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

Klaus Miller, Navdeep S. Sahni, and Avner Strulov-Shlain*

November 15, 2022

Abstract

Consumer inertia, the tendency to remain inactive, is a robust and well-documented phenomenon. However, if consumers are aware of their future inertia they can act to mitigate its effects on their outcomes. Using a large-scale randomized field experiment with a leading European newspaper we investigate consumer response to inertia inducing subscription contracts and study, in the same setting, both the actual inertia, and the inertia consumers anticipate before it actually takes place. We vary the promotional subscription price, the duration, and whether the contract automatically renews by default, or not, after the promotional period. Indeed, we find strong inertia, such that the auto-renewal contract takers have a seven times higher tendency of continuing their subscription after the promotional period, relative to the auto-cancel contract takers. However, consumers preempt inertia; 24%-36% of potential subscribers avoid taking the auto-renewal offers. Further, 9% avoid subscribing, at all, for two years due to being offered the auto-renewal contract. Overall, even though auto-renewal generates a higher revenue in the short term, auto-renewal and auto-cancel are revenue equivalent after one year, but with fewer subscribers in auto-renewal. Using a simple model we find stronger support for inertia driven by "inaction" rather than switching costs. Our estimates of heterogeneity suggest about 20-35% of the individuals are naifs while the rest are not, and would take advantage of an inertia-exploiting short term incentive. Our results highlight the often-ignored effects of potentially exploitative inertia-inducing contracts: lower take up in the short- and long-run driven by sophisticated consumers.

*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. We are thankful for the comments we received from seminar and conference participants at Booth-Berkeley Behavioral seminar, ITAM, Cornell research camp, UC Irvine, HBS AI Conference , UChicago AFE conference, 2021 Marketing Science Conference, 2021 EMAC, VQMS, Quantitative Marketing and Economics Conference 2022.

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 such as to increase inertia.

Crucially, these consequences of inertia depend not just on the degree of inertia, but also on whether consumers are aware of their inertial tendency and how they account for it in their decision making. In a world where 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 appear good initially but are worse in the long run[1]. On the other hand, if consumers are aware of their behavioral limitations, they will account for them in their decision making and avoid getting into situations where they might get exploited due to inertia or find other ways to limit its effects (see also Rodemeier (n.d.)). This consideration will discourage firms from creating situations that might be construed as exploitative by consumers. Hence, even if consumers have inertia, its negative impact is mitigated due to their self-awareness. Of course, it is plausible that consumers are heterogeneous in their future inertia awareness, which can also be taken into account by firms by creating price, or inertia, discrimination (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 these differ between consumers? And what are the effects of these forces on firm incentives and outcomes?

A prerequisite to empirically inferring whether consumers take into account their inertia while making decisions is observing their behavior before they make a choice that might put them in an adverse state due to inertia. In contrast, most of the previous literature documents inertia among individuals who have already made a choice and gotten into an inert state, and misses consumers who avoided entering an inertia-inducing 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 caused by the choices 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 in which we randomize the terms of the subscription offers received by 2.1 million readers who hit the digital paywall of a large European daily newspaper. Our experiment is a 3-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 4 weeks, or 2 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

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

for 2 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. The difference in continuation of subscription after the promo time 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 the effect of product trial on the long-term subscription rate, which is useful in interpreting the effect of serving the auto-renewal vs. auto-cancel offer. Second, simultaneously varying the promotional price and the subscription renewal terms helps us quantify in monetary terms how much individuals value not getting defaulted into the subscription after the promotion ends. Third, simultaneously varying the promotional price and duration allows us to quantify the average value of subscription, which in turn enables us to calibrate the consumers' expected inaction and cancellation costs at the time they take up the subscription.

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 group, 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 offer, relative to the auto-cancel one for about four months post promotion. After 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 group. 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 9% less likely to become paid subscribers at any time in the two years after the promotion, relative to auto-cancel. We do not observe such a push-back for other experimental factors; even though €0.99 vs. free promo and 2 weeks vs. 4 weeks both cause 9% fewer people to subscribe during the promo period, their impact in the time period of two years after the promo are precise 0%. 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 term (up to four months post 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 consumer differ, non-parameterically, by their value of the subscription. There are 3 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 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 25% are non-inert, 5% are fully-inert, and 70% are partially-inert with a 65% monthly chance of not cancelling a subscription they wish to cancel. We estimate that a large majority, 54%-73%, among the inert are sophisticated and know their type. Being sophisticated means for the fully-inert that they will not subscribe and for the partially-inert that they will only subscribe if their value is worth the anticipated risk of being subscribed for longer than wished.

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 post 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, we also find long-term subscribers in the auto-renewal group to be of higher type (those that value the newspaper more) relative to the auto-cancel group, which goes against the prediction of a stochastic costs model; if hassle costs are stochastic, marginal rational consumers are more likely to remain subscribed, in the long term, in auto-renewal relative to auto-cancel leading to lower average "type" of auto-renewal subscribers.

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; Hortaçsu et al., 2017; Madrian and Shea, 2001) that documented 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 user 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.

We contribute to a much smaller literature that examines people's 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 (n.d.) 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 whole 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 avoid the service due to future inertia. Meaning that some 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 support for partial sophistication, and can categorize consumers into different types (in the tradition of O'Donoghue and Rabin (1999, 2001)).

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 users of their low usage levels. Other papers focus on firm's personalization and targeting policies. For example, Yoganarasimhan et al. (2021) 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 user 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.

2 Model

2.1 Inertia

Before specifying the consumer problem we 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 $t+1$.

Two main mechanisms explain inertia. Firstly, inertia is as an outcome of cost-benefit analysis that

[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/`)

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" behavior (e.g. Brot-Goldberg et al. (2021); **?**). 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 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 refusal to make 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 will allow for potentially incorrect beliefs about these parameters. We denote perceived parameters with $\tilde{\cdot}$.

2.2 Setting

We consider the following 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. v_i might be unknown before becoming a subscriber. 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 user); *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 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 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

[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).

per-period probability $\tilde{\phi}_i$, then their *expected* 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 $ti\tilde{l}de$ notations, indicating the perceived values of these future inaction probability and unsubscription cost rather than the actual ones. In contrast to the expected 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.

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 print 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 users visited our publisher's digital platform.

The content on the digital platform is classified into three parts. One part is "always free" to any user. 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 user 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 users 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 users and cookies, and is in line with the European General Data Protection Regulation (GDPR). A user 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 user 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 user may have multiple cookies if they access the website from multiple devices.

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 term (lasting up to one month) promotional contracts, such as our experimental contracts described below, which are offered to new users 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) 1-year contracts, which are rare.

Canceling subscriptions Users 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 user continues to have access until then. A user 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.[4]

4 Experiment Design

The field experiment was motivated by our research questions and the publisher's desire to convert most new users into subscribers of the digital platform via randomized control trials. The experiment was conducted in three phases from April to August 2018, with followup 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.

[4]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).

Table 1: Experimental offers

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

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 Appendix Table A.1. After the trial period, every user, irrespective of the experimental assignment, has the option to pay the regular amount of €19.99 for the next two months, and €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 user who takes an auto-renewal promo contract becomes a regular paid subscriber after the trial period is over, by default, unless the user explicitly terminates the subscription. On the other hand, a user who takes an auto-cancel offer does not become a paid subscriber by default. Instead, the user 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 user with an aim of reinstating the subscription.[5] In each of these methods, the user verifies her pre-entered payment information and confirms the subscription contract.

2. Duration: The second factor is the duration of the experimental offer, which is either 2 weeks or 4 weeks.

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

The eight combinations of these factors and the corresponding experimental group name are displayed in Table 1. Due to a technical error, users in experimental group G were not required to enter their payment information leading to an invalid experimental condition in experimental phases 1 and 2. This was corrected in experimental phase 3 leading to a full orthogonal experimental design for that phase. We will consider this fact when discussing our results.

[5]Approximately 5 days before the end of the trial offer, an email with a renewal prompt is sent to the user, and a restart of the subscription can be initiated with a click on this email. If a user 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 user's standpoint, the experimental offer is presented as follows. Upon hitting the paywall, the user 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. After clicking on the experimental offer, all users have to go through the standard three steps in order to start the trial. First, the user is asked to register and provide an email address and choose a password. Second, the user enters her personal and payment information. Lastly, the user 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 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 these effects, 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.[6] The second data are customer relationship management (CRM) data on all subscriptions and contracts, both experimental and regular contracts, from April 2018 to April 2020.

5.1 Usage Data

The browsing history includes each page visited by a reader (identified by a cookie) and timestamp. Other variables are the page type (open, metered, paywalled) and their subscriber identifier if the user was logged in to her account (even if not a paying account) at the time of browsing.[7] Most importantly, another variable shows if a reader was exposed to one of the experimental offers on a page visit and to which offer.

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 Figure 1. To keep an intent-to-treat design valid, we define the duration of the different periods in reference to that first *exposure* date rather than the actual take-up date (if one exists). For

[6] While all users are tracked for 6 weeks, 14% of users (291,837) are tracked up to 23 weeks.

[7] Some users become subscribers during the time window, while others had a subscription before and are thus identified in the system. However, if a user only subscribed for the first time outside of the usage time window after their exposure, we will not know to link that subscription to the user.

Table 2: Summary Statistics

	Raw	Main	Takers
Number of users	4,131,277	2,092,846	5,914
Number of subscribers	36,816	16,339	5,914
Total revenue (in €)	1,998,352	1,331,719	218,513
Number of pages viewed	143,628,050	89,177,586	3,128,991
- open	123,081,315	76,758,421	2,803,811
- paywalled	14,545,384	8,177,812	201,019
- metered	6,001,351	4,241,353	124,161

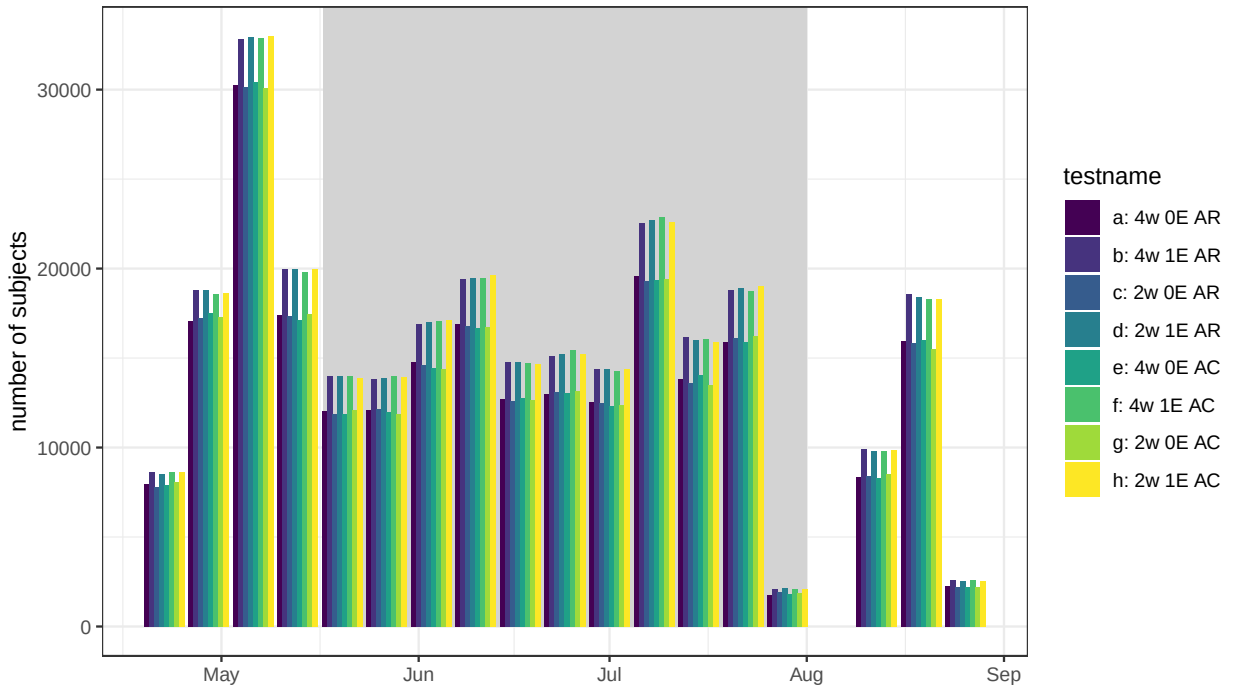
Notes: The table shows summary statistics for the raw data from the publisher, the main data of the users exposed to the experimental offers, and the subset of those who took one of the experimental offers.

example, if a reader in a 2-weeks promo treatment arm saw an offer on April 1st and took that offer on April 8th, the promo period for analysis purposes is 4/1-14 and not 4/8-21. Second, 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 date and type of first experimental exposure, and a unique subscriber identifier if they ever subscribed to the newspaper. We call that the *assignment data*. There are 2,092,861 readers in the experiment of which 26,196 (1.25%) have a subscriber identifier, and 16,339 of those has a subscription in the two weeks before, or any time after the experiment. Table 2 shows the number of users, subscribers, pages visited, and total revenue. These are presented for the raw data transfered to us from the publisher, the main data used for analysis (for participants in the experiment), and a subsample of participants who subscribed to any of the experimental offers.

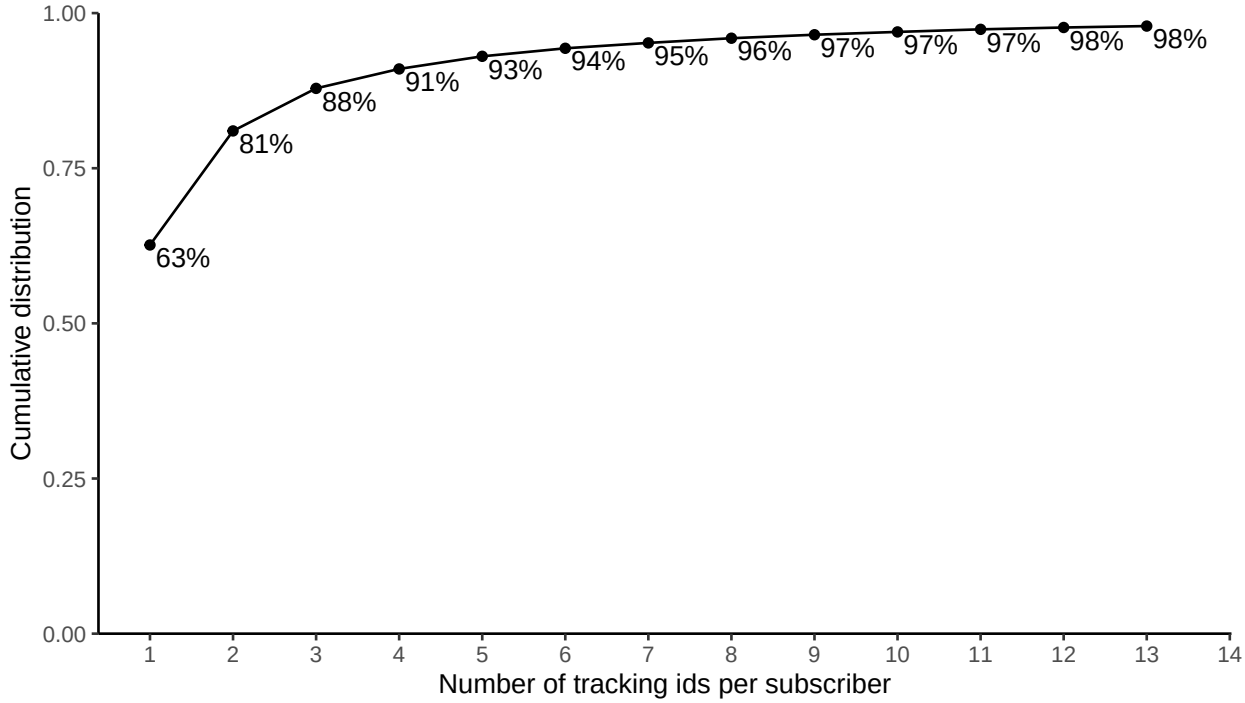
Note the common challenge in the digital world, that cookies are not people. However, we know something about the extent of the issue in our setting, and argue that it might shift the effects levels, but not in relative terms. Figure 2 shows the distribution of the number of cookies associated with each 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 users and will contribute "no subscription" and their usage to multiple treatment arms; if instead they did subscribe, then we associate all their devices 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 compressed subscription shares. However, they do not bias our results since we analyze them in relative terms.

Figure 1: Number of readers in each treatment arm by week



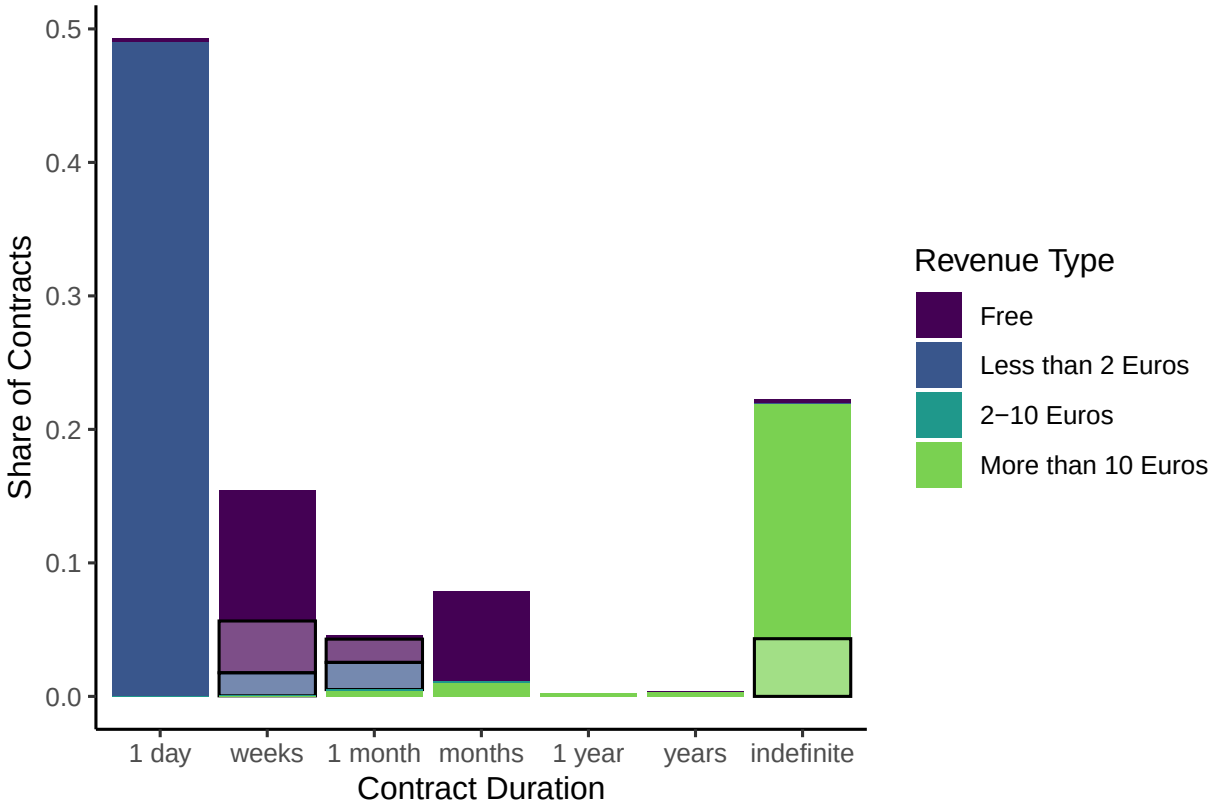
Notes: The figure shows the number of users exposed to each expereimental contract by week. The difference in shading represents the different phases of the experiment.

Figure 2: Distribution of unique number of cookies per subscriber



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

Figure 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 2 weeks or a month (4 weeks), and are either free or less than €2; the auto-renewal contracts are indefinite with a revenue above €10.

5.2 Subscription Data

The second dataset is the company'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[8]. 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. Figure 3 shows the distribution of contracts taken by the 16,339 experiment participants who subscribed at any point during the period (another 9,857 had a small 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 results, and we make it clear when we use as an outcome subscription for *any* contract, experimental or not, or focus on takers of experimental contracts only.

[8]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 users, 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 comparing 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 the user subscription rate, that is, the proportion of days subscribed to the platform through any contract within a period; the user subscription extensive margin, that is, if the user was an active subscriber within that period at all; the revenue attributed to that period; 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, the two months of €19.99 price per month, the following two months of the full price of €34.99, and then another three periods of six-months each.

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

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

where y_i represents one of the outcome measures of individual i's subscription, and Auto-renew$_i$, One-euro$_i$, Four-week$_i$ 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 1 and 2 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 3 data which has all eight groups. 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 4 plots the intent-to-treat effect of offering a promotional auto-renewal contract as opposed to an auto-cancel contract on subscription behavior at various time periods, which is the estimated coefficient β_1 in equation (1).

Figure 4a shows effects on subscription rates. As expected, the subscription rates prior to the experiment are similar across the experimental groups, so the estimate in the first time bucket is small and indistinguishable from 0. During the promotional time period, we observe a significant negative impact of auto-renewal on subscription rates, which is 28% lower relative to the auto-cancel. 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 post promotion. Subsequently, we observe a significant negative trend in the effect and eventually, about a year and a half post promo, the subscription rate is higher for the auto-cancel group. The effect on revenue shows a similar pattern, as seen in Figure A.1 in the appendix, although the revenue effect estimates are less precise.

Comparing these patterns against those in Figure 4b, we note a different pattern on the extensive margin. We do see a negative effect of offering an auto-renewal contract on the likelihood of a reader becoming a subscriber in the promo period. However, we see no positive effect post promo. Fewer readers become subscribers when they receive an auto-renewal offer relative to an auto-cancel in any time period. Overall, we see a significant 9% drop in subscribers over the entire data time period because of auto-renewal.

6.2 Other Experimental Factors

Free vs. €0.99 Figure 5 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 9% fewer readers to subscribe while the promotion was active. As expected, users 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 in any time bucket 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.

4 Weeks vs. 2 Weeks Figure 6 shows a similar pattern of the effect of increasing the trial duration. The estimates show that increasing the trial duration from 2 weeks to 4 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 time.

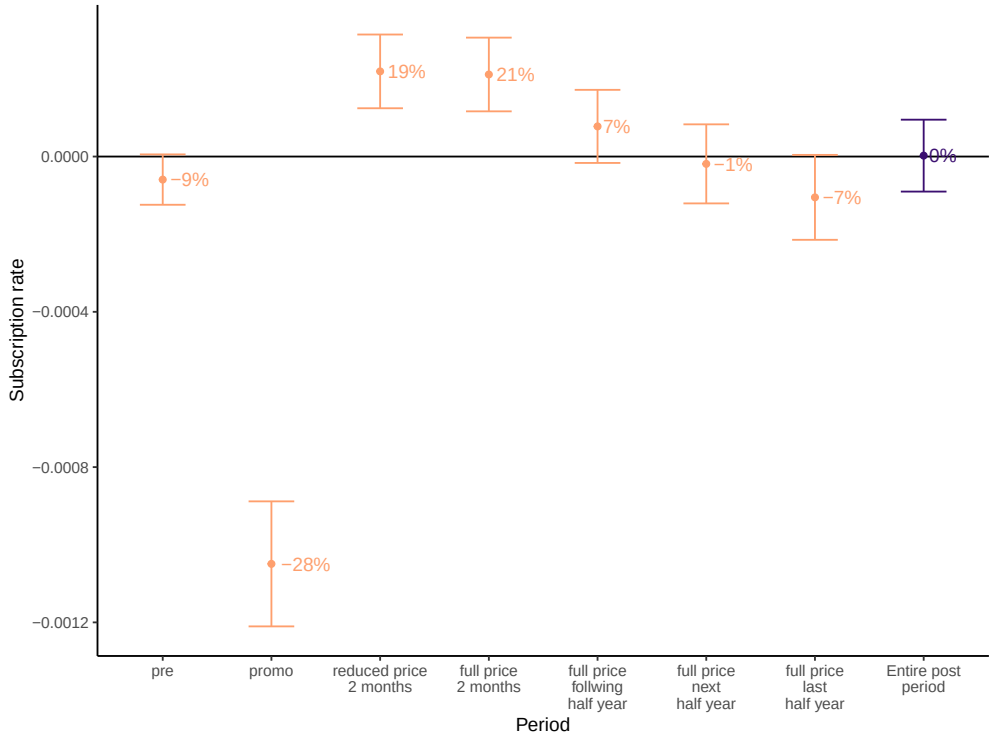
Comparing the auto-renewal vs. auto-cancel effect with the same effect of price or duration change shows the distinct consumer response to auto-renewal. While auto-renewal causes an average decline in promo take-up, similar to a price increase, it causes an increase in subscription rates few a few months after the promo, which are absent in response to price increase. At the same time, we see an overall decrease in post promo subscribers due to auto-renewal, which is also absent in response to price increase. These patterns indicate a unique consumer 'push back' to auto-renewal relative to more transparent factors such as price and trial duration.

7 Channels of inertia

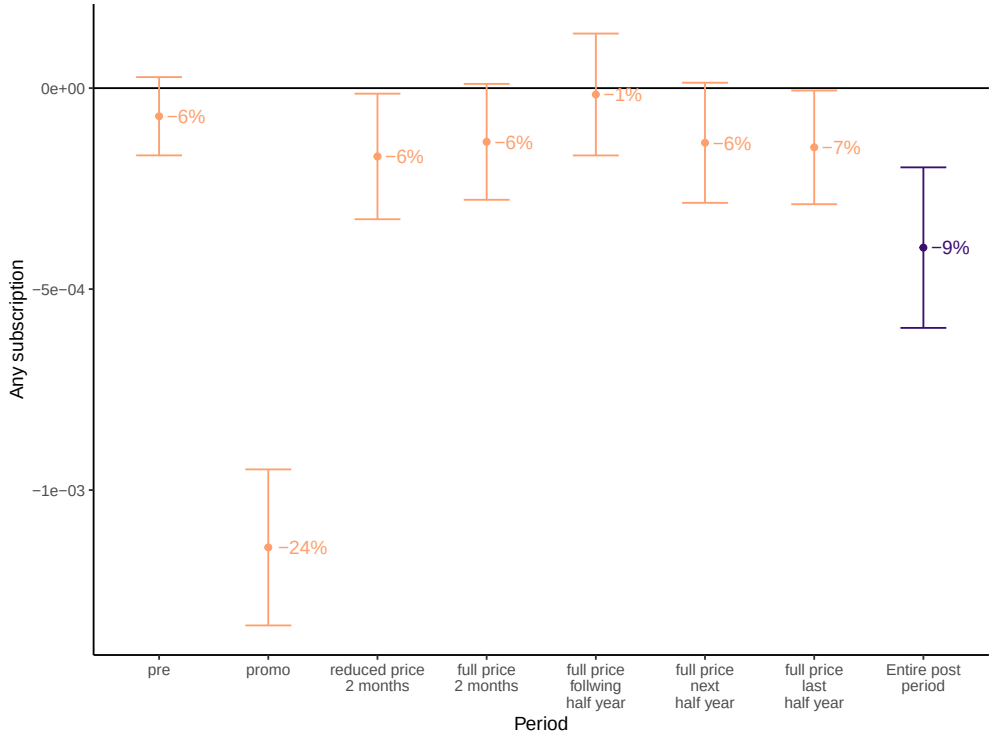
s The experiment have shown several patterns in the data. Namely, that there is lower take-up of the auto-renewal offers, implying that consumers are aware and respond to future inertia; that even with lower initial take-up, more auto-renewal subscribers stay subscribed and pay a full price in the first few months; and lastly, that there are fewer subscribers to any contract in the auto-renewal treatments in the two years post-promo.

In this section we provide further evidence on the channels of inertia, to support our modeling choices. First and foremost, we argue that auto-renewal subscribers are not using their subscription. Further, we show that treatment arms that brought more users to try the subscription did not lead to more long-run subscriptions. Together, these findings show that preferences might be uncertain yet static, with no evidence

Figure 4: Effect of Auto-renewal relative to Auto-cancel contracts on overall Subscription behavior



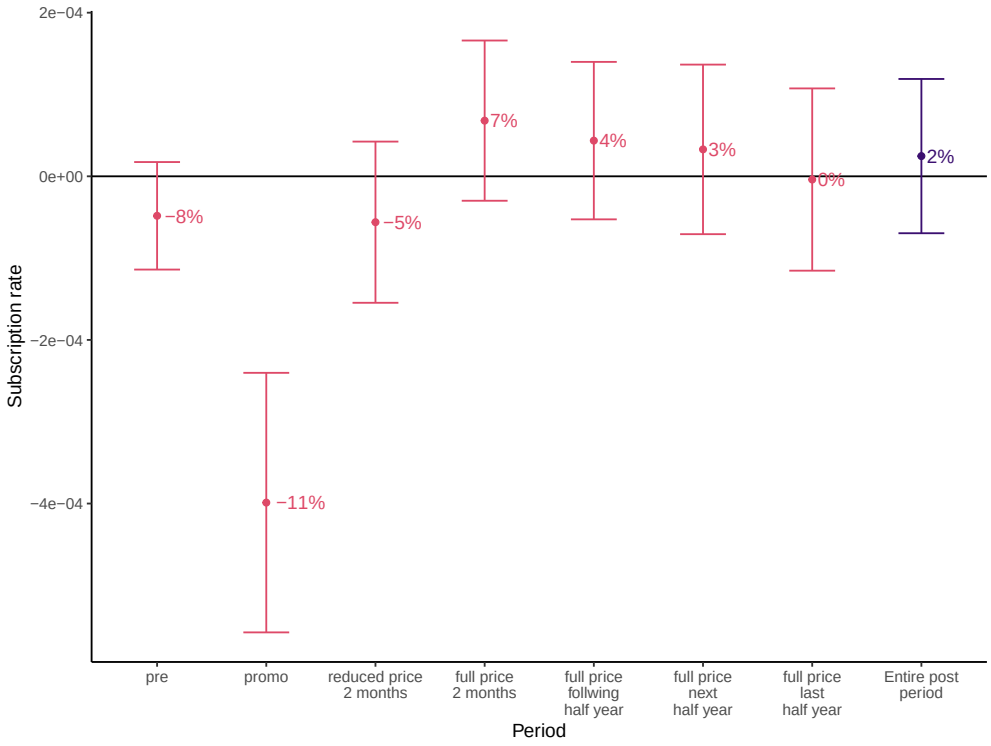
(a) Effect on Subscription rate (proportion of days an individual subscribed)



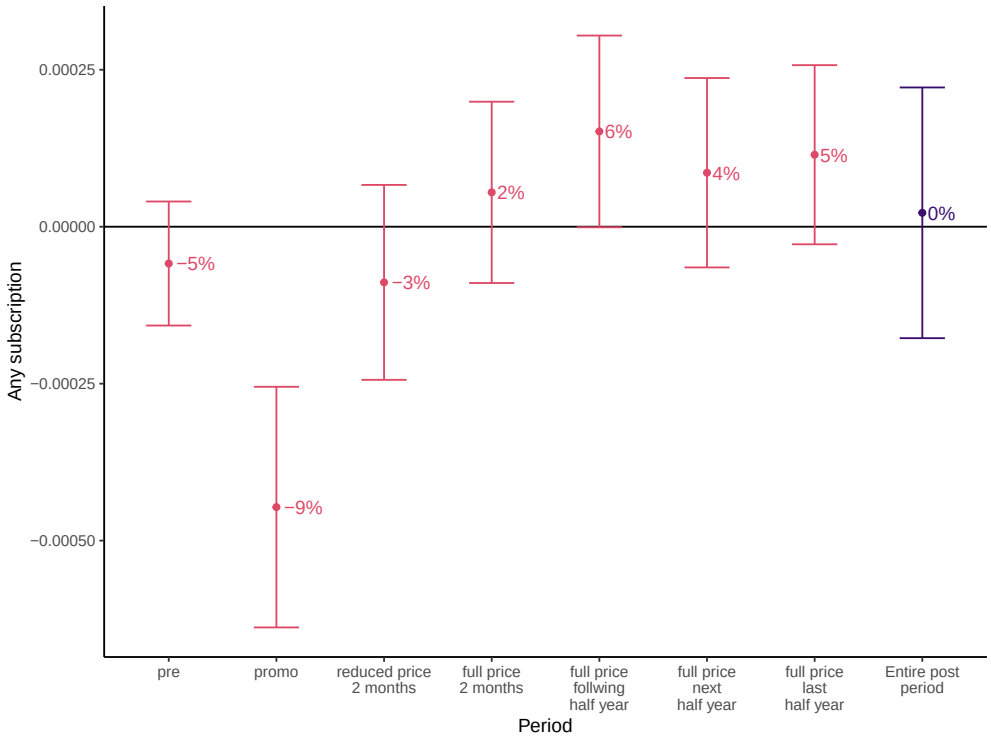
(b) Effect at the Extensive margin (whether the individual subscribed at all)

Notes: The figures plot the estimated average intent-to-treat effect of offering an Auto-renewal relative to an Auto-cancel contract on consumer subscription behavior. Specifically, we plot the estimated coefficient β_1 from equation (1) for various time periods. "pre" refers to time before the experiment started; "promo" is the during the promotional time period, the last bucket "Entire post promo" aggregates across all post 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 5: Effect of €0.99 relative to free promotional contracts on overall Subscription behavior



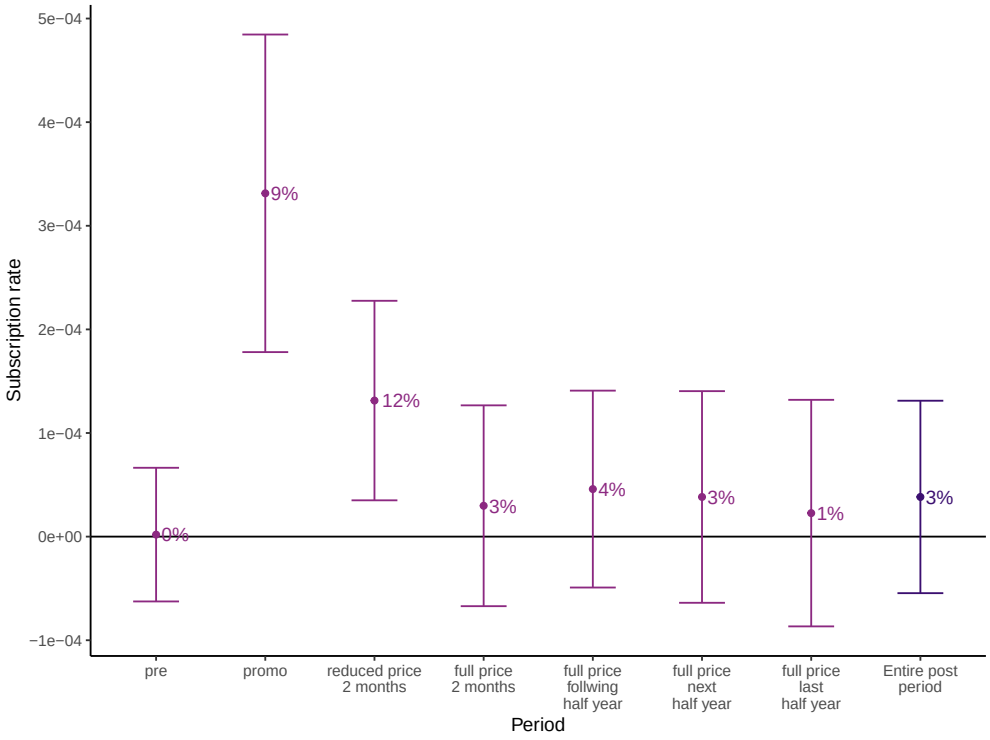
(a) Effect on Subscription rate (proportion of days an individual subscribed)



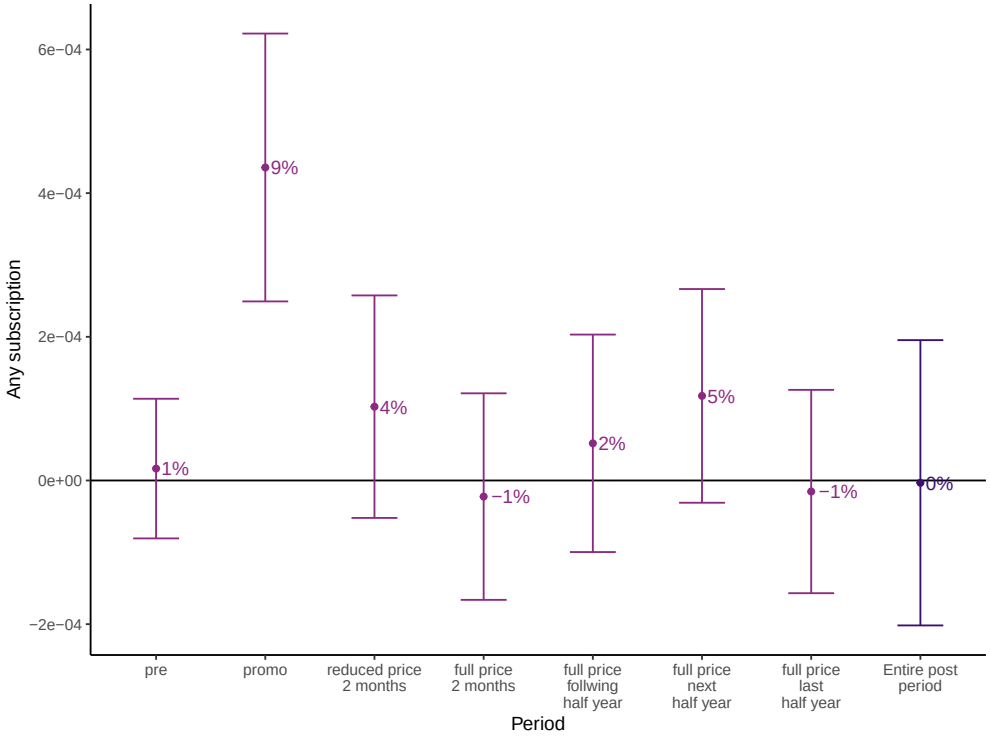
(b) Effect at the Extensive margin (whether the individual subscribed at all)

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 consumer subscription behavior. Specifically, we plot the estimated coefficient β_2 from equation (1) for various time periods. "pre" refers to time before the experiment started; "promo" is the during the promotional time period, the last bucket "Entire post promo" aggregates across all post 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 6: Effect of 4 week relative to 2 week promotional contracts on overall Subscription behavior



(a) Effect on Subscription rate (proportion of days an individual subscribed)



(b) Effect at the Extensive margin (whether the individual subscribed at all)

Notes: The figures plot the estimated average intent-to-treat effect of serving a 4 week vs. 2 week promotional contract on consumer subscription behavior. Specifically, we plot the estimated coefficient β_3 from equation (1) for various time periods. "pre" refers to time before the experiment started; "promo" is the during the promotional time period, the last bucket "Entire post promo" aggregates across all post 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.

for learning or habit formation in our context. 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 caused by auto-renewal is actually unwanted, and caused by inertia, we expect users to get little utility from their post promo subscription. We use the website usage data to gauge the utility people 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 users 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 non-subscribers.

Recall that our usage data spans 6 weeks for each user; notating day 0 as the day a user received the experimental offer, our usage data spans days -13 to 27. Figure 7 plots the average page visit and subscription rates data among promo takers for each day in this time span. For this plot, we use data for individuals who took either a 2 week, €0.99 auto-renewal promo or a 2 week, €0.99 auto-cancel promo during the first days after exposure, so we can observe the promo time ending in the middle of our 4 weeks post treatment usage data.

Figure 7 shows that auto-renewal promo takers are far more likely to be subscribed after the 2 week promo time, relative to auto-cancel subscribers. However, we do not see any difference in their website visits. 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, post promo take up.

Table 3 shows promo and post-promo usage statistics averaged across users who took either a 2 week €0.99 auto-renewal or auto-cancel promo. The sample is grouped by whether the user was also a subscriber post promo, or not. The analysis shows that more than half of the users who subscribed in the two weeks post 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 post promo and is significantly lower than those who subscribed post promo after taking an auto-cancel offer.

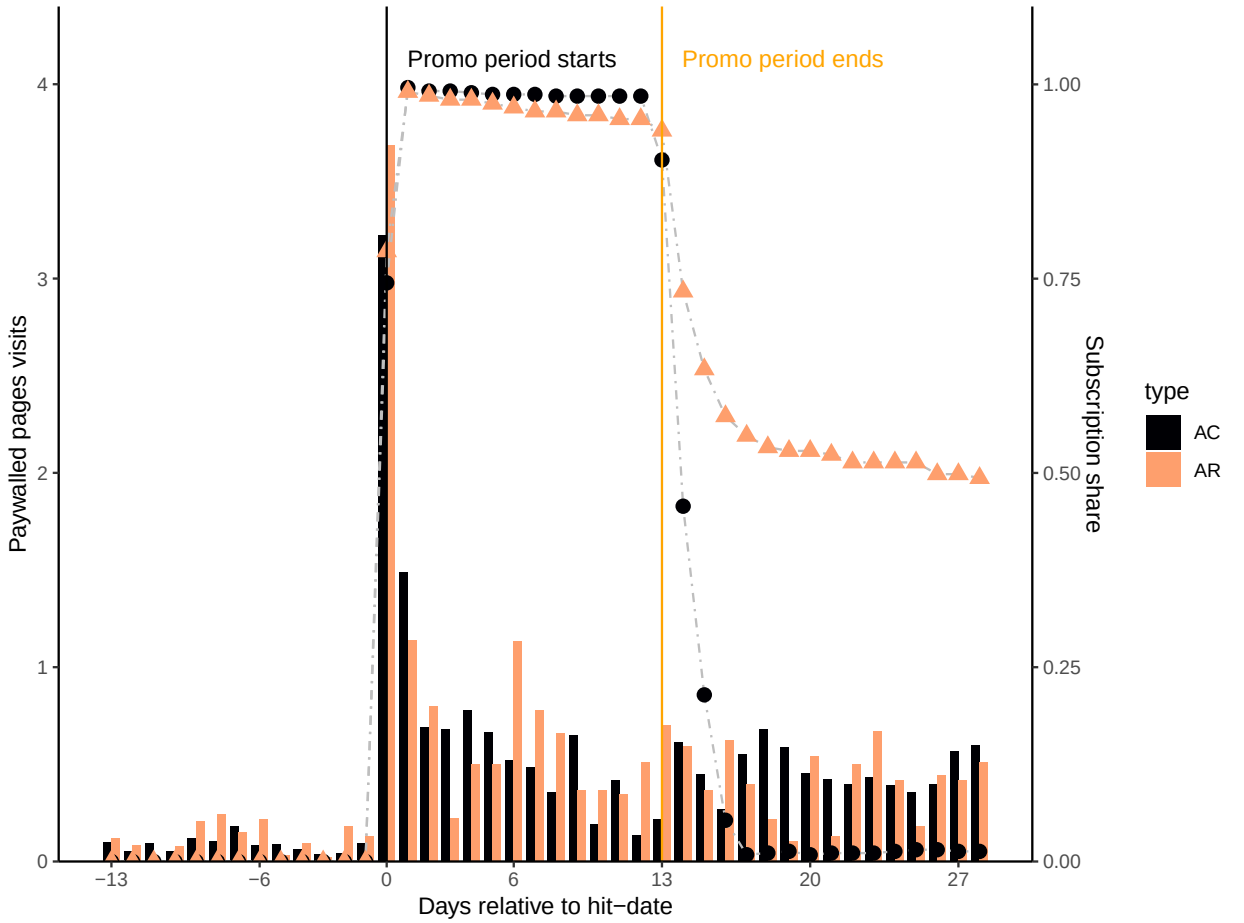
Overall, this analysis is consistent with our inference that the users who continue subscribing after taking an auto-renewal promo do not derive higher utility than those who do not. Meaning, the valuation of subscription does not grow for auto-renewal subscribers thanks to their subscription status.

7.2 Lack of Learning and Habit Formation from the Promo Trial

Reducing promotional price or increasing the trial duration increases the users' 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 5 and 6. 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 habit formation and the evidence above on the lack of usage for auto-renewal subscribers, means that consumers 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.

Figure 7: Subscription vs. Platform usage for 2 weeks, €0.99 auto-renewal promo takers vs. 2 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 2 week €0.99 auto-renewal promo and those who took the 2 week €0.99 auto-renewal promo. The time on the x-axis starts 2 weeks before the experimental offer was given to the user and covers the promotional 2 weeks and 2 weeks after that.

Table 3: Usage during vs. post promo for 2 week promo takers

	Auto-renewal 2 week €.99 subscribers		Auto-Cancel 2 week €.99 subscribers	
	Subscribed in two weeks post promo	Not subscribed in two weeks post promo	Subscribed in two weeks post promo	Not subscribed in two weeks post promo
	estimate (s.e.)	estimate (s.e.)	estimate (s.e.)	estimate (s.e.)
Promo 2 weeks: Avg. page visits	25.37 (3.66)	22.28 (2.64)	57.71 (7.10)	23.66 (1.86)
Promo 2 weeks: % users with any visit	0.78 (0.04)	0.68 (0.04)	0.90 (0.04)	0.68 (0.02)
Post promo 2 weeks: Avg. page visits	20.33 (3.57)	12.88 (2.92)	75.98 (10.71)	12.67 (1.24)
Post promo 2 weeks: % users 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 users who took the 2 week, €.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 users 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 post promo.*

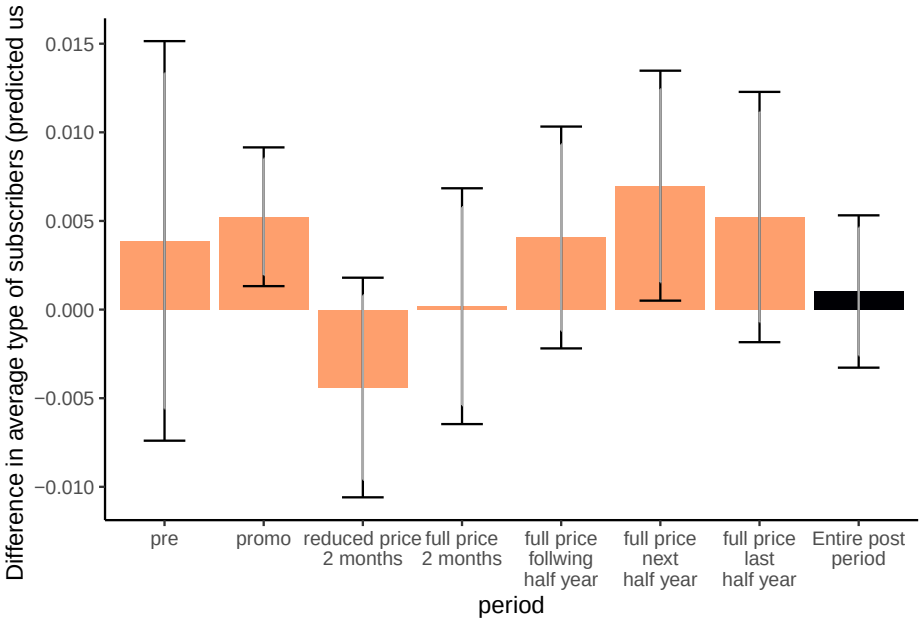
7.3 Evidence for "spite"

This final subsection studying the channels behind our results leverages a proxy for user valuation in order to understand what is driving subscription behavior. Average subscriber types at different periods, in the sense of the subscription value, speak to the mechanism driving our estimated effect of serving an auto-renewal vs auto-cancel contract. Under our model, consumers 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 users are being driven by switching costs (standard or coupled with present-bias) or random attention as described above, we should expect the types to converge from below in the long run.

We find support for the former two predictions but not for the latter. We use the pre-experimental usage data to predict post treatment usage for each user in our data. We use this predicted usage as a proxy for the user's type – those who are predicted to use the newspaper more are higher types. Figure 8 shows the by-period difference in user types between auto-renewal and auto-cancellation subscribers. The promo period subscribers in the auto-renewal group are of higher types relative to the promo period subscribers in the auto-cancel group. In the initial periods after the promo ends, lower types subscribe in the auto-renewal group (yet not significant at the 10% level), but the difference then flips sign and becomes larger again in the long-run.

These findings show that, in the long-run, users who choose to subscribe 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 term; unlike the promotional period where the auto-cancellation contract has a different continuation value, after the promo period all contracts are identical. This pattern is consistent with a psychological cost or spite against the newspaper due to the initial auto-renewal offer. This finding is consistent with the extensive

Figure 8: Difference in user types between subscribers in the auto-renewal vs auto-cancellation group



Notes: We use the pre-experimental usage data to predict post treatment newspaper usage, which we use as a proxy for user type. The figure shows the difference in average user type between subscribers in the auto-renewal group and auto-cancellation group by period. Error bars are 95% confidence intervals, and the light gray sections represent 90% confidence intervals. For example, during the promo period the subscribers in the auto-renewal group are significantly higher types compared to the subscribers in the auto-cancellation group. Standard errors are clustered at the individual level.

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.

8 Quantification of Inertia

Our results show that inertia exists and is predicted. Existence is manifested by higher retention of auto-renew takers; it is predicted, as manifested by lower take-up of auto-renew offers. In this section we turn to quantify the degrees of actual 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 users who take up an auto-renewal subscription. We look at the retention of auto-renewal promo takers in the post 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 consumers would have been subscribed if not for inertia. However, the set of consumers 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 be the set of types (combination of valuations and predicted inertia) who takes the auto-renewal promo, and C the types who takes the auto-cancellation promo. We assume $R \subset C$, 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, because the auto-cancel promo offers access to the same content without the risk of an unwanted paid subscription.[9] The evidence in section 7, showing that auto-renew takers do not use their subscription, and the lack of learning (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 taking 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 users 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 actual 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).[10]

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 post promo. This is the difference in post-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 individuals in R when the contract changes from auto-renew to auto-cancel, we need to (1) estimate their post promo subscription when they are offered an auto-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 are 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 post-promo subscription y of the consumers 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

[9] As discussed in section 4, users 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.

[10] 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.

(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 consumer in the auto-renew arm. However, as will be clearer in the next section quantifying predicted inertia, that 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 $C \setminus R$ are not 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 post-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 as 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. Appendix Table A.2 shows estimates from this exercise. The fifth column shows our final estimates.

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. 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 $\phi = 1$, will remain subscribe. 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:

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

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 2 due to the jump in costs. 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

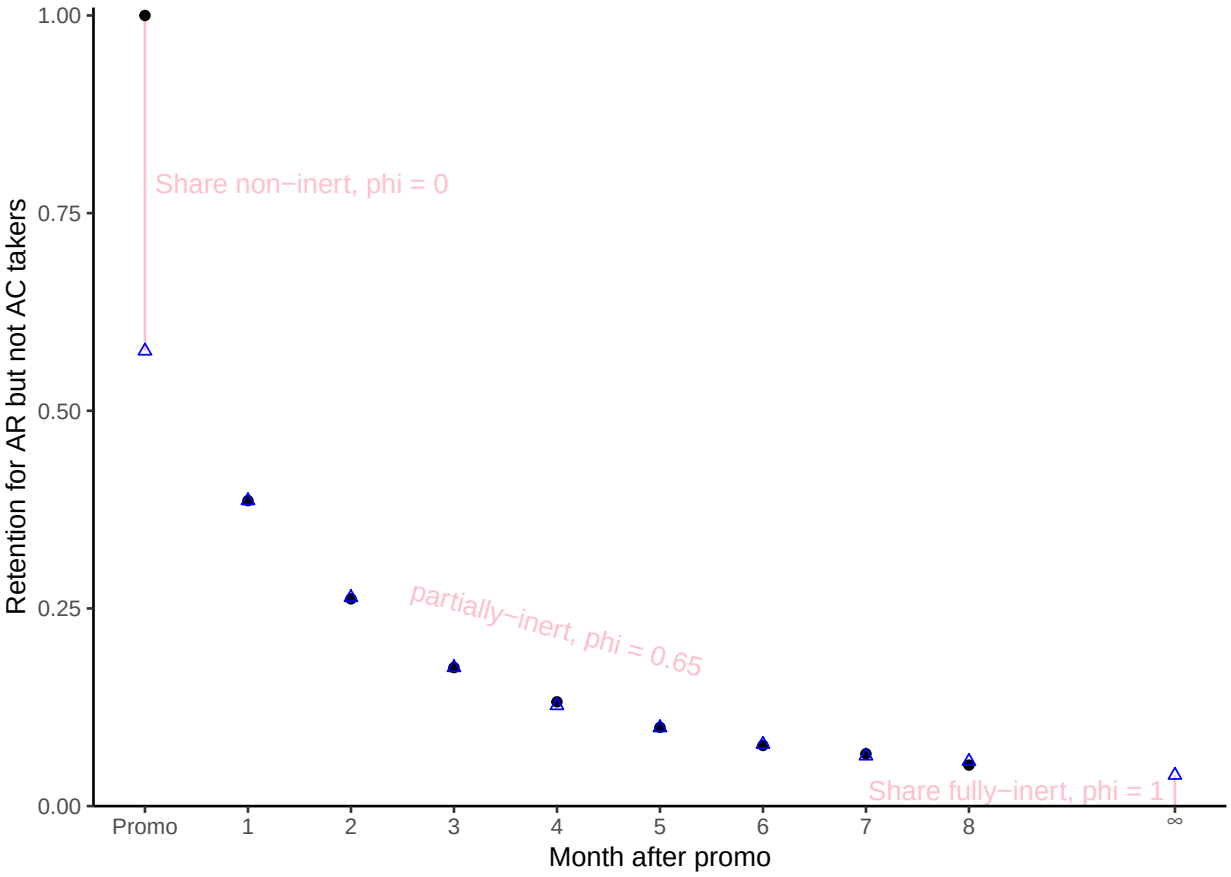


Figure 9: Data and estimation of three-types inert consumers

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 8 months following. The triangles are the predictions from the model in (2) estimated on this data. The prediction at "infinity" are the 3.9% fully inert auto-renewal takers who will not become long-run subscribers under auto-cancellation. The difference between 1 and the triangle at "promo" gives the 42.4% of non-inert subscribers who cancelled their subscription before starting to pay full prices. The remaining 53.5% are estimated to have an inaction parameter of $\phi = 0.647$.

share leaving on month 3. We predict 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.[11] 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 9 shows the results of this exercise. the dark circles are the estimated excess share of subscribers in each month post-promo, and the triangles are the predicted shares from the estimation. These are only 8 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, and the difference from 1 is the share of the non-inert; the triangle at ∞ is the estimated share of fully inert auto-renew takers $\hat{\pi}_f^R$.

We estimate that $\hat{\pi}_n^R = 42.4\%$ $(se = 2.8\%)$ 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 unsubscription costs. In contrast, there are $\hat{\pi}_f^R = 3.9\%$ (1.0%) fully-inert consumers with either $\phi = 1$ or

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

very high costs. Finally, the remaining $\hat{\pi}_p^R = 53.7\%$ (2.3%) are partially inert, with an estimated inaction of $\phi = 0.647$ (0.024) and very low costs.

We argue that the partially inert have very low unsubscription costs for the following reason. At month 3, 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 0.8% (se = 0.8%, which are 1.5% 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.

We do a similar exercise in Appendix A.2 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 $e0.99 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 actual 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 actual 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: Inertia, Sophistication Types, and Subscription Behavior under Auto-Renewal

	Non-inert ($\phi = 0$, $c^u = 0$)	**Partially-inert** ($\phi = 0.65$, $c^u = 0$)	**Fully-inert** ($\phi = 1$ or c^u is large)
Naïve ($\tilde{\phi}$=$\tilde{c}^u$=0)	Take AR promo (equivalent to AC)	Take AR promo	Take AR promo
Sophisticated ($\tilde{\phi} = \phi$, $\tilde{c}^u = c^u$)		Take only if value balances inertia risk	Take only if wants to subscribe in the long-run

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.

Within each cell we describe the auto-renew promo take-up behavior based on the type, compared to 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 them, 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 – 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.[12] The other remaining type are the sophisticated partially-inert consumers. These types know they might be subscribed for long even if not forever. Therefore, their net-present-value takes into account these higher price periods of possible subscription.

We can use the model to generate the selection into auto-renewal take-up. In the auto-renew condition for each inertial type i, there is a marginal valuation v 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.[13] 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. We will assume different values of δ, and the results are not very sensitive to their choice. We use the data to calibrate 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. To get that value, we use our 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 2 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.

We solve the above indifference condition in 3 and find that for the partially-inert with $\tilde{\phi} = 0.65$ 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

[12] 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.

[13] 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.

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 above 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 $\pi_n^R = \pi_n \cdot \frac{s^{AC} - s^{LR}}{s^{AR} - s^{LR}}$ (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 $\pi_f^R = \pi_f \cdot (1 - \pi^s) \frac{s^{AC} - s^{LR}}{s^{AR} - s^{LR}}$.

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 we do observe their sample equivalents – $F(0.83)$, $F(34.35)$, and $F(1.82)$[14] we get $a = 0.008$. However, we also just vary a to take values between $\{0.001, 0.01, 0.05, 0.1, 0.2, 0.3\}$ 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 actual 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. We find that 53.9% (se = 2.1%) of inert subscribers are sophisticated and aware of their inertia parameters. In the population, roughly 25.5% (1.7%) are non-inert, 5.1% (1.2%) are fully inert, and 69.3% (1.6%) are partially inert with average inaction of 0.647. The shares among auto-renew takers were presented in section 8.1. If we vary a 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 sophisticated inerts and shares of partially-inert in the population (and the complementary fully-inert), shown in Appendix Figure A.3. The share of partially-inert moves between 65.8%-69.4% (and fully-inert hence between 5%-8.6%). The shares of sophisticates among the inerts is more sensitive, and moves between 53.5%-72.7%. To sum, the share of exploitable naifs is thus at most 35%, and might be as little as 20%.[15]

8.3 Targetability based on Sophistication

Is it possible to target users with offers based on their sophistication? Is user sophistication predictable? We use our pre-experimental usage data—which includes the topics the users browse—to predict the heterogeneity in the effects of our treatments. While we find some predictors of higher usage, these do not

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

[15] These numbers are the shares of naifs among partially and fully inert consumers.

Table 5: Estimated shares of inertial types

	Estimate	SE
In population:		
Share sophisticates among inert	0.539	(0.021)
Inertia (of partially inert)	0.647	(0.024)
Share non inert	0.255	(0.017)
Share fully inert	0.051	(0.012)
Share partially inert	0.693	(0.016)
Among auto-renew promo takers:		
Share non inert among AR takers	0.424	(0.028)
Share fully inert among AR takers	0.039	(0.010)
Share partially inert among AR takers	0.537	(0.023)
Share responding to price increase among AR takers	0.008	(0.008)

Notes: The table shows the estimated shares of types in the population, and among auto-renew promotional offer takers.

translate to differential treatment effects. Since we do not find predictable differential treatment effects, our data is unable to predict sophistication. This is exhibited in Figure 10. The figure shows the auto-renewal vs. auto-cancellation treatment effects based on type categorization into "high" and "low." There is no predictable heterogeneity between the types.

To categorize types, we use data on usage each day before the experiment to predict if a user attempted to read any paywalled or metered article during 1 to 4 weeks after entering the experiment.[16] Specifically, as predictive variables, we use the number of always open, metered, and paywalled pages on each of the 13 to 1 days before hitting the paywall (total of 39 continuous variables) as well as the type of paywall users hit eventually when they entered the experiment (metered or always-paywalled content), as inputs to a Random Forest algorithm, cross-validated with a 10-fold procedure. We then reran the main analysis fully interacting the treatment with these predicted types.

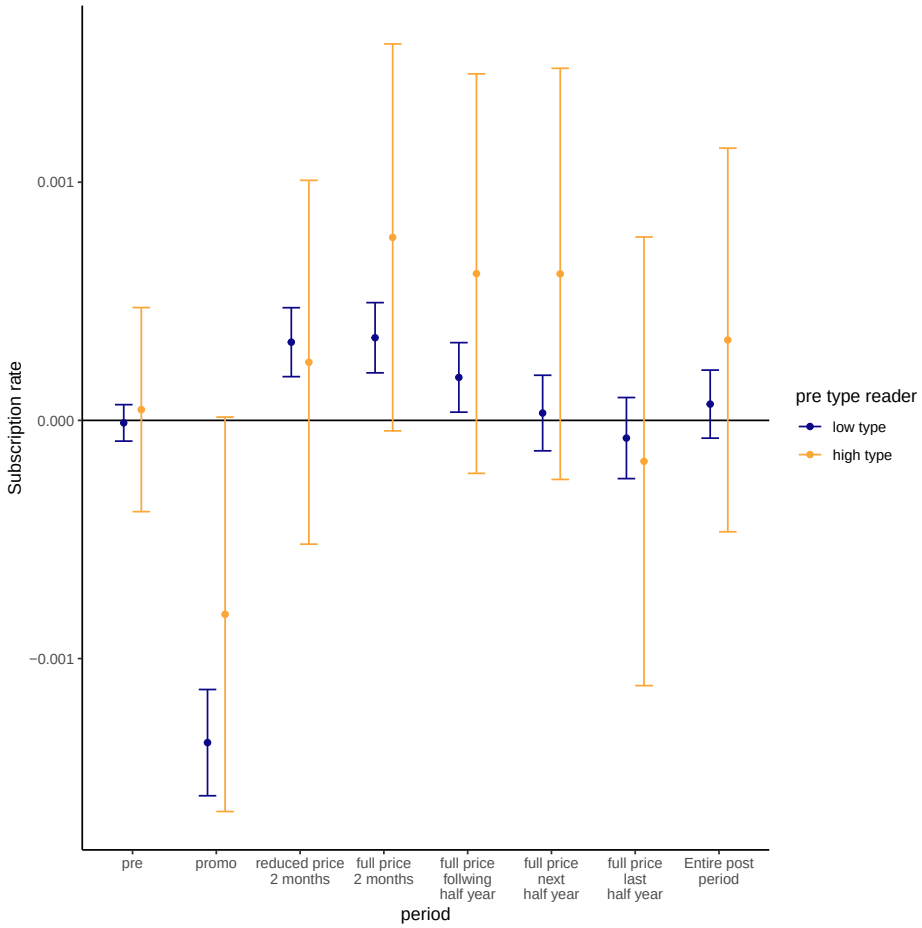
We also investigate the ability to identify observed characteristics that would predict consumer preference for the auto-renewal option. Presumably, those who would like to be long-run subscribers would rather subscribe only once, rather than having to renew once the promo period ends. However, we are not able to find a user segment that would prefer auto-renewal over auto-cancel.

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 consumers are highly inert. However, this body of knowledge relies on a selected sample of *already existing customers*. Our paper suggests that a large portion of customers, between 24% and 36%, is aware of its future inertia and avoids engaging with an exploitative contract. Furthermore, offering an exploitative contract pushes 9% of customers from engaging with the company in its entirety for the duration of our data. These new findings imply that consumers' awareness to their future inertia limits inertia exploitation.

[16] We tried different methods to ensure that the unpredictability we find is not due to our modeling choice. We used predicted usage levels rather than any usage, changed the time window of what is considered post-usage, and used pre-experiment usage to define types (regardless of post-usage). The results are similarly noisy and indistinguishable between the "low" and "high" types however we cut the data.

Figure 10: Auto-renewal effects by types of readers



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 5% of most engaged readers, low types are the bottom 95%.

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. By the half-year following the first 4 months, the revenue difference is statistically indistinguishable from 0, and the differences get smaller as time goes by. Furthermore, if the market share, or size of readership matters, then auto-renewal is worse from day 1. There are various reasons why readership matters, such as advertisement revenue, 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 Pareto improvement 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, targeting based on naivete seems limited in scope, as we were not able to find predictive features. 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 consumers 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.

We design a large-scale experiment that enables us to study inertia in consumer subscription decisions. The novel experiment design simultaneously varies the contract renewal terms along with other benefits, which allows us to quantify the inertia consumers 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 consumers do recognize and account for their inertia, but they underestimate it by a factor of six. About 30% of the consumers take up the auto-renewal subscriptions and end up paying for the subscriptions 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 term 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.

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.

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.

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."

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," `https://faculty.washington.edu/hemay/Free_Trial.pdf` 2021. Accessed: 2021-12-23.

A More Results

Table A.1: Balance of pre-experiment behavior

Dependent Variables :	Total Pages	Open	Paywalled
Model:	(1)	(2)	(3)
Test Group A	5.240***	5.048***	0.1923***
	(0.0406)	(0.0391)	(0.0061)
Test Group B	5.270***	5.075***	0.1948***
	(0.0416)	(0.0399)	(0.0073)
Test Group C	5.366***	5.167***	0.1994***
	(0.0467)	(0.0453)	(0.0064)
Test Group D	5.140***	4.947***	0.1926***
	(0.0342)	(0.0323)	(0.0072)
Test Group E	5.266***	5.065***	0.2011***
	(0.0434)	(0.0416)	(0.0074)
Test Group F	5.197***	5.012***	0.1853***
	(0.0384)	(0.0371)	(0.0053)
Test Group G	5.236***	5.039***	0.1972***
	(0.0405)	(0.0389)	(0.0060)
Test Group H	5.165***	4.968***	0.1968***
	(0.0369)	(0.0346)	(0.0086)
Fit statistics			
Observations	2,082,995	2,082,995	2,082,995
R^2	9.86×10^{-6}	1.01×10^{-5}	1.27×10^{-6}
Adjusted R^2	6.49×10^{-6}	6.69×10^{-6}	-2.09×10^{-6}

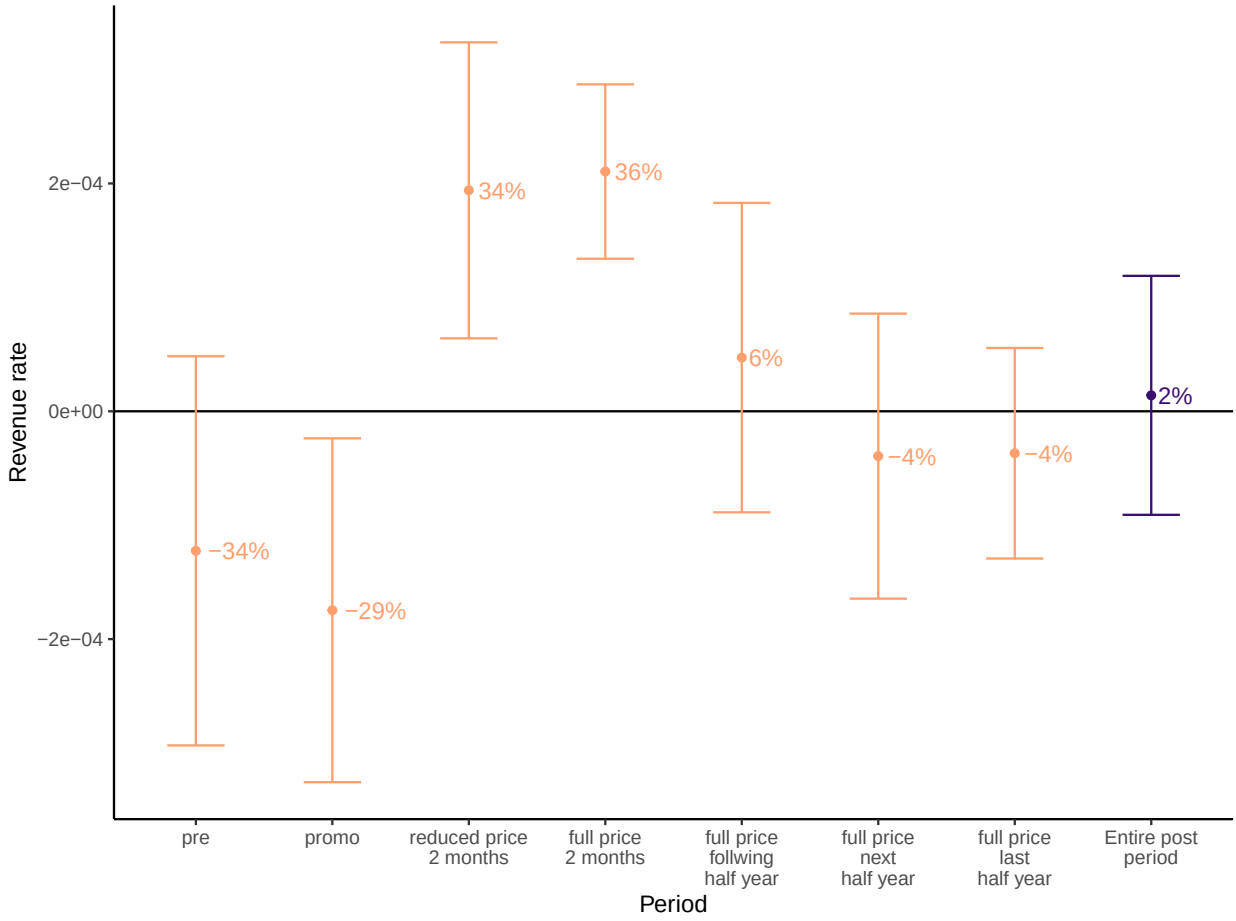
Notes: Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

A.1 Solving the problem

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

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

Figure A.1: Revenue rate when Auto-renewal contracts are served relative to Auto-cancel contracts



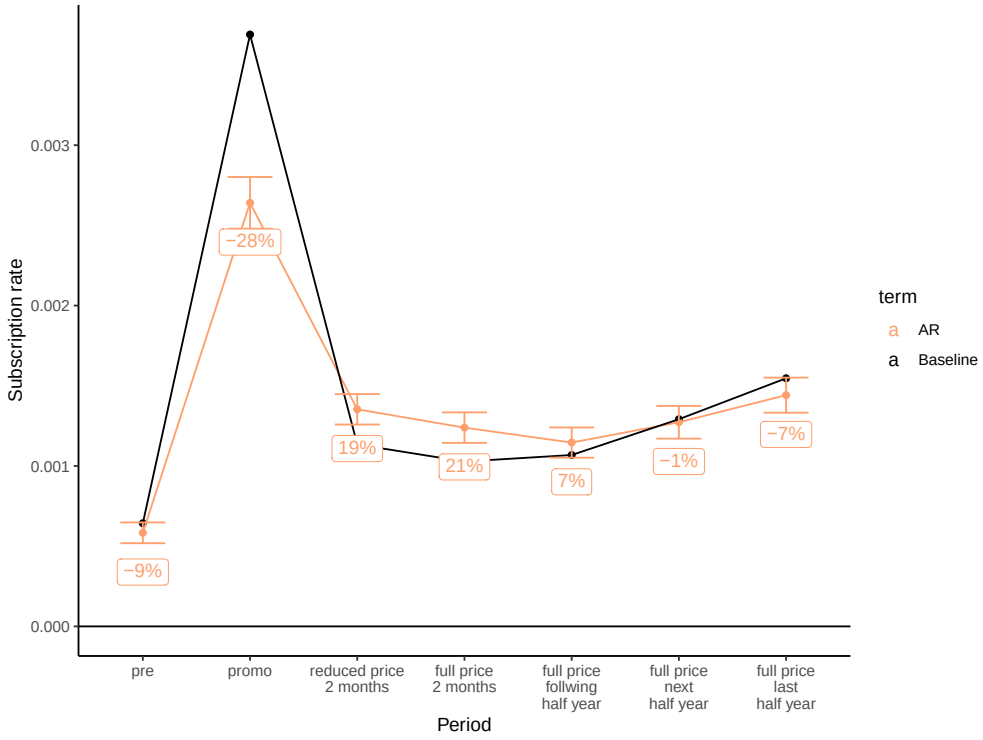
Notes: The figure plots the estimated average intent-to-treat effect of serving an Auto-renewal relative to an Auto-cancel contract on the newspaper's revenue. Specifically, we plot the estimated β_1 from equation (1) for various time periods. "pre" refers to time before the experiment started; "promo" is the during the promotional time period, the last bucket "Entire post promo" aggregates across all post promo time periods. The error bars show 95% confidence intervals of the β_1 coefficient

Table A.2: Retention of promo takers

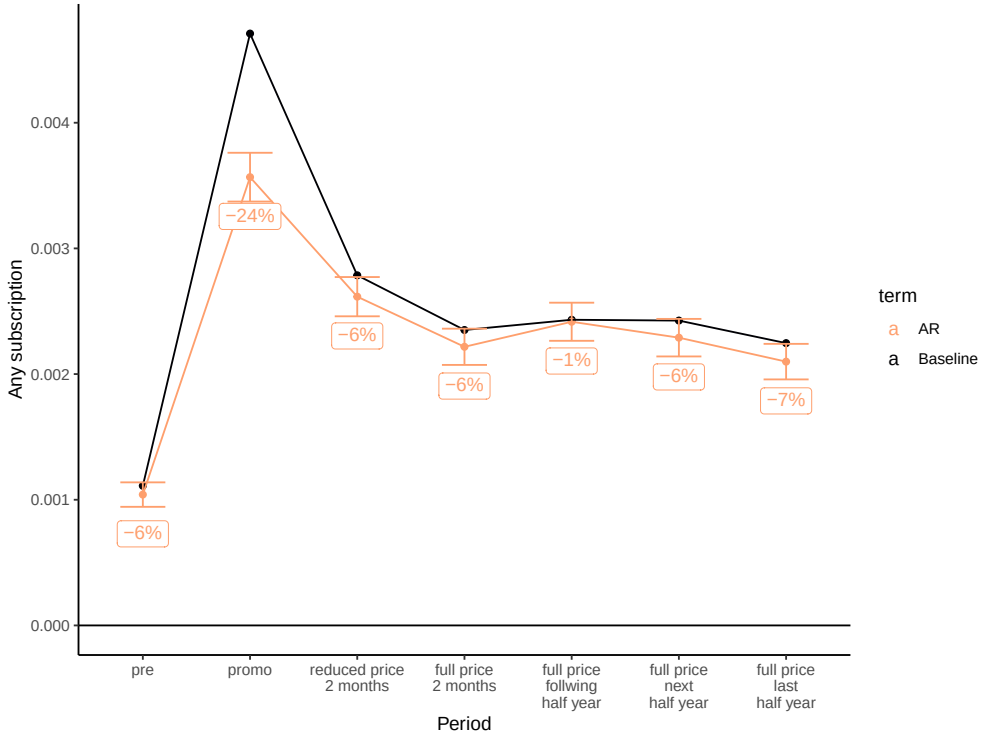
Month post promo	(1) Auto-renewal estimate (s.e.)	(2) Auto-cancel estimate (s.e.)	(1) - (2) estimate (s.e.)	(1) - scaled up (2) estimate (s.e.)
1	.491 (.014)	.063 (.006)	.428 (.016)	.388 (.017)
2	.343 (.013)	.049 (.005)	.294 (.015)	.263 (.016)
3	.264 (.012)	.055 (.006)	.208 (.014)	.173 (.016)
4	.223 (.012)	.057 (.006)	.166 (.013)	.13 (.015)
5	.188 (.011)	.055 (.006)	.133 (.012)	.098 (.014)
6	.168 (.011)	.056 (.006)	.111 (.012)	.076 (.014)
7	.155 (.01)	.055 (.006)	.1 (.012)	.065 (.014)
8	.137 (.01)	.053 (.006)	.084 (.011)	.05 (.013)

Notes: The table shows the likelihood of a promo taker subscribing to the platform (even once) in the eight months after the promo ends. Column (1) displays the proportion of promo takers in an auto-renewal experimental group who subscribed to the newspaper in a future month, and (2) does the same for those in an auto-cancel group. The fourth column displays the difference (1)-(2), and the next one does the same after scaling up column (2) by a factor of 1.64 which is the increase in the likelihood of a user taking the experimental promo going from AR to AC.

Figure A.2: Subscription Levels when Auto-renewal contracts are served relative to Auto-cancel contracts



(a) Subscription rate (proportion of days an individual subscribed)



(b) Extensive margin (whether the individual 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 $\alpha + \beta_1$ from equation (1) for various time periods. "pre" refers to time before the experiment started; "promo" is the during the promotional time period, the last bucket "Entire post promo" aggregates across all post promo time periods. The error bars show 95% confidence intervals of the β_1 coefficient.

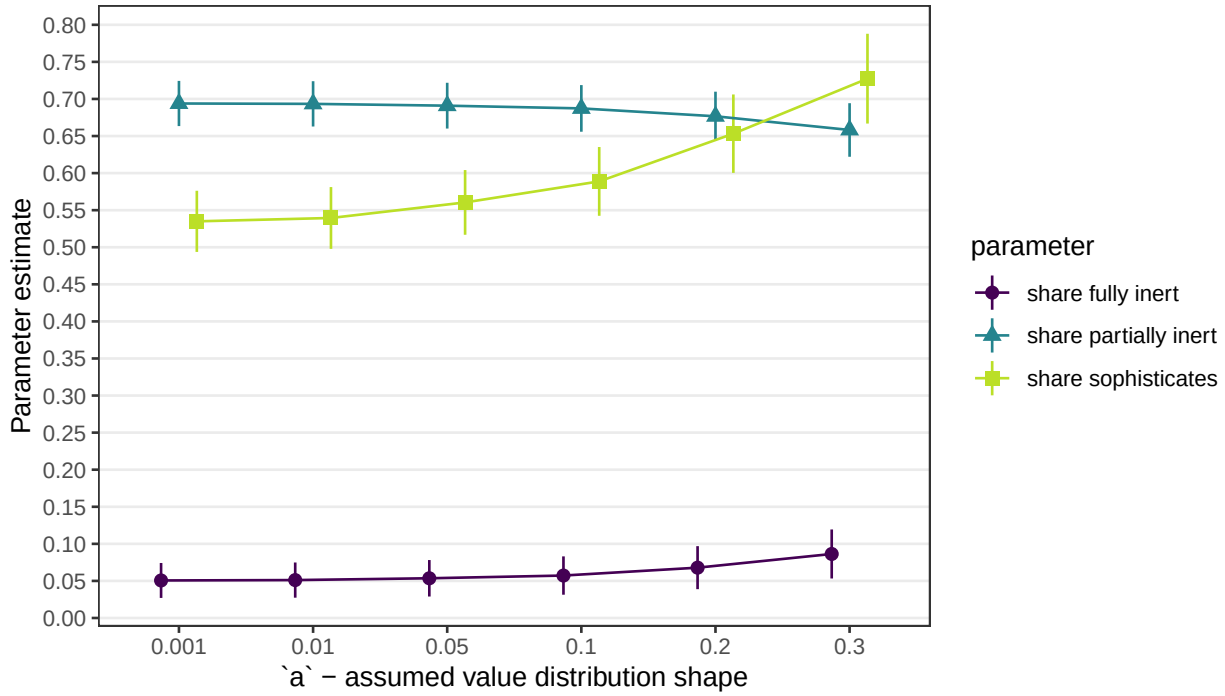


Figure A.3: Parameters' sensitivity to a, the parameter of value distribution skeweness

at all following periods if unsubscription failed, is $\sum_{\tau=0}^{\infty}\left[\hat{\phi}_i^{\tau}\delta^{\tau}\left(v_i-p-\left(1-\hat{\phi}_i\right)c^u\right)\right]=\frac{v_i-p-\left(1-\hat{\phi}_i\right)\hat{c}^u}{1-\delta\hat{\phi}_i}$.[17]

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:[18]

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

[18] 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

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

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

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

This setup highlights the different forces that affect perceived and actual inertia, and how they translate into observable subscription and usage patterns. Those who value the subscription will sign up regardless, as auto-renewal or auto-cancellation do not affect them. However, those who draw some value, enough to try but not enough to pay a full price, are possibly affected. For them, *perceived* future cancellation costs and inaction reduce take-up of an auto-renewing contract due to the risk of being locked-in paying for a product they do not like. The *actual* costs lead to an increase in the share of long-term subscribers roughly to the extent these subscribers underestimate the costs at sign-up; and actual inaction leads to a persistence in the number of medium-run subscribers to the extent that these subscribers underestimate their future inaction.[19] 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.2 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 post 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 **??**, which estimates average inertia experienced across all auto-renewal takers, this approach estimates inertia experienced by the marginal individuals—those who take an

simplification, but a realistic one.

[19] 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.

Table A.3: Effect of experimental incentives on post promo subscription

	Auto-renewal 4 weeks, Free vs. 2 weeks, €0.99	Auto-cancel 4 weeks, Free vs. 2 weeks, €0.99
	estimate (s.e.)	*estimate (s.e.)*
Effect in promo time (Δy_0)	.0013 (.0002)	.00044 (.00018)
Effect post promo month 1 (Δy_1)	.0007 (.0001)	-.00019 (.00013)
Effect post promo month 2 (Δy_2)	.0004 (.0001)	-.00011 (.00013)
Effect post promo month 3 (Δy_3)	.0001 (.0001)	-.00010 (.00012)
Effect post 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 (4 week, free) to (2 weeks, €0.99) on the promo period (Δy_0) and post promo (Δy_t) subscription rates, separately for auto-renewal and auto-cancel groups. The next four rows present our estimate of post 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 post promo and is statistically indistinguishable from zero by month 3. Under auto-cancel, the subscription rate drops immediately post promo, and is statistically insignificant in all post 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.

auto-renewal subscription only when an additional incentive is given with it.

Table A.3 shows our estimates. For individuals assigned an auto-renewal offer reducing price and increasing trial duration simultaneously, that is, going from 2 weeks, €0.99 auto-renewal to 4 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}.[20] Looking beyond the promo period, in the 4 weeks post 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 post 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 post 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 imprecise. If we assume that our imprecise estimates for $\frac{\Delta y_t^{AC}}{\Delta y_0^{AC}}$ are actually zero, we can see inertia comparable to the average effects in section **??**.

[20] 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.