Selling Subscriptions[*]

Liran Einav, Ben Klopack, and Neale Mahoney[†]

June 7, 2024

Abstract

Firms are increasingly selling products via subscriptions. In this paper, we study one benefit to firms of selling subscriptions: the prospect that consumers continue to pay for subscriptions they no longer value. We use comprehensive data from a large payment card network and focus on credit and debit cards that get replaced (e.g., due to expiration). Replaced cards require an active subscription renewal decision, and we document much higher cancellation rates in replacement months for the ten subscription services we study. We specify and estimate a stylized model of subscription renewals in which consumer inertia is driven by inattention. Relative to a counterfactual in which consumers are fully attentive, inattention raises seller revenues by 89% on average with substantial variation across the ten subscription services. We use the estimated model to explore the impact of possible regulatory remedies.

[*]We are grateful to Stefano DellaVigna and Jean-Pierre Dubé for 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.

[†]Einav: Stanford University and NBER, leinav@stanford.edu; Klopack: Texas A&M, bklopack@tamu.edu; Mahoney: Stanford University and NBER, nmahoney@stanford.edu. 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.

1 Introduction

A growing number of retail goods and services are now marketed as recurring subscriptions, which are typically billed at monthly frequencies. The goods and services are then provided indefinitely until a consumer actively cancels their subscription. While subscriptions have been pervasive 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 towards 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 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.[2] 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 company 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 a new 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.

To economically interpret these patterns and quantify the impact of inertia on seller

[1] A survey finds that 32% of US consumers "signed up to the subscription because it feels nice to receive something every month" (Emarsys 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).

[2] A recent survey found that nearly 90% of consumers underestimated their monthly spending on subscriptions, with the average respondent spending more than three times their initial estimate (West Monroe 2021).

revenues, we specify and estimate a stylized model of subscriber renewal behavior in which inertia arises from consumer inattention. In the model, subscribers are myopic and sign up for a service when their monthly flow utility from the service is greater than the monthly price. We assume that the flow utility follows an AR(1) process, and that 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."

We estimate the model separately for each subscription service. Despite having only three parameters for each service, the model replicates 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 inattention, with estimates of λ that vary from 0.04 to 0.50 across services.

We use the estimated model to perform counterfactual exercises that assess how much faster consumers would cancel their subscriptions if they were fully attentive ($\lambda = 1$). We find that seller revenues (or equivalently average subscription durations) are significantly higher due to subscriber inertia, with important heterogeneity across services. Specifically, we estimate that inertia increases seller revenues by 89% on average, with increases that range from 14% to more than 200% depending on the service.

We then use the estimated model 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. We find, for example, that requiring active choices at a 6-month frequency would reduce the excess revenue from inattention by 45%.

We explore sensitivity to our modeling choices by estimating an alternative model where inertia stems from switching costs rather than inattention. The switching cost model also closely fits the key patterns in the data and the estimated parameters are highly correlated with those from the inattention model. Counterfactual analysis with this model indicates that switching costs raise revenue by 37% on average across subscription services, which is lower than the inattention model but still substantial.

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). A more comprehensive analysis would also assess whether and to what extent consumers anticipate their future inertia when they sign up. 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, 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.[3]

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, which 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. (2022) on political contributions. Our paper also relates to a contemporaneous study by the Danish Competition and Consumer Authority (2022), which finds that subscribers whose payments are rejected are 70% more likely to cancel.[4,5]

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.

Lastly, 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 (2023).

The rest of the paper is organized as follows. In section 2 we describe the data, the selection of subscription services, and the construction of the sample. In section 3 we present descriptive evidence that motivates our key exercise and illustrates our empirical strategy. Section 4 presents the model, its estimation, and the counterfactual exercises, which allow

[3]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.

[4]Reme, Røhr, and Sæthre (2022) and Ascarza, Iyengar, and Schleicher (2016) study the dynamics of inattention and subscription attrition.

[5]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).

us to quantify how inertia affects firm revenues. The final section concludes.

2 Data and sample construction

2.1 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. Using publicly available information on total subscribers for the services in our baseline sample (described below), we estimate that our data covers approximately 30% 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.[6]

We imposed the following additional criteria to arrive at a final list of subscription ser-

[6]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.

vices. 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 4 services that are primarily sold in long-term contracts (two cell phone and two internet service providers), 6 services that were sold by merchants with many non-subscription products,[7] 2 services with average subscription length shorter than six months, 2 services that were launched toward the end of our observation period, and 2 services with non-monthly billing.

Our final sample is comprised of 10 large subscription services. We identify them throughout by letters (A through J) as our data use agreement prevents us from revealing the merchant names. The 10 services include both digital and non-digital products covering multiple merchant categories, including entertainment, security, retail goods, and newspapers.

The product offerings of these 10 services remained largely stable during our sample period with relatively small changes in pricing.[8] We investigated the consumer 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.

2.2 Sample construction

Our research design focuses on subscription renewal around card replacement. We thus limit our sample to consumer credit and debit accounts[9] 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% 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% of accounts). These restrictions leave us with about 19 million accounts and

[7] 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.

[8] 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.

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

approximately 28 million account-services.

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[10] after sign-up.[11] 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-services, representing 800,545 distinct accounts – of these 28 million account-services.[12]

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.[13] For a given service, we therefore observe a total of 57 cohorts: 19 cohorts defined by their sign-up months, each partitioned to three sub-cohorts that are based on the number of months (6, 12, or 18) at which card replacement occurs.[14]

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,[15] 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 "re-subscriptions").[16] Second, we exclude the

[10]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.

[11]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.

[12]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 non-replacement months. We therefore treat each account-service pair as an independent observation throughout the paper.

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

[14]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.

[15]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.

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

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.

3 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 its 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

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

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

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

where the $\widehat{\gamma}_{n,x}$'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 $\widehat{R}_n(x)$ for the same service ("service A") and x ($x = 12$), and 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.

Appendix Table A1 summarizes the drop in retention rates during card replacement, reporting the average monthly change in retention rates ($\widehat{R}_n(x) - \widehat{R}_{n-1}(x)$) 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, 4 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, Appendix Figure A1 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 pre-replacement 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.[17] 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, Appendix Figure A2 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% for groceries and -0.18% for gas, compared to -4.13% overall.

Cancellation rates around card replacement. Figure 2 presents the adjusted retention rates, $\widehat{R}_n(x)$, 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 a model of subscriber renewal behavior in which the sharp drops arise because card replacement forces otherwise inattentive consumers to make an active renewal choice. The model allows for multiple dimensions of heterogeneity across services, including the degree of inattention and the persistence of preferences in non-replacement months. We use the estimated model to quantify the impact of inattention on seller revenue, along with the revenue impact of counterfactual policies.

4 Quantifying the impact on revenues

Our primary objective is to estimate the revenue impact associated with the automatic renewal of subscription services. Because the focus is on the revenue impact for sellers rather than the utility impact on consumers, the model is static and highly stylized and

[17] 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.

should be viewed as a positive (rather than normative) description of renewal behavior. The key feature of the model we require for our exercise is the ability to simulate renewal behavior under alternative assumptions on consumer attention.

4.1 Baseline model

Our preferred behavioral model attributes the consumer inertia we document to inattention. Consider a specific subscription service, which is associated with a monthly subscription price p,[18] 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 $t = 0$ for the sign up month, and denote by $G(u_{i0}|u_{i0} > p)$ the cross-sectional distribution of u_{it} for new subscribers.

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.

Parameterization. We define the net flow utility as $v_{it} \equiv u_{it} - p$, and assume 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 – $G(u_{i0}|u_{i0} > p)$ or equivalently $G(v_{i0}|v_{i0} > 0)$ – is given by an exponential distribution, $v_{i0} \sim Exp(\eta)$, which has a mean and standard deviation η. Finally, we assume that the attention probability λ is the

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

same across people and over time (for a given service). We explore alternative assumptions in Section 4.4.

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 non-renewal probability of attention (λ). In 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 $\widehat{R}_n(x)$ 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.[19] 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 ($\sum_s N(s; s, x)$).

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

4.2 Model fit and parameter estimates

The parameter estimates are shown in Table 1. Appendix Figure A4 presents the model fit, 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

[19] For example, if card replacement occurs in month 6, Figure 2 shows the sharp drop in retention rates occurring over month 6 and month 7. By omitting the month-6 retention rate from the set of moments we match in estimation, we are essentially allowing the card replacement to occur in either month 6 or month 7.

parameters for each subscription service.[20]

The parameter estimates are reasonably intuitive to interpret. Consider first the inattention parameter λ. 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 0.044 (service I) to 0.503 (service G) across subscription services. For service I, almost all subscribers in a given month are inattentive and renew their subscription in a passive way, suggesting that inertia contributes significantly to the seller's revenues. This can be loosely 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% 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 1, implying that (on average) subscribers make an active decision only every two months on average.

The majority of the estimates of the ρ parameter are very close to 1, suggesting that preferences for the service are (on average) stable after initial subscription, and approximately follow a random walk. For several of the services, estimates of ρ are well below 1, implying that these services may find it difficult to retain consumers for longer durations. These are services which may benefit the most from inattentive consumers.

Finally, 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 if attentive (low η), or mostly infra-marginal 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 ($\eta = 0.004$) while service E, G, and H are associated with relatively high η estimates that are greater than 2.

4.3 Counterfactual exercises

The impact of inattention on revenues. We now use the model to quantify the impact of consumer inattention on seller revenues. To do so, we use the model and the estimated parameters described above, and simulate the renewal decisions of a cohort of initial subscribers over 10 years (that is, 120 months). Importantly, for this exercise we assume that subscribers face no card replacement throughout, so they are attentive each period with

[20] 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 a pattern that we intentionally do not aim to replicate with the model.

probability λ. We then repeat the same exercise but assume fully attentive subscribers (that is, $\lambda = 1$) and compare the results. 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 could be almost as costly as attracting "fresh" subscribers. We discount revenues at a rate of 1% per month.[21]

The results are shown in Table 2. The first column of the table shows the share of consumers that are unaffected by inattention. This group has positive valuations in every month and subscribes until the final period whether they are attentive or not. The remaining "affected" consumers have negative realizations of v_{it} at some point during the time period. The second column of the table shows the average number of months that this group subscribes 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. The last column reports the ratio of the revenues (measured by the number of subscriber-months) the seller obtains (over a horizon of 10 years) when subscribers are attentive with probability λ relative to the revenues it would obtain from the same set of initial subscribers if they were attentive with probability 1 every month.

Overall, we find that the benefits from attention (measured by the revenue ratio) are substantial, yet highly heterogeneous across subscription services. Relative to full attention, actual attention rates modestly increase revenues for some subscription services (e.g., 14% for G), but *triple* revenues for others (service B).[22] In other words, the average subscription duration for subscription B would drop from its observed duration of over a year to about 4 months if subscribers were fully attentive. It is plausible to suspect that this subscription service would not be viable from a business perspective if not for the inattention of its consumers.

Exploring the impact of possible remedies. A natural policy remedy to address subscriber inattention is to require subscribers to make an active decision every month. This could be implemented, for example, by changing the default so that subscriptions do not continue unless the consumer actively chooses to renew. Clearly, this is an extreme intervention, which would likely result in some consumer backlash; after all, a benefit of subscriptions is the convenience of *not* having to actively renew the purchase each month.

As a way to explore intermediate solutions, and in the spirit of recent Federal Trade Commission (2021) guidance, we consider the effect of possible policy interventions that re-

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

[22] Appendix Figure A5 provides a more complete mapping from the model parameters to the implied revenue ratio.

duce inattention by requiring subscribers to make an active choice periodically. The agency's Enforcement Policy Statement governing negative option marketing, which includes subscription products with automatic renewal, free trials that are converted to paid subscriptions, and similar sales strategies, requires firms to ask for informed consent for a subscription prior to billing, as well as to provide simple cancellation procedures. Specifically, we simulate the effect of making consumers fully attentive every m months, with $m \in \{1, 3, 6, 12, 18, 24\}$.

Figure 3 reports the revenue ratio associated with each subscription service, when "required" attention varies from monthly (fully attentive) to these lower frequencies. It illustrates that the revenue impact is meaningful (relative to fully attentive subscribers) even with an 18-month frequency. Requiring active choices every 6 months would reduce the overall revenue impact of inattention by 45%, nearly halfway between the revenue under observed inattention and the extreme of monthly active choice.

4.4 Robustness and heterogeneity

Robustness. To examine the sensitivity of our findings, we estimate several alternative specifications of the baseline inattention model. Panel A of Appendix Table A3 summarizes these results, reporting summary statistics for the estimates of λ and the revenue ratios across 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.

Heterogeneity across consumers. Appendix Figure A6 shows the retention rates separately for the 1.4% 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.[23] There is a sharper drop in retention rates around card replacement for consumers who take out cash advances, which is indicative of greater inattention among these consumers.

In Panel B of Appendix Table A3, we show results from re-estimating the baseline model for these two groups of consumers. Consistent with patterns shown in Appendix Figure A6,

[23] 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 is widely considered an expensive form of borrowing.

consumers who used cash advances are less attentive (lower λ) and have a revenue ratio that is nearly twice as large on average. Moreover, this pattern seems to (weakly) persist across all subscription services (Appendix Figure A7).

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 estimates vary substantially across the ten subscription services we study. This heterogeneity could reflect a combination of differences in 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 to specific factors.

To try to disentangle the role of consumer characteristics from subscription service characteristics, we separately analyze a subset of cardholders that 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 other consumers. Unfortunately, the results are mixed and do not allow us to draw strong conclusions.[24]

An alternative consumer model (switching costs). Our baseline model assumes that the inertia we document is driven by inattention. While this seems (to us) the most natural interpretation in this context, a model of switching costs could generate similar patterns. We thus explore the importance of this alternative model of consumer behavior for the quantitative results.

To do so, we hew as closely as possible to the baseline model, retaining the same parametric assumptions about the initial distribution of v_{i0} and its evolution over time. The key

[24] For two of the pairs (subscription A and subscription E, and subscription E and subscription 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 service E and service 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 Appendix Figure A8.

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. In months without card replacement, subscribers pay the switching cost to cancel, and so they renew if and only if $v_{it} + \kappa > 0$. 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 also assume that initial subscribers also face the switching cost when they sign up, which means that v_{i0} is drawn from $Exp(\eta)|v_{i0} \geq \kappa$ rather than from the unconditional distribution $Exp(\eta)$.

We estimate the switching cost model using the same moments as we did for the baseline model. The parameter estimates are reported in Appendix Table A4. As expected, the model fit is very similar (see Appendix Figure A9),[25] and the estimated parameters are highly correlated with those of the baseline model with correlation coefficients that are between 0.91 and 0.99 in absolute value (see Appendix Figure A10). Despite the similarity, the out of sample implications are a bit different, with revenue implications that are almost 60% lower than those from the baseline model (but still substantial). For example, in the baseline inattention model, subscriptions increase revenues by 89% on average (Table 2), while in the switching cost model subscriptions raise revenue by 37% (Appendix Table A5).

5 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.

Using a stylized model of consumer inattention, we estimate that consumer inertia leads to substantially higher subscription revenues, with increases ranging from 14% to more than 200% of the counterfactual revenue where subscribers were fully attentive and made active renewal decisions every month.

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.[26]

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

[26]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.

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 inattention by nearly 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 research.

We highlight two important caveats to interpreting 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 auto-renewal features of the contract. Combining the initial subscription and monthly renewal decisions would be a promising direction for future work. Second, given 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 welfare analysis that considers the equilibrium effects of reducing inertia is outside the scope of our analysis, but is another interesting area for future work.

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* 2009 (2): 51–117.

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.

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): 2718–2758.

Cabral, Luis MB, 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." *The quarterly journal of economics* 124 (4): 1639–1674.

Danish Competition and Consumer Authority. 2022. *Inactive Consumers in Subscription Markets.* Technical report. Available at https://www.kfst.dk/media/aipp013z/inactive-consumers_seneste.pdf. Last accessed November 30, 2023.

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

———. 2006. "Paying Not to Go to the Gym." *The American Economic Review,* 694–719.

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,* 1–14.

Dubé, Jean-Pierre, Günter J Hitsch, and Peter E Rossi. 2010. "State dependence and alternative explanations for consumer inertia." *The RAND Journal of Economics* 41 (3): 417–445.

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.

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

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

Emarsys. 2021. *Everything Has Changed: The Rise of the Subscription and On-Demand Economies.* Available at https://emarsys.com/learn/blog/the-rise-of-the-subscription-and-on-demand-economies/. Last accessed Feb 9, 2023.

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

Federal Trade Commission. 2021. *Enforcement Policy Statement Regarding Negative Option Marketing.* Technical report.

———. 2023. *Negative Option Rule: Notice of Proposed Rulemaking - 88 FR 24716.* Available at https://www.federalregister.gov/documents/2023/04/24/2023-07035/negative-option-rule.

Goettler, Ronald L, and Karen Clay. 2011. "Tariff choice with consumer learning and switching costs." *Journal of Marketing research* 48 (4): 633–652.

Handel, Benjamin R. 2013. "Adverse selection and inertia in health insurance markets: When nudging hurts." *American Economic Review* 103 (7): 2643–2682.

Johnen, Johannes. 2019. "Automatic-renewal contracts with heterogeneous consumer inertia." *Journal of Economics & Management Strategy* 28 (4): 765–786.

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." *The Quarterly journal of economics* 116 (4): 1149–1187.

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." *Proceedings of the ACM on Human-Computer Interaction* 3 (CSCW): 1–32.

Miller, Klaus M, Navdeep S Sahni, and Avner Strulov-Shlain. 2023. "Sophisticated Consumers with Inertia: Long-Term Implications from a Large-Scale Field Experiment." Available at SSRN: https://ssrn.com/abstract=4065098 or https://dx.doi.org/10.2139/ssrn.4065098. Last accessed Jun 7, 2023.

Posner, Nathaniel, Andrey Simonov, Kellen Mrkva, and Eric J Johnson. 2022. "Dark Defaults: How Choice Architecture Steers Campaign Donations." *Available at SSRN 4258478.*

Reme, Bjørn-Atle, Helene Lie Røhr, and Morten Sæthre. 2022. "Inattention in Contract Markets: Evidence from a Consolidation of Options in Telecom." *Management Science* 68 (2): 1019–1038.

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

Zuora. 2022. *The Subscription Economy Index.* Technical report. Available at https://www.zuora.com/resources/subscription-economy-index/. Last accessed Feb 9, 2023.

Figure 1: Aggregating retention rates

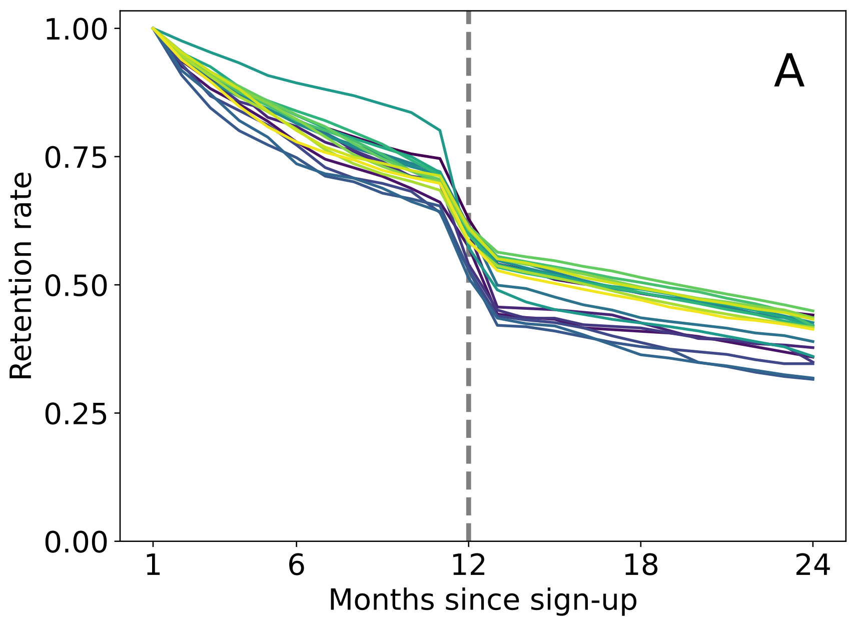

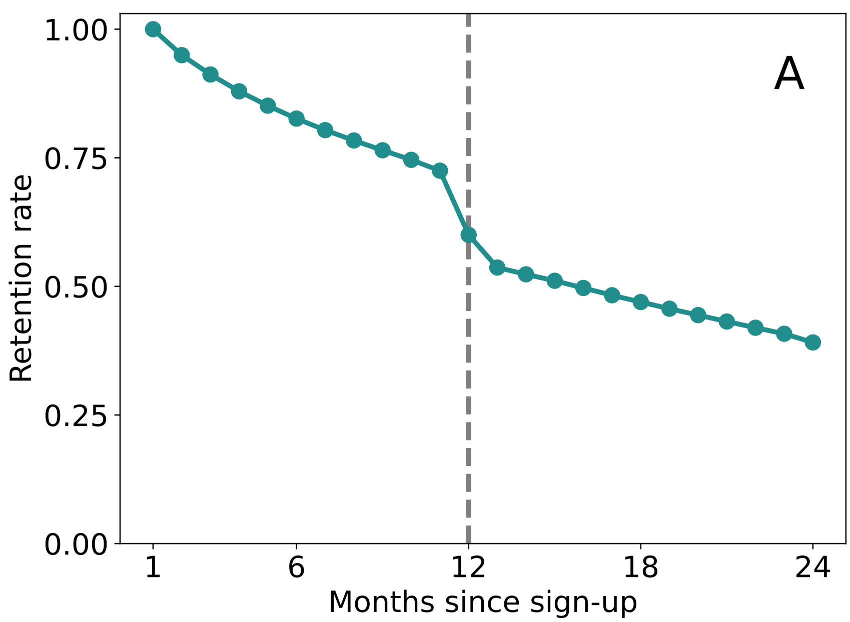

Note: Figure shows retention rates by month since sign-up for subscription service A and cohorts with card replacement 12 months after sign-up. The top panel shows the raw retention rates for all 19 cohorts. The bottom panel shows the adjusted retention rate, $\widehat{R}_n(x)$, which aggregates across cohorts netting out calendar-month fixed effects. See Section 3 for details.

Figure 2: Retention rate by months since sign-up, all subscription services and cohorts

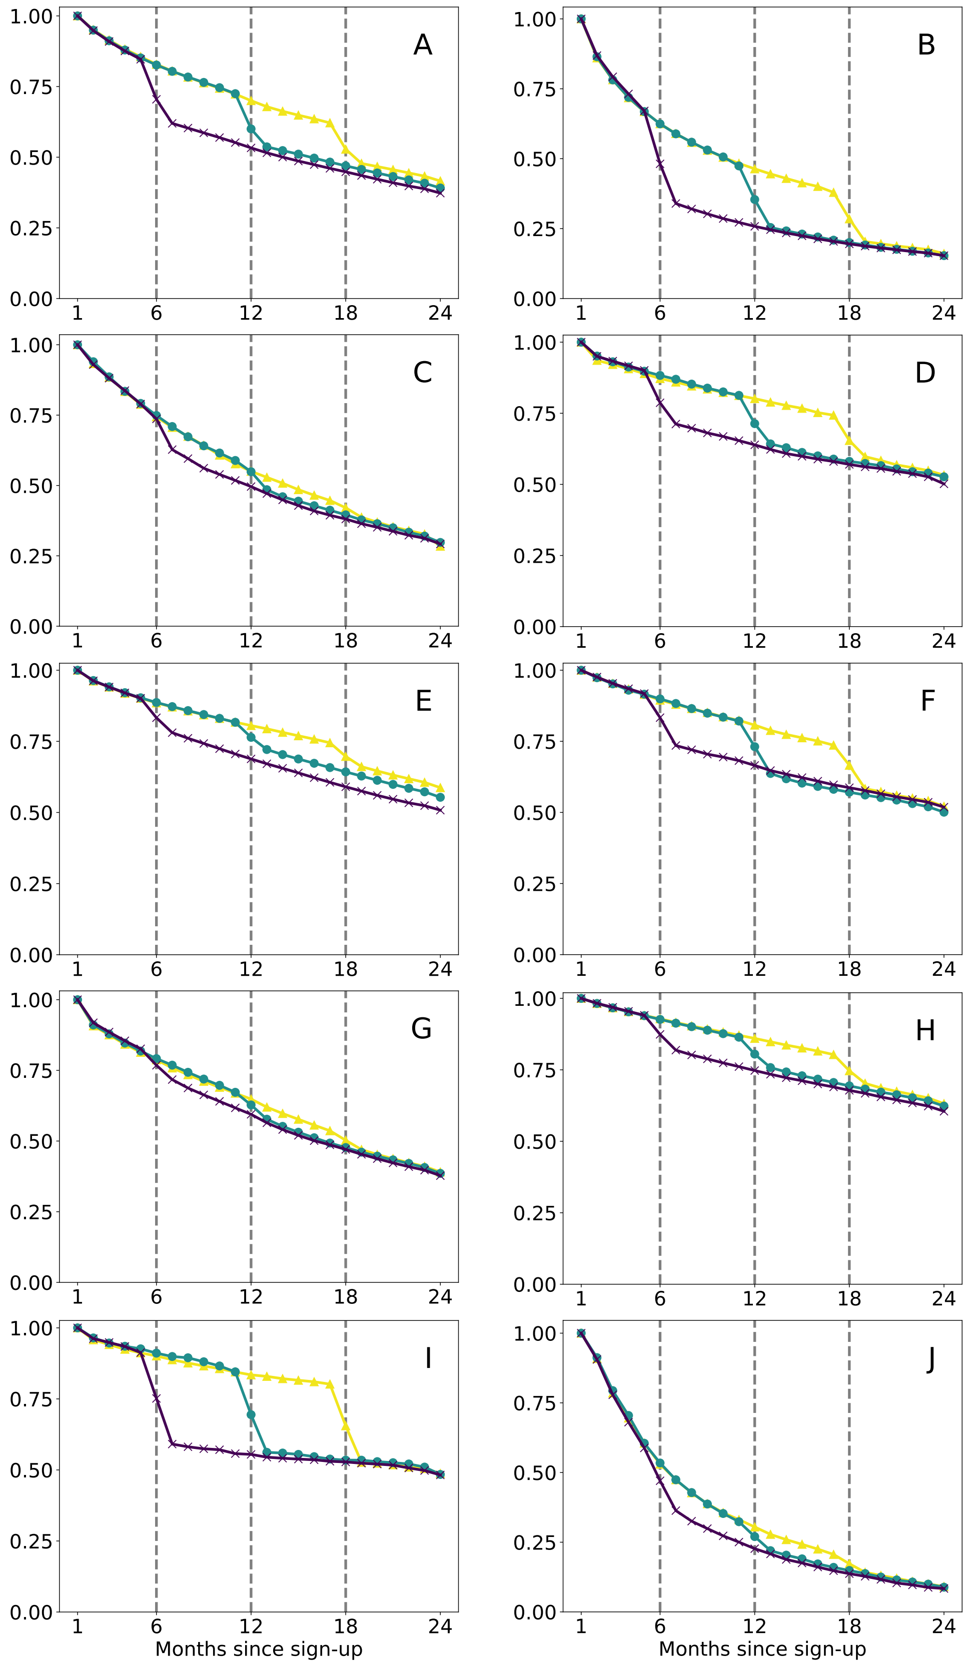

Note: Figure shows the adjusted retention rate, $\widehat{R}_n(x)$, by month since sign-up, separately by groups of cohort with card replacement at 6, 12, or 18 months after sign-up and subscription service (denoted by the letter in the top right corner of each panel). The adjusted retention rate aggregates across cohorts netting out calendar-month fixed effects. See Section 3 for details.

Figure 3: Revenue impact of required active choice

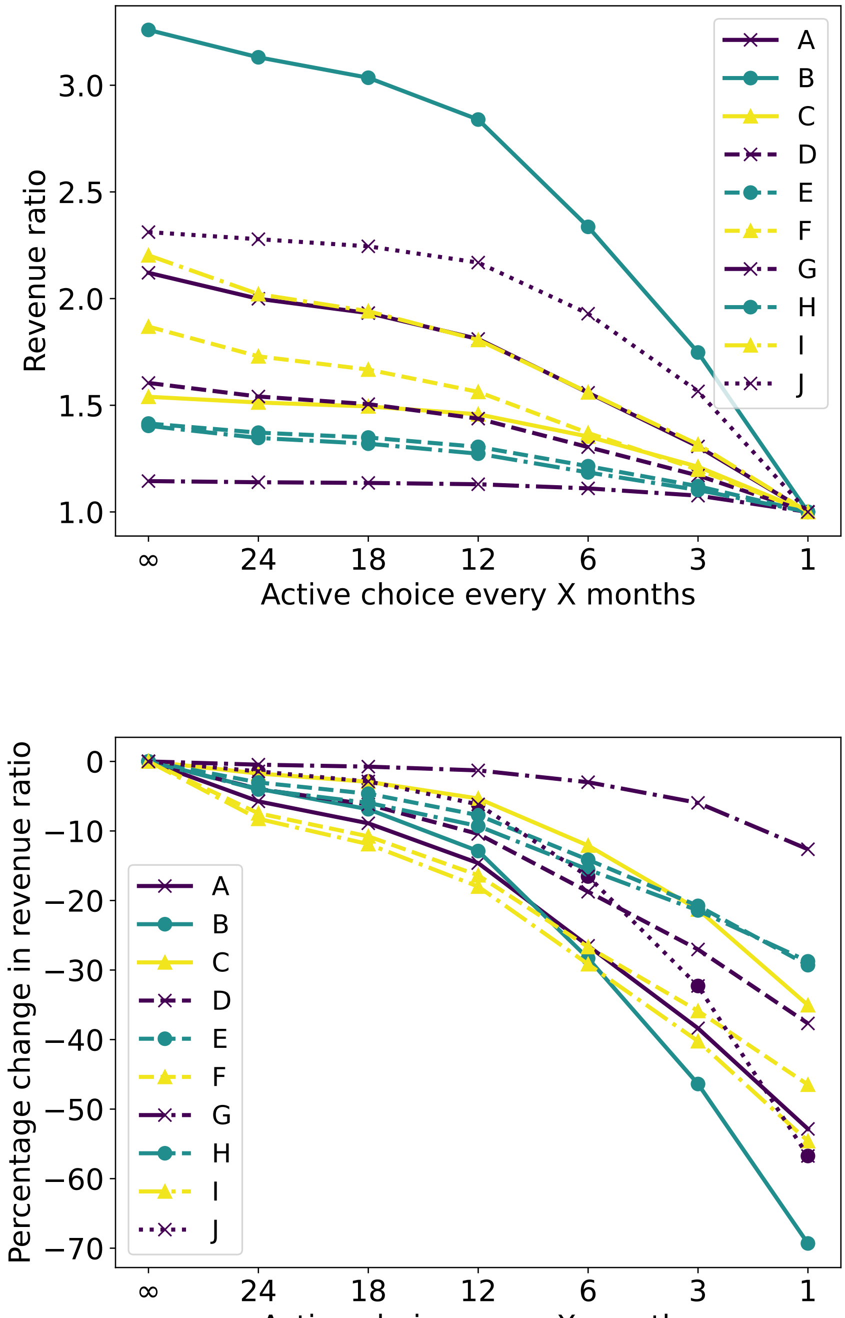

Note: Figure shows the revenue impact of requiring subscribers to make an active choice every 1, 3, 6, 12, 18, 24 months or never (∞) by subscription service. The top panel shows the ratio of revenue under a counterfactual where consumers make an active choice every X months to revenue under a counterfactual where subscribers are attentive every month. The bottom panel shows the percentage change in this ratio relative to the baseline of consumers never being required to make an active choice (∞). 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. 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%. The baseline is that subscribers are never required to make an active choice (∞) and only pay attention with probability λ every month. Appendix Table A2 provides the underlying numbers associated with the top panel.

Table 1: Parameter estimates by subscription service

Service	ρ	λ	η
A	0.985	0.110	0.858
	(0.003)	(0.002)	(0.025)
B	0.845	0.182	0.004
	(0.011)	(0.002)	(0.000)
C	0.947	0.264	1.530
	(0.006)	(0.014)	(0.098)
D	1.019	0.089	1.377
	(0.014)	(0.006)	(0.143)
E	0.998	0.128	2.500
	(0.002)	(0.003)	(0.049)
F	0.980	0.092	1.913
	(0.007)	(0.004)	(0.118)
G	0.943	0.503	3.900
	(0.003)	(0.000)	(0.072)
H	1.000	0.097	3.122
	(0.003)	(0.002)	(0.080)
I	1.053	0.044	0.467
	(0.042)	(0.003)	(0.105)
J	0.822	0.281	0.163
	(0.011)	(0.021)	(0.119)

Note: Table reports parameter estimates and bootstrapped standard errors (in parentheses) for our model of subscription renewal described in Section 4. We estimate the model separately for each of the 10 subscription services A through J. The standard errors are the standard deviations of the 1,000 bootstrap estimates.

Table 2: Revenue impact of inattention

Service	Share unaffected	Avg months subscribed		Revenue ratio
		If inattentive	if attentive	
A	0.04	36.8	13.7	2.12
B	0.00	14.0	4.0	3.26
C	0.00	22.0	13.6	1.54
D	0.28	42.1	9.3	1.60
E	0.20	42.3	20.4	1.41
F	0.04	45.3	19.9	1.87
G	0.00	23.6	20.2	1.14
H	0.26	48.4	21.8	1.40
I	0.25	54.8	4.0	2.20
J	0.00	10.0	4.1	2.31
Mean	**0.11**	**33.9**	**13.1**	**1.89**

Note: Table reports our counterfactual estimates on how inattention affects firm revenue. For each subscription service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers for 120 months after sign-up. Column (1) reports the share of subscribers not affected by inattention because they have a positive valuation in every month. Columns (2) and (3) show, for subscribers with a negative valuation in at least one month, the average number of months they are subscribed if they are inattentive (paying attention with probability λ each month) or attentive (paying attention with probability 1 each month), respectively. Column (4) shows the revenue ratio from inattention 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 are required to make an active choice every month. The numerator is the discounted sum of monthly subscribers if consumers make an active choice with probability λ. We discount future revenues at a rate of 1%.

Online Appendix

Appendix A: Comparative statics and identification

We use the exercise shown in Appendix Figure A3 to build intuition for the identification of the model and illustrate its comparative statics. The figure uses the data and the estimated parameters associated with subscription A.

The dashed lines in all panels report the model predictions for retention rates when the card replacement is in month 6 (left panels), month 12 (middle panels), and month 18 (right panels). We then change one parameter at a time (holding the other two at their estimated values) in order to hit an (arbitrary) retention rate of 0.56 by month 24, and explore how this affects the retention rates in earlier months. The two solid lines in each panel of Appendix Figure A3 show this exercise for a pair of parameters to facilitate comparison.

Consider for example the bottom left panel. In order to hit a retention rate of 0.56 in month 24 by *only* changing λ, the value of λ (that is, the attention level of subscribers) must be lower, leading to a fairly flat decline in retention rates before and after card replacement, and to a sharp drop in the retention rate in the month of card replacement. In contrast, if we wanted to hit the same level of retention rate in month 24 by *only* changing η, the value of η would have to be much higher so that more initial subscribers renew their subscription. Yet, this change in the parameter value would still predict a steep decline in retention rates before and after card replacement, and make the drop in retention rate in the month of card replacement much smaller.

Similar contrasts are illustrated in the other panels of the figure, which may help provide intuition for the separate identification of the three parameters. A flatter slope in retention rates before and after card replacement would load on λ, a steeper one would load on ρ and η, and variation in how retention rates change before versus after card replacement helps distinguish between the latter two (as illustrated in the plots in the middle row).

Appendix B: Simulating model predictions

Let $\varphi = (\rho, \lambda, \eta)$ be the set of candidate parameter values. In order to construct model predictions for retention rates as a function of these parameters, we simulate a large panel of initial subscribers and record their renewal decisions as given by the model.

For each subscription service, and for each number of months between initial subscription and card replacement $x \in \{6, 12, 18\}$, we simulate three sets of $N = 100{,}000$ subscribers, and then simulate their renewal decisions as a function of the model and the parameters given by φ.

Specifically, we start by drawing the random components of the model, which include: (i) an $N \times 1$ vector of draws that will affect initial valuations – iid draws from a $[0, 1]$-uniform distribution – denoted by $u_0(x)$; (ii) an $N \times T$ matrix of taste shocks – iid draws from a standard normal distribution – denoted by $e_0(x)$; and (iii) an $N \times T$ matrix of "attention shocks" – iid draws from a $[0, 1]$-uniform distribution – denoted by $l_0(x)$.

We then construct the vector of initial values, $v_0(x)$, by transforming u_0 into a variable with exponential distribution and mean η using $v_0(x) = -\eta \log(1 - u_0(x))$. $v_0(x)$ is then an $N \times 1$ vector of initial net flow utilities in month $s = 1$. Then, using $v_0(x)$ and $e_0(x)$, we simulate the net flow utility for all subsequent periods using ρ and the model assumption that $v_t(x) = \rho v_{t-1}(x) + e_0(x)$.

We then say that subscriber i pays attention in month s if and only if $l_0(x) \leq \lambda$. Denote by $a(x)$ the resulting $N \times T$ matrix of binary attention indicators, and assign $a(x) = 1$ for month $s = x$ for all subscribers.

Finally, we construct the $N \times T$ subscription matrix, $s(x)$. As all individuals are subscribers in $s = 1$, we have that $s(x)_{i1} = 1$ for all i. For $t > 1$, $s(x)_{it} = 0$ if $s(x)_{i,t-1} = 0$ or if $a(x)_{it} = 1$ and $v(x)_{it} < 0$. The simulated retention rates for each month s are then given by the share of 1s in each column of the subscription matrix $s(x)$, which we denote by $R(x; \varphi)$.

Appendix Figure A1: Account activity around replacement date

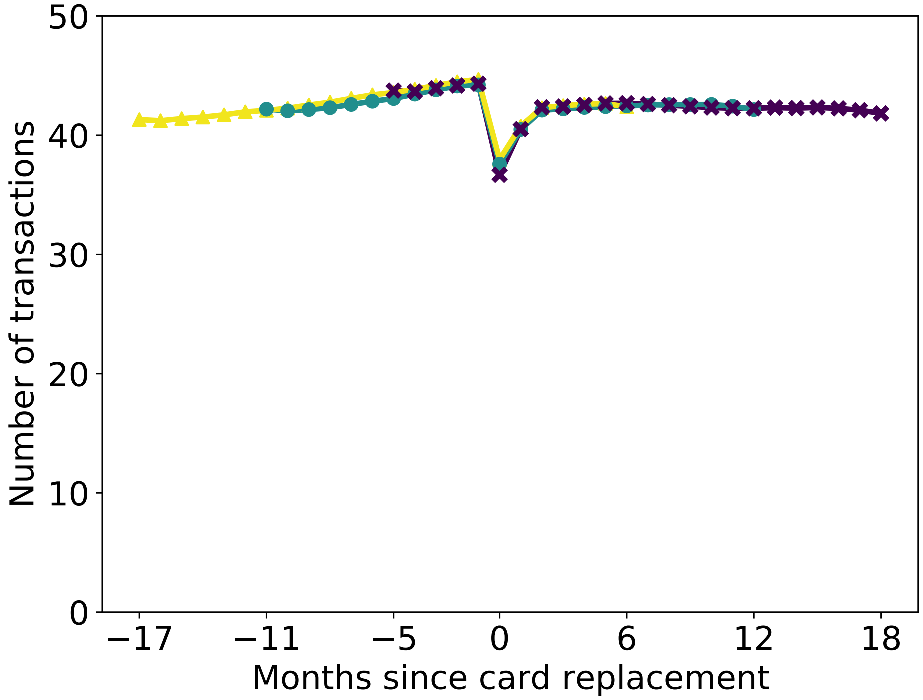

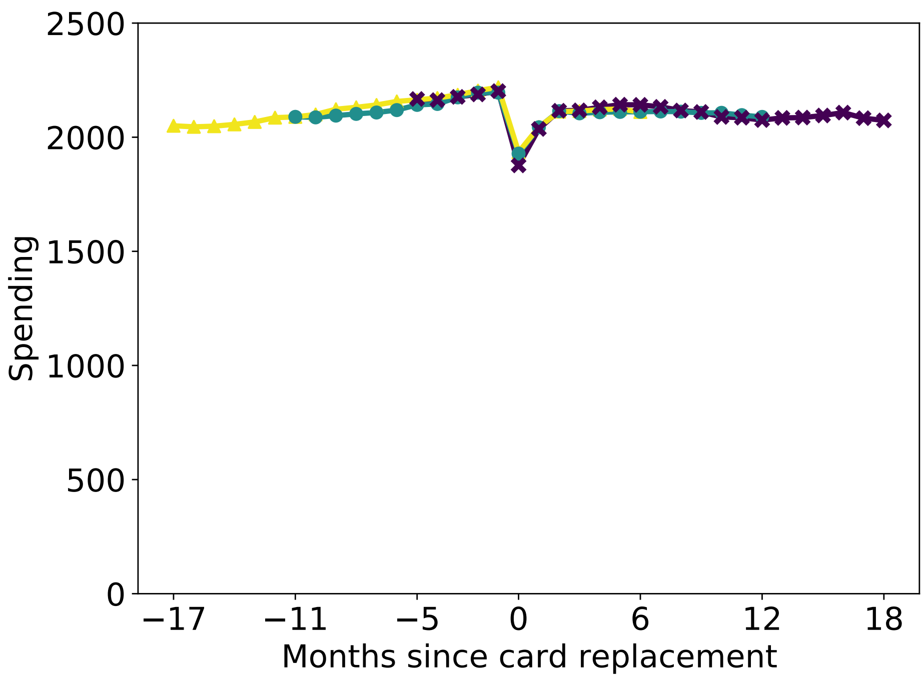

Note: Figure shows account activity, measured by number of transactions (top panel) and spending (bottom panel) per month, around the month of card replacement. To construct the figure, we calculate transactions and spending in each month, including transactions at the ten subscription services that we study. We regress this on account and month fixed effects and average the residuals by months from the date of card replacement. The plot shows the average of the residuals plus the mean number of transactions and spending per month computed across the entire sample for our three groups of cohorts (accounts that subscribe to one of the ten services 6, 12, or 18 months prior to the date of card expiration.)

Appendix Figure A2: Account activity around replacement date for groceries and gas

Groceries

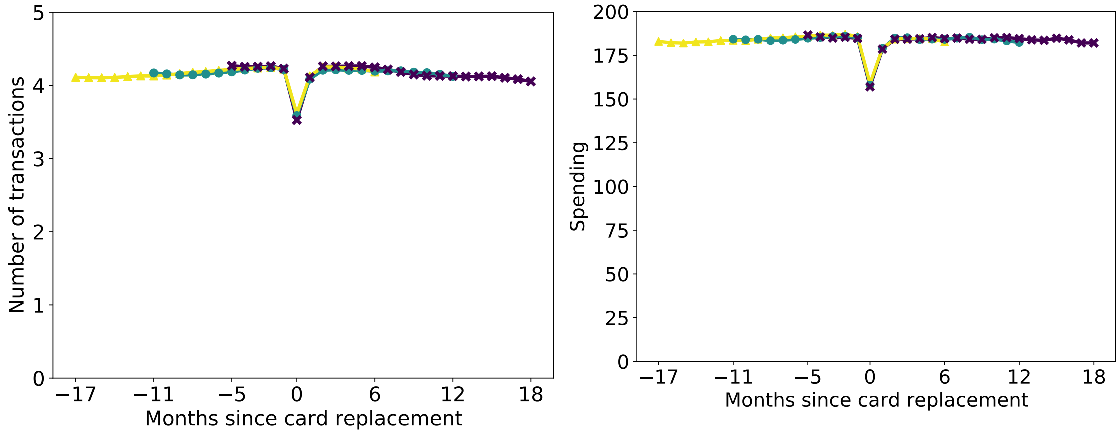

Gas

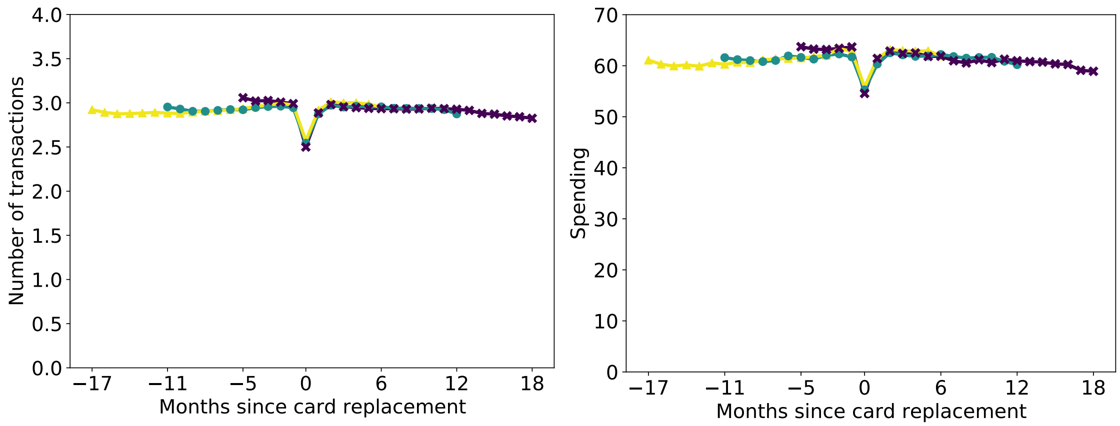

Note: Figure shows account activity, measured by number of transactions (left panel) and spending (right panel) per month, around the month of card replacement for gas and grocery purchases. To construct the figure, we calculate transactions and spending in each month by store category. We regress this on account and month fixed effects and average the residuals by months from the date of card replacement. The plot shows the average of the residuals plus the mean number of transactions and spending per month computed across the entire sample for our three groups of cohorts (accounts that subscribe to one of the ten services 6, 12, or 18 months prior to the date of card expiration.)

Appendix Figure A3: Comparative statics

Holding $\eta = 0.858$ fixed:

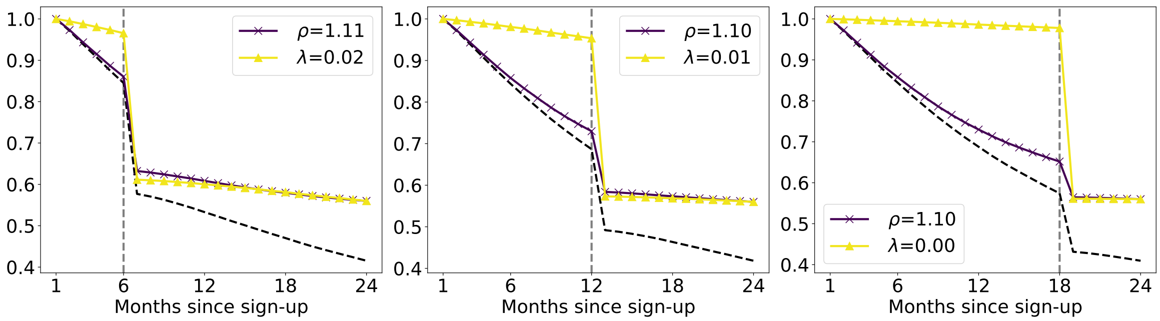

Holding $\lambda = 0.110$ fixed:

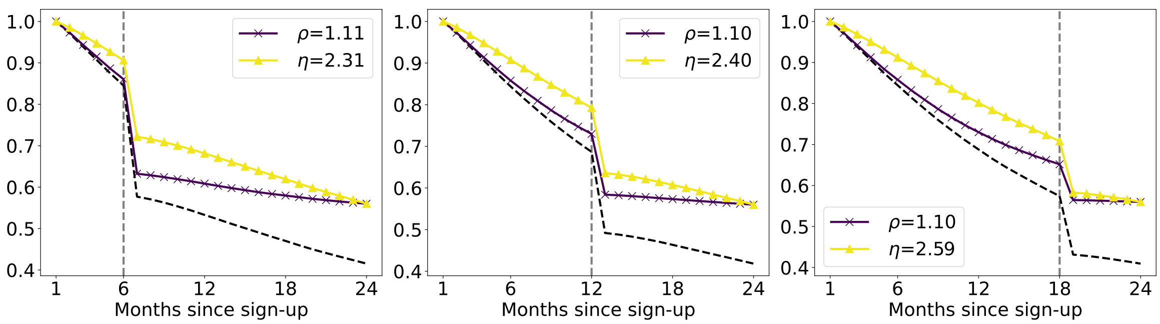

Holding $\rho = 0.985$ fixed:

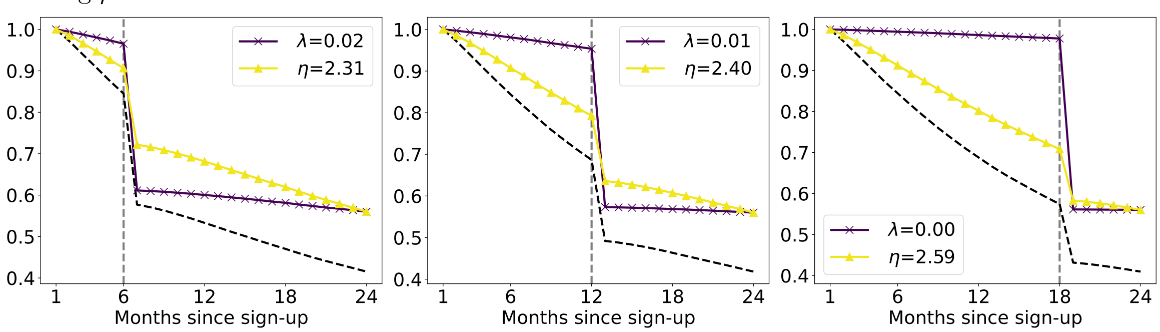

Note: Figure shows the comparative statics of our model of subscription renewal behavior using the data and estimated parameters for service A. Each panel shows model-predicted retention rates for service A on the y-axis, by month since sign-up on the x-axis. We consider the three groups of cohorts separately, with card replacements after 6 months (left panels), 12 months (middle panels), and 18 months (right panels). In each panel, the vertical line represents the month of card replacement and the black dashed line shows the model-predicted retention rates implied by our parameter estimates $\rho = 0.985$, $\lambda = 0.110$, and $\eta = 0.858$ (see Table 1). In each panel, the colored, solid lines illustrate the trade-off between pairs of model parameters, (ρ, λ), (ρ, η), and (λ, η), in matching the (arbitrary) target retention rate of 0.56 in the last month. In the top row, we keep $\eta = 0.858$ fixed and show, separately for each group of cohorts, how either ρ (also holding $\lambda = 0.110$ fixed) or λ (also holding $\rho = 0.985$ fixed) has to change to attain the target retention rate. For example, in the top-left panel, we attain this retention rate for the cohort of card replacements after 6 months by increasing ρ to $\rho = 1.11$, while holding $\lambda = 0.110$ and $\eta = 0.858$ fixed, or by decreasing λ to $\lambda = 0.02$, while holding $\rho = 0.985$ and $\eta = 0.858$ fixed. In the middle row, we hold $\lambda = 0.110$ fixed and show the trade-off between ρ and η. In the bottom row, we hold $\rho = 0.985$ fixed and show the trade-off between λ and η.

Appendix Figure A4: Model fit

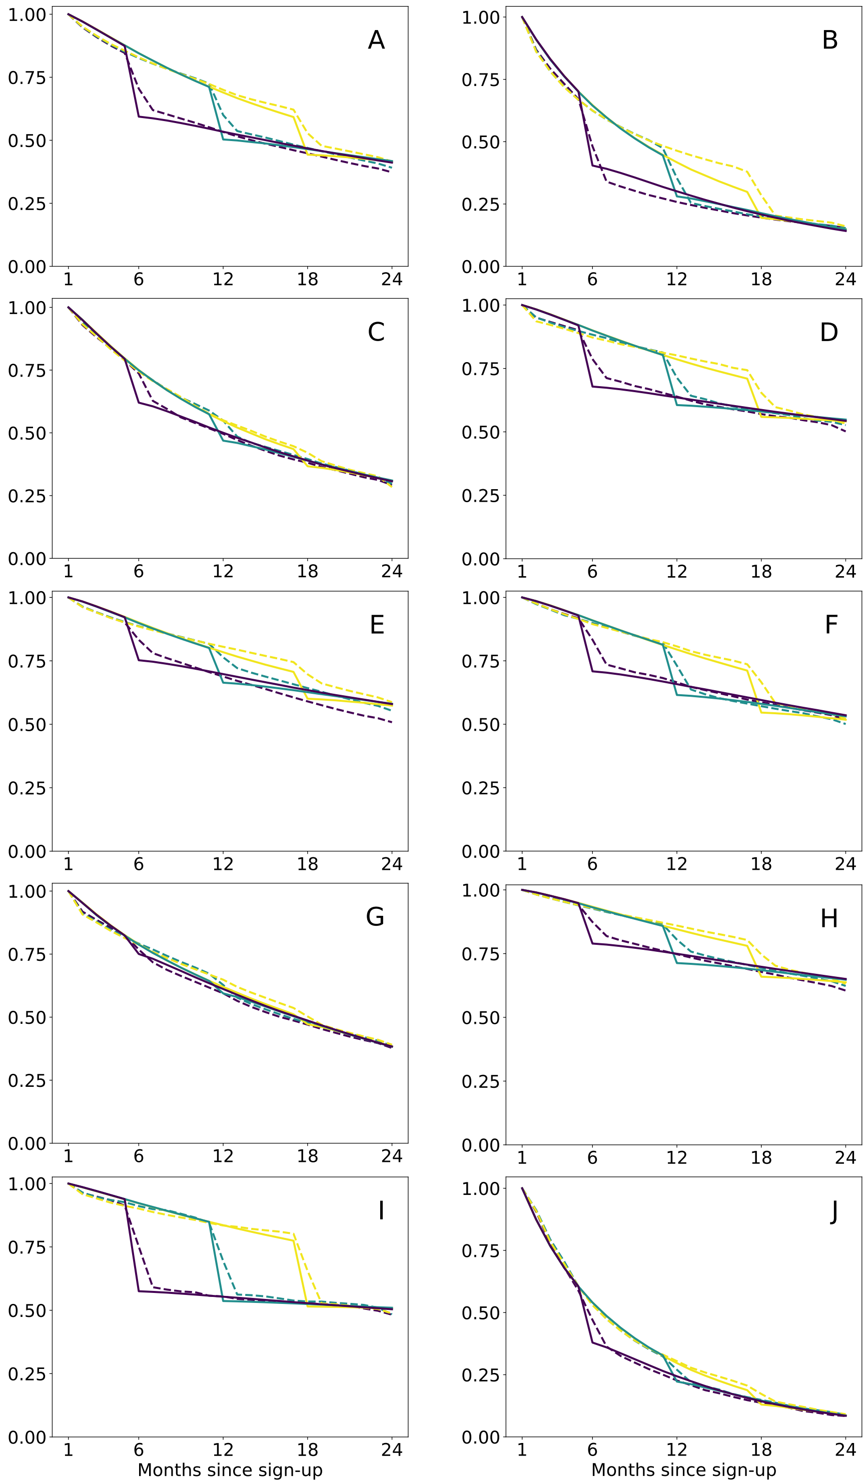

Note: Figure shows the observed (adjusted) retention rate $\widehat{R}_n(x)$ (dashed lines) and the predicted retention rate from our inattention model (solid lines) by month since sign-up, separately by group of cohort with card replacement 6, 12, or 18 months after sign-up and subscription service (denoted by the letter in the top right corner of each panel).

Appendix Figure A5: Iso-ratio curves

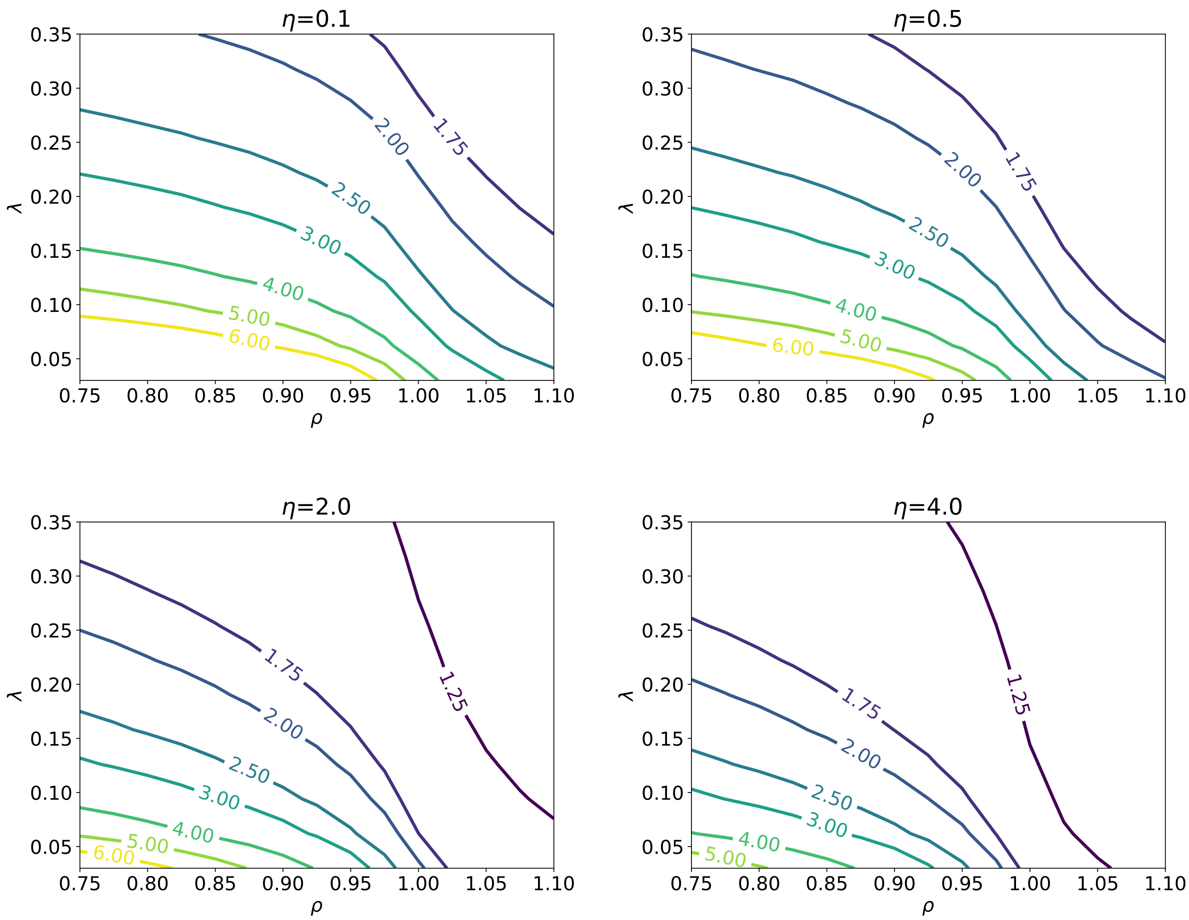

Note: Figure shows the revenue ratio as a function of the three model parameters (ρ, λ, and η). Given η, each line shows all (ρ, λ) combinations that yield the same revenue ratio. For given parameter values, we construct the revenue ratio as follows: For each subscription service, we simulate the monthly 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. The numerator is the discounted sum of monthly subscribers if consumers make an active choice with probability λ. We discount future revenues at a rate of 1%.

Appendix Figure A6: Heterogeneity by cash advance

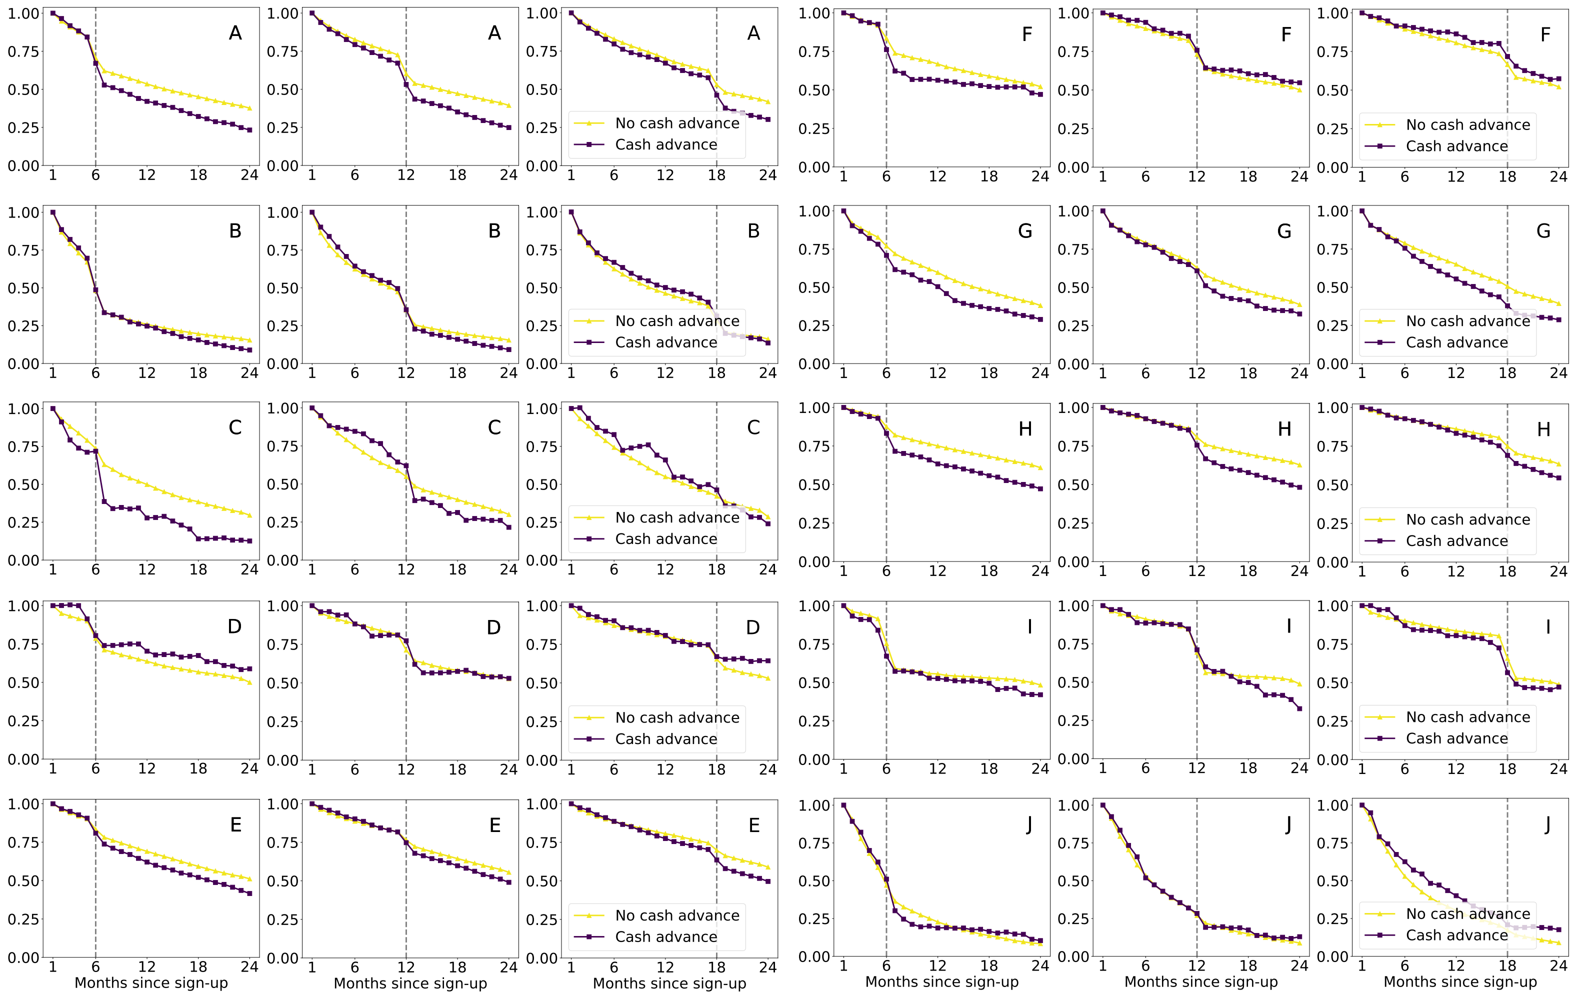

Note: Figure shows the adjusted retention rates by month since sign-up for cards with and without cash advance and by group of cohorts with card replacement 6, 12, or 18 months after sign-up. The letter in the top right corner of each panel identifies the subscription service.

Appendix Figure A7: Estimation results by cash advance

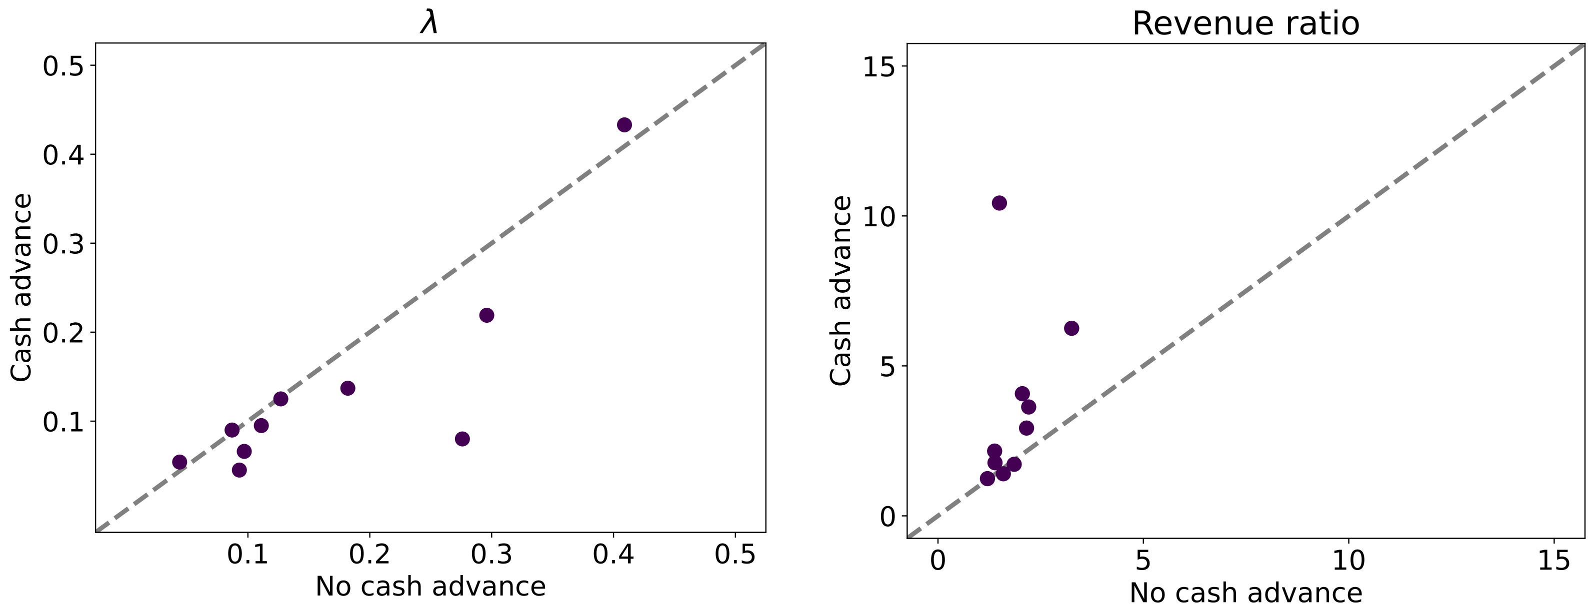

Note: The plot on the left shows the estimate of λ for each subscription service for cards with a cash advance (y-axis) against the estimate for cards without cash advance (x-axis). The plot on the right shows the revenue ratio for cards with a cash advance against cards without a cash advance. The dashed line is the 45-degree line.

Appendix Figure A8: Estimates of λ from multi-service subscribers

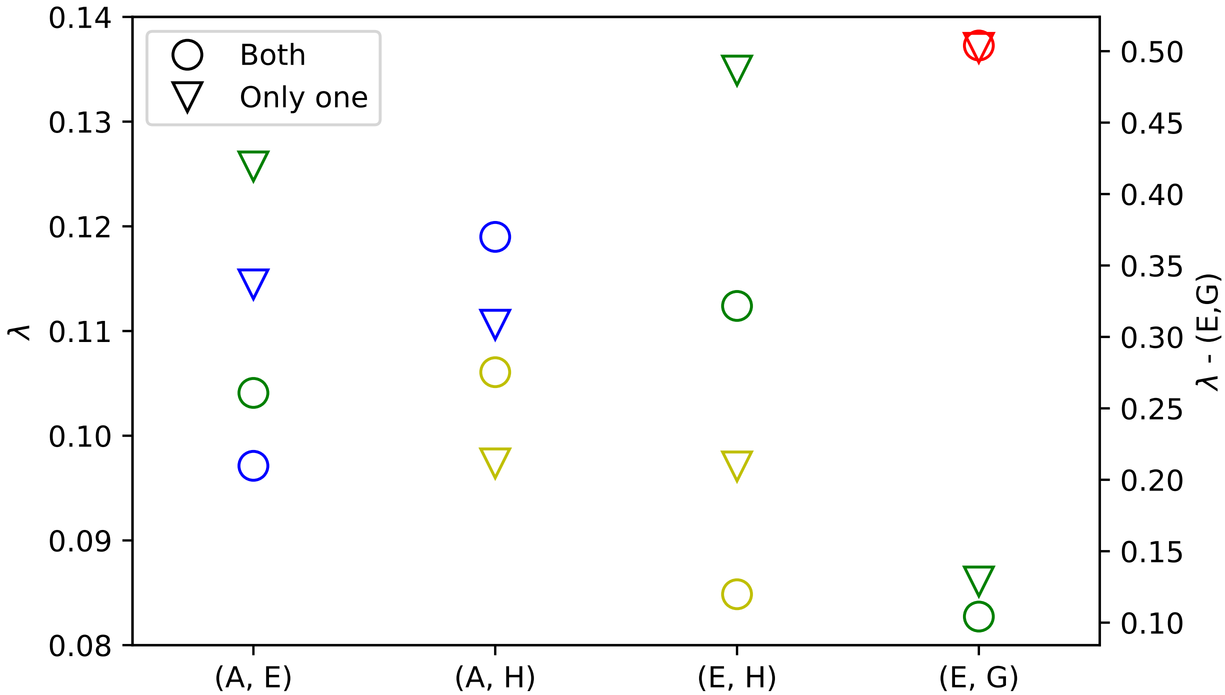

Note: Figure shows estimates of λ from consumers that subscribed to multiple services among the set we study. We construct the figure by taking the most common four pairs of services (ranked by number of subscribers that had both) and estimating the inattention model separately for cardholders that had both vs. those that had only one. Points on the graph are color-coded by service (A in blue, E in green, H in yellow, and G in red). The marker shape corresponds to the consumer group – circles denote cardholders that subscribed to both in the pair, while triangles correspond to those that had only one. The pair (E,G) are plotted on the right y-axis, while the other three pairs are plotted on the left.

Appendix Figure A9: Model fit, switching-cost model

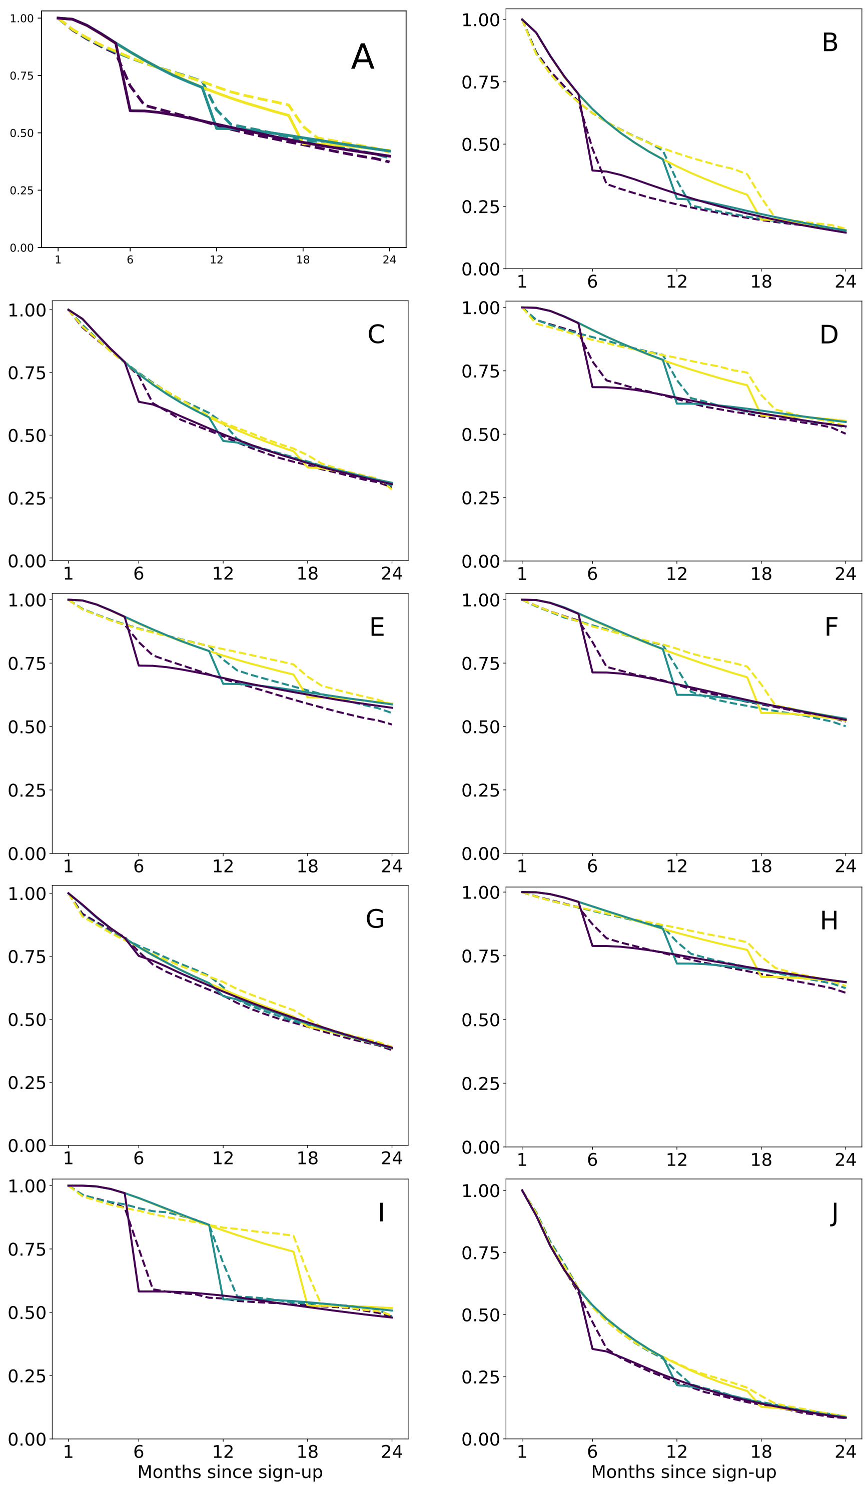

Note: Figure shows the observed (adjusted) retention rate $\widehat{R}_n(x)$ (dashed lines) and the predicted retention rates from our switching cost model (solid lines) separately by group of cohorts with card replacement 6, 12, or 18 months after sign-up and subscription service letter (shown in the top right corner of each panel).

Appendix Figure A10: Estimates comparison, inattention vs. switching costs

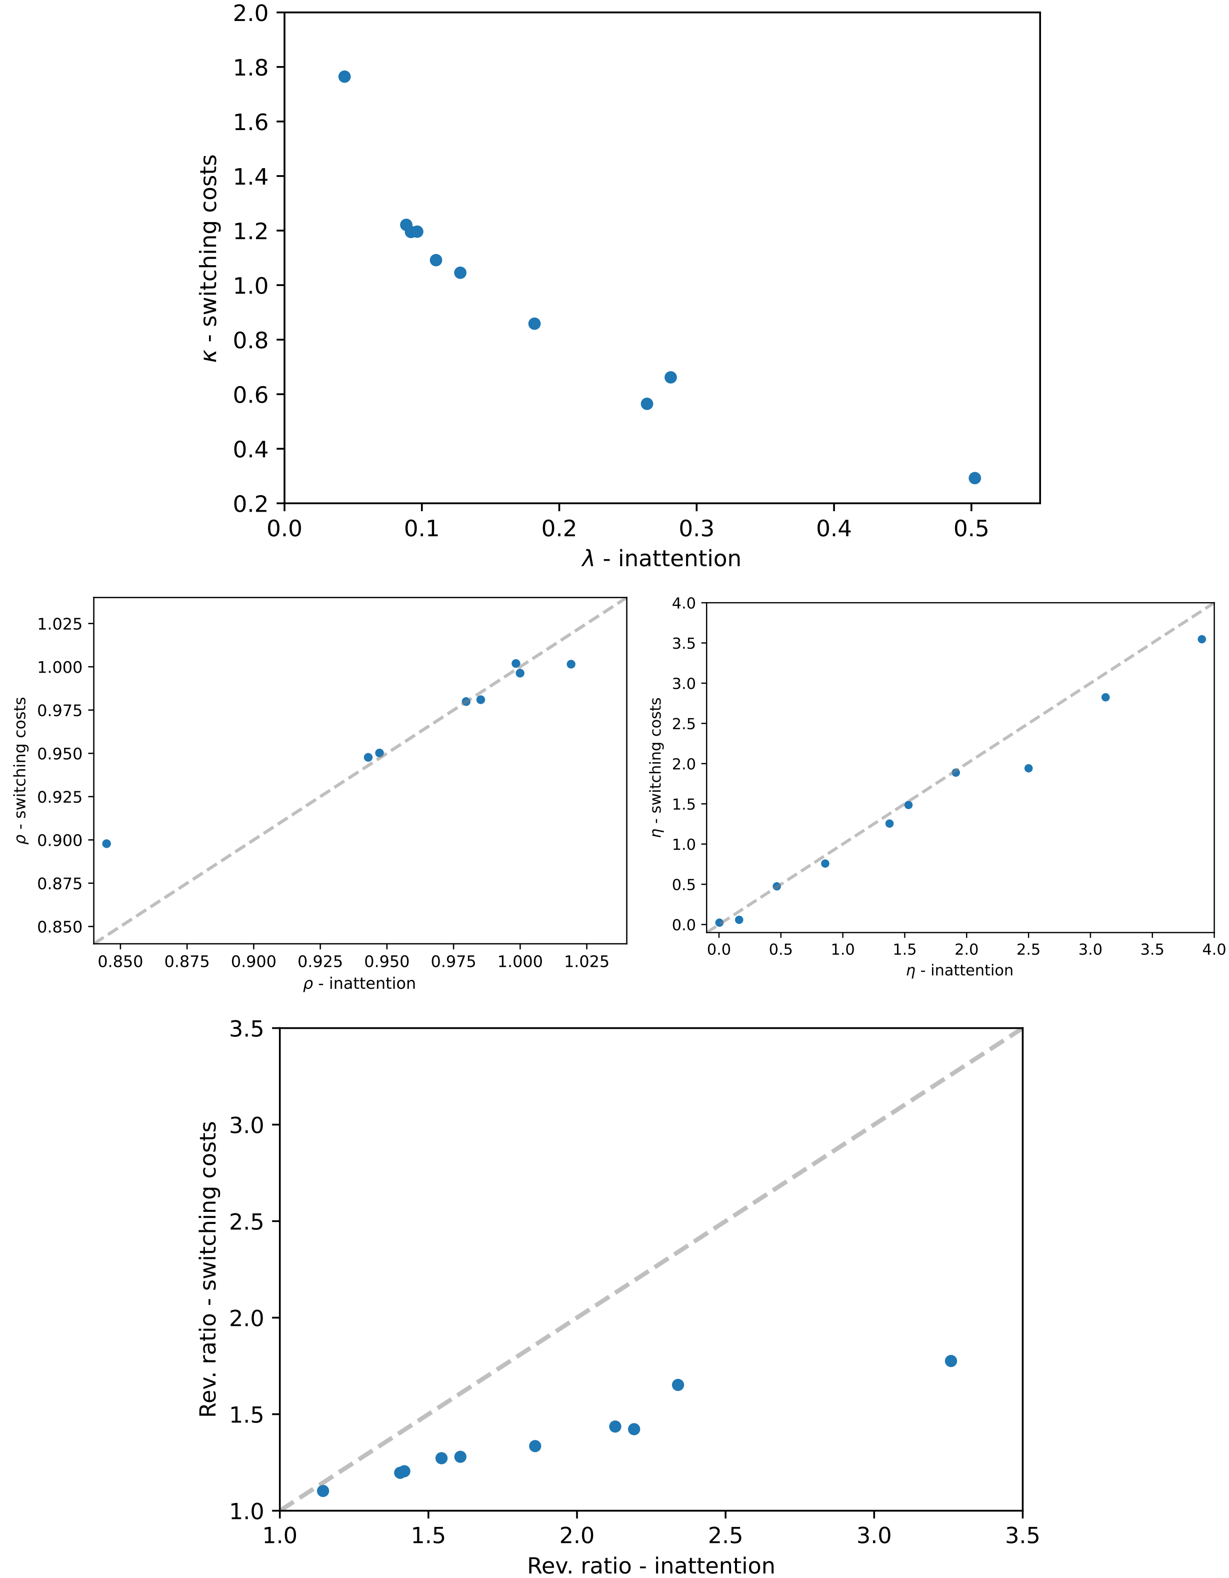

Note: Figure shows our estimates of the three parameters and the counterfactual revenue ratio from the inattention and switching cost models. Each dot represents a subscription service. The correlation coefficients in each plot are: -0.908 (λ/κ), 0.970 (ρ), 0.993 (η), and 0.970 (the revenue ratio).

Appendix Table A1: Average change in retention rates

	Average monthly change in retention rate	
Service	Replacement months	Other months
A	-0.09	-0.02
B	-0.12	-0.03
C	-0.05	-0.03
D	-0.08	-0.01
E	-0.05	-0.02
F	-0.09	-0.01
G	-0.05	-0.03
H	-0.05	-0.01
I	-0.15	-0.01
J	-0.07	-0.04
Mean	-0.08	-0.02

Note: Table shows the average monthly change in retention rates during months of card replacement and other months by subscription service. For each service and cohort ($x = 6, 12, 18$), we compute the change in adjusted retention rates, $\widehat{R}_n(x) - \widehat{R}_{n-1}(x)$. For each subscription service, then we average the change in retention rates across months and cohorts, separately for months with and without card replacement. The last line shows the average monthly change in retention rates averaged across the ten subscription services.

Appendix Table A2: Revenue impact of possible regulatory remedies

	Active choice every X months						
Service	∞	24	18	12	6	3	1
A	2.12	2.00	1.93	1.81	1.56	1.31	1.00
B	3.26	3.13	3.04	2.84	2.34	1.75	1.00
C	1.54	1.51	1.49	1.46	1.35	1.21	1.00
D	1.60	1.54	1.51	1.44	1.30	1.17	1.00
E	1.41	1.37	1.35	1.31	1.21	1.12	1.00
F	1.87	1.73	1.67	1.56	1.37	1.20	1.00
G	1.14	1.14	1.14	1.13	1.11	1.08	1.00
H	1.40	1.35	1.32	1.27	1.19	1.10	1.00
I	2.20	2.02	1.94	1.81	1.56	1.32	1.00
J	2.31	2.28	2.25	2.17	1.93	1.57	1.00
Mean	1.89	1.81	1.76	1.68	1.49	1.28	1.00

Note: Table summarizes the revenue impact of possible regulatory remedies as shown in the top panel of Figure 3. That is, the revenue impact of requiring subscribers to make an active choice every 1, 3, 6, 12, 18, 24 months or never (∞). Each column shows the ratio of revenue under a counterfactual where consumers make an active choice every X months to revenue under a counterfactual where subscribers are attentive 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. 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%. The baseline is that subscribers are never required to make an active choice (∞) and only pay attention with probability λ every month.

Appendix Table A3: Robustness and heterogeneity

	Lambda					Revenue ratio				
	Mean	St. Dev.	2nd	9th	Corr. w/ baseline	Mean	St. Dev.	2nd	9th	Corr. w/ baseline
Baseline	0.18	0.13	0.09	0.28		1.89	0.59	1.40	2.31	
A. Robustness										
Linear decay of lambda	0.16	0.09	0.09	0.28	0.94	2.88	3.44	1.26	2.69	0.90
Linear decay of lambda with reset	0.18	0.14	0.10	0.28	0.98	2.65	1.98	1.39	2.84	0.94
Lambda at card expiration = 0.75	0.18	0.14	0.07	0.36	1.00	2.34	1.26	1.49	3.68	0.68
Lambda at card expiration = 0.5	0.18	0.15	0.04	0.37	0.99	4.39	4.39	1.55	7.70	0.35
B. Heterogeneity										
Never used cash advance	0.17	0.11	0.09	0.30		1.86	0.57	1.38	2.21	
Used cash advance	0.13	0.11	0.05	0.22		3.56	2.71	1.40	6.25	

Note: Table summarizes the results of our robustness and heterogeneity analyses for the estimates of λ and the revenue ratio. The first row summarizes our baseline estimates. We report the mean and standard deviation of the respective estimates across the ten subscription services, as well as the second and ninth value of the estimates if sorted in ascending order. Panel A summarizes the results of our robustness analysis, where we modeled alternative specifications for the inattention process. The first row summarizes the results of a specification that allows λ to vary linearly in time since subscription (i.e., $\lambda_t = \lambda_0 + \theta t$). In the second row, the model again allows λ to vary linearly, but we assume that λ_t "resets" to λ_0 after card expiration (i.e., $\lambda_t = \lambda_0 + \theta t$ for $t < x$, and $\lambda_t = \lambda_0 + \theta(t - x)$ for $t > x$, where x is the month of card expiration). For λ, we compute the average "experienced" λ for each service. That is, we weight the period-specific λ_t by the share of consumers still subscribed in period t, normalized so that the weights add up to one. We do so using the observed, regression-adjusted retention rates of each cohort, omitting the period of card expiration. We average across the three cohorts of card expiration, weighing by the total number of initial subscribers. In the third and fourth rows, we estimate our baseline model with time-invariant λ, but assume that $\lambda = 0.75$ and $\lambda = 0.5$ in the month of card expiration, respectively (instead of $\lambda = 1$). We report the mean and standard deviation of the estimates of the ("experienced") λ and revenue ratio, the second and ninth value, as well as the correlation with the baseline estimates. Panel B shows model results where we split the sample of cards by whether they ever had a cash advance.

Appendix Table A4: Parameter estimates, switching-cost model

Service	ρ	κ	η
A	0.981	1.092	0.758
	(0.007)	(0.143)	(0.345)
B	0.898	0.859	0.023
	(0.013)	(0.106)	(0.176)
C	0.950	0.565	1.486
	(0.005)	(0.038)	(0.118)
D	1.002	1.221	1.255
	(0.013)	(0.272)	(0.773)
E	1.002	1.045	1.942
	(0.010)	(0.170)	(0.723)
F	0.980	1.195	1.888
	(0.007)	(0.155)	(0.515)
G	0.948	0.292	3.546
	(0.002)	(0.019)	(0.076)
H	0.996	1.196	2.825
	(0.006)	(0.098)	(0.529)
I	1.000	1.765	0.474
	(0.009)	(0.141)	(0.276)
J	0.853	0.662	0.058
	(0.005)	(0.009)	(0.002)

Note: Table reports parameter estimates and bootstrap standard errors (in parentheses) for a model in which subscribers are fully attentive but face a switching cost of κ to cancel their subscription in non-replacement months, as well as to renew it in the month of card replacement. We estimate the model separately for each of the 10 subscription services. To compute the standard errors, we estimate our model on 1,000 bootstrap samples. For each parameter estimate, the standard error is the standard deviation of the 1,000 bootstrap estimates.

Appendix Table A5: Revenue impact, switching-cost model

Service	Share unaffected	Avg months subscribed		Revenue ratio
		If switching cost	If no switching cost	
A	0.05	34.4	21.0	1.44
B	0.00	14.3	7.6	1.78
C	0.00	21.9	16.6	1.27
D	0.24	39.9	22.0	1.28
E	0.28	38.8	23.6	1.20
F	0.06	41.8	27.7	1.33
G	0.00	24.1	21.5	1.10
H	0.26	46.0	30.2	1.20
I	0.21	46.5	21.9	1.42
J	0.00	10.1	5.9	1.65
Mean	**0.11**	**31.8**	**19.8**	**1.37**

Note: Table reports analogous counterfactual estimates as in Table 2 for a model in which subscribers are fully attentive but face a switching cost of κ to cancel their subscription in non-subscription months, as well as to renew it in the month of card replacement. For each subscription service, we simulate the monthly subscription choice of 100,000 hypothetical subscribers for 120 months after initial sign-up. Column (1) reports the share of subscribers not affected by switching costs because they have a positive valuation in every month. Columns (2) and (3) show, for subscribers with a negative valuation in at least one month, the average number of months they are subscribed under switching cost κ and with zero switching cost, respectively. Column (4) shows the gain in relative revenue from switching costs 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 face no switching costs ($\kappa = 0$). The numerator is the discounted sum of monthly subscribers if consumers face a switching cost κ. We discount future revenues at a rate of 1%.