

# Personalized Versioning: Product Strategies Constructed from Experiments on Pandora

Ali Goli\*, David H. Reiley†, Hongkai Zhang‡

## Abstract

The role of advertising as an “implicit price” has long been recognized by economists and marketers. However, the impact of personalizing implicit prices on firm profits and consumer welfare has not been studied. We first conduct a set of large-scale field experiments on Pandora by exogenously shifting the ad load, that is the implicit price, for ad-supported users. We then use a state-of-the-art machine learning model to examine the heterogeneous treatment effects of firm’s interventions on listeners behavior, both in terms of listening hours and in terms of the propensity to subscribe to the ad-free version of the product. We next show that by reallocating ads across individuals, the firm can improve subscription profits by 10% without reducing total profits generated from advertising. To achieve the same subscription rate using a uniform ad-allocation policy, the firm would need to increase the number of ads served on the platform by more than 30%. Furthermore, the gains from personalization emerge quickly after implementation, as subscription behavior adapts to changing ad load relatively quickly. We also evaluate the welfare implications of personalizing implicit prices. Our results show that, on average, consumer welfare drops by 2% with the proposed personalization strategy, and the effect seems to be more pronounced for users that have a higher willingness to pay.

**Keywords**— Advertising, Personalization, Field Experiments, Heterogeneous Treatment Effects, Machine Learning

---

\*University of Washington, agoli@uw.edu

†SiriusXM Pandora, dreiley@pandora.com

‡SiriusXM Pandora, hzhang@pandora.com

---

# 1 Introduction

The abundance of free online content creates a challenge for online-content providers to monetize their platforms. In the mid 1990s and early 2000s, to attract large audiences and generate advertising revenues, many firms offered online content for free (Edgecliffe-Johnson 2009). As the industry matured, a number of content publishers experimented with subscription paywalls (Pérez-Peña and Arango 2009). Although some firms, such as Netflix, have earned substantial profits through this strategy, the transition to a subscription-only model has not been especially easy for most firms. For instance, using data on a media publisher’s website visits, Chiou and Tucker (2013) show that instituting paywalls led to a 51% drop in online visits. The trade-off between viewership and subscription profits has led a number of firms, including Hulu, YouTube, Spotify, and Pandora, to adopt a hybrid approach, offering both an ad-supported free version and an ad-free subscription version. An interesting question is how those multiple versions should be designed and priced.

In the age of the Internet, digital products have become highly customizable, both in terms of content (e.g., Pandora’s personalized radio stations) and in terms of pricing. Although academic authors have discussed the returns to personalizing subscription prices (Shiller et al. 2013) or to engaging in fine-grained group-pricing strategies (Dubé and Misra 2017), these ideas are rarely implemented in industry.<sup>1</sup> Firms such as Amazon and Staples have faced public backlash for experimenting with charging different prices to different customers (CNET 2002, Valentino-DeVries et al. 2012).

Fear of customer backlash has led many firms to instead adopt “versioning” strategies (Shapiro et al. 1998), where the seller offers each customer a menu of different product options, for example ad-supported and paid subscriptions, and allows customers to self-

---

<sup>1</sup>The only empirical investigation of fine-grained group pricing that we are aware of is Dubé and Misra (2017). Fine-grained group pricing approaches the ideal of personalized pricing as the number of groups becomes larger and the group size becomes smaller.

select into choosing one of them. These versions can further be customized based on what the publisher knows about the customer, such as when Pandora customizes music streams for listeners based on what they have thumbed up and down. The public appears to have a much more positive view of personalized product content than of personalized prices, because many people feel that one group of consumers getting different prices for the same product or service is unfair. Versioning and group pricing are two distinct strategies for the more general problem of price discrimination.<sup>2</sup>

In 2015, the White House’s Council of Economic Advisors report (CEA 2015) noted that it is unclear which of these two strategies will become more prevalent in the era of Big Data:

“It is difficult to predict how big data will influence the prevalence of versioning. If it becomes easier to predict individual customers’ willingness to pay and charge different prices for an identical product, versioning may be replaced by personalized pricing. On the other hand, versioning has the benefit of reducing concerns about inequity that arise with personalized pricing, and big data may facilitate versioning strategies based on “mass customization,” particularly for information goods that can be customized at relatively little incremental cost.”

Our goal is to demonstrate the role of advertising as an instrument for combining the two strategies into an idea that might be called “personalized versioning”: consumers choose between two versions of a product offering, one of which has quality personalization based on consumer characteristics. In this paper, we report the results and analysis of a field experiment conducted on Pandora during 2016-2017. During this period, Pandora offered two products: (i) the Pandora Plus subscription product, and (ii) the ad-supported product which was offered free of charge. Both of these products were *non-interactive radio* products,

---

<sup>2</sup>Pigou (1920) originally gave versioning and group pricing the names “second-degree” and “third-degree” price discrimination, respectively. We prefer the more modern terms, versioning and group pricing, both coined by Shapiro and Varian (Shapiro et al. 1998), because they are more evocative and do not create the confusion caused by calling two strategies “second-degree” and “third-degree” when they cannot easily be ranked.

meaning the listener could not listen to an audio track on demand but could create stations based on a favorite artist or track, and personalize her stations by thumbing songs up and down. Both products used the same music catalog and user interface; the main difference between the two was that Pandora Plus was ad-free, whereas ad-supported Pandora listeners would encounter ads between tracks, when switching between stations, or when skipping tracks.<sup>3</sup> To manage the trade-off between ad and subscription revenues between these two versions of the product, Pandora has two levers: (a) setting subscription prices (explicit price) and (b) changing the number of ads served to ad-supported listeners (implicit price). Pandora can shift the implicit price by either changing the price of ad impressions or by allocating ads across listeners. Particularly, the platform can change the number of ads served to each listener by setting the frequency of scheduled commercial interruptions (or “ad pods”) as well as the number of scheduled ads per pod.

Increasing the frequency or length of ad pods increases the opportunities to serve an ad, which we refer to as “ad capacity.” By contrast, the listener’s realized “ad load,” or actual number of ads served per hour, also depends on the consumption level<sup>4</sup> and advertisers’ demand for each listener. Our goal in this paper is to study the gains that arise from personalizing the ad capacity at the individual level while holding the overall ad inventory fixed. To solve the firm’s optimization problem, we need to address the following issues:

- (a) **Endogeneity:** Even holding fixed the ad-serving strategy, the realized ad load experienced by different listeners correlates with both their behavior and advertisers’ demand. For instance, apart from differences in advertisers’ demand for each listener, the longer listeners stay on the platform, the harder it is for Pandora to fill their full ad capacity. These correlations create an endogeneity problem, which motivates our

---

<sup>3</sup>Pandora Plus offers some additional features including offline listening, the ability to replay songs, and higher-quality audio.

<sup>4</sup>Given listener-level frequency caps standard in these ad campaigns, the longer listeners stay on the platform, the more difficulty they have filling their full ad capacity.

randomized experiment.

- (b) **Partial control over realized ad load:** The realized ad load not only depends on the firm’s ad-scheduling policy, but also on advertiser demand (some listeners are in higher demand than others) and on listener behavior (e.g., the length of a listening session, which interacts with advertiser frequency caps). This partial control causes one listener to receive more ads than another even when both receive the same policy from the firm, and this difference in realized ad load needs to be taken into account in the firm’s optimization problem.
- (c) **Extensive-margin adjustments:** Listeners assigned to higher ad load conditions in a given period may change their status by becoming inactive or by switching to a Pandora Plus subscription. Such *extensive-margin* adjustments can themselves affect treatment exposure levels.
- (d) **Intensive-margin adjustments:** The total number of ads served to each individual depends on both the ad load and the *intensive margin* of consumption, that is listening hours in a given period, which needs to be accounted for when solving the firm’s optimization problem.

We exploit a set of large-scale field experiments that exogenously shift the ad-pod frequency and length for over seven million Pandora listeners during 2016-2017. We then use a state-of-the-art machine learning model to learn the heterogeneous treatment effects of the firm’s interventions on the realized ad load, consumers’ extensive-margin decisions (switching between outside option, plus, and ad-supported options), and consumers’ intensive-margin decisions (number of ad-supported hours consumed). To learn the heterogeneous treatment effects of a firm’s interventions, we combine insights from structural estimation with neural networks.<sup>5</sup> We use split neural networks (Kim et al. 2017) to impose exclusion restrictions

---

<sup>5</sup>For a few use cases of neural networks for estimating structural models, see Wei and Jiang (2020), and

that enable the model to better learn heterogeneous treatment effects. Our architecture is similar to Shalit et al. (2017), who use neural networks to predict individual-level outcomes across different counterfactuals. Shalit et al. (2017) and Farrell et al. (2018) show neural networks are effective in learning treatment heterogeneity and achieve comparable performance to direct methods for learning heterogeneous treatment effects such as causal forests and treatment-effect projection (Wager and Athey 2018, Hitsch and Misra 2018).

Subsequently, we solve the firm’s optimization problem and evaluate the impact of the prescribed policy using an inverse probability-weighted (IPW) estimator; see Horvitz and Thompson (1952) for its use and origins in statistics and Hitsch and Misra (2018), Yoganarasimhan et al. (2020), Rafieian and Yoganarasimhan (2021) for a few recent examples of IPW estimators in marketing. Our results demonstrate that holding fixed the total number of ads served, the firm can improve subscription profits by 10% without any loss in total ad revenue. To achieve the same subscription rate with a uniform allocation strategy, the firm would have to serve 30% more ads, which would have a negative impact on hours listened. We then study the impact of the proposed policy on consumer welfare, and show that, on average, consumer welfare declines by 2%. To the best of our knowledge, this study is the first to use a field experiment to evaluate the returns to personalizing product quality. Our results inform policy makers, and firms regarding the implications and returns to personalizing price/quality of product offerings.

In analyzing the potential that arises from “personalized versioning,” our findings contribute to four strands of academic research. First, we contribute to the literature that measures the returns to personalization. Researchers have studied personalization in a wide variety of contexts, examples include prices (Rossi et al. 1996, Shiller et al. 2013, Dubé and Misra 2017), e-mail content (Ansari and Mela 2003, Sahni et al. 2018), website design (Hauser et al. 2009), search rankings (Yoganarasimhan 2020), promotions (Hitsch and Misra 2018, Kaji et al. (2020).

---

Yoganarasimhan et al. 2020), mobile advertising (Rafieian and Yoganarasimhan 2021), and ad sequencing (Rafieian 2019). While the research that involves personalizing advertisement measures returns to changing ad content or targeting ads to improve ad effectiveness or revenue from ads, our work focuses on personalizing ad load as an instrument for improving subscription revenue.

Second, we add to the literature that models product quality as an endogenous decision. In a single-product setting, Spence (1975) shows that a monopolist may offer a higher or lower quality level than the social optimum. In a multi-product regime, Mussa and Rosen (1978) and Maskin and Riley (1984) demonstrate that to attract high-type customers, the monopolist has the incentive to degrade the quality of lower-end products, which creates a negative externality on customers with lower quality valuation. This finding relates to the “damaged goods” literature, where a firm has the incentive to “damage” a developed product to build a lower-quality version (Deneckere and Preston McAfee 1996). One approach for implementing versioning is by bundling a good with a “bad” like waiting time or search cost (Salop 1977, Chiang and Spatt 1982). McManus (2007), Clerides (2002), and Verboven (2002) provide evidence of versioning in specialty coffee, book publishing, and European auto industries, respectively. Crawford and Shum (2007) measure the extent of quality degradation in cable-television subscription bundles, and Crawford et al. (2015) study the welfare effects of endogenous quality choice. We show that personalizing product quality helps limit such distortions. In particular, we find that our proposed optimal personalization of ad load produces subscription benefits equivalent to a uniform increase in ad load (a degradation in quality) of about 30%. Previous research in marketing has shown service-quality variation over time can improve profits by increasing customer retention, for some consumers, a phenomenon potentially explained by risk aversion in the consumer learning process (Sriram et al. 2015). Our findings show that another kind of variation in quality of service – across consumers – can improve profits in a product line by inducing users to

upgrade to the higher-end products. The substitution between products offered in a product line, along with switching costs between products, presents yet another opportunity for firms to leverage changes in quality of service as a screening mechanism.

Third, our results contribute to the literature that considers product-line strategy in offering free (ad-supported) and paid versions of information goods (Shapiro et al. 1998). On the theoretical side, this literature extends the versioning framework in Mussa and Rosen (1978) for information goods that rely on both advertising and subscriptions as sources of revenue. Tåg (2009) shows that introducing an ad-free subscription decreases consumer welfare because the firm has the incentive to increase advertising in the ad-supported version to earn more profits from the paid product. Researchers have studied the role of dynamics, consumer heterogeneity, competition, quality learning, and advertiser heterogeneity on the revenue model adopted by firms (Kumar and Sethi 2009, Sato 2019, Prasad et al. 2003, Godes et al. 2009, Halbheer et al. 2014, Lin 2020).

On the empirical side, Chiou and Tucker (2013) show introducing paywalls can dramatically reduce viewership. Lambrecht and Misra (2017) present evidence for counter-cyclical quality improvements to ESPN’s free service. The authors argue consumers are heterogeneous in their valuation of content, which may vary over time. This heterogeneity rationalizes a quality-discrimination mechanism along the time dimension. In this paper, we establish the trade-offs between ad and subscription revenue, and then show that by personalizing the ad schedule (quality of service), the firm could improve subscription profits. Although the idea of using ads as a screening mechanism in freemium products is not new (Tåg 2009, Sato 2019), we are not aware of any paper that has empirically investigated the personalization of product quality, especially in the advertising context.

Finally, our findings are also relevant to the price-discrimination literature. Although the amount of advertising on ad-supported media is a quality of service measure, it can also be viewed as an *implicit price* that is charged in units of time rather than money. The



idea of regarding ads as a price and their possible negative impact on demand for media is not new (Becker and Murphy 1993, Gentzkow 2007, Wilbur 2008, Goldstein et al. 2014, Huang et al. 2018). However, to the best of our knowledge, the returns to personalizing this implicit price and its welfare implications has not been studied before. Theoretically, third-degree price discrimination could improve social welfare (Varian 1985) or could even improve consumer surplus by expanding output (Cowan 2012). Dubé and Misra (2017) examine the returns to an extreme form of group pricing using a large-scale field experiment at Ziprecruiter. They show that while firm profits improve by about 10%, consumer surplus falls less than 1%. One of the main differences between our problem and a classical price-discrimination problem is the fact that the consumer has the option to pay both with time and money. Therefore, listeners are screened based on both willingness to pay and their marginal value of time.<sup>6</sup> This means that, the correlation between willingness to pay in time and money units could influence the effectiveness of our personalization algorithm. For instance, income and marginal value of time could be positively correlated (Aguiar et al. 2011, 2013). Furthermore, income is likely to be negatively correlated with price sensitivity. On the one hand, the algorithm has the incentive to move more ads towards higher-income individuals because they are less price sensitive and more likely to upgrade to the paid subscription. On the other hand, higher-income individuals may place a larger value on their time and are also more likely to churn when faced with more ads. Because of these trade-offs, the ad-allocation mechanism and its welfare implications are a priori ambiguous.

The rest of this paper is organized as follows. We first introduce a conceptual model to discuss the personalized versioning idea and illustrate trade-offs between ad and subscription revenues. We then discuss the field experiments conducted at Pandora Media and present reduced-form evidence to illustrate the impact of changing ad load on listeners’ choices.

---

<sup>6</sup>The idea of using differences in valuation of time for optimizing menu offerings and its welfare implications has been discussed in Salop (1977), and Chiang and Spatt (1982). However, we are not aware of any empirical work that has investigated a personalized policy that leverages this heterogeneity.

Subsequently, we use a state-of-the-art machine learning model to learn the heterogeneity in response to changes in ad load among listeners. Which are then leveraged to reallocate ads to improve firm profits. Finally, we discuss the prescribed policy and its welfare implications.

## 2 Conceptual Model

In this section, we illustrate the benefits of personalized versioning 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:

$$u(z, p \mid \theta, \beta, \alpha) = \max \begin{cases} 0 & \text{outside option,} \\ \theta - \alpha z & \text{ad supported,} \\ \theta - \beta p & \text{paid version,} \end{cases} \quad (1)$$

where  $\theta$  is the utility from consuming the product,  $z$  specifies the ad load<sup>7</sup>, and  $p$  is the subscription price. In the ad-supported condition, users are effectively paying with their time by listening to ads. Therefore, parameters  $\alpha$  and  $\beta$  reflect how time and money are valued by users, that is users with higher/lower values of  $\alpha$  and  $\beta$  are more/less sensitive to ads and prices, respectively. Also, let  $\gamma$  and  $c$  be the revenue per ad and marginal cost of offering the service, respectively. Finally, let  $\theta_\alpha = \frac{\theta}{\alpha}$  and  $\theta_\beta = \frac{\theta}{\beta}$ . 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 1 illustrates the decision regions for different types when price is set to  $p$  and ad load is equal to  $z$ .

---

<sup>7</sup>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.




Figure 1: Decision regions for different types  $(\theta_\alpha, \theta_\beta)$  for price vector  $(p, z)$ , where  $p$  and  $\gamma z$  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.

A monopolist that can perfectly discriminate along both ad load and price dimensions will maximize its profits for each type  $(\theta_\alpha, \theta_\beta)$ . 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 2 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:

$$\Pi^* = \int_{\theta_\beta=c}^{\infty} \int_{\theta_\alpha=0}^{\frac{1}{\gamma}\theta_\beta} (\theta_\beta - c) f(\theta_\alpha, \theta_\beta) d\theta_\alpha d\theta_\beta + \int_{\theta_\alpha=\frac{c}{\gamma}}^{\infty} \int_{\theta_\beta=0}^{\gamma\theta_\alpha} (\gamma\theta_\alpha - c) f(\theta_\alpha, \theta_\beta) d\theta_\beta d\theta_\alpha, \quad (2)$$

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




Figure 2: 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. We call this practice “personalized versioning,” which is akin to combining second- and third-degree price discrimination. The seller uses its information about the type of each customer to offer a personalized menu rather than only one product or price.

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

- **Separable types:** Consider a monopolist (he) who needs to provide service to a customer (she). The monopolist knows that with probability  $\rho_1$ , 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  $\theta_\alpha^{(1)} > \theta_\alpha^{(2)}$




Figure 3: 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.

and  $\theta_\beta^{(2)} > \theta_\beta^{(1)}$  (see Figure 3 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  $\rho_1$  and  $\rho_2 = 1 - \rho_1$ . This time  $(\theta_\alpha^{(2)}, \theta_\beta^{(2)}) > (\theta_\alpha^{(1)}, \theta_\beta^{(1)})$ ; see Figure 4 for visual illustration. In this case, the optimal menu corresponds to one of the five red dots in Figure 4. Depending on the values of  $\gamma$ ,  $c$ ,  $\rho_1$ ,  $\rho_2$ ,  $(\theta_\alpha^{(1)}, \theta_\beta^{(1)})$ , and  $(\theta_\alpha^{(2)}, \theta_\beta^{(2)})$ , either of these can be the optimal menu to offer. The only case where both




Figure 4: Inseparable condition. The uncertainty in the types cannot be perfectly dealt with by using a personalize 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.

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<sup>8</sup>.

The examples above illustrate that the combination of group pricing and versioning can make use of both the partial information that the firm may have and the customer's private information. We refer to this approach as personalized versioning. Note the application of personalized versioning is not limited to the case where uncertainty is present in parameter estimates, but also applies in random utility models even when parameter uncertainty is

<sup>8</sup>Recall from Figure 2 that if  $\frac{\theta_\beta}{\theta_\alpha} > \frac{1}{\gamma}$ , the seller is better off providing the ad-supported service.

neglected. Before moving on to the next section, we use a random utility model that corresponds to (1) to further discuss the trade-offs involved in personalizing ad load. Consider the following model:

$$u(z, p; \theta, \beta, \alpha) = \max \begin{cases} \epsilon_0 & \text{outside option,} \\ \theta - \alpha z + \epsilon_a & \text{ad-supported,} \\ \theta - \beta p + \epsilon_p & \text{paid subscription,} \end{cases}$$

where  $\epsilon_0$ ,  $\epsilon_a$ , and  $\epsilon_p$  are independent random variables that follow a type-I extreme value distribution. The rest of the parameters are defined as in (1).

Those who tend to have higher willingness to pay in money terms are likely more sensitive when paying in time units. For instance, higher-income individuals tend to have lower price elasticity but higher marginal value of time (Aguilar et al. 2011, 2013). This confound can generate a negative correlation between price ( $\beta$ ) and time sensitivity ( $\alpha$ ) in this set up. On the one hand, the monopolist has the incentive to serve fewer ads to more ad-sensitive users, that is larger  $\alpha$ . On the other hand, the same users likely have higher willingness to pay, that is smaller  $\beta$ , and are more likely to upgrade to the subscription service if they face higher prices in time terms.

Note that if the seller were to offer only the ad-supported version, customers with higher ad sensitivity would receive fewer ads. However, the correlation structure between ad sensitivity ( $\alpha$ ) and price sensitivity ( $\beta$ ) can lead to both higher or lower frequency of ads for more ad-sensitivity users. Let us consider the problem of personalizing the ad load  $z$  given a fixed price  $p$  for the subscription service. Let us assume the marginal cost of offering service is  $c$ , and revenue from serving each ad is  $\gamma$ . The problem the service provider faces is to

$$\underset{z}{\text{maximize}} \quad \underbrace{\frac{e^{\theta-\alpha z}}{1 + e^{\theta-\alpha z} + e^{\theta-\beta p}}(\gamma z - c)}_{\text{expected profits from ads}} + \underbrace{\frac{e^{\theta-\beta p}}{1 + e^{\theta-\alpha z} + e^{\theta-\beta p}}(p - c)}_{\text{expected profits from subscription}}. \quad (3)$$

As discussed above, on the one hand, a higher ad load leads to more profits from subscriptions and increases revenue conditional on being an ad-supported member; on the other hand, it lowers profits from the ad-supported users by increasing churn. Furthermore, the correlation structure between ad and price elasticity can lead to either higher or lower ad load for users with higher ad elasticity. To illustrate this trade-off, let us hold the price fixed and optimize the ad load for a set of given parameters  $(\theta, \alpha, \beta, p, \gamma, c)$  while enforcing different correlation structures between  $\alpha$  and  $\beta$ . Let us assume  $\theta = 4$ ,  $\beta = 2 - (0.1)\alpha$ ,  $\gamma = 0.1$ ,  $c = 1$ , and  $p = 5$ , and let  $\alpha$  vary between 0 and 1. The optimal ad load ( $z$ ) as a function of ad sensitivity is strictly decreasing and is plotted in panel (a) of Figure 5. However, if the correlation structure between price and ad sensitivity is stronger, say,  $\beta = 2 - 2\alpha$ , the optimal ad load could be a non-monotonic function of ad sensitivity as depicted in panel (b) of Figure 5. This example illustrates that in our multi-product setting, forming ex-ante predictions on which customer segments, for example high/low income, bear the cost of personalizing the ad load is difficult.






Figure 5: Optimal ad load as a function of ad sensitivity. (a) In this panel, price sensitivity  $\beta = 2 - (0.1)\alpha$ , the optimal ad load is a decreasing function of ad sensitivity because gains from subscriptions do not outweigh the losses. (b) The price sensitivity  $\beta = 2 - 2\alpha$ , and  $\alpha$  varies between 0 and 1. In this case, optimal ad load is a non-monotonic function of ad sensitivity, because users with higher ad sensitivity are less price sensitive, and uplift from subscriptions outweighs the losses.

### 3 Field Experiments at Pandora Media

Now that we have built a conceptual model to understand the trade-offs, we delve into the field experiments used in this study. At the time of this experiment, Pandora offered two tiers of products: (a) ad-supported and (b) plus. The ad-supported and plus versions are both “radio”<sup>9</sup> products and have the same music catalog and user interface. Whereas the plus subscription is ad free with a monthly subscription fee of \$4.99, ad-supported listeners are exposed to video/audio ads in exchange for using the service.

<sup>9</sup>The radio products offer quasi-audio-on-demand services as they personalize the radio stations to cater to listener preferences using the feedback (thumbs up/down, and skips) provided by the listeners. In the second quarter of 2017, Pandora started offering the premium service, which was an ad-free audio-on-demand product.

Before delving into the analysis, we illustrate that our randomization algorithm has achieved its goal and treatment assignment is not systematically correlated with any covariates of interest. We compare the users across treatment and control groups who were active in the first quarter of 2016 before the experiment, and compare their age, gender, and some of the other key behavioral variables in the pre-treatment period in Table 1<sup>10</sup>. Except for all hours, which is slightly higher for the treatment group at  $P < 0.05$ , the rest of the variables are not statistically different from each other. Overall, this table shows the treatment and control groups are not systematically different, and confirms the treatment assignment has been random.<sup>11</sup>

---

<sup>10</sup>Due to our agreement with Pandora we cannot reveal the actual estimates for some of these features; therefore, we have normalized those features such that the sample average of the control group is equal to 100.

<sup>11</sup>Note that within the treatment group, we also have different treatment conditions, for example 6x3 versus 4x2 group, and the assignment within these groups has also been randomized. We verified that the randomization worked as expected; however, here we are only comparing the overall treatment and control groups.

Table 1: Comparing treatment and control groups across some of the pre-treatment features calculated during the first quarter of 2016. All features except for gender and zip code mean income are normalized such that the mean of control is equal to 100. Standard errors are reported in the parenthesis.

Variable	Treatment	Control	Difference
All hours	100.437 (0.132)	100 (0.351)	0.437* (0.192)
Ad-supported hours	100.341 (0.144)	100 (0.386)	0.341 (0.21)
Thumbs (up/down)	100.058 (0.218)	100 (0.586)	0.058 (0.318)
Thumbs up	100.183 (0.214)	100 (0.574)	0.183 (0.313)
Skipped tracks	100.358 (0.216)	100 (0.578)	0.358 (0.314)
Station changes	100.07 (0.393)	100 (1.053)	0.07 (0.573)
Age	99.946 (0.032)	100 (0.086)	-0.054 (0.047)
Gender (male = 1)	0.451 (0)	0.452 (0.001)	-0.001 (0.001)
Zip code mean income	73,436.775 (24.738)	73,438.642 (66.062)	-1.867 (36.055)

The set of experiments used in this paper shift the time spent consuming ads, by changing the number of audio ads played between music tracks. Ads are delivered using a set of timers. When the user starts a session or when an ad pod is delivered, the timer is reset. At the beginning of every track, the system checks to see if the user is eligible to receive an ad pod, that is if the timer is set. The length of an ad pod determines the number of ads that can be served in an ad break. The experiments shift both *frequency* and *length* of ad pods using six experiment conditions and a control cell, which represents the current strategy employed by the firm. The experiment cells are presented in Table 2 below:

## Personalized Versioning

Table 2: Experiment conditions for audio-ad arms. The size of each treatment cell is specified as the percentage of all listeners on the platform. The control cell consists of 1% of the total listeners. Rows and columns correspond to pod frequency and size, respectively.

		Audio ads per interruption		
		1	2	3
Audio ad interruptions per hour	3	1%		
	4		2%	0.5%
	5			0.5%
	6		0.5%	0.3%

From this point forward, we refer to each cell in our experiment as FxL, where F and L are *intended* pod frequency, and length, respectively. For instance, the 3x1 condition refers to a treatment with three pods per hour, and each pod consists of one ad. The control condition is similar to the 4x2 condition, but the first ad pod within each listening session is constrained to have at most one ad. The control condition comprises 1% of the total listeners on Pandora.

As highlighted before, pod length and frequency determine *ad capacity* rather than ad load. Also note that a listener in the 6x3 condition ends up becoming eligible for an ad pod fewer than six times per hour, because the song endings do not perfectly align with the timers. In other words, the ad capacity, that is, the number of opportunities to show an ad per hour, in the FxL condition ends up being far less than  $F \cdot L$ . Although the experiments shifts ad capacity by changing pod frequency and size, the realized number of ads shown to each listener (ad load) also depends on advertisers’ demand. In other words, the experiment shifts ad capacity, that is the rate at which ads *can be shown* to users, rather than the ad load, that is *the realized rate* of ads for each user.

Figure 6 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. This finding is indicative of the fact that filling higher capacities for users tends to be more difficult, because running out of ads to




Figure 6: Realized ad load across different treatment arms

serve in the higher-capacity conditions is more probable. Table 3 reports the average realized ad load, capacity, and fill rate during the first year of the experiment. Realized ad capacity is the number of opportunities that the ad delivery system determines a listener as being eligible to receive an ad, though not all these opportunities get filled as the system may not be able to fetch ads to serve to users. The proportion of ad opportunities that were filled are referred to as the fill rate. As one would expect, the fill rate tends to fall as we move toward higher-capacity conditions; for example compare the 3x1 and 6x3 conditions. Finally, note the realized ad load depends on both the realized ad capacity and the fill rate; therefore, an X% increase in realized ad capacity does not necessarily translate into an X% increase in realized ad load.

## Personalized Versioning

Table 3: The realized ad load, capacity, and fill rate across treatment/control conditions. Note that, for example in the 6x3 condition, the realized ad capacity differs from 18 ads per hour, because the times when songs finish do not perfectly align with the times that users become eligible to see an ad. Also note that an X% increase in realized ad capacity does not necessarily translate to an X% increase in realized ad load, because the ad load also depends on the fill rate, and the system is more likely to run out of ads in higher-capacity conditions.

	Experiment condition						
	3x1	4x2	4x3	5x3	6x2	6x3	Control
Realized ad load	2.947 (0.006)	4.659 (0.007)	5.541 (0.008)	6.123 (0.011)	5.602 (0.008)	6.665 (0.023)	4.208 (0.008)
Realized ad capacity	3.512 (0.007)	6.326 (0.009)	8.289 (0.008)	9.35 (0.013)	7.789 (0.009)	10.347 (0.025)	5.56 (0.013)
Fill rate	0.853 (0)	0.738 (0)	0.676 (0.001)	0.665 (0.001)	0.723 (0.001)	0.657 (0.001)	0.763 (0)

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 7. 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.

## 4 Reduced-Form Results

Let us first illustrate the impact of ad load on the overall consumption and substitution to the subscription service. To that end, we plot the change in the realized ad load, ad revenue, active users, subscription revenue, ad-supported hours, and all hours, that is, the sum of

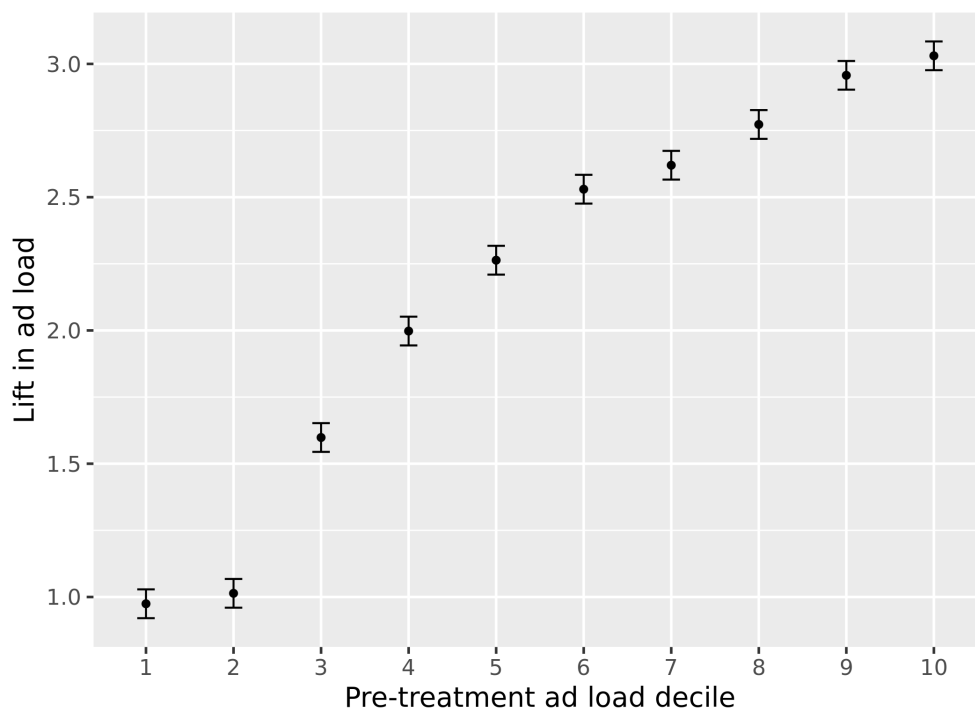


Figure 7: Realized lift in ad load (ads/hours) in the 6x3 condition relative to control as a function of pre-treatment ad load. Note the lift in ad load could vary drastically across different groups, due to differences in the attractiveness of different segments for advertisers.

ad-supported and plus hours, across the highest and lowest ad-load arms relative to the control condition in Figure 8. The figure measures each outcome of interest as a percentage change relative to the control arm. For instance, in the 3x1 condition, ad-supported hours increase by about 2.5% relative to control. Higher ad load leads to higher revenue from ads and subscriptions, but it drastically affects both extensive and intensive margins of consumption. Although in the short run, the firm can improve profits by increasing the ad load, a higher ad load can negatively affect the number of ad-supported hours, which reflects the long-run potential for ad revenue. Note the impact on the plus subscription stabilizes much faster than the change in other outcomes. This finding shows that even persistent short-run changes in quality of service in the ad-supported product can lead to substitution to the plus subscription. This finding, coupled with switching costs between plus and ad-supported products, presents an opportunity for the firm to improve subscription profits through temporary changes in the implicit prices.

Table 4 compares some of the key outcomes, including monthly active users, subscription rate, and ad-supported hours, across different treatment arms in December 2016, six months after starting the experiment. As expected, the higher-ad-load conditions led to fewer active users, ad-supported hours, but an increase in the number of subscribers to the paid version. Note the outcomes are measured relative to the control arm and as the outcome for the control condition is normalized to 100. Therefore, one can interpret these results as percentage changes relative to the control condition.

## 5 Estimation models

Our goal is to study the impact of personalizing quality of service on firm profits and consumer welfare. To optimize profits and evaluate welfare implications, we first need a demand model that reflects listeners' choice between being inactive (outside option), using the ad-



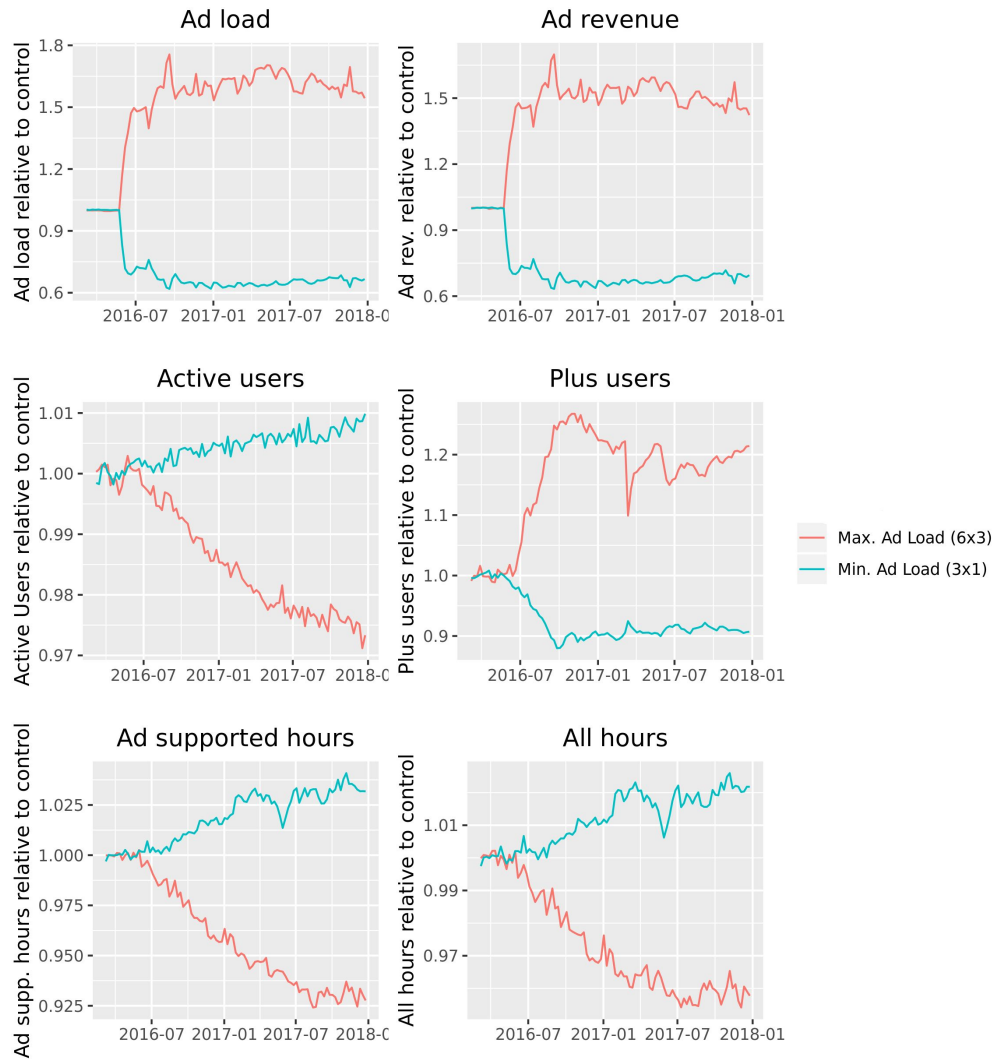


Figure 8: The trade-offs involved in optimizing ad load. The highest ad load condition seems to deliver 50% more ads compared to control, the ad revenue grows by about the same magnitude even though user churn nullifies some of it. Plus subscription revenue grows by about 25% as users substitute toward the ad-free version as the number of ads decreases. Note that the impact of the treatment on subscription rate stabilizes fairly quickly post-treatment. The number of hours spent on the platform reacts negatively to ad load.

Table 4: Change in monthly active users, subscription rate, and ad-supported hours relative to the control condition across different treatment arms.

	<i>Dependent variable:</i>		
	Active users	Subscription rate	Ad-supported hours
	(1)	(2)	(3)
Control (Baseline)	100.000*** (0.076)	100.000*** (0.735)	100.000*** (0.206)
3x1	0.138 (0.107)	-6.197*** (1.039)	1.530*** (0.291)
4x2	-0.078 (0.093)	2.234** (0.900)	-0.713*** (0.252)
4x3	-0.224* (0.132)	7.992*** (1.273)	-1.732*** (0.357)
6x2	-0.362*** (0.132)	11.212*** (1.274)	-2.198*** (0.357)
5x3	-0.340*** (0.132)	15.482*** (1.272)	-2.839*** (0.357)
6x3	-0.527*** (0.158)	17.984*** (1.529)	-3.854*** (0.429)
Observations	7,350,278	7,350,278	7,350,278
R <sup>2</sup>	0.00000	0.0001	0.00004
Adjusted R <sup>2</sup>	0.00000	0.0001	0.00004
Residual Std. Error (df = 7350271)	85.530	827.142	231.912

*Note:*

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

supported version, or paying for the subscription service. Also note the profit structure depends on listeners' subscription state. Although the profits generated by a paid subscriber are not a function of consumption intensity, the profits from an ad-supported user are a direct function of hours consumed and the *realized* ad load. In the following subsections, we first estimate a discrete-choice demand model where consumers choose between the outside option, ad-supported, and paid subscription. Then, conditional on being an ad-supported user, the number of hours consumed in a given period is estimated. Finally, as illustrated in section 3, the realized ad load depends on both the treatment condition and demand from advertisers. Hence, we construct a third model to account for the the partial control problem discussed in section 3. Combining these models enables us to optimize profits and study the welfare implications of personalizing ad load.

## 5.1 Demand model

We estimate a nested logit model to reflect users' decision between two nests of options, namely, a degenerate nest that includes the outside option and another nest where the user decides between the ad-supported and paid versions. Consider the following discrete-choice demand model:

$$u(\boldsymbol{\tau}; \mathbf{x}) = \max \begin{cases} \epsilon_0 & \text{outside option,} \\ \underbrace{\theta(\mathbf{x}) + \sum_j \eta_j(\mathbf{x}) \mathbb{1}_{\{\boldsymbol{\tau}=e_j\}} + \epsilon_a}_{v_1(\mathbf{x}, \boldsymbol{\tau})} & \text{ad-supported,} \\ \underbrace{\gamma(\mathbf{x}) + \epsilon_p}_{v_2(\mathbf{x})} & \text{paid version,} \end{cases} \quad (4)$$

where  $\mathbf{x}$  is a user-specific vector that includes exogenous and pre-experimental endogenous features at the listener level. The utility of consuming the outside option is normalized to zero, and the net utility of consuming the ad-supported and paid product in the control

condition is captured by  $\theta(\mathbf{x})$ , and  $\gamma(\mathbf{x})$ , respectively. This net utility consists of consumption utility and the disutility caused by ads and payment. One can impose further restrictions to disentangle these parts; however, our goal is to impose as few assumptions as possible. The experiment only affects the ad load across different conditions; therefore, the *treatment effect* only enters the utility of ad-supported product through  $\eta_j(\mathbf{x})$ . The treatment condition is represented by a binary vector  $\boldsymbol{\tau}$ , and  $e_j$  is a unit vector whose  $j^{th}$  element is equal to 1. The treatment effect of condition  $j$  is denoted by  $\eta_j(\mathbf{x})$ , which measures the change in utility of consumption for each treatment arm relative to the control condition. Finally,  $(\epsilon_0, \epsilon_a, \epsilon_p)$  follows a generalized extreme value (GEV) distribution that allows us to estimate a nested-logit model, where the probability of each option is as follows:

$$\mathbb{P}(Y = \text{outside, ad-supported, paid} | \boldsymbol{\tau}; \mathbf{x}) = \begin{cases} \frac{1}{1 + \exp(\lambda \cdot IV)}, \\ \frac{\exp(\lambda \cdot IV)}{1 + \exp(\lambda \cdot IV)} \frac{\exp\left(\frac{v_1(\mathbf{x}, \boldsymbol{\tau})}{\lambda}\right)}{\exp\left(\frac{v_1(\mathbf{x}, \boldsymbol{\tau})}{\lambda}\right) + \exp\left(\frac{v_2(\mathbf{x})}{\lambda}\right)}, \\ \frac{\exp(\lambda \cdot IV)}{1 + \exp(\lambda \cdot IV)} \frac{\exp\left(\frac{v_2(\mathbf{x})}{\lambda}\right)}{\exp\left(\frac{v_1(\mathbf{x}, \boldsymbol{\tau})}{\lambda}\right) + \exp\left(\frac{v_2(\mathbf{x})}{\lambda}\right)}, \end{cases}$$

where  $v_1(\mathbf{x}, \boldsymbol{\tau})$ , and  $v_2(\mathbf{x})$  are defined as in equation (4). Also,  $IV = \log\left(\exp\left(\frac{v_1(\mathbf{x}, \boldsymbol{\tau})}{\lambda}\right) + \exp\left(\frac{v_2(\mathbf{x})}{\lambda}\right)\right)$  is the inclusive value of the nest, and  $1 - \lambda \in [0, 1]$  reflects the correlation structure inside the nest.

Functions  $\gamma(\mathbf{x})$ ,  $\theta(\mathbf{x})$ , and  $\eta_j(\mathbf{x})$  are parameterized as neural networks, which allows them to be represented as flexible functional forms of pre-treatment features. Note that a neural network with a terminal softmax activation layer is effectively a flexible logit, and by restricting the values fed to the terminal layer, we can create flexible structural models. Because we are estimating a nested logit model instead of a simple logit one, we need to change the terminal layer of the neural network to reflect the probability structure imposed by the nested logit with a tunable parameter  $\lambda$ . We impose a few restrictions on the structure on

the neural network. First, treatment dummies only enter the last layer of the neural network and are multiplied by coefficients  $\eta_i(\mathbf{x})$  that are also parameterized as neural networks. This technique forces<sup>12</sup> the neural network to use the information provided by treatment dummies and has been used in Shalit et al. (2017) and Farrell et al. (2018). Second, because features that may explain heterogeneity in treatment  $\eta_i(\mathbf{x})$  may be very different from those that explain the cross-sectional heterogeneity ( $\theta(\mathbf{x})$  and  $\gamma(\mathbf{x})$ ), we use split neural networks (Kim et al. 2017) to separately fit values to each part. In particular, we let the network that learns  $\eta_j(\mathbf{x})$ , that is the “treatment effect,” be disjoint from the network that learns  $\theta(\mathbf{x})$  and  $\gamma(\mathbf{x})$ , that is the utility from consuming the ad-supported product. Split neural networks have been mainly used for parallel computing and to boost the training process. In this application, however, the issue is that the heterogeneity in the treatment effect may be explained by different types of features that would explain the cross-sectional heterogeneity, and these exclusion restrictions help the network to learn treatment heterogeneity more efficiently. Note that even though we impose a split structure, the networks are trained jointly. A schematic view of the architecture is presented in Figure 9. Note the purple part that learns  $\theta(\mathbf{x})$ , and  $\gamma(\mathbf{x})$  is separate from the green part, which is responsible for explaining the variation caused by the treatment ( $\eta_i(\mathbf{x})$ ). Furthermore, change in the ad load affects only the utility of consuming the ad-supported version.

To train this model, we minimize a weighted negative log-likelihood function similar to Shalit et al. (2017), which jointly optimizes for the treatment effect  $\tau(\cdot)$  and parameters that explain cross-sectional variation  $\theta(\cdot)$  and  $\gamma(\cdot)$ :

$$\underset{\theta, \gamma, \eta, \lambda}{\text{minimize}} \quad \frac{1}{N} \sum_i w_i \mathcal{L}(h(\theta(\mathbf{x}_i), \gamma(\mathbf{x}_i), \eta(\mathbf{x}_i)), y_i; \lambda), \quad (5)$$

---

<sup>12</sup>Note the treatment effect tends to be very small, and if the treatment dummies are inserted in the input layer along with other features, they may get regularized out by the network. The fact that we use a number of shared layers to construct  $\eta_j(\mathbf{x})$  and then spread into separate heads forces the model to use the information provided by treatment dummies and improves the statistical power of our algorithm; see Shalit et al. (2017) for more details.

where  $\mathcal{L}(.,.)$  is the negative log-likelihood for the nested-logit model,  $\theta(.)$ ,  $\gamma(.)$ , and  $\eta(.)$  are functions parameterized by the neural networks. Finally,  $N$  is the total number of users, and  $w_i$  is an inverse propensity score for each treatment condition that is equal to  $\frac{N}{\sum_{n=1}^N \mathbb{1}_{\{\tau_n=e_i\}}}$ . Note that inverse probability weighting is often used to address selection issues; however, our goal is to balance different treatment conditions. For instance, if, by design, the size of one treatment cell is significantly larger than another cell, the optimization problem will have more incentive to fit the data to improve its prediction power for the larger cell. For instance, if a treatment group is twice as large as another treatment group, the same type of error is penalized twice for the larger treatment cell relative to the smaller one. However, our goal is to better learn the *differences* across treatment cells, and this weighting balances the prediction power across different counterfactual scenarios (Shalit et al. 2017).

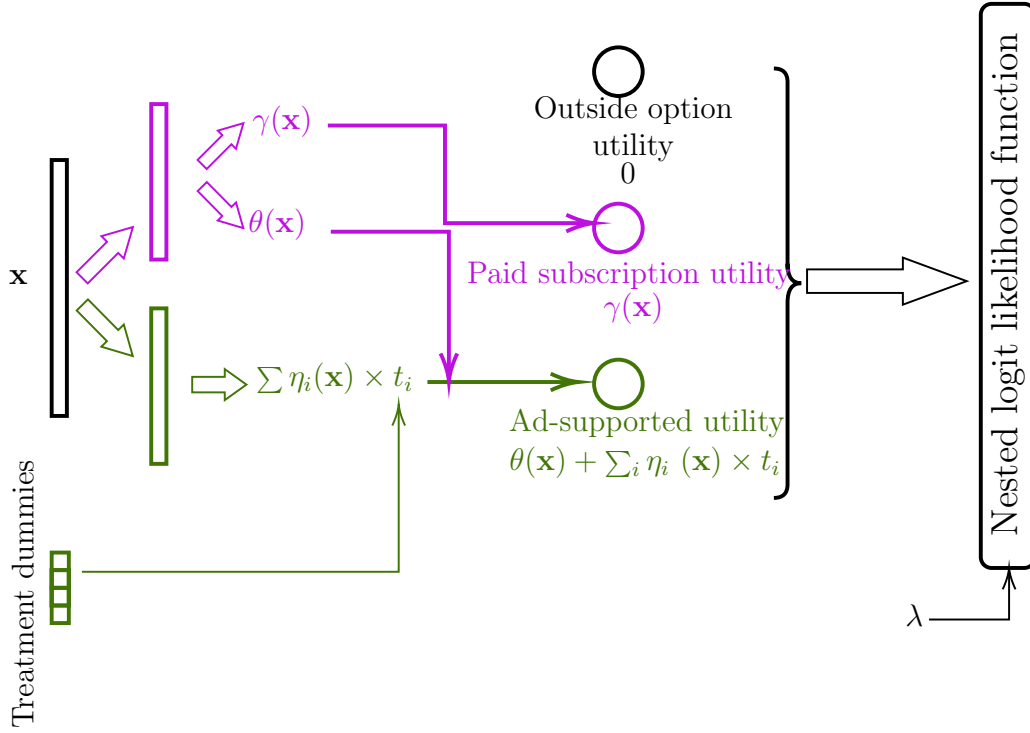


Figure 9: A schematic view of the split neural network architecture used for demand estimation. The purple and green parts of the network are separate. Also note that treatment dummies enter right before the last layer. This restriction imposes structure on the network and forces it to learn the relationship between the output and ad load even though the amount of variation explained by the treatment dummies could be very small. Furthermore, the split structure of the network allows it to separately learn features that could explain the treatment effect  $\eta_i(\mathbf{x})$  from other constructs that explain cross-sectional differences, that is  $\theta(\mathbf{x})$  and  $\gamma(\mathbf{x})$ .

## 5.2 Intensive margin

A change in ad load not only affects the extensive margin of consumption and the choice across different products within the product line, but also affects the intensive margin of consumption, that is hours spent listening to music conditional on being an ad-supported listener. Consider the following model:

$$\log(\mathbf{Y}) = \alpha(\mathbf{x}) + \sum_j \nu_j(\mathbf{x}) \mathbb{1}_{\{\tau=e_j\}} + \epsilon, \quad (6)$$

where  $\mathbf{Y}$  is the number of hours consumed conditional on being an active ad-supported user,  $\alpha(\mathbf{x})$  represents the conditional expectation of  $\log(\mathbf{Y})$  in the control condition, and the  $\nu_i(\mathbf{x})$  aim to capture the conditional average treatment effect on ad-supported hours. Finally,  $\mathbf{x}$  is a set of exogenous and user-generated features collected during the pre-treatment period, and  $\epsilon$  is a random variable with a Normal distribution.

To estimate this model, we again resort to a split neural network model, where one split learns  $\alpha(\mathbf{x})$  and the other one fits  $\nu_i(\mathbf{x})$ . Note  $\alpha(\mathbf{x})$  and  $\nu_i(\mathbf{x})$  are estimated jointly; however, the networks that estimate them do not share weights and are allowed to have different parameters similar to the architecture presented in Figure 9. This model is learned by optimizing a weighted  $\ell_2$  loss counterpart of (5).

### 5.3 Partial control over realized treatment

Apart from the control condition, we have six treatment conditions, namely, 3x1, 4x2, 6x2, 4x3, 5x3, and 6x3. As demonstrated in Figure 6, the realized ad load is not necessarily equal to the intended ad load. In other words, the experiments only shift the ad capacity, which is the number of opportunities to show ads to each listener; however, the realized ad load depends on the advertisers' interest in different demographics. Therefore, the treatment depends on both the firm's actions and the advertisers' demand, and it needs to be accounted for in our optimization problem for reallocating ads. To that end, we estimate a model similar to that in section 5.2, with  $\mathbf{Y}$  being the realized ad load conditional on being an ad-supported listener. The rest of the parameters and the estimation procedure are similar to section 5.2.

## 6 Heterogeneous Treatment Effects

In this section, we demonstrate that the models described in section 5 are able to sort users based on the magnitude of the treatment effect. Before proceeding, we discuss the training



process and introduce some notation. To train the models, we randomly divided the data set into halves. We study user outcomes in the first four weeks of December 2016, that is six months after the treatment began. We train and validate the models on one half of the data, and we use the other for evaluating the performance of the model in terms of improving profits. In this section, we show our models are able to sort users based on the treatment magnitude in the hold-out sample.

Let  $\mathbf{x}$  be the set of pre-treatment outcomes and user features used for describing each individual. Also as before, let  $\boldsymbol{\tau}$  represent each of the seven ad-load conditions. Note that in the experiment, each listener can participate in only one of the treatment cells. The models introduced in section 5 enable us to predict user outcomes across “counterfactual” ad-load conditions. We built three sets of models, and we use the following notation to refer to them:

- **Extensive margin:** Let  $P_0(\mathbf{x}, \boldsymbol{\tau})$ ,  $P_a(\mathbf{x}, \boldsymbol{\tau})$ , and  $P_s(\mathbf{x}, \boldsymbol{\tau})$  denote the conditional probability of choosing the outside option, ad-supported service, and the paid subscription as a function of pre-treatment user features  $\mathbf{x}$  and treatment condition  $\boldsymbol{\tau}$ . These conditional probabilities are the output from the estimated models in section 5.1.
- **Intensive margin:** Let  $C(\mathbf{x}, \boldsymbol{\tau})$  be the expectation of the number of ad-supported hours consumed in a given period conditional on being an ad-supported user.  $C(\mathbf{x}, \boldsymbol{\tau})$  would be the output of the model discussed in section 5.2.
- **Realized treatment:** Let  $A(\mathbf{x}, \boldsymbol{\tau})$  be the conditional expectation of the realized ad load (number of ads per hour).  $A(\mathbf{x}, \boldsymbol{\tau})$  is the output of the machine learning model discussed in section 5.3.

To demonstrate the effectiveness of our approach, we show our models are able to detect heterogeneous treatment effects in terms of change in ad-supported hours and switching to the subscription service. We first calculate the conditional average treatment effect on

ad-supported hours for each condition relative to the control condition as:

$$\zeta(\mathbf{x}_i, e_j) = P_a(\mathbf{x}_i, e_j)C(\mathbf{x}_i, e_j) - P_a(\mathbf{x}_i, e_0)C(\mathbf{x}_i, e_0), \quad (7)$$

where  $P_s(\mathbf{x}_i, e_j)$  and  $C(\mathbf{x}_i, e_j)$  are the predicted probability of user  $i$  being an active ad-supported user when exposed to treatment condition  $j$ , and the number of hours consumed conditional on being an active ad-supported user, respectively.  $P_a(\mathbf{x}_i, e_0)$ , and  $C(\mathbf{x}_i, e_0)$  are defined similarly for the control condition, and  $e_0$  is a vector of all zeros.

The histogram of  $\zeta(\mathbf{x}_i, e_j)$  for each of the six treatment conditions is plotted in Figure 10. As mentioned above, the control condition is basically the same as the 4x2 condition, with the only difference being that the first ad pod within the session is constrained to be of length one. Therefore, the only condition with fewer ads relative to control is 3x1 and the model has correctly predicted a positive effect for the majority of users in that arm. Similarly for the rest of the arms, with higher ad load relative to control, the treatment effect is negative for the majority of users.




Figure 10: The histogram of predicted change in ad-supported hours for each treatment arm relative to control in the hold-out sample. The x-axis is limited to be between -2 to 2 for aesthetics purposes. The median treatment effect across each condition is represented by the dashed line, and the solid red line represents zero.

To demonstrate that the predicted heterogeneity is indeed correlated with the realized change in ad-supported hours, we break users in the hold-out sample into five quintiles based on the predicted treatment effect for the 6x3 condition<sup>13</sup>. Then, we compare the treatment effect of the 6x3 condition relative to control across these quintiles. The results are presented in Figure 11 and show the model is indeed able to sort users based on their ad sensitivity. In particular, the negative effect on ad-supported hours in the first quintile is twice as large as that of the second quintile.

<sup>13</sup>Recall that the 6x3 condition is the most extreme treatment condition in our experiment with the largest treatment effect.

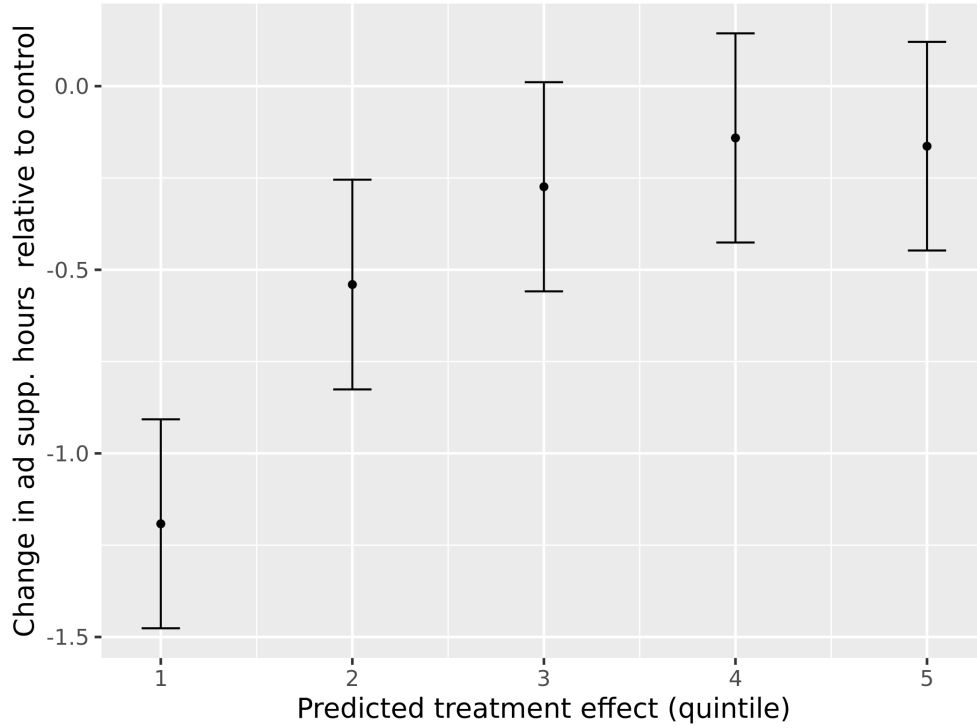


Figure 11: Realized change in ad-supported hours for different predicted treatment quintiles in the hold-out sample.

We repeat this analysis to illustrate the impact of changing ad load on subscription propensity. In particular, we measure the conditional average treatment effect on subscription propensity relative to the control condition as:

$$\xi(\mathbf{x}_i, e_j) = P_s(\mathbf{x}_i, e_j) - P_s(\mathbf{x}_i, e_0),$$

where  $\mathbf{x}_i$ ,  $e_j$ , and  $P_s(.,.)$  are defined as before. The histogram of predicted change in the propensity of subscription across individuals is plotted in Figure 12. As expected, higher ad load tends to increase subscription propensity. We also plot the realized lift in subscription propensity in the 6x3 condition relative to control as a function of the predicted treatment effect quintile in Figure 13. Figure 13 confirms the model was indeed successful in learning the heterogeneity of the effect on subscription propensity. Interestingly, the lift for the lowest

quintile is centered around zero, which means the model has identified customer segments who are very unlikely to subscribe even when facing higher ad load. Note the model used for making these prediction, that is sorting users into five group, was not trained on the hold-out sample, and this prediction power on the hold-out sample demonstrates that the model was able to learn meaningful patterns that generalize beyond the training set.




Figure 12: The histogram of predicted change in subscription propensity for each treatment arm relative to control in the hold-out sample. The x-axis is limited to be between -0.03 to 0.03 for aesthetics purposes. The median treatment effect across each condition is represented by the black dashed line, and the solid red line represents zero.




Figure 13: Realized lift in subscription propensity as a function of predicted treatment effect quintile on the hold-out sample.

## 7 Optimizing profits

To personalize the ad load, one needs to leverage the heterogeneity both in terms of substitution to the subscription service and intensive and extensive margin adjustments to ad-supported hours, which determine the profits from selling ads. To that end, we use the sets of models presented in section 5, namely, a discrete-choice model that reflects the user’s decision between the outside option, ad-supported consumption, and the subscription service, a model that determines the number of hours consumed conditional on using the ad-supported

service, and a model that estimates the realized ad load across different treatment conditions.

Throughout this section, we rely on the notation developed in section 6.

The firm's optimization problem is to maximize profits holding the total ad inventory fixed.<sup>14</sup> If the ad inventory is held fixed, this optimization problem then translates to maximizing profits from subscriptions while serving the same number of ads to the overall user base:

$$\begin{aligned} \underset{\boldsymbol{\tau}}{\text{maximize}} \quad & \sum_i m_s P_s(\mathbf{x}_i, \boldsymbol{\tau}_i) \\ & \sum_i P_a(\mathbf{x}_i, \boldsymbol{\tau}_i) C(\mathbf{x}_i, \boldsymbol{\tau}_i) A(\mathbf{x}_i, \boldsymbol{\tau}_i) = \Gamma, \end{aligned} \tag{8}$$

where  $i$  indexes users, and  $m_s$  is the margin from subscriptions. The objective function is the expected profits from subscriptions across all users.  $\Gamma$  is the total number of ads available in the inventory, and the constraint ensures all the ads in the inventory are served. The problem assigns each user to one of the six treatment conditions or control group to maximize profits while selling  $\Gamma$  ads. In other words,  $\boldsymbol{\tau}_i$  is a binary vector that takes one of the seven values  $\{e_0, e_1, \dots, e_6\}$ , where  $e_0$  is the vector of all zeros, and  $e_i$  is a unit vector whose  $i^{th}$  element is non-zero. This discrete optimization problem aims at assigning each user to one of the seven conditions, that is  $7^N$  different combinations in total.

The functions above are constructed using the estimates of models developed in section 5. For instance, note  $P_s(\mathbf{x}, \boldsymbol{\tau}) = \frac{\exp(\lambda IV)}{1 + \exp(\lambda IV)} \frac{\exp\left(\frac{v_2(\mathbf{x})}{\lambda}\right)}{\exp\left(\frac{v_1(\mathbf{x}, \boldsymbol{\tau})}{\lambda}\right) + \exp\left(\frac{v_2(\mathbf{x})}{\lambda}\right)}$ , where  $v_1(\mathbf{x}, \boldsymbol{\tau})$ ,  $v_2(\mathbf{x})$ , and  $\lambda$  are estimated parameters from the model described in section 5.1. The rest of the functions are also outputs from estimated models presented in section 5. Note (8) is a discrete non-

---

<sup>14</sup>The overall problem is more complex than this. Note that the firm can change the price of ad impressions and that would affect its ad inventory size. Unfortunately, we do not observe the closing price of those contracts. However, even with a fixed ad inventory, the firm faces an implicit pricing problem that involves allocating ads across individuals. We show the gains to personalizing this implicit price across different inventory levels.

convex optimization problem with a non-convex constraint, and even finding a local optimum is NP-Hard for general continuous non-convex problems in the worst case (Murty and Kabadi 1985). In its current form, the problem is intractable. To approach it, we use the Lagrangian relaxation of (8), which yields

$$\underset{\boldsymbol{\tau}}{\text{maximize}} \quad \sum_i m_s P_s(\mathbf{x}_i, \boldsymbol{\tau}_i) + \lambda(\Gamma) \left( \sum_i P_a(\mathbf{x}_i, \boldsymbol{\tau}_i) C(\mathbf{x}_i, \boldsymbol{\tau}_i) A(\mathbf{x}_i, \boldsymbol{\tau}_i) \right), \quad (9)$$

where  $\lambda(\Gamma)$  is the marginal impact of ads on subscriptions. Different values of  $\lambda(\Gamma)$  lead to different ad inventory sizes. Note that given  $\lambda(\Gamma)$ , the problem can now be decoupled across users:

$$\sum_i \underset{\boldsymbol{\tau}_i}{\text{maximize}} \quad \underbrace{m_s P_s(\mathbf{x}_i, \boldsymbol{\tau}_i) + \lambda(\Gamma) (P_a(\mathbf{x}_i, \boldsymbol{\tau}_i) C(\mathbf{x}_i, \boldsymbol{\tau}_i) A(\mathbf{x}_i, \boldsymbol{\tau}_i))}_{f(\mathbf{x}_i, \lambda(\Gamma), \boldsymbol{\tau}_i)}, \quad (10)$$

which is simply equivalent to evaluating  $f(\mathbf{x}_i, \lambda(\Gamma), \boldsymbol{\tau}_i)$  for each user  $i$  across different treatment conditions and choosing the maximum.

Given  $\lambda$ , the complexity of problem (10) is  $\mathcal{O}(N)$  compared to  $\mathcal{O}(7^N)$  for (8). Note the shadow price  $\lambda(\Gamma)$  corresponding to each value of  $\Gamma$  is not a priori known and one needs to resolve problem (8) for different shadow-price values  $\lambda$  to find the corresponding ad inventory level  $\Gamma$ . However, this task can be done easily with binary search techniques.

To evaluate the performance of our personalization policy, we follow a procedure similar to Hitsch and Misra (2018), Yoganarasimhan et al. (2020). However, our problem is more nuanced than Hitsch and Misra (2018), Yoganarasimhan et al. (2020). In particular, the treatment exposure itself not only depends on treatment assignment, but is also a function of user state, that is ad-supported or subscriber, the amount of consumption, and advertisers' demand. Therefore, we construct an inverse probability weighted estimator for realized ad load and profits from the subscription service:



$$\hat{\Pi}(\mathcal{A}) = \frac{1}{N} \sum_{i=1}^N \sum_{j=0}^6 w_j \mathbb{1}_{\{\tau_i=e_j\}} \mathbb{1}_{\{\mathcal{A}(i)=e_j\}} \pi_i, \quad (11)$$

where  $w_j = \frac{N}{\sum_{i=1}^N \mathbb{1}_{\{\tau_i=e_j\}}}$  is the inverse propensity for each treatment condition and is a fixed number for each of the seven treatment conditions because we have a randomized control trial. The randomized treatment assignment for user  $i$  is denoted by  $\tau_i$ , and  $\mathcal{A}$  is an assignment function that assigns each user  $i$  to one of the treatment conditions. Our goal is to evaluate the performance of an assignment rule  $\mathcal{A}$  that is created by solving (10). The product of the two indicator functions in (11) filters out observations for which the assignment rule  $\mathcal{A}$  and the randomized treatment  $\tau_i$  coincide. Finally,  $\pi_i$  is an outcome of interest for user  $i$ , for example subscription status or number of ads received. Equation (11) then provides a consistent estimator of expectation of  $\pi_i$  under a given assignment rule  $\mathcal{A}$ .

To illustrate gains to personalizing ad load, we vary  $\lambda(\Gamma)$  and solve (10) for users in the hold-out sample to get an assignment rule  $\mathcal{A}_\lambda$ . Then, we evaluate the average number of ads realized under the assignment rule  $\mathcal{A}_\lambda$  and the expected profits from the subscription service under this assignment rule using (11). Essentially, for each  $\lambda$ , we get a point on the 2D plane whose x coordinate is the average number of ads realized, and its y coordinate is the expected profits from subscriptions across individuals. We vary the shadow price  $\lambda$  to generate points across the Pareto frontier. The results are plotted in Figure 14. Each purple dot in Figure 14 reflects the performance of an assignment rule for different values of  $\lambda$ , which translates to different levels of ads served on the platform. Each pink dot represents the performance of a personalized ad-allocation strategy. Our agreement with Pandora prevents us from sharing the actual numbers in dollar terms, and the performance has been reported relative to the control condition. Note the personalized counterpart of the control condition, that is the pink dot with the same x coordinate as control, leads to the same number of subscribers as in the 6x2 and 4x3 conditions. Therefore, if the firm's goal was to achieve the same number

of subscribers using a uniform ad-load strategy, it would have to increase its ad load by more than 30%.




Figure 14: Change in subscription profits as a function of number of ads served. Each pink dot represents the performance of a personalized assignment rule. Holding fixed the number of ads served, the personalized assignment strategy dominates uniform ad-load strategies that the firm experimented with.

So far, we have examined the gains from personalization after a six-month period. To find the minimum time required for such gains to materialize, we compare the control condition with its personalized counterpart throughout time. Note (11) provides a consistent estimator of any outcome of interest in any given week. Figure 15 compares the 6x3 and 3x1 conditions, and the personalized counterpart of the control condition relative to the control; that is each outcome of interest is measured relative to control. The results demonstrate that the control and its counterpart lead to similar ad load. However, the counterpart increases the Plus subscription rate by 10%. Whereas the impact on subscriptions manifests within three months after personalizing the ad load, the effect on all and ad-supported hours in the same time period seems to be negligible. This finding indicates a dynamic optimization model

may be beneficial here. However, due to the nature of our experiments, we cannot evaluate the benefits to dynamic implicit pricing using the randomization.

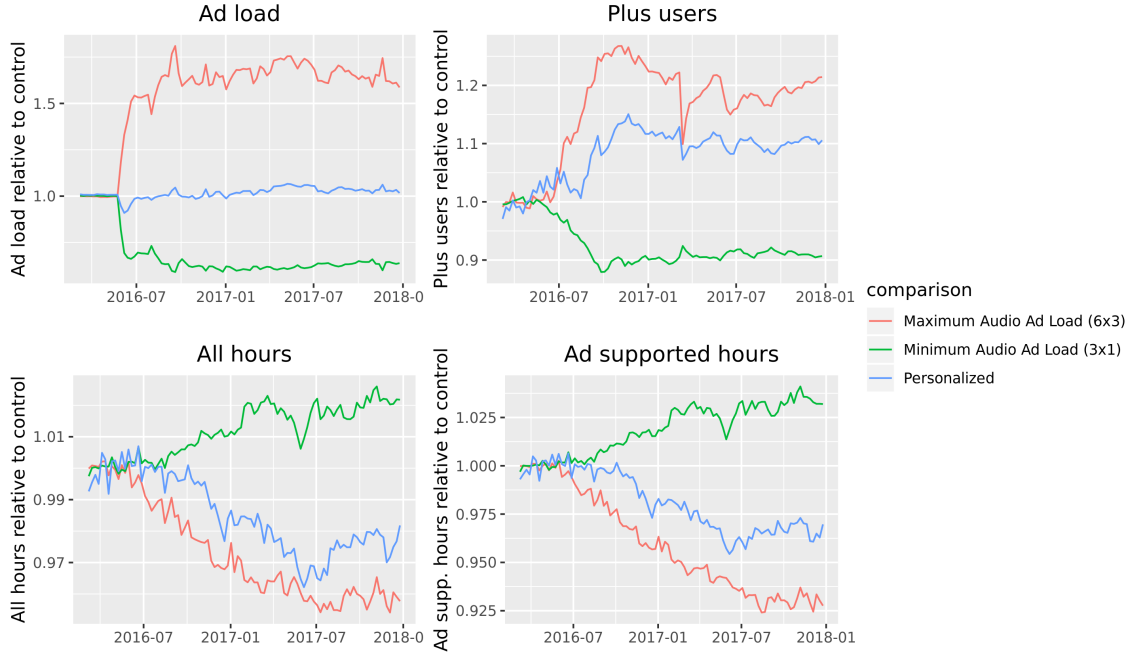


Figure 15: The effect of personalization throughout time. The top-left panel shows the personalized counterpart of the control condition is delivering approximately the same amount of ads as control, whereas the 6x3 condition is delivering 60% more ads relative to control. The top-right figure shows the personalized counterpart increases the number of subscribers by 10%, and this gain is expected to materialize within three months of implementing this policy. Note the impact on ad-supported hours or all hours within three months of implementation seems to be negligible.

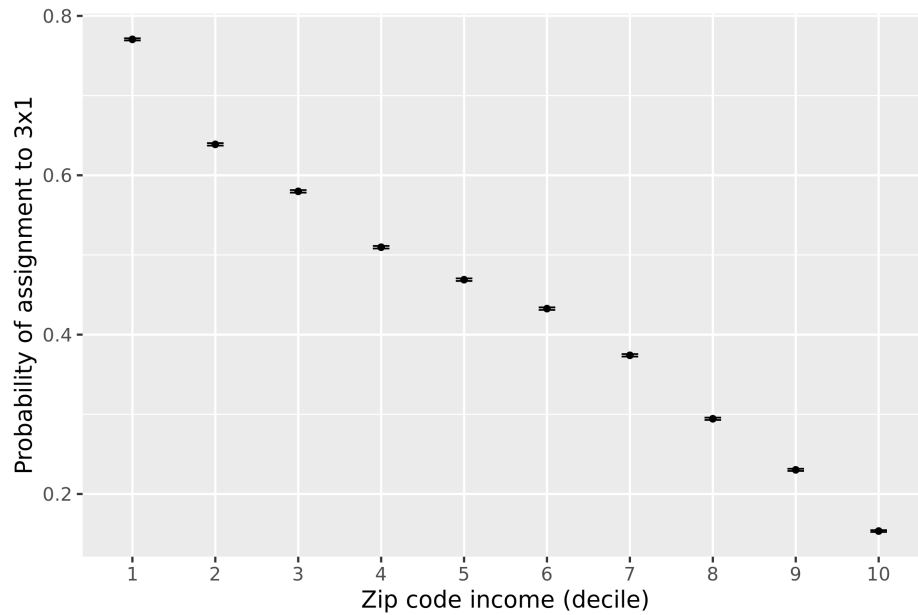
We now illustrate the underlying mechanism that enables the algorithm to improve the firm’s profits. Pandora offers two types of products: the subscription service (high-tier) and the ad-supported (low-tier) product. When the menu of products cannot be personalized, the problem is similar to the one discussed in Mussa and Rosen (1978) and Deneckere and Preston McAfee (1996). In the absence of personalization, the seller has the incentive to lower the quality of the low-tier product for everybody to make adopting the high-tier product worthwhile for those who have higher willingness to pay. However, personalizing ad

load limits the distortion to high willingness to pay customers, and the “implicit price” for other segments falls (quality improves) at the expense of this segment.


## 8 Distribution of Welfare

As we mentioned before, the welfare implications of personalizing ad load are a priori ambiguous due to possible correlation between willingness to pay in time and money units. For instance, whether higher-income users would be assigned to higher or lower-ad-load conditions would be unclear, whereas in a single-product pricing problem, one would expect low-income users to face lower prices because they are more price sensitive. To examine the proposed policy, we compare the allocation of ads in the control condition with its personalized counterpart, that is the personalized assignment rule that serves the same number of ads as the control condition. Note the optimization problem assigns each user to one of the seven conditions, and the only experiment arm that has lower ad load than control is the 3x1 condition. The propensity of being assigned to the lower-ad-load condition, that is higher quality of service for the ad-supported product, tends to be monotonically increasing as a function of zip code income and decreasing as a function of age; see Figure 16. Both younger individuals and those residing in lower-income zip codes tend to be more price sensitive, and if the algorithm were to charge prices in dollar amounts, it would likely charge them a lower price. In our case, the company is following a uniform pricing scheme for the paid service; however, our personalization algorithm induces wealthier individuals to upgrade to the paid service by personalizing the ad load of the ad-supported product, and provides a better quality of service to other demographic groups.

We now illustrate the impact of our ad-allocation algorithm on consumer welfare by comparing the control condition with its personalized counterpart, that is the personalized algorithm that delivers the same overall number of ads. The overall utility for an individual



(a)



(b)

Figure 16: Probability of assignment to lower ad load than control as a function of price sensitivity. Note the total number of ads for this assignment was set to be equal to the control condition, and those assigned to the 3x1 condition are effectively receiving higher quality of service on the ad-supported product. (a) Users from lower-income zip codes tend to be more likely to receive an ad-load “discount.” (b) Older users tend to be less likely to receive a discount. The algorithm seems to be adjusting the quality of service for users who have higher willingness to pay to make converting incentive compatible for them.

with features  $\mathbf{x}$  who is assigned to treatment condition  $\boldsymbol{\tau}$  is equal to:

$$U(\mathbf{x}, \boldsymbol{\tau}) = \log \left( 1 + \left[ \exp\left(\frac{v_1}{\lambda}\right) + \exp\left(\frac{v_2}{\lambda}\right) \right]^\lambda \right), \quad (12)$$

where  $v_1$  and  $v_2$  are the utilities associated with the ad-supported, and paid-subscription products, respectively. Recall that  $v_1$  and  $v_2$  are functions of  $\mathbf{x}$  and  $\boldsymbol{\tau}$  and were defined in (4). To study the impact of our personalization model on consumer surplus, we can compare the percentage change in the utility of users in the control condition relative to its personalized counterpart. In particular, we examine the following construct:

$$\Delta U_{\mathcal{A}} = 100 \times \frac{U(\mathbf{x}, \mathcal{A}(\mathbf{x})) - U(\mathbf{x}, \text{control})}{U(\mathbf{x}, \text{control})},$$

where  $\mathcal{A}(\mathbf{x})$  denotes the assignment rule that is the personalized counterpart of the control condition. The distribution of  $\Delta U_{\mathcal{A}}$  is plotted in Figure 17. On average, consumer utility drops by -2% and utility improves for 41.2% of users.

We now illustrate the impact of this policy on users from different age and income groups in Figure 18. The results demonstrate the loss in consumer utility is more pronounced for older users and those from higher-income zip codes. This observation is consistent with our prior findings in Figure 16 that showed younger users and those from lower-income zip codes tend to be more likely to be assigned to lower-ad-load conditions.

## 9 Discussion and Conclusions

The advent of big data and large-scale data-processing technologies has allowed firms to optimize services, ads, and prices at the individual level. Although a large body of literature has focused on the implications of personalized pricing, the impact of personalizing the product itself is overlooked. The public perceives price discrimination to be unfair and


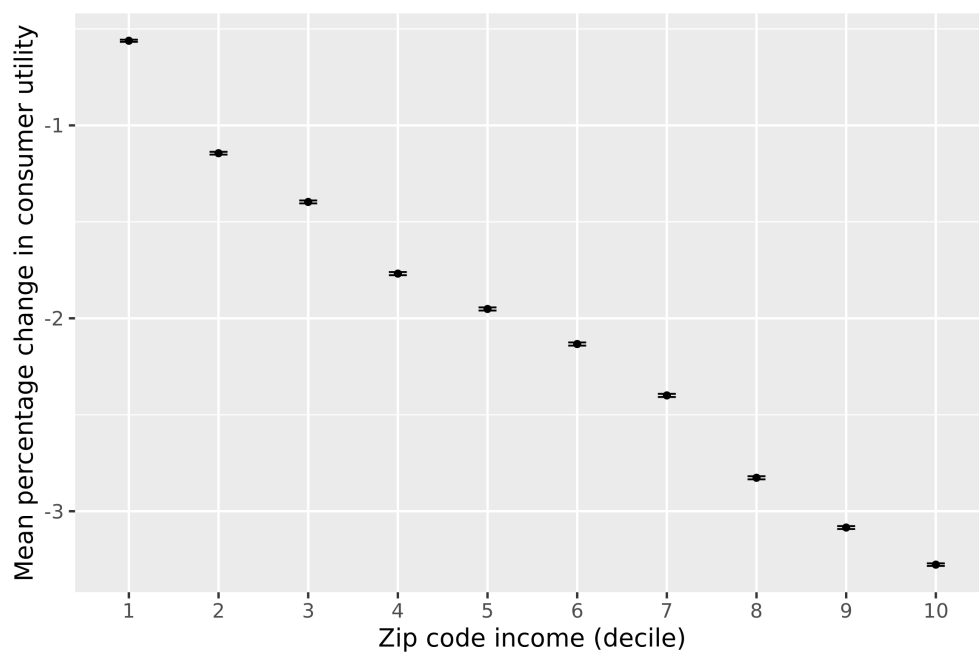


Figure 17: The impact of ad-load personalization on consumer welfare. The figure compares the percentage change in consumer utility across the control condition and its personalized counterpart. On average, personalizing ad load lowers consumer utility by 2%.

companies have largely avoided such practices fearing a consumer backlash. Given these limitations, whether big data is going to be employed for personalizing pricing or versioning instead is still unclear. To the best of our knowledge, this study is the first empirical paper to investigate returns to personalized product versioning.

Although advertising can be used as an instrument to implement versioning for many content providers, including YouTube, Spotify, or Pandora, the idea of personalized versioning applies more broadly to other freemium business models. We provide two such examples here. First, in the online newspaper industry, the number of free pages or the amount of free content available to users can be used for versioning. Second, among cloud storage services, consider Dropbox’s free plan, which offers 2GBs of free storage to users. However, users may be eligible for a wide variety of free storage promotions, including student discounts or offers available to users who purchase HP or Samsung devices (Martinez 2014). We are not sure



(a)



(b)

Figure 18: The average percentage change in consumer utility. The impact of the policy across (a) different income levels and (b) users of different age.



how targeted these strategies are, but the amount of free space offered to users on Dropbox is surely not uniform. We are not aware if companies have used these features to experiment with targeted versioning strategies, but as big data and experiments gain popularity, personalized versioning strategies may become an alternative to personalized pricing.

Our field experiments at Pandora present a unique opportunity to study the impact of personalized versioning in a product line. The availability of large-scale pre-treatment features allows us to segment users and prescribe personalized ad schedules. Our study highlights the importance of conducting field experiments and data-collection efforts for designing reliable prescriptive strategies. We also highlight challenges that take place in causal inference in two-sided platforms including partial control over realized outcomes or treatment exposure. The fact that Pandora has allowed us to share the details of their experiments and analyze the data to evaluate counter-factual strategies<sup>15</sup> is unfortunately an exception in the industry, not a norm. We hope that efforts by firms such as Pandora, Yahoo, eBay, and Ziprecruiter (Lewis and Reiley 2014, Blake et al. 2015, Dubé and Misra 2017) promote transparency of firm-sponsored research.

Our results show that to achieve the same level of subscribers in the absence of a personalized ad-scheduling strategy, the firm needs to increase its ad load by more than 30%. This finding shows personalization can both improve firm profits and the average quality of service. We also find that gains from ad-load personalization materialize quickly. In particular, within three months of implementing the personalized counterpart of the control condition, the profits from subscriptions increase by 10%. Interestingly, the short-term impact of this strategy on the overall consumption of ad-supported service is negligible. This finding, combined with switching costs between products, presents an opportunity for firms to investigate returns to dynamic optimization of implicit prices. Although some evidence

---

<sup>15</sup>The personalized versioning algorithm developed in this paper is not adopted by Pandora and we used inverse probability weighting to evaluate its performance using the randomization in the data.

shows firms change their quality of service in time due to demand seasonality (Lambrecht and Misra 2017), we believe studying the trade-offs between personalized and time-varying quality of service adjustments is a fruitful area for future research. Finally, changing ad load could affect the click-through rate of ads or, in general, their effectiveness. This effect adds an additional layer of complexity for platforms that are compensated based on conversions or click-through rates. Although studying how ad effectiveness changes as a function of the number of ads in online platforms is beyond the scope of this paper, we acknowledge it could play an important role in the firm’s decision to adopt the personalization algorithm discussed here.

---

**References**

- Big data and differential pricing / executive office of the president of the united states, council of economic advisors. 2015.
- Mark Aguiar, Erik Hurst, and Loukas Karabarbounis. Time use during the great recession. *American Economic Review*, 103(5):1664–96, 2013.
- Mark A Aguiar, Erik Hurst, and Loukas Karabarbounis. Time use during recessions. Technical report, National Bureau of Economic Research, 2011.
- Asim Ansari and Carl F Mela. E-customization. *Journal of marketing research*, 40(2):131–145, 2003.
- Gary S Becker and Kevin M Murphy. A simple theory of advertising as a good or bad. *The Quarterly Journal of Economics*, 108(4):941–964, 1993.
- Thomas Blake, Chris Nosko, and Steven Tadelis. Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1):155–174, 2015.
- Raymond Chiang and Chester S Spatt. Imperfect price discrimination and welfare. *The Review of Economic Studies*, 49(2):155–181, 1982.
- Lesley Chiou and Catherine Tucker. Paywalls and the demand for news. *Information Economics and Policy*, 25(2):61–69, 2013.
- Sofronis K Clerides. Book value: intertemporal pricing and quality discrimination in the us market for books. *International Journal of Industrial Organization*, 20(10):1385–1408, 2002.
- CNET. Now showing: random dvd prices on amazon, Jan 2002. URL <https://www.cnet.com/news/now-showing-random-dvd-prices-on-amazon/>.
- Simon Cowan. Third-degree price discrimination and consumer surplus. *The Journal of Industrial Economics*, 60(2):333–345, 2012.
- Gregory S Crawford and Matthew Shum. Monopoly quality degradation and regulation in cable television. *The Journal of Law and Economics*, 50(1):181–219, 2007.

- Gregory S Crawford, Oleksandr Shcherbakov, and Matthew Shum. The welfare effects of endogenous quality choice in cable television markets. 2015.
- Raymond J Deneckere and R Preston McAfee. Damaged goods. *Journal of Economics & Management Strategy*, 5(2):149–174, 1996.
- Jean-Pierre Dubé and Sanjog Misra. Scalable price targeting. Technical report, National Bureau of Economic Research, 2017.
- Andrew Edgecliffe-Johnson. Media wants to break free, May 2009. URL <https://www.ft.com/content/d0960f18-4303-11de-b793-00144feabdc0?fbclid=IwAR06SW3pqxG14RVQVKqjg02Nap-xF9gNlpVcycTuGRstKVdWi6aL1Bb2VHU>.
- Max H Farrell, Tengyuan Liang, and Sanjog Misra. Deep neural networks for estimation and inference. *arXiv preprint arXiv:1809.09953*, 2018.
- Matthew Gentzkow. Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, 97(3):713–744, 2007.
- David Godes, Elie Ofek, and Miklos Sarvary. Content vs. advertising: The impact of competition on media firm strategy. *Marketing Science*, 28(1):20–35, 2009.
- Daniel G Goldstein, Siddharth Suri, R Preston McAfee, Matthew Ekstrand-Abueg, and Fernando Diaz. The economic and cognitive costs of annoying display advertisements. *Journal of Marketing Research*, 51(6):742–752, 2014.
- Daniel Halbheer, Florian Stahl, Oded Koenigsberg, and Donald R Lehmann. Choosing a digital content strategy: How much should be free? *International Journal of Research in Marketing*, 31(2):192–206, 2014.
- John R Hauser, Glen L Urban, Guilherme Liberali, and Michael Braun. Website morphing. *Marketing Science*, 28(2):202–223, 2009.
- Günter J Hitsch and Sanjog Misra. Heterogeneous treatment effects and optimal targeting policy evaluation. *Available at SSRN 3111957*, 2018.
- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.

Jason Huang, David Reiley, and Nick Riabov. Measuring consumer sensitivity to audio advertising:

A field experiment on pandora internet radio. *Available at SSRN 3166676*, 2018.

Tetsuya Kaji, Elena Manresa, and Guillaume Pouliot. An adversarial approach to structural estimation. *arXiv preprint arXiv:2007.06169*, 2020.

Juyong Kim, Yookoon Park, Gunhee Kim, and Sung Ju Hwang. Splitnet: Learning to semantically split deep networks for parameter reduction and model parallelization. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1866–1874. JMLR. org, 2017.

Subodha Kumar and Suresh P Sethi. Dynamic pricing and advertising for web content providers. *European Journal of Operational Research*, 197(3):924–944, 2009.

Anja Lambrecht and Kanishka Misra. Fee or free: When should firms charge for online content? *Management Science*, 63(4):1150–1165, 2017.

Randall A Lewis and David H Reiley. Online ads and offline sales: measuring the effect of retail advertising via a controlled experiment on yahoo! *Quantitative Marketing and Economics*, 12(3):235–266, 2014.

Song Lin. Two-sided price discrimination by media platforms. *Marketing Science*, 39(2):317–338, 2020.

Juan Martinez. Acer and hp to preload dropbox on all laptops, pcs, Oct 2014. URL <https://www.techradar.com/news/internet/cloud-services/acer-and-hp-to-preload-dropbox-on-all-laptops-pcs-1270862>.

Eric Maskin and John Riley. Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2):171–196, 1984.

Brian McManus. Nonlinear pricing in an oligopoly market: The case of specialty coffee. *The RAND Journal of Economics*, 38(2):512–532, 2007.

Katta G Murty and Santosh N Kabadi. Some np-complete problems in quadratic and nonlinear programming. Technical report, 1985.

- Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317, 1978.
- Richard Pérez-Peña and Tim Arango. They pay for cable, music and extra bags. how about news?, Apr 2009. URL <https://www.nytimes.com/2009/04/08/business/media/08pay.html>.
- Arthur Cecil Pigou. The economics of welfare. *New Brunswick & Londres: Transaction Publishers*, 4, 1920.
- Ashutosh Prasad, Vijay Mahajan, and Bart Bronnenberg. Advertising versus pay-per-view in electronic media. *International Journal of Research in Marketing*, 20(1):13–30, 2003.
- Omid Rafieian. Optimizing user engagement through adaptive ad sequencing. Technical report, Technical report, Working paper, 2019.
- Omid Rafieian and Hema Yoganarasimhan. Targeting and privacy in mobile advertising. *Marketing Science*, 40(2):193–218, 2021.
- Peter E Rossi, Robert E McCulloch, and Greg M Allenby. The value of purchase history data in target marketing. *Marketing Science*, 15(4):321–340, 1996.
- Navdeep S Sahni, S Christian Wheeler, and Pradeep Chintagunta. Personalization in email marketing: The role of noninformative advertising content. *Marketing Science*, 37(2):236–258, 2018.
- Steven Salop. The noisy monopolist: Imperfect information, price dispersion and price discrimination. *The Review of Economic Studies*, 44(3):393–406, 1977.
- Susumu Sato. Freemium as optimal menu pricing. *International Journal of Industrial Organization*, 63:480–510, 2019.
- Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3076–3085. JMLR. org, 2017.
- Carl Shapiro, Shapiro Carl, Hal R Varian, et al. *Information rules: a strategic guide to the network economy*. Harvard Business Press, 1998.

- Benjamin Reed Shiller et al. *First degree price discrimination using big data*. Brandeis Univ., Department of Economics, 2013.
- A Michael Spence. Monopoly, quality, and regulation. *The Bell Journal of Economics*, pages 417–429, 1975.
- S Sriram, Pradeep K Chintagunta, and Puneet Manchanda. Service quality variability and termination behavior. *Management Science*, 61(11):2739–2759, 2015.
- Joacim Tåg. Paying to remove advertisements. *Information Economics and Policy*, 21(4):245–252, 2009.
- Jennifer Valentino-DeVries, Jeremy Singer-Vine, and Ashkan Soltani. Websites vary prices, deals based on users’ information, Dec 2012. URL <https://www.wsj.com/articles/SB10001424127887323777204578189391813881534>.
- Hal R Varian. Price discrimination and social welfare. *The American Economic Review*, 75(4):870–875, 1985.
- Frank Verboven. Quality-based price discrimination and tax incidence: evidence from gasoline and diesel cars. *RAND Journal of Economics*, pages 275–297, 2002.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Yanhao’Max’ Wei and Zhenling Jiang. Estimating parameters of structural models using neural networks. *USC Marshall School of Business Research Paper*, 2020.
- Kenneth C Wilbur. A two-sided, empirical model of television advertising and viewing markets. *Marketing science*, 27(3):356–378, 2008.
- Hema Yoganarasimhan. Search personalization using machine learning. *Management Science*, 66(3):1045–1070, 2020.
- Hema Yoganarasimhan, Ebrahim Barzegary, and Abhishek Pani. Design and evaluation of personalized free trials. *Available at SSRN 3616641*, 2020.