American Economic Review 2025, 115(5): 1650–1671
https://doi.org/10.1257/aer.20231612

Selling Subscriptions†

By Liran Einav, Ben Klopack, and Neale Mahoney*

We study one benefit to firms of selling subscriptions: the prospect that consumers will continue to pay for subscriptions they no longer value. We use comprehensive data from a large payment card network to document that months during which cards are replaced, when active renewal is required, are associated with much higher rates of cancellation. Using two stylized models of consumer inertia—driven by inattention or switching costs—we estimate that these cancellation frictions roughly double seller revenues on average, holding fixed initial subscribers. We use the estimated models to explore the impact of possible regulatory remedies. (*JEL* D12, D18, L81, L88)

A growing number of retail products are now sold as subscriptions, which are typically billed on a monthly basis and automatically renewed unless a consumer actively cancels. While subscriptions have been common in some product categories for years (e.g., newspapers and gym memberships), their use has recently expanded to digital products (e.g., media streaming services and software licenses), home security systems, consumer products (e.g., clothing, shaving products, and makeup), and ingredients for home-cooked meals. According to some estimates, the "subscription economy" more than quadrupled in size over the last decade (Zuora 2022).

This rapid growth is often attributed to two factors. On the supply side, digital products have become a larger share of the retail sector, and such products may lend themselves more naturally toward a subscription model. On the demand side, there seems to have been an increased emphasis on convenience, and subscriptions are often associated with more convenient, hassle-free transactions.[1]

In this paper, we explore the potential for a third factor to play a quantitatively important role in the growth of the subscription economy. Because subscriptions are automatically renewed, consumers who are inertial may continue to pay for

* Einav: Stanford University and NBER (email: leinav@stanford.edu); Klopack: Texas A&M (email: bklopack@tamu.edu); Mahoney: Stanford University and NBER (email: nmahoney@stanford.edu). Katja Seim was the coeditor for this article. We are grateful to Stefano DellaVigna, Jean-Pierre Dubé, and three anonymous referees for many helpful suggestions, and to seminar participants at the Chicago Fed, Chicago Harris, Copenhagen, FTC, HBS, IIOC, Jerusalem, LMU, Montreal IO conference, NBER Digitization, Nemmers Prize conference, Stanford, Tel Aviv, TXio conference, UCLA, UPenn, and WU Vienna for useful comments. We thank Chris Xue and especially Katja Hofmann for superb research assistance. Mahoney was on leave from Stanford and NBER at the White House in 2022–2023; all his work on this paper was done while he was at Stanford. We received IRB approval for the full study, including interventions and data collection, from Stanford IRB, eProtocoal 53506.

† Go to https://doi.org/10.1257/aer.20231612 to visit the article page for additional materials and author disclosure statement(s).

[1] A survey finds that 32 percent of US consumers "signed up to the subscription because it feels nice to receive something every month" (Godderidge 2021). A decade ago, a similar emphasis on increased convenience may have helped explain the transition on eBay from auctions to fixed prices (Einav et al. 2018).

subscriptions they no longer value. Indeed, there are now multiple new companies whose business model (marketed as a subscription!) is to help subscribers find and cancel unwanted subscriptions. If consumers do not fully anticipate their inertia at sign-up, this may create supply-side incentives to offer subscriptions to exploit inertial consumers, amplifying the growth of subscription offerings.

To quantify these incentives, we use transaction-level data from a large payment card network to analyze consumer renewal and cancellation behavior for ten popular subscription services in the United States. Our research design takes advantage of our ability to observe card replacement (when cards expire, are lost, or stolen) and to link new cards to the cards they replaced. Because the replacement card is associated with new card information (such as the expiration date, security code, and sometimes the card number), consumers typically need to update their billing information with the subscription provider, inducing an active renewal decision. We document a sharp drop in subscriber retention rates during the month of card replacement. The sharp drop is clearly present for all the subscription services we study although it varies substantially in magnitude.

The patterns we observe could result from some combination of inattention or switching costs. Because our data and variation do not allow us to tease apart the relative importance of these factors, we separately specify and estimate both inattention and switching cost models to economically interpret these patterns and quantify the impact of inertia on seller revenues.

In both models, subscribers are myopic and sign up for a service when their flow utility (plus sign-up costs) is greater than the price, with utility evolving according to an AR(1) process in future periods. In the inattention model, we assume a fully attentive subscriber would cancel their subscription as soon as their flow utility falls below the monthly charge. Yet we assume that in most periods (with the exception of the month of card replacement), subscribers are imperfectly attentive and make an active choice with probability $\lambda < 1$. Consumers thus sometimes continue paying for a subscription even when they "shouldn't." In the switching cost model, we assume that subscribers are fully attentive each month but face a (symmetric) switching cost κ. In months without card replacement, subscribers pay the switching cost to cancel, while in months with replacement, they pay the switching cost to renew.

We estimate each model separately for each subscription service. Despite having only three parameters for each service, the models replicate the key patterns in the data remarkably well. Consistent with the sharp drop in retention rates during the month of card replacement, we estimate a fairly large degree of inertia. In the inattention model, the average estimate of the attention parameter λ is 0.18, with a range of 0.04 to 0.50 across subscriptions. The estimates of switching costs κ are also large relative to the variation in preferences. For instance, the median ratio of switching cost to the standard deviation of initial consumer surplus is 0.8, again with substantial variation across subscriptions. Naturally, the estimates of inattention and switching costs are strongly correlated across services.

We use the estimated model to perform counterfactual exercises that assess how much more quickly consumers would cancel their subscriptions if there was no inertia, which corresponds to fully attentive consumers (inattention model) or default cancellation every month (switching cost model). We find that seller revenues (or equivalently average subscription durations) are significantly higher due to subscriber inertia with important heterogeneity across services. Specifically, in the inattention model, we find that inertia increases seller revenues by 87 percent on average, with increases that range from 14 percent to more than 200 percent depending on the service. In the switching cost model, inertia raises revenue by 120 percent on average, with a range of 17 percent to 259 percent.

The optimal policy response depends on the underlying source of inertia. If it stems from inattention, policies that require active choice may be welfare improving. If it derives from switching costs, then the optimal policy may depend on whether the switching costs are welfare relevant and whether they can be reduced through policy actions (such as the recent Federal Trade Commission click-to-cancel rule (Federal Trade Commission 2024)). A full analysis would also require modeling the impacts of the policy on the pool of initial subscribers (extensive margin) and effects on prices or product offerings, which are outside the scope of our stylized models.

Our view is that inattention is probably the more dominant force although the evidence is circumstantial, and both mechanisms likely play a nonnegligible role. The limited research that directly compares inattention-type mechanisms to switching-cost-type mechanisms indicates that inattention-type mechanisms are quantitatively more important (Andersen et al. 2020; Heiss et al. 2021).[2] Consistent with inattention being important, a recent survey found that nearly 90 percent of consumers underestimated their monthly spending on subscriptions, with the average respondent spending more than three times their initial estimate (West Monroe 2021). The Rocket Money subscription management service emphasizes consumer inattention in its marketing with advertisements featuring consumer testimonials on hundreds of dollars in forgotten subscriptions.[3] This uncertainty about mechanisms notwithstanding, we use the estimated models to explore the potential impact of simple policy remedies, which—in the spirit of recent policy guidance from the Federal Trade Commission (2021)—would require firms to provide consumers with an active renewal decision at regular frequencies. In the inattention model, we find that requiring active choices at a six-month frequency would reduce the excess revenue from inattention by 45 percent. The switching cost model makes a similar quantitative prediction; moving from default renewal to default cancellation once every six months would reduce excess revenue by 48 percent.

As mentioned, our study focuses on the renewal or cancellation decision (intensive margin), taking as given the set of subscribers who initially sign up for the service (extensive margin). The counterfactual revenue effects would naturally be smaller if consumers anticipated their future inertia when they signed up and adjusted their enrollment decision in response to a policy change. This question is studied in important complementary work by Miller, Sahni, and Strulov-Shlain (2023) in the context of a newspaper subscription in Europe. Using a field experiment that varies whether a subscription automatically renews or cancels, they find that consumers are less likely to sign up under automatic renewal, indicating that consumers partially anticipate their future inertia. In addition to our focus on the intensive margin,

[2] Evidence and discussions in de Silva (2023), Giglio et al. (2021), and Ascarza, Iyengar, and Schleicher (2016) also point to a more prominent role for inattention-type mechanisms, but these findings are not central to these papers.

[3] See https://www.youtube.com/watch?v=ykap7HHGUZI for an example. The service also addresses switching costs by helping consumers cancel their subscriptions.

our work also differs from Miller, Sahni, and Strulov-Shlain (2023) in that we study multiple subscription services and document significant heterogeneity across both services and consumers.[4]

Our paper builds on a large literature on consumer inertia in consumer product markets. This includes the seminal DellaVigna and Malmendier (2006) study on gym memberships that finds substantial cancellation lags for customers with automatically renewing contracts. It also includes studies by Esteves-Sorenson and Perretti (2012) on television channel choices, Handel (2013) and Brot-Goldberg et al. (2023) on health insurance enrollment, and Posner et al. (2023) on political contributions. Our paper also relates to a contemporaneous study by Jespersen (2022) that finds that subscribers whose payments are rejected are 70 percent more likely to cancel.[5,6]

Our study connects to a literature on optimal contract design for "behavioral" consumers who exhibit inattention or limited self-control (e.g., Eliaz and Spiegler 2006, 2008). DellaVigna and Malmendier (2004) show that firms may exploit the overconfidence of present-biased consumers by designing contracts with back-loaded pricing and automatic renewal. Johnen (2019) studies how firms trade off the exploitation of naïvely inattentive consumers and the adverse selection of sophisticated consumers who make an active decision about contract renewal and can avoid high renewal prices.

Finally, our paper relates to a growing literature in computer science and law that studies the prevalence and impact of deceptive user interfaces on websites and smartphone apps, sometimes referred to as dark patterns (e.g., Mathur et al. 2019; Di Geronimo et al. 2020). Using experiments, Luguri and Strahilevitz (2021) show that dark patterns induce consumers into signing up for subscriptions they would otherwise avoid. Our paper is more broadly related to policy efforts to reduce inertia in subscription plans, such as the "click to cancel" rule proposed by the Federal Trade Commission (2024).

The rest of the paper is organized as follows. In Section I we describe the data, the selection of subscription services, and the construction of the sample. In Section II we present descriptive evidence that motivates our key exercise and illustrates our empirical strategy. Section III presents the model, its estimation, and the counterfactual exercises, which allow us to quantify how inertia affects firm revenues. The final section concludes.

I. Data and Sample Construction

A. *Data*

Data Source.—Our primary data source is transaction data from a large payment card network in the United States between August 2017 and December 2021 (Anonymous Firm 2017-2021). Using publicly available information on total

[4] Goettler and Clay (2011) analyze consumer choices between flat fee and pay-per-use plans. They show that rational, forward-looking consumers who face uncertain utility and cancellation costs may inefficiently sign up for the flat fee and fail to switch.

[5] Reme, Røhr, and Sæthre (2021) and Ascarza, Iyenga, and Schleicher (2016) study the dynamics of inattention and subscription attrition.

[6] Our study is more tangentially related to the literature on default effects and active versus passive choices in retirement savings (e.g., Madrian and Shea 2001; Carroll et al. 2009).

subscribers for the services in our baseline sample (described below), we estimate that our data cover approximately 30 percent of subscribers.

An observation in our data is a transaction, and the information on each transaction is similar to the typical information one would find on monthly credit card statements: the name of the merchant, a unique card identifier, a transaction amount, and a date. Importantly, there is no information on the specific goods or services that were purchased nor their prices. The sample is depersonalized and does not contain the name, address, or any other personally identifiable information about the cardholder.

Critically for the empirical strategy we describe below, a card in the data is associated with a unique account identifier, so that multiple cards within the same account can be linked. Thus, if a card expires or is lost or stolen, its replacement card will have a new card identifier but retain the same account identifier. Because the quality of the account identifier variable is low for cards that were replaced in the early part of our sample, we focus our analysis on cards that were replaced in July 2018 and after. We note that we cannot link multiple accounts held by the same consumer, so we treat accounts as independent of each other.

Identifying Subscription Services.—We use an industry report as a starting point to identify the set of subscription services for our study. Specifically, we start with a list of 21 categories and 49 specific subscription services used by West Monroe (2021) for their consumer survey of subscription spending. We augmented this list by searching for industry reports for each category and adding any additional subscription services with more than 500,000 subscribers (as reported by public sources). This process yielded a list of 69 subscription services. Of these 69 services, we were able to identify 57 via "manual" name searching in the payment card data.[7]

We imposed the following additional criteria to arrive at a final list of subscription services. First, we required that subscription services had a minimum of 500,000 average monthly subscribers in our data, which eliminates 31 of the 57 services. We then dropped four services that are primarily sold in long-term contracts (two cell phone and two internet service providers), six services that were sold by merchants with many nonsubscription products,[8] two services with average subscription length shorter than six months, two services that were launched toward the end of our observation period, and two services with nonmonthly billing.

Our final sample is comprised of ten large subscription services. We identify them throughout by letters (A–J) as our data use agreement prevents us from revealing the merchant names. The ten services include both digital and nondigital products covering multiple merchant categories, including entertainment, security, retail goods, and newspapers.

The product offerings of these ten services remained largely stable during our sample period with relatively small changes in pricing.[9] We investigated the consumer

[7] We can provide additional details on the process that led to selecting these 69 subscription services upon request and subject to review by the data provider.

[8] Because we are unable to observe which items were included in a transaction, it is hard for us to identify subscribers separately from other customers for a merchant that sells both subscriptions and other products.

[9] Two of the ten services made one-time price increases of \$1–\$2 to the monthly rate of the base subscription package, a third service increased the monthly price of their family package by \$1, and a fourth service reduced the price of their base package while increasing the price of their premium services.

response to these price changes in the data and found it to have a negligible impact on consumer cancellation rates, so we abstract from these (small) changes for the rest of the analysis. Moreover, we use service-specific month fixed effects in our empirical specification, which should absorb any impact of changes in prices or product offerings.

B. *Sample Construction*

Our research design focuses on subscription renewal around card replacement. We thus limit our sample to consumer credit and debit accounts[10] that had their cards replaced exactly once between July 2018 and January 2021. There are about 23 million accounts (and about 35 million account-service pairs) that meet these criteria and transacted with at least one of the ten subscription services analyzed.

We organize the data at the monthly level and use the last transaction made on the old card to define the last month in which the old card was available and the subsequent month as the first month in which the new card is active. We drop accounts for which there was a gap of more than one month between the last time the old card was used and the first time the new card was used (7 percent of the sample of 23 million accounts mentioned above) and accounts where the old card continued to be used after the replacement card was issued (an additional 12 percent of accounts). These restrictions leave us with about 19 million accounts and approximately 28 million account account-service pairs.

To construct our final sample, we make further restrictions to ease the analysis and graphical exposition of our results. Specifically, we include an account service in the final sample if card replacement occurs exactly 6, 12, or 18 months[11] after sign-up.[12] Once included in the sample, we track each account service for 25 months and require that the account is active, that is, that the account is associated with at least one transaction (any transaction) in each of the 25 months. Accounts that do not satisfy this activity requirement are excluded from the analysis sample. The resulting analysis sample is a relatively small subset—870,358 account-service pairs, representing 800,545 distinct accounts—of these 28 million account-service pairs.[13]

Our final sample thus contains a collection of cohorts of initial subscribers to each subscription service. A cohort of subscribers to a particular service is defined by a sign-up month and a card expiration at month x, where x is equal to 6, 12, or 18 months. The initial subscription month runs from January 2018 through July 2019.[14] For a given service, we therefore observe a total of 57 cohorts: 19 cohorts

[10] This excludes card types intended for business use and prepaid cards.

[11] The choice of 6, 12, and 18 is our (admittedly arbitrary) attempt to evenly span the 25-month panel structure to facilitate graphical presentation of the data.

[12] Recall that the earliest month in our data is August 2017; to focus on initial subscriptions, we keep subscriptions that start in January 2018 and after.

[13] Some consumers sign up for more than one subscription service. While card replacement could in principle lead to correlated cancellation decisions, we do not detect such a pattern in the data. Cancellation decisions are no more correlated during card replacement months than nonreplacement months. We therefore treat each account-service pair as an independent observation throughout the paper.

[14] July 2019 is the latest month for which we can observe a card replacement that occurs 18 months after sign-up.

defined by their sign-up months, each partitioned to three subcohorts that are based on the number of months (6, 12, or 18) at which card replacement occurs.[15]

We apply two final "data cleaning" steps that facilitate the subsequent analysis. First, we guarantee that each account-service observation follows a simple data structure that would fit a hazard model: if we observe an account transacting with a service but "skipping" a single month, we "fill in" that month,[16] and if we observe two months or more without transactions with the service, we assume that the cardholder unsubscribed to the service regardless of any subsequent transactions (which we interpret as "resubscriptions").[17] Second, we exclude the first month we observe a transaction for each cohort. For most services, we observe a much larger drop in subscriptions after the first month than after subsequent months. We thus view this first month as "special" (e.g., a "trial period"), and in what follows we consider the "sign-up month" as the second month in which we observe a transaction in our data. This restriction reduces the number of account-service pairs from 870,358 to 635,021.

II. Descriptive Evidence

Empirical Constructs.—Consider a cohort (s,x), which is associated with a given service, a sign-up month s, and card replacement which occurs in month $s+x$. Denote the number of subscribers in each month $t \geq s$ by $N(t;s,x)$, and define the cohort-specific retention rate as

$$R_n(t;s,x) \equiv N(t;s,x)/N(s;s,x), \tag{1}$$

where $n \equiv t-s$ is the age of the cohort in months. That is, the retention rate is the share of initial subscribers that remain subscribed at age n.

The top panel of Figure 1 presents the data in their most granular form. It is focused on one subscription service ("service A") and only on the 19 cohorts whose cards are replaced 12 months after sign-up ($x = 12$). For those 19 cohorts, we plot the retention rate, R_n, in each month, throughout the 24-month observation period. The pattern is quite similar across the cohorts, revealing a smooth decline in retention rates over time, with a sharp drop in retention rates around the card replacement month (month 12).

Although the raw patterns across cohorts are quite similar, it seems natural to aggregate across cohorts to adjust for any possible differences in cohort sizes, service popularity, seasonal variation, and (relatively small, as mentioned earlier) changes in the product offering and monthly subscription prices.

To do so, we estimate the following regression separately for each service:

$$R_n(t;s,x) = \beta_t + \gamma_{n,x} + \varepsilon_{t,s,x}, \tag{2}$$

[15] For four (out of the ten) services, we observe only 56 (rather than 57) cohorts because the relatively small service size and small number of card replacements we observe in 2018 lead to no observations associated with January 2018 subscribers whose card is replaced in July 2018.

[16] That is, if we observe no transaction in month $s+t$ for a given service, we still consider this account as "subscribed" as long as there are transactions in months $s+t-1$ and $s+t+1$. This adjustment is quantitatively small and raises the average subscription duration in the entire sample from 17.2 to 17.6 months.

[17] About 20 percent of accounts that unsubscribe for two months or more return to the service within the 25-month window.

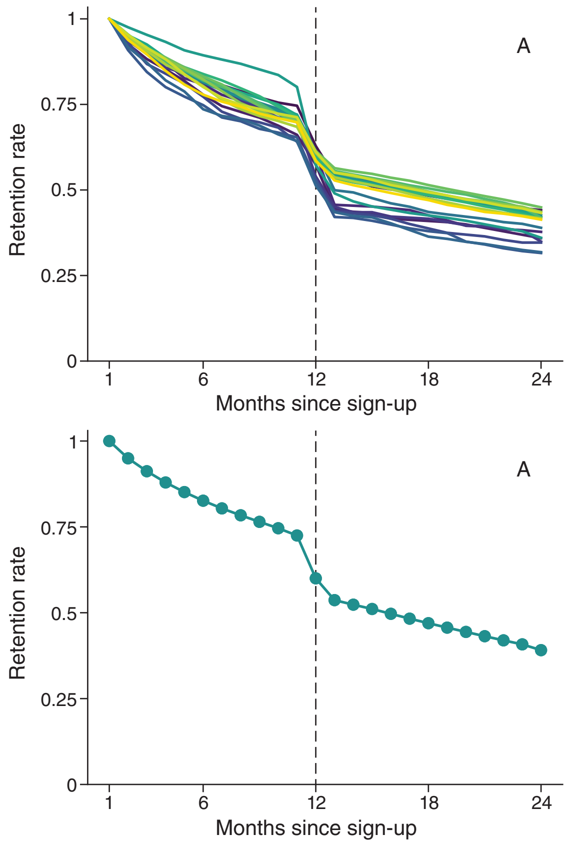

Figure 1. Aggregating Retention Rates

Notes: Figure shows retention rates by month since sign-up for subscription service A and cohorts with card replacement 12 months after sign-up. The top panel shows the raw retention rates for all 19 cohorts. The bottom panel shows the adjusted retention rate, $\hat{R}_n(x)$, which aggregates across cohorts netting out calendar-month fixed effects. See Section II for details.

where β_t is a calendar-month fixed effect and $\gamma_{n,x}$ is a fixed effect for the number of months since sign-up, which is allowed to flexibly vary with x. We weight observations by cohort size, $N(s;s,x)$, to reflect the behavior of the average subscriber. With these estimates in hand, we define the adjusted retention rate as

$$(3) \qquad \hat{R}_n(x) \equiv \hat{\gamma}_{n,x}/\hat{\gamma}_{1,x},$$

where the $\hat{\gamma}_{n,x}$s are the estimated coefficients from equation (2).

The bottom panel of Figure 1 displays this empirical construct for the 19 cohorts shown in the top panel of the figure. That is, it plots $\hat{R}_n(x)$ for the same service ("service A") and x ($x = 12$), and it shows how the adjustment aggregates across cohorts and smooths out some of the cohort-specific noise.

Figure 1 reveals our inability to perfectly time the date of card replacement. Specifically, it shows that the sharp drop in retention rates does not happen in a single month but instead occurs over two consecutive months. This pattern will repeat itself throughout, and we will explain below how we adjust for it when estimating the model.

Supplemental Appendix Table B1 summarizes the drop in retention rates during card replacement, reporting the average monthly change in retention rates $(\hat{R}_n(x) - \hat{R}_{n-1}(x))$ during the two-month replacement window and outside of it for each subscription service. On average, the monthly drop in retention is 0.08 during the replacement window, four times larger than the 0.02 drop during other months. This understates the difference since, as discussed above, we cannot pin down the exact month of card replacement.

Account Activity around Card Replacement.—Our economic interpretation of the sharp drop in retention rate around card replacement (discussed in more detail below) is that it reflects a change in the default choice faced by the cardholder. Doing nothing prior to card replacement results in subscription renewal, while doing nothing after the card is replaced leads to cancellation.

A potential threat to this interpretation is that card replacement may have a more general impact on consumers' spending across their portfolio of credit and debt card accounts. For instance, if replacement is associated with an interruption in card access (e.g., while waiting for the new card to arrive in the mail), then the cardholder may switch their subscription to a different account. In this case, the sharp drop in retention rate would simply reflect substitution of the subscription to this other account rather than cancellation.

To assess this concern, Supplemental Appendix Figure B1 uses the entire analysis sample and presents the variation in account activity around the month of card replacement. The top panel shows the number of monthly transactions associated with the account, that is, the number of monthly transactions on the old card prior to replacement and the number of monthly transactions on the new card after replacement. The bottom panel repeats the same exercise but uses total monthly spending on the account.

The plots show some disruption in account activity during the month of card replacement followed by a quick and almost full recovery to prereplacement levels. The average number of monthly transactions falls from 44.2 in the month before card replacement to 36.8 in the month of replacement before recovering to 42.1 two months later. Average monthly spending falls from \$2,202 to \$1,928 before recovering to \$2,114.[18] The small decline in activity could be driven by subscription lapses. While the plots look very similar if we exclude transactions at the ten subscription services we study, to fully eliminate any renewal effects we would need to exclude all subscription services, which we cannot do because we cannot identify all subscriptions. Consistent with this explanation, Supplemental Appendix Figure B2 shows a complete recovery for the gas and groceries categories of spending where there are no subscriptions. The difference in monthly transactions in the three months before and after card replacement (excluding the two months around replacement) is 0.00 percent for groceries and −0.18 percent for gas, compared to −4.13 percent overall.

[18] The reason that it takes two months rather than one to recover to the original level is due to our inability to perfectly time the month of card replacement, as discussed earlier.

Cancellation Rates around Card Replacement.—Figure 2 presents the adjusted retention rates, $\hat{R}_n(x)$, for the analysis sample. Each panel shows retention rates for a separate service (indicated by the letter in the top right corner of the plot) with separate lines for cohorts of accounts with card replacement at 6, 12, and 18 months.

The retention patterns are quite heterogeneous across the ten different services, but the common theme across them is a sharp drop in retention rates around card replacement. The drop is noticeable but relatively small for some services (e.g., C, G, and J) and is much larger in some of the others (e.g., A, B, D, and I). Across the services, the magnitude of the drop at card replacement is quantitatively similar for subscribers whose card is replaced 6, 12, and 18 months after sign-up.

To economically interpret and quantify these patterns, we next specify two models of subscriber renewal behavior. We then use the estimated models to quantify the impact of subscriptions on seller revenues along with the revenue impact of counterfactual policies.

III. Quantifying the Impact on Revenues

Our primary objective is to estimate the revenue impact associated with the automatic renewal of subscription services. To do so, we specify and estimate two alternative models. In the first model, inertia is driven by inattention, while in the second, inertia stems from switching costs. The descriptive patterns are consistent with either underlying source of inertia, and our data and variation do not allow us to tease apart the relative importance of these mechanisms. Of course, the underlying source of inertia—inattention or switching costs—may have implications for policy. We return to this issue in our discussion of policy responses below.

Because the focus is on the revenue impact for sellers rather than the utility impact on consumers, the models are static and highly stylized and should be viewed as a positive (rather than normative) description of renewal behavior. Other than the source of inertia, we try to keep the models as similar as possible.

A. *Inattention Model*

Consider a specific subscription service, which is associated with a monthly subscription price p,[19] and a potential subscriber i, whose flow utility from the service during month t is denoted by u_{it}.

We assume that u_{it} follows a Markov process, such that $u_{it} \sim F(\cdot \,|\, u_{i,t-1})$, and that consumers—once they are already subscribed—do not take into account any dynamic considerations, so their renewal decision only relies on the comparison between the flow utility u_{it} and the price p. This latter assumption is consistent with consumers being myopic or alternatively with consumers being forward looking but failing to anticipate their future inattention.

Given these assumptions, all new subscribers must have $u_{it} > p$ in the month in which they subscribe to the service for the first time, so we normalize the sign-up

[19] The estimation below is carried out on a service-by-service basis, so we omit service subscripts for expositional clarity throughout.

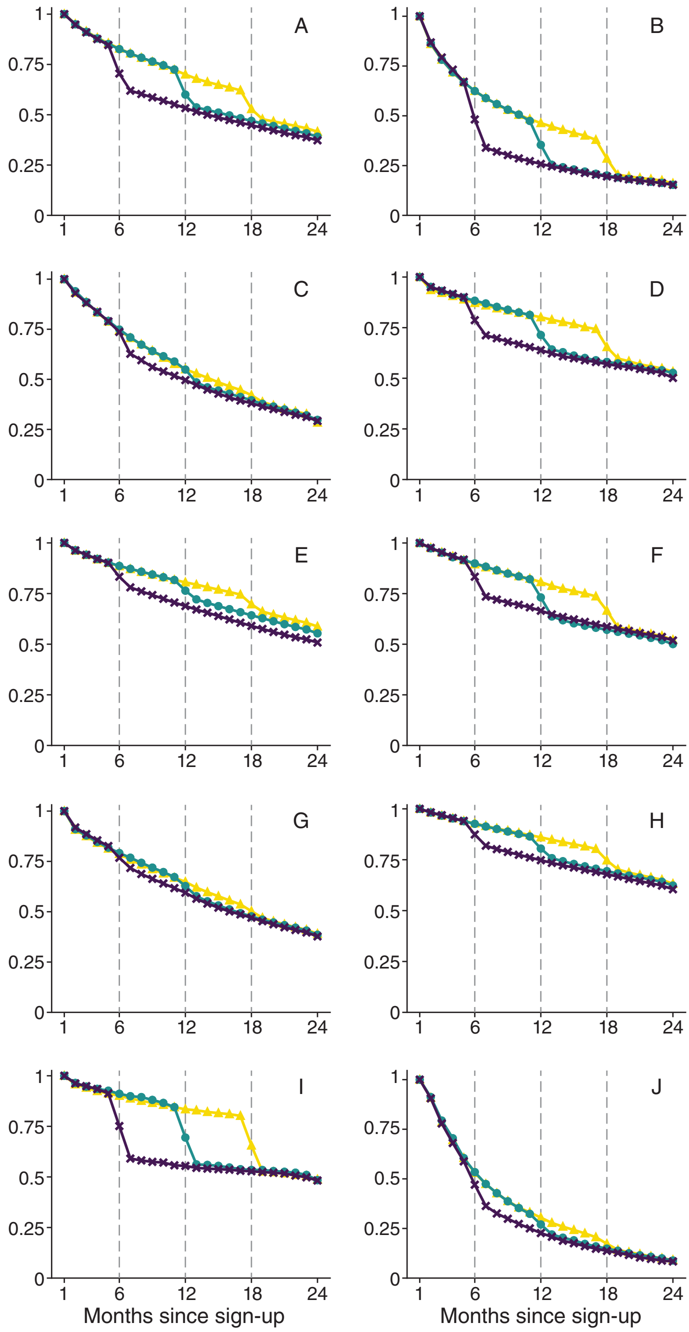

FIGURE 2. RETENTION RATE BY MONTHS SINCE SIGN-UP, ALL SUBSCRIPTION SERVICES AND COHORTS

Notes: Figure shows the adjusted retention rate, $\hat{R}_n(x)$, by month since sign-up, separately by groups of cohort with card replacement at 6, 12, or 18 months after sign-up and subscription service (denoted by the letter in the top right corner of each panel). The adjusted retention rate aggregates across cohorts net out calendar-month fixed effects. See Section II for details.

month to $t = 0$ and denote the cross-sectional distribution of u_{it} for new subscribers by $G\left(u_{i0} \mid u_{i0} > p\right)$.

In a typical month t, a subscriber can be either attentive or inattentive. If inattentive, the subscriber automatically renews the subscription. If attentive, the subscriber renews if and only if $u_{it} > p$. Subscribers are attentive in a given month with probability $\lambda_{it} \in (0,1]$. Importantly, in the first month after card replacement, we assume subscribers are perfectly attentive ($\lambda = 1$) because they are asked to actively enter the details of their new card.

Parametrization.—We define the net flow utility as $v_{it} \equiv u_{it} - p$ and assume that it follows an AR(1) process (without a constant),

$$v_{it} = \rho v_{i,t-1} + \varepsilon_{it}, \tag{4}$$

where ε_{it} follows a mean-zero normal distribution with a standard deviation that is normalized to one. We assume that the distribution of initial net utilities—$G(u_{i0} \mid u_{i0} > p)$ or equivalently $G(v_{i0} \mid v_{i0} > 0)$—is given by an exponential distribution, $v_{i0} \sim \text{Exp}(\eta)$, which has a mean and standard deviation η. Finally, we assume that the attention probability λ is the same across people and over time (for a given service). We explore alternative assumptions in Section IIIE.

Taken together, the model can be summarized by three service-specific parameters: the trajectory of flow utility from the subscription service (ρ), the extent to which new subscribers are close to the renewal margin (η), and the monthly probability of attention (λ). In Supplemental Appendix A we illustrate some of the model's comparative statics and provide intuition about the way the descriptive patterns map to the model's parameters.

Estimation.—We estimate the model separately for each subscription service using the method of simulated moments. Specifically, we focus on matching the moments in Figure 2: the adjusted retention rates $\hat{R}_n(x)$ for each subscription service, which vary by month since sign-up n and months between card replacement and sign-up x (which takes on values of 6, 12, or 18 months). To account for the fact (mentioned earlier) that we cannot perfectly time the month of card replacement, we omit the month of card replacement from the set of moments we try to match.[20] Overall, for each subscription service, we have 66 moments: for each of the three values of x (6, 12, and 18 months after initial subscription), we have 23 monthly retention rates, and we use all of them except the month in which the card is replaced. In estimating the parameters, we weight each moment by its corresponding cohort size $\left(\sum_s N(s;s,x)\right)$.

To construct model predictions for a given set of parameter values, we use the model to simulate retention rates as a function of the three model parameters ρ, λ, and η (see Supplemental Appendix B for more details) and estimate the parameters by minimizing the quadratic distance between the simulated moments and their empirical counterparts. While the parameters are allowed to vary flexibly across

[20] For example, if card replacement occurs in month six, Figure 2 shows the sharp drop in retention rates occurring over month six and month seven. By omitting the month-six retention rate from the set of moments we match in estimation, we are essentially allowing the card replacement to occur in either month six or month seven.

subscription services, we require them to be the same for a given subscription service across the three values of x (6, 12, and 18).

B. *Switching Costs Model*

The switching cost model hews as closely as possible to the inattention model, retaining the same parametric assumptions about the initial distribution of v_{i0} and its evolution over time.

The key modification we make is to the treatment of inertia. Instead of modeling it as arising from inattention, we assume that subscribers are fully attentive every month but face a (symmetric) switching cost κ.[21] In months without card replacement, subscribers pay the switching cost to cancel, and so they renew if and only if $v_{it} > -\kappa$. In months with card replacement, subscribers pay the switching cost to renew, and so they remain subscribed if and only if $v_{it} - \kappa > 0$. For internal consistency, we assume that initial subscribers also face the switching cost when they sign up, which means that v_{i0} is drawn from the $\text{Exp}(\eta)$ distribution that is truncated from below at κ rather than from the unconditional distribution $\text{Exp}(\eta)$.

Like the inattention model, the switching cost model is summarized by three service-specific parameters: the trajectory of flow utility from the subscription service (ρ), the extent to which new subscribers are close to the renewal margin (η), and the switching costs (κ). As before, we estimate the model separately for each subscription service, matching the moments shown in Figure 2, and we allow the parameters to vary flexibly across subscription services but require them to be the same for a given subscription service across the three values of x (6, 12, and 18).

C. *Model Fit and Parameter Estimates*

The parameter estimates for the inattention and switching cost models are shown in Table 1. Supplemental Appendix Figures B4 and B5 show the fit of the two models, plotting the predicted retention rates from the estimated model against their empirical counterparts, service by service. In general, the fit of the model is quite good, especially when taking into account the stylized nature of the model and the fact that it only has three parameters for each subscription service.[22,23]

The estimates of the inattention parameter λ are straightforward to interpret. The natural benchmarks are $\lambda = 1$, when consumers are fully attentive every period, and $\lambda = 0$, when consumers are fully *in*attentive every period. The estimates of λ range between 0.045 (service I) to 0.502 (service G) across subscription services. For service I, almost all subscribers in a given month are inattentive and renew their

[21] In principle, switching costs could be asymmetric. For instance, Fraccaroli, Mahoney, and Thabet (2024) find that many newspapers impose "sludge" that makes it difficult to cancel subscriptions, and the recently finalized Federal Trade Commission click-to-cancel rule is designed to reduce unnecessarily large cancellation costs. However, the cancellation patterns we observe only allow us to identify the sum of renewal and cancellation costs.

[22] We note again that the model, by design, does not try to fit the fact that in the data the drop in retention rates covers two months rather than one. As discussed before, this is an artifact of our inability to perfectly time the card replacement in the data, and it is a pattern that we intentionally do not aim to replicate with the model.

[23] Because we use the same moments for estimation, the objective function is comparable across the inattention and switching cost models. For half of the subscription services the fit is essentially identical, while for the other half the inattention model performs a little better.

Table 1—Parameter Estimates by Subscription Service

	Inattention			Switching Costs		
Service	ρ	λ	η	ρ	κ	η
A	0.987	0.111	0.864	0.981	1.081	0.761
	(0.003)	(0.002)	(0.023)	(0.002)	(0.011)	(0.023)
B	0.841	0.180	0.004	0.898	0.848	0.004
	(0.011)	(0.002)	(0.000)	(0.003)	(0.009)	(0.000)
C	0.947	0.271	1.596	0.949	0.567	1.485
	(0.006)	(0.014)	(0.099)	(0.005)	(0.034)	(0.129)
D	1.019	0.090	1.407	1.001	1.213	1.276
	(0.014)	(0.006)	(0.141)	(0.007)	(0.055)	(0.181)
E	0.999	0.131	2.560	0.998	0.986	2.111
	(0.002)	(0.003)	(0.049)	(0.001)	(0.012)	(0.044)
F	0.979	0.093	1.961	0.980	1.215	1.839
	(0.007)	(0.004)	(0.121)	(0.004)	(0.032)	(0.137)
G	0.945	0.502	3.829	0.944	0.234	3.909
	(0.003)	(0.000)	(0.072)	(0.002)	(0.007)	(0.063)
H	0.999	0.095	3.208	0.997	1.177	2.825
	(0.003)	(0.002)	(0.081)	(0.002)	(0.017)	(0.093)
I	1.058	0.045	0.475	1.001	1.769	0.444
	(0.041)	(0.003)	(0.104)	(0.006)	(0.060)	(0.131)
J	0.819	0.288	0.229	0.853	0.662	0.058
	(0.011)	(0.006)	(0.023)	(0.005)	(0.010)	(0.015)

Notes: Table reports parameter estimates and bootstrapped standard errors (in parentheses) for the inattention model (first three columns) and switching cost model (last three columns) described in Section III. We estimate the model separately for each of the ten subscription services, A–J. The standard errors are the standard deviations of the 1,000 bootstrap estimates.

subscription in a passive way, suggesting that inertia contributes significantly to the seller's revenues. This can loosely be seen in the shape of the retention rates for service I: a flat pattern before and after card replacement and a very sharp decline at card replacement, during which more than 30 percent of subscribers are lost. For service G, the empirical pattern is quite different: a fairly steep decline in retention rates over time and only a small incremental decrease in the month of card replacement. Yet even for this case, our estimate of λ is well below one, implying that (on average) subscribers make an active decision only every two months.

The estimates of the switching cost parameter κ are harder to interpret. Our preferred approach is to compare them to the η parameter that represents the mean and the standard deviation of the initial consumer surplus (since it is exponentially distributed). Like other papers with switching-cost-type parameters (Handel 2013; Berger et al. 2021), our estimates of switching costs are quite large and sometimes implausibly big. The median ratio of switching costs to consumer surplus is 0.8, and three services (B, I, and J) have ratios of three or larger. As shown in the top panel of Supplemental Appendix Figure B6, the attention and switching cost estimates are strongly negatively correlated across services (correlation of -0.92). Naturally, services that are rationalized by low attention in the inattention model are rationalized by high switching costs in the switching cost model.

The majority of the estimates of the ρ parameter are very close to one, suggesting that preferences for the service are (on average) stable after sign-up, and approximately follow a random walk. For several of the services, estimates of ρ are well below one, implying that these services may find it difficult to retain consumers for longer durations. These are services that may benefit the most from inertial consumers. Supplemental Appendix Figure B6 shows that the estimates of ρ are very similar across models (correlation of 0.96).

The η estimates reflect the extent to which new subscribers who sign up for the service are mostly marginal subscribers who are at risk of quickly unsubscribing absent inertia (low η) or mostly inframarginal subscribers who would require a sequence of negative preference shocks before they cancel (high η). The estimates are quite heterogeneous across services. Service B draws almost entirely marginal subscribers ($\eta = 0.004$ for both models), while services E, G, and H are associated with relatively high η estimates that are greater than two. Again, as shown in Supplemental Appendix Figure B6, the η estimates are very similar across models (correlation of 0.99).

D. *Counterfactual Exercises*

The Impact of Inertia on Revenues.—We now use the models to quantify the impact of consumer inertia on seller revenues. To do so, we use the two models and their estimated parameters described above and simulate the renewal decisions of a cohort of initial subscribers over ten years (120 months). For the inattention model, we assume a baseline where subscribers face no card replacement throughout, so they are attentive each period with probability λ. We then repeat the same exercise but assume fully attentive subscribers ($\lambda = 1$) and compare the results. For the switching cost model, we assume a baseline where the default is automatic renewal and so the subscriber has to pay a cost κ to cancel, and we then repeat the exercise where the default is automatic cancellation and the subscriber has to pay a cost κ to renew. Throughout, we make the simplifying assumption that subscribers who do not renew are lost "forever." This is a strong assumption but can be motivated by the observation that getting old subscribers to resubscribe to the subscription service could be almost as costly as attracting "fresh" subscribers. We discount revenues at a rate of 1 percent per month.[24]

The results are shown in Table 2, with panel A displaying results from the inattention model and panel B displaying results from the switching cost model. The first column of the table shows the share of consumers who are unaffected by the counterfactual exercise because they remain subscribed even without inertia. For the inattention model, the second column of the table shows the average number of months that the remaining affected consumers subscribe when they are attentive with probability λ every period, and the third column shows the average number of months they subscribe when they are fully attentive in every month. For the switching cost model, the second column shows the average number of months if the default is renewal and consumers have to pay the switching cost to cancel, and the third column shows the average number of months if the default is cancellation and consumers have to pay the switching cost to renew.[25] The last column reports the ratio of the revenues (measured by the number of subscriber-months) the seller obtains (over a horizon of ten years) between these two scenarios.

[24] The revenue benefits from inattention are not very sensitive to discounting.

[25] For this counterfactual, we hold fixed the initial valuation draws v_{i0} and assume that subscribers are defaulted into cancellation, rather than renewal, and remain subscribed if $v_{it} + \kappa > 0$.

Table 2—Revenue Impact of Inertia

Service	Share unaffected	Avg. months subscribed: If inattentive	Avg. months subscribed: If attentive	Revenue ratio
Panel A. Inattention				
A	0.04	37.0	13.7	2.10
B	0.00	14.2	3.9	3.30
C	0.00	21.8	13.7	1.51
D	0.28	42.0	9.4	1.59
E	0.21	41.9	20.3	1.40
F	0.04	45.2	20.0	1.87
G	0.00	23.8	20.4	1.14
H	0.24	49.3	22.9	1.41
I	0.26	54.7	3.9	2.16
J	0.00	9.9	4.2	2.25
Mean	0.11	34.0	13.2	1.87
	Share unaffected	Avg. months subscribed: If default renew	Avg. months subscribed: If default cancel	Revenue ratio
Panel B. Switching costs				
A	0.02	36.8	12.1	2.49
B	0.00	14.2	3.6	3.59
C	0.00	21.9	12.2	1.69
D	0.14	48.1	14.2	1.92
E	0.17	45.8	19.0	1.56
F	0.03	45.2	18.4	2.04
G	0.00	23.5	19.7	1.17
H	0.19	52.5	22.6	1.54
I	0.09	56.8	9.2	3.13
J	0.00	10.1	3.3	2.90
Mean	0.06	35.5	13.4	2.20

Notes: Table reports our counterfactual estimates on how inertia affects firm revenue. Panel A shows results using the model of inattention, while panel B shows results using a model with switching costs. For each subscription service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers for 120 months after sign-up. Column 1 in each panel reports the share of subscribers not affected by inertia because they never cancel even if they make an active choice in every month. Columns 2 and 3 show (for subscribers that cancel at some point absent inertia) the average number of months they are subscribed with and without inertia, respectively. Column 4 shows the revenue ratio with and without inertia, aggregating over both affected and unaffected subscribers. We construct the revenue ratio as follows: For each subscription service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers for 120 months after sign-up. The denominator is the discounted sum of monthly subscribers if consumers make an active choice in every month (i.e., if consumers are fully attentive in the inattention model and default into cancellation in the switching costs model). The numerator is the discounted sum of monthly subscribers if subscribers are never required to make an active choice and only pay attention with probability λ every month (inattention) or face default renewal (switching costs). We discount future revenues at a rate of 1 percent.

Overall, we find that the revenue effects from inertia are substantial and highly heterogeneous across subscription services. Under the inattention model, estimated attention—relative to the counterfactual of full attention—raises revenue by 87 percent on average, with a range of 14 percent for service G to 230 percent for service B.[26] In other words, if consumers were fully attentive, the average subscription duration would drop from 23.8 months to 20.4 months for service G and from 14.2 months to 3.9 months for service B. It is plausible to suspect that, from a business

[26] Supplemental Appendix Figure B7 illustrates the mapping between the model parameters and the implied revenue ratio.

perspective, the inattention of their consumers has a first-order relevance for subscription services like B.

The switching cost model predicts even larger revenue effects, with the average service seeing an increase in revenue of about 120 percent, compared to 87 percent under inattention. There is similar heterogeneity, with revenue increases ranging from 17 percent to 259 percent across services. As shown in Figure B6, the revenue ratios for each service are highly correlated across models (correlation of 0.94).

Exploring the Impact of Possible Policy Responses.—Our modeling framework does not allow us to conduct a full welfare analysis of potential policy responses. Due to limitations of our data and variation, we take as fixed the pool of initial subscribers, prices, and product offerings. Under policy counterfactuals, these features could change in ways that may offset our revenue estimates.

Moreover, as discussed, we are unable to tease out the relative importance of inattention and switching costs in driving subscriber behavior. If inattention is the predominant factor, there may be a rationale for policy that induces consumers to make active decisions or alerts them of their subscriptions. If switching costs are key, then the policy rationale is less clear. If switching costs are welfare-relevant, immutable features of the environment, then policy intervention may not be justified. However, if switching costs are not welfare relevant or can be reduced through appropriate interventions, then there may be a case for policy intervention.

These limitations notwithstanding, it is useful to conduct some simple counterfactual exercises, both to provide a qualitative sense of how subscription patterns might evolve in our stylized setting and because these estimates could be combined with outside estimates of unmodeled factors to build toward a more comprehensive policy assessment.

A natural policy remedy to consider, in the spirit of recent Federal Trade Commission (2021) guidance, is one that reduces inertia by requiring subscribers to make an active choice on a periodic basis. During these active choice periods, the subscription would terminate unless the consumer actively chooses to renew it (inattention model) or pays the switching cost of renewing it (switching cost model).

Figure 3 reports the revenue ratio by subscription service when consumers are counterfactually induced to make an active choice every X months, with $X \in \{1, 3, 6, 12, 18, 24, \infty\}$. In the inattention model, requiring an active choice every six months would reduce the excess revenue from inattention by 45 percent, nearly halfway between the extremes of full inattention (active choice every ∞ months) and full attention (active choice every 1 month). In the switching cost model, defaulting consumers into cancellation every six months would reduce excess revenue by 48 percent.

E. *Robustness and Heterogeneity*

Robustness.—To examine the sensitivity of our findings, we estimate several alternative specifications of the baseline inattention and switching cost models. Panels A and B of Supplemental Appendix Table B3 summarize these results, reporting summary statistics for the estimates of λ, κ, and the revenue ratios across

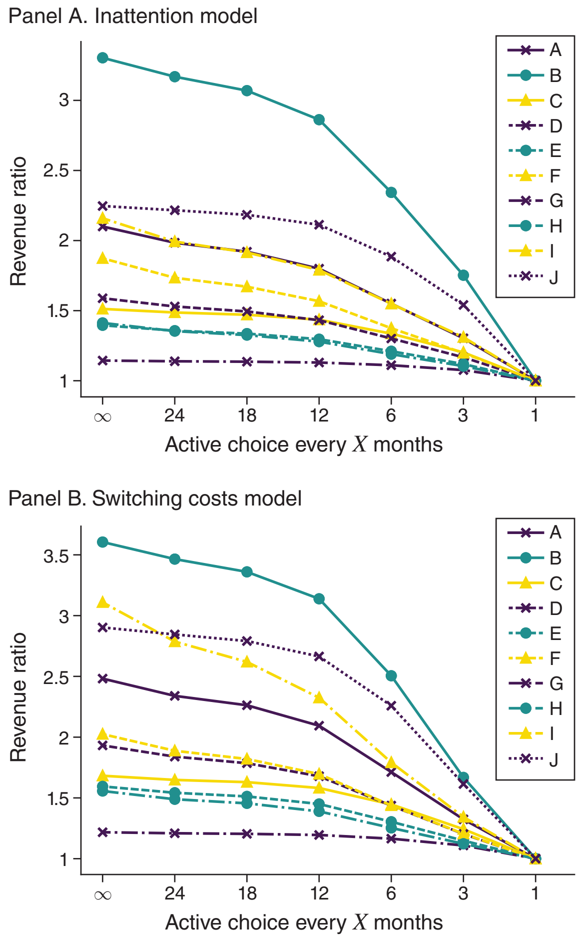

Figure 3. Revenue Impact of Required Active Choice

Notes: Figure shows the revenue impact of requiring subscribers to make an active choice every 1, 3, 6, 12, 18, 24 months, or never (∞) by subscription service under a model of consumer inattention (top panel) and switching costs (bottom panel). The plot shows the ratio of revenue under a counterfactual where consumers make an active choice every X months to revenue under a counterfactual where subscribers make an active choice in every month. We construct the revenue ratio as follows: For each subscription service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers for 120 months after sign-up. The denominator is the discounted sum of monthly subscribers if consumers are required to make an active choice every month (i.e., if consumers are fully attentive in the inattention model and default into cancellation in the switching costs model). The numerator is the discounted sum of monthly subscribers if consumers make an active choice every X months. We discount future revenues at a rate of 1 percent. The baseline is that subscribers are never required to make an active choice (∞) and only pay attention with probability λ every month (inattention) or face default renewal (switching costs). Supplemental Appendix Table B2 provides the underlying numbers in both plots.

subscription services, as well as how these service-level objects correlate with those from the baseline model.

The first two rows of panel A report results from specifications that allow inattention to vary over time since sign-up. The results from both specifications are similar to the baseline results. The third and fourth rows of panel A consider a case of $\lambda < 1$ at card replacement, which accounts for the possibility that when a card is replaced, some merchants may be able to update payment information automatically for some consumers. Our estimates of λ remain almost the same, although the revenue ratio results become even more striking.

Panel B reports results from a similar exercise using the switching cost model. Specifically, we assume that only a subset of consumers are defaulted into cancellation when their card is replaced, while the rest default to renewal. This alternative assumption tends to increase our estimates of both κ and the revenue ratio, consistent with the findings from panel A.

Heterogeneity across Consumers.—Supplemental Appendix Figure B8 shows the retention rates separately for the 1.4 percent of consumers who used their card at some point for a cash advance, which is a proxy for low financial sophistication (Agarwal et al. 2009), and consumers who never took out a cash advance.[27] There is a sharper drop in retention rates around card replacement for consumers who take out cash advances, which is indicative of greater inertia among these consumers.

In panels C and D of Supplemental Appendix Table B3, we show results from reestimating the baseline model for these two groups of consumers separately for the inattention and switching cost model. Consistent with patterns shown in Supplemental Appendix Figure B8, consumers who used cash advances are less attentive (lower λ) and have a revenue ratio that is 75 percent larger on average. Although we find it less natural to interpret this heterogeneity through the lens of the switching cost model, the switching cost estimates are also larger for the cash advance group. These patterns are (weakly) persistent across all subscription services (Supplemental Appendix Figure B9).

We have also explored heterogeneity along other cardholder attributes but do not detect meaningful heterogeneity by demographics of the zip code where the account holder has the largest number of transactions (not reported).

Heterogeneity across Subscription Services.—The λ and κ estimates vary substantially across the ten subscription services we study. This heterogeneity could reflect a combination of differences in the nature of the services as well as heterogeneity in the consumers who purchase them. Quantifying the role of subscription service characteristics is challenging. The services differ along many dimensions, including their product categories, whether they provide digital or physical goods, and their salience to inactive users. With only ten subscription services and multiple dimensions of heterogeneity, it is hard to attribute the differences in λ or κ to specific factors.

To try to disentangle the role of consumer characteristics from subscription service characteristics, we separately analyze a subset of cardholders who signed up for multiple subscription services. For four of the most common pairs of services, we estimate the baseline model separately for the cardholders who subscribed to both services and for those who subscribed to only one. If consumer characteristics drive most of the heterogeneity, the λ estimates should be more similar when we hold the set of subscribers fixed (that is, when we estimate the model using the population of those who subscribe to both services) relative to the λ estimates recovered from

[27] Cash advances allow cardholders to borrow a certain amount of cash against their credit card's balance. It is often associated with fees and high interest rates, so it is widely considered an expensive form of borrowing.

other consumers. Unfortunately, the results are mixed and do not allow us to draw strong conclusions.[28]

IV. Conclusions

In this paper, we use payment card replacement to examine the consequences of consumer inertia for subscription revenue. Focusing on ten large subscription services that we can reliably identify in our data, we document a sharp drop in the monthly renewal rate when the payment card is replaced relative to other months when the subscription is automatically renewed without an active decision by the consumer.

We specify and estimate stylized models where inertia either arises from inattention or switching costs. In both models, we find that inertia roughly doubles subscription revenue relative to counterfactuals with full attention (inattention model) or default cancellation (switching costs model). We also estimate substantial heterogeneity, with the excess revenue from inertia ranging from less than 20 percent to more than 200 percent across subscription services. Our counterfactuals hold fixed the pool of initial subscribers (extensive margin) and prices or product offerings; endogenizing these factors may lead to smaller revenue effects.

While this inertia raises concerns about firms exploiting "behavioral" consumers, automatic renewal also conveys convenience benefits, suggesting that the extreme policy of requiring consumers to make an active renewal choice every month may not be optimal.[29] We quantify the revenue impact of more balanced remedies, which would require active renewal at intermediate frequencies. We find, for example, that requiring consumers to make an active choice every six months cuts the excess revenue from inertia by about half. For some types of subscription services, such as those for digital goods, requirements could depend on account activity (e.g., firms could be required to obtain active renewal for accounts that have been dormant for a certain period of time). Exploring these policies would require data with product use and is an interesting topic for future work.

We highlight two important caveats, which we referenced earlier, to the interpretation of our results. First, a key limitation of our study is that we condition on the initial subscription decision. As Miller, Sahni, and Strulov-Shlain (2023) document in the context of newspaper subscriptions, the propensity to subscribe is also affected by the autorenewal features of the contract. Combining the initial subscription and monthly renewal decisions would be a promising direction for future work. Second, given the limited variation in our data, we do not model the pricing behavior of sellers, which could be impacted by policy interventions that reduce consumer inertia (Dubé, Hitsch, and Rossi 2010; Cabral and Villas-Boas 2005). A complete

[28] For two of the pairs (services A and E, and services E and H), we estimate λs that are more similar for the cardholders who subscribed to both. Yet, for the pair of service A and service H, we find that consumers who subscribe to both have λs that are equally far apart as those who subscribed to a single service. For the pair of services E and G (a pair that displays large difference in the estimate of λ), we find that the estimates of λs for those who subscribed to both closely resemble the baseline estimates. See Supplemental Appendix Figure B10.

[29] For example, if inertia stems from switching cost, a "forced" active renewal would result in many consumers paying their switching cost every month to renew their subscription.

welfare analysis that considers the equilibrium effects of reducing inertia is outside the scope of our analysis but is another interesting area to explore.

REFERENCES

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson. 2009. "The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation." *Brookings Papers on Economic Activity* 40 (2): 51–117.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai. 2020. "Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market." *American Economic Review* 110 (10): 3184–230.

Anonymous Firm. 2017–2021. *Payment and Transaction Data.* Unpublished data (accessed January 20, 2025).

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

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra. 2021. "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy." *American Economic Review* 111 (9): 2829–78.

Brot-Goldberg, Zarek, Timothy Layton, Boris Vabson, and Adelina Yanyue Wang. 2023. "The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D." *American Economic Review* 113 (10): 2718–58.

Cabral, Luis M. B., and Miguel Villas-Boas. 2005. "Bertrand Supertraps." *Management Science* 51 (4): 599–613.

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2009. "Optimal Defaults and Active Decisions." *Quarterly Journal of Economics* 124 (4): 1639–74.

DellaVigna, Stefano, and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence." *Quarterly Journal of Economics* 119 (2): 353–402.

DellaVigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review* 96 (3): 694–719.

de Silva, Tim. 2023. "Insurance versus Moral Hazard in Income-Contingent Student Loan Repayment." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4614108.

Di Geronimo, Linda, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. "UI Dark Patterns and Where to Find Them: A Study on Mobile Applications and User Perception." In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, edited by Regina Bernhaupt, Florian Mueller, David Verweij, and Josh Andres, 1–14. Association for Computing Machinery.

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi. 2010. "State Dependence and Alternative Explanations for Consumer Inertia." *RAND Journal of Economics* 41 (3): 417–45.

Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. 2018. "Auctions versus Posted Prices in Online Markets." *Journal of Political Economy* 126 (1): 178–215.

Einav, Liran, Ben Klopack, and Neale Mahoney. 2025. *Data and Code for "Selling Subscriptions."* American Economic Association; distributed by Inter-university Consortium for Political and Social Research. https://doi.org/10.3886/E216781V1.

Eliaz, Kfir, and Ran Spiegler. 2006. "Contracting with Diversely Naive Agents." *Review of Economic Studies* 73 (3): 689–714.

Eliaz, Kfir, and Ran Spiegler. 2008. "Consumer Optimism and Price Discrimination." *Theoretical Economics* 3 (4): 459–97.

Esteves-Sorenson, Constança, and Fabrizio Perretti. 2012. "Micro-costs: Inertia in Television Viewing." *Economic Journal* 122 (563): 867–902.

Federal Trade Commission. 2021. "Enforcement Policy Statement Regarding Negative Option Marketing." *Federal Register*, 2021-24094 (86 FR 60822). https://www.federalregister.gov/documents/2021/11/04/2021-24094/enforcement-policy-statement-regarding-negative-option-marketing.

Federal Trade Commission. 2024. "Negative Option Rule." *Federal Register*, 2024-25534 (89 FR 90476). https://www.federalregister.gov/documents/2024/11/15/2024-25534/negative-option-rule.

Fraccaroli, Giacomo, Neale Mahoney, and Zahra Thabet. 2024. "How Big Is the Subscription Cancellation Problem?" *Briefing Book*, September 9. https://www.briefingbook.info/p/how-big-is-the-subscription-cancellation.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021. "Five Facts about Beliefs and Portfolios." *American Economic Review* 111 (5): 1481–522.

Godderidge, Chris. 2021. "Everything Has Changed: The Rise of the Subscription and On-Demand Economies." Emarsys Blog, July 2. https://emarsys.com/learn/blog/the-rise-of-the-subscription-and-on-demand-economies/.

Goettler, Ronald L., and Karen Clay. 2011. "Tariff Choice with Consumer Learning and Switching Costs." *Journal of Marketing Research* 48 (4): 633–52.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.

Heiss, Florian, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou. 2021. "Inattention and Switching Costs as Sources of Inertia in Medicare Part D." *American Economic Review* 111 (9): 2737–81.

Jespersen, Andreas Maaløe. 2022. *Inactive Consumers in Subscription Markets*. Danish Competition and Consumer Authority.

Johnen, Johannes. 2019. "Automatic-Renewal Contracts with Heterogeneous Consumer Inertia." *Journal of Economics and Management Strategy* 28 (4): 765–86.

Luguri, Jamie, and Lior Jacob Strahilevitz. 2021. "Shining a Light on Dark Patterns." *Journal of Legal Analysis* 13 (1): 43–109.

Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116 (4): 1149–87.

Mathur, Arunesh, Gunes Acar, Michael J. Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. "Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites." In *Proceedings of the ACM on Human-Computer Interaction*, Vol. 3, edited by Airi Lampinen, Darren Gergle, and David A. Shamma, 1–32. Association for Computing Machinery.

Miller, Klaus M., Navdeep S. Sahni, and Avner Strulov-Shlain. 2023. "Sophisticated Consumers with Inertia: Long-Term Implications from a Large- Scale Field Experiment." Preprint, SSRN. https://dx.doi.org/10.2139/ssrn.4065098.

Posner, Nathaniel, Andrey Simonov, Kellen Mrkva, and Eric J. Johnson. 2023. "Dark Defaults: How Choice Architecture Steers Campaign Donations." *Proceedings of the National Academy of Sciences* 120 (40): e2218385120.

Reme, Bjørn-Atle, Helene Lie Røhr, and Morten Sæthre. 2021. "Inattention in Contract Markets: Evidence from a Consolidation of Options in Telecom." *Management Science* 68 (2): 1019–38.

West Monroe. 2021. *The State of Subscription Services Spending*. West Monroe.

Zuora. 2022. *The Subscription Economy Index: February 2022*. Zuora. https://www.zuora.com/wp-content/uploads/2022/09/Zuora_SEI_2022_BX-99-Update_to_SEI_2022.pdf.