Online appendix

A. a conceptual model for user heterogeneity and ad load personalization

In this section, we illustrate the benefits of personalizing ad load with a set of examples. First, we consider a simple utility model that reflects the discrete choice between the outside options, consuming the ad-supported product, or using the paid subscription. Consider the following model:

$$u(z, p \mid \theta, \beta, \alpha) = \max \begin{cases} 0 & \text{outside option,} \\ \theta - \alpha z & \text{ad supported,} \\ \theta - \beta p & \text{paid version,} \end{cases} \tag{17}$$

where θ is the utility from consuming the product, z specifies the ad load[26], and p is the subscription price. In the ad-supported condition, users are effectively paying with their time by listening to ads. Therefore, parameters α and β reflect how time and money are valued by users, that is users with higher/lower values of α and β are more/less sensitive to ads and prices, respectively. Also, let γ and c be the revenue per ad and the marginal cost of offering the service, respectively. Finally, let $\theta_\alpha = \frac{\theta}{a}$ and $\theta_\beta = \frac{\theta}{b}$. If $\theta_\beta > z$ and $\frac{\theta_\beta}{\theta_\alpha} > \frac{z}{p}$, the user picks the ad-supported version. And if $\theta_\alpha > p$ and $\frac{\theta_\beta}{\theta_\alpha} < \frac{z}{p}$, the paid version is purchased; otherwise, the outside option is preferred. Figure 14 illustrates the decision regions for different types when price is set to p and ad load is equal to z.

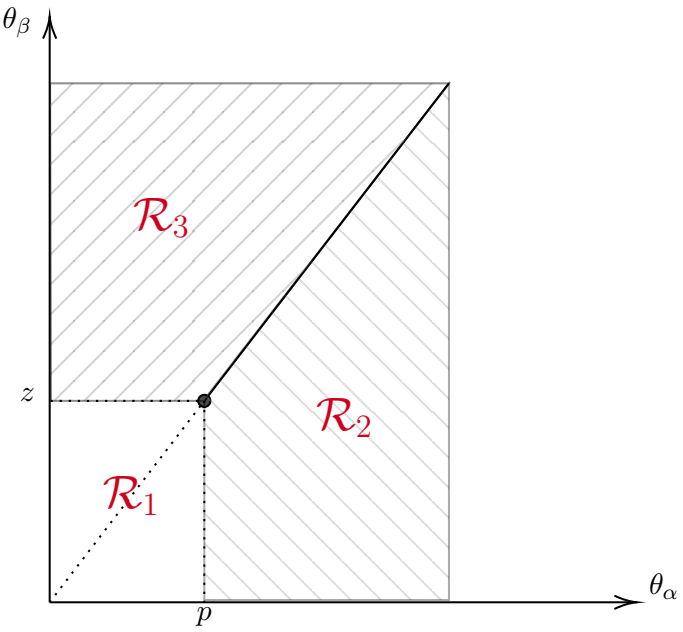

Figure 14: Decision regions for different types $(\theta_\alpha, \theta_\beta)$ for price vector (p, z), where p and γz are assumed to be larger than c. Types that lie in $\mathcal{R}_1$ pick the outside option, types in $\mathcal{R}_2$ subscribe for the paid service, and those in $\mathcal{R}_3$ choose the ad-supported version.

[26] In this simplified model, we assume users can consume the service in exchange for listening to z ads. In our case study, we account for the fact that the intensive margin of consumption (number of hours) could vary across users, and that possibility factors into the ad revenue.

A monopolist that can perfectly discriminate along both ad load and price dimensions will maximize its profits for each type (θ_a, θ_b). Note that a listener of type $(\theta_\alpha, \theta_\beta)$ has a willingness to consume at most $z = \theta_\alpha$ ads, and pay price $p = \theta_\beta$. For a listener with $\gamma\theta_\alpha > \max(\theta_\beta, c)$, the monopolist will only offer the ad-supported service with $(p, z) = (\infty, \theta_\alpha)$, and for those with $\theta_\beta > \max(\gamma\theta_\alpha, c)$, only the subscription service is offered with $(p, z) = (\theta_\beta, \infty)$, and when none of these conditions are satisfied serving the customer would not be worthwhile and $(p, z) = (\infty, \infty)$. Figure 15 demonstrates the decision regions for a monopolist based on the type of products sold, and the profits over regions $\mathcal{R}_2$ and $\mathcal{R}_3$ are equal to:

$$\Pi^* = \int_{\theta_\beta=c}^{\infty} \int_{\theta_\alpha=0}^{\frac{1}{\gamma}\theta_\beta} (\theta_\beta - c) f(\theta_\alpha, \theta_\beta) d_{\theta_\alpha} d_{\theta_\beta} + \int_{\theta_\alpha=\frac{c}{\gamma}}^{\infty} \int_{\theta_\beta=0}^{\gamma\theta_\alpha} (\gamma\theta_\alpha - c) f(\theta_\alpha, \theta_\beta) d_{\theta_\beta} d_{\theta_\alpha}, \tag{18}$$

where $f(\theta_\alpha, \theta_\beta)$ is the joint density of $(\theta_\alpha, \theta_\beta)$.

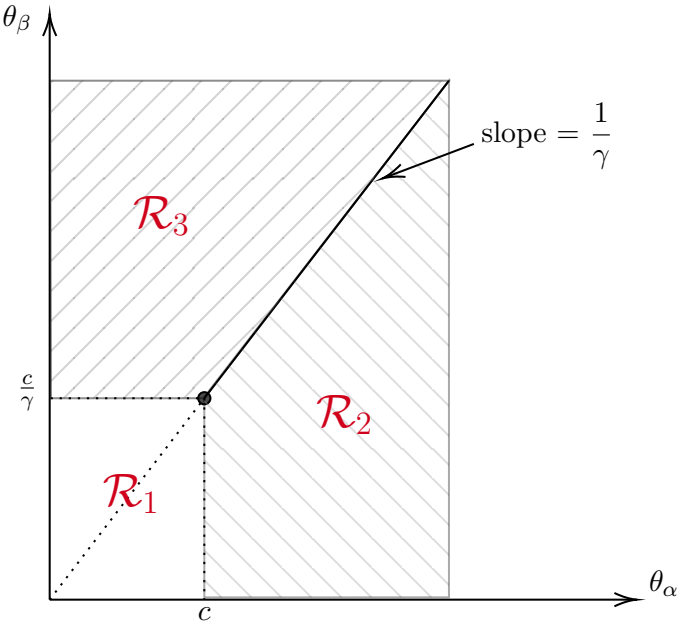

Figure 15: Decision regions for different types $(\theta_\alpha, \theta_\beta)$ for a price-ad load-discriminating monopolist. Serving customers in region $\mathcal{R}_1$ is not worthwhile; those in $\mathcal{R}_2$ will purchase the paid subscription and the rest will use the ad-supported service.

The results above demonstrate that when the monopolist has full information, he will only make one of the products available to each customer. However, these results rely on the crucial assumption that the monopolist can accurately observe the type of each listener $(\theta_\alpha, \theta_\beta)$. Our goal is to illustrate the benefits of offering a personalized menu of products when the seller has partial information about the type of consumers or randomness exists in choices made by the customers. In this case, personalization can be achieved by customizing the ad load. The seller uses their information about each customer's type to offer a personalized ad load. This ad load then serves as a screening mechanism, influencing the substitution patterns among ad-supported, paid, and outside options.

Let us consider a few simple examples to illustrate this idea:

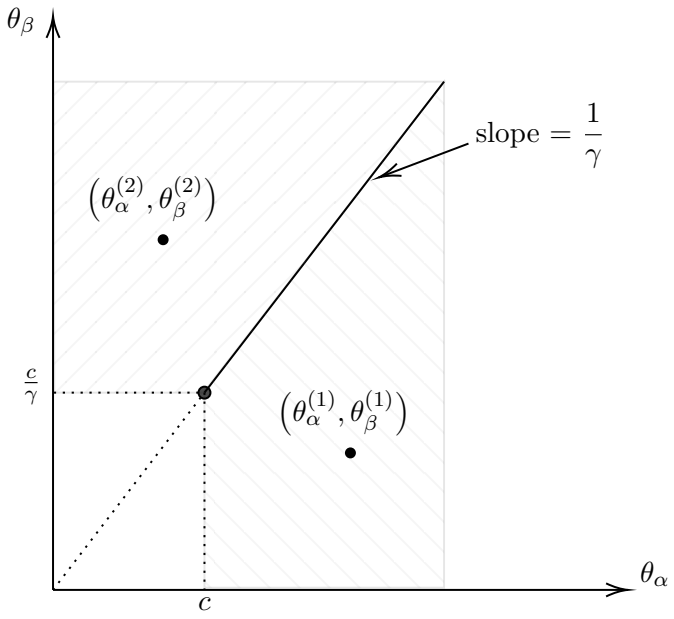

Figure 16: Perfectly separable condition. The seller is uncertain if the consumer is of type 1 or 2, but offering a menu with $(p, z) = (\theta_\alpha^{(1)}, \theta_\beta^{(2)})$ yields profits that are equal to the case where the seller has full information.

- **Separable types**: Consider a monopolist (he) who needs to provide service to a customer (she). The monopolist knows that with probability ρ_1, the customer is of type $(\theta_\alpha^{(1)}, \theta_\beta^{(1)})$, and with probability $\rho_2 = 1 - \rho_1$, she is of type $(\theta_\alpha^{(2)}, \theta_\beta^{(2)})$. If $\theta_\alpha^{(1)} > \theta_\alpha^{(2)}$ and $\theta_\beta^{(2)} > \theta_\beta^{(1)}$ (see Figure 16 for visual illustration), it is optimal for the monopolist to offer the paid version when the realized type is 1 and offer the ad-supported product when the realized type is 2. In this case, by offering a menu $(p, z) = (\theta_\alpha^{(1)}, \theta_\beta^{(2)})$, the firm can extract monopoly profits regardless of the type of customer. In particular, if the customer is of type 1, she purchases the paid version, whereas if she is of type 2, she uses the ad-supported product. In other words, in this case, using a menu can fully separate types from each other and resolves the uncertainty.

- **Inseparable types**: Let us now consider a more nuanced case. The customer can be of one of two types with probabilities ρ_1 and $\rho_2 = 1 - \rho_1$. This time $(\theta_\alpha^{(2)}, \theta_\beta^{(2)}) > (\theta_\alpha^{(1)}, \theta_\beta^{(1)})$; see Figure 17 for visual illustration. In this case, the optimal menu corresponds to one of the five red dots in Figure 17. Depending on the values of γ, c, ρ_1, ρ_2, $(\theta_\alpha^{(1)}, \theta_\beta^{(1)})$, and $(\theta_\alpha^{(2)}, \theta_\beta^{(2)})$, either of these can be the optimal menu to offer. The only case where both products are offered is when $(p, z) = \left(\theta_\alpha^{(1)}, \theta_\alpha^{(1)} \frac{\theta_\beta^{(2)}}{\theta_\alpha^{(2)}} - \epsilon\right)$. Note that in this case, the existence of type 1 imposes a positive externality on the ad load, that is the quality of service when type 2 is realized. In other words, if the seller were certain the customer is of type 2, he would have the incentive to increase the ad load. However, in this case, because the customer is served under both conditions the ad load cannot be increased to more than $\theta_\alpha^{(1)} \frac{\theta_\beta^{(2)}}{\theta_\alpha^{(2)}}$ to make it incentive compatible for the type 2 customer to use the ad-supported service[27].

[27] Recall from Figure 15 that if $\frac{\theta_\beta}{\theta_\alpha} > \frac{1}{\gamma}$, the seller is better off providing the ad-supported service.

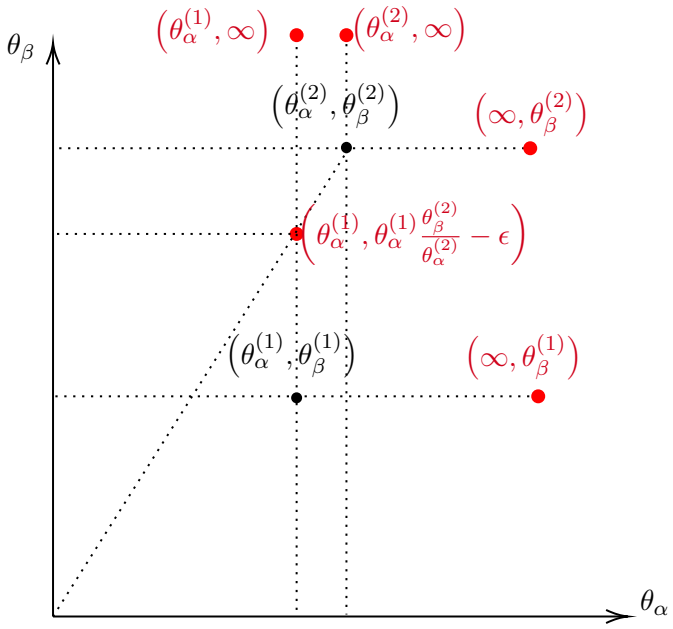

Figure 17: Inseparable condition. The uncertainty in the types cannot be perfectly dealt with by using a personalized menu. The optimal menu (p, z) is one of the five red dots, and depending upon the types and realization probabilities either one can be optimal.

The examples above illustrate that personalizing implicit prices or non-price attributes in a product line, where users can pay with both price and a non-pecuniary method such as ad load, can make use of both the partial information that the firm may have and the customer's private information. Note that these applications are not limited to the case where uncertainty is present in parameter estimates but also apply in random utility models even when parameter uncertainty is neglected. In section 3, we used a random utility model that corresponds to (17) to illustrate the trade-offs involved in personalizing ad load.

B. Ad delivery mechanism and partial control over realized ad load/capacity

In section 4, we discussed the field experiment and presented results on randomization checks, and realized changes in ad load across experiment cells. In this section, we discuss the realized effect on ad load and capacity in more detail.

Our experiment consists of seven cells: 3x1, 4x2, 6x2, 4x3, 5x3, 6x3, and control. The experiment shifts the timers in the ad delivery system. For instance, in the 3x1 condition, the timer is set to 20 minutes, whereas it is set to 15 minutes in the control, 4x2, and 4x3 conditions. The only difference between the 4x2 and the control condition is that the first ad pod[28] within each listening session in the control condition is constrained to be of length one. Note that a listener in the 6x3 condition ends up becoming eligible for an ad pod fewer than six times per hour (every 15 minutes), because the song endings do not perfectly align with the timers, see Figure 3. For example, the ad capacity, that is, the number of opportunities to show an ad per hour, in the 6x3 condition ends up being far less than 6×3=18. Figure 18 presents the realized ad capacity across different experiment conditions. Note that the ad capacity for the FxL condition is often much smaller than F·L, since the song endings do not align with the timers. There is only one exception and that is the 3x1 condition. As highlighted in Figure 3, the listeners across all conditions become eligible for the first ad pod within the first five minutes of each listening session, therefore, the realized ad capacity may end up being larger than three in the 3x1 condition, especially for listeners that have short listening sessions.

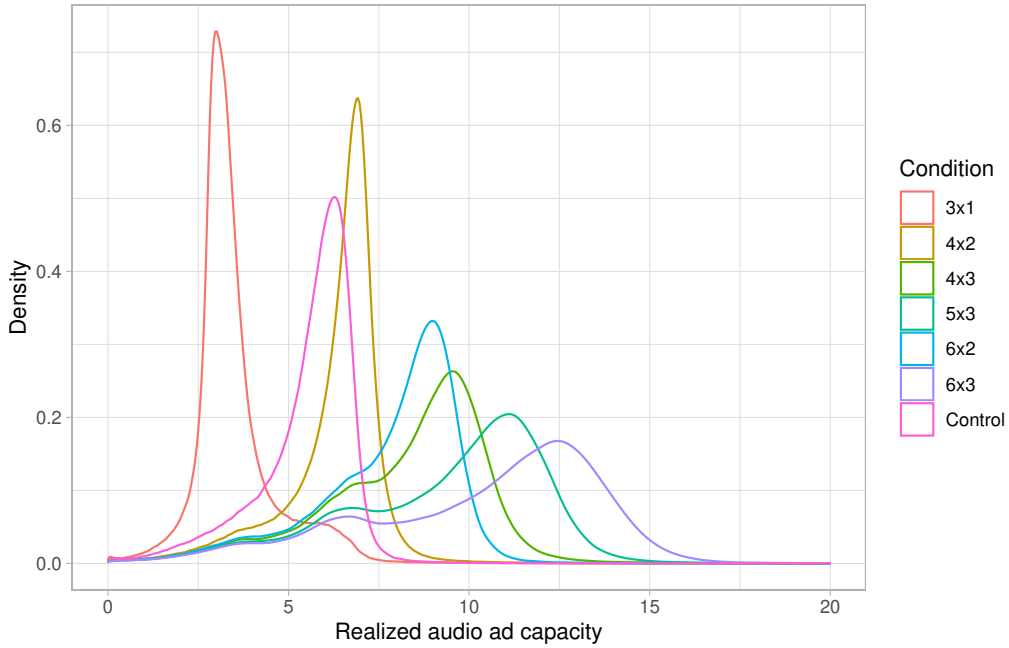

Figure 18: Realized audio ad capacity across different treatment arms

When a listener (she) becomes eligible to receive an ad, the ad delivery system makes a request to fetch an ad for her. If she belongs to a demographic group that is attractive to advertisers it is easier

[28]The listeners across all experiment conditions become eligible for the first ad pod after the first five minutes of each listening session.

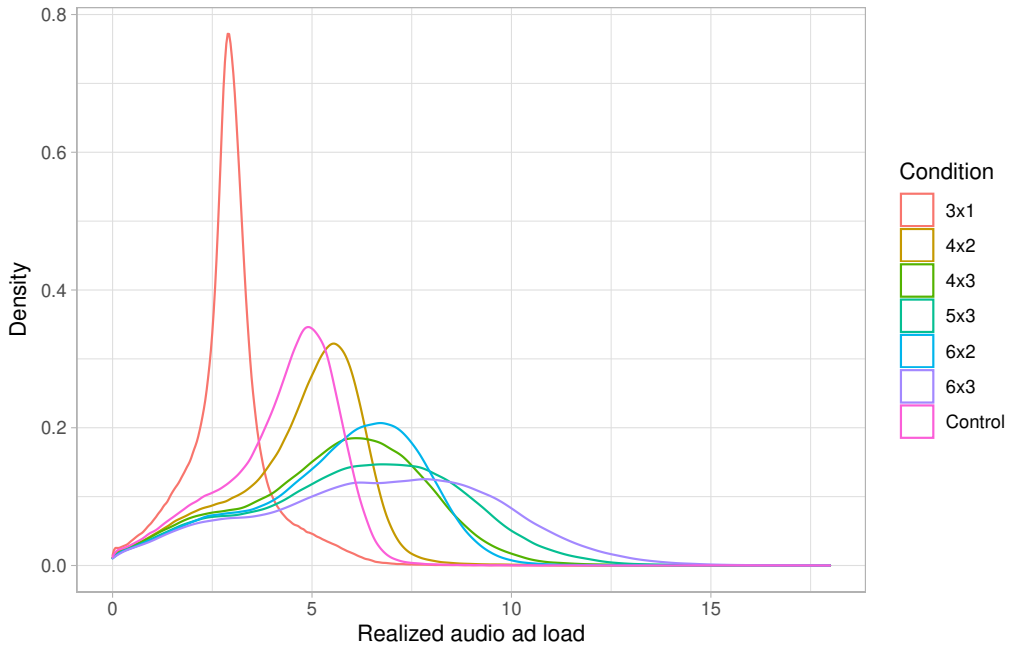

Figure 19: Realized ad load across different treatment arms

to fetch ads for her. However, if the system runs out of ads to show to her those opportunities (ad capacity) are not filled. For instance, a listener may be eligible to receive eight ads in a given hour, but the system may end up fetching only five ads in the ad inventory that could be played for her. In our data, we record both the number of opportunities (ad capacity) and the realized number of ads (ad load) delivered in an hour for every user.

To summarize, pod length and frequency determine *ad capacity* rather than ad load. Figure 3 and Figure 18 illustrate that the realized ad capacity for the FxL condition ends up being far less than F·L as the song endings do not perfectly align with the timers. Furthermore, the realized number of ads shown to each listener (ad load) also depends on advertisers' demand. As we move to more extreme conditions such as the 6x3 condition it becomes more difficult to shift both ad capacity and ad load, see Table 3. Figure 19 depicts the density of realized ad load for users in different treatment cells. Although higher-ad-capacity conditions have a higher realized ad load, the distribution becomes more dispersed as the capacity increases, compare Figures 18 and 19. This finding is indicative of the fact that filling higher capacities for users tends to be more difficult because running out of ads to serve in the higher-capacity conditions is more probable. Table 3 in the body illustrates these findings by reporting the average realized ad load, capacity, and fill rate during the first year of the experiment.

To further demonstrate the partial control problem, we plot the lift in ad load between the control and 6x3 condition across different consumer groups based on their pre-treatment ad load in Figure 20. The figure shows the increase in ad load in the 6x3 condition relative to the control condition is not uniform across different listener groups. The lift in ad load is more pronounced for consumers who received more ads in the pre-treatment period. This heterogeneity in the lift reflects the role of advertisers' demand in the realized ad load and shows that the additional capacity is more likely to

be filled for those consumers who are more attractive to advertisers. This demonstrates the fact that firms need to account for the discrepancy between the *intended* and *realized* change in the implicit price, which leads to an additional layer of complexity relative to the traditional pricing problems.

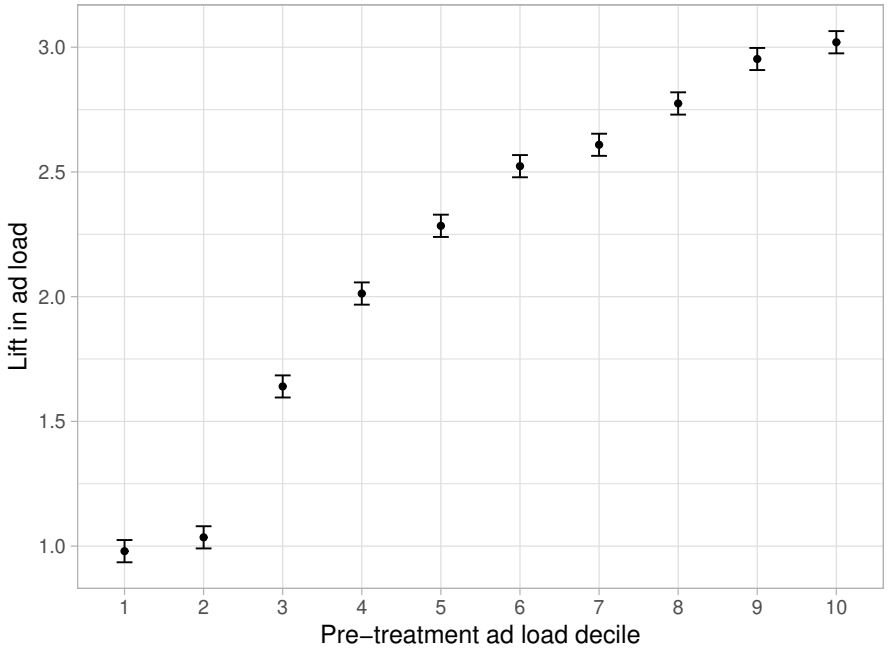

Figure 20: Realized lift in ad load (ads/hours) in the 6x3 condition relative to control as a function of pre-treatment ad load. Note the lift in ad load could vary drastically across different groups, due to differences in the attractiveness of different segments for advertisers.

C. Reduced-form regressions

In Figure 5 of section 5, we illustrated the average treatment effect of changing ad load on consumption and subscription status using a series of instrumental variable regressions. In this section, we present the reduced-form regressions that correspond to those IV regressions. In particular, we regress the normalized outcomes (3) directly on the treatment dummies using the following specification:

$$\tilde{\mathbf{Y}}_i = \alpha + \sum_{j=1}^{6} \beta_j \cdot \mathbb{1}_{\{\boldsymbol{\tau}=e_j\}} + \epsilon_i, \tag{19}$$

where i indexes listeners, and $\tilde{\mathbf{Y}}_i$ is a normalized outcome of interest, i.e., listening hours, activity dummy, and plus subscription dummy, for each user in a given week. We present the results of these analyses for the last week of each quarter after the experiment kicked off in Tables 5-10. Similar to our pooled IV results in Figure 5, our results here show that the treatment effect of ad load changes remains stable after Q4 2016, however, the impact on consumption takes longer than a year to stabilize.

Table 5: Activity and subscription status in the last week of Q3-2016 across different treatment arms relative to control.

	Dependent variable:			
	All hours	Ad-supported hours	Active user	Subscription rate
	(1)	(2)	(3)	(4)
Control	100.000***	100.000***	100.000***	100.000***
	(0.218)	(0.234)	(0.110)	(0.809)
3x1	0.288	0.814**	0.364**	−4.188***
	(0.308)	(0.330)	(0.156)	(1.143)
4x2	−0.375	−0.522*	−0.106	3.913***
	(0.267)	(0.286)	(0.135)	(0.990)
6x2	−1.144***	−1.546***	−0.413**	6.959***
	(0.377)	(0.405)	(0.191)	(1.400)
4x3	−0.878**	−1.185***	−0.430**	8.956***
	(0.377)	(0.405)	(0.191)	(1.401)
5x3	−0.862**	−1.794***	−0.471**	12.886***
	(0.377)	(0.404)	(0.190)	(1.399)
6x3	−1.415***	−2.169***	−0.536**	14.129***
	(0.453)	(0.486)	(0.229)	(1.682)
Observations	7,350,278	7,350,278	7,350,278	7,350,278
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 6: Activity and subscription status in the last week of Q4-2016 across different treatment arms relative to control.

	Dependent variable:			
	All hours	Ad-supported hours	Active user	Subscription rate
	(1)	(2)	(3)	(4)
Control	100.000***	100.000***	100.000***	100.000***
	(0.219)	(0.238)	(0.110)	(0.706)
3x1	1.003***	1.544***	0.462***	−5.418***
	(0.309)	(0.336)	(0.155)	(0.998)
4x2	−0.615**	−0.785***	−0.210	2.293***
	(0.268)	(0.291)	(0.134)	(0.865)
6x2	−1.499***	−2.196***	−0.447**	8.126***
	(0.379)	(0.412)	(0.190)	(1.223)
4x3	−1.410***	−2.163***	−0.564***	11.858***
	(0.379)	(0.412)	(0.190)	(1.223)
5x3	−1.834***	−3.156***	−0.843***	14.945***
	(0.379)	(0.412)	(0.190)	(1.222)
6x3	−2.478***	−3.888***	−1.051***	17.941***
	(0.455)	(0.495)	(0.228)	(1.469)
Observations	7,350,278	7,350,278	7,350,278	7,350,278
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 7: Activity and subscription status in the last week of Q1-2017 across different treatment arms relative to control.

	Dependent variable:			
	All hours	Ad-supported hours	Active user	Subscription rate
	(1)	(2)	(3)	(4)
Control	100.000***	100.000***	100.000***	100.000***
	(0.215)	(0.233)	(0.109)	(0.591)
3x1	1.998***	2.946***	0.644***	−5.581***
	(0.303)	(0.329)	(0.154)	(0.835)
4x2	−0.468*	−0.691**	−0.460***	2.738***
	(0.263)	(0.285)	(0.134)	(0.723)
6x2	−1.794***	−2.512***	−0.993***	8.267***
	(0.372)	(0.403)	(0.189)	(1.023)
4x3	−1.861***	−2.955***	−1.233***	10.620***
	(0.372)	(0.403)	(0.189)	(1.023)
5x3	−2.617***	−4.516***	−1.419***	14.838***
	(0.372)	(0.403)	(0.189)	(1.022)
6x3	−3.156***	−5.027***	−1.867***	17.215***
	(0.447)	(0.484)	(0.227)	(1.229)
Observations	7,350,278	7,350,278	7,350,278	7,350,278
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 8: Activity and subscription status in the last week of Q2-2017 across different treatment arms relative to control.

	Dependent variable:			
	All hours	Ad-supported hours	Active user	Subscription rate
	(1)	(2)	(3)	(4)
Control	100.000***	100.000***	100.000***	100.000***
	(0.219)	(0.243)	(0.112)	(0.548)
3x1	2.134***	3.401***	0.678***	−5.615***
	(0.310)	(0.344)	(0.158)	(0.774)
4x2	−0.215	−0.334	−0.345**	2.795***
	(0.268)	(0.298)	(0.137)	(0.671)
6x2	−1.928***	−3.040***	−1.077***	9.398***
	(0.379)	(0.421)	(0.194)	(0.949)
4x3	−2.031***	−3.096***	−1.475***	10.937***
	(0.380)	(0.421)	(0.194)	(0.949)
5x3	−2.939***	−4.673***	−1.764***	15.182***
	(0.379)	(0.421)	(0.194)	(0.948)
6x3	−3.759***	−6.054***	−2.267***	16.315***
	(0.456)	(0.506)	(0.233)	(1.140)
Observations	7,350,278	7,350,278	7,350,278	7,350,278
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 9: Activity and subscription status in the last week of Q3-2017 across different treatment arms relative to control.

	Dependent variable:			
	All hours	Ad-supported hours	Active user	Subscription rate
	(1)	(2)	(3)	(4)
Control	100.000***	100.000***	100.000***	100.000***
	(0.225)	(0.249)	(0.114)	(0.566)
3x1	2.179***	3.392***	0.554***	−5.457***
	(0.318)	(0.352)	(0.161)	(0.800)
4x2	−0.532*	−0.835***	−0.759***	3.955***
	(0.276)	(0.305)	(0.140)	(0.693)
6x2	−1.510***	−2.599***	−1.298***	10.684***
	(0.390)	(0.431)	(0.197)	(0.981)
4x3	−1.758***	−3.082***	−1.528***	12.532***
	(0.390)	(0.431)	(0.198)	(0.981)
5x3	−2.814***	−5.043***	−2.185***	16.985***
	(0.390)	(0.431)	(0.197)	(0.980)
6x3	−3.527***	−6.174***	−2.349***	18.201***
	(0.469)	(0.518)	(0.237)	(1.178)
Observations	7,350,278	7,350,278	7,350,278	7,350,278
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 10: Activity and subscription status in the last week of Q4-2017 across different treatment arms relative to control.

	Dependent variable:			
	All hours	Ad-supported hours	Active user	Subscription rate
	(1)	(2)	(3)	(4)
Control	100.000***	100.000***	100.000***	100.000***
	(0.234)	(0.262)	(0.119)	(0.575)
3x1	2.114***	3.173***	0.897***	−5.802***
	(0.331)	(0.370)	(0.168)	(0.813)
4x2	−0.569**	−1.238***	−0.635***	4.100***
	(0.287)	(0.321)	(0.145)	(0.705)
6x2	−1.411***	−2.491***	−1.072***	10.695***
	(0.405)	(0.454)	(0.206)	(0.997)
4x3	−1.761***	−3.447***	−1.738***	13.386***
	(0.406)	(0.454)	(0.206)	(0.997)
5x3	−2.330***	−5.082***	−2.166***	17.077***
	(0.405)	(0.453)	(0.205)	(0.996)
6x3	−3.604***	−6.424***	−2.369***	18.932***
	(0.487)	(0.545)	(0.247)	(1.197)
Observations	7,350,278	7,350,278	7,350,278	7,350,278
Note:				*p<0.1; **p<0.05; ***p<0.01

D. Persistence in heterogeneous treatment effects

In section 7 we showed that our models are able to sort users based on the magnitude of the treatment effect on ad and subscription revenues using data from December 2016. In this section, we illustrate that the heterogeneous treatment effects detected in this time period are persistent. In particular, we show that if one sorts users in the hold-out sample based on the predicted treatment effect on subscription and ad revenues, the ordering does not only explain the lift in December 2016, see Figure 8, but it also has explanation power in December 2017.

In section 7 we divided the listeners in the hold-out sample into five quintiles based on the predicted treatment effects of the 6x3 relative to the control condition for both subscriptions and ads. We then demonstrated that the *realized lift* in subscriptions and ads were indeed correlated with predicted treatment effects in Figure 21. We replicated this analysis but instead of using December 2016 data for calculating we used data from a year later, that is during December 2017. We present these results in Figure 21. Note that the magnitude of the effects could be slightly different from those in Figure 8, but the ordering remains consistent. This means that the ordering of users based on treatment effects predictions using 2016 data is still valid a year later in 2017. This is very interesting and shows that it is possible to use the short-run 6 months data to detect heterogeneous treatment effects that are persistent.

These results are important and interesting for two reasons. First, they demonstrate that the heterogeneous treatment effects identified using short-run data in this case have explanatory power in the long run. Second, given our reliance on only a snapshot of data from December 2016 for training and targeting, these results suggest that our approach is likely to yield the anticipated outcomes over a longer time horizon as well. This is confirmed by our inverse propensity weighted estimates of the performance of the personalized policy. As depicted in Figure 10, using this snapshot for personalization achieves the desired results over a longer-term horizon (18 months). In particular, the delivered ad load remains stable and similar to that of the control condition, and the increase in plus users under the personalized policy stabilizes within 4-6 months after the experiment begins.

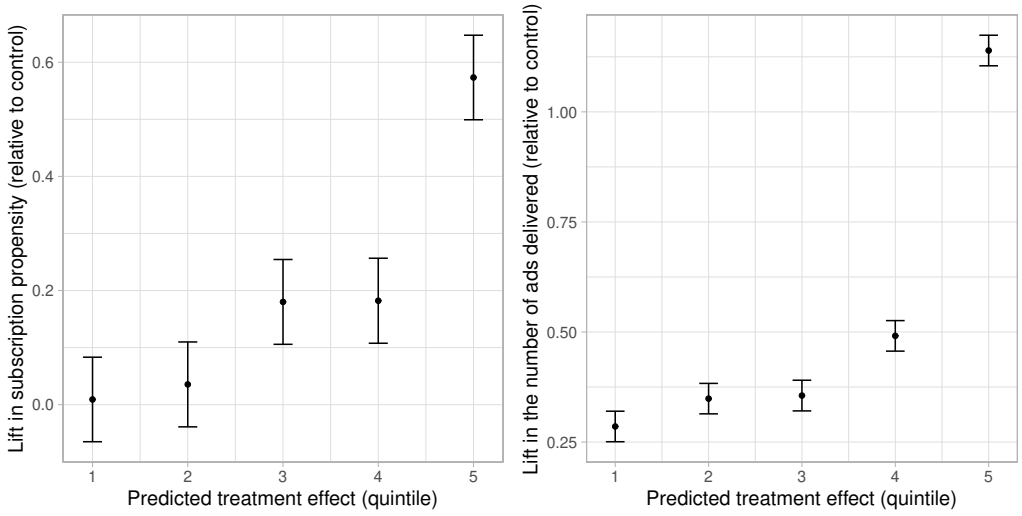

(a) Subscriptions

(b) Ads

Figure 21: Realized lift in subscription propensity and ads delivered in the 6x3 condition relative to control in December 2017 on the hold-out sample as a function of predicted treatment effect quintile using December 2016 as training.

E. A direct estimation model for ads delivered

In section 7, we combined three models to predict the number of ads played to each listener under different treatment conditions (counterfactual states). A model that predicts the probability of being an active ad-supported listener $P_a(\mathbf{x}, \boldsymbol{\tau})$, a model that predicts the number of listening hours conditional on being active and ad-supported $C(\mathbf{x}, \boldsymbol{\tau})$, and finally a model that predicts the ad load conditional on being active and ad-supported $A(\mathbf{x}, \boldsymbol{\tau})$ across different conditions. The predicted number of ads for an individual with pre-treatment features $\mathbf{x}$ and in treatment condition, $\boldsymbol{\tau}$ was calculated as:

$$\text{predicted \# of ads delivered} = P_a(\mathbf{x}, \boldsymbol{\tau}) \cdot C(\mathbf{x}, \boldsymbol{\tau}) \cdot A(\mathbf{x}, \boldsymbol{\tau}). \tag{20}$$

Since models $C(\mathbf{x}, \boldsymbol{\tau})$ and $A(\mathbf{x}, \boldsymbol{\tau})$ are trained on listeners conditional on being ad-supported and active, one might be concerned about selection. However, note that we only use these models to solve for the allocation policy (12), and we *do not* use the predictions of the models to evaluate the realized number of ads or subscriptions in the counterfactuals. In particular, the performance of the model in terms of lift in subscriptions and delivered ads is calculated using inverse propensity weighted estimates in the hold-out sample, which is similar to the approach used by Hitsch et al. (2023), Rafieian and Yoganarasimhan (2021), Simester et al. (2020), and Yoganarasimhan et al. (2022). This means that even if the errors are correlated across the first model that learns the probability of being ad-supported and the models used for predicting listening hours and ad load, this correlation does not affect the counterfactual estimates of model performance, e.g., Figures 9-10, because those results are based on inverse propensity weighted estimates and do not depend on the predicted values themselves. Nevertheless, here we show that we obtain similar performance if we use a model that directly predicts the number of ads delivered instead of the approach used in (20).

Let $\mathbf{Y}_i$ be the number of ads played to listener i during the training period (December 2016). Note $\mathbf{Y}_i = 0$ for paid listeners or inactive users in that time period. We model $\mathbf{Y}_i$ as follows:

$$\mathbf{Y}_i = \mathbb{1}_{f(\mathbf{x}_i, \boldsymbol{\tau}_i) > 0} \cdot f(\mathbf{x}_i, \boldsymbol{\tau}_i), \tag{21}$$

where $f(\mathbf{x}, \boldsymbol{\tau})$ is parameterized as a neural network. A schematic view of the neural network's structure is presented in Figure 22. We use a weighted objective similar to (6) with mean-squared error as the loss to train the model. Let $\mathcal{A}(\mathbf{x}, \boldsymbol{\tau})$ be the model that predicts the number of ads served to each individual. Using this model instead to predict the number of ads served to each individual transforms the optimization problem (13) into:

$$\sum_i \quad \underset{\boldsymbol{\tau_i}}{\text{maximize}} \quad \underbrace{m_s \cdot P_s(\boldsymbol{x}_i, \boldsymbol{\tau}_i) + \lambda \cdot \mathcal{A}(\mathbf{x}_i, \boldsymbol{\tau}_i)}_{f(\mathbf{x}_i, \lambda, \boldsymbol{\tau}_i)}. \tag{22}$$

Similar to our approach in section 8, we shift λ in equation (13) to obtain different personalized policies albeit using the joint model $\mathcal{A}(\mathbf{x}_i, \boldsymbol{\tau}_i)$ for predicting the number of ads served. The counterpart of the Pareto frontier in Figure 9 for the new set of personalized policies using the joint model is pre-

sented in Figure 23. The performance remains similar to our approach in section 8 and the personalized counterpart of control achieves about a 7% gain in subscription profits relative to control.

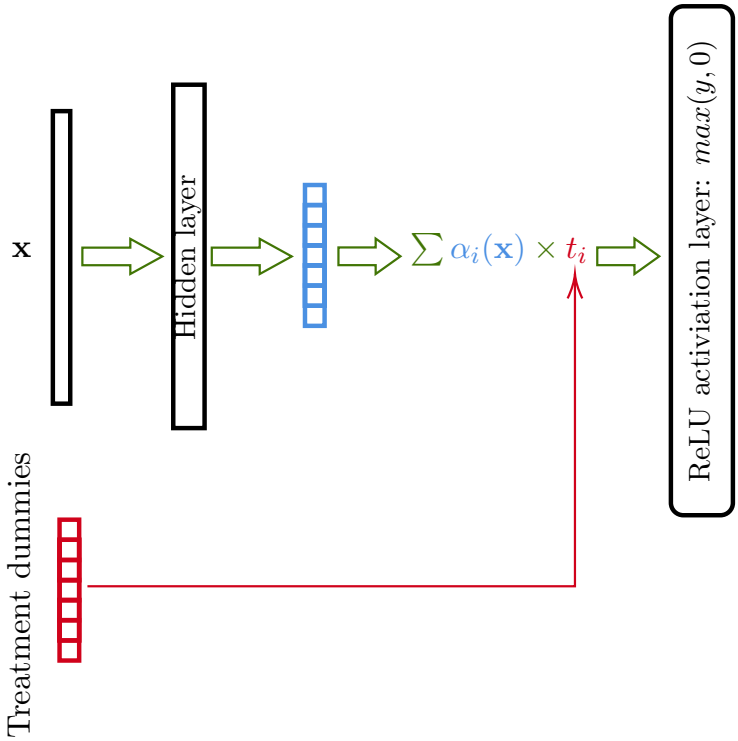

Figure 22: A schematic view of the neural network architecture used for estimating the realized number of ads delivered across experimental conditions.

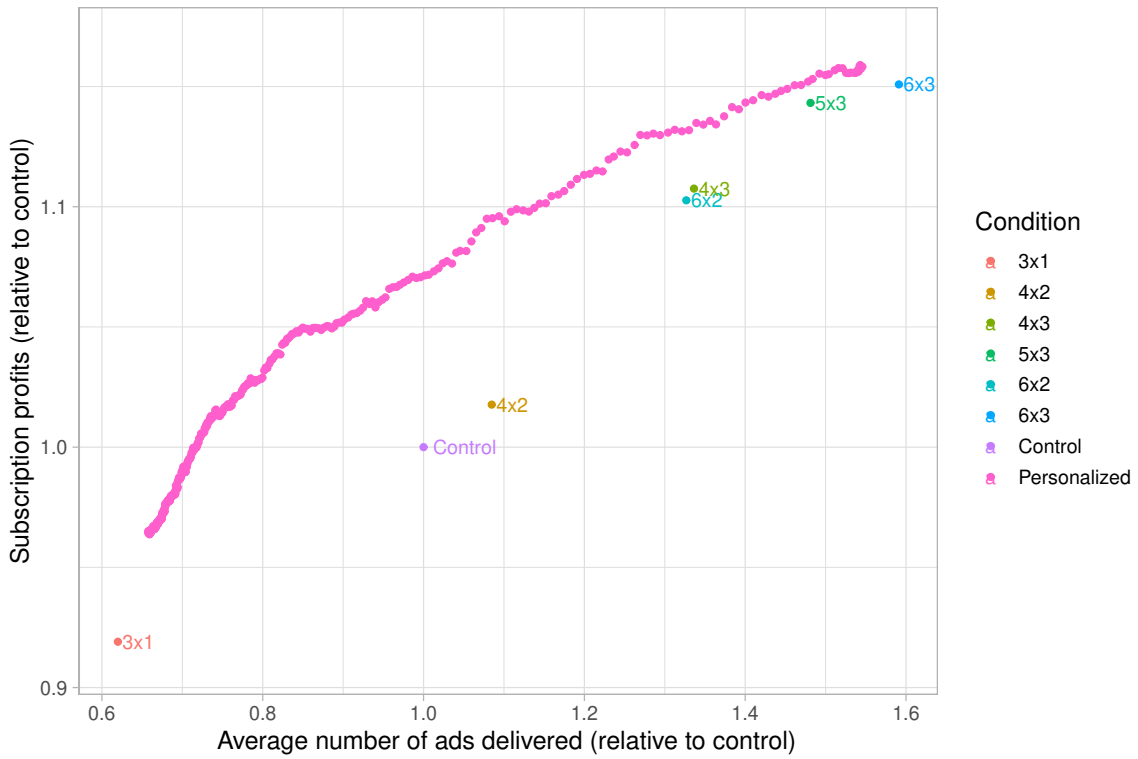

Figure 23: Change in subscription profits as a function of the number of ads served using a joint prediction model for the number of ads served. The performance of the personalized policies is similar to Figure 9.

F. Using long-term outcomes for training

In section 7, we used data from December 2016 to train our models and investigate their performance. Additionally, we employed separate models to learn about consumption hours (intensive margin), realized ad load, and user state (extensive margin). As demonstrated in Appendix E, a direct method of learning the number of ads served and user state can achieve a performance similar to our initial approach. However, one caveat is that these models rely solely on a snapshot of data from December 2016 and therefore may not account for the fact that different types of users may remain in the subscription state for varying durations.

As discussed in section 7, our selection of December 2016 data was based on two factors: both subscription and ad revenues appeared to have stabilized by this point, and from a practical standpoint, the firm might prefer to use short-term outcomes (6 months, rather than 18 months) to personalize results. Nonetheless, this raises the question of whether the firm could achieve better results by using all available information from the 18-month post-experiment period.

To address this question, we aggregated the outcomes for the 18-month period following the experiment, from June 2016 through December 2017. Specifically, we calculated the number of weeks each user was a plus subscriber and the total number of ads they listened to during this period. Let $\mathcal{S}_i$ denote the number of weeks a user was a subscriber, and let ζ_i represent the total number of ads they listened to. Using an architecture and objective function similar to that in Appendix E, we trained two models to predict the number of ads played and the number of weeks a user remained a subscriber in the 18-month period following the experiment, as a function of pre-treatment variables $\mathbf{x}_i$ and the treatment condition $\boldsymbol{\tau}_i$. Then, as per our notation in section 8, the firm's optimization problem translates to:

$$\begin{aligned} \underset{\boldsymbol{\tau_i}}{\text{maximize}} \quad & \sum_i m_s \mathcal{S}(\mathbf{x}_i, \boldsymbol{\tau}_i) \\ & \sum_i \zeta(\mathbf{x}_i, \boldsymbol{\tau}_i) = \Gamma, \end{aligned} \tag{23}$$

Where Γ is the total number of ads that must be served, and m_s is the profit margin from subscribers each week. We employ the same approach as in section 8 to solve the Lagrangian relaxation of the problem.

We first present the Pareto frontier for the outcomes in December 2016 in Figure 24, which is the counterpart to Figures 9 and 23. As expected, using 18-month outcomes does push out the Pareto frontier, achieving better performance. For instance, in both cases, using one month's outcomes in Figures 9 shows a 7% increase in subscription profits relative to the control condition, whereas using aggregated 18-month outcomes improves profits by about 9%. Note that aggregation reduces noise and incorporates information on the length of subscriptions, which are both useful for more effective sorting of individuals based on their response to changes in ad load.

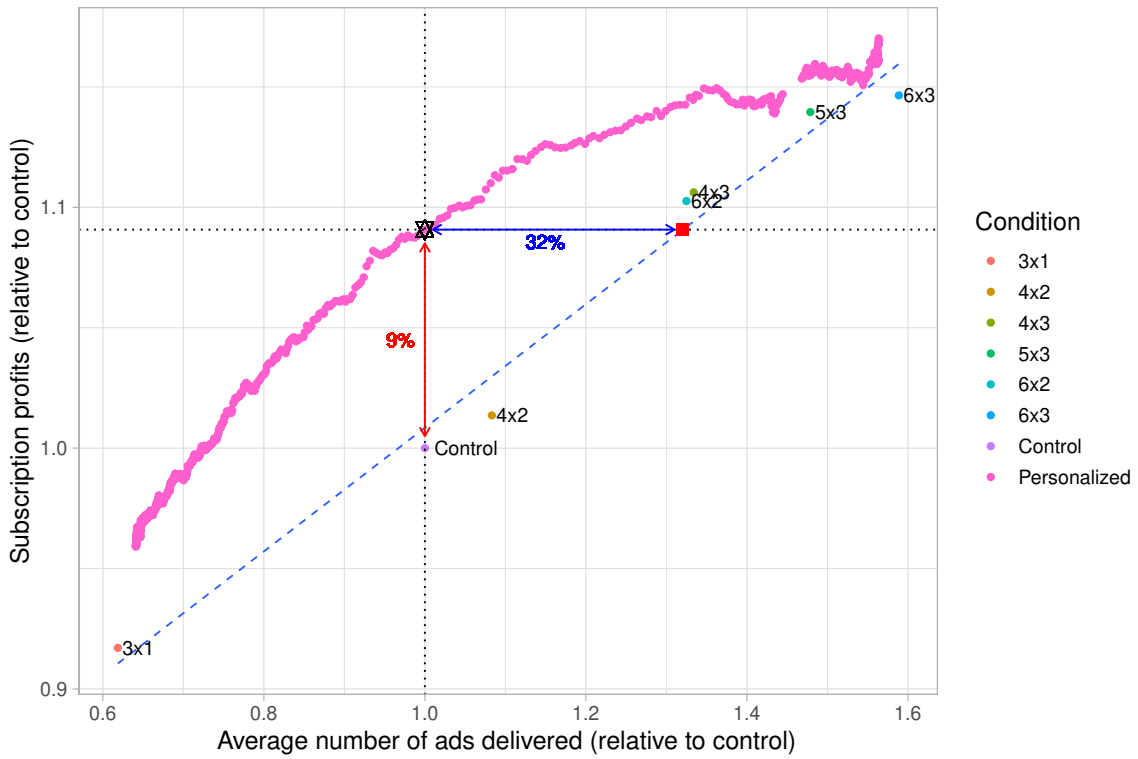

Figure 24: Change in subscription profits as a function of the number of ads served using a joint prediction model and long-term 18-month outcomes. Using long-term outcomes improves the performance of the personalized ad load policy, compared gains in Figure 9.

G. Adjusted standard errors for the analyses in section 5

To measure the average treatment effect of ad load on consumption and subscriptions in section 5, we normalize the outcomes relative to their average in the control and scale them to measure differences relative to the control conditions. These average outcomes in the control cell are calculated based on roughly 1 million users. Given the substantial sample size, the standard errors around these averages are quite tight, so they are unlikely to significantly influence the standard errors of our estimates. Nevertheless, to address concerns regarding standard errors, we will outline the analysis in section 5 below and subsequently offer a correction for the standard errors. In equation (4), we considered:

$$\tilde{\mathbf{Y}}_i = 100 \cdot \frac{\mathbf{Y}_i}{\frac{\sum_{j \in \mathcal{C}} \mathbf{Y}_j}{\mathcal{N}_\mathcal{C}}}, \tag{24}$$

where $\mathcal{C}$ and $\mathcal{N}_\mathcal{C}$ are the set of users in the control condition, and the number of users in control, respectively. We then consider the following instrumental variable regression:

$$\tilde{\mathbf{Y}}_i = \alpha + \beta \cdot \mathcal{A}_i + \epsilon_i, \tag{25}$$

The issue is that $\tilde{\mathbf{Y}}_i$ is an estimate derived from observations in the control cell. Without our data agreement restrictions to share the value of actual outcomes, the regression we would estimate would be:

$$\mathbf{Y}_i = \alpha + \beta^* \cdot \mathcal{A}_i + \epsilon_i^*, \tag{26}$$

Note that the reported β in (25) is essentially:

$$\beta = \frac{\beta^*}{m_\mathbf{Y}}$$

where $m_y = \frac{\sum_{j \in \mathcal{C}} \mathbf{Y}_j}{\mathcal{N}_\mathcal{C}}$ represents the average outcome in the control group. In the equation above, we have omitted the scaling factor (100) for simplicity. Our goal is to do inference on β. Note that we can estimate β^* using (26); let $\beta^* \sim \mathcal{N}(\mu_{\beta^*}, \sigma_{\beta^*})$ and let $m_\mathbf{Y} \sim \mathcal{N}(\mu_m, \sigma_m)$ be the estimate for the mean of the outcomes in the control cell. Our goal is to do inference on $\beta = \frac{\beta^*}{m_\mathbf{Y}}$. Following the delta method, we consider $g(x, y) = \frac{x}{y}$ and evaluate $g(\beta, m)$. First, observe that:

$$\frac{\partial}{\partial \beta} g(\beta, m) = \frac{1}{m},$$

and

$$\frac{\partial}{\partial m} g(\beta, m) = -\frac{\beta}{m^2}.$$

Using the Taylor expansion of g(.,.), we have

$$\mathbf{E}\left[g(\beta^*, m_\mathbf{Y})\right] \approx \frac{\mu_{\beta^*}}{\mu_m}$$

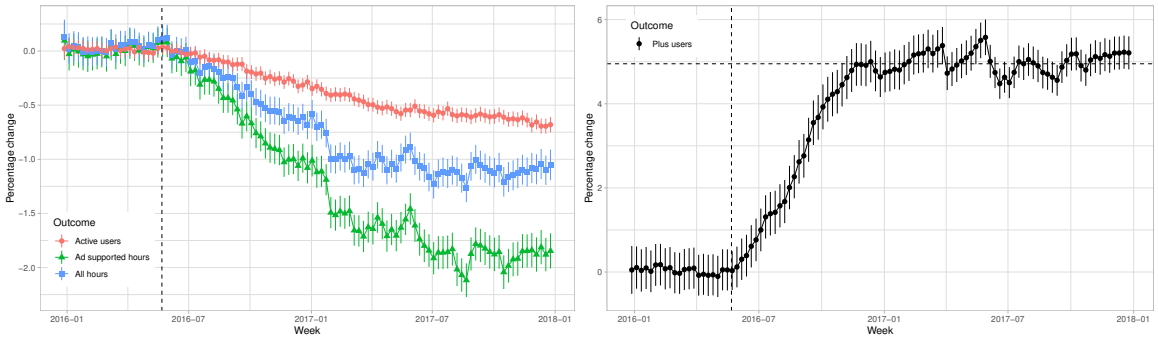

(a) Effect of ad load on consumption

(b) Effect of ad load on subscriptions

Figure 25: This figure presents the instrumental variable regression estimates for the marginal effect of an additional ad per hour on consumption and plus subscriptions. Outcomes are normalized based on their average in the control condition. It serves as a counterpart to Figure 5 in the manuscript, with corrections for standard errors taking into account that the average in the control cell itself is an estimate.

and

$$
\begin{aligned}
Var(g(\beta^*, m_{\mathbf{Y}})) = Var\left(\frac{\beta^*}{m_{\mathbf{Y}}}\right) &\overset{(1)}{\approx} \left(\frac{\mu_{\beta^*}}{\mu_m}\right)^2 \left(\frac{\sigma^2_{\beta^*}}{\mu^2_{\beta^*}} + \overbrace{\frac{\sigma^2_m}{\mu^2_m} - 2\frac{Cov(\beta^*, m_{\mathbf{Y}})}{\mu_m \cdot \mu_{\beta^*}}}^{\text{adjustments}}\right) \\
&\overset{(2)}{\leq} \left(\frac{\mu_{\beta^*}}{\mu_m}\right)^2 \left(\frac{\sigma^2_{\beta^*}}{\mu^2_{\beta^*}} + \frac{\sigma^2_m}{\mu^2_m} + 2\left|\frac{\sigma_{\beta*} \cdot \sigma_m}{\mu_m \cdot \mu_{\beta^*}}\right|\right),
\end{aligned}
\tag{27}
$$

where (1) follows from the Taylor expansion of $g(\cdot, \cdot)$ around $(\beta^*, m_{\mathbf{Y}})$ and the delta method, and (2) is a consequence of the triangle inequality and the Cauchy-Schwarz inequality. To avoid dealing with the covariance of β^* and $m_{\mathbf{Y}}$ we bound it above using the Cauchy-Schwartz inequality. This calculation provides a conservative estimate for the standard error of β. We employ this formula and re-compute the confidence intervals depicted in Figure 5 of the revised manuscript. It is worth noting that the confidence intervals that treat the normalizing factor as a constant and not as an estimate would omit the terms labeled as "adjustments" above in equation (27). As previously mentioned, since the average $m_{\mathbf{Y}}$ is derived from a large user sample, σ_m is quite small. This adjustment accounts for roughly 5% (or less) of the unadjusted standard error across various outcomes. We present the counterpart to Figure 5 in the manuscript, with adjusted standard errors, in Figure 25.