model parameters, first note a slight difference exists in the type of rebate that is offered in the experiment and the rebate modeled in the theory section. The theoretical model involves a lump-sum rebate, whereas in the experiment, the company offers an ad valorem rebate whose absolute monetary value depends on the purchase value. I use a lump-sum rebate in the main part of the paper for clarity because it is slightly less complicated to model. However, I show in Appendix A.3 that an extended model with an ad valorem rebate yields the exact same proposition as a model with a lump sum rebate. All structural parameters can be identified by aggregate demand responses to an ad valorem rebate, as well, but s from the theory section translates to $s = t \times y$ in the experiment, where $t \in \{10\%, 15\%\}$ is the ad valorem rebate and y is the average purchase value in the treatment group that receives the automatically applied discount (i.e., group A).[13]

Perceived inattention, $\hat{\theta}$, can be identified by comparing the difference in the buying probability between the 10%- and 15%-rebate groups with and without inattention. Specifically, a linear approximation of the numerator in equation 8 is identified by comparing group B.1 with C.1, whereas for the denominator, we need to compare group B.2b with C.2b. To identify the buying probability under fully anticipated attention, that is, under $\hat{\theta} = 1$, we need to use the empirical moments from the treatment groups in which the reminder is explicitly announced. Otherwise, it would be unclear whether the buying probability fully incorporates the anticipation of the reminder.

We can identify perceived hassle costs by analyzing how the buying probability changes as we move from an automatically applied (hassle-free) discount to a rebate that needs to be actively claimed but that consumers remember during check-out. This change is the difference in the buying probability between group A and group B.2b. Scaling this treatment effect by the effect of a 1 EUR increase in the rebate value on the buying probability yields an approximation of $\hat{c}$, as shown in equation 9.

Identifying true inattention and true hassle costs is more complicated and crucially depends on the treatment effects on the extensive margin. An empirical challenge with the identification of the true redemption frictions is that I only observe the redemption probability for individuals who have self-selected into the pool of buyers. This results in a classical sample selection problem in which intensive-margin choices are only observed conditional on a particular extensive-margin decision. In Section 5, I address this selection issue empirically by using a sample selection model with an arguably credible exclusion restriction.

For expositional purposes, assume for a moment that the treatments do not cause systematic selection on the extensive margin. In this case, differences in the redemption probability have a causal interpretation and also reflect average differences for the entire population of website visitors. True inattention, θ, can be identified by the redemption probability with and without a reminder. Since I observe these moments at two different rebate values, the design overidentifies inattention. Inattention could be identified by taking the ratio of the empirical intensive-margin moments in groups B.1 and B.2b, or by the moments in groups C.1 and C.2b. As I show in the appendix, a third

[13] For the estimation, I use the median instead of the mean to adjust for outliers in the revenue distribution. The median purchase value in group A is 96 EUR.

way of identifying inattention is through a comparison of redemption elasticities with respect to the rebate value. Since the system is, therefore, overidentified, I approximate the efficient weighting matrix by a two-step GMM estimator as described in Section 6.

To identify true hassle costs, c, we first estimate the average treatment effect of reducing hassle costs to zero on the redemption probability, that is, the difference between group A and group B.2b. We then scale this effect by the treatment effect of increasing the rebate value by 1 EUR under full attention, as can be identified from a comparison between group B.2b and group C.2b.

4.6 Sample

I observe a total of 816,662 website visits by 601,471 individually identified subjects. Table 1 reports summary statistics for each of the eight experimental groups. Each group consists of around 75,000 subjects who visit the website using one of three possible devices. Approximately 35% visit the website using a desktop, 56% use a mobile phone, and 9% use a tablet. These fractions are balanced across all experimental groups and provide confidence of successful randomization.

The four variables at the bottom of the table do not need to balance, because they are potentially endogenous to the treatment variation. The average subject in the control group makes around 1.35 website visits. As the discount size increases, the number of website visits tends to slightly increase. The average buying probability in the control group equals 1.8% and is substantially larger in any of the treatment groups. The redemption probability is zero in the control group by construction and close to 90% in group A, in which the rebate is redeemed automatically. Redemption probabilities are dramatically lower once the rebate needs to be actively claimed as they fall to between 53% and 72% depending on the experimental condition.

In the next section, I estimate and discuss average treatment effects on a number of outcomes capturing consumer behavior and on firm profits.

Table 1: Summary Table

Variable	A 10%, automatic	B.1 10%, w/o reminder	B.2a 10%, w/ reminder	B.2b 10%, w/ reminder + announcement	C.1 15%, w/o reminder	C.2a 15%, w/ reminder	C.2b 15%, w/ reminder + announcement	D Control
Desktop user (Yes=1)	0.351	0.351	0.352	0.352	0.350	0.352	0.349	0.353
	(0.477)	(0.477)	(0.477)	(0.478)	(0.477)	(0.478)	(0.477)	(0.478)
Mobile phone user (Yes=1)	0.563	0.563	0.562	0.563	0.564	0.562	0.566	0.562
	(0.496)	(0.496)	(0.496)	(0.496)	(0.496)	(0.496)	(0.496)	(0.496)
Tablet user (Yes=1)	0.087	0.086	0.087	0.085	0.085	0.086	0.085	0.085
	(0.281)	(0.280)	(0.281)	(0.279)	(0.279)	(0.280)	(0.278)	(0.278)
Number of sessions	1.354	1.354	1.353	1.354	1.372	1.361	1.359	1.350
	(1.140)	(1.143)	(1.126)	(1.101)	(1.194)	(1.123)	(1.117)	(1.098)
Made purchase (Yes=1)	0.022	0.021	0.022	0.021	0.023	0.024	0.024	0.018
	(0.147)	(0.144)	(0.148)	(0.145)	(0.149)	(0.154)	(0.153)	(0.133)
Number of purchases	1.041	1.024	1.059	1.032	1.027	1.028	1.026	1.025
	(0.288)	(0.157)	(0.461)	(0.213)	(0.182)	(0.177)	(0.178)	(0.179)
Redeemed rebate (Yes=1)	0.877	0.530	0.630	0.681	0.569	0.691	0.723	0.000
	(0.328)	(0.499)	(0.483)	(0.466)	(0.495)	(0.462)	(0.448)	(0.000)
Observations	75,256	75,428	74,731	75,368	74,911	75,835	75,403	74,873

Note: This table presents the mean of observable variables in different treatment conditions. Standard deviations are reported in parentheses.

5 Reduced-Form Estimates

5.1 Consumer Behavior

Figure 6 plots the redemption and buying probability for each treatment condition. For expositional purposes, I first pool the two reminder groups with and without announcement of the reminder into one group for each rebate size.[14] The right ordinate depicts the buying probability, while the left ordinate shows the redemption probability conditional on buying.

Table 2 complements Figure 6 by showing estimated average differences in the buying and redemption probability across treatments. All coefficients are multiplied by 100 to ease readability. The empirical specification that produces the coefficients in column 1 for the buying probability is a linear probability model of the following form:

$$Buy_i = \alpha + \beta' \mathbf{T}_i + \xi_i, \tag{14}$$

where Buy_i is an indicator equal to 1 if subject i made at least one purchase, and 0 otherwise. The column vector $\mathbf{T}_i$ is a vector of indicators for the treatment groups. Each element of the vector takes the value of 1 if the subject is assigned to the respective treatment group, and 0 otherwise. The vector of average treatment effects is given by β. The buying probability in the control group equals α, and ξ_i is the residual.

In column 1, I pooled the reminder treatments into group B.2 and C.2, whereas column 2 reports results for each subgroup to capture the incremental effect of adding an explicit announcement to the reminder.

Estimated differences in the redemption probability are shown in column 3 and 4. The empirical specification is the following linear probability model:

$$Redeem_i = \phi + \tau' \mathbf{X}_i + \nu_i, \tag{15}$$

where $Reedem_i$ equals 1 if subject i redeemed the rebate, and 0 otherwise. The vector $\mathbf{X}_i$ indicates whether subject i is in one of the four treatment groups in which the rebate needs to be actively claimed. I exclude control group subjects from the estimation because they could not redeem the rebate by construction of the experimental design. The regression constant, ϕ, is the redemption probability in group A, in which the discount is automatically applied. ν_i denotes the unobserved residual of subject i. The coefficient τ measures average differences in the redemption probability across treatments. Column 3 reports pooled effects for the reminders and column 4 shows the effect for each subgroup. Since the rebate is only redeemed if the subject decides to buy, only the subsample of observations who ended up making a purchase is included in the estimation.

The control group has a buying probability of 1.8% and a rebate redemption probability of

[14] The disaggregated effects of the reminder treatments can be found in Figure E4 in the Appendix.

0. The treatment group that receives an automatically applied 10% discount has a substantially larger buying probability of 2.2%. The difference constitutes an increase of 22% relative to control and is highly statistically significant ($p < 0.01$). The redemption probability is close to 90% and implies that idiosyncratic noise reduces the redemption probability by around 10 percentage points. As previously discussed, the reason the redemption probability does not exactly equal 1 may be attributable to device- or browser-specific technical issues. Because these unobservables should be balanced by random treatment assignment, they only affect the *levels* of but not the *differences* between the redemption probabilities across experimental conditions.

Figure 6: Buying and Redemption Probabilities

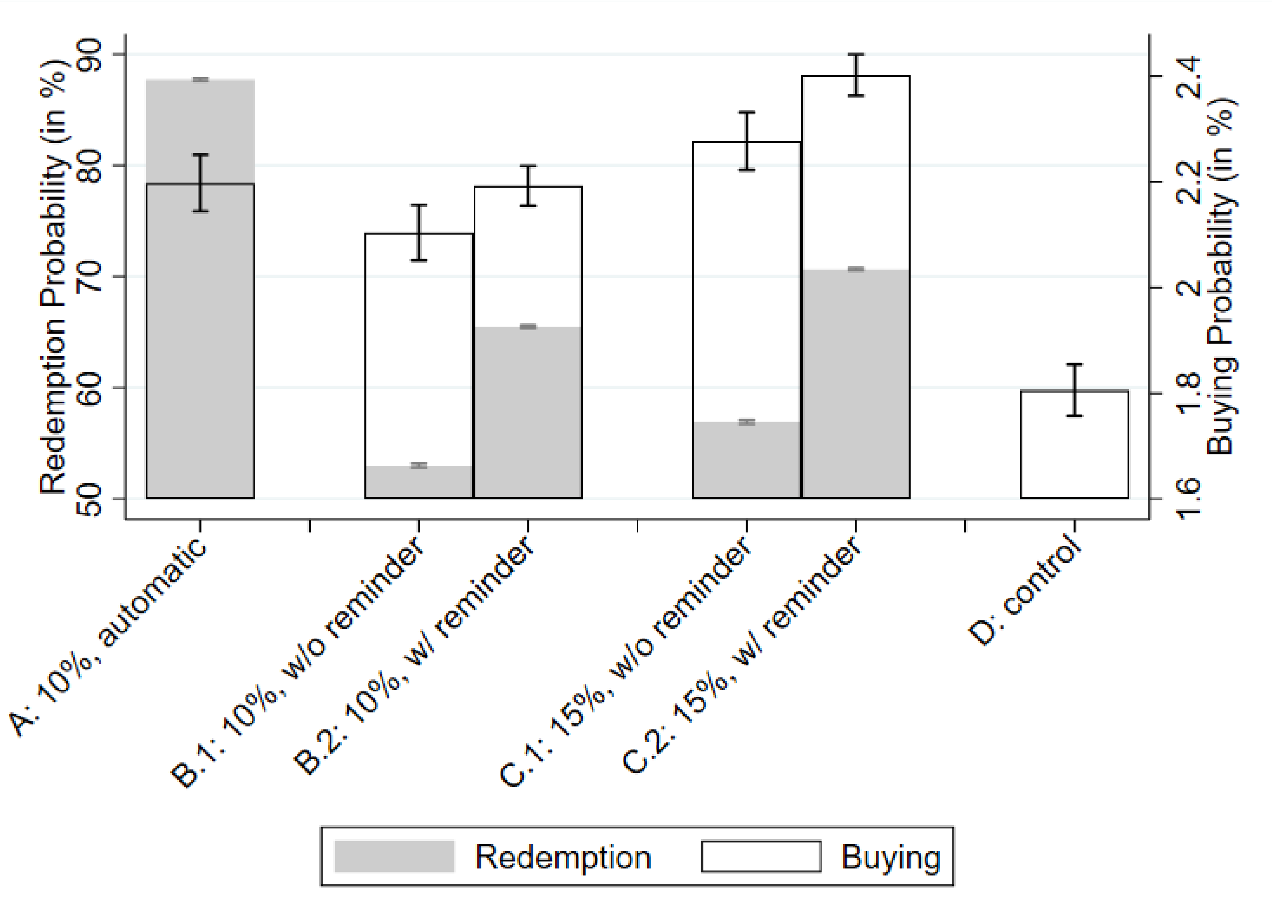

Note: This figure shows the buying probability for the entire sample and the redemption probability conditional on buying. The error bars represent standard errors. Reminder treatments with and without announcement are pooled for each discount size.

Table 2: Differences in Buying and Redemption Probabilities

	Buying Probability (in %)		Redemption Probability (in %)	
	(1)	(2)	(3)	(4)
A: 10%, automatic	0.392***	0.392***	87.727***	87.727***
	(0.075)	(0.075)	(0.844)	(0.844)
B.1: 10%, w/o reminder	0.298***	0.298***	-34.734***	-34.734***
	(0.073)	(0.073)	(1.519)	(1.519)
B.2: 10%, w/ reminder	0.387***		-22.245***	
	(0.064)		(1.202)	
B.2a: 10%, w/ reminder		0.432***		-24.748***
		(0.076)		(1.500)
B.2b: 10%, w/ reminder+announce		0.342***		-19.660***
		(0.074)		(1.444)
C.1: 15%, w/o reminder	0.472***	0.472***	-30.810***	-30.810***
	(0.075)	(0.075)	(1.477)	(1.478)
C.2: 15%, w/ reminder	0.597***		-17.069***	
	(0.064)		(1.141)	
C.2a: 15%, w/ reminder		0.613***		-18.643***
		(0.076)		(1.384)
C.2b: 15%, w/ reminder+announce		0.580***		-15.464***
		(0.076)		(1.364)
D: Control	1.806***	1.806***		
	(0.050)	(0.050)		
Regression constant	D	D	A	A
N	601,804	601,804	11,872	11,872

Note: The table reports average treatment effects from the OLS regressions specified in equations 14 and 15. Robust standard errors are in parentheses and clustered on the subject level. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

Inattention

The introduction of redemption frictions by requiring consumers to actively claim the rebate is associated with a sharp and statistically highly significant ($p < 0.01$) drop in the redemption probability of 35 percentage points (group B.1). If the firm offers a reminder at checkout, the redemption probability only falls by 25 percentage points. This result is consistent with the notion that cognitive inattention accounts for a large share of redemption frictions: simply reminding consumers of redemption during checkout is associated with an increase in redemption rates by 10 percentage points. The difference does not necessarily constitute the causal treatment effect, because the reminder may also affect the type of consumers sorting into the pool of buyers. I address the concern of sample selection in more detail in Section 5.3.

In fact, a comparison of the buying probability with and without a reminder reveals a remarkable degree of consumer anticipation on the extensive margin. The movement in the buying probabilities in groups B.1 and B.2 in Figure 6 mirrors the movement in the redemption probabilities. Making the rebate more attractive by offering a reminder causes an increase in the buying probability of 0.89 percentage points—a large relative increase of 30%. The treatment effect is direct evidence of consumer sophistication: consumers are aware they might forget to redeem the rebate and reduce their buying probability relative to a scenario in which they will be attentive due to a reminder. The same behavioral pattern is replicated for the treatment groups that receive a 15% discount. The reminder causes an increase in the buying probability of 1.25 percentage points—a relative treatment effect of 26%—and is associated with an increase in the redemption probability of around 13.7 percentage points.[15]

Overall, the results provide us with strong reduced-form evidence that consumers anticipate their inattention when responding to a price reduction in the form of a claimable rebate. Note that this result also highlights the importance of observing the buying probability—a behavioral margin that is typically not observed with transaction-level data in many observational studies. Studies that do not observe the buying probability frequently make the implicit assumption that the treatments have no extensive-margin effects. This assumption implicitly amounts to assuming $\hat{\theta} = 1$ and $\hat{c} = 0$, and thereby overestimates the distortions arising from consumer mistakes.

Hassle Costs

The effect of hassle costs on extensive-margin behavior is identified by the difference in the buying probability between group A and group B.2. Hassle costs account for a relatively small drop in the buying probability of 0.05 percentage points. This effect is statistically indistinguishable from zero.

If consumers fully anticipate hassle costs, we would also expect no substantial difference in redemption rates between these two groups. If sophisticated consumers do not reduce their buying probability, the effort of actively claiming the rebate must not affect the probability of obtaining the discount. However, the active redemption requirement causes a large reduction in the redemption rate by approximately 22 percentage points in comparison to the automatically applied discount, which is a relative decrease of 25%. Because we see no indication of systematic sorting on the extensive margin between these two groups, the difference on the intensive margin should be the causal treatment effect of hassle costs on redemption rates.

The reduction in the redemption probability caused by hassle costs is more than twice as large

[15] Another remarkable result is that the absolute effect of adding a reminder does not depend on the stakes of the discount. The addition of the reminder boosts the buying probability by an additional 0.14 percentage points both for the 15% and the 10% discount. A model of rational inattention would typically imply that the reminder is less effective as the stakes increase because consumers already pay more attention at higher stakes.

as the reduction associated with removing the reminder. Differences in the buying probability would instead suggest that inattention is a more important barrier to rebate redemption than hassle costs. These results are suggestive evidence of a perverse correlation between perceived and true redemption barriers: consumers are less sophisticated about redemption frictions that cause more disutility. This result also replicates in both pilot studies.

Announcement of the Reminder

Next, I analyze the effect of the information treatment on consumer behavior. As previously discussed, explicitly announcing the reminder to consumers may be important in order to ensure they are aware of the reminder at the start of their shopping experience. As documented in column 2 of Table 2, the announcement reduces the treatment effect of the reminder by 0.09 and 0.03 percentage points for the 10% and 15% discount, respectively.

Recall that the effect of announcing the reminder on the buying probability is governed by the consumers' prior beliefs regarding the type of rebate the firm is offering. If consumers believe in the absence of the announcement that the rebate needs to be actively claimed, informing them about the reminder should increase the buying probability because consumers learn that they will not be inattentive. On the other hand, if consumers' belief is that the rebate is automatically redeemed, the announcement also implicitly tells them that the rebate needs to be actively claimed. In this case, we would expect a negative effect on the buying probability because consumers learn rebate redemption requires effort. Thus, the effect of announcing the rebate is ambiguous and depends on the beliefs of consumers in the absence of the announcement. The reported interaction effects imply that the latter effect dominates and that at least some consumers in group B.2a think the rebate is automatically redeemed, such that the information treatment is negative news.

An interesting observation is that even though the announcement tends to decrease the buying probability, it is associated with a considerable increase in redemption rates by 5.1 and 3.2 percentage points for the 10% and 15% rebate, respectively. This finding is consistent with the notion that subjects who select into the pool of buyers based on the provided information have lower hassle costs. On the extensive margin, the effect of the announcement moves behavior further away from the naive benchmark: the buying probability decreases relative to the buying probability in group A. This finding indicates that a larger share of the true hassle costs is taken into account at the extensive margin.

The incremental effect of the announcement also tells us that extensive-margin behavior under fully anticipated attention ($\hat{\theta} = 1$) is identified by observations in group B.2b rather than observations in B.2a, because the latter group involves subjects with incorrect expectations regarding the details of the rebate. To further investigate whether the effect of the announcement operates through a change in information, I analyze differences in the browsing behavior of subjects. Figure 7 plots the probability of visiting the rebate page at least once that involves the necessary information about

Figure 7: Probability to Visit Rebate Page

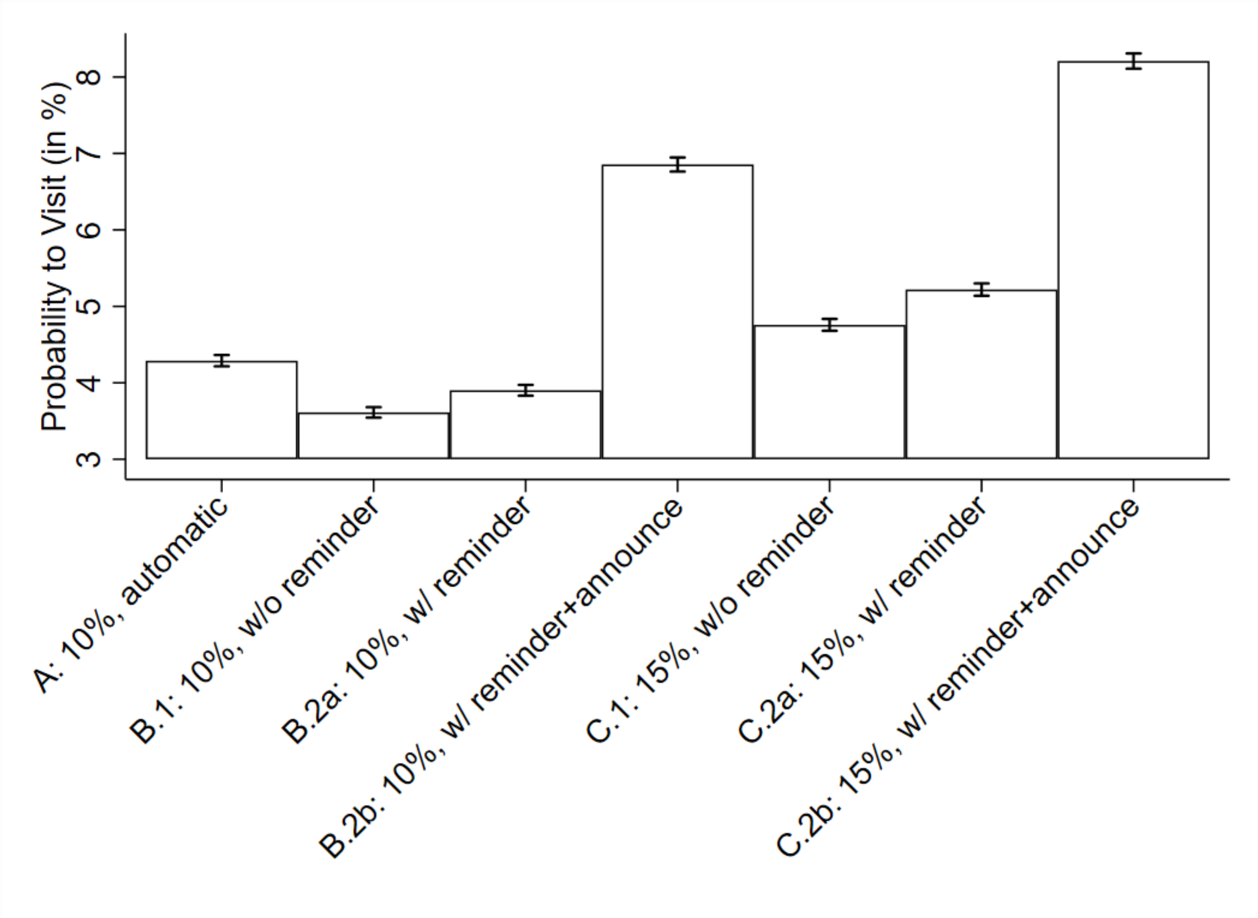

Note: This figure shows the probability of visiting the rebate page at least once. The error bars represent standard errors.

how to redeem the rebate, including the rebate code. For groups B.1, B.2a, and B.2b, this is the page shown in Figure 3. For groups C.1, C2.a, and C2.b, the page looks the same but the rebate size is larger. Group A can also visit a rebate page, but the page only informs them again that the rebate is automatically redeemed, and therefore involves no additional information in comparison to the information button already provided in the banner. Table 3 reports complementary results from a linear probability model regressing the outcome variable on the treatment vectors.

Announcing the reminder causes consumers to engage substantially more with the details of the rebate. Whereas around 3.9% of website visitors in group B.2a visit the rebate page, this number increases to 6.9% due to the announcement. Consumers are also more likely to engage with the rebate as the stakes increase. A comparison between group B.1 and C.1 shows the probability of visiting the rebate page increases by 0.54 percentage points as the firm raises the discount value by another 5 percentage points. The announcement also increases the probability of visiting the rebate page at the 15% discount. The probability increases from 5.23% to 8.21% in response to the announcement. The absolute treatment effect of the announcement is highly statistically significant

($p < 0.01$) at both rebate sizes and remarkably similar in absolute size across stakes (around +2.9 percentage points for both rebate values).

Subjects who receive the announcement are substantially more likely to engage and understand the details of the rebate.[16] Extensive-margin responses are accordingly more informed for subjects in group B.2b than for those in group B.2a. One implication is that only the empirical moments in group B.2b should be used for the structural estimation of behavior under fully anticipated attention.

Table 3: Probability to Visit Rebate Page

	Probability to Visit Rebate Page (in %)	
	(1)	(2)
B.1: 10%, w/o reminder	-0.693***	-0.693***
	(0.100)	(0.100)
B.2: 10%, w/ reminder	1.081***	
	(0.094)	
B.2a: 10%, w/ reminder		-0.396***
		(0.102)
B.2b: 10%, w/ reminder+announce		2.546***
		(0.118)
C.1: 15%, w/o reminder	0.465***	0.465***
	(0.107)	(0.107)
C.2: 15%, w/ reminder	2.423***	
	(0.098)	
C.2a: 15%, w/ reminder		0.938***
		(0.110)
C.2b: 15%, w/ reminder+announce		3.915***
		(0.124)
A: 10%, automatic (constant)	4.296***	4.296***
	(0.074)	(0.074)
N	526,919	526,919

Note: The table reports average treatment effects from an OLS regression of the probability of visiting the rebate page on the treatment indicators. Robust standard errors are in parentheses and clustered at the subject level. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

[16]This interpretation is also supported by complementary results shown in Figure E5 in the Appendix. The announcement substantially increases the probability of clicking on the checkout page and viewing the reminder.

5.2 Heterogeneity

As documented in the pre-analysis plan, I analyze heterogeneous treatment effects for different income groups. In particular, I obtain publicly available state-level income data and merge it to each subject based on the state from which they are visiting the website. Since around 3% of subjects use browser settings that hide their origin, the analysis is restricted to 582,630 subjects.

I construct a dummy variable equal to one if a subject visits from a region with an income equal or above the sample median, and zero otherwise.[17] I then extend the regressions in equation 14 and 15 by adding the dummy variable and interacting it with the treatment variables. Results are documented in Table 4.

The first important observation is that there is no statistically significant heterogeneity in the demand response to an automatically applied price reduction. While low-income consumers are 0.065 percentage points more likely to respond to the price reduction than high-income consumers, this difference is both economically and statistically small. The redemption rates are 3.5 percentage points lower for high-income consumers, suggesting that idiosyncrasies affecting redemption are slightly more relevant for high income consumers.

Based on these results, there is no clear rationale to price-discriminate between income groups, as price elasticities are roughly homogeneous. For classical price discrimination, we would expect to observe correlated heterogeneity in redemption probabilities and preferences, such that the firm could charge higher prices from consumers with lower price elasticities. This would require that less price-elastic consumers face larger redemption barriers. Conversely, rebates cannot price discriminate between preference types if redemption barriers and preferences are independent. In this case, all preference types have the same average redemption probability. Understanding heterogeneity in treatment effects is, therefore, crucial to separate the role of price discrimination from naiveté-based exploitation.

Turning to the redemption barriers, the addition of a reminder that is announced at the beginning of the website visit (B.2b) is associated with a 6.96 percentage points larger increase in the redemption rate for high- than for low-income consumers ($p < 0.05$). The demand response to this treatment is also 0.19 percentage points larger for high-income consumers. While the difference in the demand response is not statistically significant at conventional levels, the t-statistic of 1.25 is fairly large.

There is no evidence of substantial heterogeneity of the effect of the announced reminder for the 15% rebate (C.2b). However, the unannounced reminder has a 0.34 percentage point larger effect on demand for high-income consumers ($p < 0.05$). The difference in the redemption rate is only 2.95 percentage points larger and not statistically significant.

The differences in the behavioral patterns between the 10% and 15% rebate do not have an

[17] As noted in the pre-analysis plan, I calculate the median based on the income distribution of the *sample*, not based on the income distribution of the German population.

obvious interpretation. Since any heterogeneity analysis may involve false positives resulting from sampling variation, I am careful not to draw stark conclusions from these patterns.

Overall, there is relatively little robust evidence of significant heterogeneity in treatment effects based on income groups. Importantly, there is no clear correlation between price elasticities and redemption frictions that would provide a motive for classical price discrimination. Instead, rebates are more likely to be a profitable promotion because consumers are partially naive about their redemption probability.

It is important to highlight that there may be other idiosyncratic correlations between preferences and redemption probabilities that this analysis does not capture. To fully understand the role of heterogeneity for rebate promotions I would require knowledge of the joint distribution of price elasticities and redemption probabilities. Since this distribution obviously remains unobserved, I have to resort to observable characteristics of consumers. This approach, therefore, relies on the strong assumption that observables can be used to separate preference types.

Table 4: Heterogeneity by Income

	(1) Buying Probability $\times 100$	(2) Redemption Probability $\times 100$
A: 10%, automatic	0.429***	89.568***
	(0.103)	(1.059)
B.1: 10%, w/o reminder	0.283***	-35.956***
	(0.101)	(2.068)
B.2a: 10%, w/ reminder	0.414***	-24.105***
	(0.103)	(1.961)
B.2b: 10%, w/ reminder+announce	0.261***	-23.333***
	(0.101)	(2.005)
C.1: 15%, w/o reminder	0.497***	-33.209***
	(0.104)	(1.999)
C.2a: 15%, w/ reminder	0.443***	-20.166***
	(0.103)	(1.901)
C.2b: 15%, w/ reminder+announce	0.588***	-16.670***
	(0.105)	(1.819)
D: Control	1.802***	
	(0.069)	
Above median income	0.039	-3.489**
	(0.099)	(1.629)
$\times$ A: 10%, automatic	-0.065	
	(0.148)	
$\times$ B.1: 10%, w/o reminder	0.042	2.564
	(0.146)	(3.013)
$\times$ B.2a: 10%, w/ reminder	0.052	-0.704
	(0.149)	(2.887)
$\times$ B.2b: 10%, w/ reminder+announce	0.185	6.958**
	(0.147)	(2.856)
$\times$ C.1: 15%, w/o reminder	-0.002	4.463
	(0.150)	(2.915)
$\times$ C.2a: 15%, w/ reminder	0.335**	2.945
	(0.151)	(2.731)
$\times$ C.2b: 15%, w/ reminder+announce	0.008	1.687
	(0.151)	(2.691)
N	582,630	11,613

Note: The table reports treatment effects for consumers from regions with an income below and above the sample median. Robust standard errors are in parentheses and clustered on the subject level. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

5.3 Effect on Firm Profits

Are rebates with larger redemption frictions more profitable for the firm? If consumers are naive, profits must increase, because the number of transactions remains the same while the number of issued discounts decreases. If consumers are at least partially aware of redemption frictions, effects on profits are ambiguous. Increasing the difficulty of rebate redemption may still reduce redemption rates, but it also attracts fewer consumers. The effect on total profits therefore crucially depends on the degree of consumer sophistication.

The firm provides me with rich data on markups for each product offered in the online store. Based on the data, I calculate profits for each transaction I observe in the experiment. Total profits are simply calculated by multiplying average profits per transaction by the total number of transactions. Figure 8 plots the change in total profits in percent relative to the control group for each treatment condition.

The automatically applied rebate increases total profits by 4.23% relative to control. Introducing hassle costs (B.2a) increases the rate of return of the rebate to 15.37% and is the most profitable rebate policy among the 10% discounts. This constitutes a dramatic increase in the profitability of *263%* relative to the automatically applied rebate. The underlying mechanism is that the firm can reduce expenses from giving out discounts without affecting the buying probability, because consumers are naive about hassle costs.

A comparison with group B.2b confirms almost all of this increase in profits is driven by consumers not anticipating hassle costs. The effect of hassle costs on profits is only one third the size when consumers are explicitly informed at the start of their website visit that the rebate needs to be actively claimed. Profits decrease because the announcement deters some subjects from buying at the shop but is associated with a larger redemption probability.

Without the announcement of the reminder, behavior also suggests some path dependency: consumers first think the rebate is automatically applied, and then realize during the checkout that they have to actively claim the rebate. At this point, subjects have already invested in the shopping experience by searching for goods and putting suitable products into the shopping cart. As a consequence, some of them do *not* revoke their buying decision even though they would not have purchased at the shop had they known from the beginning that the rebate needed to be actively claimed.

A typical rebate that requires active redemption (group B.1) causes an 8.97% increase in profits. Note that this increase is larger than the effect of the rebate that involves an explicitly announced reminder (B.2b), even though the latter leads to a slightly larger buying probability. The reason is that rebate redemption rates are substantially lower in group B.1 than in B.2b.

Increasing the rebate value increases profits slightly, but the profit function appears to be concave in the rebate size on the empirically observed interval.

In conclusion, simple and costless changes in the rebate feature have dramatic affects on firm returns. These results rationalize firm practices referred to as buy baits or “sludges” (Thaler 2018) more generally, and provide empirical evidence for the theoretical literature on naiveté-based exploitation (e.g., Heidhues and Kőszegi 2017).

The positive effect of the reminder on profits is also interesting for the emerging literature on ”behavioral firms” that studies situations in which firms might fail to maximize profits (for an overview, see DellaVigna 2018). Recall that the status-quo rebate policy of the firm equals the rebate offered in treatment B.1. The data at hand suggests that the firm has not chosen the optimal rebate policy, because profits can be increased by adding an unannounced reminder during checkout. Since this finding comes from only one observation from the universe of existing firms, it does not prove that the average or modal firm chooses suboptimal rebate policies. However, it is suggestive evidence that managers do not fully take into account the degree of consumer sophistication when designing promotions.

Figure 8: Change in Total Profits

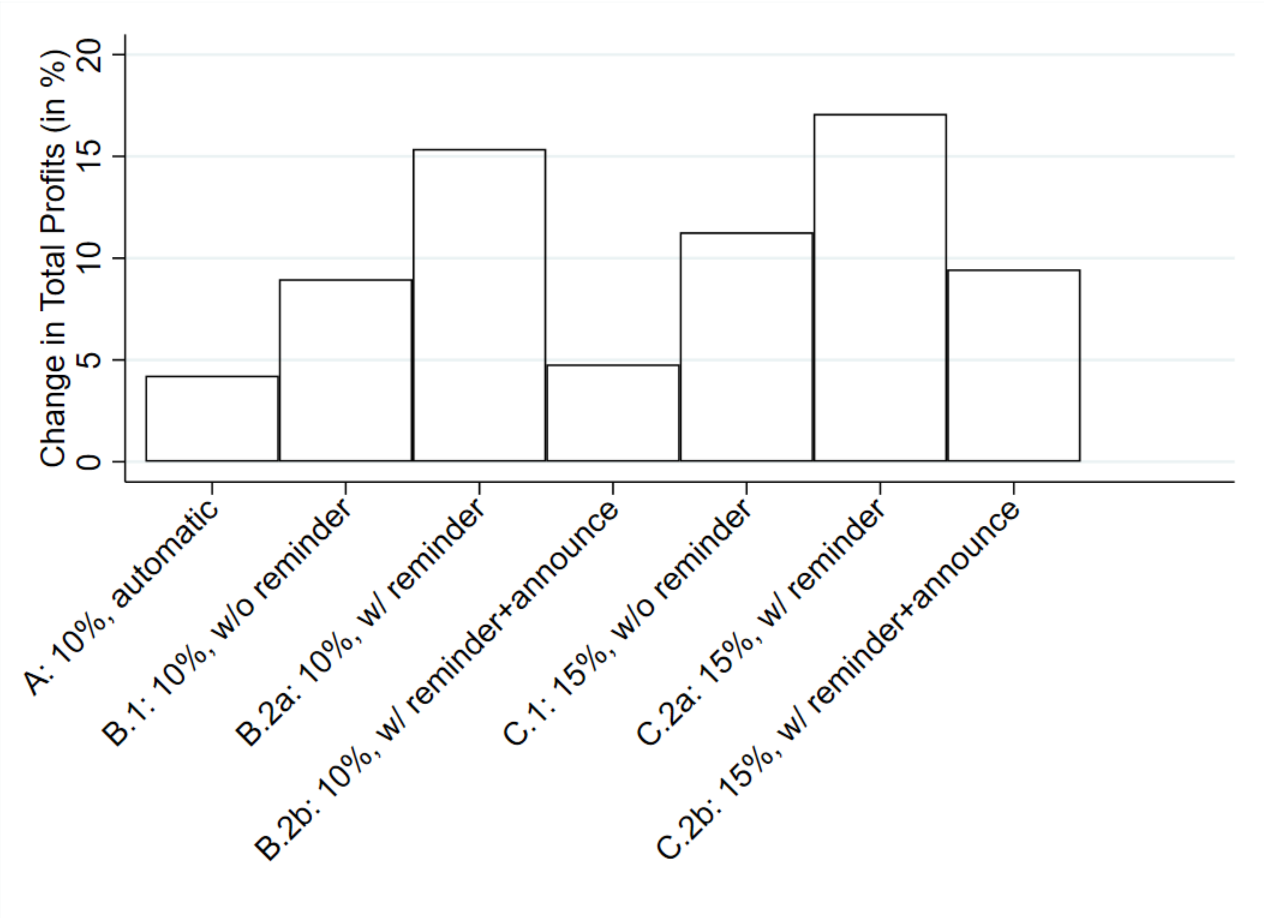

Note: This figure shows the change in total profits as a share of control group profits for each experimental group.

An important caveat of the previous analysis should be made explicit. The data only allows us to draw conclusions for short-term effects of the treatments on profits. This limitation is a potential concern because exploiting consumer biases may have negative long-term effects on the probability of returning to the shop. Even consumers who are naive in the short term might realize they are being exploited in the long term and, as a result, decide not to return to the store. In Appendix C, I investigate dynamic effects of redemption frictions on the probability of returning to the store. Although I find no significant negative effects on customer loyalty, the time period in which I observe the sample is relatively short and, as a result, follow-up purchases tend to be rare in the data.

However, although long-term effects on the repurchase rate may potentially be large, they do not seem sufficiently large to drive business out of the market. As previously mentioned, both the cooperating firm and the majority of large online retailers frequently offer these type of rebates, indicating the use of buy baits persists in a long-run equilibrium.

5.4 Addressing Incidental Truncation

Before moving to the estimation of the structural parameters, I address a potential concern related to systematic sorting of subjects on the extensive margin. Although the causal effect of the treatment on the buying probability is identified through random assignment, I only observe the redemption probability for those subjects who decide to buy. Therefore, the unconditional empirical distribution of redemption choices may not be the one I observe in the experiment. Because the treatments may affect the type of subjects that select into the pool of buyers, differences in the redemption probability may not be due only to the treatment effect, but also to unobservables that correlate with the decision to buy. This possibility is not a concern for the identification of the treatment effect of hassle costs on redemption rates, because there is no evidence that hassle costs affect the buying probability. However, the effect of the reminder on redemption rates may not be identified because consumers' extensive-margin choice responds to the reminder.

A large literature in econometrics has developed techniques to address selection bias resulting from "incidental truncation," building on the seminal work in Heckman (1976) and Heckman (1979). A consensus in the literature is that convincing identification in these models requires a credible exclusion restriction: a variable that does not directly affect the outcome of interest but affects whether subjects select into the sample.

I address the potential bias resulting from sample selection by estimating a fully parametric selection model with an arguably credible exclusion restriction: regional and temporal variation in sudden internet outages. The exclusion restriction requires that internet outages affect the redemption probability only indirectly through its effect on the buying probability. I view this assumption as a plausible one: an internet outage reduces the probability of redeeming the rebate because it cuts people off from access to the online shop, but not because of other channels.

Using this exclusion restriction, I estimate a selection model with normally distributed residuals and a binary intensive and extensive margin as first formulated by Van de Ven and Van Praag (1981). Monte Carlo simulations show that when these distributional assumptions are violated, the model still performs well in many cases as long as a valid exclusion restriction exists (Cook and Siddiqui 2020). I provide a detailed discussion about this model in Appendix D, where I elaborate on the dataset on internet outages, the construction of the exclusion restriction, and the estimation of the selection model.

A summary of the results in the Appendix is that the treatment effects are fairly similar in magnitude to the ones estimated from the two independent linear probability models presented previously. Internet outages have a large negative effect on the buying probability, consistent with intuition. Overall, there is little indication of systematic sorting on the extensive margin. However, the precision of the treatment coefficient drops substantially, which is a well-known problem resulting from high levels of collinearity in sample-selection models. Collinearity is a particular challenge in my application, because all treatments must be included in both the selection and the outcome equation, such that a high level of collinearity exists by construction. In upcoming versions of this paper, I plan to extend the current approach by including additional exclusion criteria and by using alternative estimators as additional robustness tests.

6 Structural Parameter Estimates

The reduced-form results tell us that consumers exhibit some degree of sophistication regarding their inattention, but are fully naive about their hassle costs. To quantify the *degree* of perceived and true redemption barriers, we require structural estimates of the underlying choice parameters. For this purpose, I estimate the sufficient statistics derived in Proposition 1.

To obtain values of the perceived redemption frictions, I estimate the structural parameters by substituting the regression coefficients in equation 14 with the structural parameters and then solve a set of moment conditions.

Denote the buying probability in the control group by β_D. Let the treatment effect on the buying probability by treatment $t \in \{A, B.1, B.2b, C.1, C.2b\}$ be denoted by β_t. In Appendix A.4, I show that we can rewrite these reduced-form treatment effects in terms of the structural parameters. This reformulation results in the following six moment conditions for behavior on the extensive margin:

$$\mathbb{E}\Bigg[\mathbf{I}_i\Bigg(Buy_i - \beta_A \times A_i - \beta_{B_1} \times B.1_i - \beta_{B_2.b} B.2b_i - \left(\frac{\hat{\theta}\Delta s}{\hat{c}}(\beta_A - \beta_{B.2b}) + \beta_{B.1}\right) \times C.1_i - \left(\frac{\Delta s}{\hat{c}}(\beta_A - \beta_{B.2b}) + \beta_{B.2b}\right) \times C.2b_i - \beta_D\Bigg)\Bigg] = 0, \quad (16)$$

where $\mathbf{I}_i = (A_i, B.1_i, B.2b_i, C.1_i, C.2b_i, D_i)$ is the 6×1 vector of instruments indicating the experimental group of subject i. The change in the rebate value is given by $\Delta s = 4.80$EUR.[18] Since the number of parameters to be estimated is also six, the model is exactly identified.

Next, true inattention and hassle costs are estimated using moments on the intensive margin. Again assuming independence between the residuals on the intensive and extensive margin, I can rewrite the reduced-form treatment effect parameters in terms of the underlying structural choice parameters, which yields the following five moment conditions:

$$\mathbb{E}\Bigg[\mathbf{J}_i\Bigg(Redeem_i - \tau_A - \left(\theta(\tau_A - c\frac{\partial R(s,1,c)}{\partial s}) - \tau_A\right) \times B.1_i + c\frac{\partial R(s,1,c)}{\partial s} \times B.2b_i$$
$$- \left(\theta\left[\tau_A + \frac{\partial R(s,1,c)}{\partial s}(\Delta s - c)\right] - \tau_A\right) \times C.1_i - \left(\frac{\partial R(s,1,c)}{\partial s}(\Delta s - c)\right) \times C.2b_i\Bigg)\Bigg] = 0, \tag{17}$$

with $\mathbf{J}_i = (A, B.1, B.2b, C.1, C.2b_i)$ denoting the vector of instruments, excluding the control group. With five moments and four parameters, the model is overidentified and I use a two-step GMM estimator to find the optimal weight matrix.

Table 5 reports the estimation results. The first column shows estimates from equations 16 and the second column from equations 17. Differences in the buying probability across treatments imply consumers' subjective probability of paying attention to rebate redemption is 73%. Their true probability of paying attention is 78% and is therefore remarkably close to consumers' expectations. Recall that Proposition 1 only requires these point estimates to be *approximately* equal. We can therefore conclude that consumers may well be sophisticated about their inattention, because their extensive- and intensive-margin responses are approximately consistent.

By contrast, consumers vastly underestimate the hassle costs of redeeming the rebate. Since the buying probability hardly responds to an introduction of hassle costs, the perceived hassle only equals 1 EUR. However, the strong drop in the redemption probability implies hassle costs are a significant redemption barrier equal to approximately 21.16 EUR. Although this number may seem large, note that the hassle of redeeming the rebate may come in various forms. For instance, the hassle of finding out where to find the rebate code and how to exactly redeem it may be substantial for certain consumers. Especially people with a low level of experience with online shopping may struggle to understand how to redeem the rebate and give up after some time spent engaging with the feature. Hassle costs therefore represent not only the process of copying and pasting the rebate code into the respective field, but also the time and effort required to understand how redemption is done and where to find the necessary rebate code. This process may be more challenging for

[18] The incremental rebate value is calculated by multiplying the additional five percentage points in the rebate by the purchase value of 96 EUR in group A. See Appendix A.3 for a formal proof.

subjects with lower levels of digital literacy, such as more senior citizens. In addition, people with larger opportunity costs of time might be more willing to forgo discounts when required to complete the technical redemption process themselves.

There may, however, be a potential bias that causes the estimated hassle costs to be too large, and that is attributable to the nature of the sufficient statistics approach. In particular, the estimated parameter is only a linear approximation of hassle costs. If demand is highly convex in rebate value as we move from a 10% to a 15% discount, then the presented estimate is too large and represents an upper bound of the true hassle costs.

Overall, the large difference between perceived and actual hassle costs yields the conclusion that the average consumer does not fully anticipate the struggle of redemption when deciding to make a purchase, and underestimates the hassle of rebate redemption by up to 20.16 EUR. This lack of sophistication results in a demand response that is larger than the response that would be optimal for consumers.

As discussed in Section 2, marketers from PMA suggest that an appropriate approach to understand redemption behavior is to calculate redemption rates as the number of redemptions divided by *incremental* (instead of total) sales. This approach would mitigate the issue that infra-marginal consumers simply do not redeem the rebate. Using the reduced-form results in Table 2, I calculate this ratio. Claimable rebates offered in group B.1 would cause 1,793 additional sales, but 6,292 of all buyers would redeem the rebate. The ratio suggested by PMA would therefore equal 351% and could yield the incorrect conclusion that all marginal consumers managed to redeem the rebate (plus a large share of inframarginal consumers). Instead, the theoretically founded sufficient statistics indicate a substantial share of marginal consumers failed to claim the rebate and, as a result, the demand response is excessively high.

The presented results provide a motive for consumer protection laws that limit the use of claimable rebates. Consumer protection agencies are typically concerned with situations in which consumers are not fully aware of the consequences of their choice. Both reduced-form and structural results point to a substantial misperception of the effort related to rebate redemption. This misperception lures consumers into making a purchase even though the hassle of obtaining the discount turns out to be too large to be worthwhile.

Table 5: Structural Estimates

	Extensive Margin	Intensive Margin
Inattention and Hassle Costs:		
$\hat{\theta}$	0.730 *	
	(0.390)	
$\hat{c}$ (in EUR)	1.002	
	(1.380)	
θ		0.784 ***
		(0.014)
c (in EUR)		21.156 ***
		(5.289)
Other Parameters:		
β_A	0.004 ***	
	(0.001)	
β_{B1}	0.003 ***	
	(0.001)	
β_{B2}	0.003 ***	
	(0.001)	
β_D	0.018 ***	
	(0.001)	
τ_A		0.877 ***
		(0.011)
$\frac{dR(s,1,c)}{ds}$		0.009 ***
		(0.003)
N	451,239	8,365

Note: The table reports estimation results from the GMM estimations specified in equations 16 and 17. Standard errors are in parentheses. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

7 Additional Mechanisms

I discuss a number of alternative mechanisms that may affect the interpretation of the data, as well as the identification of the structural parameters.

Risk preferences and loss aversion. Demand may be more elastic to price reductions than to rebates with the same expected value because consumers are averse to risk and losses. First, consumers' utility of income may be concave, which implies classical risk-averse preferences. Second, consumers may be loss-averse, meaning the disutility of losing a monetary amount significantly exceeds the utility of a monetary gain of equal size.

Loss aversion is arguably a more relevant factor in my empirical setting than risk aversion, because the rebate value is relatively small. For reasonable degrees of risk aversion, even risk-averse consumers should behave risk neutral over small gambles (Rabin 2000).

Both risk and loss aversion would have the same directional effect on the empirical estimates. In particular, introducing risk and loss aversion provides two additional reasons for why consumers should respond *even less* to the rebate than to the automatically-applied discount. The previously presented estimates imply that the demand response to a rebate is already excessively high under the assumption of risk neutrality and no loss aversion. If consumers are risk- and loss-averse, then the demand response of sophisticated consumers should be even lower than previously discussed. Thus, extending the model to capture an aversion to risk and losses would further strengthen the qualitative conclusion that the demand response to rebates is too large because consumers overestimate their redemption probability. As a consequence, the implication for consumer protection regulation would remain unchanged.

An important limitation is that the structural estimates would change. The presented estimates in the previous section would overstate the degree of consumer sophistication, because the model estimates a counterfactual demand response under full sophistication that is too large for risk- and loss-averse consumers.

Social preferences. Existing evidence in literature on social preferences indicates that subjects exhibit altruistic and reciprocal preferences (e.g., Fehr and Gächter 1998). Even sophisticated consumers who would benefit from rebate redemption might decide not to buy at the store, because they consider rebates an unfair marketing practice. Consumers may also receive direct disutility from a firm's attempt to exploit their own or other consumers' behavioral tendencies.

Introducing a distaste for exploitation to the model would have the same directional effect as the introduction of risk and loss aversion: it provides another reason for why consumers should respond less to a rebate than to an automatically-applied discount. This model extension would, therefore, not affect the qualitative conclusion that demand overreacts to rebates. However, it would increase the magnitude of that overreaction.

8 Conclusion

This paper studies behaviorally-motivated price discrimination and the role of consumer sophistication in the context of large-scale rebate promotions. I develop novel theoretical predictions of sophisticated behavior that can be empirically tested by observing a small set of aggregate demand elasticities. I then take this model to a series of natural field experiments and estimate consumer sophistication and its economic implications, using choices from hundreds of thousands of consumers.

I find that consumers exhibit a remarkable degree of sophistication regarding their inattention but are almost fully naive with respect to hassle costs. Claimable rebates cause an excessive increase in the buying probability that is above the benchmark response if consumers were sophisticated. Exploiting consumer naiveté is an impressive lever of profit and increases rebate returns by up to 260%.

Results have important implications for the policy debates around the regulatory framework that limits the use of rebate promotions in many countries across the world. I provide the first evidence from a natural field experiment that tests whether rebates harm consumers by causing them to make systematically distorted buying decisions. The evidence indicates the regulators' qualitative intuition is correct and that consumer protection laws may have large positive effects on consumer welfare.

A limitation of the study is that it can only draw narrow conclusions regarding supply-side responses to regulatory interventions. For instance, it is unclear whether a ban on claimable rebates would result in a long-run equilibrium in which firms do not offer price promotions at all instead of offering frictionless price reductions. Although the estimated profit function would not suggest such perverse supply-side responses, it may not capture long-run adjustments in a competitive general equilibrium.

Future research can make important contributions by studying firm responses to various policy counterfactuals in order to obtain a complete picture of the economic implications of consumer protection regulation.

References

Allcott, Hunt, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman. 2020. "Are high-interest loans predatory? Theory and evidence from payday lending." Working Paper.

Augenblick, Ned, and Matthew Rabin. 2019. "An experiment on time preference and misprediction in unpleasant tasks." *Review of Economic Studies* 86 (3): 941–975.

Bai, Liang, Benjamin Handel, Edward Miguel, and Gautam Rao. 2017. "Self-control and demand for preventive health: Evidence from hypertension in India." NBER Working Paper No. 23727.

Bernheim, B. Douglas, and Dmitry Taubinsky. 2019. "Behavioral Public Economics." In *Handbook of Behavioral Economics*, edited by B. Douglas Bernheim, Stefano Della Vigna, and David I. Laibson, Volume 1. Elsevier.

Bronchetti, Erin T, Judd B Kessler, Ellen B Magenheim, Dmitry Taubinsky, and Eric Zwick. 2020. "Is Attention Produced Rationally?" NBER Working Paper No. 27443.

Brown, Jennifer, Tanjim Hossain, and John Morgan. 2010. "Shrouded attributes and information suppression: Evidence from the field." *The Quarterly Journal of Economics* 125 (2): 859–876.

Caplin, Andrew, Mark Dean, and John Leahy. 2017. "Rationally inattentive behavior: Characterizing and generalizing Shannon entropy." NBER Working Paper No. 23652.

Carrera, Mariana, Heather Royer, Mark Stehr, Justin Sydnor, and Dmitry Taubinsky. 2019. "How are preferences for commitment revealed?" NBER Working Paper No. 26161.

Chaloupka IV, Frank J, Matthew R Levy, and Justin S White. 2019. "Estimating Biases in Smoking Cessation: Evidence from a Field Experiment." NBER Working Paper No. 26522.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *The American Economic Review* 99 (4): 1145–1177.

Connecticut Unfair Trade Practices Act Regulations. 2004. Connecticut Unfair Trade Practices Act Regulations, Sec. 42-110b-19.

Cook, Jonathan A, Noah Newberger, and Joon-Suk Lee. 2020. "On identification and estimation of Heckman models." Working paper.

Cook, Jonathan A, and Saad Siddiqui. 2020. "Random forests and selected samples." *Bulletin of Economic Research* 72 (3): 272–287.

Cummings, Ronald G, Glenn W Harrison, and E Elisabet Rutström. 1995. "Homegrown values and hypothetical surveys: is the dichotomous choice approach incentive-compatible?" *The American Economic Review* 85 (1): 260–266.

DellaVigna, Stefano. 2018. "Structural behavioral economics." In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 1, 613–723. Elsevier.

DellaVigna, Stefano, and Ulrike Malmendier. 2004. "Contract design and self-control: Theory and evidence." *The Quarterly Journal of Economics* 119 (2): 353–402.

———. 2006. "Paying Not to Go to the Gym." *American Economic Review* 96 (3): 694–719 (June).

Dertwinkel-Kalt, Markus, Mats Köster, and Matthias Sutter. 2020. "To buy or not to buy? Shrouding and partitioning of prices in an online shopping field experiment." *European Economic Review*, vol. 130. Artciel 103593.

Eliaz, Kfir, and Ran Spiegler. 2006. "Contracting with diversely naive agents." *The Review of Economic Studies* 73 (3): 689–714.

———. 2008. "Consumer optimism and price discrimination." *Theoretical Economics* 3 (4): 459–497.

Enke, Benjamin, and Florian Zimmermann. 2019. "Correlation Neglect in Belief Formation." *The Review of Economic Studies* 86 (1): 313–332.

Ericson, Keith M Marzilli. 2011. "Forgetting we forget: Overconfidence and memory." *Journal of the European Economic Association* 9 (1): 43–60.

Falk, Armin, and Florian Zimmermann. 2018. "Information Processing and Commitment." *The Economic Journal* 128 (613): 1983–2002.

Farhi, Emmanuel, and Xavier Gabaix. 2020. "Optimal Taxation with Behavioral Agents." *American Economic Review* 110 (1): 298–336.

Federal Trade Commission. 2020. Consumer Information: Rebates. URL: https://www.consumer.ftc.gov/articles/0096-rebates (Retrieved: 09/20/2020).

Fehr, Ernst, and Simon Gächter. 1998. "Reciprocity and economics: The economic implications of homo reciprocans." *European Economic Review* 42 (3-5): 845–859.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded attributes, consumer myopia, and information suppression in competitive markets." *The Quarterly Journal of Economics* 121 (2): 505–540.

Gilpatric, Scott M. 2009. "Slippage in rebate programs and present-biased preferences." *Marketing Science* 28 (2): 229–238.

Grubb, Michael D. 2009. "Selling to overconfident consumers." *American Economic Review* 99 (5): 1770–1807.

Harrison, Glenn W., and John A. List. 2004. "Field Experiments." *Journal of Economic Literature* 42 (4): 1009–1055.

Heckman, James J. 1976. "The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models." In *Annals of economic and social measurement, volume 5, number 4*, 475–492. NBER.

———. 1979. "Sample selection bias as a specification error." *Econometrica* 47 (1): 153–161.

Heidhues, Paul, and Botond Kőszegi. 2010. "Exploiting naivete about self-control in the credit market." *American Economic Review* 100 (5): 2279–2303.

———. 2017. "Naivete-based discrimination." *The Quarterly Journal of Economics* 132 (2): 1019–1054.

Hossain, Tanjim, and John Morgan. 2006. "... plus shipping and handling: Revenue (non) equivalence in field experiments on ebay." *The BE Journal of Economic Analysis & Policy* 6, no. 2.

Jolson, Marvin A, Joshua L Wiener, and Richard B Rosecky. 1987. "Correlates of rebate proneness." *Journal of Advertising Research* 27 (1): 33–43.

Kőszegi, Botond. 2014. "Behavioral Contract Theory." *Journal of Economic Literature* 52 (4): 1075–1118.

List, John A. 2003. "Does market experience eliminate market anomalies?" *The Quarterly Journal of Economics* 118 (1): 41–71.

———. 2011. "Does market experience eliminate market anomalies? The case of exogenous market experience." *American Economic Review* 101 (3): 313–17.

List, John A, and Craig A Gallet. 2001. "What experimental protocol influence disparities between actual and hypothetical stated values?" *Environmental and Resource Economics* 20 (3): 241–254.

Löschel, Andreas, Matthias Rodemeier, and Madeline Werthschulte. 2020. "When Nudges Fail to Scale: Field Experimental Evidence from Goal Setting on Mobile Phones." CESifo Working Paper No. 8485.

Müller, Karsten, and Carlo Schwarz. 2020. "Fanning the flames of hate: Social media and hate crime." *Journal of the European Economic Association*.

PMA. 2005. Mail-in Rebate Benchmarking Study. The Promotion Marketing Association.

Rabin, Matthew. 2000. "Risk aversion and expected-utility theory: A calibration theorem." *Econometrica* 68 (5): 1281–1292.

Rhode Island General Laws. 2016. Commercial Law – General Regulatory Provisions. Chapter 6-13.1. Deceptive Trade Practices. Section 6-13.1-1-.

Rodemeier, Matthias, and Andreas Löschel. 2020. "The Welfare Effects of Persuasion and Taxation: Theory and Evidence from the Field." CESifo Working Paper No. 8259.

Silk, Tim. 2004. Examining Purchase and Non-redemption of Mail-in Rebates: The Impact of Offer Variables on Consumers' Subjective and Objective Probability of Redeeming. PhD dissertation, University of Florida 2004.

Tasoff, Joshua, and Robert Letzler. 2014. “Everyone believes in redemption: Nudges and overoptimism in costly task completion.” *Journal of Economic Behavior & Organization* 107:107–122.

Taubinsky, Dmitry, and Alex Rees-Jones. 2017. “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment.” *The Review of Economic Studies* 85 (4): 2462–2496.

Thaler, Richard H. 2018. “Nudge, not sludge.” *Science* 361 (6401): 431.

Van de Ven, Wynand PMM, and Bernard MS Van Praag. 1981. “The demand for deductibles in private health insurance: A probit model with sample selection.” *Journal of Econometrics* 17 (2): 229–252.

Online Appendix: Not for Publication

A Mathematical Appendix

A.1 Proof of Proposition 1

The probability of buying at the store can be written as

$$B(s,\hat{\theta},\hat{c}) = \int^{s-\hat{c}} \int^{\hat{\theta}(s-\hat{c})+(1-\hat{\theta})\kappa} f(\epsilon|\kappa)d\epsilon h(\kappa)d\kappa + \int_{s-\hat{c}} \int^{\kappa} f(\epsilon|\kappa)d\epsilon h(\kappa)d\kappa. \tag{18}$$

For convenience, let $Q(s,\hat{\theta},\hat{c},\kappa) = \int^{\hat{\theta}(s-\hat{c})+(1-\hat{\theta})\kappa} f(\epsilon|\kappa)d\epsilon h(\kappa)$ and $M(\kappa) = \int^{\kappa} f(\epsilon|\kappa)d\epsilon h(\kappa)$. Then,

$$\begin{aligned}
\frac{\partial}{\partial s}B(s,\hat{\theta},\hat{c}) &= Q(s,\hat{\theta},\hat{c},s-\hat{c}) + \int^{s-\hat{c}} \frac{\partial Q(s,\hat{\theta},\hat{c},\kappa)}{\partial s}d\kappa - M(s-\hat{c}) \\
&= \underbrace{Q(s,\hat{\theta},\hat{c},s-\hat{c}) - M(s-\hat{c})}_{=0} + \int^{s-\hat{c}} \frac{\partial Q(s,\hat{\theta},\hat{c},\kappa)}{\partial s}d\kappa \\
&= \int^{s-\hat{c}} \frac{\partial Q(s,\hat{\theta},\hat{c},\kappa)}{\partial s}d\kappa \\
&= \hat{\theta}\int^{s-\hat{c}} f(\hat{\theta}(s-\hat{c}) + (1-\hat{\theta})\kappa|\kappa)h(\kappa)d\kappa \\
&\approx \hat{\theta}\int^{s-\hat{c}} f(s-\hat{c}|\kappa)h(\kappa)d\kappa,
\end{aligned}$$

which implies

$$\hat{\theta} \approx \frac{\frac{\partial}{\partial s}B(s,\hat{\theta},\hat{c})}{\frac{\partial}{\partial s}B(s,1,\hat{c})}. \tag{19}$$

The approximation requires that $f(\epsilon|\kappa)$ is roughly constant on $[\hat{\theta}(s-\hat{c}) + (1-\hat{\theta})\kappa, s-\hat{c}]$ for all κ.

Next, I derive sufficient statistics for perceived hassle costs. To a first-order approximation,

$$\Delta_{\hat{c}}B(s,\hat{\theta},\hat{c}) \approx \Delta\hat{c}\frac{\partial}{\partial\hat{c}}B(s,\hat{\theta},\hat{c}).$$

If the treatment fully eliminates hassle costs, then $\Delta\hat{c} = 0 - \hat{c}$, and to first order:

$$\hat{c} \approx -\frac{\Delta_{\hat{c}} B(s,\hat{\theta},\hat{c})}{\frac{\partial}{\partial \hat{c}} B(s,\hat{\theta},\hat{c})} \tag{20}$$

$$= \frac{\Delta_{\hat{c}} B(s,\hat{\theta},\hat{c})}{\frac{\partial}{\partial s} B(s,\hat{\theta},\hat{c})}. \tag{21}$$

To go from the first to the second line, I have used the fact that

$$\frac{\partial}{\partial \hat{c}} B(s,\hat{\theta},\hat{c}) = Q(s,\hat{\theta},\hat{c},s-\hat{c}) + \int^{s-\hat{c}} \frac{\partial Q(s,\hat{\theta},\hat{c},\kappa)}{\partial \hat{c}} d\kappa - M(s-\hat{c})$$
$$= -\frac{\partial}{\partial s} B(s,\hat{\theta},\hat{c}).$$

This proves the first part of the proposition. To derive the sufficient statistics for the true redemption frictions, recall that the unconditional redemption probability is given by

$$R(s,\theta,c) = \theta \int^{s-c} dH(\kappa). \tag{22}$$

It immediately follows that

$$\theta = \frac{R(s,\theta,c)}{R(s,1,c)}. \tag{23}$$

An alternative way to identify θ relies on a comparison of redemption elasticities with and without inattention. Note that a very small change in the rebate size changes the redemption probability by

$$\frac{\partial R(s,\theta,c)}{\partial s} = \theta h(s-c),$$

which implies that

$$\theta = \frac{\frac{\partial R(s,\theta,c)}{\partial s}}{\frac{\partial R(s,1,c)}{\partial s}}. \tag{24}$$

Hassle costs can be approximated to first order by

$$c \approx -\frac{\Delta_c R(s,\theta,c)}{\frac{\partial R(s,\theta,c)}{\partial c}} \tag{25}$$

$$= \frac{\Delta_c R(s,\theta,c)}{\frac{\partial R(s,\theta,c)}{\partial s}}, \tag{26}$$

where I have used the fact that

$$\begin{aligned}\frac{\partial R(s,\theta,c)}{\partial c} &= -\theta h(s-c) \\ &= -\frac{\partial R(s,\theta,c)}{\partial s}.\end{aligned}$$

Recall that consumers are sophisticated if and only if $\hat{\theta}=\theta$ and $\hat{c}=c$. Comparing equation 19 with equation 23, and equation 21 with equation 26, implies consumers are sophisticated if and only if

$$\frac{\frac{\partial}{\partial s}B(s,\hat{\theta},\hat{c})}{\frac{\partial}{\partial s}B(s,1,\hat{c})} \approx \frac{R(s,\theta,c)}{R(s,1,c)} \tag{27}$$

and

$$\frac{\Delta_{\hat{c}}B(s,1,\hat{c})}{\frac{\partial}{\partial r}B(s,1,\hat{c})} \approx \frac{\Delta_{c}R(s,1,c)}{\frac{\partial}{\partial s}R(s,1,c)}. \tag{28}$$

This completes the proof. □

A.2 Model with Heterogeneity in Redemption Frictions

In the main part of the paper, behavioral frictions are homogeneous. In this section, I introduce heterogeneity in perceived and true inattention and hassle costs, respectively. With heterogeneity, the main prediction of the model remains the same: if consumers are sophisticated, the aggregate demand elasticity with respect to a rebate must equal the aggregate demand elasticitiy with respect to an equivalent price reduction times the (average) rebate redemption probability.

However, it is more difficult to identify the underlying mechanisms of the lower redemption probabilities. I show that perceived and true inattention are still identified by the same aggregate demand elasticities in Proposition 1. By contrast, hassle costs are no longer identified by aggregate elasticities. The predictions of Proposition 1 are therefore robust to the introduction of heterogeneity for inattention but not for hassle costs.

To show this formally, let $L_{\hat{\theta}}(\hat{\theta})$ and $P_{\theta}(\theta)$ denote the marginal distributions of perceived and true inattention, respectively. Assume that both distributions are smooth and that perceived and true inattention are independent of the idiosyncratic taste parameters, κ and ϵ. $B(\hat{\theta})$ and $R(\theta)$ are now the buying and redemption probability for a given realization of $\hat{\theta}$ and θ, respectively.

The effect of a small change in the rebate value on aggregate demand is therefore

$$\mathbb{E}\left[\frac{\partial B(s,\hat{\theta},\hat{c})}{\partial s}\right] = \int \frac{\partial}{\partial s} B(s,\hat{\theta},\hat{c}) dL_{\hat{\theta}}(\hat{\theta}). \tag{29}$$

Using the same derivation to arrive at equation 19, it follows that the expectation of perceived inattention can be identified by aggregate demand elasticities:

$$\mathbb{E}[\hat{\theta}] \approx \frac{\mathbb{E}\left[\frac{\partial}{\partial s} B(s,\hat{\theta},\hat{c})\right]}{\frac{\partial}{\partial s} B(s,1,\hat{c})}. \tag{30}$$

Similarly, using equation 23, it immediately follows that the expectation of true inattention is identified by aggregate redemption probabilities:

$$\mathbb{E}[\theta] = \frac{\mathbb{E}[R(s,\theta,c)]}{R(s,1,c)}. \tag{31}$$

These results show that perceived and true inattention are identified by the same aggregate buying and redemption behavior as in Proposition 1.

Next, consider the case in which hassle costs are heterogeneous. Let $L_{\hat{c}}(\hat{c})$ and $P_c(c)$ denote the marginal distribution of $\hat{c}$ and c, respectively, and assume that both distributions are smooth and independent to the idiosyncratic taste parameters. The aggregate demand response to a change in perceived hassle costs is approximated to first order by

$$\mathbb{E}\left[\Delta_{\hat{c}} B(s,\hat{\theta},\hat{c})\right] \approx \int \Delta\hat{c} \frac{\partial}{\partial \hat{c}} B(s,\hat{\theta},\hat{c}) dP_{\hat{c}}(\hat{c}), \tag{32}$$

which is generally not equal to $\mathbb{E}[\Delta\hat{c}]\mathbb{E}[\frac{\partial}{\partial \hat{c}} B(s,\hat{\theta},\hat{c})]$. The demand response for consumer types with a given $\hat{c}$ depends on both the type-specific change in their perceived hassle costs and the type-specific buying elasticity. Since both $\Delta\hat{c}$ and $\frac{\partial}{\partial \hat{c}} B(s,\hat{\theta},\hat{c})$ vary with $\hat{c}$, the expectation of the product is not equal to the product of the individual expectations. Thus, we cannot re-arrange terms and use the same identification strategy as in equation 21. An analogous argument can be made when true hassle costs are heterogeneous by taking the expectation of both sides of equation 26: changes in aggregate redemption probabilities are not sufficient to identify expected hassle costs.

In sum, the empirical identification strategy of perceived and true inattention is robust to the introduction redemption frictions, but hassle costs are only identified if they are approximately homogenous. These results hold as long as redemption frictions are independent of the idiosyncratic taste parameters, i.e. of preferences. In future versions of the paper, I will further address challenges resulting from heterogeneity by estimating heterogeneous treatment effects using the empirical strategy specified in the pre-analysis plan.

A.3 Model with Ad Valorem Rebate

Proposition 1 was derived using a lump sum rebate of value s, whereas the experimental design uses an ad valorem rebate. In this section, I show that the same predictions from Proposition 1 can be derived with an ad valorem rebate. The difference between the two types of rebates is that it is slightly more complicated to model an ad valorem rebate because the rebate value depends on the endogenous purchase value of the consumer.

Let t denote an ad valorem rebate. The value of the rebate is given by $t\mathbf{p}'\mathbf{x}$ where $\mathbf{p} = (p^1, p^2, ..., p^J)$ is the vector of prices and $\mathbf{x} = (x^1, x^2, ..., x^J)$ the consumption vector. Different to a lump-sum rebate, an ad valorem rebate changes the optimal consumption vector because it effectively changes the price of each good. We therefore need to model a third margin where the consumption vector is a function of the rebate.

Let $\mathbf{x}_r$ be the chosen consumption vector given redemption choices r. Given the consumer buys at the store and is attentive, she chooses

$$\mathbf{x}_r = \arg\max_{\mathbf{x}} \{v(\mathbf{x}) - \mathbf{p}'\mathbf{x} + r\,(t\mathbf{p}'\mathbf{x} - \hat{c})\}$$

If she is not attentive, she chooses the same consumption vector as if she was attentive but decided not to redeem the rebate, i.e. $\mathbf{x}_0$. The first-order conditions are

$$\frac{\partial v}{\partial x^j} - p^j + rtp^j = 0$$

for every good j.

Given the consumer buys at the store and is attentive, she chooses $r = 1$ if and only if

$$\begin{gathered} v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c} \geq v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0} + \kappa \\ \Leftrightarrow u(t, \hat{c}) \geq \kappa \end{gathered}$$

with $u(t, \hat{c}) = v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c} - (v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0})$.

She chooses to buy at the store if and only if

$$\hat{\theta}\,\{r\,(v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c}) + (1 - r)\,(v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0} + \kappa)\} + \left(1 - \hat{\theta}\right)\{v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0} + \kappa\} \geq \epsilon.$$

For convenience, let $w_1(t, \hat{\theta}, \hat{c}, \kappa) = \hat{\theta}\,\{(v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c})\} + \left(1 - \hat{\theta}\right)\{v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0} + \kappa\}$

and $w_0(\kappa) = v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0} + \kappa$. The probability to buy at the store can be expressed by

$$B(t, \hat{\theta}, \hat{c}) = \int^{u(t,\hat{c})} \int^{w_1(t,\hat{\theta},\hat{c},\kappa)} f(\epsilon|\kappa) d\epsilon h(\kappa) d\kappa + \int_{u(t,\hat{c})} \int^{w_0(\kappa)} f(\epsilon|\kappa) d\epsilon h(\kappa) d\kappa.$$

Let $Q(t, \hat{\theta}, \hat{c}, \kappa) = \int^{w_1(t,\hat{\theta},\hat{c},\kappa)} f(\epsilon|\kappa) d\epsilon h(\kappa)$ and $M(\kappa) = \int^{w_0(\kappa)} f(\epsilon|\kappa) d\epsilon h(\kappa)$. The effect of a very small change in the rebate on the buying probability is given by

$$\begin{aligned}
\frac{\partial}{\partial t} B(t, \hat{\theta}, \hat{c}) &= \frac{\partial u}{\partial t} Q(t, \hat{\theta}, \hat{c}, u) + \int^u \frac{\partial Q(t, \hat{\theta}, \hat{c}, \kappa)}{\partial t} d\kappa - \frac{\partial u}{\partial t} M(u) \\
&= \frac{\partial u}{\partial t} \Big(Q(t, \hat{\theta}, \hat{c}, u) - M(u) \Big) + \int^u \frac{\partial Q(t, \hat{\theta}, \hat{c}, \kappa)}{\partial t} d\kappa.
\end{aligned}$$

Note that

$$\begin{aligned}
w_1(t, \hat{\theta}, \hat{c}, u) &= \hat{\theta} \left\{ (v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c}) \right\} + \left(1 - \hat{\theta} \right) \left\{ v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0} + u \right\} \\
&= \hat{\theta} \left\{ (v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c}) \right\} + \left(1 - \hat{\theta} \right) \left\{ [v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c}] \right\} \\
&= v(\mathbf{x_1}) - \mathbf{p}'\mathbf{x_1} + t\mathbf{p}'\mathbf{x_1} - \hat{c} \\
&= v(\mathbf{x_0}) - \mathbf{p}'\mathbf{x_0} + u \\
&= w_0(u).
\end{aligned}$$

Therefore, $Q(t, \hat{\theta}, \hat{c}, u) - M(u) = 0$ and

$$\begin{aligned}
\frac{\partial}{\partial t} B(t, \hat{\theta}, \hat{c}) &= \int^u \frac{\partial Q(t, \hat{\theta}, \hat{c}, \kappa)}{\partial t} d\kappa \\
&= \int^u \frac{\partial w_1(t, \hat{\theta}, \hat{c}, \kappa)}{\partial t} f(w_1(t, \hat{\theta}, \hat{c}, \kappa)|\kappa) h(\kappa) d\kappa \\
&= \int^u \hat{\theta} \left((v_{\mathbf{x}_1} - \mathbf{p}') \frac{\partial \mathbf{x}_1}{\partial t} + (t\mathbf{p}' \frac{\partial \mathbf{x}_1}{\partial t} + \mathbf{p}'\mathbf{x}_1) \right) f(w_1(t, \hat{\theta}, \hat{c}, \kappa)|\kappa) h(\kappa) d\kappa \\
&= \hat{\theta} \mathbf{p}'\mathbf{x}_1 \int^u f(w_1(t, \hat{\theta}, \hat{c}, \kappa)|\kappa) h(\kappa) d\kappa.
\end{aligned}$$

If f is roughly constant on the interval $[w_1(t, \hat{\theta}, \hat{c}, \kappa), w_1(t, 1, \hat{c}, \kappa)]$, then

$$\hat{\theta} \approx \frac{\frac{\partial}{\partial t} B(t, \hat{\theta}, \hat{c})}{\frac{\partial}{\partial t} B(t, 1, \hat{c})}.$$

To derive the sufficient statistics for perceived hassle costs, first note that a small change in perceived hassle costs changes the buying probability by

$$\begin{aligned}\frac{\partial}{\partial \hat{c}}B &= \frac{\partial u}{\partial \hat{c}}Q(t,\hat{\theta},\hat{c},u(\cdot)) + \int^{u} \frac{\partial Q(t,\hat{\theta},\hat{c},\kappa)}{\partial \hat{c}}d\kappa - \frac{\partial u}{\partial \hat{c}}M(u(\cdot)) \\ &= -\hat{\theta}\int^{u} f(w_1|\kappa)h(\kappa)d\kappa.\end{aligned}$$

This implies that

$$\frac{\partial}{\partial t}B(t,\hat{\theta},\hat{c}) = -\frac{\partial}{\partial \hat{c}}B(t,\hat{\theta},\hat{c})\mathbf{p}'\mathbf{x}_1.$$

To a first-order approximation,

$$\begin{aligned}\Delta_{\hat{c}}B(t,\hat{\theta},\hat{c}) &\approx \frac{\partial}{\partial \hat{c}}B(t,\hat{\theta},\hat{c})\Delta\hat{c} \\ \Delta\hat{c} &\approx \frac{\Delta_{\hat{c}}B(t,\hat{\theta},\hat{c})}{\frac{\partial B(t,\hat{\theta},\hat{c})}{\partial \hat{c}}}.\end{aligned}$$

If $\Delta\hat{c} = 0 - \hat{c}$, then:

$$\begin{aligned}\hat{c} &\approx -\frac{\Delta_{\hat{c}}B(t,\hat{\theta},\hat{c})}{\frac{\partial B(t,\hat{\theta},\hat{c})}{\partial \hat{c}}} \\ &= \frac{\Delta_{\hat{c}}B(t,\hat{\theta},\hat{c})}{\frac{\partial B(t,\hat{\theta},\hat{c})}{\partial t}}\mathbf{p}'\mathbf{x_1}.\end{aligned}$$

As we can see, both $\hat{\theta}$ and $\hat{c}$ are identified in the same way as in Proposition 1 but s is replaced by $t\mathbf{p}'\mathbf{x}_1$.

Next, I derive the sufficient statistics for the true redemption frictions. The redemption probability is given by:

$$R(t,\theta,c) = \theta\int^{u} dH(\kappa)$$

such that inattention is identified by

$$\theta = \frac{R(t,\theta,c)}{R(t,1,c)}.$$

To identify ture hassle costs first note that

$$\frac{\partial R(t,\theta,c)}{\partial c} = -\theta h(u)$$
$$= -\frac{\partial R(t,\theta,c)}{\partial t}(\mathbf{p}'\mathbf{x_1})^{-1}$$

which implies that, to first order,

$$\Delta_c R(t,\theta,c) \approx \frac{\partial R(t,\theta,c)}{\partial c}\Delta c$$
$$= -\frac{\partial R(t,\theta,c)}{\partial t}(\mathbf{p}'\mathbf{x}_1)^{-1}\Delta c$$
$$\Leftrightarrow c \approx \frac{\Delta_c R(t,\theta,c)}{\frac{\partial R(t,\theta,c)}{\partial t}}\mathbf{p}'\mathbf{x}_1$$

with $\Delta c = -c$.

A.4 Derivation of Moment Conditions

In this section, I derive the moment conditions used to estimate the structural parameters in Section 6. Perceived hassle costs are approximated by

$$\hat{c} \approx \frac{\Delta_{\hat{c}} B(r,1,\hat{c})}{\frac{\partial}{\partial r}B(r,1,\hat{c})} \approx \frac{\beta_A - \beta_{B.2b}}{\frac{\beta_{C.2b}-\beta_{B.2b}}{\Delta s}} \tag{33}$$
$$\Leftrightarrow \beta_{C.2b} \approx \frac{\Delta s}{\hat{c}}\left(\beta_A - \beta_{B.2b}\right) + \beta_{B.2b}. \tag{34}$$

Perceived inattention is approximated by

$$\hat{\theta} \approx \frac{\beta_{C.1} - \beta_{B.1}}{\beta_{C.2b} - \beta_{B.2b}} \tag{35}$$
$$\Leftrightarrow \beta_{C.1} \approx \hat{\theta}\left(\beta_{C.2b} - \beta_{B.2b}\right) + \beta_{B.1} \tag{36}$$
$$= \hat{\theta}\frac{\Delta s}{\hat{c}}\left(\beta_A - \beta_{B.2b}\right) + \beta_{B.1}, \tag{37}$$

where the last line follows from substituting $\beta_{C.2b}$ from equation 34.
Under this assumption, true hassle costs are identified by

$$c = \frac{\Delta_c R(s,1,c)}{\frac{\partial R(s,1,c)}{\partial s}} = \frac{-\tau_{B2b}}{\frac{\partial R(s,1,c)}{\partial s}} \tag{38}$$

$$\Leftrightarrow \tau_{B2b} = -c\frac{\partial R(s,1,c)}{\partial s}. \tag{39}$$

$$\tag{40}$$

We can insert this into the expression for τ_{C2b}:

$$\begin{aligned}\tau_{C2b} &= \tau_{B2b} + \Delta s\frac{\partial R(1,c,s)}{\partial s}\\ &= \frac{\partial R(s,1,c)}{\partial s}(\Delta s - c)\end{aligned}$$

This yields the moment conditions in equation 16.

True inattention can be identified in multiple ways. The first identification strategy relies on the comparison between redemption probabilities with and without inattention:

$$\begin{aligned}\theta &= \frac{R(s,\theta,c)}{R(s,1,c)} \approx \frac{\tau_{B1} + \tau_A}{\tau_A + \tau_{B2b}}\\ \Leftrightarrow \tau_{B1} &\approx \theta(\tau_A + \tau_{B2b}) - \tau_A\\ &= \theta(\tau_A - c\frac{\partial R(s,1,c)}{\partial s}) - \tau_A.\end{aligned}$$

The second identification strategy of inattention relies on the comparison of demand derivatives:

$$\theta = \frac{\frac{\partial}{\partial s}R(s,\theta,c)}{\frac{\partial}{\partial s}R(s,1,c).}$$

We can insert this condition into the expression for τ_{C1}:

$$\begin{aligned}\tau_{C1} &= \tau_{B1} + \Delta s\frac{\partial R(s,\theta,c)}{\partial s}\\ &= \tau_{B1} + \theta\Delta s\frac{\partial R(s,1,c)}{\partial s}\\ &= \theta(\tau_A - c\frac{\partial R(s,1,c)}{\partial s}) - \tau_A + \theta\Delta s\frac{\partial R(s,1,c)}{\partial s}\\ &= \theta\left[\tau_A + \frac{\partial R(s,1,c)}{\partial s}(\Delta s - c)\right] - \tau_A,\end{aligned}$$

where in the last line, I have substituted for τ_{B1}. Rewriting τ_{B1} and τ_{C1} as above produces the moment conditions in equation 17.

B Pilot Studies

The first pilot study was implemented between July and August 2018 for a period of three weeks and with a smaller sample of 13,204 website visitors in the United Kingdom. Different from the main experiment, the study did not include users with mobile devices and tablets. Only subjects who used a desktop were randomized into one of the experimental groups. The experiment included the following three treatment groups and one control group: group A, B.1, B.2a, and D.

The second pilot study took place in August 2019 for a period of one week and with a larger sample of 52,302 consumers in the German online shop of the company. Just as in the first pilot study, the study included only desktop users, and the experimental design consisted of the groups A, B.1, B.2a, and D.

The actual design of the banners and the rebate code differed somewhat for both studies in comparison to the one presented in the main part of the paper. The reason for the divergence is that the marketing department of the company frequently changes the promotion design. The experiments I ran effectively changed certain features of these promotions but not the entire visual design. I am currently waiting for clearance by the company to display the banners of the pilot studies in this paper. Overall, the conceptual design of the banners was very similar to the banners presented in the main body of the paper.

Table B1 documents the results from linear probability models of the outcomes of interest on the treatment indicators. The baseline buying probability in the first pilot study with UK customers is around 3.4%, and therefore higher than for the sample analyzed in the main experiment of the paper. The automatically applied 10% discount increases the buying probability by around 1.04 percentage points. Introducing hassle costs lowers this effect to 0.86 percentage points, but the difference in treatment coefficients is not statistically significant. Increasing inattention by removing the reminder further reduces the positive effect of the rebate down to 0.42 percentage points. Thus, even though the coefficients are noisier than for the larger main experiment, the qualitative results are the same: consumers make large adjustments on the extensive margin in response to the introduction of redemption barriers. Inattention is perceived to be a larger barrier than hassle costs.

Looking at intensive-margin responses, I observe a redemption probability of 95% for subjects who receive the automatically applied discount. Redemption rates fall much more steeply than in the main experiment of the paper. The introduction of hassle costs is associated with a decrease of 73 percentage points. Removing the reminder is associated with an additional decrease of 5 percentage points.

An interesting observation is that even though redemption barriers appear to be much larger in this experiment, extensive-margin responses are also subsantially larger than in the main experiment: the introduction of hassle costs and inattention reduces the positive effect on the buying probability of the rebate by 60% (from 1.04 to 0.42 percentage points). Although differences across experiments are only correlational, this behavior is highly consistent with the notion of consumer sophistication: rebates with larger redemption barriers have smaller positive effects on demand.

In the second pilot study with subjects in Germany, the baseline buying probability equals 4.9%, and an automatically applied discount increases demand by 0.73 percentage points. Again, introducing hassle costs has no significantly different effect on the extensive margin. The coefficient is even slightly larger than the one in group A, but the difference in coefficients is likely a result of sampling variation. Removing the reminder substantially lowers the demand response to the rebate. The positive effect on the buying probability is 86% lower than the effect of the automatically applied discount. Once more, this result confirms the importance of observing extensive-margin behavior in order to draw a correct conclusion regarding the effects of rebates on consumer behavior and welfare.

On the intensive margin, the baseline redemption probability is 77% and falls by 30 percentage points with the introduction of hassle costs. Removing the reminder has an additional negative effect, but hassle costs appear to be the more substantial redemption barrier.

Although the effects in the pilot studies are quantitatively not the same as the ones presented in the main experiment, the qualitative results are the same. Differences in point estimates are to be expected because the samples differed along many dimensions, such as the type of device consumers used to visit the website, the user's country of origin, and the year of the experiment. Other idiosyncratic differences are the visual designs of the promotions created by the marketing department. However, the qualitative results are identical to the ones discussed in the main body of the paper: both hassle costs and inattention are substantial barriers to rebate redemption, and consumers make large adjustments on the extensive margin in response to these frictions. Despite the fact that hassle costs appear to be a larger redemption barrier than inattention in all experiments, consumers respond more to inattention than hassle costs on the extensive margin.

Table B1: Buying and Redemption Probabilities in Pilot Studies

	Pilot Study 1		Pilot Study 2	
	(1)	(2)	(3)	(4)
	Buying Probability ×100	Redemption Probability ×100	Buying Probability ×100	Redemption Probability ×100
A: 10%, automatic	1.040**	95.172***	0.728***	77.166***
	(0.481)	(1.787)	(0.277)	(1.558)
B.1: 10%, w/o reminder	0.417	-77.850***	0.099	-43.376***
	(0.460)	(3.815)	(0.268)	(2.416)
B.2a: 10%, w/ reminder	0.862*	-72.795***	0.873***	-30.341***
	(0.472)	(3.928)	(0.277)	(2.393)
D: control	3.403***		4.889***	
	(0.317)		(0.189)	
Regression constant	D	A	D	A
N	13,204	415	52,302	2,140
Country	United Kingdom		Germany	
Year	2018		2019	
Sample	Desktop users only		Desktop users only	

Note: The table shows average treatment effects for the two pilot studies that preceded the main experiment in the paper. Average treatment effects are estimated from a linear probability model of the buying and redemption probability on the treatment indicators. The control group is excluded in columns 2 and 4 because control group subjects could not redeem the rebate by construction. Robust standard errors are in parentheses and clustered at the subject level. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

C Dynamic Effects

The effect of redemption frictions on consumer behavior may vary with experience. Some consumers who are naive at the first purchase might learn from their mistake and, for instance, be more attentive during follow-up purchases. They may also decide not to return to the store that exploited their initial naiveté.

Evaluating these effects with the data at hand is difficult. Most of the products the company offers are durable goods that are purchased infrequently. For instance, new furniture is often bought every five or 10 years. The longevity of the goods may actually increase the scope of exploitation, because consumers are more likely to forget they made a mistake as time passes.

I provide some suggestive evidence on the dynamic effects by analyzing two additional outcomes. First, I estimate the effects of the treatment on the probability of buying more than once during the experimental period. The outcome variable is a dummy equal to 1 if the consumer purchased twice or more, and 0 otherwise. Second, I evaluate whether subjects have gained experience with rebates, by estimating the treatment effects on the probability of redeeming the rebate during the follow-up purchases. In both regression models, the constant represents the first moment of group A that received the automatically applied discount. All treatment coefficients are therefore interpreted relative to a standard price reduction.

Table C1 reports the results. Overall, the probability of buying again is low for all experimental groups, due to the durable nature of the goods offered by the firm. In the group with the automatically applied discount, 2.8% of all buyers make a second purchase during the experimental period. Almost all rebates that involve redemption frictions have a negative coefficient. For instance, the probability of buying again is 0.51 percentage points for the typical rebate that was offered in group B.1. Even though this difference is a fairly large effect of 17.9%, the estimate is too noisy to conclude that this difference is significant at any conventional levels. The same applies for the other rebate coefficients: most of them are large in size, but statistical power is too low due to the low baseline probability of buying more than once. It is important to note that the most negative coefficient is the one of the control group where subjects did not receive a discount. This observation suggests that even though rebates have negative effects on customer loyalty relative to an automatically applied discount, they still increase the repurchase rate relative to offering no price promotion.

The redemption probability in group A on the follow-up purchases is lower and equal to 70%. Redemption frictions are still associated with lower redemption rates for all groups except group C.2a. Interestingly, all coefficients are less negative than on the first purchase. This may be suggestive evidence of consumers learning over time. It may also be a result of selection: more attentive consumers with low hassle costs are more likely to return to the shop because they are less affected by the barriers introduced by the firm.

Overall, the directional effects on follow-up purchases and redemption behavior are in line with economic intuition, but the data only allows us to draw limited conclusions.

Table C1: Repurchase Probabilities and Redemption Probabilities on Follow-Up Purchases

	(1) Probability of follow-up purchase (in %)	(2) Redemption probability on follow-up purchase (in %)
B.1: 10%, w/o reminder	-0.510	-16.159
	(0.551)	(10.689)
B.2a: 10%, w/ reminder	0.029	-16.046
	(0.544)	(9.981)
B.2b: 10%, w/ reminder+announce	-0.186	-12.073
	(0.549)	(10.264)
C1: 15%, w/o reminder	-0.438	-6.798
	(0.542)	(10.393)
C.2a: 15%, w/ reminder	-0.279	0.000
	(0.532)	(10.033)
C.2b: 15%, w/ reminder+announce	-0.563	-6.798
	(0.535)	(10.393)
D: control	-0.623	
	(0.575)	
A: 10%, automatic (constant)	2.842***	70.213***
	(0.386)	(7.094)
N	13,223	304

Note: Column 1 reports average treatment effects on the probability of purchasing more than once. Column 2 shows average treatment effects on the probability of redeeming a rebate on a follow-up purchase, conditional on buying more than once. In both columns, the regression constant is the mean of group A. Robust standard errors are in parentheses. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

D Sample Selection Model

The selection model uses regional and temporal variation in internet outages as an exclusion restriction. I use publicly available data on internet outages from Heise Online, a platform that documents user complaints about internet outages received by phone across the country. The dataset includes, among other variables, the area code and the duration of the outage. For the experimental observations, I only observe the city of each website visitor and not the area code. To merge internet outages with the dataset from the experiment, I use geo data from OpenGeoDB to assign each area code to a respective city. This approach allows me to assign internet outages collected from Heise Online to website visitors in the experiment.

One could use various approaches to construct a dummy variable that indicates whether a city experienced a major internet outage. In constructing the variable, I closely follow Müller and Schwarz (2020), who have used outages as exogenous variation in a different setting. Specifically, they study the effect of social media utilization on hate crime, and use internet outages as exogenous variation for access to social media. Following their approach, I count the total number of internet outages that occurred in the city of the website visitor. Because larger cities will have more internet outages mechanically, the authors normalize the number of internet outages by the number of inhabitants of each city, and I follow their approach. I then create a dummy variable that indicates whether the subject's area experienced a major internet outage. I define major internet outages as the 90th percentile of total internet outages normalized by the number of inhabitants. Because internet outages may also affect whether subjects even appear in my dataset (another level of sample selection), I only count internet outages that happened after the subject's *first* visit to the website during the experimental period. To avoid that subjects who visit at a later point in time have a lower number of outages mechanically, I count internet outages for each subject seven days after their first visit. Thus, even for subjects whose first visit was during the last day of the experiment, the following seven days are accounted for in terms of outages.

The sample selection model follows the standard setup introduced by Van de Ven and Van Praag (1981) when both the intensive- and extensive-margin variables are binary. With some abuse of notation, I denote the buying decision of subject i by b_i and her rebate redemption choice by r_i. The utility from buying at the shop is given by

$$u_i = \gamma \mathbf{Z}_i + \eta_i, \tag{41}$$

where $\mathbf{Z}_i$ is a vector that includes an indicator for each treatment and the instrument indicating whether the city of subject i experiences a major internet outage.

The latent utility component is denoted by η_i. Utility from rebate redemption equals

$$v_i = \omega \mathbf{T}_i + \zeta_i, \tag{42}$$

where ζ_i is the unobserved utility from rebate redemption.

Subject i's buying decision is given by

$$b_i = \begin{cases} 1 \text{ if } u_i > 0 \\ 0 \text{ otherwise }. \end{cases}$$

Her redemption choice is determined by the intensive margin utility and only observed if she buys:

$$r_i = \begin{cases} 1 \text{ if } v_i > 0 \text{ and } b_i = 1 \\ 0 \text{ if } v_i \leq 0 \text{ and } b_i = 1 \\ 0 \text{ if } b_i = 0. \end{cases} \tag{43}$$

Selection arises when $cov(\eta, \zeta) \neq 0$. I make the standard assumption that each error term follows a standard normal distribution, $\eta \sim N(0,1)$ and $\zeta \sim N(0,1)$, with correlation between the residuals given by $\rho = corr(\eta, \zeta)$. Monte Carlo simulations show that when these distributional assumptions are violated, the model still performs well in many cases, as long as a valid exclusion restriction exists (Cook and Siddiqui 2020).

To estimate the parameters of interest, I maximize the well-known form of the log-likelihood function that can be derived from the model above:

$$\begin{aligned} lnL = \sum_{i=1}^{N} \{ & b_i r_i ln\Phi_2(T\omega, Z\gamma, \rho) + b_i(1-r_i)ln[\Phi(T\omega) - \Phi_2(T\omega, Z\gamma;\rho)] + (1-b_i)r_i ln[\Phi(Z\gamma) \\ & - \Phi_2(T\omega, Z\gamma;\rho)] + (1-b_i)(1-r_i)ln[1 - \Phi(T\omega) - \Phi(Z\gamma) - \Phi_2(T\omega, Z\gamma;\rho)]\}, \end{aligned} \tag{44}$$

where I denote the standard normal distribution by Φ and the joint distribution by Φ_2. If the correlation between residuals is zero, this likelihood simply equals the sum of the likelihoods of two independent probit models.

Given the structure of the model, the differences in redemption rates between experimental conditions, that is, the coefficients in ω, have a causal interpretation.

Before estimating the selection model, I first analyze whether internet outages have a significant effect on the buying probability and whether this effect is orthogonal to the treatment variation. Table D1 provides results from a linear probability model of the buying decision on internet outages and the treatments. Column 1 only includes the instrument, whereas column 2 adds the experimental treatments. Major internet outages cause an economically large and highly statistically significant decrease in the buying probability by 8.3%, or 0.18 percentage points. The addition of experimental treatments in column 2 does not affect the coefficient of the instrument. This finding is reassuring because it indicates the exclusion restriction is not correlated with the experimental treatments—a result we would expect through random treatment assignment.

Next, I maximize the likelihood function in 44. To ensure I have found the global, instead of a local, maximum, I estimate the model for various *given* values of the correlation between residuals, ρ, and then compare the log likelihood values with the one when ρ is estimated. I estimate the log likelihood for given values of the correlation between the residuals using the code developed by Cook, Newberger, and Lee (2020). Figure D1 reports results by plotting the log-likelihood value for given values of the correlation. Visual inspection suggests the global maximum lies between 0.3 and 0.4.

Table D2 reports the main estimation results and shows that the global maximum has been found. The log-likelihood value is –63,934.67, and the correlation between residuals is 0.35. The latter would imply that unobservables that increase the buying probability also increase the redemption probability. Two independent linear probability models would then overestimate the redemption probability, because subjects with a systematically larger likelihood of redeeming have selected into the subsample of buyers. However, there is no indication for significant sample selection bias: the coefficient showing Fisher's Z-transformation of the correlation between residuals is not statistically significantly different from zero.

The treatment coefficients are similar to the ones of the two independent OLS models. I find that the effect of hassle costs, as identified by B.2b, equals a reduction in the redemption probability by 22 percentage points, and is therefore virtually the same as in the OLS model. The (announced) reminder increases the redemption rate by 9 percentage points, which is around 6 percentage points less than in the model with independent residuals. Overall, the treatment coefficients are fairly similar to the treatment effects estimated in the main part of the paper, indicating the degree of selection bias is small if the model is correctly specified.

A main difference is that the standard errors are substantially larger. Inflated standard errors are a well-known issue with many sample selection models and a result from a high level of collinearity between the treatment regressors and the correction term. Collinearity is a particular limitation in my application, because all treatments need to appear both in the selection and outcome equation. Even though the exclusion restriction reduces the resulting noise through a strong effect in the selection equation, high collinearity remains a serious concern. Results should accordingly be interpreted with this caveat in mind. In future drafts of this paper, I will include alternative selection models to test the robustness of the conclusion that can be drawn from the present estimation results.

Table D1: Effect of Internet Outages on Buying Probability

	(1) Buying Probability (in %)	(2) Buying Probability (in %)
Internet outage	-0.184***	-0.185***
	(0.051)	(0.051)
A: 10%, automatic		0.393***
		(0.076)
B.1: 10%, w/o reminder		0.298***
		(0.076)
B.2a: 10%, w/ reminder		0.432***
		(0.076)
B.2b: 10%, w/ reminder+announce		0.343***
		(0.076)
C1: 15%, w/o reminder		0.472***
		(0.076)
C.2a: 15%, w/ reminder		0.613***
		(0.076)
C.2b: 15%, w/ reminder+announce		0.580***
		(0.076)
D: control	2.228***	1.836***
	(0.021)	(0.054)
N	601,805	601,805

Note: The table reports average treatment effects from a linear probability model of internet outages and treatment indicators on the buying probability. Column 1 only includes internet outages as a regressor, and column 2 adds the experimental treatments. Robust standard errors are in parentheses. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

Figure D1: Log Likelihood Values for Given Values of ρ

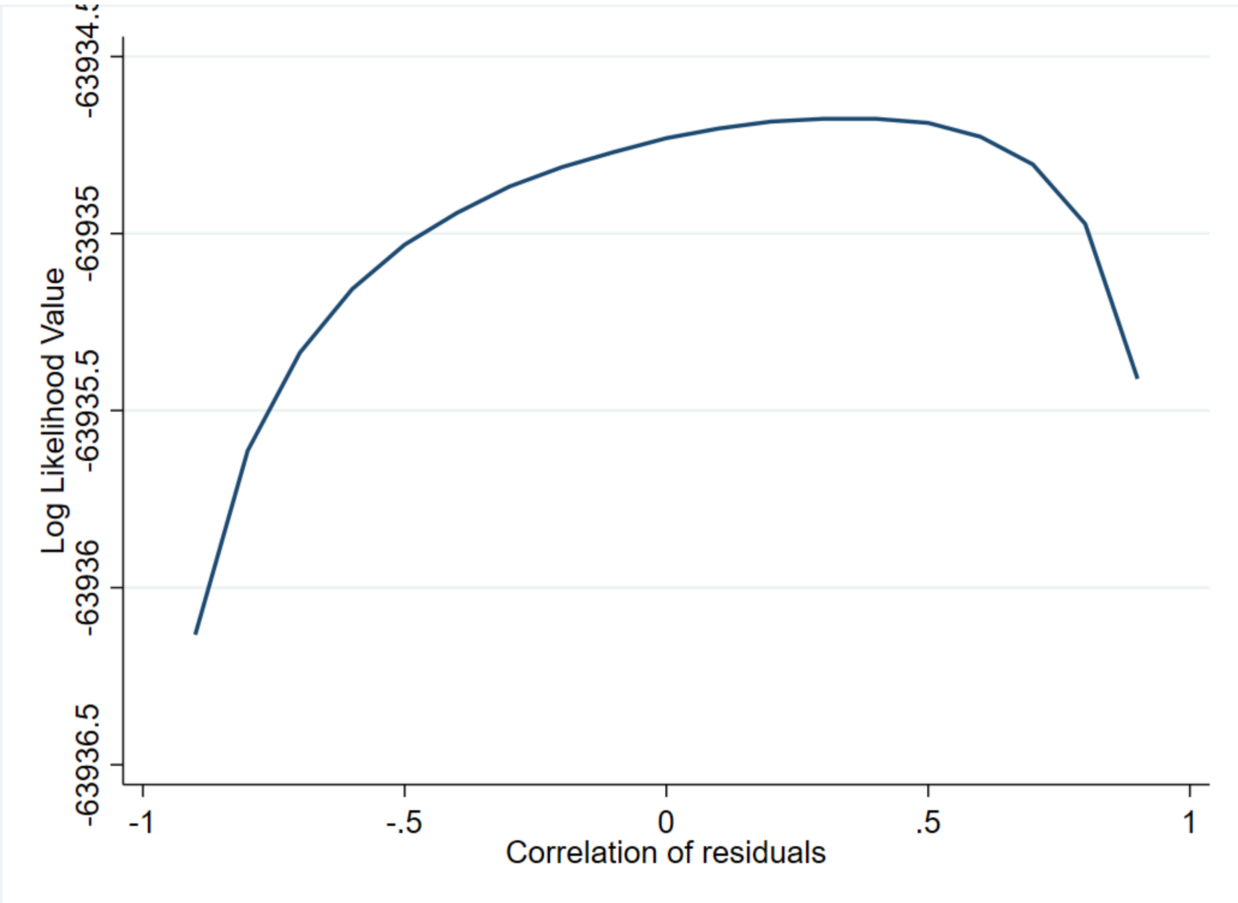

Table D2: Intensive Margin Treatment Effects

	Redemption Probability (in %)
B.1: 10%, w/o reminder	-31.181***
	(8.636)
B.2a: 10%, w/ reminder	-25.302***
	(6.481)
B.2b: 10%, w/ reminder+announce	-22.008***
	(5.314)
C.1: 15%, w/o reminder	-28.849***
	(7.802)
C.2a: 15%, w/ reminder	-20.534***
	(5.084)
C.2b: 15%, w/ reminder+announce	-18.023***
	(4.352)
ρ	0.344
	(0.853)
Fisher's Z-transformation	0.359
	(0.968)
Log likelihood	-63934.67
N	526,933

Note: This table reports estimation results from the sample selection model in 44. The correlation between intensive and extensive margin residuals is denoted by ρ. Fisher's Z transformation is the inverse hyperbolic tangent of ρ and asymptotically normally distributed. Standard errors are in parentheses. *,**,***: significant at $p < 0.1$, $p < 0.05$, $p < 0.01$, respectively.

E Additional Figures

Figure E1: Banner in Groups C.1 and C.2a: 15% Rebate, Active Redemption

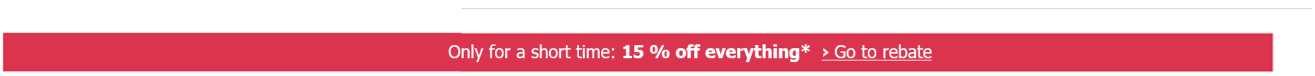

Note: This figure shows an English translation of the banner displayed in experimental groups C.1 and C.2a.

Figure E3: Subpage with Rebate Code in Group C.1, C.2a and C.2b

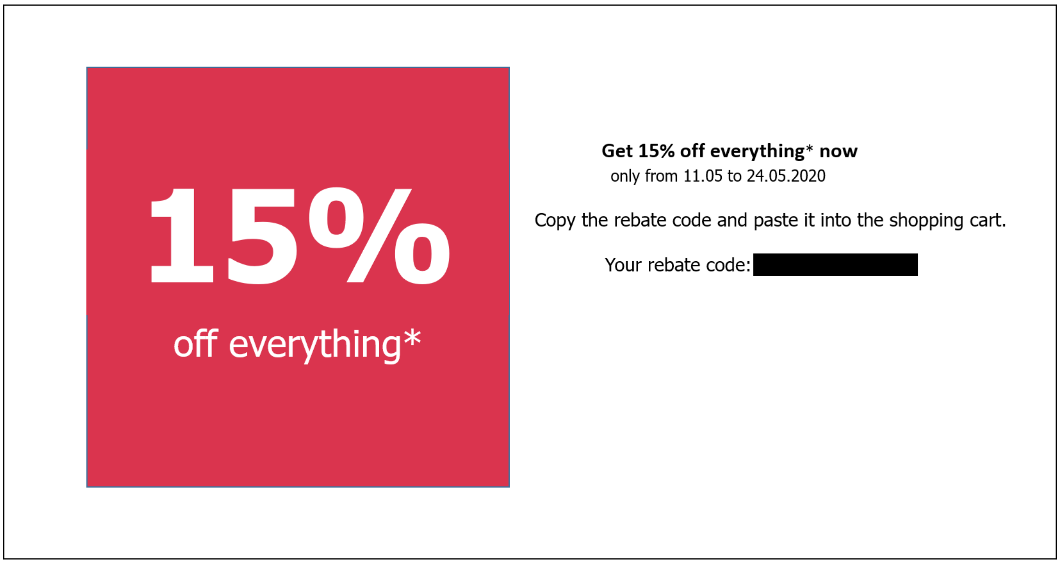

Note: This figure shows an English translation of the subpage showing the rebate code in experimental groups C.1, C2.a, and C2.b.

Figure E2: Banner in Groups C.2b: 15% Rebate, Active Redemption with Announcement of Reminder

Note: This figure shows an English translation of the pop-up box that explicitly announces the reminder upon visiting the website in experimental group C2.b.

Figure E4: Buying and Redemption Probabilities: Disaggregated Treatments

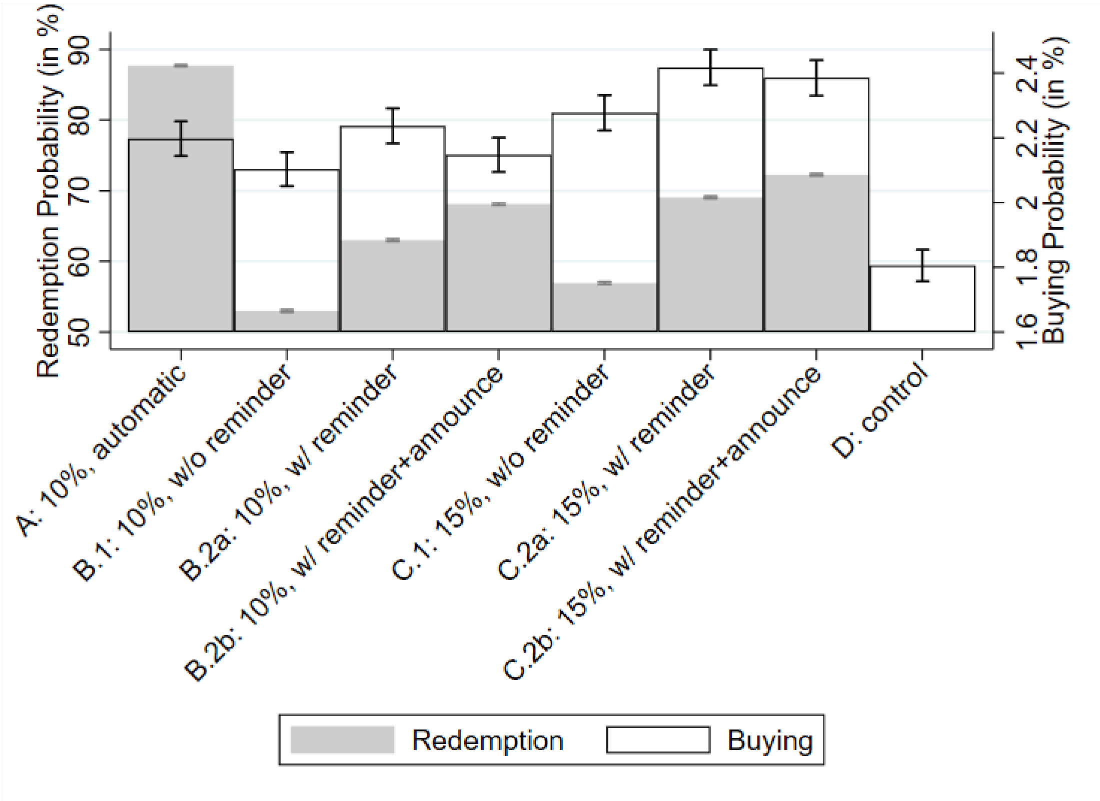

Note: This figure shows the buying probability for the entire sample and the redemption probability conditional on buying. The error bars represent standard errors.

Figure E5: Probability to Visit Checkout Page

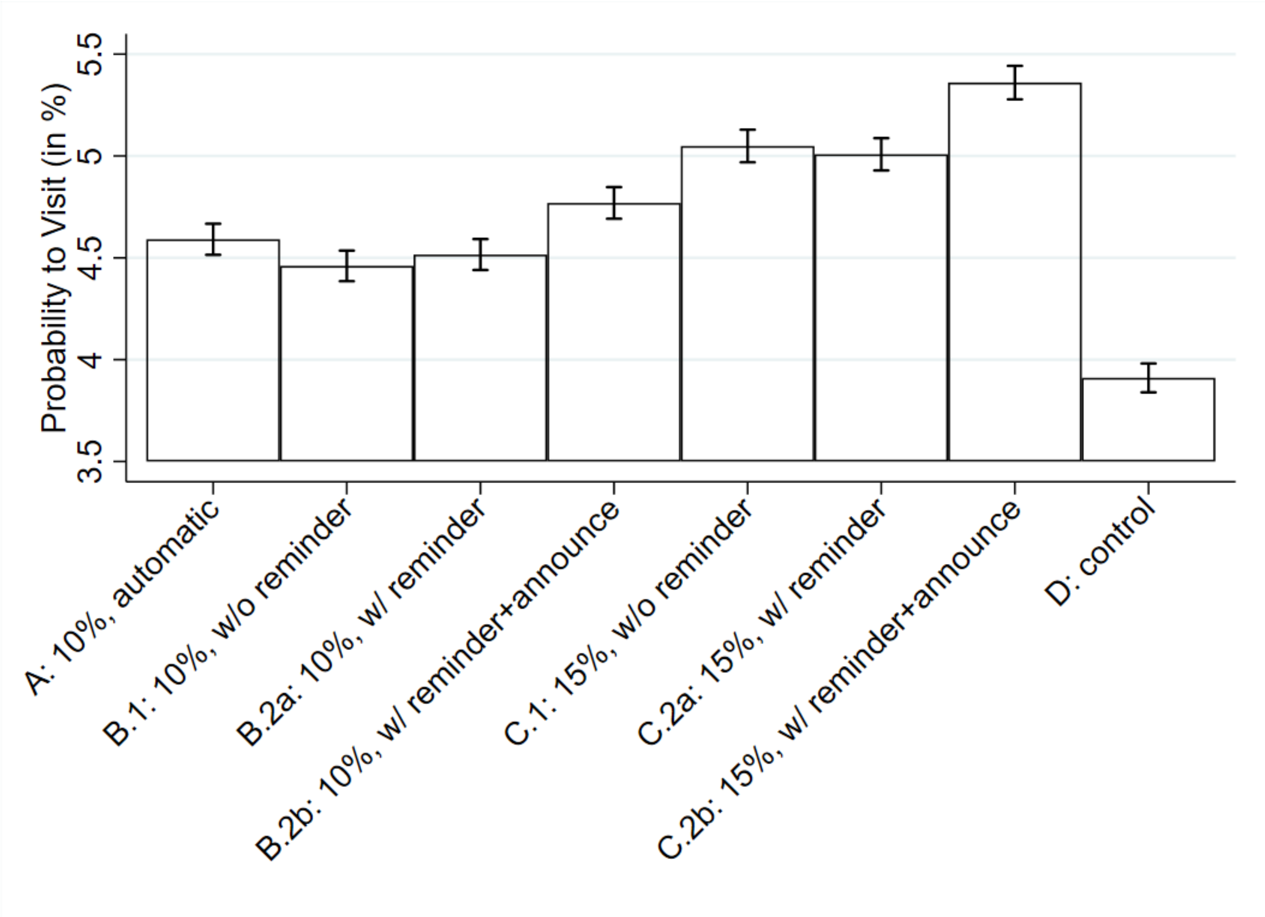

Note: This figure shows the probability of visiting the checkout page at least once. The error bars represent standard errors.