

Selling Subscriptions

Liran Einav, Benjamin Klopach, Neale Mahoney
Supplemental Appendix

APPENDIX A: COMPARATIVE STATICS AND IDENTIFICATION

We use the exercise shown in Appendix Figure B3 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 B3 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 $R(x; \varphi)$.

We follow an analogous approach to simulate behavior under the switching cost model, making small modifications as needed based on differences in the models.




Figure B1. 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).






Figure B2. Account activity around replacement date for groceries and 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).

Holding $\eta = 0.864$ fixed:



Holding $\lambda = 0.111$ fixed:



Holding $\rho = 0.987$ fixed:




Figure B3. Comparative statics for inattention model

Note: Figure shows the comparative statics of our model of subscription renewal behavior under inattention 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 $\rho = 0.987$, $\lambda = 0.111$, and $\eta = 0.864$ (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 $\eta = 0.864$ fixed and show, separately for each group of cohorts, how either ρ (also holding $\lambda = 0.111$ fixed) or λ (also holding $\rho = 0.987$ 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 $\lambda = 0.111$ and $\eta = 0.864$ fixed, or by decreasing λ to $\lambda = 0.02$, while holding $\rho = 0.987$ and $\eta = 0.864$ fixed. In the middle row, we hold $\lambda = 0.111$ fixed and show the trade-off between ρ and η . In the bottom row, we hold $\rho = 0.987$ fixed and show the trade-off between λ and η .




Figure B4. Model fit, inattention model

Note: Figure shows the observed (adjusted) retention rate $\hat{R}_n(x)$ (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).




Figure B5. Model fit, switching-cost model

Note: Figure shows the observed (adjusted) retention rate $\hat{R}_n(x)$ (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).




Figure B6. 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: -0.920 (λ/κ), 0.969 (ρ), 0.992 (η), and 0.944 (the revenue ratio).




Figure B7. Iso-ratio curves - inattention

Note: Figure shows the revenue ratio as a function of the three inattention 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%.




Figure B8. 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.




Figure B9. 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.




Figure B10. 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.

Table B1— Average change in retention rates

Service	Average monthly change in retention rate	
	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, $\hat{R}_n(x) - \hat{R}_{n-1}(x)$. 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.

Table B2— Revenue impact of possible regulatory remedies

(a) Inattention

Service	Active choice every X months						
	∞	24	18	12	6	3	1
A	2.10	1.98	1.92	1.80	1.55	1.30	1.00
B	3.30	3.17	3.07	2.86	2.34	1.75	1.00
C	1.51	1.49	1.47	1.44	1.34	1.20	1.00
D	1.59	1.53	1.49	1.43	1.30	1.17	1.00
E	1.40	1.36	1.34	1.30	1.21	1.12	1.00
F	1.87	1.73	1.67	1.57	1.37	1.20	1.00
G	1.14	1.14	1.14	1.13	1.11	1.08	1.00
H	1.41	1.35	1.33	1.28	1.19	1.10	1.00
I	2.16	1.99	1.92	1.79	1.55	1.31	1.00
J	2.25	2.22	2.18	2.11	1.89	1.54	1.00
Mean	1.87	1.80	1.75	1.67	1.49	1.28	1.00

(b) Switching costs

Service	Active choice every X months						
	∞	24	18	12	6	3	1
A	2.49	2.35	2.28	2.11	1.72	1.33	1.00
B	3.59	3.44	3.34	3.12	2.49	1.66	1.00
C	1.69	1.66	1.64	1.59	1.45	1.25	1.00
D	1.92	1.83	1.78	1.67	1.43	1.20	1.00
E	1.56	1.51	1.49	1.43	1.29	1.14	1.00
F	2.04	1.90	1.83	1.70	1.44	1.21	1.00
G	1.17	1.16	1.16	1.15	1.13	1.09	1.00
H	1.54	1.48	1.44	1.38	1.25	1.12	1.00
I	3.13	2.80	2.64	2.34	1.80	1.35	1.00
J	2.90	2.84	2.79	2.66	2.26	1.61	1.00
Mean	2.20	2.10	2.04	1.92	1.63	1.30	1.00

Note: Table summarizes the revenue impact of possible regulatory remedies as shown in Figure 3 under the inattention (panel (a)) and switching cost models (panel (b)) described in Section III. 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 inattentive (panel (a)) or default to cancellation (panel (b)) 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%.

Table B3— Robustness and heterogeneity

	Lambda/kappa					Revenue ratio				
	Mean	St. Dev.	2nd	9th	Corr. w/ baseline	Mean	St. Dev.	2nd	9th	Corr. w/ baseline
Baseline - inattention	0.18	0.13	0.09	0.29		1.87	0.59	1.40	2.25	
Baseline - switching costs	0.97	0.40	0.57	1.22		2.20	0.75	1.54	3.13	
A. Robustness - inattention										
Linear decay of lambda	0.16	0.09	0.09	0.28	0.93	2.89	3.45	1.27	2.74	0.88
Linear decay of lambda with reset	0.18	0.14	0.10	0.28	1.00	2.66	1.98	1.39	2.86	0.91
Lambda at card expiration = 0.75	0.17	0.13	0.07	0.29	1.00	2.08	0.68	1.50	2.50	0.99
Lambda at card expiration = 0.5	0.17	0.14	0.04	0.28	0.99	3.08	2.01	1.54	5.75	0.33
B. Robustness - switching costs										
Share default cancel at card expiration = 0.75	1.02	0.46	0.57	1.26	1.00	2.47	1.04	1.65	3.58	0.91
Share default cancel at card expiration = 0.5	1.13	0.54	0.58	1.68	0.97	3.22	1.82	1.71	4.51	0.71
C. Heterogeneity - inattention										
Never used cash advance	0.17	0.11	0.09	0.28		1.87	0.59	1.37	2.27	
Used cash advance	0.14	0.11	0.06	0.22		3.27	2.21	1.44	5.74	
D. Heterogeneity - switching costs										
Never used cash advance	0.98	0.40	0.56	1.21		2.20	0.75	1.55	3.12	
Used cash advance	1.16	0.44	0.79	1.65		3.09	1.30	1.62	4.53	

Note: Table summarizes the results of our robustness and heterogeneity analyses for the estimates of λ (in the attention model), κ (in the switching costs model), and the revenue ratio. The first two rows summarize our baseline estimates for each model. 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 for the inattention model, where we estimated 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 robustness results for the switching cost model where we assume that only a subset of consumers are into cancellation in the month of card replacement (in the baseline, we assume all consumers are defaulted into cancellation in the replacement month). Panel B row 1 assumes that 75% of consumers default to cancellation, while row 2 assumes that 50% of consumers default to cancellation. Panel C shows results from the inattention model where we split the sample of cards by whether they ever had a cash advance. Panel D shows analogous estimates from the switching cost model by cash advance usage.