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

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

April 28, 2023

Abstract

Are inert consumers aware of their future inertia? To answer this we run a field experiment that offers two million readers of a European newspaper auto-renewing or auto-canceling contracts. We find consumers are inert yet anticipate and account for their inertia: offering auto-renewing contracts lowers subscriptions by 24% and reduces subscribers by 10% over two years; most of the inert readers preempt inertia. Consumers' inertia 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, 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, and the Hebrew university. The study is registered at the AEA RCT Registry, AEARCTR-0011311.

1 Introduction

One of the most researched and widely documented characteristics of consumer behavior is inertia—the tendency of an individual to take no action and stay in the same state as before. For example, an individual is likely to pay a higher price for a subscription if they previously enrolled in it, 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 worsening of an option they previously chose, 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. Further, they will design their products to increase inertia.

Crucially, these consequences of inertia depend both on the degree of inertia and whether consumers are aware of it and how they factor it into their decision making. If consumers are not aware of their inertia, or are myopic about their future inertial behavior, they will not preempt it and get stuck with choices that seem good initially but turn out to be worse in the long run[1]. However, if consumers are aware of their behavioral limitations, they can avoid situations where they might be exploited due to inertia or find ways to limit its effects (see also Rodemeier (2022)). This awareness can discourage firms from creating situations that might be construed as exploitative by consumers. Thus, while inertia can have negative effects, self-awareness can mitigate its impact. Of course, consumers may differ in their future inertia awareness, and firms may create price, or inertia, discrimination in response to this heterogeneity (Eliaz and Spiegler, 2006).

In this paper, we empirically assess how inertia affects consumer decisions in the context of digital newspaper subscriptions contracts. We ask the following specific questions: What is the degree of inertia in consumer subscription choices? What is the degree of awareness to future inertia and how does it affect subscription choices? How do inertia and awareness differ between consumers? And what are the effects of these forces on firm incentives and outcomes?

To empirically infer whether consumers take their inertia into account while making decisions, 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., Handel (2013); Drake et al. (2022)). Additionally, to assess consumer sensitivity to inertia, we need variation in the degree of future inertia consumers face, which is rarely observed. Further, we need the variation in inertia to be exogenous, which is challenging to obtain.

We overcome these challenges by running a large-scale field experiment. Specifically, we randomized the terms of subscription offers received by 2.1 million readers who hit the digital paywall of a large European

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

daily newspaper. Our experiment is a three-way full factorial ($2 \times 2 \times 2$) design; a reader in our experiment is offered a subscription promo that (1) either automatically renews, by default, into a paid subscription for those who take the promotion unless they explicitly cancel it, or does *not* automatically renew but requires the promo taker to click to enroll into a paid subscription (which we call an auto-cancel offer), (2) has a promotional trial period for either four weeks, or two weeks, (3) has a promotional price of either €0, or €0.99. Importantly, all other aspects of the contract, including the information consumers need to provide to take up the offers are the same across the eight experimental groups. We then follow these potential subscribers for two years and observe their interaction with the platform and use the treatment arms to learn about inertia and responses to it.

Comparing the subscription take-up behavior during the promo period between those who receive the auto-renewal promo and those who receive the auto-cancel promo tells us whether consumers are sensitive to the future possibility of being defaulted into the paid subscription. We expect no differences between the two groups if consumers overlook the future outcomes, or believe (e.g., due to overconfidence) that they would costlessly cancel the subscription before it renews if they do not want the paid subscription anymore. The difference in continuation of subscription after the promo period helps us assess the actual degree of inertia caused by taking up the auto-renewal contract.

The experimental variation in price and promo duration serves the following purposes. First, it enables us to estimate "learning or habit formation", or the effect of product trial on the long-term subscription rate. This is useful in interpreting the effect of serving the auto-renewal vs. auto-cancel offer. It also informs us about the channels through which inertia operates. Second, varying the promotional price and the subscription renewal terms helps us estimate incidental parameters, such as the hassle cost of subscribing and the subscription valuation of marginal subscribers. In turn, it enables us to characterize the distribution of inertial types and their inertia expectations.

Our first main finding is that consumers are less likely to take a future-inertia-exploiting contract. We find that 24% fewer readers take up any newspaper subscription during the promotional time period when *offered* an auto-renewal offer, relative to an auto-cancel offer. Thus indicating that some readers recognize and adapt their behavior to future auto-renewal terms and, overall, they prefer the promo that does not convert into a paid subscription by default.

Second, we find that some consumers are more inert than they anticipate. While the initial take-up is lower for the auto-renewal groups, we find that the subscription-rate (the proportion of days a reader subscribes to the newspaper) is higher by 20% among those who received the auto-renewal offers, relative to the auto-cancel ones for about four months after the promotion period ends. Beyond this time, the difference in subscription rates declines. A year after the end of the promo, the subscription rate is higher in the auto-cancel relative to the auto-renewal groups. Among those who take up an auto-renewal promo and do not cancel, we quantify the actual inaction that causes inertia to be 0.72 – a 72% monthly chance that a consumer does not cancel a subscription they would rather not have.

Examining the actual individual-level usage of the newspaper's website, we see that auto-renewal subscribers rarely read the newspaper, further establishing that auto-renewal subscribers do not use their subscription for consumption.

Third, offering inertia-inducing contracts discourages readers from engaging with the newspaper. On the extensive margin, the readers who were assigned an auto-renewal offer are 10% less likely to become paid subscribers at any time in the two years after the promotion, relative to those offered an auto-cancel offer. We do not observe such a push-back for other experimental factors; even though €0.99 vs. free promo, and two weeks vs. four weeks both cause 9-10% fewer people to subscribe during the promo period, they have no impact in the time period of two years after the promo. This pattern indicates that the negative impact on the extensive margin is the direct effect of the auto-renewal contract term, and not due to lower trial caused by it in the promo period. It also suggests that the medium-run (up to six months after the promo) increase in subscription-rates experienced by the newspaper is coming from *few* individuals who end up paying more on the intensive margin.

We then use a simple choice model to estimate anticipated and actual inertia types. In the model, inertia is driven by either inaction (e.g. due to forgetfulness or procrastination) or switching costs, and consumers differ, non-parameterically, by their value of the subscription. There are three types of actual inertia – some consumers are fully-inert and will not cancel their subscription, some are non-inert and act as if there are no costs or frictions, and the rest are partially-inert who with some probability will not take an action they would wish to take. Independently, we allow each inert consumer to be either sophisticated, i.e. to know their future inertia parameters, or naive, and to think they will be non-inert. We use the difference in per-period subscription rates to estimate the actual inertia of the takers, and the share of sophisticates.

We find that in the population, about 30% are non-inert, 2% are fully-inert, and 68% are partially-inert with a 72% monthly chance of not cancelling a subscription they wish to cancel. We estimate that a large majority, 58%-67%, among the inert are sophisticated and know their inertial type. Being sophisticated means for the fully-inert that they will not subscribe to an auto-renew contract. Sophistication means that inert consumers only subscribe 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.

We also shed light on the mechanisms that cause inertia in our setting. We find the most support for inaction, and limited roles for switching costs or habit formation. Our findings cannot be explained by classic switching costs alone, regardless of whether consumers have perfect foresight about these costs (Klemperer, 1995), are completely myopic (Dubé et al., 2010), or due to stochastic switching costs: In contrast with our results, perfect foresight implies that auto-renewal subscribers should remain subscribed after the promo at similar rates to auto-cancel subscribers. Also in contrast with our results, myopia about switching costs implies no effect on initial take-up. Finally, stochastic costs and full attention imply a change in hazard rates when prices change, a variation we have in the data. Yet, we estimate a precise null response to those price increases. Finally, habit formation is ruled out by the finding that additional promo takers enticed by the price and duration treatment arms are not retained.

We conclude by studying in our setting the practicality of the common prescription of behavioral IO theory, calling for naivete-based discrimination (Heidhues and Kőszegi, 2010). We show that treatment effects are indeed heterogeneous in a predictable way. We predict consumer valuations of the subscription, out of sample, based on their baseline expected usage in the few weeks after hitting the paywall. Usage predicts promo and after-the-promo take-up, revenues, and subscription rates. Unlike most readers who never subscribe, the readers of highest predicted value actually appreciate the auto-renewal structure over auto-cancel.

We predict the firm's targeting strategy if it were to maximize either total revenues, subscriptions, or short-term paying subscriptions. Then we study the share of sophisticates under each of these strategies. We find overall small differences in sophisticates shares. Even when maximizing short-term paying subscribers, which should theoretically target naives the most, we find fewer naives are being assigned auto-renew than auto-cancel offers. These results highlight that with the targetable variables the newspaper and us have, discriminating on sophistication seems minor compared to targeting based on expected value.

We add three new findings—that consumers predict their inertia and push-back in the long-run, and the quantification of the type distribution—to a large literature (Brot-Goldberg et al., 2021; Choi et al., 2002; Della Vigna and Malmendier, 2006; Grubb and Osborne, 2015; Handel, 2013; Heiss et al., 2022; Ho et al., 2017; Hortaçsu et al., 2017; Kong et al., 2022; Madrian and Shea, 2001) that documents high degree of inertia among takers who appear to be naive about their tendency to procrastinate. We differ by considering consumers who are able to avoid the inertia inducing engagement altogether (here, contract). While we also find subscription takers to exhibit substantial inertia in our context, our study highlights the importance of considering the entire population of consumers who considered the contract in assessing the overall impact of inertia in the marketplace. For instance, if we follow the literature and compare the likelihood of a reader converting to a paid subscriber *conditional* on taking up the promo, we find the conversion rate to be 2000% higher for auto-renew takers relative to auto-cancel takers. However, accounting for all consumers, we see that there are actually fewer subscribers in auto-renew for any time horizon, and even the differences on the intensive margin are two orders of magnitude weaker. Via the lens of the model, we argue that while there are 40% non-inert consumers and that most inert consumers are sophisticated, focusing on the sample of long-run subscribers who took auto-renew leads to a sample consisted of almost entirely naive inerts. Meaning, counterfactuals that use estimates from takers are unlikely to generalize well to the population.

We also contribute to a much smaller literature that examines consumers' response to future inertia, and how it affects companies' decision making. For example, Reme et al. (2021) find that notifying existing subscribers of a mobile company about future plan changes leads to increased churn, even before prices change and even if their prices decrease. Meaning, some existing consumers are already inert and dormant, and the notification of future change draws their attention and potentially makes them aware that they might be inattentive again in the future. Rodemeier (2022) finds that consumers are aware of their lower likelihood of redeeming a rebate, focusing on short-term interaction between a retailer and its consumer base. Like these papers, we find that future inertia is a factor that consumers take into account, but we focus on assessing the overall role of inertia by analyzing the longer-term behavior and considering the type heterogeneity in the population (not just the takers) of consumers exposed to the contract. Further, our experiment is unique in eliciting consumer response to contracts that induce varying degrees of inertia. Indeed, the above papers find that for existing consumers exploitation of inertia is beneficial even if some of them are aware of it, while we find significant adverse consumer reactions to inertia inducing contracts. Finally, our paper also speaks to the conceptual way of incorporating inertia in models and empirical work. In the industrial organization tradition, 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)). In contrast, we find that a substantial share of potential subscribers are sophisticated about their future inertia. 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). Sophistication or partial sophistication regarding these forces may lead consumers to respond to future inertia. We do not distiniguish between every possible source of inertia, but as mentioned above, we find the most support for inattention (which we call inaction), 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 closely relates to the literature focused on firm marketing policies in contractual settings. Goettler and Clay (2011) show how learning and switching costs can interact to generate inertia in take-up of multi-part tariffs. Ascarza et al. (2016) show using a field experiment that a telecommunications company's proactive churn prevention initiatives backfire, possibly because such interventions reduce inertia, for example, by reminding readers of their low usage levels. Other papers focus on firm's personalization and targeting policies. For example, Yoganarasimhan et al. (2022) assess the effect of free-trial duration on customer acquisition using a field experiment and show that policies that maximize short-run also perform well in the long run. Datta et al. (2015) show that customers acquired by promo subscriptions have a lower lifetime value to the firm. Focusing on newspaper reader subscription discounts, Yang et al. (2020) show the predictability of long-term outcomes based on short-term outcomes. Our paper differs in that we explicitly vary inertia-related contractual terms and assess the degree of consumer sophistication.

Our findings are relevant 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.[2] 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 use a simple model to precisely define what we mean by inertia. An individual is inertial if being in a state, for example, being subscribed to a service, at period t causes them to be in the same state at time $t+1$, conditional on their preferences.

Two main mechanisms generate inertia. Firstly, inertia is as an outcome of a cost-benefit analysis that is driven by the costs incurred by the consumer (e.g., effort) for taking a state-changing action, versus the benefits of changing it. Here, state-dependence arises because past choices have lingering effects on the current costs or benefits. Some examples are switching and hassle costs which make it harder to change states; or conversely, habit formation that increases preferences toward an action previously taken reducing the desire to change states. Secondly, inertia is driven by naive inaction due to inattention or "autopilot"

[2]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 (see, for example, FTC May 2021, and Washington Post, June 2021 `https://www.consumer.ftc.gov/articles/getting-and-out-free-trials-auto-renewals-and-negative-option-subscriptions` and `https://www.washingtonpost.com/business/2021/06/02/automatic-renewals-ftc-subscriptions/`)

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 do the cost-benefit analysis.

A main empirical threat to showing and estimating inertia is preference heterogeneity (Dubé et al., 2010); we may observe a person continuing to stick to their past choices because they chose at t, and continue to choose at $t+1$, the option that is best for them. In other words, they would have chosen the same option at $t+1$ regardless of their t period choice and the persistence in choice simply reflects underlying preferences. This is what Heckman (1981) refers to as spurious state dependence. We do not refer to that as inertia.

In what follows we will entertain both sources of inertia and use our experiment to alleviate concerns about preference heterogeneity. First, we assume preferences are indeed heterogeneous and a main driver of take-up. Thanks to our randomization we have a comparable set of consumers who are exposed to different offers. Second, we assume the existence of *costs* for taking actions - to subscribe, to cancel, or to renew. Next, we model naive *inaction* as the probability of not taking an action at any given period by consumer i. This is a purely descriptive parameter, and not one that reflects underlying reasons for not taking an action. Namely, it 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.[3]

Finally, we assume that these parameters are fixed at the individual level, but allow for potentially incorrect beliefs about the future value of these parameters. We denote perceived parameters with $\tilde{}$.

2.2 Setting

We consider the following setup which closely resembles our empirical setting. Time is discrete: $t = 1, ..., \infty$. A customer faces a choice whether to subscribe to a service at period 1, and then whether to renew or cancel the subscription at later periods. At each period subscription is priced with non-decreasing prices $p_t \geq p_{t-1} \geq 0$ and at some period T the price becomes constant $p_t = p$ for $t \geq T$. We assume that each consumer i has some fixed per-period value from the subscription, denoted by v_i, which is drawn from an arbitrary distribution F.[4] In this setting there are three possible actions - subscribing initially, renewing, and unsubscribing - and which action is relevant depends on the state of the customer and the contract they are offered. 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 at the end of the period, incurs a

[3] One force we do not incorporate into the model is learning and habit formation. This is done for two reasons: Simplicity, and since our results suggest that this is inconsequential in our setting (see section 7.2).

[4] We assume that the value is known before subscribing for simplicity, and because absence habit formation, which we also assumed away and will later will find no support for, the assumption is inconsequential. We acknowledge that if there is no uncertainty of value nor habit formation the motives for promotions or penetration pricing are limited. Yet a company may offer a trial period as a way to introduce price discrimination via a temporary price reduction without affecting long-term prices.

cost c^r (e.g., clicking "renew" on an email or browser pop-up) which for simplicity, we assume is costless, i.e. $c^r = 0$.

A key feature of a contract is a single period z at which the contract cancels automatically. That is, if the consumer takes no action at period z, the contract will be terminated. If we set $z = \infty$ it means that the contract never cancels. Of course, the contract can be renewed at period z or afterwards.

To summarize, a consumer's per-period value of a subscription is $v_i - p_t$. We assume that time is discounted with a discount factor δ. Therefore, if a consumer plans to subscribe at period 1, cancel at period k (and z, the auto-cancellation period is either ∞ or $z < k$), and acknowledges some future inaction per-period probability $\tilde{\phi}_i$, then their *predicted* net present value is

$$\tilde{U}_i^k = -c^s + \sum_{t=1}^{k-1} \delta^{t-1} (v_i - p_t) + \delta^k \sum_{\tau=0}^{\infty} \left[\tilde{\phi}_i^\tau \delta^\tau \left(v_i - p_{\tau+k} - \left(1 - \tilde{\phi}_i\right) \tilde{c}^u \right) \right]$$

The first term is the cost of subscription; the sum to period $k-1$ is the net present value from the subscription; the final sum is the expected value due to a decision to unsubscribe at period k taking into account that the unsubscription takes effect in the next time period and there is a per-period probability $\tilde{\phi}_i$ that unsubscription will actually not happen.

Note the *til̃de* notations, indicating the perceived values of these future inaction probability and unsubscription cost rather than the actual ones. In contrast to the predicted value of unsubscribing at k, the actual valuation of the plan, if followed, is

$$U_i^k = -c^s + \sum_{t=1}^{k-1} \delta^{t-1} (v_i - p_t) + \delta^k \sum_{\tau=0}^{\infty} \left[\phi_i^\tau \delta^\tau \left(v_i - p_{\tau+k} - (1 - \phi_i) c^u \right) \right]$$

We will return to this model when estimating the distribution of types ϕ_i, and c_i^u. For now, notice that if a contract ends at a finite period z and that period coincides with when a consumer wants to cancel, inertia does not matter in the sense that a consumer will be happy to take the contract even if they anticipate to be highly inert. Also notice how changes in prices upfront might have differential response than changes at a non-z period, due to inaction and switching costs.

To summarize, we model the consumers subscription decision as being driven first and foremost by their value of subscribing and prices, while taking into account the perceived risk of remaining subscribed for longer than desired due to cancellation costs or innate inaction. Similarly, conditional on being subscribed, a consumer will cancel or remain subscribed based on their value, the price, and the actual unsubscription costs and the innate probability of inaction.

3 Empirical Setting

Our study was conducted in cooperation with a large European publisher that wishes to stay 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 in the United States or the Guardian in the United Kingdom. It publishes daily news in the main categories of politics, economics and business, sports, local news, culture, society, science, digital, working life, and travel. In addition to the printed newspaper, which started in 1945, the publisher has a digital platform which provides daily online news on its news website and mobile platforms. In 2018, approximately 12 million unique readers visited our publisher's digital platform.

The content on the digital platform is classified into three parts. One part 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 which are also available for free elsewhere. Another part of the content is "always paid", that is, it is available only to the platform's paid subscribers. This part includes high quality exclusive content from the printed newspaper and commentaries. The rest of the content is "metered" and subject to a metered paywall. Readers are allowed to consume 10 news articles per week for free and then hit the paywall where they are prompted to purchase a subscription in order to be able to continue reading the metered articles. The metered articles are specifically produced for the digital news channels and are generated by a dedicated digital editorial team. Traffic referred from online search platforms (e.g., Google or Bing) and social media platforms (e.g., Facebook or Twitter) receives no special treatment, that is, a reader referred by these platforms are subject to the same rules as any other.

Overall, such a 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).

In addition to subscription revenue, the publisher earns revenue from displaying ads to its readers. Paid subscribers generally see fewer ads (e.g., no performance ads) and are allowed to use their ad blocker, if they wish to do so. Non-paying readers see all ads and are not allowed to access the content using an ad blocker.

Tracking on the digital platform takes place via logins of registered readers and cookies, and is in line with the European General Data Protection Regulation (GDPR). A reader is assigned a cookie id once she hits the platform for the first time and is tracked on repeated visits as long as the cookie persists. Cookie-based tracking is not foolproof: A reader can decide at any time to to delete some or all cookies (i.e., active cookie deletion by clearing the cookies in her browser), and the same reader may have multiple cookies if they access the website from multiple devices (see Lin and Misra (2022)).

Pricing and Contracts The newspaper offers multiple subscription options to its readers. The most commonly bought contract is a daily pass, which provides reader access to paid content for one day for €2. The second most common are short-run (lasting up to one month) promotional contracts, such as our experimental contracts described below, which are offered to new readers who have never been paid subscribers before. Third are regular subscription contracts that continue for an unlimited time until explicitly terminated by the subscriber. The regular subscription prices are €19.99 for the first two months, and €34.99 per month thereafter. Additionally, the publisher has pre-committed (full lock in) one-year contracts, which are rare. A distribution of take-up is presented in Figure A.3.

Canceling subscriptions readers are notified of the subscription terms and conditions and of the technicalities of cancelling before they start their subscription. A subscriber can terminate their subscription at any time, which takes effect in the next billing cycle and the reader continues to have access until then. A reader can cancel their subscription by calling the publisher's call center, through the website using the "contact the publisher" page by entering their contract details, or by sending a cancellation letter by mail or email in response to the monthly invoice.[5]

4 Experimental Design

The field experiment was motivated by our research questions and the publisher's desire to convert most new readers into subscribers of the digital platform via randomized control trials. The experiment was conducted in three phases from April to August 2018, with follow-up data collected until April 2020. The experiment allows to document and quantify inertia and perceived inertia, and to learn about the drivers of inertia.

4.1 Participants and Randomization

Any "new" potential subscriber who hits the paywall either by exhausting their quota of free metered articles or by clicking on an always paid article enters the experiment. The reader is randomly assigned to one of eight experimental treatment groups outlined below, and receives the corresponding experimental subscription offer. The newspaper defines a new subscriber as someone who did not pay a full monthly price (€34.99) in the past.

Randomization is induced on the cookie-level and the assigned experimental group persists over time. A balance of the average number of pages visited before hitting the paywall by experimental group is shown in

[5] Overall, the modes of cancellation 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).

Appendix Table A.1. After the trial period, every reader, irrespective of the experimental assignment, has the option to pay €19.99 for the next two months, and the regular amount of €34.99 per month thereafter.

4.2 Experimental Contracts

Our experiment simultaneously varies three factors of the subscription offer. Each factor has two levels leading to a $2 \times 2 \times 2$ experimental design.

1. Subscription Renewal after the Promo: The first factor is the subscription renewal after the end of the promotional trial, which is either auto-renewal or auto-cancellation. A reader who takes an auto-renewal promo contract becomes a regular paid subscriber after the trial period is over, by default, unless the reader explicitly terminates the subscription. On the other hand, a reader who takes an auto-cancel offer does not become a paid subscriber by default. Instead, the reader can actively choose to resume the subscription access the next time she hits the paywall, through a pop-up on the platform's home page, or by clicking on a link in any one of several emails the platform sends the reader with an aim of reinstating the subscription.[6] In each of these methods, the reader verifies her already entered payment information and confirms the subscription contract.

2. Duration: The second factor is the duration of the experimental offer, which is either two weeks or four weeks.

3. Promotional Price: The third factor is price for the promo period, which is either €0.99 or €0. The price after the experimental offer is identical across individuals, and so is the set of contracts they can choose from.

The eight combinations of these factors and the corresponding experimental group name are displayed in Table 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 in experimental phases one and two. This was corrected in experimental phase three leading to a full orthogonal experimental design for that phase. We will consider this fact when discussing our results.

[Table 1 about here.]

[6] Approximately 5 days before the end of the trial offer, an email with a renewal prompt is sent to the reader, and a restart of the subscription can be initiated with a click on this email. If a reader does not respond to this email, she will be targeted in several follow-up emails as part of the standard process.

4.3 Taking up an experimental offer

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.[7] 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.

4.4 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. As the model makes clear, if perceived future inaction and unsubscription costs are small, there will be no differential take-up. Further, 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, which means that they also must have under-predicted their inertia. Comparing the auto-renew treatment effects on subscriptions at different time periods, and leveraging the price and duration treatments, informs us about the nature of inertia. Finally, the price and duration treatments allow us, under some assumptions, to quantify inertia, as we 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, giving us an observation window of at least 42 days of browsing history per cookie id.[8] 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.

[7] 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).

[8] 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.

5.1 Usage Data

The browsing history includes each page visited by a reader (identified by a cookie) along with its timestamp. Additional variables captured in the browsing history are the page type (open, metered, paywalled), and the subscriber identifier if the reader was logged in to their account (even if the subscription is free).[9] Moreover, another critical variable in the browsing history is 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 about here.]

From that data we learn a few things. First, we observe and define for each reader their first experimental exposure, and all subsequent exposures. For each reader we use the first exposure as their treatment group and define that date as day 0 of being in the experiment. The number of readers assigned to each treatment group is shown in Appendix Figure A.1. To keep an intent-to-treat design valid, we do two things. First, we measure time relative to that first *exposure* date rather than the actual take-up date (if one exists). Second, the reader is assigned to the promo offer that they saw first, regardless of the offers they subscribed to. For example, if a reader in a two-weeks promo treatment arm saw an offer on April 1st and took any offer (the experimental offer, a four-weeks offer, or any other subscription) on April 8th, the promo period for analysis purposes is 4/1-14 and not 4/8-21. In addition, we have information on readers' usage two weeks before first exposure and four weeks after. That data allow us to compare behavior across treatment arms, and of subscribers and non-subscribers. Finally, we use the data to consolidate multiple cookies associated with the same subscriber, and to consolidate multiple "subscribers" using the same cookie.

After these consolidations we are left with one line for each reader (cookie), which includes their first experimental exposure, and a unique subscriber identifier if they ever subscribed to the newspaper. We call that the *assignment data*. There are 2,088,910 readers in the experiment of which 16,237 had an active subscription during the experimental period (from two weeks before exposure until the end of data). Table 2 shows the number of readers, subscribers, pages visited, and total revenue. These are presented for the data transferred to us from the publisher which includes readers that did not get any promotional offers, the main data used for analysis (all participants in the experiment), and a sub-sample of participants who eventually subscribed to any of the experimental offers.

Note the common challenge in the digital world, which is that a person may appear with multiple cookies. However, we know something about the extent of the issue in our setting, and argue that it might shift the effects levels, but not in relative terms. Appendix Figure A.2 shows the distribution of the number of cookies

[9] Some readers become subscribers 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 usage time window after their exposure, we are unable to link that subscription to the reader.

associated with each subscriber. Once a subscriber logs in to their account, we link all the cookies to that subscriber. 63% of all subscribers are associated with only one cookie, and another 18% have two cookies associated with them. While some subscribers regularly clear their cookies, this is a small minority (less than 3% of subscribers have 10 or more cookies associated with them). Yet, the prevalence of multiple cookies per readers who subscribed suggests that non-subscribers will also show up in the data with multiple cookies and might be exposed to multiple treatments. Because we cannot defragment different cookies for the never-subscribed, this fact leads to inflation of the number of zeros across the treatment arms. For example, the same reader might have been exposed to several treatments accessing the newspaper from different devices. If they did not take any offer, they would appear as separate readers and will contribute "no subscription" and their usage to multiple treatment arms; if instead they did subscribe, then we associate all their devices ("readers") to the same subscriber with their first exposure determining "day 0" and accumulate all their usage from different cookies together. Therefore, fragmented never-subscribers may lead to attenuated subscription shares. However, they do not bias our results since we focus on treatment effects in relative terms.

5.2 Subscription Data

The second dataset is the publisher's customer relationship management (CRM) data which reports all signed contracts between April 2018 and April 2020 with their revenue, start date, and end date. Each contract is associated with a subscriber identified with a "contractor id" which is the subscriber identifier[10]. The main variable of interest beyond a contract's start and end time and collected revenue, is the contract code and description. Each of the contracts offered by the newspaper, including the 8 experimental contracts, has a unique code and description. We use these codes to see if readers took an experimental contract or others. Appendix Figure A.3 shows the distribution of contracts taken by the 16,339 experimental participants who subscribed at any point during the period (another 9,857 had a low-value subscription before the experiment and did not choose a new one over these 2 years). A contract is characterized by its maximal potential duration and revenue. The experimental contracts are highlighted with black boxes. As can be seen in the figure, there are many other different contracts being taken and offered. The abundance of possible products matters for the interpretation of the results. The main ITT results use as an outcome a subscription for *any* contract, experimental or not; For the quantification we add some explicit assumptions that allow us to focus on takers of experimental contracts only.

[10] 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.

5.3 Merging the Data Sets

Finally, we merge the datasets for analysis purposes. We merge the assignment data with the subscription data to construct at each day, relative to the exposure date, if a reader is subscribed and the average price they paid that day. We then aggregate the days to longer periods as we describe in the next section.

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. After showing the results, we will interpret what they imply for the existence and quantification of inertia and perceived inertia. In later sections, we analyze the take up of our experimental contracts.

Our main measures are subscription rate, that is, the proportion of days subscribed to the platform through any contract within a period; and the readers' subscription extensive margin, that is, if the reader was an active subscriber within that period at all. We also study the cumulative revenue accrued at different periods, and the numbers of visited pages.

For ease of interpretation, we divide our subscription data time span of two years into smaller time periods as follows. We use the two weeks before the promotional period as a placebo to test balance, the first two weeks of the promotional period (to make the 2 and four weeks promos comparable), and then move to a monthly frequency after the promo ends and until 24 months after-the-promo. We also consider a separate period of the entire time span past the promo end-date.

We set up the analysis described below in the form of the following regression

$$y_i = \alpha + \beta_1 \text{Auto-renewal}_i + \beta_2 \text{One-euro}_i + \beta_3 \text{Four-weeks}_i + \epsilon_i, \tag{1}$$

where y_i represents one of the outcome measures of reader i's subscription, and Auto-renewal_i, One-euro_i, Four-weeks_i are dummy indicators of i being assigned to an experimental group with auto-renewal (as opposed to auto-cancel), €0.99 (as opposed to free) and four weeks (as opposed to two weeks) contract terms, respectively. The β coefficients estimate the marginal effects of the experimental factors.

Recall that the experimental group G was incorrectly implemented in phases one and two of the experiment. So for the main analysis we exclude group G data for consistency across the three experimental phases, and verify that our results do not change when we separately analyze phase three data which has all eight groups. We use group G to train a regression forest to predict readers usage types purely out of sample in the final parts of the paper. Further, since the experimental assignment probabilities varied across experiment phases, we weigh each observation equal to the inverse of the assignment probability so that each experimental group receives the same overall weight. Our empirical results are not sensitive to this.

6.1 Auto-renewal vs. Auto-cancel

Figure 1 plots the intent-to-treat per-period effects of offering a promotional auto-renewal contract as opposed to an auto-cancel contract on subscription behavior at various time periods, which are the estimated coefficients β_1 in equation (1).

Figure 1a shows effects on subscription rates in absolute and relative terms. As expected, the subscription rates prior to the experiment are similar across the experimental groups, so the estimate in the first period is small and indistinguishable from 0. During the promotional time period, we observe a significant negative impact of auto-renewal on subscription rates of 0.1 percentage points[11], which is 28% lower relative to the auto-cancel subscription rate average. Meaning, there are 28% fewer subscription days during the promo period among those offered the auto-renewal versus the auto-cancellation offers. After the promo, however, the effect changes sign, and we see a positive effect of auto-renewal on subscription rate for a few months after the promotion. Subsequently, we observe a significant negative trend in the effect and eventually, about a year after the promo, the subscription rate is higher for the auto-cancel group and significantly so after 20 months. The effects on cumulative revenue show a similar pattern, as seen in Appendix Figure A.4. In the first few months after the promo ends cumulative revenue is higher by 20%, but the effect decreases. By eight months we cannot reject no effect on revenues, and after two years the point estimate is 1% higher for the auto-renew group and non-significant.

Comparing the intensive margin subscription rate patterns against those in Figure 1b, we note a different pattern on the extensive margin. We do see a similarly negative effect, -24%, of offering an auto-renewal contract on the likelihood of a reader becoming a subscriber at all in the promo period. However, we see no positive effect after the promo. Meaning, fewer readers become subscribers when they receive an auto-renewal offer relative to an auto-cancel in any time period. The increase in subscription rate is likely coming from those who remained subscribed due to them being inert. Overall, we see a significant 10% fewer subscribers over the entire two years after-the-promo because of offering an auto-renewing promotion.

[Figure 1 about here.]

[11] A 0.1pp lower subscription rate means that if there are 1000 potential subscription days during that period, readers offered a promotion are subscribed for one day less out of the thousand.

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 11% and causes 10% 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 over time; we do not observe any effect of the promotional price starting with the second month after the promotion on the extensive margin or the subscription rate. This implies that increasing subscription trial by decreasing price does not lead to long term subscriptions.

Four weeks vs. two weeks Figure 3 shows a similar pattern of the effect of increasing the trial duration. The estimates show that increasing the trial duration from two weeks to four weeks increases the subscription rate and the number of subscribers by 9%. However, similar to the effect of price, this difference also fades away over two months. Importantly, there is no effect from a longer trial on more subscribers in the long-run.

[Figure 2 about here.]

[Figure 3 about here.]

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, it causes an opposite effect on subscription rates a few months after the promo. Even same sign effects are short-lived in response to benefits such as a price reduction or trial duration extension.

At the same time, we see an overall decrease in after the promo subscribers due to auto-renewal, which is also absent in response to to the other treatments. These patterns indicate a unique consumer 'push back' to auto-renewal relative to factors such as price and trial duration.

7 Channels of Inertia

The experiment have shown several patterns in the data. Namely, that there is a lower take-up of the auto-renewal offers, implying that consumers are aware and respond to future inertia; that even with lower initial take-up, auto-renewal subscribers stay subscribed for longer and pay a full price in the first months, implying inertia is large enough to over-come the initial reduction in take-up; and lastly, that there are fewer subscribers to any contract in the auto-renewal treatments in the two years after-the-promo, implying that inertia's effects are not long-lived.

In this section we provide further evidence on the channels of inertia, to support our modeling assumptions and estimation. First and foremost, we show that auto-renewal subscribers are not using their subscription. Further, we show that treatment arms that brought more readers to try the subscription did not lead to more long-run subscriptions. Together, these findings show that even if the valuation of the subscription is ex-ante uncertain for the reader there is no evidence for learning or habit formation in our context. i.e., assuming constant per-period subscription value is a reasonable assumption. Next, we provide evidence that the reduction in subscribers in the long-run is coming from a reduction in valuation of the newspaper due to the offer of an auto-renewing contract.

7.1 Subscription versus Usage

If the increased subscription behavior caused by auto-renewal is actually unwanted, for example if caused by inaction or cancellation costs, then we expect readers to get little utility from their after-the-promo subscription. We use the website usage data to gauge the utility subscribers receive through reading the news articles and empirically assess this explanation.

Using the website usage click-stream data, we estimate the average daily number of pages visited by readers who took an experimental auto-renewal or auto-cancel offer. We then compare the trends in subscriptions with those of actual website usage. If auto-renewal takers receive utility from keeping their subscription, we expect their subscription usage, as evident by visits to paywalled pages, to be larger than the usage of non-subscribers.

Recall that our usage data spans six weeks for each reader; notating day zero as the day a reader received the experimental offer, our usage data spans days -13 to 27. Figure 4 plots the average page visits (bars) and subscription rates (dots) among promo takers for each day in this time span. For this plot, we use 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 usage data. Note that this is a different sample than the results above, as we condition on takers rather than include all those exposed to the treatment.

Figure 4 shows that auto-renewal promo takers are orders of magnitude more likely to be subscribed after the two weeks promo time, relative to auto-cancel promo-takers who overwhelmingly do not continue their plan nor take on new subscriptions. However, we do not see any difference in the unconditional groups' website visits in the last two weeks, after the promo period ends. If they were to remain subscribed because they use the subscription to access paywalled articles, the average usage should have correspondingly be orders of magnitude higher than the auto-cancel group. This indicates that the auto-renewal takers who continue to subscribe do not visit the website more often. Compared to pre-treatment days, we see that both groups use the website more, after the promo take up.

Table 3 shows promo and after-the-promo usage statistics averaged across readers who took either a two weeks €0.99 auto-renewal or auto-cancel promo. The sample is grouped by whether the reader was also a subscriber after the promo, or not. i.e., this is conditional on being subscribed in a period. The analysis shows that more than half of the readers who subscribed in the two weeks after the promo after taking an auto-renewal promo did not even visit the newspaper's portal. This proportion is similar to those that did not subscribe after the promo and is significantly lower than those who subscribed after the promo after taking an auto-cancel offer.

Overall, this analysis is consistent with our inference that the readers who continue subscribing after taking an auto-renewal promo do not derive higher utility than those who do not. Meaning, the valuation of subscription does not grow for auto-renewal subscribers thanks to their subscription status.

[Figure 4 about here.]

[Table 3 about here.]

7.2 Lack of Learning and Habit Formation from the Promo Trial

Habit formation in our setting is the idea that once someone subscribes to the newspaper they would learn to like it, their utility from it increases, and therefore we will observe inertial behavior. Namely, that subscribing causes long-term subscription. Are higher subscription rates in the auto-renew treatment arms due to habit formation? We can address this channel by examining the effects of the free and long duration promos.

Indeed, reducing the promotional price and increasing the trial duration increases the readers' initial subscription rates, leading more people to try the product. However, this increase in take-up does not significantly change their future likelihood of subscribing to the platform, as indicated by Figures 2 and 3. This finding indicates that the learning from trial experience is not significant enough to change longer term subscription behavior. Meaning, there is *no evident habit formation*, in the sense of an increase in the benefits from being a subscriber, as a driver of inertia in our setting.

The absence of learning and habit formation and the evidence above on the lack of usage for auto-renewal subscribers, means that readers valuation of the service does not change due to usage. Meaning, while there might be resolution of uncertainty thanks to subscription, there is no persistent increase in its value over time.

7.3 Evidence for "Spite"

Who are the readers who subscribe to auto-renew contracts versus auto-cancel ones? We define a reader type with a proxy for their consumption value in order to understand the heterogeneity behind subscription behavior for different contract types. Average subscriber types at different periods speak to the mechanism driving our estimated effect of serving an auto-renewal vs auto-cancel contract. Under our model, readers with higher valuation of the product are more likely to subscribe. Therefore, we expect the average type of auto-renewal subscribers to be *higher* than auto-cancellation subscribers during the promo period because the marginal reader that does not take an auto-renewal will take an auto-cancellation promo. However, after the promotional offer ends, auto-cancellation subscribers are those who actively subscribe while the large share of auto-renewal subscribers remain due to inertia. Meaning, after the promotional period, the average type of auto-renewal subscribers should be *lower* than for auto-cancellation subscribers. Over time, irrespective of whether readers are being driven by switching costs (standard or coupled with present-bias) or inaction, we should expect the types to converge from below in the long-run.

We cannot reject the former two predictions but can reject the latter. We train a regression forest on pre-experimental usage data to predict post-treatment usage for readers in the mis-implemented G group. We then predict out-of-sample the expected usage for each reader in the other experimental groups in our data. We use this predicted usage as a proxy for the reader's type – those who are predicted to use the newspaper more are higher types. Figure 5 shows the by-period difference in reader types between auto-renewal and auto-cancellation subscribers at these periods. We cannot reject that the promo period subscribers in the auto-renewal group are of higher types relative to the promo period subscribers in the auto-cancel group. We find statistically significant support that in the initial periods after the promo ends, lower types subscribe in the auto-renewal group. But, the difference then flips sign and becomes significantly positive in the long-run. Meaning, the average predicted value of long-run auto-renew subscribers is significantly higher than that of long-run auto-cancel subscribers.

To dig in deeper, Figure 6 shows the full predicted usage distributions of all readers and of subscribers during three specific periods. The top-left panel shows the distribution of predicted usage of all readers in the experiment for comparison. Moving to subscribers, during the promo period (top-right panel) the distributions are similar between auto-cancel and auto-renew, with many low predicted usage types subscribing. In particular, and as expected, the auto-cancel subscribers skew more toward the low predicted usage types. In contrast, during the first month after the promo period (bottom-left panel), the auto-renew subscribers are more likely to be lower types compared to the auto-cancel which are skewed to high types. Finally, two years out (bottom-right panel), auto-renew types density is shifted up. The density of types in auto-renew is lower for low types and higher for the very high types.

These findings imply that, in the long-run, there is an implicit penalty induced by the auto-renew 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. This implies that auto-renewal deters even those who wish to remain subscribed. 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. This finding is consistent with the extensive margin result of fewer subscribers after the promo period (shown in section 6), and it further suggests that some of these missing subscribers are high value subscribers.

[Figure 5 about here.]

[Figure 6 about here.]

8 Quantification of Inertia

Our results show that inertia exists and is predicted and avoided by some readers. Inertia's existence is manifested by higher retention of auto-renew takers despite the lower initial take-up. In turn, the lower take-up of auto-renew offers demonstrates future inertia being predicted. In this section we turn to quantify the degrees of experienced inertia, of predicted inertia, and their heterogeneity. We do so guided by the model with minimal added assumptions.

As a reminder, our model includes two terms that may generate inertia: probability of inaction, and cost of unsubscription. For an *existing*-subscriber, both forces prevent one from canceling on time; for the *potential*-subscriber, their expected magnitude will determine if they will subscribe or not.

8.1 Experienced Inertia and Heterogeneity

In this section, we quantify the degree and heterogeneity of actual inertia experienced and exhibited by readers who take up an auto-renewal subscription. We look at the retention of auto-renewal promo takers in the after the promo period. We use the comparable set of auto-cancel takers and renewers that tells us, under a monotonicity assumption and assuming negligible renewal costs, how many readers would have been subscribed if not for inertia. However, the set of readers who take an auto-renewal promo is different than those who take the auto-cancel promo, and in what follows we describe how we account for that selection.

To fix notation, let R designate the set of types (combination of valuations and predicted inertia) who take the auto-renewal promo, and C the set of types who take the auto-cancellation promo. We assume that $R \subset C$, that is, our main assumption is monotonicity — individuals who take up the auto-renewal promo would also take the auto-cancel promo, everything else held constant. We justify this assumption with the idea that an auto-cancel contract offers access to the same content without the risk of an unwanted paid subscription.[12] The evidence in section 7, showing that auto-renew takers do not use their subscription, and the lack of learning and habit formation (subscribing does not change the value of the service), support this assumption – those who remain subscribed are indeed left at an *unwanted* subscription. Meaning, auto-renewal takers who are predicted to not take the full price contract under auto-cancel all strive to cancel their subscription if it were not for inertia. We expect them to try and cancel at every period unless they are inert.

To quantify inertia we assume that renewing the contract conditional on subscribing in the promotional period is determined by the model. That is, it is driven by a trade-off between the net-present-value of remaining subscribed and the cost of unsubscribing. Prices are known, and we assume homogeneous subscription cost and time discount factor.

We assume readers are heterogeneous in their valuations v of the subscription, and $v \sim F$, which may be a distribution of any shape. Further, we assume consumers differ along two dimensions when it comes to inertia: in their experienced inertia (parameters related to inaction and costs), and in their sophistication about it. We assume there are 3 discrete inertia types and 2 sophistication types, and that type-dimensions are independent of one another. The actual inertia types are: fully-inert, partially-inert, or non-inert. The sophistication types are: completely naive (believing they will not be inert in the future), or correctly calibrated (knowing their individual level of inaction and unsubscription cost).[13]

In this section, we quantify the distribution of the actual inertia types, as well as the parameter of inaction for the partially-inert. The sophistication types distribution will be estimated in the next subsection.

Estimation We are interested in the causal effect of auto-renewal on the likelihood of being subscribed after the promo. This is the difference in after-the-promo subscription of those who take an auto-renew contract when offered it (our group R), versus, if they were instead offered an auto-cancel contract. To estimate the change in subscription behavior of readers in R when the contract changes from auto-renew to auto-cancel, we need to (1) estimate their after the promo subscription when they are offered an auto-

[12]As discussed in section 4, readers in our sample chose not to try the subscription at regular price so their prior expected value from the subscription is low. This supports our assumption that they would prefer an auto-cancel contract to auto-renewal which might enroll them into paying for the subscription.

[13]For the non-inert, sophistication or naivete is inconsequential. Hence, it might be simpler to think of all non-inert consumers as sophisticated, and the naivete only applies to the partial- and fully-inert.

renewal contract ($y_R^{AR} = \mathbb{E}[y_i|t = AR, i \in R]$) and (2) estimate the same when they are offered an auto-cancel contract ($y_R^{AC} = \mathbb{E}[y_i|t = AC, i \in R]$).

Estimating (1) from the data is straightforward because the auto-renewal promo takers identify the set R and the share of them renewing the contract is the object of interest. We can estimate (1) with the sample equivalent

$$\hat{y}_R^{AR} = \frac{1}{|R|}\sum_{i\in R}\mathbf{1}(y_i = 1|t = AR) = \frac{\sum_i \mathbf{1}(y_i = 1|t = AR, promo = 1)}{\sum_i \mathbf{1}(promo = 1|t = AR)}$$

where we look at the after-the-promo subscription y of the readers who took the auto-renewing contract when assigned to it. In contrast, the observed average subscription behavior of auto-cancel promo takers, C, represents a combination of types from R and from $C \setminus R$. Thus, we do not have a sample equivalent of (2), but we do have $y_C^{AC} = \mathbb{E}[y_i|t = AC, i \in C]$.

Under the assumption of independence of naivete and inertia, the marginal types who take an auto-cancel promo and an auto-renew promo have the same valuation – it will be a fully naive or non-inert reader in the auto-renew arm. However, as will be clearer in the next section quantifying predicted inertia, the same marginal taker type does not mean that the distributions of taker types are identical in inertia. In fact, takers will be on average less inert than non-takers. The reason is that auto-renew non-takers do so due to low valuation combined with predicted inertia. It is the risk of paying for a subscription they do not want, or paying the cancellation costs of that subscription that keeps them from taking the promo. Therefore, those in $C \setminus R$ will not be renewing their subscription under auto-cancel – $\sum_{i\in C\setminus R}\mathbf{1}(y_i = 1|t = AC) = 0$. Therefore,

$$\begin{aligned}
\hat{y}_C^{AC} &= \frac{1}{|C|}\sum_{i\in C}\mathbf{1}(y_i = 1|t = AC) \\
&= \frac{1}{|C|}\left(\sum_{i\in R}\mathbf{1}(y_i = 1|t = AC) + \sum_{i\in C\setminus R}\mathbf{1}(y_i = 1|t = AC)\right) \\
&= \frac{1}{|C|}\left(\sum_{i\in R}\mathbf{1}(y_i = 1|t = AC)\right) \\
&= \frac{1}{|C|}\left(|R|\cdot \hat{y}_R^{AC}\right) \\
\implies \hat{y}_C^{AC} \times \frac{|C|}{|R|} &= \hat{y}_R^{AC}
\end{aligned}$$

Hence, the excess share at each after-the-promo period (time subscript omitted) is

$$s = \hat{y}_R^{AR} - \hat{y}_R^{AC} = \hat{y}_R^{AR} - \hat{y}_C^{AC} \times \frac{|C|}{|R|}.$$

The RHS of this equation is estimable with the appropriate sample equivalents: the per period share of auto-renew subscribers minus the share of auto-cancel subscribers times the ratio of auto-cancel promo takers to auto-renew promo takers.

Recall that all those shares are the excess subscribers due to inertia, driven by either unsubscription costs or inaction. Let s_t designate the excess share at month t, with 0 being the promo period, and 1 being the first month after the promo. Recall that we have three inertia types – fully inert, partially inert, and non-inert. By definition, the non-inert will not be part of the excess mass, so they will be part of s_0 but will drop out by s_1. Starting at s_1, if the cost of unsubscription is not too high, the partially inert will try to unsubscribe and will succeed at rate $1-\phi$. The fully inert, due to very high costs or due to full inaction, will remain subscribed. Further, we might think that going from month 2 to 3, as the monthly price increases by €15, a mass of inert consumers driven by cancellation costs will leave. Therefore, we have the following series:

$$s_{t+1} = \phi \cdot \left(s_t - \pi_f^R\right) + \pi_f^R + \alpha^c \cdot \mathbf{1}\left\{t=2\right\} \tag{2}$$

where ϕ is the parameter of inaction of the partially-inert, π_f^R is the share of fully inert subscribers among the auto-renew promo takers, and α^c is the excess share of auto-renew promo takers who leave on month two due to the increase in the monthly price. We therefore regress s_{t+1} on s_t and a dummy for month 2, and get that the coefficient on s_t is an estimate of ϕ, the intercept equals $(1+\phi)\pi_f^R$, and the dummy estimates the excess share leaving on month three. We extrapolate to obtain the share of partially inert and fully inert consumers at s_0 and the difference between the actual share of promo-takers and the predicted share gives us an estimate of the share of non-inert consumers among auto-renew takers π_n^R.[14] Thus, we have estimates of the shares of all types, the share leaving at month 3, and the average inertia of the partially inert.

Figure 7 shows the moments of excess subscribers and the results of the estimation. The dark circles are the estimated excess share of subscribers in each month after-the-promo, and the triangles are the predicted shares from the estimation. These are only 12 dots, but notice that this three parameters process fits the data very well. The first triangle at *Promo* is the projected share of the partial and fully-inert subscribers among excess auto-renew takers. The difference from that share and 1 is the share of the non-inert. The triangle at ∞ is the estimated share of fully inert auto-renew takers $\hat{\pi}_f^R$.

[Figure 7 about here.]

We estimate that $\hat{\pi}_n^R = 50.6\%$ ($se = 2.0\%$) of the auto-renew promo takers are non-inert – they take the promo offer and unsubscribe before paying. We interpret them as having low inaction $\phi \approx 0$ and very low

[14] $\hat{\pi}_n^R = s_0 - \hat{s_0} = s_0 - \left(\frac{1}{\hat{\phi}} s_1 - \frac{1-\hat{\phi}}{\hat{\phi}} \hat{\pi}_f^R\right)$

unsubscription costs. In contrast, there are $\hat{\pi}_f^R = 1.3\%$ (0.9%) fully-inert consumers with either $\phi = 1$ or very high costs. Finally, the remaining $\hat{\pi}_p^R = 48.2\%$ (1.8%) are partially inert, with an estimated inaction of $\phi = 0.718$ (0.020) and very low costs.

We argue that the partially inert have very low unsubscription costs for the following reason. At month three, the monthly price increases from €19.99 to €34.99. To the extent that unsubscription costs are what driving the partially-inert, we expect to see a distinctive drop in retention. Yet, when we estimate the excess share of subscribers who leave before the second price increase we find that only 1.7% (se = 0.8%, which are 3.4% of the partially-inert) leave due to the €15 increase. Therefore, we conclude that for the partially-inert, unsubscription costs are not a major component behind their decisions.[15]

We do a similar exercise in Appendix A.3 where instead of comparing auto-renew takers to auto-cancel takers, we compare the excess subscribers between auto-renew takers who join due to the longer promo duration auto-renew to the shorter promo, and the free versus €0.99 promo auto-renew. This exercise lets us look at a different population, but we lack power and while we find qualitatively similar patterns, the estimates are too noisy to be informative.

8.2 Predicted Inertia

We now turn to calculate the predicted inertia of the different types. As mentioned above, we assume that independently of the experienced inertia type, each partially or fully-inert consumer is either naive or sophisticated. Therefore, their beliefs about the parameters of inertia (inaction and costs), are that those are either zero or perfectly calibrated, respectively. We have the estimates of the inertia-types shares among those who take auto-renewal promo. Hence, what is left to estimate is the share of sophisticates, and to recalculate the types' shares in the population. The recalibration is needed, because of differential selection into taking the auto-renewal promo given each reader's valuation and *predicted* inertia type.

Table 4 describes the 5 types and their predicted behavior. Each column is a different experienced inertia type, and the rows are whether they are sophisticated or naive. For example, the bottom-right cell are the sophisticated fully-inert – they are fully inert and they know it.

[Table 4 about here.]

Within each cell, we describe the auto-renew promo take-up behavior based on the type, compared to

[15] If instead we assume that costs are stochastic and there is no inaction at all, we expect the share of cancellers at month one to drop distinctly. If costs are distributed according to some distribution G. The difference in retention is $G(\bar{c} + e15) - G(\bar{c})$. Where $\bar{c}$ is a steady-state cutoff cost under which a subscriber cancels. In this model without inaction and only stochastic costs $G(\bar{c}) = 0.72$. To get a sense of what $\bar{c}$ might be, if we assume $G \sim U[0, k]$ then $\bar{c}$ is on the range of 3-30 (depending on the value and time discounting factor which are unidentified). Yet, $G(\bar{C} + e15) - G(\bar{C}) \approx 0$, meaning that the distribution is flat close to a region where it is 0.72. Therefore, any standard distribution should change significantly from such a shift in its argument. But there is no drop, implying the stochastic costs are inconsistent with the data.

auto-cancel. The take-up behavior of auto-renew promo is the same as auto-cancel promos for the naives and for the non-inert. For these types, rightfully or not, the contracts seem equivalent because they predict cancellation will be effectively costless and frictionless. Meaning, there are only two types who will not take the auto-renew promo but will take the auto-cancel contract – these are the sophisticated fully-inert, and the sophisticated partially-inert. The sophisticated fully-inert know that if they take the promo they will convert to become paying subscribers regardless of their valuation. Therefore, they will only take the promo if they also have high-enough valuation, equivalent to the value of the long-run renewing subscribers of the auto-cancel.[16] The remaining type is the sophisticated partially-inert consumers. Consumers of this type know they might be subscribed for long even if not forever. Therefore, their net-present-value takes into account the higher price periods of possible subscription.

We use the model to pin down selection into auto-renewal take-up. In the auto-renew condition for each inertial type i, there is a marginal valuation $v(\tilde{\phi}_i, \tilde{c}_i^u)$ that satisfies the following condition such that if the valuation is higher than that v they will take the auto-renew promo:

$$0 = v - c^s + \sum_{\tau=1}^{\infty} \left[\tilde{\phi}_i^\tau \delta^\tau \left(v - p_\tau - \left(1 - \tilde{\phi}_i\right) \tilde{c}_i^u \right) \right] \tag{3}$$

where v is the valuation of the per-period subscription, c^s is the subscription hassle cost, $\tilde{\phi}_i$ is type i's predicted inaction, $\tilde{c}_i^u$ is the predicted unsubscription cost, δ is the time discount factor, and p_τ are the per-period (month) prices which are known.[17] From section 8.1 we know the predicted inertia parameters for the partially-inert. So we are left with two unknowns – δ and c^s.

To get the subscription cost notice that for auto-cancel, the summation term in the RHS of equation 3 drops out, and we are left with $0 = v - c^s$. Or in other words, the value of the marginal auto-cancel taker is exactly the subscription cost. We identify that marginal value with the price and duration treatment arms. The extensive margin effect during the promo period of β_2 in equation (1) shows that a €0.99 difference in price leads to 0.0479 percentage points fewer subscribers. Meaning, $F(v + €0.99) - F(v) = 0.0479pp$. Taking a first order approximation, we get that $€0.99 \cdot F'(v) = 0.0479$. Similarly, we compare the estimates of the extensive margin effects during the promo period of β_2 and β_3 in equation (1). Taking a first order approximation for both suggests that for those on the margin of subscribing, the average value of additional two weeks of subscription is equivalent to $0.84 \cdot €0.99$. Hence, we estimate the value of additional two weeks of subscription, for those on the margin, is €0.83. Meaning, that the cost of subscription is also $c^s =€\,0.83$.

[16] To be precise, the long-run auto-cancel subscribers have higher value. The marginal type there has a positive net present value starting at the full-price period, while for the fully inert auto-renew promo, the net present value includes the lower priced promo period.

[17] We could have added a utility shifter for those offered an auto-renew contract to capture the long-run extensive margin reduction in take-up. However, that will make the model unidentified without further restrictions, so we abstract away from that for now.

We solve the above indifference condition in 3 and find that for the partially-inert with $\tilde{\phi} = 0.72$ and $\tilde{c}^u = 0$, v ranges between €16.9-17.4 for δ between $0.98 - 0.999$. Meaning, the marginal type has a valuation that is close to the first two months' price of 19.99.

Therefore, the difference in take-up of the promotional offers between auto-renew and auto-cancel is coming from these two sub-populations: the sophisticated partial naifs with value lower than 17.4 (and higher than 0.83, i.e. high enough to take the auto-cancel), and the sophisticated fully naifs who would otherwise take the subscription.

We can then compare the share reduction in take-up of auto-renew promo versus auto-cancel promo, and that share will equal the sum of these above types, weighted by their shares. Note that their shares in the population is not the same as the shares estimated among takers, exactly because of the selection here. For example, the fully-inert among the auto-renew takers are only the naive fully-inerts. Overall, the total share of auto-renew takers in the data equals all of the naives with value above 0.83, plus the sophisticated non-inert, the partially inert with value above 17.4, and the sophisticated with value above 34.35 (long-run subscribers). We designate π^s of consumers as sophisticated, and $1 - \pi^s$ are naive, independently of their actual inertia. Designate the shares of fully-inert, partially-inert, and non-inert types in the population as π_f, π_p, π_n respectively. Then, the share of auto-renew takers who are not long-run subscribers under auto-cancel is:

$$\begin{aligned}
&\pi_n \cdot (F(34.35) - F(0.83)) + \\
&\pi_f \cdot (1 - \pi^s) \cdot (F(34.35) - F(0.83)) + \\
&\pi_p \left[(1 - \pi^s) \cdot (F(34.35) - F(0.83)) + \pi^s \cdot (F(34.35) - F(17.4))\right]
\end{aligned} \tag{4}$$

We estimate the shares with the sample equivalents - s^{AC}, the share who take the auto-cancel promo gives us an estimate of $1 - F(0.83)$; s^{AR}, the share who take the auto-renew promo; and $s^{AC,long-run}$, the share of long-run subscribers in auto-cancel gives us $F(34.35)$. Finally, the expression in 4 equals $s^{AR} - s^{AC,long-run}$.

To get the shares of types in the population we rely on their selection into their estimated shares among auto-renew takers. Namely, that $\pi_n^R = \pi_n \cdot \frac{s^{AC} - s^{LR}}{s^{AR} - s^{AC,long-run}}$ (the share of non-inert among excess auto-renew takers is their share in the population times those who take it but not renew, divided by the total number of takers) and similarly $\pi_f^R = \pi_f \cdot (1 - \pi^s) \frac{s^{AC} - s^{LR}}{s^{AR} - s^{AC,long-run}}$.

The remaining unknown parameter is $F(17.4)$, the share of consumers below the cutoff valuation of sophisticated inertia taking. We do not have a way to pin this parameter down, so we take two approaches. We write down $F(17.4)$ as $F(17.4) = F(34.35) - a \cdot (F(34.35) - F(0.83))$. Meaning, $a \in (0, 1)$ is a measure of how close to $F(34.35)$ is that share. If we fit a normal distribution to the moments of the CDF for which we observe their sample equivalents – $F(0.83)$, $F(34.35)$, and $F(1.82)$[18] we get $a = 0.008$. To gain some intuition, we observe the tail of the distribution and therefore the CDF in this region is highly concave. However, we also vary a to take values between $\{0.001, 0.01, 0.05, 0.1, 0.15\}$ and test the sensitivity to these choices.

Finally, to get standard errors, we bootstrap the entire procedure of estimating the inertia-type shares and the experienced inertia for the partially inert among takers, and then propagate these estimates to the identification of shares in the population and sophistication given equation 4. We use Bayesian bootstrapping, drawing random weights for the entire sample 1000 times, and recalculating the parameters, and then use the standard deviation of the 1000 estimates to provide standard errors.

The results for $a = 0.01$ are shown in Table 5. Our first finding is that the share of sophisticates is substantial. We find that 58.2% (se = 1.9%) of inert readers are sophisticated and aware of their inertia parameters. This awareness leads, by assumption, to strong selection if one is to look at takers only. Among auto-renew promo takers, the share of sophisticates drops to 6% (se = 1%).

There is also substantial heterogeneity of inertial types. In the population, roughly 30% (1.3%) are non-inert, 1.8% (1.3%) are fully inert, and 68% (1.5%) are partially inert with average inaction of 0.72 (0.02).[19] The shares among auto-renew takers were presented in section 8.1. Note, that by definition the apart from a small share (about 6%, equivalent to the share of auto-cancel subscribers who extend their contract), auto-renewers who are still subscribed after the promo period ends are naive (as highlighted in the previous paragraph), and inert.

To put the results in context notice the following implication: In the population we have about 30% non-inerts, 40% sophisticated inerts and 30% naive inerts; however, if we condition on takers who pay full price we get a selected sample of almost 95% naive inerts. The reason is that the non-inerts take the promo but do not remain subscribed, and most of the sophisticates avoid taking the contract altogether. This result explains why most of the literature, which focuses on takers, finds very strong support for highly naive highly inert population. Yet, it also highlights why these samples are highly selected on sophistication and inertia, and therefore will not generalize well to the population as a whole.

[Table 5 about here.]

[18] the latter one is the shares of takers of the promo with the 0.99 price compared to the free one.

[19] If we vary a, our non-identified parameter, the estimates among takers are not affected, and neither are the shares of non-inert (since their share in the population is not a function of a), but we do get some variation in the share of sophisticates, and minor variation in the shares of partially-inert in the population (and the complementary fully-inert), shown in Appendix Figure A.6. The shares of sophisticates among the inerts varies between 57.7%-67.4% and the share of sophisticated inerts among the promo takers is the most sensitive and varies between 4.8%-27.4%).

8.3 Targetability Based on Sophistication

Is it possible to target readers with offers based on their sophistication? Is reader sophistication predictable? We use our pre-experimental usage data—which includes the timing, page-views, and topics the readers browse—to predict the heterogeneity in the effects of our treatments.

8.3.1 Heterogeneity Based on Readers' Valuation of the Subscription

We examine how the effect of giving an auto-renewal vs auto-cancel contract varies with individual's valuation for the paid subscription. As a proxy for a reader's valuation of the subscription (v_i) we use out-of-sample usage. We run a regression forest on the omitted group, test-group "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. These data are the same as the newspaper's first party data, which makes the exercise business relevant. Then, we predict out-of-sample on the other test groups and assign each reader their predicted usage score, which is their predicted number of page impressions.

The pre-expreiment usage consists of the number of pages impressions by number of days before hitting the paywall (five or more before, four, three, two, one, and all impressions on the day-of until hitting the paywall and entering the experiment), category (e.g., homepage, sports, culture, politics), and page type (is it open, metered, or always paywalled). In addition, we use the total pages impressions by day and page type. This way we construct 54 variables for every reader.

We validate this measure 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 8. 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, are more likely to subscribe, and subscribe for longer. Each point in the figure is one percent of readers, showing that the most predictably avid readers are those for whom the take-up is higher.

[Figure 8 about here.]

Given the skewed predicted usage, and since the overall take-up during our promo period is about 0.4%, we classify readers as "high type" if they are in the top 0.4% of predicted usage. Next, we find that readers who are predicted to be of the highest value exhibit different pattern of treatment effects. They are more likely to subscribe, and subscribe for longer, when offered auto-renewal vs auto-cancel. Figure 9 shows the auto-renewal vs. auto-cancel treatment effects interacted with the usage-type categorization.[20] The auto-renew offer treatment effects on the per-period subscription rate for the majority of readers are, unsurprisingly, the same as the main results – 29% lower during the promo period, then about 20% higher initially after, and becoming 12% lower at the end of the period. In stark contrast, for the highest value types, the promo treatment effect is insignificant and positive, and then becoming persistently and increasingly positive, from 20% higher subscription rate immediately after the promo ends, to about 40% two years later. The effects are similar for any subscriptions (including non-experimental) and cumulative revenue (Appendix Figures A.8 and A.9, respectively). Meaning, the treatment effect of auto-renewal on high value readers is positive — they might even appreciate the value that an automatically on-going subscription provides. This is not surprising, as the logic behind the subscription business model may not necessarily be exploitation, but rather convenience for those who are indeed willing to pay.

[Figure 9 about here.]

8.3.2 Who Gets Targeted?

After establishing that substantial heterogeneity exists, we ask what does that imply for the ability to target specific sophistication types — theoretical work implies that the firm may wish to discriminate based on naivete, is that feasible? Can the firm identify sophistication types?

We run several causal forests to estimate heterogeneous treatment effects of giving auto-renewal vs auto-cancel offers. As above, we use the same pre-experiment browsing behavior as covariates that feed into the causal forest. We estimate auto-renewal treatment effect on three outcome variables the firm might focus on (1) total revenue, (2) the probability of subscribing at all after the promo ends, and (3) the probability of being a subscriber on the first month after the promo ends. Meaning, we are asking how well will the firm target sophisticates if it is maximizing revenue, long-term subscriptions, or initial inertia exploitation.

For each outcome, the heterogeneous treatment effects may be positive for some readers and negative for others. We "assign" readers to auto-renew if the effect is positive, creating a sub-sample of readers who should have been targeted with auto-renew offer to maximize the outcome; and similarly to auto-cancel, creating a sub-sample of those that should have been assigned to auto-cancel.

Within each sub-sample, however, there was random assignment to auto-renew versus auto-cancel. Therefore, we can re-estimate the types classification as in Section 8.1.

The results from this analysis are quite insightful and shown in Table 6. Maximizing revenue (column 1) or first month after-the-promo subscriptions (column 3) leads to offering most readers an auto-renew

[20] Because the magnitude of the effects in percentage points is so different, we are now plotting the per-period effects in relative terms (i.e., within type auto-renew relative to auto-cancel baseline). The absolute value ITT effects are in Appendix Figure A.7.

contract (78% and 72% respectively). Those offered auto-renew are also somewhat more inert. However, the strongest differentiator is the predicted value, where those offered auto-renew have much lower predicted value on average. It is curious that the numbers are similar if trying to maximize first month after-the-promo subscription rather than revenue. First month after the promo subscription is what we thought as the best proxy for inertia and naivete exploitation, yet the share of sophisticates is not distinguishable and if anything slightly higher among those offered auto-cancel. Notably however, the auto-cancellation is targeting the highest value consumers, meaning that perhaps what the assignment does is assigning auto-renewal to most, except those who would sign up anyway based on their valuation.

Finally, trying to maximize the total probability of subscription leads to a different assignment altogether. Maximizing total subscribers leads to only 20% being assigned to auto-renew. Here, those assigned to auto-renew are far more likely to be naive, are less inert, and value differentiation is weaker than in the other two assignment rules.

These results show that targeting offers of auto-renew or auto-cancel heavily depend on the objective function. Further, even for objectives that seem closely tied with sophistication, such as maximizing first month after-the-promo subscriptions, the targeting scheme seems to pick up something quite different. Namely, the effectiveness of auto-renew depends on inertia but also on valuation. Therefore, segmentation picks up both differences in valuation and in sophistication, yet the former may swamp the latter making sophistication-based discrimination quite limited. We conclude that either the data the newspaper and us have, of pre-experiment browsing behavior, is not a strong predictor of sophistication or inertia, or that other factors are more important for assignment.

[Table 6 about here.]

9 Conclusion

The common wisdom in the academic literature, as well as in the industry, is that consumers are highly inert. Once a firm gained a consumer, the argument goes, the firm 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, consumers are highly inert. However, this body of knowledge relies on a selected sample of *already existing consumers*. Our paper suggests that a large portion of consumers in our setting, of at least 58%, is aware of its future inertia and avoids engaging with an exploitative contract. Furthermore, offering an exploitative contract pushes 10% of consumers from engaging with the company for the duration of our data. These new findings imply that consumers' awareness to 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.

In our setting, if the firm's horizon is a few months of profits, then indeed offering an auto-renewal contract would have been beneficial. However, if longer term profits matter, then there is no benefit for the auto-renewal contract if offered uniformly to all. The cumulative revenue advantage of auto-renewal is maximized after 8 months, but at that stage is already statistically indistinguishable from 0, and shrinks toward 0 as time goes by. Furthermore, if the market share, or size of readership matters, then auto-renewal is worse from day one. There are various reasons why readership matters, such as advertisement revenue, investor metrics, or the potential for word-of-mouth and social media engagement to expand readership further. Finally, as our usage analysis suggests, those who remain subscribed due to the auto-renewal nature of the contract do not use their subscription, meaning that the contract is indeed exploitative and does not bring value to consumers. Overall, at the medium- and long-run, if the firm can only choose one type of contract to offer, auto-cancellation contracts seems like a better choice on average for the firm and consumers.

In theory, the firm might be able to benefit from "sophistication discrimination." Either by offering different readers different contracts based on their naivete, as in a third-degree price discrimination, or by designing a contract menu to exploit naivete (e.g., Eliaz and Spiegler (2006)). In our setting, while targeting based on predicted value is useful, and there are substantive heterogeneous treatment effects, targeting that is based on naivete is infeasible even ex-post (to our best attempt). Further, while the newspaper already offers a host of contracts, including some not exploitative such as a one day pass, our results suggest that the mere offer of an exploitative contract as part of the menu deters some readers from participation. This notion, of consumers making inference about the firm from the set of contracts it offers, should be taken into account in contract design.

To summarize, we design a large-scale field experiment that enables us to study inertia in consumer subscription decisions. The experimental design simultaneously varies the contract renewal terms along with other benefits, which allows us to quantify the inertia readers anticipate from taking up the subscription, before they actually take it. Their subsequent subscription behavior enables us to quantify the actual inertia they experience. Overall, we find that readers do recognize and account for their inertia. At least 58% of readers are sophisticated about their future inertia, and will not enter a contract they do not wish to take. At the same time, about half of those who do take up the auto-renewal subscription are inert and end up paying for a subscription they do not want. Overall, in the long term, consumers behavior disincentivizes the newspaper to present auto-renewal offers, even though auto-renewal leads to higher firm revenue in the medium-run because of inertial subscribers.

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.

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.

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

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.

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.

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.

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: Theory and Field Evidence from Large-Scale Rebate Promotions," 2022.

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.

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

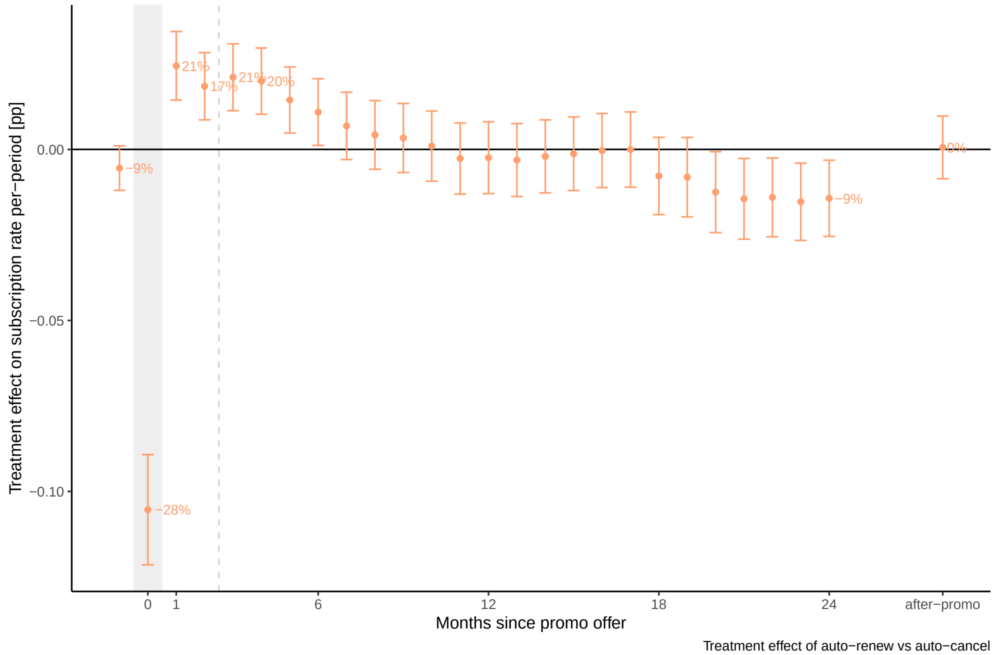

(a) Effect on Subscription Rate (Proportion of Days a Reader Subscribed)

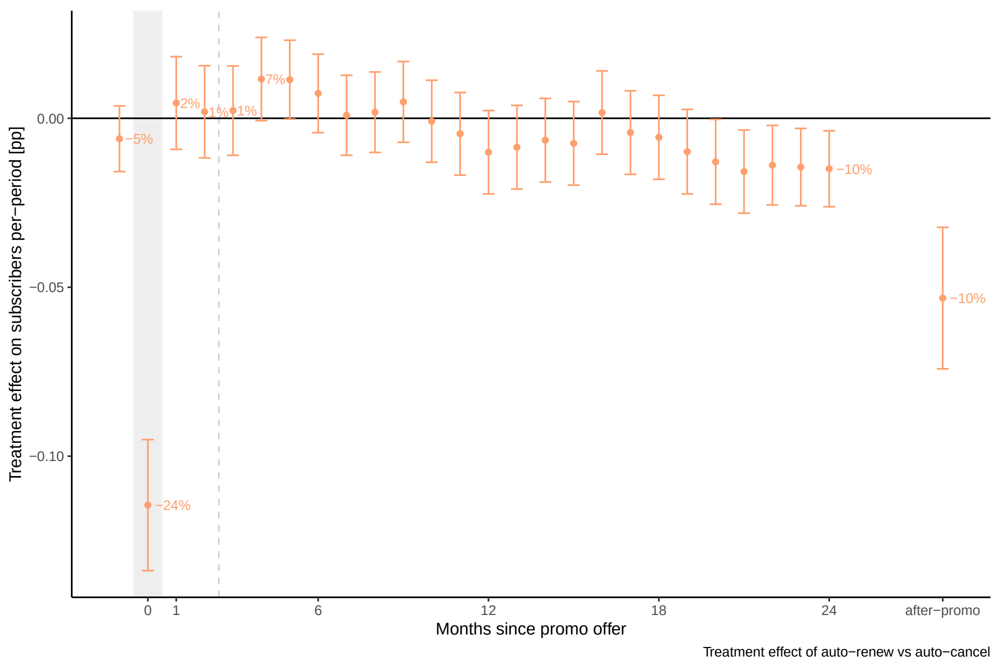

(b) Effect at the Extensive Margin (Whether the Reader Subscribed At All)

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* (1) *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 on 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 2: Effect of €0.99 Relative to Free Promotional Contracts on Overall Subscription Behavior

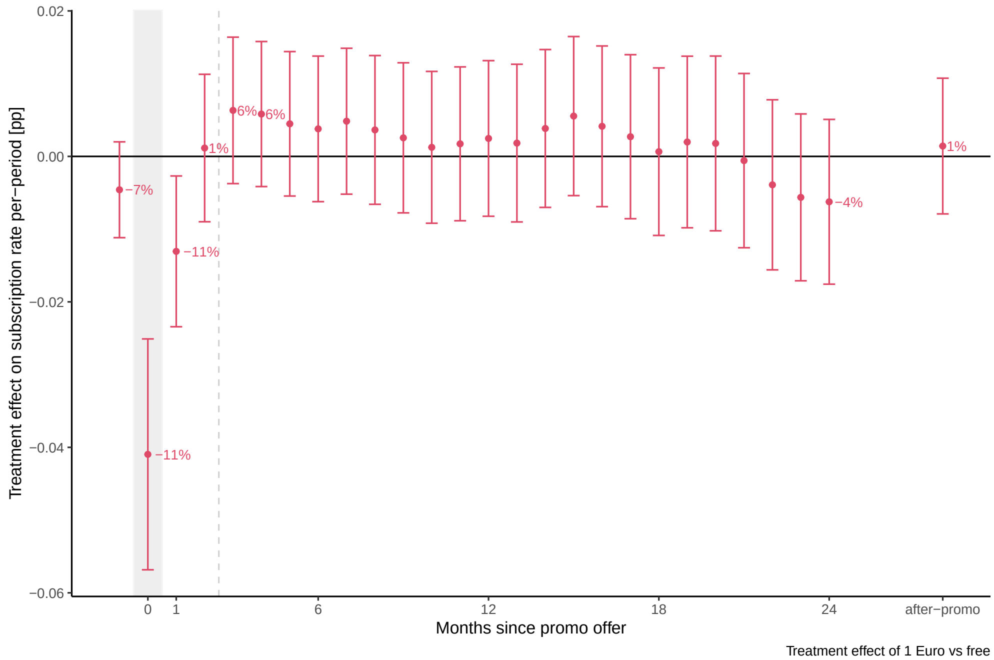

(a) Effect on Subscription Rate (Proportion of Days a Reader Subscribed)

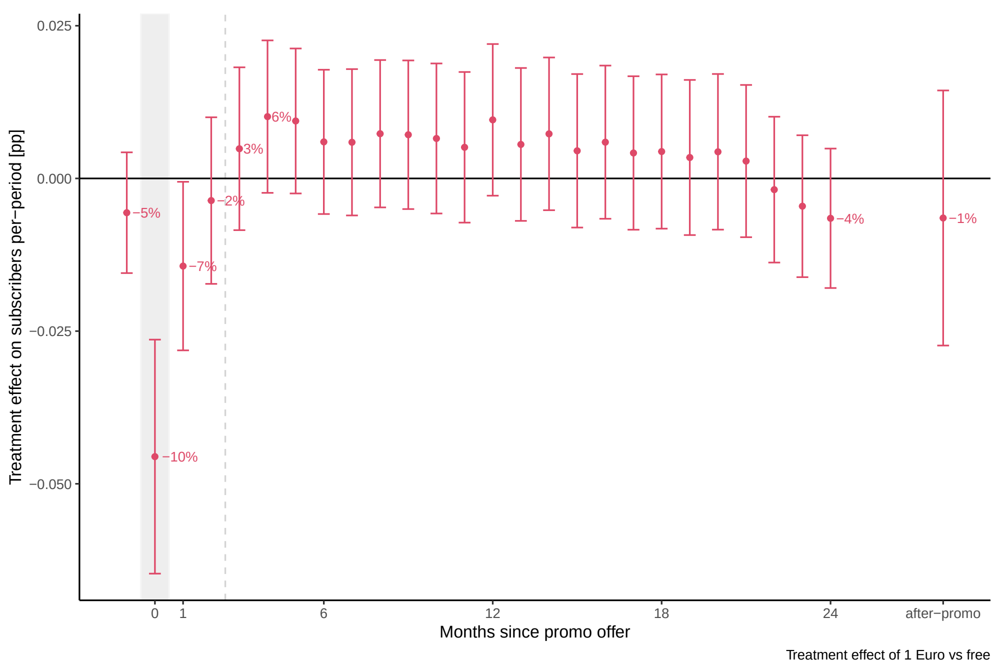

(b) Effect at the Extensive Margin (Whether the Reader Subscribed At All)

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 (1) 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 on 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 3: Effect of Four Weeks Relative to Two Weeks Promotional Contracts on Overall Subscription Behavior

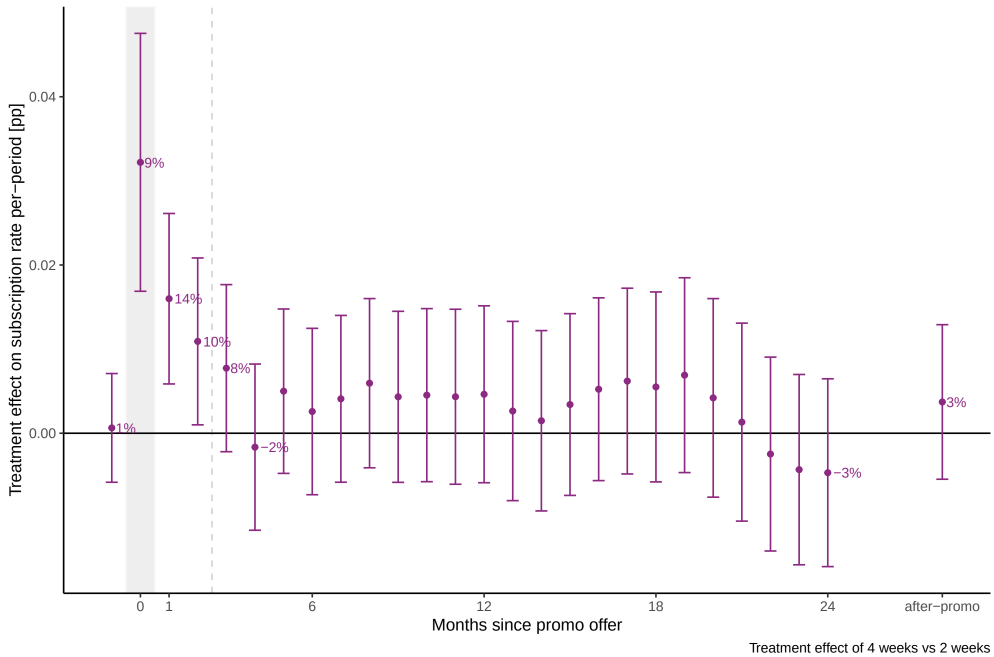

(a) Effect on Subscription Rate (Proportion of Days a Reader Subscribed)

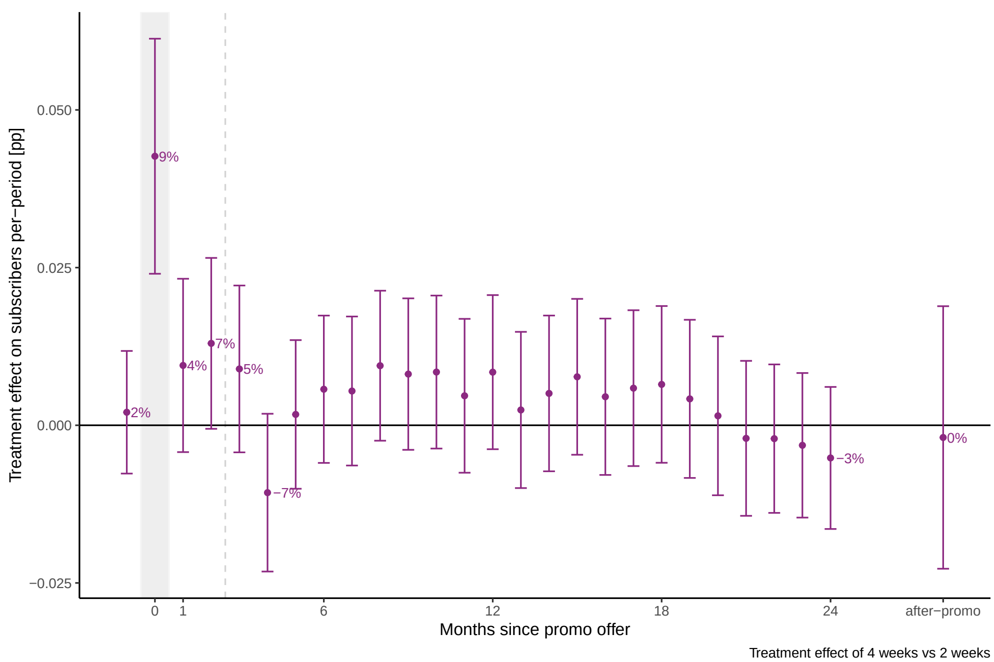

(b) Effect at the Extensive Margin (Whether the Reader Subscribed At All)

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 (1) 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 on 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 4: Subscription vs. Platform Usage for Two Weeks €0.99 Auto-Renewal Promo Takers vs. Two Weeks €0.99 Auto-Cancel Promo Takers

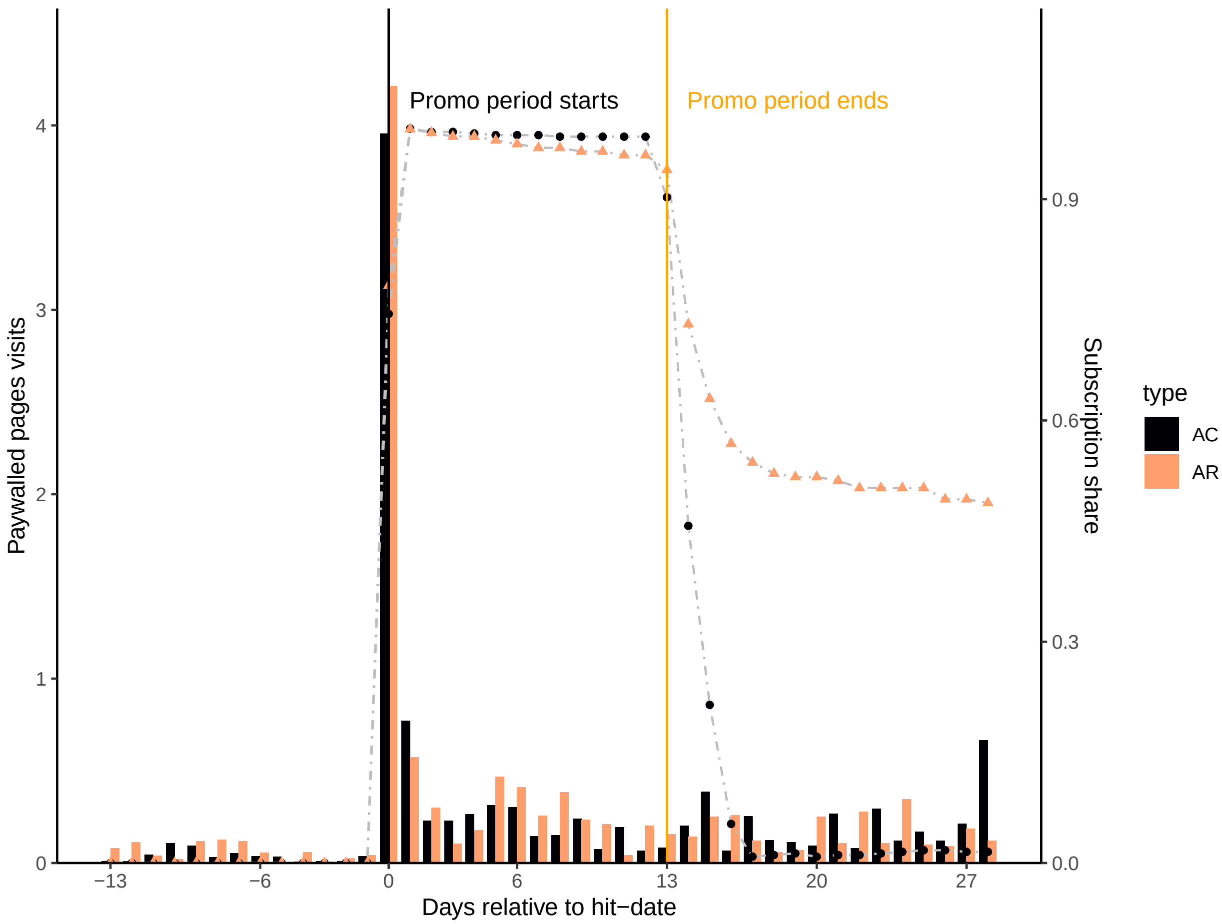

Notes: The figure plots the daily average subscription rate (dots and triangles) and average newspaper consumption—measured by number of website page visits (bars)—separately for those who took the two weeks €0.99 auto-renewal promo (AR, orange) and those who took the two weeks €0.99 auto-cancel promo (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

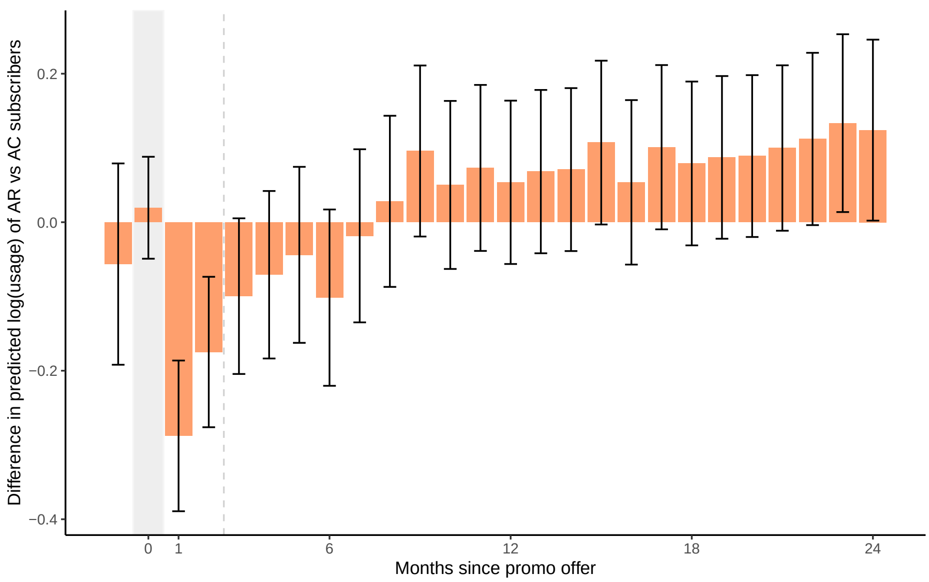

Notes: We use group G readers' pre and post-experimental usage data to predict post out-of-sample newspaper usage for the main sample. We use predicted usage as a proxy for reader type. The figure shows the difference in average reader type between subscribers in the auto-renewal groups and auto-cancellation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

Figure 6: Distributions of Reader Types for Subscribers in the Auto-Renewal and Auto-Cancel Group by Period

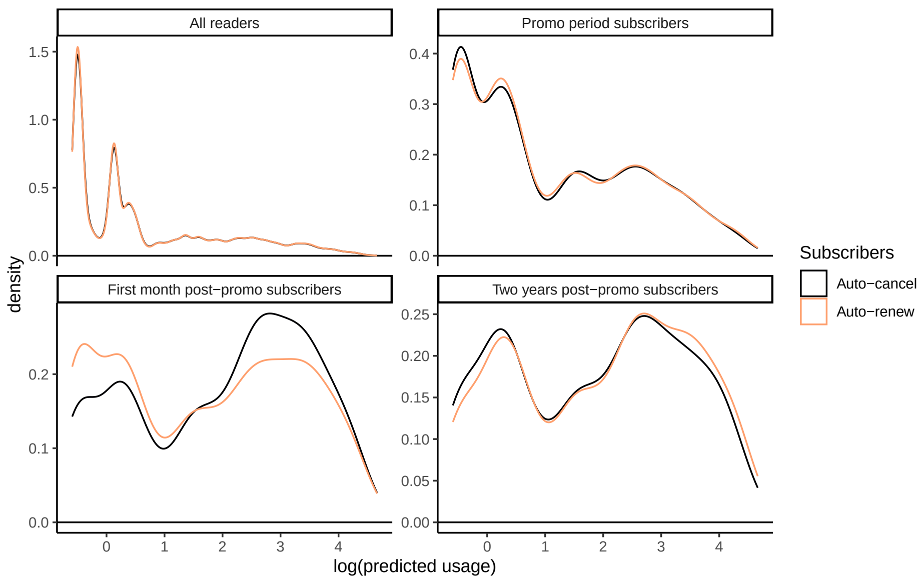

Notes: We use group G readers' pre and post-experimental usage data to predict post out-of-sample newspaper usage for the main sample. We use predicted usage as a proxy for reader type. The figure shows the distributions of predicted reader types for subscribers in the auto-renewal group and auto-cancellation 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.

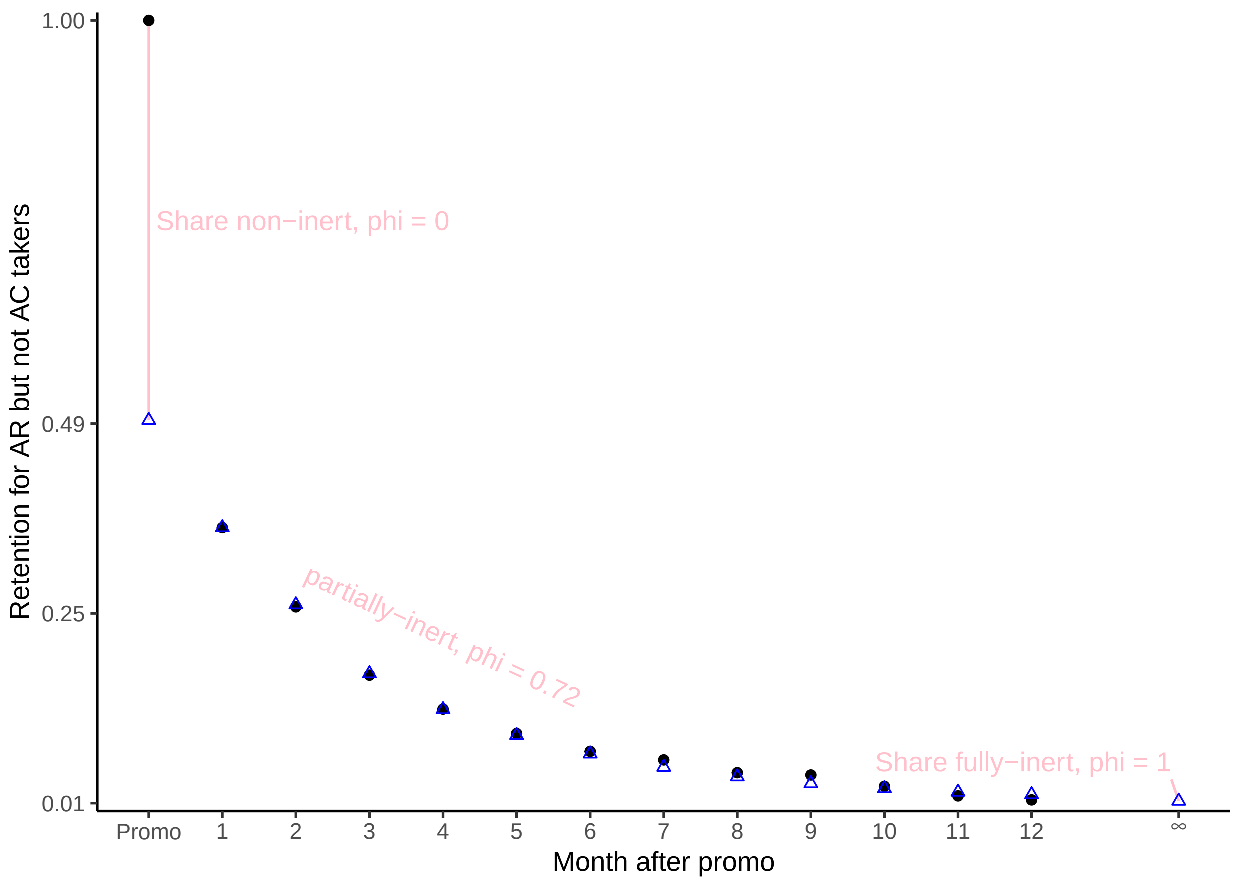

Figure 7: Data and Estimation of Three-Types Inert Subscribers

Notes: The figure shows the excess share of auto-renew promo takers who stay subscribed after the promo ends overlayed with the three-type model's prediction. The points are the data, starting with the full sample (y = 1) as the share of promo takers, and showing the survival of this sample in the eight months following. The triangles are the predictions from the model in (2) estimated on this data. The prediction at "infinity" are the 1.3% fully inert auto-renewal takers who will not become long-run subscribers under auto-cancellation. The vertical difference between 1.00 and the triangle at "promo" gives the 50.6% of non-inert subscribers who cancelled their subscription before starting to pay full prices. The remaining 48.2% are estimated to have a monthly inaction parameter of $\phi = 0.718$.

Figure 8: Validation of Predicted Usage as Predicting Subscriptions

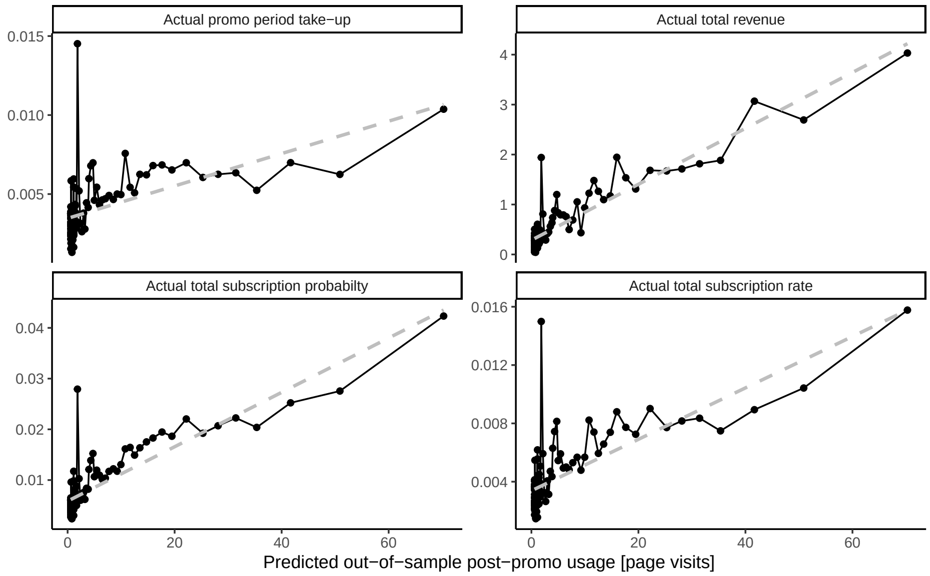

Notes: The figure shows the ITT effects of auto-renewal separating readers into types based on their predicted usage given their pre-experiment usage. High types are those predicted to be at the 0.4% of most engaged readers, low types are the bottom 99.6%.

Figure 9: Auto-Renewal Effects by Types of Readers

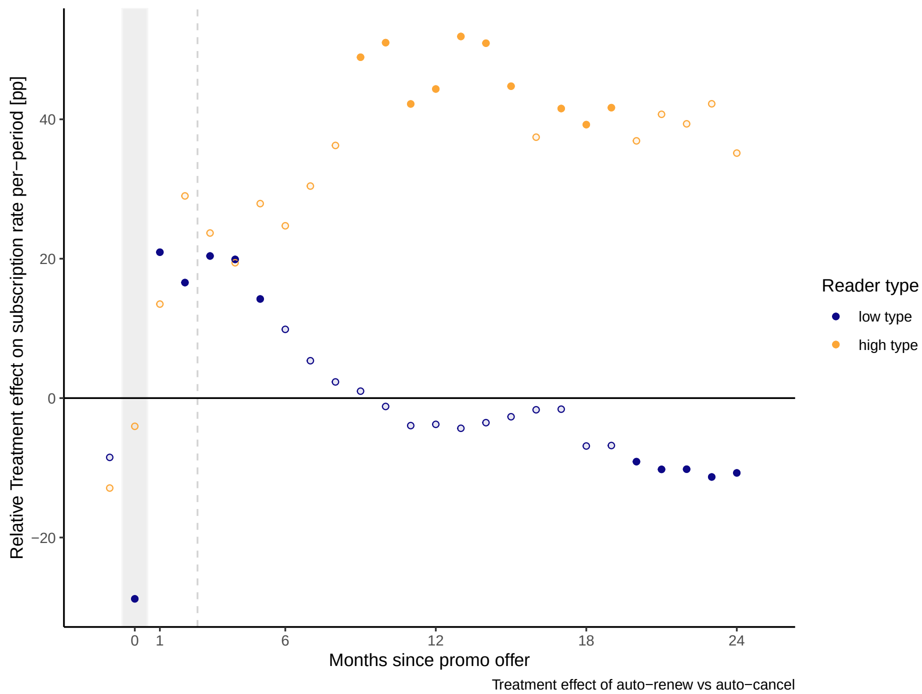

Notes: The figure shows the relative ITT effects of auto-renewal separating readers into types based on their predicted usage given their pre-experiment usage. To calculate the relative effect, we regress for each period the outcome on a fully interacted set of treatments (auto-renewal, price, duration) with type classification, and then divide the treatment effect of auto-renew by the type's baseline level. High types are those predicted to be at the 0.4% of most engaged readers, low types are the bottom 99.6%. Full circles are statistically significant at the 95% level, hollow circles are not.

Table 1: Experimental offers

Experimental group	Renewal	Duration	Price
A	Auto-renewal	four weeks	€0
B	Auto-renewal	four weeks	€0.99
C	Auto-renewal	two weeks	€0
D	Auto-renewal	two weeks	€0.99
E	Auto-cancel	four weeks	€0
F	Auto-cancel	four weeks	€0.99
G	Auto-cancel	two weeks	€0
H	Auto-cancel	two weeks	€0.99

Table 2: Summary Statistics

vars	Raw	Main	Takers
Number of readers	4,131,277	2,088,910	5,882
Number of subscribers	36,816	16,237	5,882
Total revenue (in Euros)	1,998,352	1,325,036	216,514
Number of pages viewed	143,628,050	87,429,810	3,074,297
- Always free	123,081,315	75,297,261	2,755,315
- Always paid	14,545,384	8,091,819	198,338
- Metered	6,001,351	4,040,730	120,644

Table 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 (s.e.)	Not subscribed in two weeks after the promo estimate (s.e.)	Subscribed in two weeks after the promo estimate (s.e.)	Not subscribed in two weeks after the promo estimate (s.e.)
Promo two weeks: Avg. page visits	25.37 (3.66)	22.28 (2.64)	57.71 (7.10)	23.66 (1.86)
Promo two weeks: % readers with any visit	0.78 (0.04)	0.68 (0.04)	0.90 (0.04)	0.68 (0.02)
after the promo two weeks: Avg. page visits	20.33 (3.57)	12.88 (2.92)	75.98 (10.71)	12.67 (1.24)
after the promo two weeks: % readers with any visit	0.49 (0.04)	0.43 (0.04)	0.93 (0.03)	0.48 (0.02)
N	134	136	58	567

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. The next two rows do the same for the subsequent two weeks. The results show that, on average, the auto-renewal contract takers who are still subscribed after the promo period use the newspaper with the same intensity as those who did not subscribe; their usage is lower than those in auto-cancel who subscribe after the promo.*

Table 4: Inertia, Sophistication Types, and Subscription Behavior under Auto-Renewal

<table>
<tr><th></th><th>Non-inert
($\phi = 0$, $c^u = 0$)</th><th>Partially-inert
($\phi = 0.72$, $c^u = 0$)</th><th>Fully-inert
($\phi = 1$ or c^u is large)</th></tr>
<tr><td>Naïve
($\tilde{\phi}$=$\tilde{c}^u$=0)</td><td rowspan="2">Take AR promo
(equivalent to AC)</td><td>Take AR promo</td><td>Take AR promo</td></tr>
<tr><td>Sophisticated
($\tilde{\phi} = \phi$, $\tilde{c}^u = c^u$)</td><td>Take only if value
balances inertia risk</td><td>Take only if wants
to subscribe in the long-run</td></tr>
</table>

Notes: The table shows a schematic breakdown of the 5 different types of consumers in our model. The rows describe naivete or sophistication, and the columns are the degree of inertia. The left-most column, non-inert, is not split by naivete since naivete and sophistication are equivalent. Within each cell the table shows the predicted take-up of an auto-renewal promotional offer of that type compared to if they were offered an auto-cancel offer. For example, the behavior of all naive types is equivalent between auto-renew and auto-cancel.

Table 5: Estimated Shares of Inertial Types

	Estimate	**SE**
Inertia (of partially-inert)	0.718	(0.020)
In population:		
Share sophisticated inerts	0.582	(0.019)
Share non-inert	0.302	(0.013)
Share partially-inert	0.680	(0.015)
Share fully-inert	0.018	(0.013)
Among auto-renew promo takers:		
Share sophisticated inerts	0.060	(0.010)
Share non-inert	0.506	(0.020)
Share partially-inert	0.482	(0.018)
Share fully-inert	0.013	(0.009)
Share responding to price increase	0.017	(0.008)

Notes: The table shows the estimated shares of inertial types in the population, and among auto-renew promotional offer takers. Dashed lines separate between independently distributed types. For example, in the population, there are 58% sophisticated among the inerts (and 42% are naifs), and sophistication is assumed to be independent of the inertial type. Therefore, if there are 68% partially-inert readers, then $68\% \times 58\% = 39\%$ *in the population are sophisticated-partial-inerts.*

Table 6: Who is Being Targeted

	Total revenue	Post-subscription	First month post-subscription
Share assigned to AR	77.7%	19.8%	72.4%
Share sophisticates in AR vs AC	62.1% vs 48.4%	44.8% vs 62.1%	58.6% vs 59.8%
Average predicted value in AR vs AC	3.7 vs 13.2	1.3 vs 6.9	1.4 vs 17.4
Actual inertia for partial-inerts AR vs AC	0.77 vs 0.68	0.68 vs 0.77	0.76 vs 0.74

Notes: The table shows some main characteristics of the readers who would have been targeted with auto-renewal (AR) or auto-cancel (AC) offers. Each column describes a different target function to maximize, from left to right – maximizing total revenue, maximizing the probability of any subscription after the promo period ends, and maximizing the probability of subscription shortly after the promo period ends. The rows show the baseline probability of being offered auto-renewal contract, and among those who are assigned what share of them are sophisticated, and what are their inertia level if they are partially-inert.

A ONLINE APPENDIX

A.1 Additional Tables and Figures

[Table 7 about here.]

[Figure 10 about here.]

[Figure 11 about here.]

[Figure 12 about here.]

[Figure 13 about here.]

[Figure 14 about here.]

[Figure 15 about here.]

[Figure 16 about here.]

[Figure 17 about here.]

[Figure 18 about here.]

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 at all following periods if unsubscription failed, is $\sum_{\tau=0}^{\infty}\left[\hat{\phi}_i^{\tau}\delta^{\tau}\left(v_i - p - \left(1-\hat{\phi}_i\right)c^u\right)\right] = \frac{v_i - p - \left(1-\hat{\phi}_i\right)\hat{c}^u}{1-\delta\hat{\phi}_i}$.[21]

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):

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

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:[22]

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

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,

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

[22] 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.

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.[23] 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.

A.3 Using Experimental Incentives to Quantify Inertia

In this subsection, we estimate inertia by comparing the differential treatment effects of price reduction and trial duration across auto-renewal and auto-cancel contracts. The rationale is as follows. An experimental incentive—price reduction or an increase in the trial duration—causes some people assigned to an auto-renewal group to take up a subscription during the promo time period. Let Δy_0^{AR} denote this effect. Since the marginal type who should renew is higher than the marginal type who takes the auto-renewal promo, those who are encouraged to take the promo should not become full-price subscribers. For example, in the duration and price reduction treatments indeed additional auto-cancel subscribers do not become full-price subscribers. However, for auto-renewal, some of these additional promo takers do stay subscribed. The proportion of this effect that lasts after the promotional time $\Delta y_t^{AR} = (\lambda + \phi^t) \times \Delta y_0^{AR}$, where $\phi^t \times \Delta y_0^{AR}$ continue because of the inertia caused by auto-renewal, and $\lambda \times \Delta y_0^{AR}$ are those who decide to continue the subscription (e.g., due to them learning that they value it more than the price).

The corresponding effect of the experimental incentives within the auto-cancel group will be similar except that there will be no inertia, i.e. $\Delta y_t^{AC} = \lambda \times \Delta y_0^{AC}$. Hence, we estimate the effect of inertia in any month t as

$$\phi^t = \frac{\Delta y_t^{AR}}{\Delta y_0^{AR}} - \frac{\Delta y_t^{AC}}{\Delta y_0^{AC}}. \tag{5}$$

In contrast to the approach in section 8.1, which estimates average inertia experienced across all auto-renewal takers, this approach estimates inertia experienced by the marginal individuals—those who take an auto-renewal subscription only when an additional incentive is given with it.

Table A.2 shows our estimates. For individuals assigned an auto-renewal offer reducing price and in-

[23] 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.

creasing trial duration simultaneously, that is, going from two weeks, €0.99 auto-renewal to four weeks, free auto-renewal increases the likelihood of an individual subscribing during the promo period by 0.0013, which is our estimate for Δy_0^{AR}.[24] Looking beyond the promo period, in the four weeks after the promo the difference Δy_1^{AR} is 53.26% $\times \Delta y_0$. This suggests that about half of the immediate increase in subscribers due to the experimental incentives extends beyond the time when the incentives are applicable. Beyond the first month after the promo, we see a gradual drop in Δy_t^{AR}, which is detectable up to month 3.

The same incentive for those assigned to the auto-cancel group also increases subscriptions during the promo period by 0.00044, which is smaller relative to the auto-renew group. However, we do not see this increase extending beyond the promo time period. If anything, we see lower subscription after the promo, which could just be due to imprecision. For the auto-cancel group our estimate for $\frac{\Delta y_1^{AC}}{\Delta y_0^{AC}}$ is imprecise but significantly lower than the corresponding estimate for auto-renewal group.

Overall, these estimates indicate the presence of inertia on the marginal individuals. However, the ϕ_t estimates are very imprecise.

[Table 8 about here.]

[24] For this exercise, we consider the largest increase in incentives within our experiment for most precise estimation of relative increases. Considering only price changes gives similar findings.

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

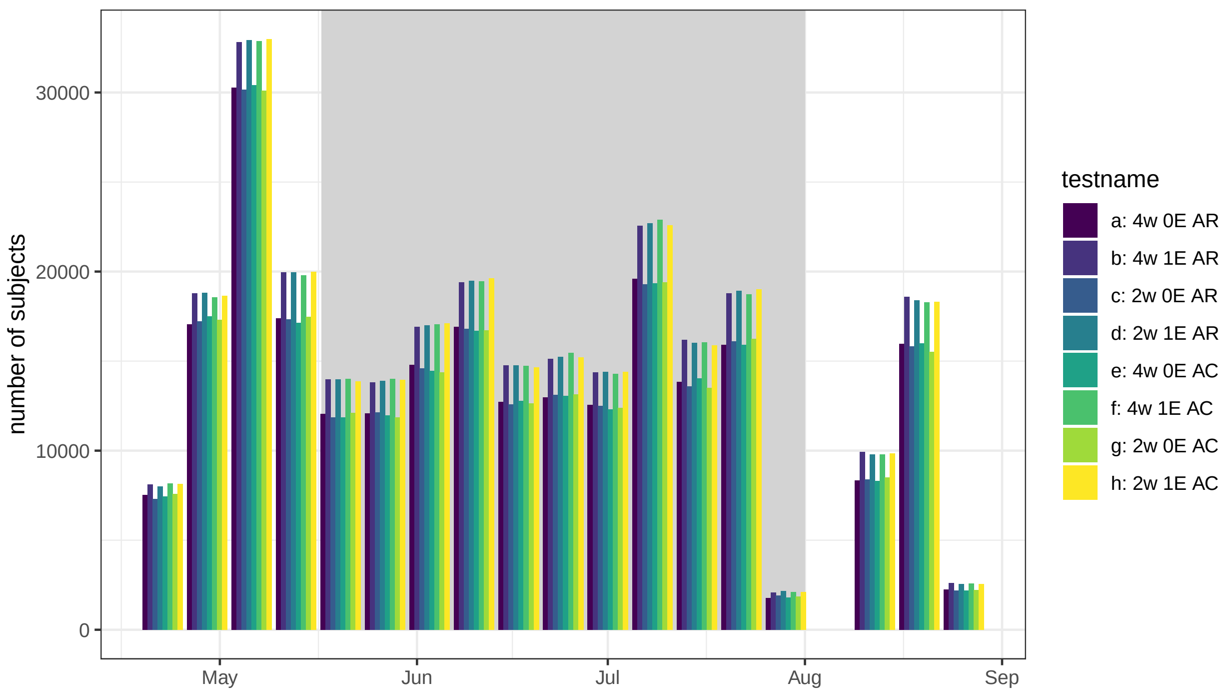

Notes: The figure shows the number of readers exposed to each experimental contract by week. The difference in shading represents the different phases of the experiment.

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

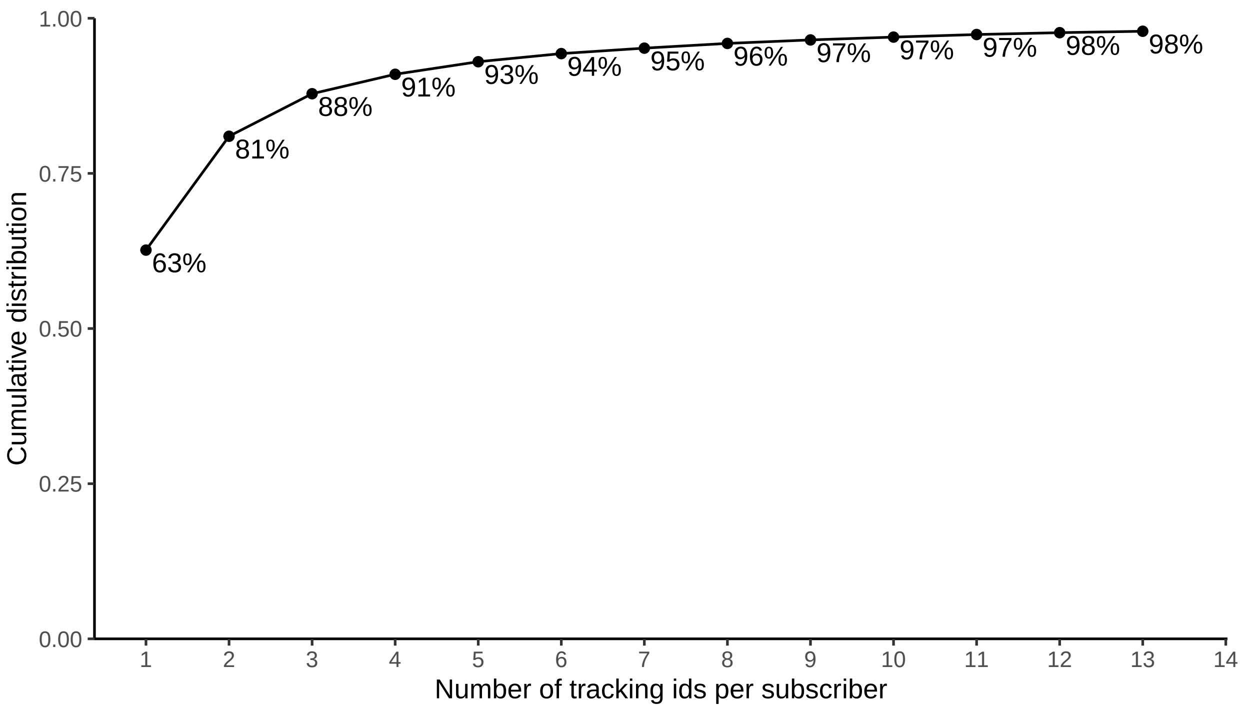

Notes: The figure shows the cumulative distribution of the unique number of cookies for each subscriber.

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

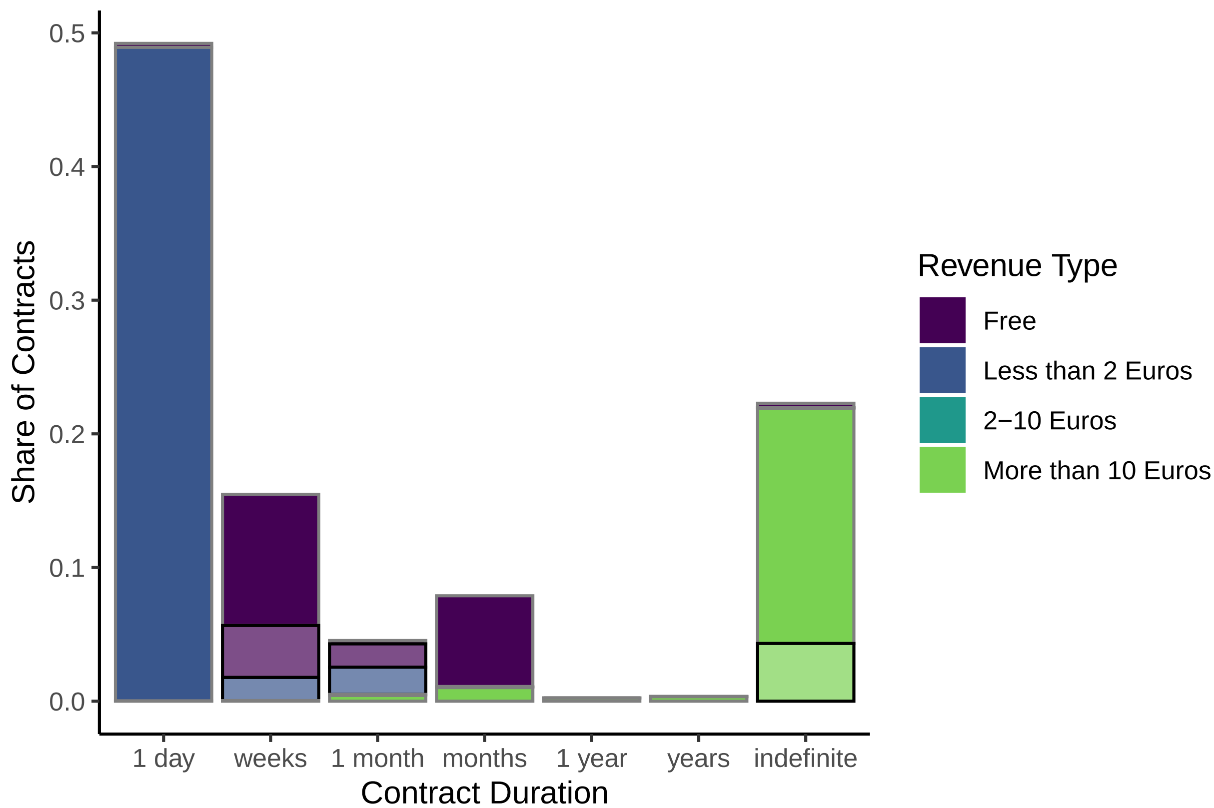

Notes: The figure shows shares of contracts taken, characterized by their maximal duration (horizontal axis) and revenue (color). For example, almost half of all contracts are daily passes that cost €1.99. The dark rectangles highlight the experimental contracts—the auto-cancellation contracts are either 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

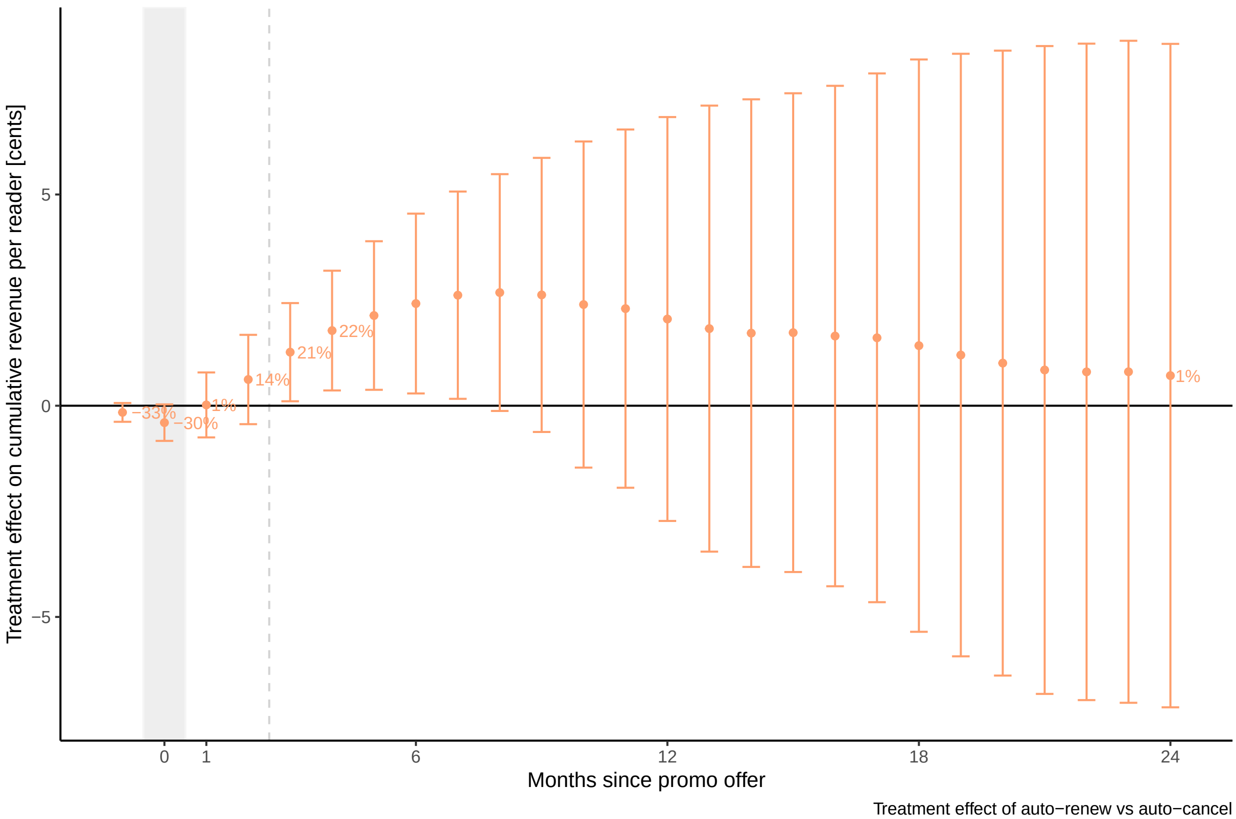

Notes: The figure plots the estimated average intent-to-treat effect of serving an auto-renewal relative to an auto-cancel contract on the newspaper's cumulative revenue.Specifically, we plot the estimated coefficient β_1 from equation (1) for every month with cumulative revenue as the dependent variable. 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 on 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

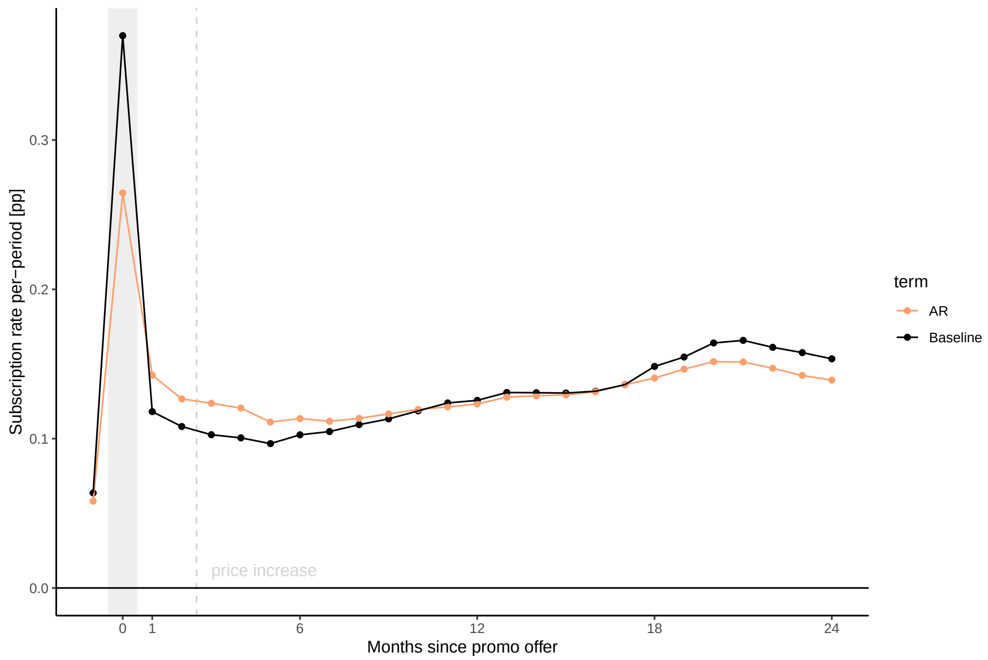

(a) Subscription Rate (Proportion of Days a Reader Subscribed)

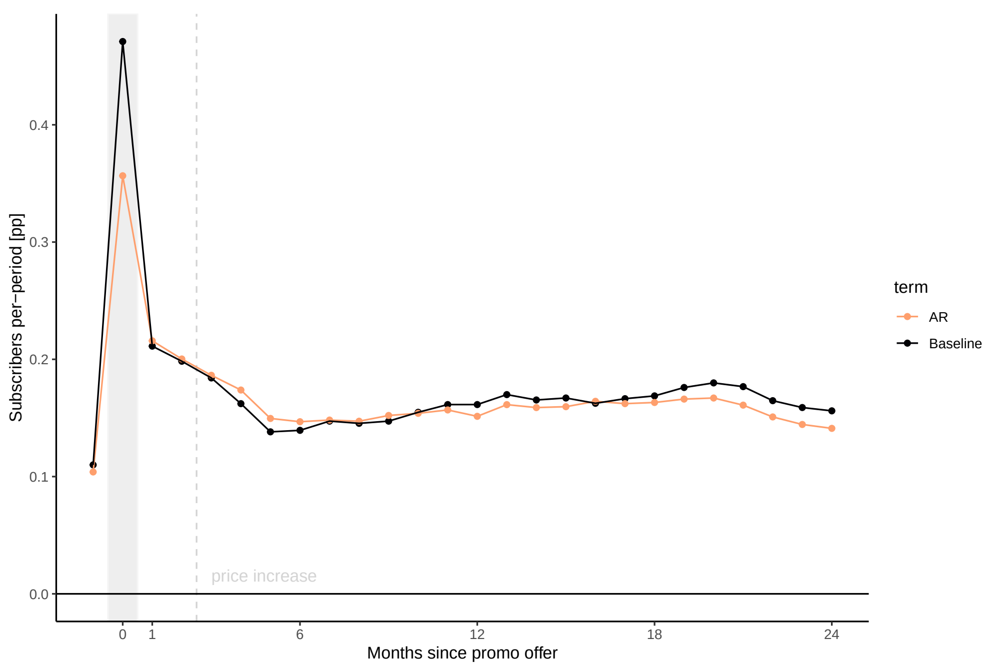

(b) Extensive Margin (Whether the Reader Subscribed At All)

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* (1) *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 on and so forth. 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.*

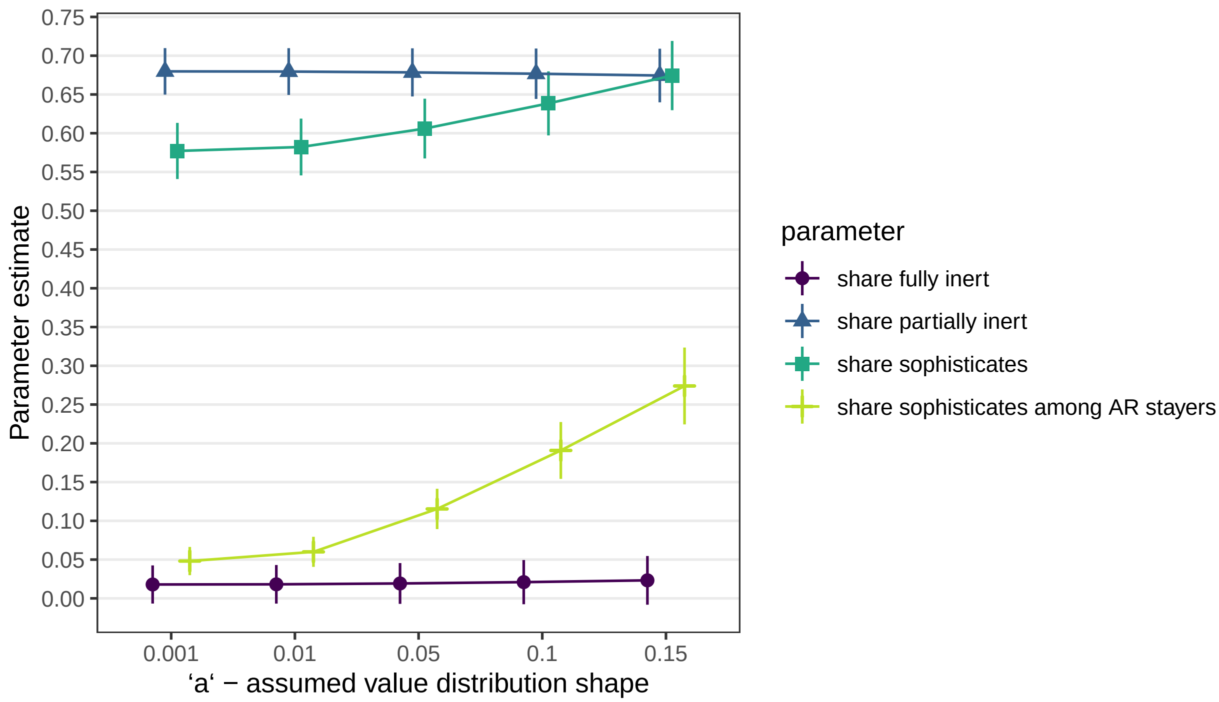

Figure A.6: Parameters’ sensitivity to a, the parameter of value distribution skeweness

Figure A.7: Auto-Renewal Effects on Subscription Rate by Types of Readers

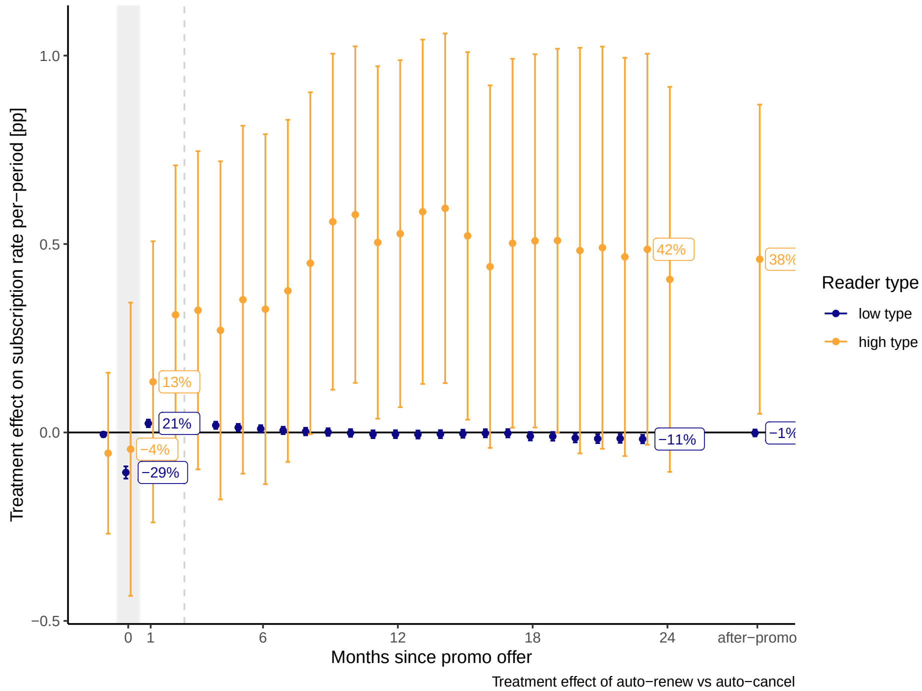

Notes: The figure shows the absolute intent-to-treat (ITT) effects of auto-renewal separating readers into types based on their predicted usage given their pre-experiment usage. We regress for each period the outcome on a fully interacted set of treatments (auto-renewal, price, duration) with type classification, and plot here the treatment effect of auto-renew. High types are those predicted to be at the 0.4% of most engaged readers, low types are the bottom 99.6%.

Figure A.8: Auto-renewal Effects on Subscribers by Types of Readers

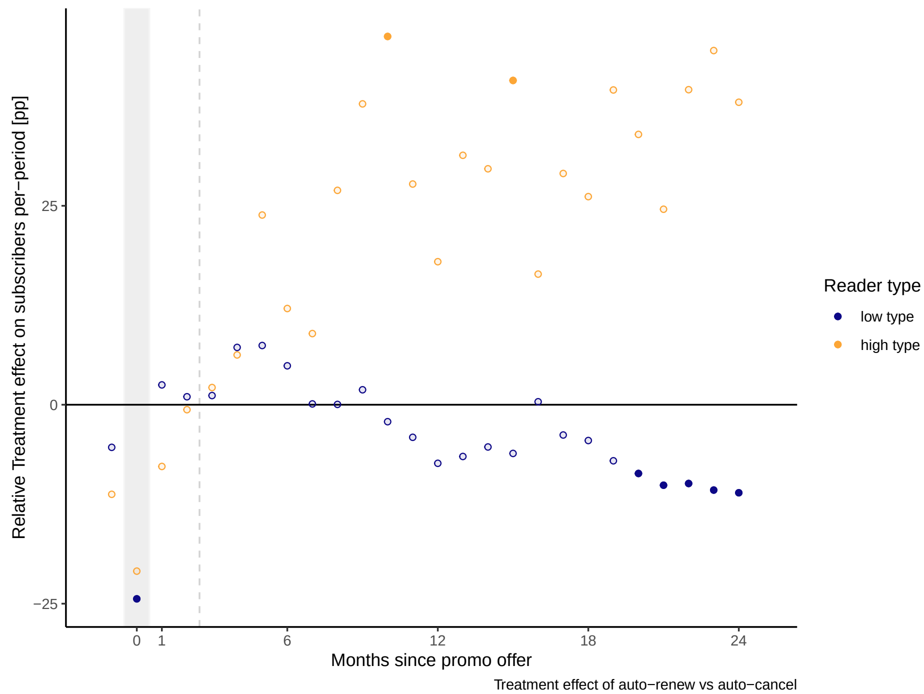

Notes: The figure shows the relative intent-to-treat (ITT) effects of auto-renewal separating readers into types based on their predicted usage given their pre-experiment usage. To calculate the relative effect we regress for each period the outcome on a fully interacted set of treatments (auto-renewal, price, duration) with type classification, and then divide the treatment effect of auto-renew by the type's baseline level. High types are those predicted to be at the 0.4% of most engaged readers, low types are the bottom 99.6%. Full circles are statistically significant at the 95% level, hollow circles are not.

Figure A.9: Auto-renewal effects on cumulative revenue by types of readers

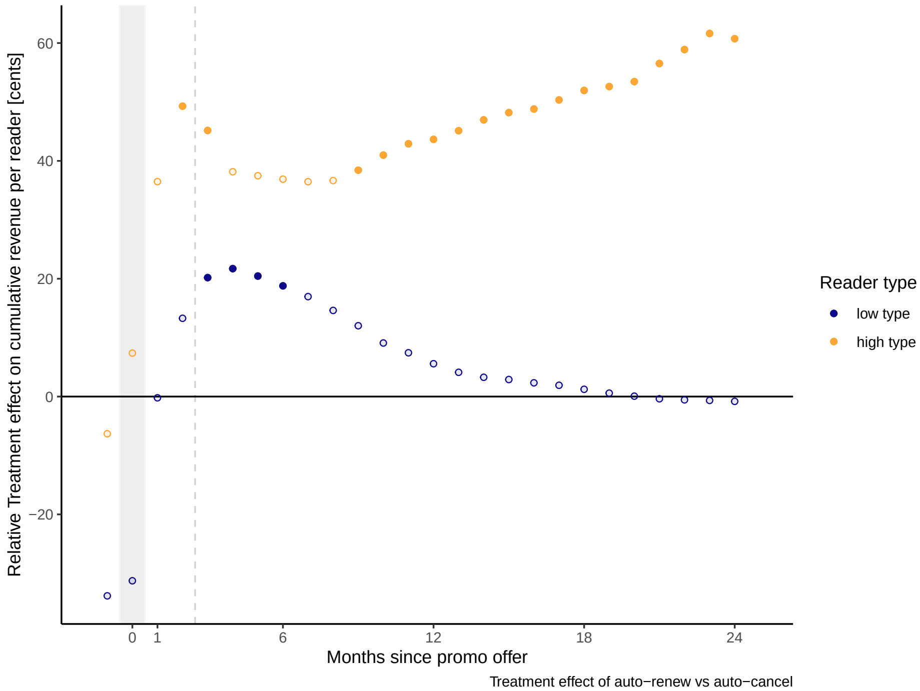

Notes: The figure shows the relative ITT effects of auto-renewal separating readers into types based on their predicted usage given their pre-experiment usage. To calculate the relative effect we regress for each period the outcome on a fully interacted set of treatments (auto-renewal, price, duration) with type classification, and then divide the treatment effect of auto-renew by the type's baseline level. High types are those predicted to be at the 0.4% of most engaged readers, low types are the bottom 99.6%. Full circles are statistically significant at the 95% level, hollow circles are not.

Table A.1: Balance of pre-experiment behavior

Dependent Variables:	Total_Pages	Open	Paywalled
Model:	(1)	(2)	(3)
Variables			
Test Group A	5.179***	4.999***	0.1798***
	(0.0403)	(0.0389)	(0.0058)
Group B vs A	0.0117	0.0087	0.0031
	(0.0560)	(0.0539)	(0.0092)
Group C vs A	0.1169*	0.1077*	0.0092
	(0.0611)	(0.0592)	(0.0085)
Group D vs A	-0.1155**	-0.1177**	0.0022
	(0.0514)	(0.0492)	(0.0089)
Group E vs A	0.0279	0.0169	0.0110
	(0.0589)	(0.0568)	(0.0092)
Group F vs A	-0.0335	-0.0295	-0.0040
	(0.0555)	(0.0537)	(0.0077)
Group H vs A	-0.0706	-0.0798	0.0092
	(0.0543)	(0.0518)	(0.0103)
Fit statistics			
Observations	1,837,309	1,837,309	1,837,309
R^2	1.17×10^{-5}	1.21×10^{-5}	2.16×10^{-6}
Adjusted R^2	8.46×10^{-6}	8.88×10^{-6}	-1.11×10^{-6}

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

Table A.2: Effect of Experimental Incentives on After-the-Promo Subscription

	Auto-renewal four weeks, Free vs. two weeks, €0.99	Auto-cancel four weeks, Free vs. two weeks, €0.99
	estimate (s.e.)	*estimate (s.e.)*
Effect in promo time (Δy_0)	.0013 (.0002)	.00044 (.00018)
Effect after the promo month 1 (Δy_1)	.0007 (.0001)	-.00019 (.00013)
Effect after the promo month 2 (Δy_2)	.0004 (.0001)	-.00011 (.00013)
Effect after the promo month 3 (Δy_3)	.0001 (.0001)	-.00010 (.00012)
Effect after the promo month 4 (Δy_4)	-.00018 (.00011)	-.00025 (.00011)
$\frac{\Delta y_1}{\Delta y_0}$	.5326 (.0906)	-.4434 (.4047)
$\frac{\Delta y_2}{\Delta y_0}$	.3394 (.0900)	-.2666 (.3338)
$\frac{\Delta y_3}{\Delta y_0}$	.1112 (.0889)	-.2299 (.3143)
$\frac{\Delta y_4}{\Delta y_0}$	-.1389 (.0939)	-.5696 (.3976)
$\phi_1 = \frac{\Delta y_1^{AR}}{\Delta y_0^{AR}} - \frac{\Delta y_1^{AC}}{\Delta y_0^{AC}}$	.9760 (.4147)	
$\phi_2 = \frac{\Delta y_2^{AR}}{\Delta y_0^{AR}} - \frac{\Delta y_2^{AC}}{\Delta y_0^{AC}}$	.6061 (.3458)	
$\phi_3 = \frac{\Delta y_3^{AR}}{\Delta y_0^{AR}} - \frac{\Delta y_3^{AC}}{\Delta y_0^{AC}}$	.3411 (.3267)	
$\phi_4 = \frac{\Delta y_4^{AR}}{\Delta y_0^{AR}} - \frac{\Delta y_4^{AC}}{\Delta y_0^{AC}}$	.4307 (.4085)	

Notes: The first four rows of the table present the effect of changing the promotional terms from (four weeks, free) to (two weeks, €0.99) on the promo period (Δy_0) and after the promo (Δy_t) subscription rates, separately for auto-renewal and auto-cancel groups. The next four rows present our estimate of after the promo depreciation of subscription relative to promo time. These estimates show that, under auto-renewal, the subscription rate drops to 53% of Δy_0 in the first month after the promo and is statistically indistinguishable from zero by month 3. Under auto-cancel, the subscription rate drops immediately after the promo, and is statistically insignificant in all after the promo months. The next four rows present our estimate of the difference in subscription depreciation in auto-renew minus auto-cancel groups. These numbers are large – implying significant inertia – but imprecise because the auto-cancel estimates are large and imprecise.