Sophisticated Consumers with Inertia: Long-Term Implications from a Large-Scale Field Experiment

Klaus M. Miller
Navdeep S. Sahni
Avner Strulov-Shlain*

January 22, 2026

Abstract

Are consumers aware of their future inertia? We run a field experiment that offers over a million readers of a European newspaper auto-renewing or auto-canceling contracts. Many consumers are inert yet most anticipate and account for their inertia: though offering auto-renewing contracts benefits the firm in the short-term, it lowers subscriptions take-up by 35% and total subscribers by 23% over 20 months. Inertia's impact on market outcomes depends on consumers' overall awareness of it, which is often ignored by the literature, firms, and policy makers. In our context, consumer sophistication limits the firm from exploiting their behavioral limitations.

*Miller: HEC Paris, millerk@hec.fr; Sahni: Graduate School of Business, Stanford University, navdeep.sahni@stanford.edu; Strulov-Shlain: University of Chicago, Booth School of Business, Avner.Strulov-Shlain@ChicagoBooth.edu (corresponding author). We are thankful for the comments we received from seminar and conference participants at NBER IO 2023 Winter Meeting, Booth-Berkeley Behavioral seminar, ITAM, Cornell research camp, UC Irvine, HBS AI Conference, UChicago AFE conference, Marketing Science Conference, VQMS, 2022 QME Conference, EMAC, Goethe University, ESADE, AMA CBSIG, INSEAD-ESSEC-HEC Seminar, Digital Economics Paris Seminar 2022, Digital Economics Workshop Lausanne, BASS conference, Tel-Aviv University, the Hebrew university, Berkeley, and SITE psychology and economics session among others. The study is registered at the AEA RCT Registry, AEARCTR-0011311.

1 Introduction

Inertia, the tendency of consumers to take no action and remain in the same state, is a well-documented feature of economic decision-making. For example, an individual might not cancel a subscription after a price increase, but will not subscribe under this price if they were not already enrolled.

Inertia has consequences for firms and policy makers trying to assess the functioning of markets. If consumers are unresponsive to declining quality of previously chosen options, it might give incumbents undue advantage. This behavior incentivizes firms to offer choices that are better in the short run but worse in the long run and to design products that intentionally amplify consumer inertia.

The consequences of inertia depend not only on its magnitude but also on consumers' awareness of it. If consumers are unaware of their inertia or are myopic about their future inertial behavior, they will not preempt it and stick with choices that seem good initially but turn out to be worse in the long run[1]. If aware, they can preemptively avoid exploitative situations, retaliate against exploiting firms, or find ways to limit inertia's effects. Awareness can discourage firms from appearing exploitative, mitigating inertia's negative effects. Heterogeneity in future inertia awareness, opens the door to price- or inertia-discrimination in response (Eliaz and Spiegler, 2006).

In this paper we empirically assess consumers' sophistication regarding their future inertia and their responses to inertia exploitation. Specifically: How inert are consumers in subscription settings? How aware are they of future inertia, and how does this shape subscription choices? And what are the effects of these forces on firm incentives and outcomes?

To test whether consumers take their inertia into account, we must first observe their behavior before they choose an option that could lead to an adverse inert state. However, most previous literature focuses on individuals who have already made a choice and become inert, while missing those who avoided such a situation (e.g., Della Vigna and Malmendier (2006); Drake et al. (2022)). Additionally, to assess consumer sensitivity to inertia, we need (exogenous) variation in the degree of future inertia consumers face, which is rarely observed.

We overcome these challenges by running a large-scale field experiment. We randomized subscription offers to 1.4 million readers hitting a large European newspaper's paywall. A reader in our (2 × 2 × 2) experiment is offered a promotional subscription contract that varies along three dimensions: (1) renewal policy — automatically renews into a paid subscription unless canceled vs. auto-cancel, requiring active renewal; (2) trial length — four weeks vs. two weeks; (3) trial price — €0 vs. €0.99. All other aspects of the contract are held constant, including the information consumers need to provide to take up the offers.

[1] Such suggestive evidence is by Shui and Ausubel (2004) showing that consumers are more likely to take low introductory rate credit card offers.

We tracked potential subscribers over two years to observe their engagement with the platform and use the treatment arms to learn about inertia and responses to it.

Comparing take-up during the promo period between those offered auto-renewal and auto-cancel promo reveals if consumers are sensitive to the future possibility of being defaulted into the paid subscription. If they overlook future outcomes, or believe (e.g., due to overconfidence) that they would cancel the subscription on time, take-up should be similar; if they anticipate it, take-up should be lower. Differences in subscription retention after the promo period informs us on actual inertia caused by taking auto-renewal contract. Long-term differences may reveal any persistent pushback consumers might have against the contractual terms.

Varying price and duration helps us isolate mechanisms such as learning or habit formation (does product trial increase long-term subscriptions?), which helps to interpret the effects. Second, it allows us to estimate nuisance parameters, such as the hassle cost of subscribing and the subscription valuation of marginal subscribers, enabling us to map out inertia type distribution and their expectations about inertia.

Our main findings compare the effects of auto-renewal offers to auto-canel. The first key empirical finding is that consumers are less likely to take a contract when offered a future-inertia-exploiting contract. We find that 35% fewer readers take up any newspaper subscription during the promotional time period when *offered* an auto-renewal offer. This indicates that some readers recognize and adapt their behavior to future auto-renewal terms and, overall, prefer the promo that does not convert into a paid subscription by default.

Our second main finding is that some consumers are inert to the extent they stay subscribed longer than they wish. While take-up is lower for the auto-renewal groups, we find that the initial post-trial subscription rate (the proportion of days a reader subscribes to the newspaper) is higher by 20%-38%. Over time, the difference in subscription rates declines and flips sign to become lower for auto-renew after a year. Among those who take up an auto-renewal promo and become full-price paying subscribers, we quantify the actual inaction that causes inertia to be 0.85 – an 85% monthly chance that a consumer does not cancel a subscription they would rather not have (consistent with Einav et al. (2023)). Examining the actual individual-level usage of the newspaper's website, we observe that auto-renewal subscribers engage minimally with the newspaper, further establishing that auto-renewal subscribers do not use their subscription for consumption.

Third, offering auto-renewal contracts discourages readers from ever subscribing with the newspaper throughout the data period. Auto-renewal groups are 23% less likely to ever become paid subscribers, including the promo effects. Even focusing purely on the 20 months post-promotion, they are 7% less likely to subscribe. This effect primarily arises from fewer long-term, non-experimental subscriptions, suggesting consumers are less likely to become loyal readers long after the initial promotion. We do not observe similar long-term exposure effects or dynamic patterns — initial decrease, subsequent increase, and eventual rebound — for other experimental factors. Auto-renewal has a unique long-term deterring effect on potential subscribers decision to engage.

We then use a simple choice model to estimate anticipated and actual inertia types, as well as parameters of pushback and learning. In the model, inertia is driven by either inaction (e.g. due to forgetfulness or procrastination) or switching costs, and consumers further differ by their value of the subscription. Some consumers are non-inert and act as if there are no frictions, and the rest are inert, who with some probability will not take an action they would wish to take. Each inert consumer is either sophisticated, i.e. knows their future inertia parameters, or naive, and thinks they will be non-inert. Therefore, there are three types of readers: Non-inert, inert-naive, and inert-sophisticated.

We find a mixture of types, with a vast majority of sophisticates among the inert. In the population, about 35-55% are non-inert, and the rest are inert with a 81-85% monthly chance of not canceling a subscription they wish to cancel. We estimate that the vast majority, 83-92%, among the inert are sophisticated and know their inertial type. Sophistication means that inert consumers only subscribe in auto-renewal if the added value due to the promotional terms is worth the anticipated risk of being subscribed for a full price for longer than wished, while naifs only think of the promo value.

Firms offer trials to encourage consumers to try the service and learn if they like it. We allow for a simple form of learning, where by having a subscription trial, consumers may learn that they like the product enough to pay its full price. Variations in price and duration within the auto-cancel group identify this, since they cause additional readers to try, and potentially to remain subscribed. We estimate that a four week trial leads about 0.1% of the trial-takers to update their valuation up enough to remain subscribed, which is a relatively small proportion compared to the inerts, indicating that learning has a lower impact on post-trial subscription in our context.

We shed some light on the drivers of inertia. We can rule out switching costs alone as an explanation when consumers have perfect foresight about these costs (Klemperer, 1995), or are completely myopic about them (Dubé et al., 2010).[2] We can also rule out that there is strong habit formation: only a small share of experimentally-induced added subscribers in auto-cancellation remain subscribed; those who remain subscribed due to auto-renew do not use the service; and eventually all the added full-price auto-renew subscribers end up canceling their subscriptions. Our model assumes Klempererian switching costs together with naive and/or sophisticated inaction. The prevalence of sophisticated inerts, which is our main finding, is robust to other assumptions (see Appendix A.6).

[2]A consumer could have stochastic switching costs, that is, she is uncertain about the costs she will face on any given day, but knows the distribution. Appendix A.6 shows that stochastic switching costs with correct expectations do not explain our data well.

We conclude with an attempt to implement naivete-based discrimination, a common prediction of behavioral IO theory (Heidhues and Kőszegi, 2010). We show that treatment effects are indeed heterogeneous and predictable to some extent. We simulate the firm's optimal targeting strategies if it were to maximize either total revenues or subscriptions. We predict for each consumer their valuation and sophistication based on pre-intervention covariates and then estimate the share of sophisticates that are differentially targeted in each strategy. We find that predicted naivete, conditional on predicted valuation, is a positive but weak predictor of optimal assignment to auto-renewal, suggesting a limited significance of contract targeting based on sophistication type in our context.

Our contribution adds three elements to a large literature on consumer inertia (Brot-Goldberg et al., 2021; Choi et al., 2002; Della Vigna and Malmendier, 2006; Einav et al., 2023; Handel, 2013; Heiss et al., 2022; Ho et al., 2017; Hortaçsu et al., 2017; Kong et al., 2022; Madrian and Shea, 2001, among others): (1) direct measurement of sophistication, (2) heterogeneity in inertial types, and (3) long-run push-back effects.

Closely related work by Einav et al. (2023) documents high inertia among *existing* subscribers using exogenous payment card expiration. Our contribution is to study consumers *before* they subscribe, which reveals a strikingly different picture: auto-renew contracts reduce, not increase, total subscribers at any time horizon. This finding stands in sharp contrast to what conditional comparisons would suggest — among those who take the promo and stay as paid subscribers for at least a month, the conversion rate is 1200% higher for auto-renew takers relative to auto-cancel takers. However, this conditional comparison is misleading for two reasons: fewer consumers take the auto-renew offers, and auto-cancel groups are far more likely to subscribe through a different contract after declining the promotional offer. Consistent with this, our model estimates strong selection: auto-renew takers are up to 30% more likely to be inert and up to five-fold overrepresented as naifs compared to their population shares.

We add field evidence to a literature studying consumers' sophistication about their biases, mostly in lab experiments or using surveys. Augenblick and Rabin (2019); Chaloupka et al. (2019); Bai et al. (2021); Ericson (2011) study mostly present-bias and some projection-bias. Rodemeier (2025) studies consumer sophistication with respect to buy baits. Rodemeier (2025) makes an important contribution by studying consumer sophistication with respect to buy-bait offers in a field experiment. He shows that consumers are aware of their lower likelihood of redeeming a rebate, emphasizing the short-term dynamics, and also documents sophistication through reduced engagement in the short run. Our paper complements this work by focusing on the long-term effects of inertia, also in a naturalistic setting. We also contribute to a much smaller literature that examines consumers' response to future inertia, and how it affects companies' decision making (e.g., Reme et al., 2021). We consider longer-term behavior and find significant adverse consumer reactions to inertia-inducing contracts.

Finally, our paper speaks to the conceptual way of incorporating inertia in models and empirical work. Often inertia is operationalized as a transitory utility term to which consumers are fully naive (e.g., a brand coefficient as in Dubé et al. (2010)), while we find forward looking behavior. In the behavioral economics literature, inertia is an outcome of preferences that include either present-bias (DellaVigna and Malmendier, 2004), over-confidence (Grubb and Osborne, 2015), inattention (Brot-Goldberg et al., 2021; Hortaçsu et al., 2017), or habit formation (Allcott et al., 2021). We do not distinguish between every possible source of inertia, but find the most support for inattention, and can categorize consumers into different types (in the tradition of O'Donoghue and Rabin (1999, 2001)). We are also able to empirically address the (low) possibility for naivete-based discrimination within our setting and data, a key theoretical prescription that has yet to be tested empirically.

Our paper also contributes to the marketing literature on firm policies in contractual settings (Goettler and Clay, 2011; Ascarza et al., 2016; Yoganarasimhan et al., 2022; Datta et al., 2015; Yang et al., 2020). We differ by explicitly varying inertia-related contractual terms and assessing the degree of consumer sophistication. Our findings could be of relevance for businesses and regulators. While many companies try to make it harder for consumers to leave their services thinking that it increases their profits ("sludges" in Thaler and Sunstein (2021) language), we provide evidence that such practices, even if mild, can backfire due to two reasons. First, exploiting future inertia reduces initial take-up; second, exploiting future inertia pushes new consumers to disengage from the company completely. Our finding of an economically significant negative reaction to auto-renewal contracts is relevant for regulatory agencies such as the FTC who worry about deceptive practices in subscription selling.[3] Admittedly being one specific instance, our evidence stands against the common wisdom and findings in the past literature which has assumed that people "passively" accept defaults (Benartzi et al., 2017). People in our study are susceptible to defaults, but most are also aware of these effects and successfully avoid them. Our analysis suggests that rather than exacerbating inertia exploitation, businesses that can credibly promise easy cancellation and timely reminders might end up with more consumers and larger revenues.

2 Model

2.1 Inertia

Before specifying the consumer problem, we define what we mean by *inertia* and *sophistication*.

[3] In the policy literature such practices are referred to as negative options, and the regulatory concerns about consumers getting deceived and being economically harmed by selling of negative options are widely discussed. In October 2024 the FTC announced its rule of "click to cancel" (*Negative Option Rule: Final Rule*, 2024) requiring companies to be truthful and transparent, and to allow consumers to cancel subscriptions with as much ease as signing up for it.

An individual is *inertial* if being in a state at period t increases the probability of remaining in that state at $t + 1$ conditional on preferences. For instance, someone subscribed in period t is more likely to be subscribed in $t + 1$ than someone who is not.

Several main mechanisms generate inertia. First, inertia may arise from a cost-benefit analysis driven by the costs incurred for taking a state-changing action (e.g., effort), versus the benefits of changing it. Past choices can have lingering effects on current costs. Some examples are switching and hassle costs which make it *harder* to change states. Conversely, preferences can be state-dependent. For example, addiction lowers the preference for the alternative state, habit formation increases preferences toward an action previously taken, or information is revealed through trying a state. Finally, inertia may be driven by *inaction* due to forgetting, inattention, or "autopilot" behavior (e.g. Brot-Goldberg et al. (2021); Camerer et al. (2018)). For example, forgetting to act or being reluctant to devote any thought to actually doing the cost-benefit analysis due to its "complexity".

A main empirical challenge in showing and estimating inertia is preference heterogeneity (Dubé et al., 2010): because individual persistence in choices may reflect stable preferences rather than state-dependence, we cannot exclude they would have chosen the same option at $t + 1$ regardless of their choice at period t. Heckman (1981) refers to this as spurious state dependence. Our experiment allows us to overcome this challenge by randomizing contract offers, allowing identification of inertia independent of preferences.

In what follows, we incorporate both sources of inertia of costs and inaction. Preferences are heterogeneous and a main driver of take-up, but are comparable across treament groups. We assume the existence of *costs* for taking actions – to subscribe, cancel, or renew. We model *inaction* as the probability of not taking an action at any given period by consumer i. This is a descriptive parameter that may reflect different mechanisms. Namely, inaction may be due to a psychological barrier to making a decision, due to forgetting to act, or due to a time-inconsistent desire to postpone an action to a later period driven by present bias.

Further, we also allow for two forms of changing preferences due to experience and exposure: (i) learning: Each day of subscription and access to premium content may lead consumers to discover that they wish to become full-price subscribers and (ii) a utility penalty ("spite") from being *offered* an auto-renewing contract (even if not taken). We interpret it as a form of inference about the business as a whole, lowering the value of engagement.

Finally, inert consumers can be *sophisticated* or naive, the main focus of this paper. Consumers are sophisticated if they accurately know that they might be inactive in the future; consumers are naive if they erroneously think that they will be non-inert in the future.[4] Therefore, we allow for potentially incorrect

[4]We exclude partial naivete (O'Donoghue and Rabin, 2001) for simplicity and practicality reasons. The variation in the data will not allow to separate between predicted and actual inaction which is necessary for estimation of partial naivete.

beliefs about the future value of the inaction and action-cost parameters at the individual level. We denote the perceived parameters with adding a tilde sign (e.g., $\tilde{\phi}$) and allow them to differ from the actual realized inertia-driving parameters.

2.2 Setting

We now describe the formal model, which aligns with our empirical context.

Time is discrete and infinite (because contracts never expire): $t = 1, 2, \ldots$. Each consumer i has a fixed per-period value from the subscription, denoted by v_i, drawn from a distribution with CDF, F. Subscription prices are weakly increasing: $p_t \geq p_{t-1} \geq 0$, stabilizing at $p_t = p$ for $t \geq T$. This price trajectory reflects the common practice of introductory promotional prices, which also appear in our empirical setting.

There are three possible actions – subscribing, renewing, and unsubscribing. Initially, the consumer can subscribe; a subscribed consumer can unsubscribe; an unsubscribed consumer can renew. We assume that initial *subscribing* incurs a cost c^s (e.g., giving credit card details and setting up a reader account); *unsubscribing* has a cost c^u (e.g., finding out how to unsubscribe or some true hassle); and *renewal*, if one is needed in case the contract otherwise terminates, incurs a cost c^r (e.g., clicking "renew" on an email or browser pop-up) which for simplicity, we assume is costless, i.e. $c^r = 0$.[5]

A contract is defined by an auto-cancel period z: if no action is taken at z, the subscription terminates. Setting $z = \infty$ corresponds to auto-renewal. The contract can still be renewed in period z or afterwards. Offering an auto-renew contract can shift utility by a negative push-back "spite" factor α.

During any period (except z), a subscriber fails to act with probability of inaction $\phi \in [0, 1)$. Therefore, they may remain subscribed even if they wish to cancel. When planning ahead, a consumer *perceives* their probability of inaction as $\tilde{\phi} \in \{0, \phi\}$.[6]

In summary, a consumer's temporal net value of a subscription is $v_i - p_t$ plus any associated costs of actions and a continuation value from remaining subscribed or not. We assume that time is discounted with a discount factor δ.

We find the consumer's plan using backward induction. Since prices are nondecreasing and known, and the value is stable, the consumer only needs to ask themselves if they would cancel or renew at period t if they were subscribed at the previous period. In period 1, consumers choose whether to subscribe. In accordance with our focal newspaper's policy, cancellation means that the current period is still paid and the contract

[5] We also find it realistic – a renewal only requires a single prompted click since user details, prior agreement to terms, and payment information are already on file. In addition, firms actively try to make renewing as easy and salient as possible to renew.

[6] We assume that inaction does not apply for renewal: someone who wants to read the paper, will go to the website and be prompted to renew their subscription if it lapsed. A user who does not go to read the newspaper likely also does not want to renew. This assumption may matter for our results if it were to lead to consumers not returning after auto-cancel. However, for the above reasons we think the assumption is realistic.

terminates at the end of the period. If continuing, their utility is $\frac{v_i - p_T}{1-\delta}$. If canceling, the (perceived) utility is:

$$\sum_{\tau=0}^{\infty} \tilde{\phi}_i^{\tau} \delta^{\tau} \left(v_i - p_{T+\tau} - (1 - \tilde{\phi}_i) \tilde{c}^u \right)$$

reflecting delay due to inaction. Consumers choose the option with higher expected value (if active).

Solving backwards, the consumer faces the same choice every period t – to remain subscribed, or to try to cancel. In an auto-renewing contract, the outcome is that, conditional on subscribing, the consumer will remain subscribed forever or until a specific price increase takes place. The behavior under an auto-canceling contract is different. At the auto-cancel z period, subscribers ask themselves if they wish to renew or cancel for free. This implies that the choice is between a value of 0, or a renewal and canceling at a later time. As such, an auto-cancel contract has two benefits: it saves the lower-value consumer the cost of cancellation and eliminates the risk of being subscribed for a longer-than-desired time.

This solution means that at $t = 1$, each consumer has a plan for when to cancel (or start trying to) if they were to subscribe. We call that optimal period k, with an expected value (net of subscription costs) $U(v_i)$. This period k may reflect a plan to cancel during the initial promo period, staying longer and trying to cancel at a certain price increase, canceling automatically and not renewing, or staying subscribed forever (at which case $k = \infty$).

In addition, we allow for learning or habit formation by subscribing. We assume the consumers know that with probability λ, if they subscribe to the promo, they might learn that their value of the subscription equals $\bar{v}$, which is high enough to remain subscribed at full price.[7] This gives a net value of $U(\bar{v}) = \sum_{t=1}^{\infty} \delta^{t-1}(\bar{v} - p_t)$. A consumer i will subscribe at time 1 if the expected value of subscribing is greater than the subscription cost:

$$(1 - \lambda)\, U(v_i) + \lambda U(\bar{v}) \geq c^s \tag{1}$$

Finally, plans and reality can diverge. Subscribers who plan to cancel at k, given $\tilde{\phi}$, might end up acting differently because their actual inaction ϕ is higher. For example, someone with a low value v_i who perceived herself as always active, with $\tilde{\phi} = 0$, might choose to take a promo offer, thinking she will cancel before the promo ends. However, if her true inaction is a higher $\phi > 0$, she may end up subscribed for longer by failing to act in time. Had she known her true ϕ, she would have not subscribed in the first place.

In summary, subscription paths are determined by the value of subscription v_i, time preferences δ, spite α, the probability of learning λ, the costs of subscription and cancellation c^s and c^u, and the perceived and

[7] This implies that before trying the subscription, the consumer believes his valuation is $\bar{v}$ with probability λ and v_i with probability $1 - \lambda$, where v_i is distributed with the CDF F in the population.

actual inaction $\tilde{\phi}$ and ϕ.

We estimate this model in Section 8.2 to recover structural parameters.

3 Empirical Setting

Our study was conducted in cooperation with a large European publisher that wishes to remain anonymous. The publisher is one of the largest daily newspapers in its market with strong readership in several European countries. The publisher represents a highly reputed quality news outlet similar to the New York Times or the Washington Post or the Guardian. It publishes daily coverage of politics, economics and business, sports, local news, culture, society, science, digital, working life, and travel. In addition to the printed newspaper, the publisher has a digital platform which provides daily online news. In 2018, approximately 12 million unique readers visited our publisher's digital platform.

Content on the digital platform is divided into three categories. The first is "always free" to any reader. This content includes the main homepage, as well as the separate section homepages, agency news, breaking news, and also other commodity news. The second is "always paid", which is exclusive to subscribers and includes premium content from the printed newspaper and commentaries. The third is "metered" content, which falls under a metered paywall. Readers can access up to 10 news articles per week for free and then are prompted to subscribe. Users referred from search engines and social media are subject to the same rules as direct visitors. In addition to subscription revenue, the publisher earns revenue from displaying ads to its readers.

Tracking on the digital platform occurs via reader logins and cookies, and complies with the European General Data Protection Regulation (GDPR). When a reader first accesses the platform, they are assigned a cookie ID, which persists across visits unless deleted. However, cookie-based tracking is imperfect: users can delete cookies at any time, and a single reader may have multiple cookies if by accessing the website from multiple devices (see Lin and Misra (2022)).

Pricing and Contracts The newspaper offers multiple subscription options. The most commonly purchased is a daily pass, which grants one-day full access for €2. The second most common are short-term promotional contracts (lasting up to one month), including the experimental contracts described below. These are offered to readers with no prior paying subscriptions. Third are standard ongoing subscription contracts, which auto-renew unless terminated. The regular subscription prices are €19.99 for the first two months, and €34.99 per month thereafter. Additionally, the publisher offers one-year contracts, though those are rarely chosen. Figure A.3 presents the market shares of contract types.

Canceling subscriptions Readers are notified of subscription terms and cancellation procedures prior to sign-up. A subscriber may cancel at any time, with cancellation taking effect in the next billing cycle. Until then, access remains active. Cancellation is non-trivial and can be made via the publisher's call center, using the "contact the publisher" website page, or by mailing a letter or emailing.[8]

4 Experimental Design

The field experiment was designed to address our research questions and to support the publisher's interest in using randomized control trials to convert new readers into digital subscribers. The experiment ran from May to August 2018, with follow-up data collected until April 2020. It allows us to document and quantify inertia and perceived inertia, and to examine its underlying drivers.

4.1 Participants and Randomization

Half of "new" potential subscribers who hit the paywall, either by exhausting their quota of free metered articles or by clicking on an "always paid" article, entered the experiment.[9] These readers were randomly assigned to one of eight experimental treatment groups and shown a corresponding subscription offer. The publisher defines a new subscriber as someone who has not previously paid a full monthly price (€34.99).

Randomization was implemented at the cookie level, and a reader's assigned experimental group remains constant across visits. Table A.1 shows balance in pre-treatment engagement. After the promo trial period, every reader, irrespective of treatment, could subscribe for €19.99 for two months, followed by the regular monthly price of €34.99.

4.2 Experimental Contracts

The experiment varied three contract features in a full factorial $2 \times 2 \times 2$ design:

1. **Subscription renewal:** After the trial ends, the subscription either auto-renews by default or auto-cancels. Under auto-renewal, readers stay subscribed unless they cancel. Under auto-cancel, contracts are terminated by default after the trial but they can choose to subscribe again the next time they hit the paywall, via a homepage pop-up or an email link.[10] In all cases, readers confirm their existing payment details, entered at promo take-up, to renew a subscription.

[8]Cancellation modes in our context are very similar to The New York Times, as seen here: `https://help.nytimes.com/hc/en-us/articles/115014893968-Terms-of-sale#cancel` (accessed on Jan 11, 2022).

[9]The other half were shown a status-quo promotional offer of two weeks for free auto-cancel promo, which did not require credit card details (group "I").

[10]Approximately five days before the trial ends, an email prompts renewal via a one-click link. If no action is taken, several follow-up emails are sent.

Table 1: Experimental offers

Experimental group	Renewal	Promo Duration	Promo price
A	Auto-renewal	4 weeks	€0
B	Auto-renewal	4 weeks	€0.99
C	Auto-renewal	2 weeks	€0
D	Auto-renewal	2 weeks	€0.99
E	Auto-cancel	4 weeks	€0
F	Auto-cancel	4 weeks	€0.99
G*	Auto-cancel	2 weeks	€0
H	Auto-cancel	2 weeks	€0.99

2. **Duration:** The trial lasts either two or four weeks.

3. **Promotional Price:** The trial period costs either €0.99 or €0. Prices and available contracts after the trial period are identical across groups.

The eight treatment combinations and corresponding experimental group labels are shown in Table 1. Offers were shown at the paywall; checkout steps were identical across arms and detailed in Appendix A.3.1. Due to a technical error, readers in experimental group G were not required to enter their payment information, leading to an invalid experimental condition.[11]

4.3 What the experiment identifies

By varying offers between auto-renewal and auto-cancellation, and observing the effects on take-up, we capture the effects of participants foreseeing their future inertia. If perceived future inaction and unsubscription costs are small, there will be no differential take-up. Focusing on the subscription patterns after the promo, we capture the actual inertia of those who take an auto-renewing contract despite not valuing it at full price. Comparing the auto-cancel treatment effects on subscriptions at different time periods, and leveraging the price and duration treatments, informs us about learning. Finally, the full experimental variation allows to estimate our model and quantify sophistication and inertia terms, as we explain and do in section 8.

5 Data

We have two data sources provided by the newspaper. The first is every cookie's browsing history 14-days prior to being introduced to an experimental treatment and 27-days after leaving the experimental treatment,

[11]Effectively, these readers receive the status-quo offer. This error was corrected in the final phase of the experiment; we use this subset for robustness checks as mentioned in Section 6. Separately, during an early phase of the experiment's implementation, some users were non-randomly reassigned to different treatment offers after getting exposed to one. We exclude from our analysis data from this phase.

giving us an observation window of at least 42 days of browsing history per cookie id.[12] The second data are customer relationship management (CRM) data on all subscriptions and contracts, both experimental and regular contracts, from April 2018 to April 2020.

5.1 Usage Data

The browsing history includes each page visited by a reader (identified by a cookie) along with its timestamp, page type (open, metered, paywalled), and the subscriber identifier if the reader was logged in to their account (even if the subscription is free).[13] Moreover, the browsing history records whether a reader was exposed to one of the experimental offers during a page visit and, if so, which specific offer was presented.

Table 2: Summary Statistics

	Raw	All Subscribers	Promo Period Subscribers	Exp Takers
Number of readers	4,131,277	1,389,869	1,389,869	1,389,869
Number of subscribers	36,545	6,196	3,844	2,434
Total revenue (Euros)	1,972,198	403,477	165,140	91,859
Promo period sub share (%)		0.277	0.277	0.144
After promo sub share (%)		0.279	0.11	0.078
Total sub share (%)		0.446	0.277	0.175

From this data, we construct several key elements:

Experimental exposure For each reader, we identify the first exposure to an experimental offer and define that date as "day 0" in the experiment. Readers may be exposed again later, but assignment and measurement are based solely on the first exposure. Appendix Figure A.1 shows treatment group sizes over time.

Intent-to-treat (ITT) assignment Is identical to experimental exposure to maintain a valid ITT design. For example, if a reader in a two-week treatment group first saw an offer on April 1 but subscribed on April 8, the promotional period is coded as April 1–14, not April 8–21.

Cookie consolidation We consolidate multiple cookies belonging to the same subscriber using co-occurrences in the usage data, and merge multiple subscribers using the same cookie when they appear to be the same individual.

[12]While all readers are tracked for 6 weeks, 14% of readers (291,837) are tracked up to 23 weeks though the reason and selection for longer tracking is unclear to us.

[13]Some readers become subscribers first during the time window, while others had a subscription before and are thus identified in the system. However, if a reader only subscribed for the first time outside of the four-week usage data time window after their exposure, we are unable to link that subscription to the reader.

After these steps, we obtain a unique *reader identifier* associated with a unique subscriber identifier (if applicable). Table 2 presents summary statistics. Column 1 includes all readers in the raw data, including those not exposed to any promotional offer. Columns 2–4 restrict to readers exposed to an experimental offer: "All Subscribers" includes every subscription made by an exposed reader and is the main sample used to estimate intent-to-treat effects; "Promo Period Subscribers" retains the same sample, but only counts subscriptions by those who subscribed during the defined promotional period; "Experimental Takers" counts subscription by those who subscribed to one of the experimental offers at any time. We use these samples in the structural estimation part. Note that we keep all exposed readers in each of these three groups as the "total market" denominator in all analyses we conduct.

Our full experimental sample contains 1,389,869 readers. Of these, 6,807 (0.45%) held a subscription at some point during the tracking window (from 14 days before exposure up to 21 months after). A total of 0.28% subscribed during the promo weeks, and 40% of those (0.11%/0.28%) subscribed after the promotional period. Table A.2 provides additional breakdowns of readers, subscriptions, usage, and revenue by treatment arm. Cookie-based tracking can fragment users across devices; we consolidate for subscribers but not for non-subscribers. This may attenuate levels but should not bias relative treatment effects; see Appendix A.3.2.

5.2 Subscription Data and Consolidation

The second dataset is the publisher's customer relationship management (CRM) data, which records all signed contracts between April 2018 and April 2020. Each contract is linked to a subscriber via a unique contractor ID'[14] and includes the start date, end date, and revenue collected. Crucially, contract codes allow us to distinguish experimental contracts from the many other subscription products the publisher offers. Appendix Figure A.3illustrates this multiplicity: among the 6,196 experimental participants who subscribed, many chose non-experimental contracts (experimental contracts are highlighted with black boxes). This abundance of alternative products matters for the interpretation of the results: our main ITT results use as an outcome a subscription for *any* contract, not just the experimental one. Therefore, we also capture substitution to or away from other products.

We merge the assignment data with the CRM data at the subscriber level. We construct a daily panel indicating whether a reader is subscribed on a given day and the effective price paid. We aggregate these outcomes into longer time periods (e.g., monthly). We detail this in Section 6, where we use this merged dataset to estimate policy relevant overall intent-to-treat effects using all subscription activity. In Section 8,

[14]less than 1.2% of subscribers have multiple contractor identifiers. We identify those by observing two contractor ids with a shared cookie. That can happen if someone creates multiple readers, for example associated with different email addresses. We consolidate those and assign them a single subscriber id.

we focus on model estimation, restricting to readers who subscribed during the promotional period or other more restricted samples.

6 Results

We begin our analysis by estimating intent-to-treat measures of readers' overall subscription to the newspaper across the experimental groups, by time period. The sample includes all consumers exposed to one of our experimental contracts, and we count any subscription after this exposure, regardless of when it is taken or contract type. Doing so allows us to accurately estimate the impact of various contract terms on consumer behavior and assess the firm's overall outcomes. This broad focus also enables us to avoid spurious effects that might arise if individuals in different experimental groups switch to other non-experimental contracts differentially.[15]

Our main outcome measures are:

Subscription rate– the proportion of days a reader is subscribed during a period; **Extensive margin**– whether the reader is subscribed at all during the period; **Cumulative revenue**– total revenue from the reader by period; and **Pageviews**– number of articles read.

To structure the analysis, we divide time into discrete periods: the two weeks before exposure (placebo), the first two weeks of the promotional period (to align two- and four-week offers), followed by monthly intervals through 20 months after the promo. We also analyze the cumulative post-promo period as a whole and the total time horizon effects (including promo).

We set up the analysis described below in the form of the following regression for each period t:

$$y_{it} = \alpha_t + \beta_{1t}\text{Auto-renewal}_i + \beta_{2t}\text{One-euro}_i + \beta_{3t}\text{Four-weeks}_i + \epsilon_{it} \tag{2}$$

where y_{it} is one of the outcome variables for reader i in period t, and the β coefficients estimate the per-period marginal effects of the experimental factors: auto-renewal (vs. auto-cancel), €0.99 (vs. free), and four weeks (vs. two weeks).

We exclude group G from the main analysis due to its implementation issues.[16] Group G is later used together with the status quo offer for out-of-sample value prediction. Because treatment assignment probabilities varied across experimental phases, we weight each observation by the inverse of the assignment probability to ensure equal group representation. Results are not sensitive to this weighting.

[15] Another practical reason is that we cannot consider just the experimental contracts, since the auto-cancel contracts renew as different contract types which cannot be precisely tracked.

[16] Focusing only on data from the fully implemented phase 3, with group G, yields consistent results in sign and magnitude, though noisy. A total of 228,416 readers are only 16% of the total sample.

6.1 Auto-renewal vs. Auto-cancel

Figure 1 plots the per-period ITT effects of offering a promotional auto-renewal contract as opposed to an auto-cancel contract on subscription behavior, using the estimated coefficients β_{1t} in equation (2). Consistent patterns are visible in the raw data on subscription levels, presented in Figure A.5.

Panel A (Figure 1a) shows effects on the **subscription rate**. Pre-treatment estimates are near zero, confirming balance. During the promo period, offering auto-renewal reduces the subscription rate by 0.12 percentage points—a 41% decrease relative to the auto-cancel mean.[17] After the promo, the effect reverses: subscription rates are higher for auto-renewal groups for several months, consistent with inertia.[18] Over time, this effect declines, and by month 12, the auto-cancel group again shows higher subscription rates. By month 20, this difference is statistically significant. Appendix Figure A.4 shows cumulative revenue is initially (post-promo) significantly higher under auto-renewal, but statistically insignificant by month 16. Nonetheless, cumulative revenue is always higher under auto-renewal by about 2.5c per potential reader.

Panel B (Figure 1b) shows different effects on the **extensive margin** (whether a reader subscribes at all). During the promo, auto-renewal reduces the likelihood of subscription by 35%. Unlike the subscription rate, this margin does not show positive effects post-promo. Over the entire 20-month post-promo horizon, the auto-renewal contract leads to a 7% reduction in the number of unique subscribers. We also observe a broader reduction: the total number of subscribers (across all time periods) is 23% lower in the auto-renewal arms. In other words, roughly one-quarter of potential subscribers who would subscribe in the years following an auto-canceling promo offer do not subscribe when offered an auto-renewing one.

6.2 Other Experimental Factors

Free vs. €0.99 Figure 2 shows the corresponding effects of changing promo price. The estimates show that increasing the price from free to €0.99 reduces subscription-rate during the promotional time period by 9% and causes 9% fewer readers to subscribe during the promotion period. As expected, readers are more likely to take up a subscription if it costs less. However, this difference fades away quickly; we do not observe any significant effect of the promotional price starting after the promotion on the extensive margin or the subscription rate. This implies that increasing subscription trial by decreasing price does not lead to significantly higher long-term subscriptions.

[17] A 0.1pp lower subscription rate implies one fewer subscription day per 1,000 potential days.

[18] The rates may be higher due to inertia overcoming the initial lower subscriptions, or from readers postponing their subscription. Focusing on promo period subscribers, shows that the higher rates are mostly among those who first subscribed during the promo period (Figure A.5 dashed lines).

Figure 1: Effect of Auto-Renewal Relative to Auto-Cancel Contracts on Overall Subscription Behavior

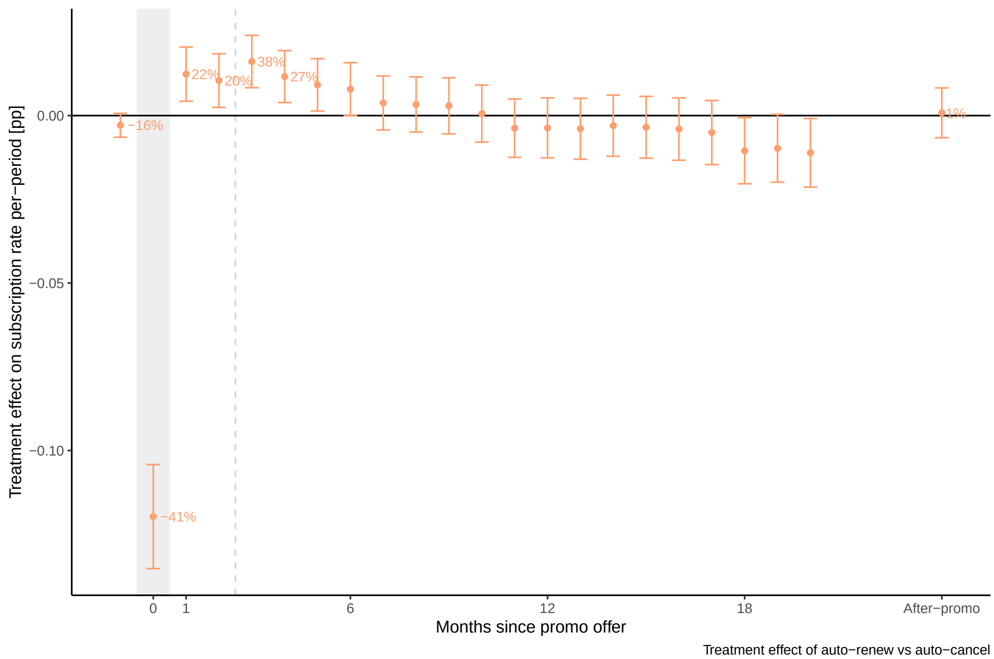

(a) Effect on Subscription Rate (Proportion of Days a Reader Subscribed)

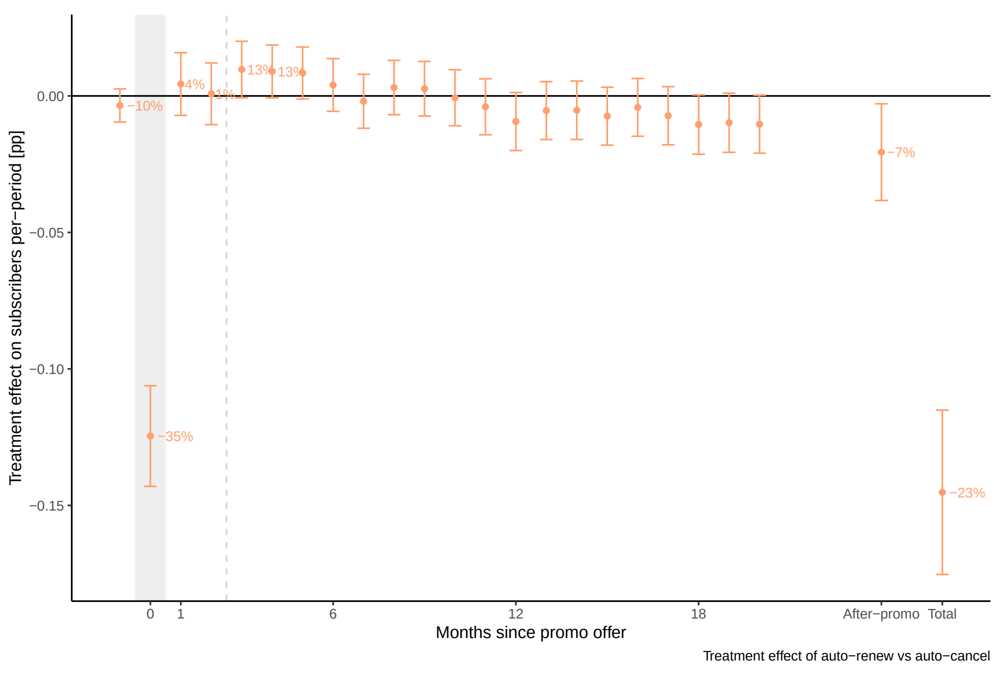

(b) Effect at the Extensive Margin (Whether the Reader Subscribed At All)

Notes: The figures plot the estimated average per-period intent-to-treat effects of offering an auto-renewal relative to an auto-cancel contract on readers' subscription behavior. Specifically, we plot the estimated coefficient β_1 *from equation (2) for every month. Month 0, shaded in gray, is the promo period (two weeks). The left-most points at month -1 are before readers hit the paywall. Month 1 is the first month after the promo ends, and so forth. The point "After-promo" aggregates across all after the promo time periods; "Total" aggregates from 0 until the end. Percentages next to the effect size compare the effect to the mean level of the omitted group. The error bars show 95% confidence intervals. Standard errors are clustered at the individual reader level.*

Four weeks vs. two weeks Figure 3 shows the effects of increasing the trial duration from two to four weeks. The estimates indicate a 4–5% increase in both the subscription rate and the number of subscribers during the promotional period, though these effects are not statistically significant. Unlike the price effect, however, this increase in subscription rate does not fade over time. While the effect on the extensive margin is weaker, there is some evidence of a persistent (though modest) long-run effect of a longer trial. The stronger effect on subscription rate relative to subscriber count suggests that some subscribers stay subscribed longer, consistent with learning about higher product value, an issue we return to in Section 7.3.

Discussion Comparing the auto-renewal vs. auto-cancel effect with the same effect of price or duration change shows the distinct consumer response to auto-renewal. While auto-renewal causes an average decline in promo take-up, similar to a price increase,[19] it causes an opposite effect on subscription rates a few months after the promo. The price and duration treatments lead to same-sign effects across the time periods. In addition, we observe that auto-renewal has a unique impact on take-up behavior, as it simultaneously reduces extensive margin take-up both during and after the promotional period and lowers overall extensive margin subscription across both periods. These patterns suggest a unique consumer 'push back' response that is not triggered by other experimental treatments.

7 Channels of Inertia

To interpret the above patterns and discipline our model, we examine three mechanisms in this section: post-promo usage (consumption value vs. inertia), a long-run valuation penalty from being offered auto-renewal ("spite"), and learning from trial (price and duration variation).

7.1 Subscription versus Usage

If the increased post-promo subscription behavior caused by auto-renewal is actually unwanted, we expect readers to get little utility from continued access. Using website usage data, we examine how much subscribers use their subscriptions, which is our proxy for consumption utility.

Recall that our usage data spans six weeks for each reader. Figure 4 plots the average page visits and subscription rates among promo takers for each day in this time span. It uses data for readers who subscribed to either a two weeks €0.99 auto-renewal promo or a two weeks €0.99 auto-cancel promo during the first days after exposure, so we can observe the promo time ending in the middle of our four weeks post treatment

[19] A back-of-the-envelope calculation comparing the take-up response of auto-renew vs auto-cancel to a €0.99 to a free promo, gives that the auto-renew take-up effect is equivalent to about a €4 promo price increase.

Figure 2: Effect of €0.99 Relative to Free Promotional Contracts on Overall Subscription Behavior

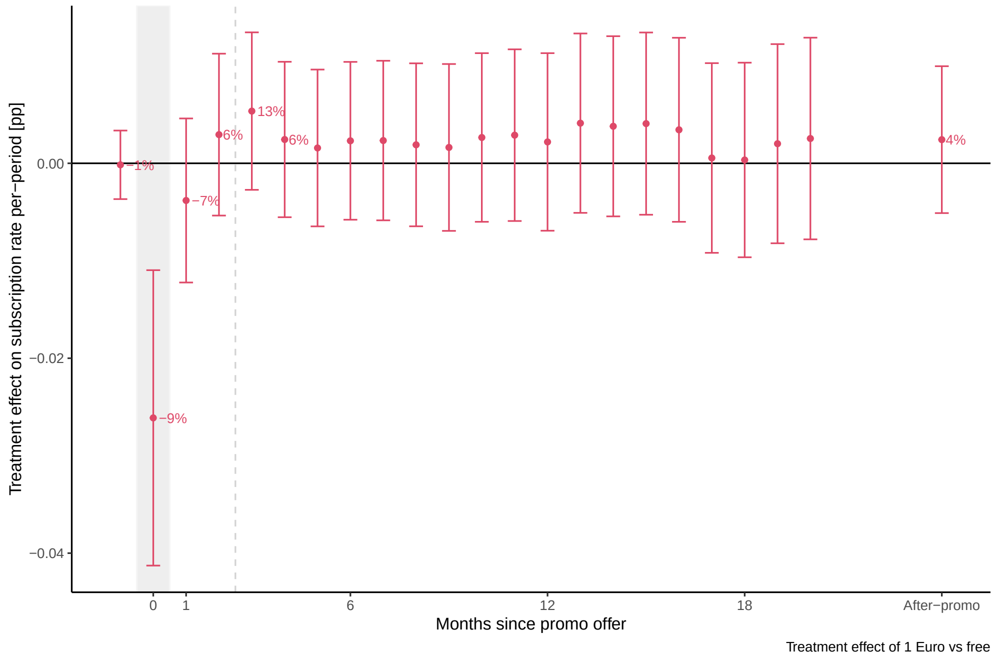

(a) Effect on Subscription Rate (Proportion of Days a Reader Subscribed)

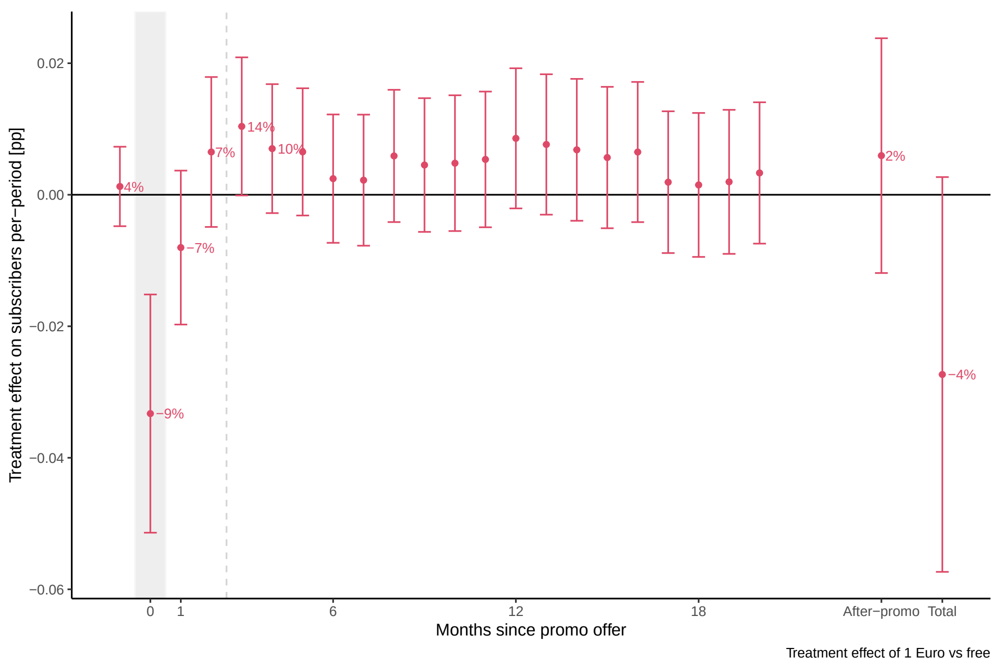

(b) Effect at the Extensive Margin (Whether the Reader Subscribed At All)

Notes: The figures plot the estimated average intent-to-treat effect of serving a promotional contract costing €0.99 relative to a free contract on readers' subscription behavior. Specifically, we plot the estimated coefficient β_2 from equation (2) for every month. Month 0, shaded in gray, is the promo period (two weeks). The left-most points at month -1 are before readers hit the paywall. Month 1 is the first month after the promo ends, and so forth. The point "After-promo" aggregates across all after the promo time periods; "Total" aggregates from 0 until the end. Percentages next to the effect size compare the effect to the mean level of the omitted group.The error bars show 95% confidence intervals. Standard errors are clustered at the individual reader level.

Figure 3: Effect of Four Weeks Relative to Two Weeks Promotional Contracts on Overall Subscription Behavior

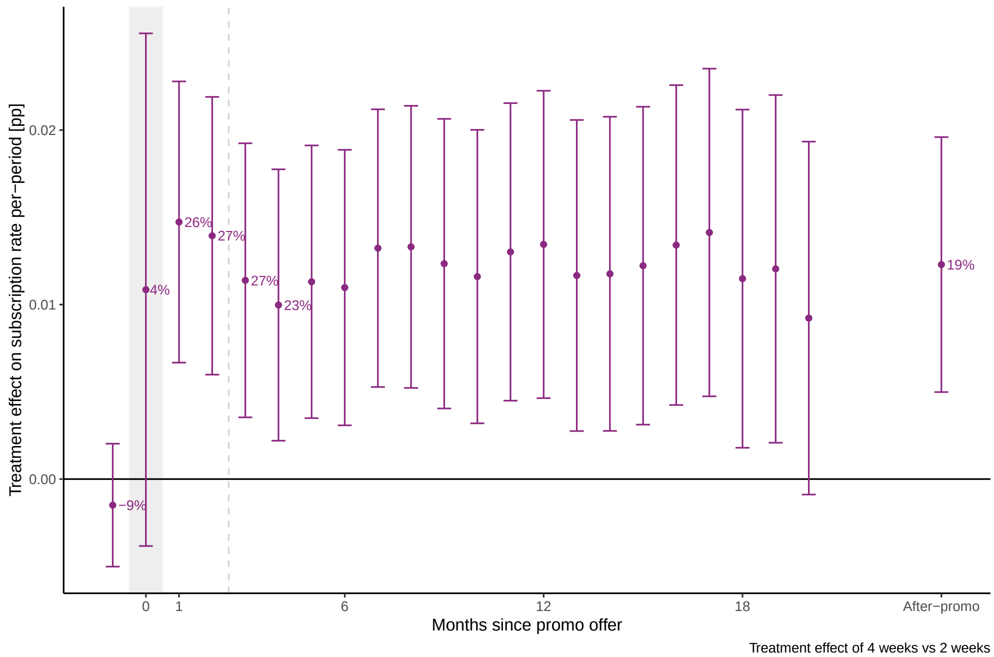

(a) Effect on Subscription Rate (Proportion of Days a Reader Subscribed)

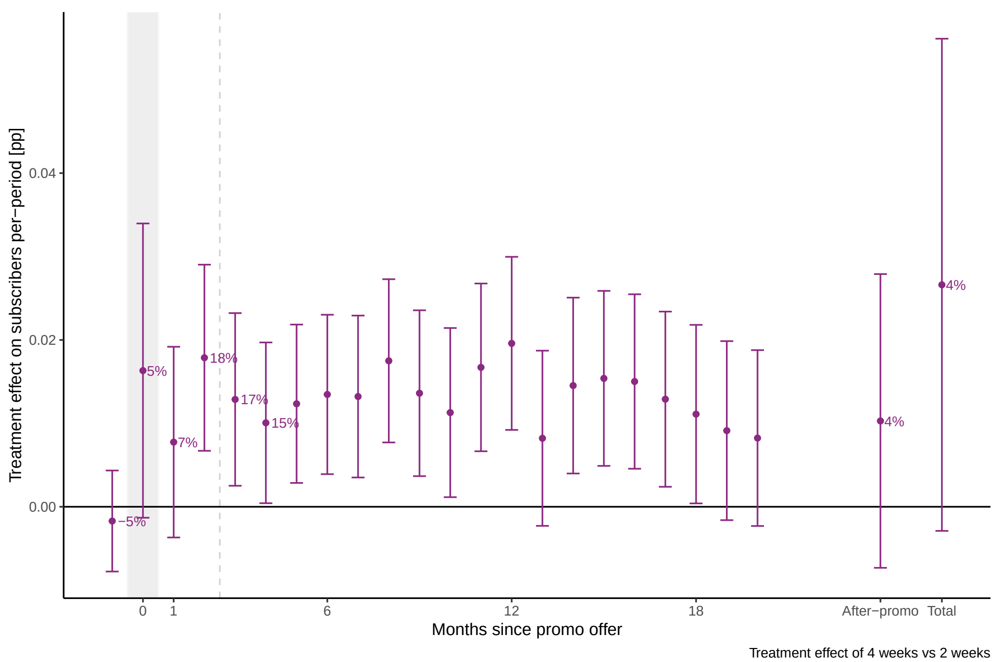

(b) Effect at the Extensive Margin (Whether the Reader Subscribed At All)

Notes: The figures plot the estimated average intent-to-treat effect of serving a four weeks vs. two weeks promotional contract on readers' subscription behavior. Specifically, we plot the estimated coefficient β_3 from equation (2) for every month. Month 0, shaded in gray, is the promo period (two weeks). The left-most points at month -1 are before readers hit the paywall. Month 1 is the first month after the promo ends, and so forth. The point "After-promo" aggregates across all after the promo time periods; "Total" aggregates from 0 until the end. Percentages next to the effect size compare the effect to the mean level of the omitted group.The error bars show 95% confidence intervals. Standard errors are clustered at the individual reader level.

usage data. The sample is different than above: we include "takers" and the sample is fixed throughout, including if someone canceled their subscription.

Figure 4 shows that auto-renewal promo takers are orders of magnitude more likely to be subscribed after the two weeks promo time. However, we do not see any difference in the groups' website visits in the last two weeks after the promo period ends, let alone orders of magnitude differences. This indicates that the auto-renewal takers who continue to subscribe do not visit the website more often. Table A.3 supports this finding. It compares usage across four groups: auto-renewal and auto-cancel promo takers, split by whether they remained subscribed after the promo. More than half of the auto-renewal subscribers do not visit the platform at all post-promo, matching the behavior of auto-cancel non-renewers. Moreover, those who remain subscribed under auto-renew visit 62% fewer pages than post-promo auto-cancel subscribers.

Overall, the results of this analysis are consistent with our inference that the valuation of subscription does not grow for auto-renewal subscribers due to their subscription status. Note that this analysis is based on two weeks post-promo usage data, and does not capture any future increase in usage auto-renewal subscribers might have.

7.2 Evidence for "Spite"

We examine heterogeneity among subscribers and ask: Which reader types, characterized by their consumption value, subscribe to auto-renew contracts versus auto-cancel contracts? If the mechanism is purely inertia and sophistication, we expect three patterns (see Section 2): (i) during the promo, auto-renewal takers will have *higher* value on average, since only higher types are willing to accept the inertia-inducing contract; (ii) after the promo, inertia will retain more low-value types in auto-renewal than the active renewal high types in auto-cancel; (iii) in the long run, the distribution of types should converge (from below) as inerts exit.[20]

We find consistent evidence for the first two predictions but reject the third. Instead, auto-renewal subscribers have *higher* predicted value in the long run, consistent with a "spite" effect. Under a model with spite, some potentially high-WTP users actively avoid subscribing when offered auto-renewal, making the surviving subscribers higher value.

We use predicted post-promo usage (from pre-exposure behavior) as a proxy for reader value; see Section A.5. Figure 5 plots the average predicted usage difference (auto-renewal minus auto-cancel) among each period's subscribers. The sample is hence different at each period based on who is currently subscribed. During the promo period, average subscriber types are similar or slightly higher in auto-renewal. Immediately after the promo, subscriber types are lower in auto-renewal, consistent with inertia. But in the long run, types in auto-renewal are significantly higher. Figure A.7 shows this pattern: during the promo, type

[20]If some readers are forever inert, we expect the value of auto-renew subscribers to remain lower.

Figure 4: Subscription vs. Platform Usage for Two Weeks €0.99 Auto-Renewal Promo Takers vs. Two Weeks €0.99 Auto-Cancel Promo Takers

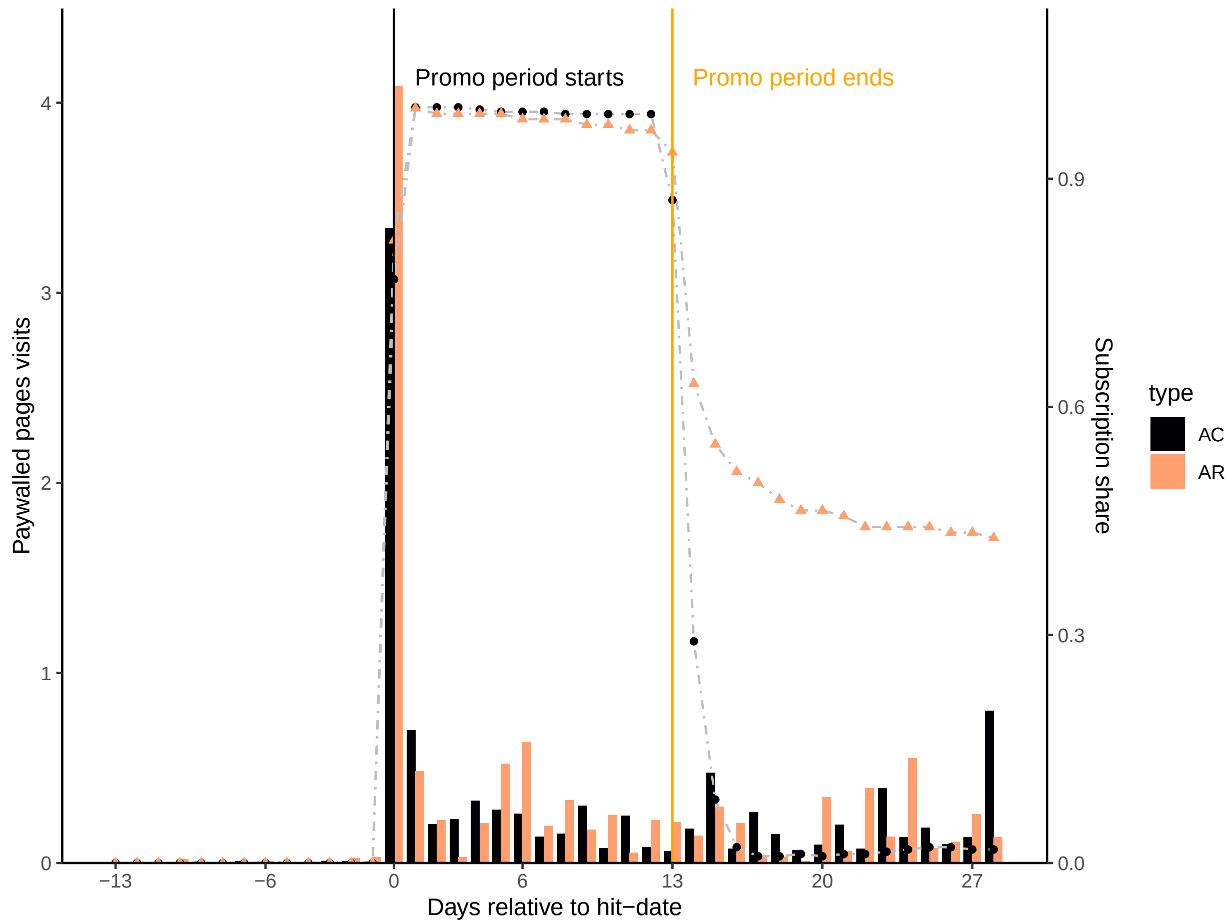

Notes: The figure plots the daily average subscription rate (dots and triangles) and average newspaper consumption—measured by number of website page visits (bars)—separately for those who took the two weeks €0.99 auto-renewal promo (AR, orange) and those who took the two weeks €0.99 auto-cancel promo (AC, black). The time on the x-axis starts two weeks before the experimental offer was given to the reader and covers the promotional two weeks and two weeks after that.

Figure 5: Difference in Reader Types Between Subscribers in the Auto-Renewal vs Auto-cancel Groups

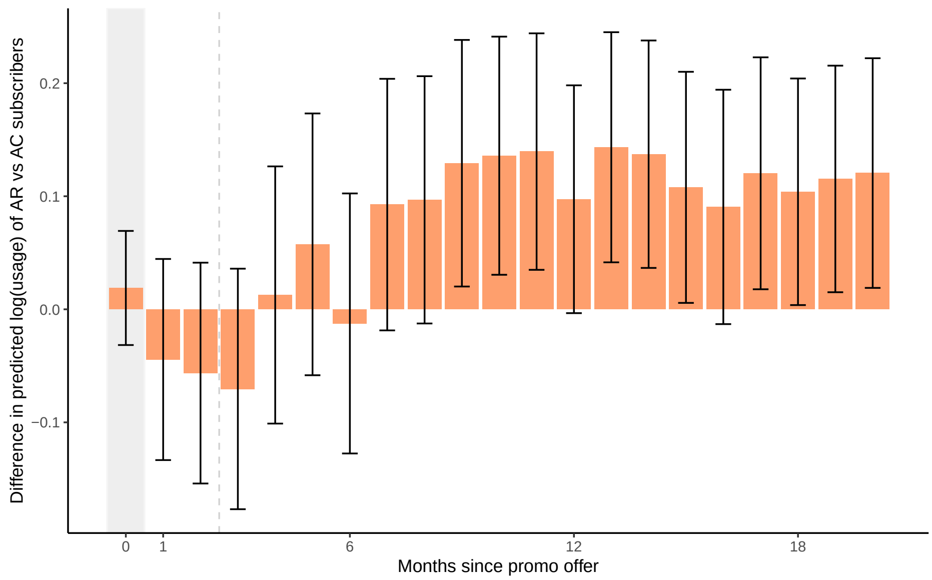

Notes: We use group G readers' pre and post-experimental usage data to predict post out-of-sample newspaper usage for the main sample. We use predicted usage as a proxy for reader type. The figure shows the difference in average reader type between subscribers in the auto-renewal groups and auto-cancellation groups by period. Error bars represent 95% confidence intervals. Standard errors are clustered at the individual level.

distributions are similar; post-promo, auto-renewal retains lower types (inertia); but in the long run, auto-renewal subscribers have higher predicted value (spite).

Overall, our findings are consistent with an implicit penalty induced by the auto-renew offer relative to the auto-cancel offer. Some readers who could have been long-run subscribers in the auto-cancel group decide not to subscribe when assigned to the auto-renewal group. Note that the contracts offered to both groups are equivalent in the long-run; unlike the promotional period where the auto-cancellation contract has a different continuation value, after the promo period all contracts are identical. The lower long-run subscription of high types is therefore consistent with a psychological cost, or spite against the newspaper, due to the initial auto-renewal offer.

7.3 Learning about Subscription Value from the Promo Trial

A trial could cause some consumers to learn that their subscription valuation is above the full price. This mechanism is unlikely to be a dominant driver of inertia in our context, given lower take-up and higher post-promo subscription in the auto-renewal group shown in section 6. The effects of price and duration on subscription shed further light on this channel. As shown in Figures 2–3, price reduction has no detectable long-run effect, while longer trials have a modest persistent effect.[21]

Together, our data patterns suggest that while some readers may learn about their valuation through usage, such effects are small. Most post-promo subscription behavior under auto-renewal appears driven by inaction and switching costs rather than learning.[22]

8 Quantification of Inertia

In this section, we quantify the extent of actual (experienced) inertia, anticipated inertia, and their heterogeneity using the structural model introduced in Section 2, imposing minimal additional assumptions.

To estimate this model coherently with our data, in this section, we focus on outcomes of those individuals who took up subscriptions during the promo period in both the auto-renewal and auto-cancel groups (column 3 of Table 2). We track and model their behavior throughout the post-promo time period to assess their long-term behavioral responses. This differs from Section 6 where we also counted subscriptions of those who subscribed for the first time after the promo period ended to estimate the intent-to-treat effects.[23] We discuss the robustness of our results to this estimation strategy in section 8.3 below.

Recall that our model includes two terms that may generate inertia: the probability of inaction and the cost of unsubscription. For an *existing*-subscriber, both forces prevent one from canceling on time; for the *potential*-subscriber, their expected magnitude will determine the decision to subscribe.[24]

8.1 Identification

8.1.1 Excess subscriptions due to inertia

Comparing the post-promotion subscription rates among AR versus AC provides a lower bound in inertia-induced subscription under the assumption that AR promo takers would also take up AC. Let R and C be the set of people who would take up AR and AC promos, respectively. The assumption implies that $R \subseteq C$. We want to estimate the per-period excess subscription due to AR among the population in R, $y_R^{AR} - y_R^{AC}$, where y_R^{AR} is the share among consumers in R who take the promo when offered AR who remain subscribed at the later period; y_R^{AC} is the share among consumers in R who would take the AC promo and

[21]Separately, we confirmed that the carry-over effect of longer duration occurs within auto-cancel groups by contrasting treatment groups “F” and “H”.

[22]This result may be particularly context-dependent. Because readers can engage with free and metered content before subscribing, the scope for learning is more limited here than for novel products that are difficult to evaluate without access.

[23]Appendix Figure A.5 shows the decomposition of subscribers based on whether they signed up first during the promo period, or after it.

[24]In our main specification, unsubscription cost is deterministic, common to all types, and known in advance. As such, it may deter some from subscribing to an auto-renew contract but it is not generating time-varying behavior. Time-varying behavior is driven only through inaction and sophistication. In Appendix A.6 we allow for stochastic unsubscription costs and naivete instead of inaction as alterative assumptions.

remain subscribed at the later period. While y_R^{AR} is observed, y_R^{AC} is not: we do not know what share among R would have renewed their subscription had they been offered auto-cancel instead. However, we can decompose the number of auto-cancel renewers to be those who renew from R and those who renew from $C \setminus R$:

$$\begin{aligned}
|C| y_C^{AC} &= |R| \times y_R^{AC} + |C \setminus R| \times y_{C \setminus R}^{AC} \\
|C| y_C^{AC} &\geq |R| \times y_R^{AC} \\
\implies y_R^{AR} - y_R^{AC} &\geq y_R^{AR} - y_C^{AC} \times \frac{|C|}{|R|}.
\end{aligned}$$

This estimable lower bound on the share of excess subscribers after the promo period among auto-renew promo takers. Any share greater than 0 implies that there is inertia. The rate of decline of these excess shares further tells us the average rate of inertia among these excess subscribers.

Appendix Figure A.6 shows that this excess subscription in the month after the subscription starts at 23.4% (s.e. 1.4%) and gradually declines over time. Consistent with the evidence above for spite, the excess becomes negative after 17 months (these are lower bounds on excess mass, and are thus more likely to find negative values). Fitting a single rate of decline, we get an inertia rate of 86.3% (s.e. 3.7%). Meaning, there is a monthly 86% chance that an unwilling subscriber fails to cancel.

This excess inertia-induced subscription results from an evolving mix of heterogeneous consumers, with potentially differing levels of sophistication. The next section delves into teasing those apart under some additional assumptions.

8.1.2 Heterogeneity in inertia and perceived inertia

To disentangle actual and perceived inertia and incorporate price and duration variation, we estimate the structural model introduced in Section 2. The model includes per-period valuations, subscription and cancellation costs, discounting, inaction, and beliefs about inaction. We make the following assumptions to make the parameters identifiable:[25]

- Valuations are drawn from a distribution F
- Subscription and cancellation costs are homogeneous and known.
- Consumers may experience a utility penalty ("spite") when offered auto-renewal.
- Inertia is discrete: individuals are either inert ($\phi > 0$) or not ($\phi = 0$).

[25] Appendix A.6 considers a model with stochastic cancellation costs and no inaction.

- Inerts are either sophisticated ($\hat{\phi} = \phi$) or naive ($\hat{\phi} = 0$).

Figure 6 illustrates the behavior of each type when offered an auto-renewal or auto-cancel contract. The Y-axis represents the per-period subscription value, with labeled threshold values of behavior change. The six bars represent the inertia types assigned to AR and AC groups.

The colored and patterned areas represent individuals who take up the promo and wish to cancel later, and the (gray) shaded areas represent those who would subscribe in the long term.

Starting with the non-inerts, those in colored regions take the promo and cancel; they do not pay the full price. The very high types stay subscribed in the long run.

The naive inerts behave like the non-inert in take-up and long-run: In take-up, they wrongfully think they will be non-inert; in the long-run in AR the excess subscribers will eventually leave. The difference between naifs and the non-inerts is being inert after the promo due to inaction.

Sophisticate inerts behave like the other types when offered an AC contract. However, when offered an AR contract, some avoid it due to the risk of being inactive and paying above their willingness to pay. Those who do take the AR contract might end up paying more than their value, but they are aware of that risk and still choose to take the AR contract. Spite is modeled as a downward shift in utility under AR. As a result, the marginal taker under AR has a higher valuation than under AC.

8.1.3 Short, interim, and long-term subscription comparisons

Though not depicted in Figure 6, experimentally varied promo price and trial duration shift the threshold values of auto-cancel takers identifying moments of the subscription value distribution, F. For the model estimation exercise in the next section, we assume this distribution follows a Gamma distribution.[26]

The difference in long-term subscription rates between AC and AR gives us the impact of spite. Given spite and the value distribution, the difference in promo takeup between AR and AC estimates the quantity marked as $A \times \pi_{\text{sophisticate}} \times \pi_{\text{inert}}$, where $\pi_{\text{sophisticate}}$ and π_{inert} represent the population proportion of inert and sophisticated types, and $A = F(v^*) - F(\underline{v} + \alpha)$. This share is coming from the sophisticates' anticipated future cost of having an AR subscription, which is a function of the inaction probability ϕ, and time discount-rate δ, along with the known prices.

These unknowns are estimated from the interim post promo attrition of inertia induced AR subscribers described in the previous section. The role of unsubscription cost is pinned down by the change in this attrition with the increase in price that occurs after month two.

[26] This is a two parameters distribution. Einav et al. (2023) use an Exponential distribution for values of subscribers. The Exponential distribution is a specific case of the Gamma distribution when the shape parameter equals 1.

Figure 6: Heterogeneity in Comparisons across AR and AC

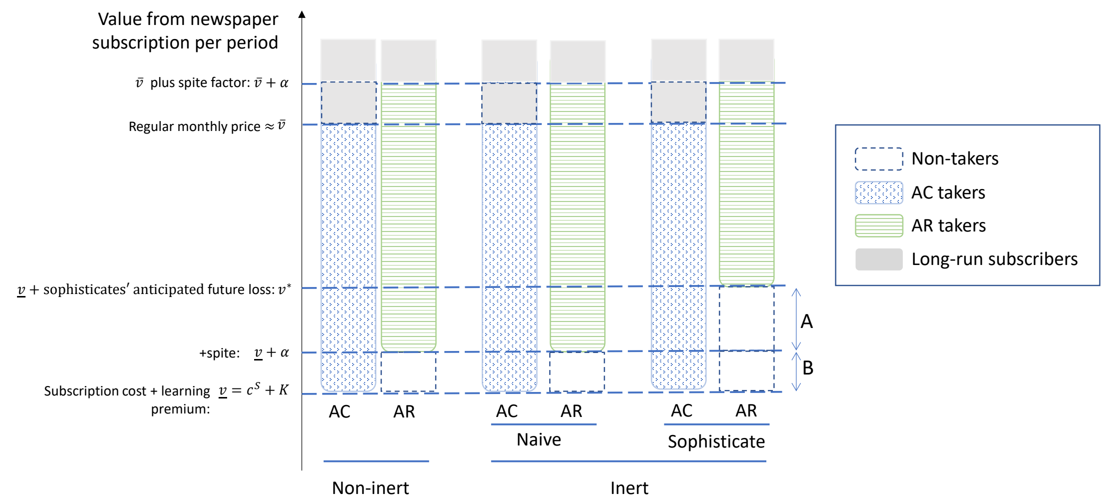

Notes: Illustration of behavior by inertia type and valuation, across AR and AC treatment groups. Colored regions denote promo takers; gray regions denote long-term subscribers. Labels on the Y-axis highlight cutoff valuations at which subscription behavior changes vis-a-vis promo take-up or the long-term. For example, non-inerts with valuations higher than the subscription cost net of the learning premium (potential benefit from learning about the true value) will take up an auto-cancel promo contract (blue pattern); those with value greater than the regular monthly price will subscribe in the long term (gray).

8.2 Estimating the Dynamic Choice Model

We solve the model described in Section 2 and generate simulated moments of subscription shares per period by treatment arm. Thus, accounting for contract-specific terms (length of promo period, price, and AR vs. AC). We solve the model backwards from the steady state price and get each type's planned subscription path. A consumer type is defined by its per-period value, actual inertia, and perceived inertia.

The estimation minimizes the sum of squared deviations between simulated and empirical subscription shares per period, across 21 periods and 7 treatments (147 moments), using equal weights and a BFGS optimization algorithm. We estimate the following parameters, summarized in Table 3: Share of inerts in the population (46%) and their average monthly inaction (85%), the share of naifs (8% of inerts), a "spite" utility AR penalty (€0.03 per month with 95% CI upper bound at €0.10), and the probability of learning of being a high type (0.1%). We also estimate the subscription cost and the cancellation cost, and assume a monthly discount factor of 0.995. Figure A.12 in the appendix section A.4 compares the empirical and simulated moments in our data and shows that the model is able to fit most data patterns.

Table 3: Structural Estimation Results

Name	Parameter	Estimate	95% Confidence Interval
Monthly inaction	ϕ	85.20%	[73.6% - 90.3%]
Share of inerts		46.20%	[37.1% - 52.2%]
Share of naifs in population		3.85%	[2.54% - 6.06%]
Share of naifs among inerts		8.33%	[5.41% - 14.8%]
Prob of high value	λ	0.08%	[0.0837% - 0.122%]
Monthly AR penalty (€-cent)	α	2.79	[0.11 - 10.06]
Unsubscription cost (€-cent)	c_u	21.63	[2.09 - 75.52]
Subscription cost (€-cent)	c_s	91.16	[67.63 - 125.4]
Gamma dist (shape)		0.00105	[0.000934 - 0.00116]
Gamma dist (scale)		36.8	[34.2 - 40.7]
Monthly discount factor	δ	0.995	—-

8.3 Discussion

Our structural estimates imply that inertia in subscription cancellation is large, but largely anticipated. Among promo-period subscribers, we estimate a high monthly probability of inaction (about 0.85) and a substantial share of inert consumers (roughly one half). At the same time, most inert consumers are sophisticated about their own inertia: naivete is present but rare in the population (only a few percent), even though it can matter disproportionately for short-run retention under auto-renewal. The estimated "spite" penalty from being offered auto-renewal is small in levels (on the order of a few euro-cents per month), and learning from trial experience is quantitatively limited. Together, these estimates rationalize the reduced-form patterns: auto-renewal deters take-up because many consumers anticipate the risk of unwanted continuation, while those who do take auto-renewal exhibit substantial excess retention driven by inaction rather than higher consumption value.

In an extension, we relax the baseline restriction that naive and sophisticated inerts share the same inaction rate. Intuitively, this is identified from the *time path* of post-promo attrition: the observed month-to-month cancellation rate is a mixture over the types who remain subscribed, and that mixture changes over time as different types sort into (and out of) paid subscription. This shifting composition provides leverage to separately pin down inaction for naifs and sophisticates. We estimate inaction for sophisticates that is similar to the baseline (95% CI 0.73–0.94), while naifs exhibit a noisier and, if anything, lower inaction rate (95% CI 0.14–0.84). Appendix Figure A.9 and Table A.5 report more details. One suggestive interpretation is that sophistication is partly learned: consumers who are more prone to inaction may be more likely to recognize it after exposure and behave as sophisticates going forward.

Finally, the model highlights why subscriber-based samples can overstate naivete. Selection implies that the pool of medium-term paying subscribers is not representative: low-value non-inerts cancel promptly, and low-value sophisticated inerts avoid auto-renewal in the first place, leaving a subscriber pool that disproportionately contains inert consumers and can temporarily overrepresent naive inerts.

These conclusions are robust to alternative definitions of promo-period takers and to allowing additional heterogeneity in inaction; the full robustness exercises and composition decompositions are reported in the appendix A.4.

8.4 Targetability Based on Sophistication

We ask whether firms can feasibly and profitably target auto-renewal (vs. auto-cancel) contracts based on consumers' sophistication. Theoretical work shows that firms may wish to discriminate based on consumer naivete. We examine whether such targeting is both feasible and desirable. In practice, naivete could be irrelevant for auto-renewal – if most subgroups that can be targeted with auto-renewal yield higher long-term subscriptions when offered such contracts – or it could be the sole determinant, if the firm benefits only through short-term post-promotion revenues from offering auto-renew contracts to naive consumers.

Using only pre-exposure clickstream behavior (a firm's first-party data), we first show that subscription value is meaningfully predictable: an out-of-sample predicted-usage score strongly correlates with take-up, retention, and revenue. We construct this score and estimate heterogeneous treatment effects of auto-renewal using generalized random-forest methods (Athey et al., 2019).

We then use these predicted heterogeneous effects to simulate simple targeting rules under two objectives—maximizing post-promo subscribers versus total revenue—and evaluate them using the experiment's random assignment (an "off-policy" validation approach; Hitsch et al., 2024). Depending on the objective, the implied optimal share assigned to auto-renewal ranges from a small minority (subscriber-maximizing) to a large majority (revenue-maximizing). Across objectives, optimal targeting is driven primarily by predicted value, not by sophistication. We classify readers as "naive" based on a distinctive response pattern—little take-up sensitivity to auto-renewal but higher short-run retention under auto-renewal—and find that naivete is at most a modest, incremental predictor of being assigned auto-renewal conditional on predicted value. Thus, while contract targeting based on observables is feasible, the scope for *pure* naivete-based discrimination appears limited in our setting. Appendix A.5 reports the full implementation details and results.

9 Conclusion

The common wisdom in the academic literature, as well as in the industry, is that consumers are highly inert. Once a firm gets a consumer to buy, the argument goes, it can increase prices or change terms and the consumer is insensitive to those. A large body of evidence, including this paper, supports the view that existing, retained, customers are highly inert. However, our findings highlight the limitations of this body of knowledge as it relies on a selected sample of already *existing* customers. Our paper suggests that the vast majority of inert consumers in our setting are aware of their future inertia and avoid engaging with an exploitative contract. Furthermore, offering an exploitative contract pushes 23% of consumers from engaging with the company for the duration of our data. These new findings imply that consumers' awareness of their future inertia limits inertia exploitation. They also imply that counterfactuals based on the inferred inertia of existing consumers will not generalize well to the population.

These results have important implications. If a firm focuses on short-term profits, offering an auto-renewal contract may appear advantageous since it resulted in higher revenue. However, auto-renewal reduces the number of unique subscribers, which can affect key business metrics such as reach, advertising revenue, and engagement-driven growth. Usage data further indicate that many subscribers retained due to auto-renewal use the product insignificantly, suggesting the contract extracts payment without delivering significant value. Overall, the benefits of auto-renewal erode over time. If the firm must choose a single contract type, our evidence suggests that auto-cancellation may better serve both the firm and its consumers in the long run.

In our setting, we found limited incentive for naivete-based targeting. Moreover, the mere offer of an exploitative contract as part of a menu, even if better alternatives are included, deters some readers from participation. Conversely, offering a very good promotional offer creates some good-will. This suggests that contract menus themselves can shape consumer trust and long-term engagement.

References

Allcott, Hunt, Matthew Gentzkow, and Lena Song, "Digital Addiction," June 2021.

Ascarza, Eva, Raghuram Iyengar, and Martin Schleicher, "The Perils of Proactive Churn Prevention Using Plan Recommendations: Evidence from a Field Experiment," *J. Mark. Res.*, February 2016, *53* (1), 46–60.

Athey, Susan, Julie Tibshirani, and Stefan Wager, "Generalized random forests," *aos*, April 2019, *47* (2), 1148–1178.

Augenblick, N and M Rabin, "An experiment on time preference and misprediction in unpleasant tasks," *Rev. Econ. Stud.*, 2019.

Bai, Liang, Benjamin Handel, Edward Miguel, and Gautam Rao, "Self-control and demand for preventive health: Evidence from hypertension in India," *Rev. Econ. Stat.*, December 2021, *103* (5), 835–856.

Benartzi, Shlomo, John Beshears, Katherine L Milkman, Cass R Sunstein, Richard H Thaler, Maya Shankar, Will Tucker-Ray, William J Congdon, and Steven Galing, "Should Governments Invest More in Nudging?," *Psychol. Sci.*, August 2017, *28* (8), 1041–1055.

Brot-Goldberg, Zarek C, Timothy Layton, Boris Vabson, and Adelina Yanyue Wang, "The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D," Technical Report w28331, National Bureau of Economic Research January 2021.

Camerer, Colin, Peter Landry, and Ryan Webb, "The Neuroeconomics of Habit," August 2018.

Choi, James J, David Laibson, Brigitte C Madrian, and Andrew Metrick, "Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance," *Tax Policy and the Economy*, January 2002, *16*, 67–113.

Datta, Hannes, Bram Foubert, and Harald J Van Heerde, "The challenge of retaining customers acquired with free trials," *Journal of Marketing Research*, 2015, *52* (2), 217–234.

DellaVigna, S and U Malmendier, "Contract design and self-control: Theory and evidence," *Q. J. Econ.*, May 2004, *119* (2), 353–402.

Drake, Coleman, Conor Ryan, and Bryan Dowd, "Sources of inertia in the individual health insurance market," *J. Public Econ.*, April 2022, *208*, 104622.

Dubé, Jean-Pierre, Günter J Hitsch, and Peter E Rossi, "State dependence and alternative explanations for consumer inertia," *Rand J. Econ.*, August 2010, *41* (3), 417–445.

Einav, Liran, Benjamin Klopack, and Neale Mahoney, "Selling Subscriptions," Technical Report w31547, National Bureau of Economic Research, Cambridge, MA August 2023.

Eliaz, Kfir and Ran Spiegler, "Contracting with Diversely Naive Agents," *Rev. Econ. Stud.*, July 2006, *73* (3), 689–714.

Ericson, Keith M, "Forgetting we forget: Overconfidence and memory," *J. Eur. Econ. Assoc.*, February 2011, *9* (1), 43–60.

Goettler, Ronald L and Karen Clay, "Tariff Choice with Consumer Learning and Switching Costs," *J. Mark. Res.*, August 2011, *48* (4), 633–652.

Grubb, Michael D and Matthew Osborne, "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," *Am. Econ. Rev.*, January 2015, *105* (1), 234–271.

Handel, Benjamin R, "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *Am. Econ. Rev.*, December 2013, *103* (7), 2643–2682.

Heckman, James J, "Statistical models for discrete panel data," *Structural analysis of discrete data with econometric applications*, 1981, *114*, 178.

Heidhues, Paul and Botond Kőszegi, "Exploiting Naïvete about Self-Control in the Credit Market," *Am. Econ. Rev.*, December 2010, *100* (5), 2279–2303.

Heiss, Florian, Carmine Ornaghi, and Mirco Tonin, "Inattention Matters: An Analysis of Consumers' Inaction in Choosing a Water Tariff," *J. Eur. Econ. Assoc.*, 2022.

Hitsch, Günter J, Sanjog Misra, and Walter W Zhang, "Heterogeneous treatment effects and optimal targeting policy evaluation," *Quantitative Marketing and Economics*, 2024, *22* (2), 115–168.

Ho, Kate, Joseph Hogan, and Fiona Scott Morton, "The impact of consumer inattention on insurer pricing in the Medicare Part D program," *Rand J. Econ.*, December 2017, *48* (4), 877–905.

Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L Puller, "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market," *American Economic Journal: Economic Policy*, November 2017, *9* (4), 192–226.

IV, Frank J Chaloupka, Matthew R Levy, and Justin S White, "Estimating Biases in Smoking Cessation: Evidence from a Field Experiment," Technical Report w26522, National Bureau of Economic Research December 2019.

Klemperer, Paul, "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Rev. Econ. Stud.*, October 1995, *62* (4), 515–539.

Kong, Xinyao, Jean-Pierre Dubé, and Øystein Daljord, "Non-parametric Estimation of Habitual Brand Loyalty," March 2022.

Lin, Tesary and Sanjog Misra, "Frontiers: The Identity Fragmentation Bias," *Marketing Science*, May 2022, *41* (3), 433–440.

Madrian, B C and D F Shea, "The power of suggestion: Inertia in 401(k) participation and savings behavior," *Q. J. Econ.*, November 2001, *116* (4), 1149–1187.

Negative Option Rule: Final Rule

***Negative Option Rule: Final Rule**, https://www.ftc.gov/legal-library/browse/federal-register-notices/negative-option-rule-final-rule October 2024. Accessed: 2024-11-18.*

O'Donoghue, Ted and Matthew Rabin*, "Doing It Now or Later," Am. Econ. Rev., March 1999, 89 (1), 103–124.*

_ and _ *, "Choice and procrastination," Q. J. Econ., 2001, 116 (1), 121–160.*

Reme, Bjørn-Atle, Helene Lie Røhr, and Morten Sæthre*, "Inattention in Contract Markets: Evidence from a Consolidation of Options in Telecom," Manage. Sci., May 2021.*

Rodemeier, Matthias*, "Buy baits and consumer sophistication: Field evidence from instant rebates," Am. Econ. J. Econ. Policy, May 2025, 17 (2), 30–59.*

Shui, Haiyan and Lawrence M Ausubel*, "Time Inconsistency in the Credit Card Market," May 2004.*

Thaler, Richard H and Cass R Sunstein*, Nudge: The Final Edition, Penguin, August 2021.*

Vigna, Stefano Della and Ulrike Malmendier*, "Paying not to go to the gym," Am. Econ. Rev., May 2006, 96 (3), 694–719.*

Yang, Jeremy, Dean Eckles, Paramveer Dhillon, and Sinan Aral, *"Targeting for long-term outcomes," arXiv preprint arXiv:2010.15835, 2020.*

Yoganarasimhan, Hema, Ebrahim Barzegary, and Abhishek Pani, *"Design and evaluation of optimal free trials," Manage. Sci., August 2022.*

A ONLINE APPENDIX

A.1 Additional Tables and Figures

Table A.1: Balance of pre-experiment behavior

Dependent Variables:	Total_Pages	Open	Paywalled
Model:	(1)	(2)	(3)
Variables			
Test Group A	3.711***	3.656***	0.0541***
	(0.0313)	(0.0308)	(0.0033)
Group B vs A	0.0771*	0.0750*	0.0021
	(0.0459)	(0.0452)	(0.0048)
Group C vs A	0.1028*	0.1063*	-0.0035
	(0.0564)	(0.0559)	(0.0041)
Group D vs A	-0.0280	-0.0291	0.0011
	(0.0414)	(0.0407)	(0.0047)
Group E vs A	0.0222	0.0244	-0.0022
	(0.0542)	(0.0537)	(0.0047)
Group F vs A	0.0410	0.0396	0.0014
	(0.0456)	(0.0450)	(0.0044)
Group H vs A	-0.0068	-1.27×10^{-5}	-0.0068*
	(0.0457)	(0.0452)	(0.0039)
Fit statistics			
Observations	1,385,002	1,385,002	1,385,002
R^2	7.51×10^{-6}	7.59×10^{-6}	5.06×10^{-6}
Adjusted R^2	3.18×10^{-6}	3.26×10^{-6}	7.32×10^{-7}

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A.2: By-group summary statistics

Variable	A	B	C	D	E	F	H	TOTAL
Auto renew	1	1	1	1	0	0	0	0.566
Long duration	1	1	0	0	1	1	0	0.565
High price	0	1	0	1	0	1	1	0.609
N	182,079	211,360	181,189	211,563	180,824	211,535	211,319	1,389,869
Subscribers during_promo	469	472	409	415	671	704	704	3,844
Subscribers during_promo_pct	0.258	0.223	0.226	0.196	0.371	0.333	0.333	0.277
Subscribers post_promo	482	561	500	545	531	638	617	3,874
Subscribers post_promo_pct	0.265	0.265	0.276	0.258	0.294	0.302	0.292	0.279
Average revenue	0.300	0.300	0.280	0.300	0.300	0.300	0.240	0.290
Total revenue	54,720	63,857	50,517	62,877	53,717	64,367	49,983	403,477
Total subscription_rate_during_promo_pct	0.184	0.154	0.162	0.147	0.299	0.272	0.262	0.211
Subscription rate_post_promo_pct	0.074	0.077	0.069	0.069	0.079	0.082	0.064	0.073
Takers during_promo	254	186	189	153	380	437	396	1,995
Takers during_promo_pct	0.139	0.088	0.104	0.072	0.210	0.207	0.187	0.144
Takers post_promo	142	153	122	114	173	209	174	1,087
Takers post_promo_pct	0.078	0.072	0.067	0.054	0.096	0.099	0.082	0.078
Takers overall	299	232	227	182	479	540	475	2,434
Takers overall_pct	0.164	0.110	0.125	0.086	0.265	0.255	0.225	0.175
Takers subscription_rate_during_promo_pct	0.121	0.078	0.090	0.068	0.195	0.192	0.173	0.131
Mean pages_pre	3.711	3.788	3.813	3.683	3.733	3.751	3.704	3.749
Mean pages_post			4.400	4.378			4.447	

Table A.3: Usage during vs. after the promo for two weeks promo takers

	Auto-renewal two weeks €.99 subscribers		Auto-Cancel two weeks €.99 subscribers	
	Subscribed in two weeks after the promo estimate	Not subscribed in two weeks after the promo estimate	Subscribed in two weeks after the promo estimate	Not subscribed in two weeks after the promo estimate
Promo two weeks: Avg. page visits	18.78	19.49	37.76	22.74
Promo two weeks: % readers with any visit	63.5%	74.2%	82.7%	68.8%
Post promo two weeks: Avg. page visits	23.58	10.25	62.63	12.16
Post promo two weeks: % readers with any visit	49.4%	37.1%	90.7%	49.3%
N	85	97	75	400

Notes: We focus on the readers who took the two weeks, €.99 experimental contract and separate them by (1) whether they took the auto-renewal or auto-cancel contract and (2) whether they were subscribed in the two weeks post promo. *The first two rows present the average number of page visits, and the proportion of readers who had any visit to the newspaper in the first two weeks (excluding exposure day). The next two rows do the same for the subsequent two weeks.*

Table A.4: Parameter estimates by sample

Name	(1) promo period taker	(2) promo period exp taker	(3) exp taker any period	(4) promo period new taker
Gamma dist (scale)	36.8	23.6	28.8	32.2
	[34.2 - 40.7]	[21.1 - 24.7]	[24 - 31.7]	[30.2 - 35.9]
Gamma dist (shape)	0.00105	0.000887	0.000832	0.00106
	[0.000934 - 0.00116]	[0.000801 - 0.00103]	[0.000702 - 0.00108]	[0.000916 - 0.00114]
Monthly AR penalty (€ cent)	2.79	0.51	1.62	2.48
	[0.11 - 10.03]	[0.1 - 6.22]	[0.12 - 8.18]	[0.1 - 13.28]
Monthly inaction	85.2%	81.9%	80.8%	84.5%
	[73.6% - 90.3%]	[73.8% - 86.8%]	[67.2% - 87.6%]	[72.5% - 90.4%]
Prob of high value	0.0846%	0.122%	0.0994%	0.094%
	[0.0837% - 0.122%]	[0.084% - 0.123%]	[0.084% - 0.123%]	[0.0842% - 0.124%]
Share of inerts	46.2%	59.7%	64.3%	51.5%
	[37.1% - 52.2%]	[55.3% - 65.5%]	[51.6% - 75.9%]	[39% - 54.1%]
Share of naives among inerts	8.33%	12.5%	17.3%	7.97%
	[5.41% - 14.8%]	[10.8% - 21%]	[13.8% - 28.8%]	[4.96% - 14.2%]
Share of naives in population	3.85%	7.49%	11.1%	4.1%
	[2.55% - 6.05%]	[6.6% - 12.4%]	[8.89% - 17.4%]	[2.45% - 6.72%]
Subscription cost (€ cent)	91.16	109.57	111.33	98.93
	[67.66 - 125.39]	[85.79 - 143.46]	[75.13 - 175.47]	[68 - 120.9]
Unsubscription cost (€ cent)	21.63	15.46	23.17	4.58
	[2.09 - 75.47]	[1.67 - 40.15]	[2.15 - 90.47]	[1.6 - 77.48]

Table A.5: Parameter Estimates by Model

Parameter	Inaction (3 types)	Inaction (2 phis)	Costs (3 types)	Costs (sophisticates)	Costs (soph+partial)
Monthly inaction	85.20%		0	0	0
	[73.6% - 90.3%]		[—]	[—]	[—]
Monthly inaction (Naifs)		48.30%			
		[14% - 84.1%]			
Monthly inaction (Sophisticates)		90.30%			
		[73.4% - 91.4%]			
Share of inerts	46.20%	52.10%	37%	100%	100%
	[37.1% - 52.2%]	[41.4% - 60.2%]	[14.4% - 86%]	[—]	[—]
Naifs among inerts	8.33%	10.40%			
	[5.41% - 14.8%]	[5.93% - 23.3%]			
Naifs in pop.	3.85%	5.41%			
	[2.55% - 6.05%]	[3.05% - 12.8%]			
Cost naifs			4.95%	0%	77.40%
			[3.45% - 51.2%]	[—]	[12.8% - 79.6%]
Underestimation			0%		63.90%
			[—]		[18.4% - 81.6%]
Unsubscription cost	21.63	14.07	84.57	206.63	101.49
(€ cent)	[2.09 - 75.47]	[2.18 - 68.45]	[64.87 - 287.93]	[47.09 - 326.2]	[52.73 - 299.91]
Range of costs	0	0	1717.91	40.45	267.51
(€ cent)	[—]	[—]	[63.21 - 5882.86]	[44.21 - 518.82]	[46.78 - 612.96]
Subscription cost	91.16	101.25	54.06	110.2	67.26
(€ cent)	[67.66 - 125.39]	[67.89 - 122.27]	[40.18 - 137.65]	[36.9 - 168.69]	[39.83 - 138.11]
Spite	2.79	2.47	9.45	0.29	0.99
(€ cent/month)	[0.11 - 10.03]	[0.1 - 14.07]	[0.14 - 9.67]	[0.1 - 5.52]	[0.11 - 8.52]
Prob of high value	0.08%	0.09%	0.09%	0.12%	0.09%
	[0.0837% - 0.122%]	[0.084% - 0.123%]	[0.0838% - 0.124%]	[0.0836% - 0.123%]	[0.0845% - 0.122%]
Gamma dist	36.8	33.2	44.6	38.2	38.2
(scale)	[34.2 - 40.7]	[31 - 36.6]	[34.1 - 48.4]	[34 - 49.8]	[31.5 - 41.8]
Gamma dist	0.00105	0.00106	0.00085	0.0011	0.00086
(shape)	[0.00093 - 0.0012]	[0.00091 - 0.0012]	[0.00078 - 0.0012]	[0.00077 - 0.0013]	[0.00077 - 0.0011]
Time discount factor	0.995	0.995	0.995	0.995	0.995
	[—]	[—]	[—]	[—]	[—]

Table A.6: Validation of heterogeneous treatment effect predictions

Outcome and assignment	All	Assigned to AR	Assigned to AC
Subscription in month 0	-0.00122 (0.00010)	0.04581 (0.00233)	-0.00180 (0.00009)
Subscription in month 1	0.00012 (0.00004)	0.00035 (0.00004)	-0.00413 (0.00036)
Revenue post	0.02136 (0.01886)	0.23740 (0.01874)	-0.78566 (0.05578)
Any subscription post	-0.00030 (0.00009)	0.00652 (0.00034)	-0.00118 (0.00010)

Note: The table provides a validation of policy assignment based on causal regression forests predictions of auto-renewal treatment effects. Each row signifies a different outcome to optimize, each column is a different counterfactual subsample of the data, and each number (parentheses) is the actual estimated treatment effect of auto-renewal on that outcome for that subsample (standard error). Each cell is thus an estimate of the effect of AR from a separate regression. For example, the effect of auto-renewal on subscriptions during month 0 (the promo period) is -0.00122. If we estimate the effect on the subsample of readers that are predicted to have a positive treatment effect, and thus counterfactually assigned to AR, their CATE is indeed positive at 0.0458 (0.0023). That all effects are positive in the column "Assigned to AR" and negative in "Assigned to AC" is a validation of the HTE estimates.

Figure A.1: Number of Readers in Each Treatment Arm by Week

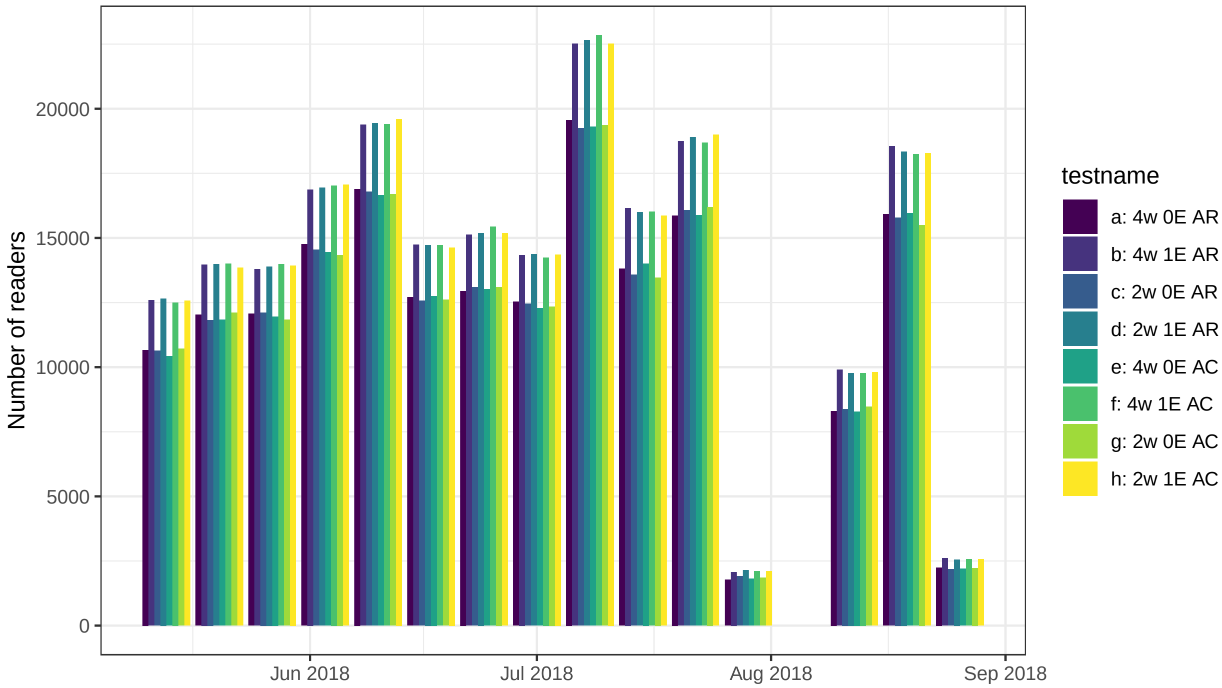

Notes: The figure shows the number of readers exposed to each experimental contract by week. The difference in shading represents the different phases of the experiment.

A.2 Solving the Problem

We solve the model in Section 2 with backward induction from the perspective of a subscriber. Since prices are non-decreasing over time, if a subscriber wishes to become unsubscribed at some period t, they will also want to unsubscribe at every period after t. Therefore, the problem reduces to finding the earliest period t^* of unsubscription. We can represent never-subscribers with $t^* = 0$ and always-subscribers with $t^* = \infty$. Since we allow for potentially incorrect beliefs, we need to solve for the *perceived* utility from subscription and unsubscription when we solve the dynamic problem backwards. The reason is that when a reader makes a plan on if and when to unsubscribe if they were to subscribe, they make these decisions based on their beliefs about future costs and future inertia.

The problem becomes stationary at period T since at that point prices are fixed and an auto-cancellation period, z, if it exists, is sooner than that ($z < T$). At period T the subscriber's problem is whether to unsubscribe or remain subscribed forever. The perceived utility of remaining subscribed is $\sum_{\tau=0}^{\infty} \delta^{\tau}(v_i - p) = \frac{v_i - p}{1-\delta}$. In contrast, the perceived utility from unsubscribing is $v_i - p - \hat{c}^u$ if the subscriber is able to unsubscribe and is not inert. Yet, the subscriber believes that with per-period probability $\hat{\phi}_i$ they will fail to unsubscribe and have to try again at a later period. Therefore, the perceived utility from unsubscribing at T, and trying

Figure A.2: Distribution of Unique Number of Cookies ("Readers") Per Subscriber

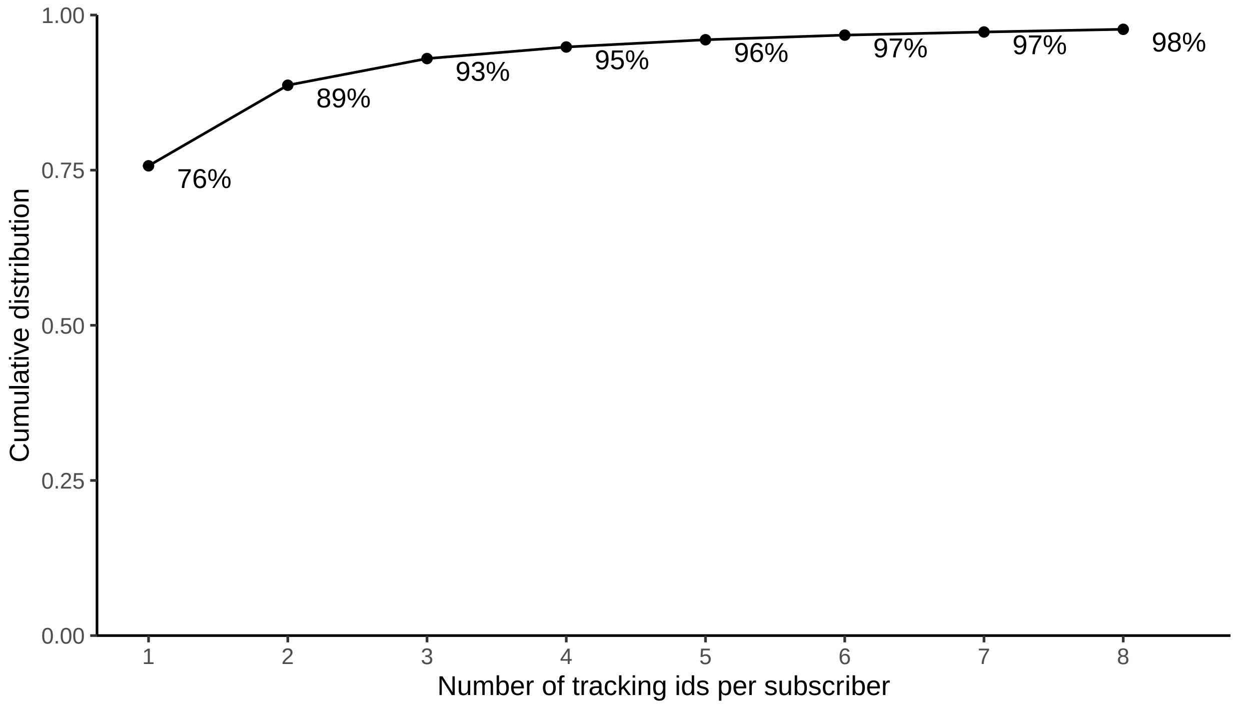

Notes: The figure shows the cumulative distribution of the unique number of cookies for each subscriber.

at all following periods if unsubscription failed, is $\sum_{\tau=0}^{\infty}\left[\hat{\phi}_i^{\tau}\delta^{\tau}\left(v_i - p - \left(1-\hat{\phi}_i\right)c^u\right)\right] = \frac{v_i - p - \left(1-\hat{\phi}_i\right)\hat{c}^u}{1-\delta\hat{\phi}_i}$.[27]

Therefore, the perceived value in period T from the perspective of an earlier period is the max of attempted cancellations and remaining subscribed

$$\hat{V_i}^T = \max\left\{\frac{v_i - p - \left(1-\hat{\phi}_i\right)\hat{c}^u}{1-\delta\hat{\phi}_i}, \frac{v_i - p}{1-\delta}\right\}$$

It is worth noting the effects of perceived inertia. If the subscriber expects to be non-inert, $\hat{\phi}_i = 0$, then we get the standard case of immediate cancellation versus remaining subscribed forever. If, in contrast, the subscriber expects to be fully inert, $\hat{\phi}_i = 1$, then both terms are identical since in either case the subscriber remains subscribed forever.

Using that value function we can solve backwards for $t < T$, as in any period except two ($t = z$ and $t = 1$), the decision is between trying to cancel (left) or remaining subscribed (right):

$$\hat{V}_i^t = \max\left\{v_i - p_t - \left(1-\hat{\phi}_i\right)\hat{c}^u + \hat{\phi}_i\delta\hat{V_i}^{t+1}, v_i - p_t + \delta\hat{V}_i^{t+1}\right\}$$

[27]If perceived inertia is $\hat{\phi}_i = 0$, we take the non-consensual convention that $\hat{\phi}_i^0 = 1$

Figure A.3: Types of Contracts Taken by Experiment Participants

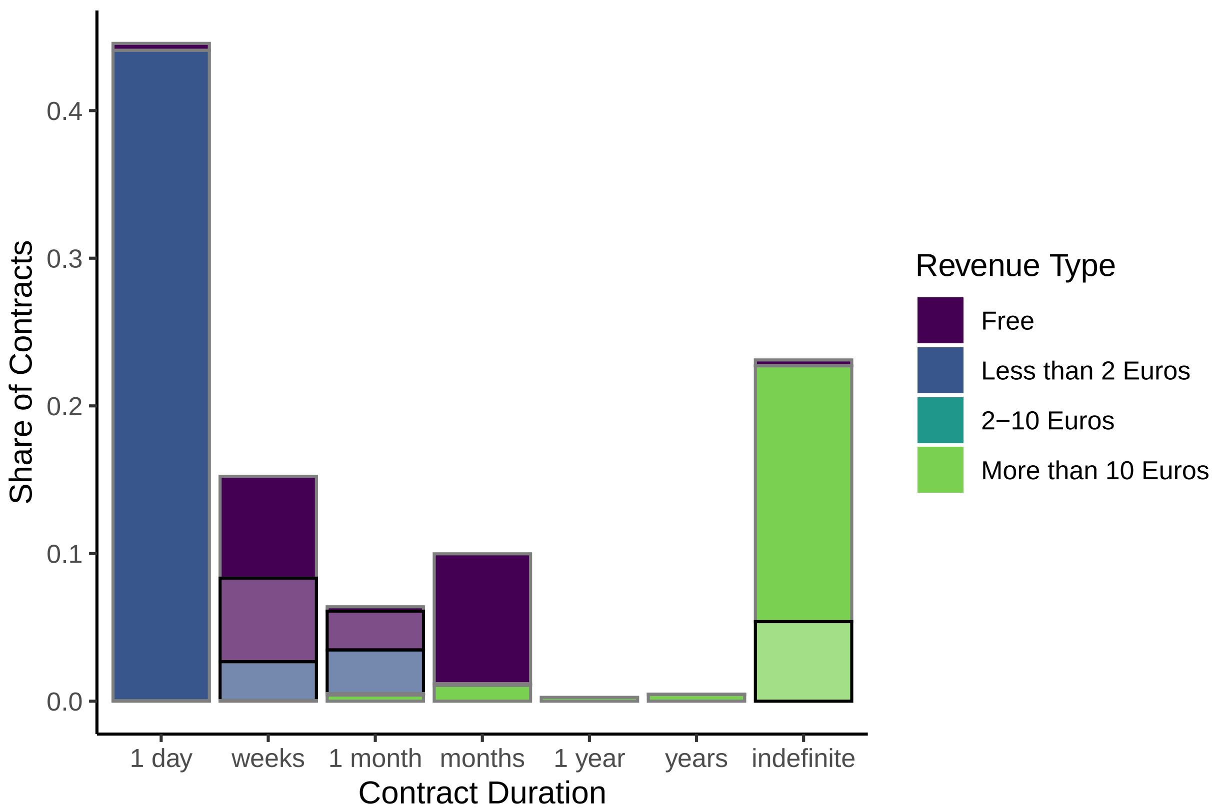

Notes: The figure shows shares of contracts taken, characterized by their maximal duration (horizontal axis) and revenue (color). For example, almost half of all contracts are daily passes that cost €1.99. The dark rectangles highlight the experimental contracts—the auto-cancellation contracts are either two weeks or a month (four weeks), and are either free or less than €2; the auto-renewal contracts are indefinite with a revenue above €10.

Figure A.4: Cumulative Revenue when Auto-renewal Contracts are Served Relative to Auto-cancel Contracts

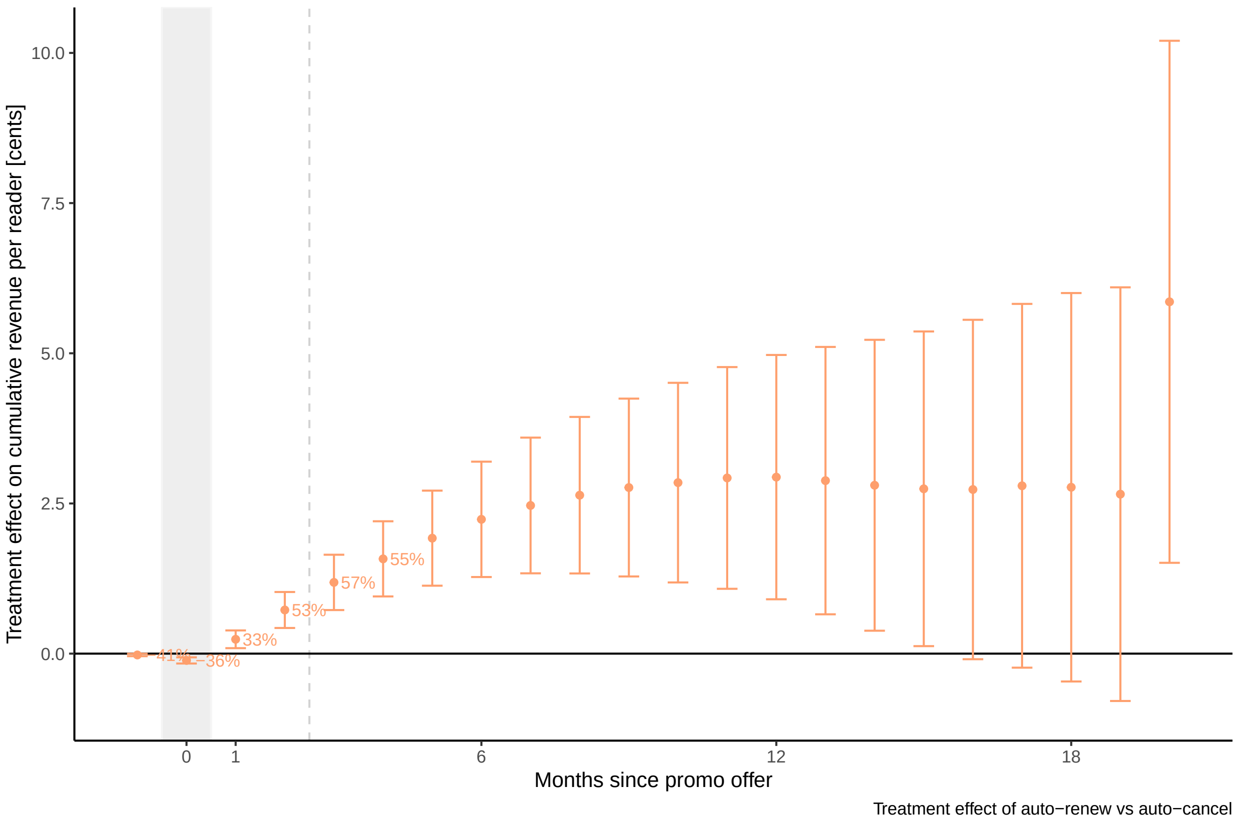

Notes: The figure plots the estimated average intent-to-treat effect of serving an auto-renewal relative to an auto-cancel contract on the newspaper's cumulative revenue.Specifically, we plot the estimated coefficient β_1 *from equation* (2) *for every month with cumulative revenue as the dependent variable. Month 0, shaded in gray, is the promo period (two weeks). The left-most points at month -1 are before readers hit the paywall. Month 1 is the first month after the promo ends, and so forth. The last point, "after-promo", aggregates across all after the promo time periods. Percentages next to the effect size compare the effect to the mean level of the omitted group.The error bars show 95% confidence intervals. Standard errors are clustered at the individual reader level.*

Figure A.5: Subscription Levels when Auto-Renewal Contracts are Served Relative to Auto-Cancel Contracts

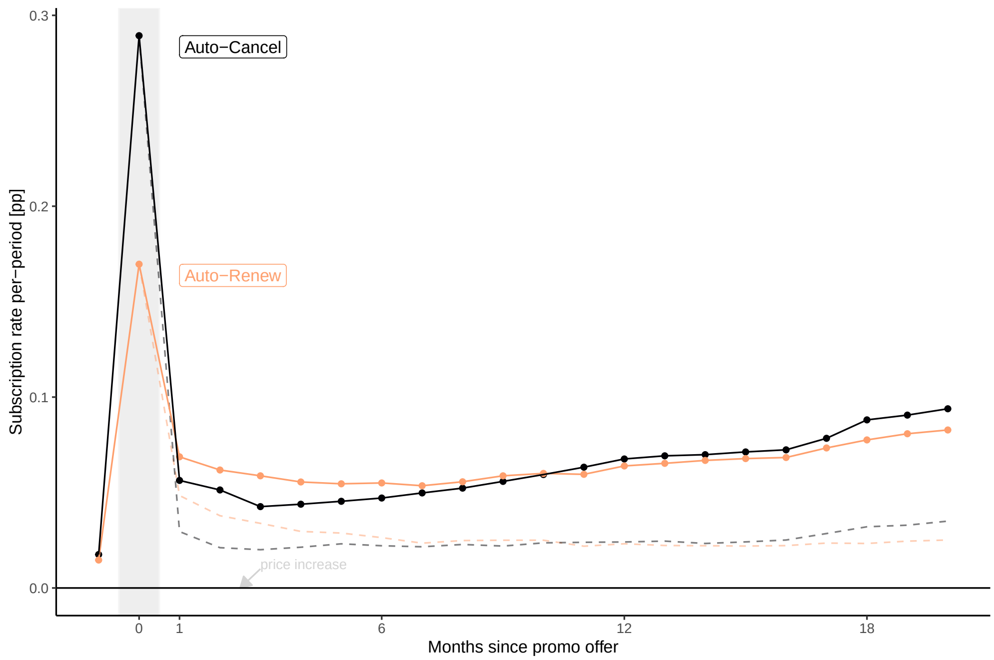

(a) Subscription Rate (Proportion of Days a Reader Subscribed)

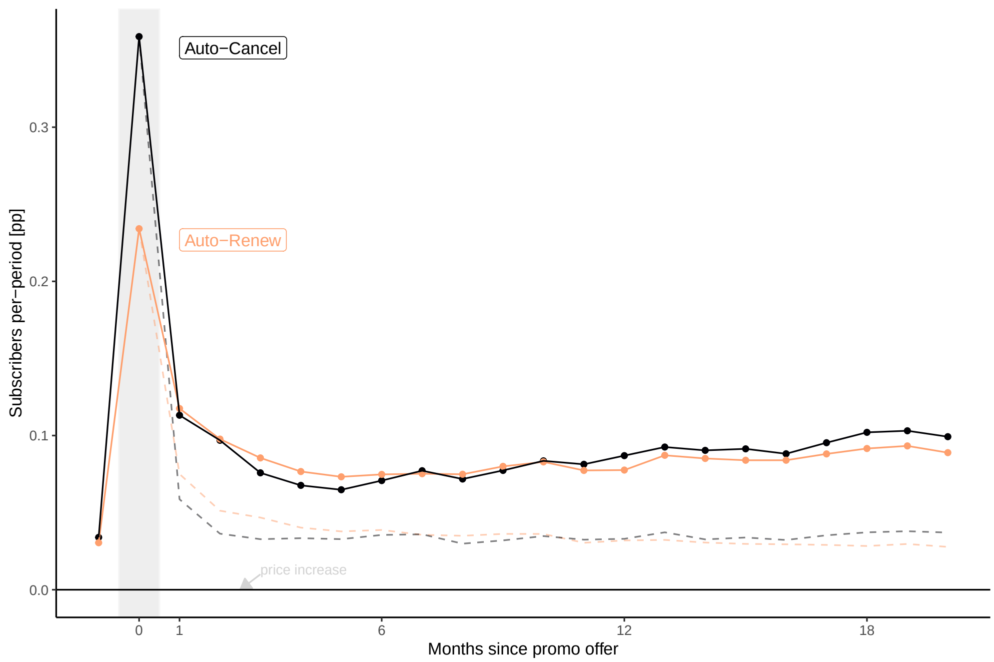

(b) Extensive Margin (Whether the Reader Subscribed At All)

Notes: The figures plot the levels along with estimated average intent-to-treat effect of serving an Auto-renewal relative to an Auto-cancel contract on consumer subscription behavior. Specifically, we plot the estimated coefficient $\alpha + \beta_1$ from equation (2) for every month. Month 0, shaded in gray, is the promo period (two weeks). The left-most points at month -1 are before readers hit the paywall. Month 1 is the first month after the promo ends, and so forth. The dashed lines show the results when a post-promo subscriber is counted only when she also subscribed during the promo period – this is the sample used in Section 8

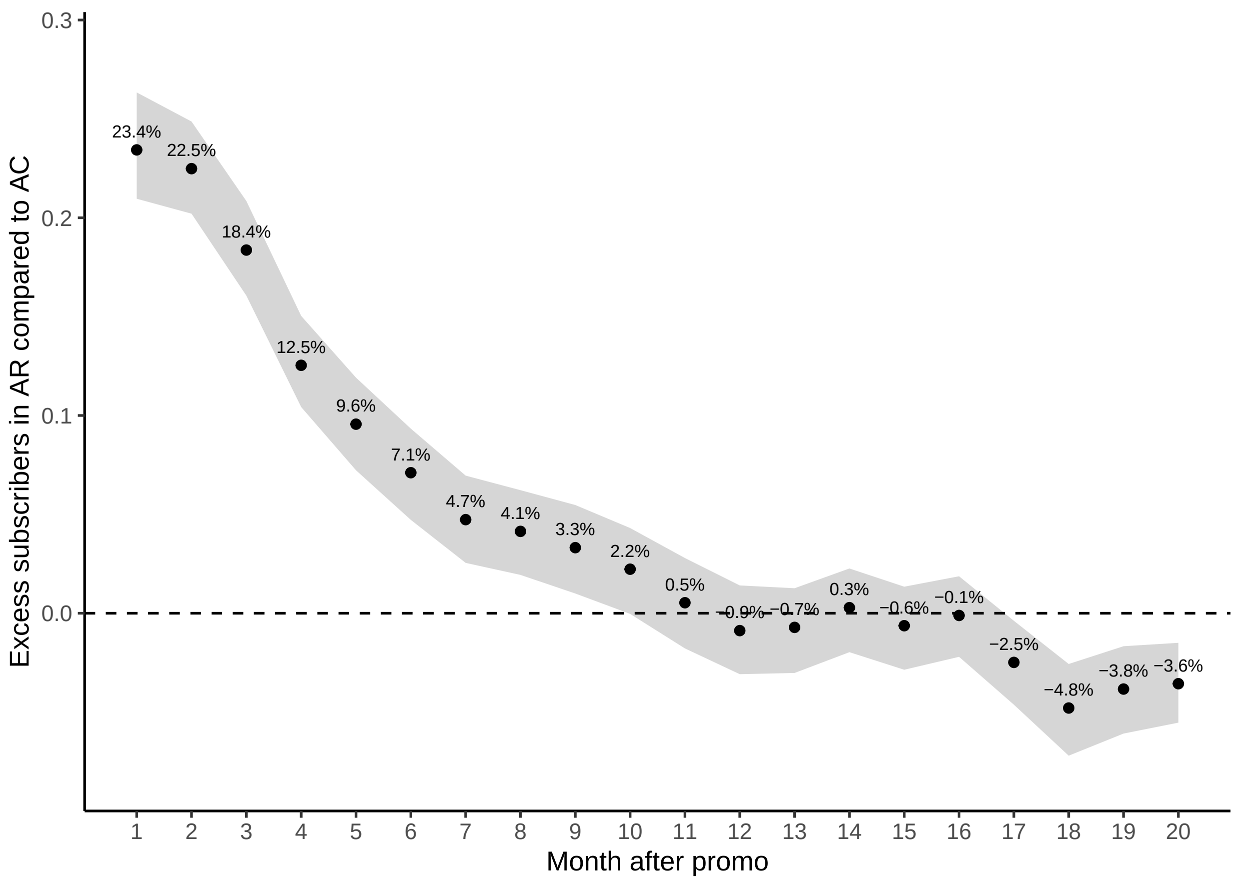

Figure A.6: Excess subscribers in AR vs AC

Notes: The figure shows lower bounds on excess subscribers in auto-renew after the promo. Excess subscribers are those who are subscribed but wish they hadn't.

Figure A.7: Distributions of Reader Types for Subscribers in the Auto-Renewal and Auto-Cancel Group by Period

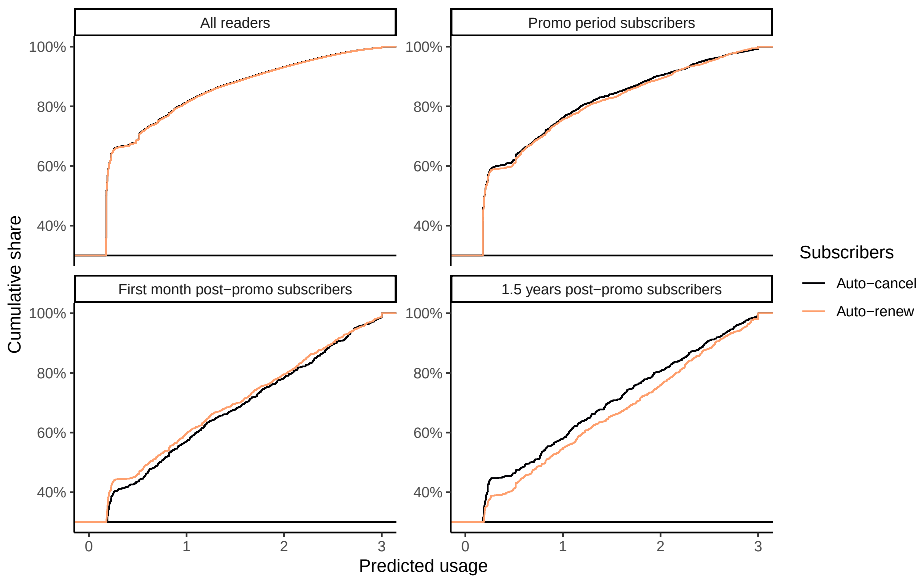

Notes: We use group G readers' pre and post-experimental usage data to predict post out-of-sample newspaper usage for the main sample. We use predicted usage as a proxy for reader type. The figure shows the distributions of predicted reader types for subscribers in the auto-renewal group and auto-cancellation group. Each panel contains a different sample of subscribers - the top-left panel shows balance of value across all readers in the experiment; the top-right panel shows only subscribers during the promo-period; the bottom-left panel shows subscribers during the first month after the promotional period ends; and the bottom-right panel shows those subscribed two years after the promo ends.

The subscriber will wish to remain subscribed if the future value is not too negative, $\hat{V}_i^{t+1} \geq -\frac{\hat{c}^u}{\delta}$. Note that inertia cancels out because it affects both the cancellation cost and the chance of continuation.

In period $t = z$, when the contract automatically cancels, the decision is slightly different since inertia nor costs come into play:[28]

$$\hat{V}_i^{t=z} = \max\left\{v_i - p_t, v_i - p_t + \delta\hat{V}_i^{t+1}\right\}$$

Here, a subscriber will only renew for a strictly positive continuation value, $\hat{V}_i^{t+1} > 0$, because there are no cancellation costs.

Finally, at period 1 the reader decides if to subscribe at all given the subscription costs against the net present value of a subscription with planned or attempted cancellation at a later stage. So will subscribe if $v_i - p_1 + \delta\hat{V}_i^2 - c^s \geq 0$ (we assume that subscription costs are "paid" at the time a contract starts and are known).

This setup highlights the different forces that affect perceived and actual inertia, and how they translate into observable subscription and usage patterns. Those who value the subscription will sign up regardless, as auto-renewal or auto-cancellation do not affect them. However, those who draw some value, enough to try but not enough to pay a full price, are possibly affected. For them, *perceived* future cancellation costs and inaction reduce take-up of an auto-renewing contract due to the risk of being locked-in paying for a product they do not like. The *actual* costs lead to an increase in the share of long-term subscribers roughly to the extent these subscribers underestimate the costs at sign-up; and actual inaction leads to a persistence in the number of medium-run subscribers to the extent that these subscribers underestimate their future inaction.[29] As mentioned above, habit formation or learning – some consumers start to like the product after trying it (or learn that they like it) – can also be a force that creates inertia. We can think of that as a shift to v_i due to subscribing, and will address that in the empirical section.

[28] We can think of inertia also tampering the choice to renew. However, we assume that renewal costs are minuscule and once a subscriber comes back to the newspaper website they are prompted to renew with a single click anyway. This is a simplification, but a realistic one.

[29] Readers can be marginal in their valuation, which might lead some to accept the subscription even if they value it less than the full price and know they might get locked-in. We will address what might be the measure of these potential subscribers later.

A.3 Experiment details

A.3.1 Taking up an experimental offer

Readers get exposed to the experimental offers when they reach the newspaper's paywall. Our newspaper's content arrangement is sometimes referred to a soft-paywall which stands in contrast to a so-called hard-paywall whereby a reader needs to pay for reading any content (e.g., academic journals, Financial Times).

From the reader's standpoint, the experimental offer is presented as follows. Upon hitting the paywall, the reader is presented one of eight experimental treatment offers in a banner and a reduced teaser version of the article that the reader intended to read.[30] After clicking on the experimental offer, all readers have to go through the standard three steps in order to start the trial. First, the reader is asked to register and provide an email address and choose a password. Second, the reader enters her personal and payment information. Lastly, the reader can view the terms and conditions of the selected offer, and click on the check-out button to complete the purchase and enter a legally binding contract with the publisher. Both the email address and payment information are verified before the subscription starts. Importantly, these steps are identical across experimental groups.

A.3.2 Cookie Fragmentation

As is common in digital settings, one individual may be associated with multiple cookies across devices or sessions. We assess the scale of this fragmentation problem using subscriber login data. Figure A.2 shows the distribution of cookies per subscriber. Roughly 76% of subscribers are linked to one cookie, and another 13% have two. Fewer than 5% have four or more. While we can consolidate cookies for subscribers, we cannot do so for non-subscribers, who never log in. As a result, some individuals may appear multiple times in the dataset assigned to different treatments and contribute usage and non-subscription outcomes to multiple groups. This fragmentation inflates the number of observations with zero subscriptions. However, this zero inflation operates symmetrically across arms and does not affect estimated *relative* treatment effects. Importantly, if a fragmented user subscribes, we successfully link all their cookies and aggregate usage under their first-exposure treatment assignment. As a result, fragmentation may attenuate absolute levels of subscription and usage but should not bias comparisons of subscriptions across groups.

[30] The presentation of the experimental offers follows the newspaper's standard format and is determined by them. We tested the extent to which offers' terms were salient and clear. We showed Amazon MTurkers in the newspaper's market examples of the offers as displayed by the newspaper, and asked them to report back the main terms of the offers and to classify if an offer is an auto-renew or an auto-cancel offer. We found that 98 out of 101 participants classified correctly (two participants said "I don't know", and one participant misclassified).

A.4 Model estimation and Robustness

A.4.1 Model fit

Figure A.12 shows the empirical and simulated moments in light circles and dark triangles, respectively. The model is able to fit most of the data patterns with a response to the promotional price and duration, and with the differences in gradual decline of subscriptions in the AR treatments and a sharp drop in the AC treatments.[31]

A.4.2 Discussion

Robustness The above estimation is focused on outcomes of all promo-period takers, including the takers of non-experimental contracts whose characteristics (e.g., price schedule) may not match the assigned experimental contract, which may lead to inaccurate model estimates. Further, because we exclude the outcomes of those who did not take-up subscriptions in the promo-period, we are not capturing the full impact of spite which also influences later subscription choices. However, doing so keeps the model simpler because we avoid modeling time-varying subscription valuations while focusing on individuals who deal with inertia, keeps the promo period take-up rates consistent with the intent-to-treat analysis in section 6, and allows us to capture some of the impact of spite that occurs through exposure to AR vs AC which leads to lower future subscriptions for takers.

To examine the robustness of our estimates, we consider several alternative criteria: (i) restricting to subscribers as those who never subscribed to any contract before seeing the experimental promo; (ii) only count those who subscribed to one of the experimental contracts at any point after the first exposure; or (iii), in the most restrictive case, only consider as subscribers those who subscribed to an experimental contract and did so within the promo window as counted from the first exposure.[32]

By estimating these alternative model versions, we observe the sensitivity of the estimates to this selection process. Figure A.8 and Table A.4 report the results. In the most restrictive sample of trial period experimental takers, we estimate a higher share of inerts (60%) and a larger fraction of naifs among them (12.5%, or 7.5% of the population). Importantly, the other structural parameters—especially inaction rates, unsubscription costs, and the spite penalty—remain stable across specifications. Overall, the main conclusion that a large share of inerts are sophisticated is robust.

[31] Our model captures most but not all patterns in the data. For example, it mispredicts the extent of drop after month 2 when the monthly price increases from €19.99 to €34.99. In order to have a lack of response to the price change in our model, we need a non-standard value distribution (with a mass of subscribers with monthly WTP between €0 and €2 where we get response in the promo price and duration; with a mass below €20 and above €35, so they would want to leave; but not between 20 and 35).

[32] In all cases we still count any subscription after the promo period, not just the original one. The main reason is that the auto-cancel renewals always receive a new contract identifiers with different designations.

So far, we assumed that sophisticates and naifs share the same inaction parameter. In principle, we can separately identify the inaction rates since the aggregate month-to-month inaction rate is an average across the population of survivors, which varies over time in its composition of naifs and sophisticates. Estimating this model, we find similar rates of inaction for the sophisticates as in the overall sample (confidence interval 0.73-0.94) and a noisy but *lower* inaction rate for the naifs (0.14-0.84). See the distributions of the separate inaction parameters in Appendix figure A.9 and the estimates in column 2 of Table A.5. One possible explanation for this is that sophistication is a learned trait, and more likely to be learned by those exposed and affected the most by it – those who are more likely to be inactive.

In summary, while many consumers are inert, the vast majority are also sophisticated about their inertia. These sophisticates cause the lower take-up of auto-renewal offers. We estimate that consumers learn their subscription valuation from promo trials, but this process's impact is limited. We estimate a spite factor about an order of magnitude lower than unsubscription costs.

Composition of subscribers It is clear that non-inerts not being exploited will lead to over-representation of inertial consumers among subscribers, which are a mix of high-value non-inerts and many lower value inerts. A similar argument explains where the common wisdom that most consumers are naive comes from. One possible explanation lies in selection into medium-term subscriptions. Naive inerts excessively become paying subscribers: low-value non-inerts do not get exploited and do not become paying subscribers; high-value non-inerts with high valuations stay subscribed in the long-run, due to their valuation; low-value sophisticates do not join, and only a subset of medium-value inert sophisticates are retained due to their inertia. Thus, excess subscriptions are heavily concentrated among naive inerts.

To quantify this selection, we decompose the estimated subscriber pool by type over time. Figure A.13 shows the share of each type among current subscribers in each month; Appendix Figure A.10 shows type shares in the full population. The top panel shows that among auto-renewal subscribers, naifs are over-represented up to a factor of five: although they are 4% of the consumer population, they represent up to 20% of consumers in the first months paying a full price. They gradually leave and after two years the distribution of takers is the same as their shares in the population. However, since new cohorts of subscribers get added each period, naive inerts remain disproportionately represented in the cross-section of subscribers.

Notice that the figure describes the shares among all subscribers, including those who are willing to pay the full price. Had we plotted the type shares of excess subscribers only, those whose willingness to pay is lower than the price, it would have shown almost entirely naive inerts, as implied by Appendix Figure A.10.

Is it possible to target readers with offers based on their sophistication? Specifically, is reader sophistication predictable, and if so, will it lead to naifs being targeted with more exploitative contracts?

Theoretical work shows that firms may wish to discriminate based on consumer naivete. We examine whether such targeting is both feasible and desirable. In practice, naivete could be irrelevant for auto-renewal – if most subgroups that can be targeted with auto-renewal yield higher long-term subscriptions when offered such contracts – or it could be the sole determinant, if the firm benefits only through short-term post-promotion revenues from offering auto-renew contracts to naive consumers.

We use our pre-experimental clickstream data—which includes the timing, page-views, and newspaper sections names—to predict the heterogeneity in the effects of our treatments.

A.4.3 Predictable Heterogeneity in Readers' Subscription Valuation

We first establish the predictability of individuals' subscription value in our context, and later use this predicted value to examine how the targeting of an auto-renewal vs auto-cancel contract varies with these valuations.

We construct a proxy for a reader's subscription value using the predicted out-of-sample usage with an approach similar to the one in Section 7.2. We run a regression forest (Athey et al., 2019) on the omitted groups, test-groups "G" and the status-quo group (which is about half the consumers, and equivalent to the misimplemented "G"), to predict total usage in the last three weeks of our data (starting from a week after first hitting the paywall to four weeks after). This total usage is predicted using the pre-experimental browsing behavior described below. These data are the same as the newspaper's first party data, which makes the exercise business relevant. Then, we predict and assign each reader in other test-groups their predicted usage score.

Covariates For the purpose of this prediction exercise, we construct 54 continuous variables for every reader based on their pre-experiment usage: The number of pages browsed by number of days before hitting the paywall (five or more before, four, three, two, one, and all page impressions on the day-of until hitting the paywall and entering the experiment); by category (e.g., politics); and by page type (is it open, metered, or always paywalled); in addition, we use the total pages browsed by day and page type.

Validation We validate our measure of consumers' valuations by predicting other variables that we expect to be correlated and consistent with our model. Namely, we predict that higher value readers will be more likely to subscribe and willing to pay more. Indeed, that is what we find as shown in Figure A.14. The figure shows that those readers who are predicted, out of sample, to consume the newspaper more regardless of the contract terms, are more likely to sign up during the promo period, bring in more revenue, to subscribe at all, and subscribe for longer. Each point in the figure represents one percent of readers.

A.5 Who Gets Targeted with Autorenewal?

Given this predictable heterogeneity, we examine whether contract targeting, auto-renewal versus auto-cancel, can be based on sophistication. We define a reader as naive based on predicted auto-renewal heterogeneous treatment effects on promo and initial subscription, and then examine if naifs would be assigned to auto-renewal or auto-cancel if the newspaper tries to maximize total subscriptions or revenue.

Targeting with AR vs. AC We run several causal forests to estimate heterogeneous treatment effects (HTE) of giving auto-renewal vs auto-cancel offers, with the same pre-experiment browsing behavior as above as covariates. We estimate auto-renewal treatment effects on two outcome variables the firm might focus on (1) total revenue, (2) the probability of subscribing at all after the promo ends.

We "assign" readers who were not a part of the training dataset to be targeted with auto-renew if the individually predicted effect of AR relative to AC is positive, and with auto-cancel otherwise. A comparison of the treatment effects across the targeted subsamples, as shown in Table A.6, validates our targeting strategy and demonstrates significant predictable heterogeneity in AR effects by showing that those assigned to be targeted with AR have a positive AR effect and those assigned to AC have a negative AR effect.[33]

Classifying readers as Naive After establishing the validity of the HTE, we estimate two more heterogeneous treatment effects that help us classify a reader as naive: their promo take-up is unaffected by auto-renew; but they remain subscribed in month 1 due to auto-renew though wouldn't be subscribed at that period under auto-cancel. We operationalize these conditions as having an estimated zero treatment effect of auto-renew on subscription in the promo period, and a positive effect in the first month after the promo.[34]

Who Gets What Contract? With these classifications, we examine targeting outcomes. Specifically, we compute Pr(Naive | AR), the share of naifs among readers targeted with auto-renewal; and Pr(AR | Naive), the share of naifs assigned to auto-renewal. Pure naivete-based discrimination implies over-representation of naifs among AR targeted, and as a perfect predictor of AR assignment (Pr(AR | Naive) = 1).

The empirical results are subtle and depend on the outcome the firm chooses to maximize, as shown in Table A.7. When the firm maximizes the number of subscribers after the promo, only 11% of readers would

[33]Since the estimates are the predicted effects based on leave-out estimation and each subsample is based on counterfactual treatment assignment, there is still de facto random assignment to auto-renew versus auto-cancel in our data within each sub-sample. Therefore, we can validate the predicted effects by estimating the treatment effects separately for each sub-sample (this is what Hitsch et al. (2024) calls "off-policy evaluation using data with randomized targeting"). For example, the effect of auto-renewal on total revenue should be positive for the subsample for which it is predicted to be positive and negative for the remaining subsample. Indeed, this is what we find for the four outcomes and their eight corresponding sub-samples.

[34]What constitutes a zero treatment effect? Since standard errors of individual-level effects are wide and hard to estimate well, we instead create a tolerance window around 0. We calibrate the window size such that the share of naifs matches the structural estimation. Reassuringly, the results are insensitive to the window size (see Figure A.11).

be assigned to auto-renew. 10% of those assigned to auto-renew are naifs, compared to 3% among those assigned to auto-cancel. That is, naifs are over-represented among auto-renewal targeted readers. However, due to the different base rate of assignment, only 30% of naifs will be assigned to auto-renewal. The results flip if the firm is maximizing revenue. In that case, 79% of readers would be assigned to auto-renew. 3.4% of those assigned to auto-renew are naifs, compared to 5.4% of those assigned to auto-cancel. So naifs are under-represented among auto-renewal targeted readers. Due to the different base rate of assignment, a majority of 70% of naifs are assigned to auto-renew.

Table A.7: Who is Being Targeted

Firm's maximization objective	Pr(AR)	Pr(Naive | AR)	Pr(Naive | AC)	Pr(AR | Naive)
After promo subscribers	11.4%	10.0%	3.1%	29.6%
Total revenue	78.8%	3.4%	5.4%	70.3%

Notes: The table shows targeting of naifs. A reader is defined as naive if their promo takeup is predicted to be unaffected by auto-renewal, but predicted to stay for the first full price month in auto-renew but not auto-cancel. The first column described the target function to maximize; the second shows that total share of readers assigned to auto-renew; the third and fourth columns show the share of naifs among those targeted with auto-renew or auto-cancel, respectively; the fifth column show that probability of a naive being assigned to auto-renewal.

Naivete Positively Predicts Assignment (Conditional on Value) To further examine this, consider Table A.8, which expands the last column of Table A.7. The table shows that, conditional on predicted value (via usage percentile fixed effects), naivete is a significant positive predictor of auto-renew assignment.[35] At the same time, the coefficients and the explained variation in targeting in terms of the R^2 are small, implying that while naivete gets targeted, not all naifs get targeted with auto-renew and other factors matter much more, especially for revenue maximization.

Limitations The above analysis relies on the assumption that the HTEs correctly identify naifs. Errors in classification should attenuate targeting on naivete, even if it is a strong predictor. The errors can arise from the algorithm or from the data itself. We cannot verify if browsing behavior before the experiment is a good predictor of naivete.

Overall, naivete, as defined by HTEs, is a statistically significant but not dominant driver of contract targeting. It predicts assignment to auto-renewal conditional on value, but most targeting is driven by predicted subscription behavior more broadly.

[35]It is worth noting that while naivete predicts lower likelihood of assignment to auto-renewal to maximize revenue (Table A.7 and column 3 in Table A.8), controlling for predicted usage makes naivete a positive predictor of higher auto-renew assignment.

Table A.8: Does Naivete Predict Auto-Renewal Assignment?

	Firm Objective: Max subscriptions		Firm Objective: Max revenue	
Model:	(1)	(2)	(3)	(4)
Constant	0.1066		0.7910	
	(0.0003)		(0.0004)	
Naive	0.1892	0.1896	-0.0882	0.0321
	(0.0020)	(0.0020)	(0.0020)	(0.0021)
Fixed-effects				
Predicted usage percentile		Yes		Yes
Observations	1,389,869	1,389,869	1,389,869	1,389,869
R^2	0.01316	0.46602	0.00172	0.37090

The table shows if naivete predicts being targeted with auto-renewal, which is the dependent measure of the regressions. Columns 1+2 predict targeting with auto-renewal if the firm's goal is to maximize post-promo subscribers; columns 3+4 do so when the firm's goal is to maximize total revenue. Standard errors clustered at the tracking ID level.

A.6 Stochastic cost models and estimation

In this section we keep the same model, but make different assumptions about costs and inaction. We are trying to understand if a model with no inaction, but with stochastic cancellation costs, would need naivete to describe the data patterns.

Here, we assume that there is no inaction ($\phi = 0$), but that the unsubscription costs c^s are stochastic instead of deterministic. We assume that they are drawn from a uniform distribution $c \sim U[\psi_0,\ \psi_0 + \psi]$.

We estimate three variations on that theme. First, we estimate a standard stochastic costs model, where all readers are sophisticated and share the same cost distribution. Second, we estimate a model with two types with stochastic costs – some who know the cost shock distribution, and some who underestimate it and think the stochastic costs are distributed over a lower range. Finally, we estimate a model with three types that mirror the types we have in the main model with inaction, equating the number of parameters we have in the main model. We focus on the last model in the exposition below.

We assume that ψ_0 is the same for all readers. Types differ in ψ, where 'non-inerts' have $\psi = 0$, and inerts have $\psi > 0$. That is, the non-inerts have a fixed cost ψ_0, while inerts have some stochastic cost drawn each period. Sophisticate inerts know their distribution correctly, $\hat{\psi} = \psi$; naive inerts think that $\hat{\psi} < \psi$. That is, the assumptions are similar to our main model, but we replaced inaction and beliefs about inaction with the cost distribution. Table A.9 summarizes the differences. All other parameters in the model remain the same.

To solve the model, we find the Bellman equation of the value of a subscription given the expectations of future actions in the steady state, when prices are stable. We then derive the reservation cost. The reader

Table A.9: Comparison of Inaction and Stochastic Costs Models

Type	**Inaction model**				**Stochastic costs model**			
	Inaction	**Lowest cost**	**Highest cost**	**Belief**	**Inaction**	**Lowest cost**	**Highest cost**	**Belief**
Sophisticate	$\phi > 0$	ψ_0	ψ_0	$\hat{\phi} = \phi$	$\phi = 0$	ψ_0	$\psi_0 + \psi$	$\hat{\psi} = \psi$
Naïve	$\phi > 0$	ψ_0	ψ_0	$\hat{\phi} = 0$	$\phi = 0$	ψ_0	$\psi_0 + \psi$	$\hat{\psi} < \psi$ or $\hat{\psi} = 0$
Non-inert	$\phi = 0$	ψ_0	ψ_0	$\hat{\phi} = 0$	$\phi = 0$	ψ_0	ψ_0	$\hat{\psi} = 0$

remains subscribed if the draw she gets is above the reservation cost, and cancels if the draw is below it. We then solve backwards to find the reservation cost at each earlier period before the steady state.

Define the value-function V by

$$V \;=\; b \;+\; \mathbb{E}\big[\max\{-c,\; \delta V\}\big], \qquad c \sim U[\psi_0,\; \psi_0 + \psi].$$

Where $b = v_i - p$ is the net benefit of a subscription. An optimal "cutoff" policy chooses a threshold c^* so that

$$\text{stop if } c \leq c^*, \quad \text{continue if } c > c^*.$$

Indifference at c^* requires

$$b - c^* \;=\; b + \delta V \;\;\Longrightarrow\;\; c^* = -\,\delta V.$$

Hence the Bellman equation is

$$V \;=\; \frac{1}{\psi}\left[\int_{\psi_0}^{-\,\delta V} (b - c)\, dc \;+\; \big(\psi_0 + \psi + \delta V\big)(b + \delta V)\right].$$

Evaluate the integral:

$$\int_{\psi_0}^{-\,\delta V} (b - c)\, dc = b\big(-\delta V - \psi_0\big) \;-\; \frac{1}{2}\big(\delta^2 V^2 - \psi_0^2\big).$$

Multiply both sides of the Bellman equation by ψ:

$$\psi\, V = b\psi \;+\; \delta V(\psi_0 + \psi) \;+\; \tfrac{1}{2}\,\delta^2 V^2 \;+\; \tfrac{1}{2}\,\psi_0^2.$$

Rearrange into a quadratic in V:

$$\tfrac{1}{2}\,\delta^2 V^2 \;+\; \big[\delta(\psi_0 + \psi) - \psi\big]\, V \;+\; \tfrac{1}{2}\,\psi_0^2 \;+\; b\psi = 0.$$

Multiply by 2:

$$\delta^2 V^2 \;+\; 2\big[\delta(\psi_0 + \psi) - \psi\big]\, V \;+\; \psi_0^2 \;+\; 2b\psi = 0.$$

Solve for V (choose the economically relevant root):

$$V = \frac{-\big[\delta(\psi_0 + \psi) - \psi\big] - \sqrt{\big[\delta(\psi_0 + \psi) - \psi\big]^2 - \delta^2(\psi_0^2 + 2b\psi)}}{\delta^2}.$$

Then the cutoff cost is

$$c^* = -\delta V = \frac{\delta(\psi_0 + \psi) - \psi + \sqrt{\big[\delta(\psi_0 + \psi) - \psi\big]^2 - \delta^2(\psi_0^2 + 2b\psi)}}{\delta}.$$

We then estimate parameters that minimize the squared distance between simulated and empirical monthly subscription rates, as in Section 8.2. Figure A.15 shows the empirical and simulated moments in light circles and dark triangles, respectively. The three types costs model is able to fit some data patterns with a response to the promotional price and duration, but struggles to generate the gradual decline of subscriptions in the AR treatments.

The results of the three types model are reported in Table A.10. We find higher unsubscription costs than a model with inaction, of at least €0.85 (for comparison, with inaction that was €0.22). 37% of consumers are (more) inert in that they face stochastic cost range of additional cost shock up to €17. Of the inerts, 13% are naive about the future cost shocks (or 5% in the population) in the sense that they think the cost is always €0.85, and the other 87% are sophisticated. All other parameters are similar to their estimates from a model with inaction. We therefore estimate a similar share of naifs in a model where gradual hazard is driven by stochastic cancellation costs

The other cost models are reported in Table A.5. In the other two models, everyone is inert by assumption. Everyone is exposed to stochastic cost shocks. If we assume that everyone is also sophisticated, a la Klemperer, we find similar cost estimates to the 3 types model, though noisier (column 4). If instead, we assume that some are sophisticated, and some are naive in the sense that they underestimate the cost shock range, we find that 77% are naive in the sense that they think the shocks will be uniformly distributed up to 64% of the maximal cost range.

The qualitative and quantitative fit of the stochastic costs models is worse than a model with inaction, as shown in Figure A.16. The figure shows the distribution of the sum of squared errors between the simulated and empirical moments of the different models. The distribution comes from the bootstrap procedure as every bootstrapped sample has its own best fit. The fit of the main model (in solid line) is to the left of the stochastic costs models (in dashed lines), and of those the three types model fits the best.

We conclude that a model of perfect foresight of stochastic costs struggles to fit the data, and also finds a positive share of inerts who are overly optimistic and have wrong beliefs about their future actions.

Table A.10: Structural Estimation Results

Name	Parameter	Estimate	95% Confidence Interval
Monthly inaction	ϕ	0	—
Lowest unsubscription cost (€-cent)	ψ_0	84.57	[64.87 - 287.93]
Range of unsubscription cost (€-cent)	ψ	1717.91	[63.21 - 5882.86]
Share of inerts		37%	[14.4% - 86%]
Share of naifs in population		4.95%	[3.45% - 51.2%]
Share of naifs among inerts		13.4%	[11.4% - 76.5%]
Prob of high value	λ	0.0943%	[0.0838% - 0.124%]
Monthly AR penalty (€-cent)	α	9.45	[0.14 - 9.67]
Subscription cost (€-cent)	c_s	54.06	[40.18 - 137.65]
Gamma dist (shape)		0.000849	[0.000782 - 0.00122]
Gamma dist (scale)		44.6	[34.1 - 48.4]
Monthly discount factor	δ	0.995	—

Figure A.8: Estimated Parameters For Different Samples

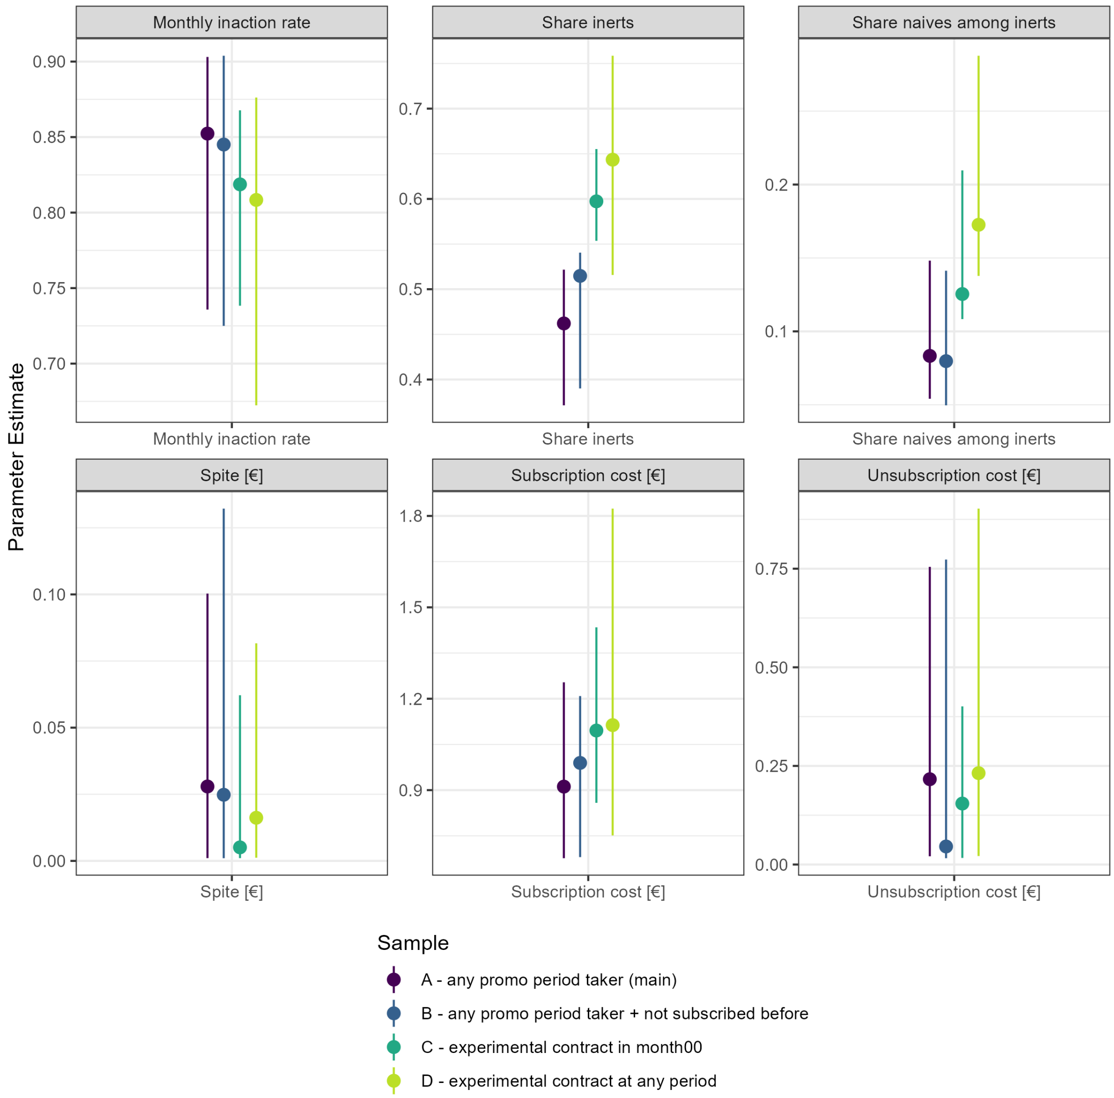

Notes: The figure shows the sensitivity of the main parameters of interest estimated on four different samples, together with 95% confidence intervals.

Figure A.9: Estimated Inaction Parameter for Sophisticates and naifs

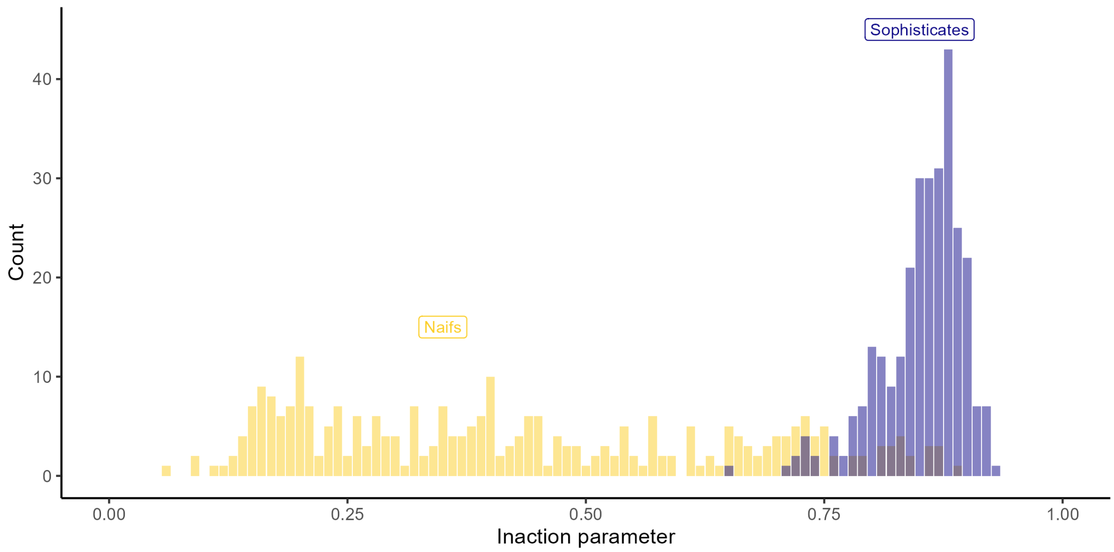

Notes: The figure shows histograms of bootstrapped estimates of the inaction parameter estimated flexibly by type of inert consumer: $\phi^{Naive}, \phi^{Soph}$.

Figure A.10: Estimated Type Shares for Subscribers, Unconditional

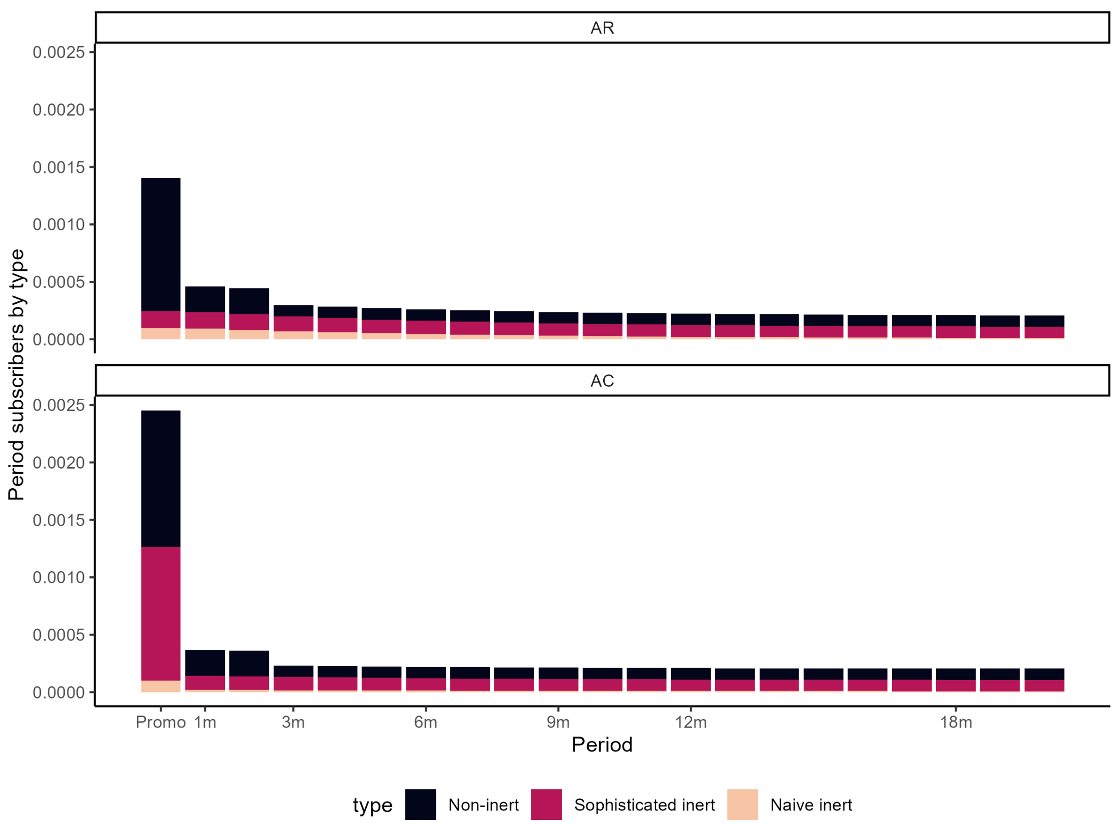

Notes: The figure shows the predicted shares of different types subscribed at different periods, out of the total population of consumers.

Figure A.11: Sensitivity of auto-renewal assignment to naivete definition

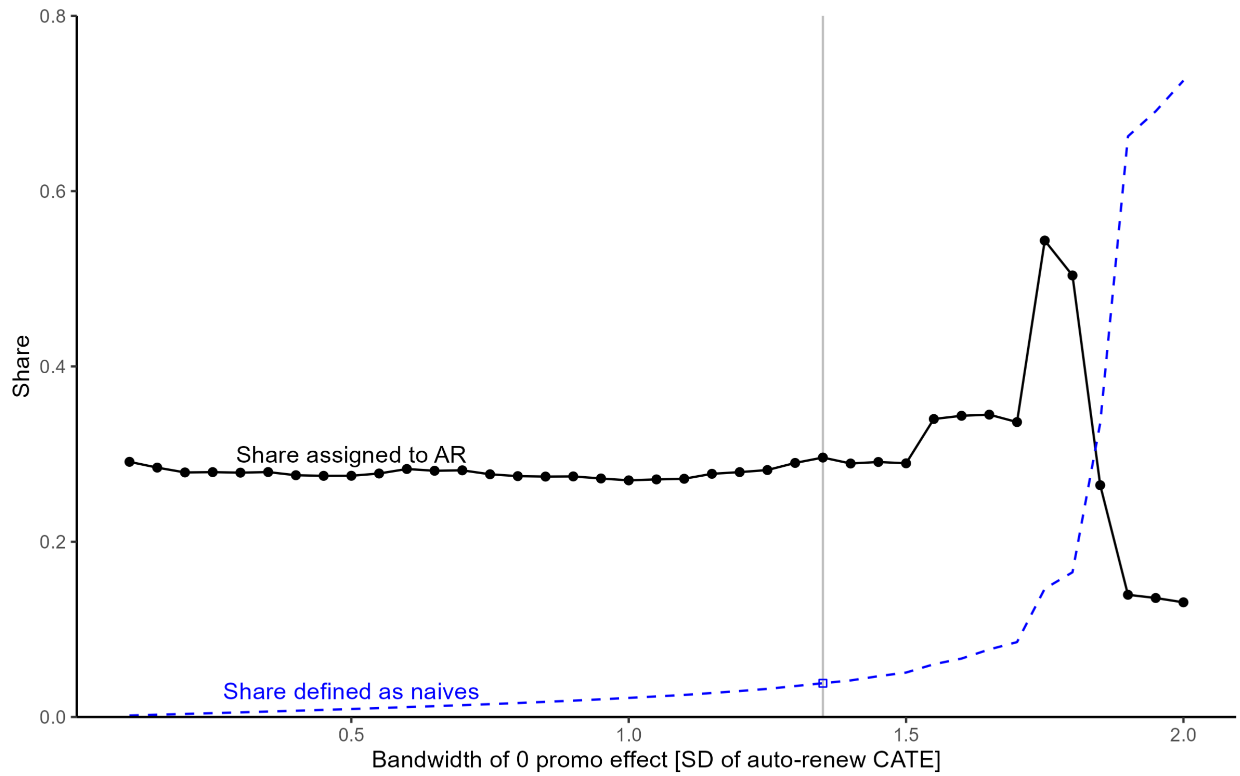

Notes: The figure shows the share of naifs counterfactually assigned to auto-renewal (black dots and line) and the share of readers classified as naifs (dashed), as a function of predicted auto-renew HTE size on promo period subscription that is considered a 0 effect. The vertical line represents the size chosen for the main estimation (1.35SD), giving 3.85% of naifs as in Table 3

Figure A.12: Empirical and simulated moments

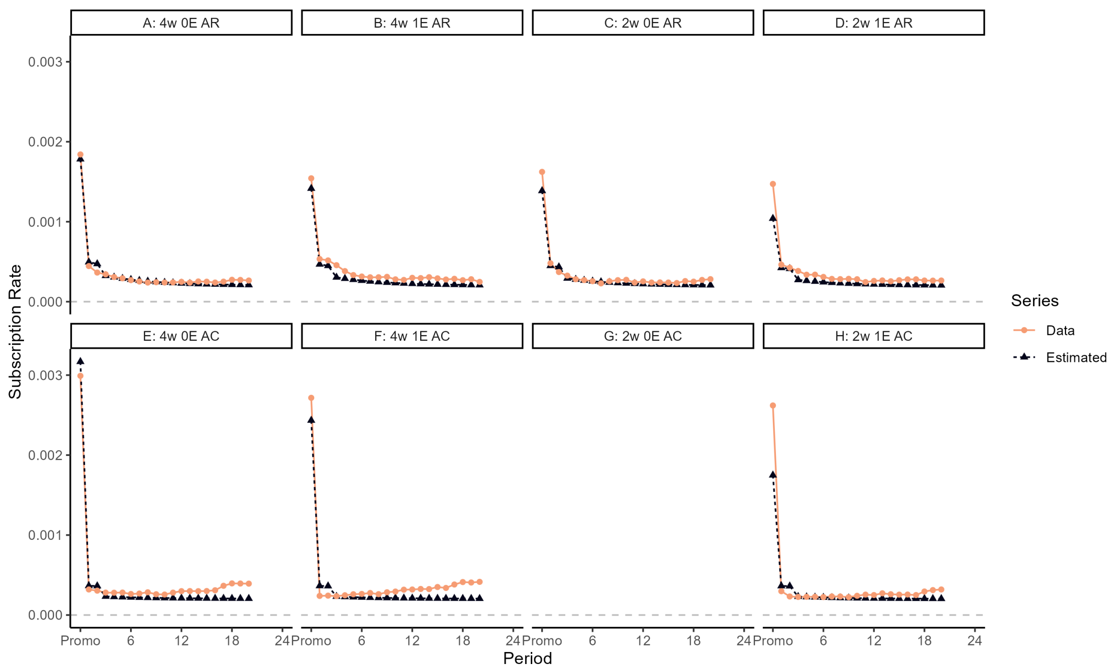

Notes: The figure shows the empirical moments used for estimation in light circles, and the simulated moments from the estimated model in dark triangles.

Figure A.13: Estimated Type Shares Conditional on Being Subscribed

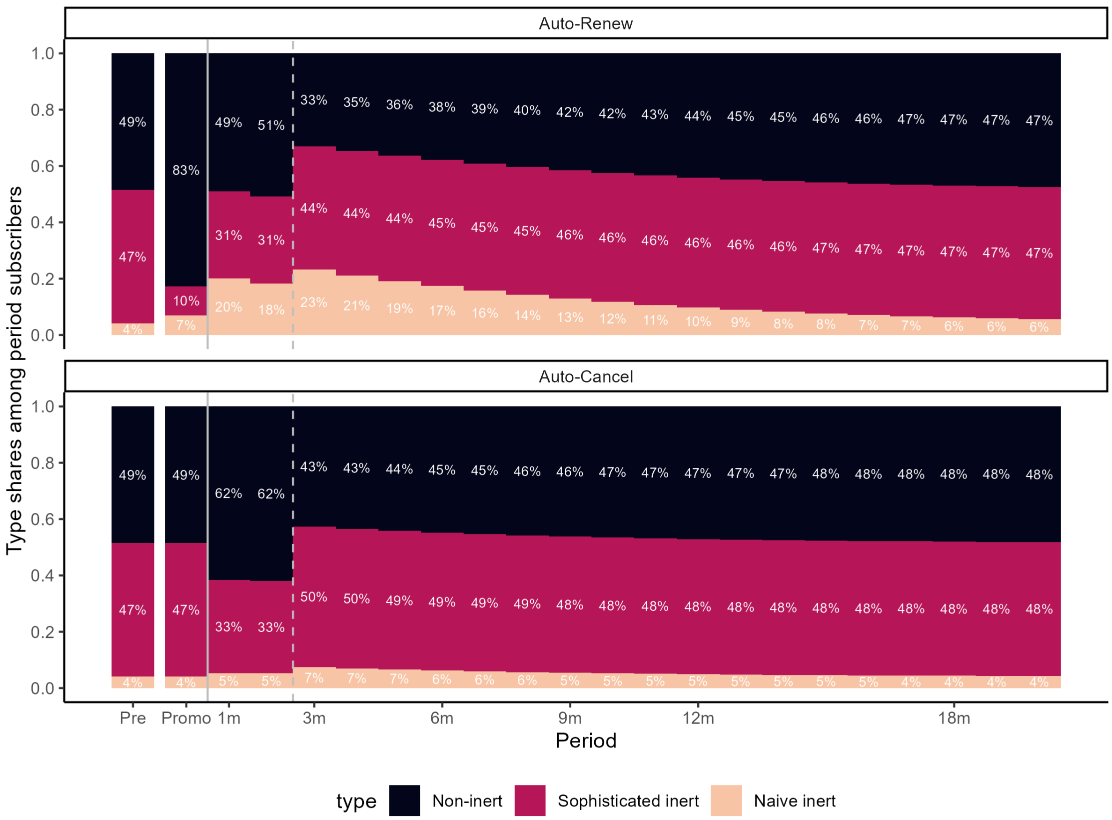

Notes: The figure shows the predicted shares of different types at different periods, conditional on being subscribed.

Figure A.14: Validation of Predicted Usage as Predicting Subscriptions

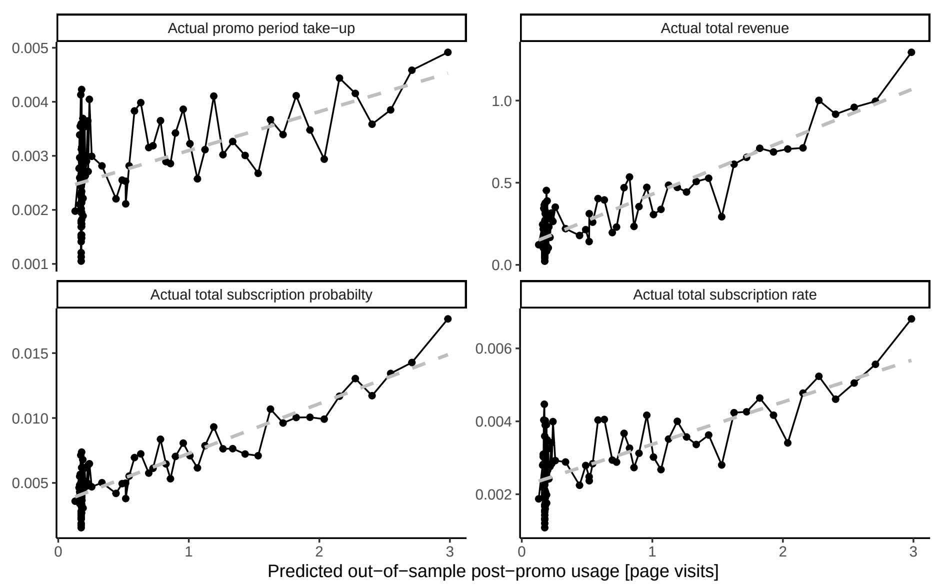

Notes: The figure shows the correlation between subscription behavior and the out-of-sample predicted usage. Each dot represents one percent of readers aggregated by their predicted behavior.

Figure A.15: Empirical and simulated moments - stochastic costs model

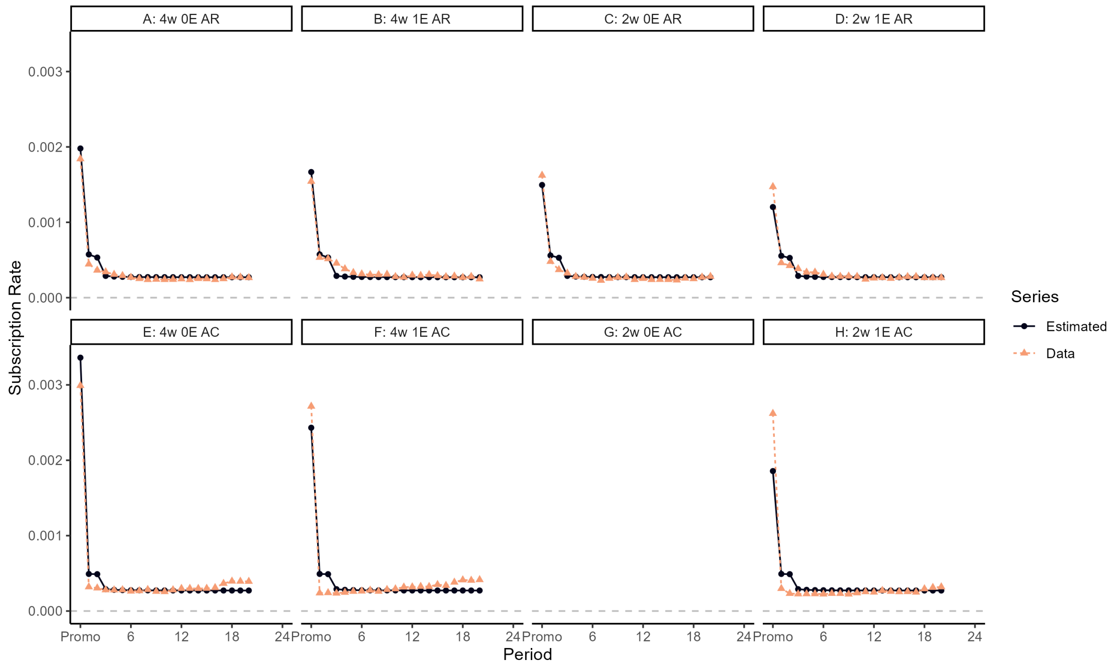

Notes: The figure shows the empirical moments used for estimation in light circles, and the simulated moments from the estimated costs model in dark triangles.

Figure A.16: Comparison of models fit

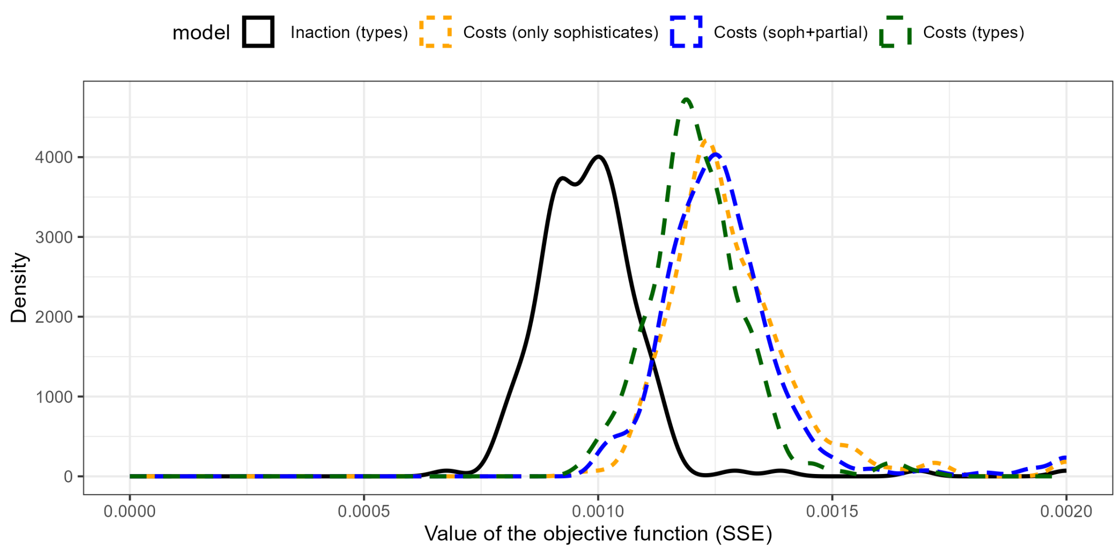

Notes: The figure shows the distribution of sum of squared errors between simulated moments under different models and the empirical moments.