Check for updates

Article

AMERICAN MARKETING ASSOCIATION

Journal of Marketing Research
2022, Vol. 59(6) 1101-1119
© American Marketing Association 2022
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/00222437221080163
journals.sagepub.com/home/mrj
SAGE

The Impact of Subscription Programs on Customer Purchases

Raghuram Iyengar, Young-Hoon Park, and Qi Yu

Abstract

Subscription programs have become increasingly popular among a wide variety of retailers and marketplace platforms. Subscription programs give members access to a set of exclusive benefits for a fixed fee up front. In this article, the authors examine the causal effect of a subscription program on customer behavior. To account for self-selection and identify the individual-level treatment effects, they combine a difference-in-differences approach with a generalized random forests procedure that matches each member of the subscription program with comparable nonmembers. The authors find that subscription leads to a large increase in customer purchases. The effect of subscription is economically significant, persistent over time, and heterogeneous across customers. Interestingly, only one-third of the effect on customer purchases is due to the economic benefits of the subscription program, and the remaining two-thirds is attributed to the noneconomic effect. Evidence supports that members experience a sunk cost fallacy due to the up-front payment that subscription programs entail. Finally, the authors illustrate how firms can calculate the profitability of a subscription program and discuss the implications for customer retention and subscription programs.

Keywords

subscription program, retailing, e-commerce, causal inference, machine learning, generalized random forest, sunk cost fallacy

Online supplement: https://doi.org/10.1177/00222437221080163

In an effort to retain and develop customers, retailers and marketplace platforms are increasingly turning to subscription programs, which are designed to keep customers engaged by giving access to exclusive benefits for a fee up front.[1] For example, Amazon Prime offers members free shipping, audio content, and video content, as well as exclusive member discounts, for an up-front payment of $119 per year. Many other retailers, such as Barnes & Noble, Sephora, and Alibaba, have similar programs with benefits that range from unlimited free shipping to member-only discounts and additional loyalty points.

As the relevance and popularity of subscription programs grows, it is of managerial interest to examine their effect on customer purchase behavior. However, there are two lacunae in the research on subscription programs. First, while anecdotal evidence suggests that subscription programs are successful, the empirical evidence of their causal impact on customer purchases is limited. According to an industry study, Amazon Prime members spend twice the average nonmember's annual spending (Reisinger 2017). However, because members self-select into the program, it is unclear whether the reported difference in spending between members and nonmembers reflects the causal impact of the program or the differences in their preferences prior to joining the program. Second, it is usually challenging to identify the underlying drivers of the change in customer purchases because multiple mechanisms may be at work.

[1] A subscription program is a paid membership program initiated by an existing noncontractual business and provides members exclusive benefits beyond those available to regular customers (i.e., nonmembers). We distinguish subscription programs from stand-alone subscription services that provide subscribers new items or personalized experiences periodically (e.g., Stitch Fix, Birchbox) and media streaming subscriptions that allow consumers unlimited access to a catalog of content (e.g., Spotify). For recent research on stand-alone subscription services and media streaming subscriptions, see McCarthy, Fader, and Hardie (2017) and Datta, Knox, and Bronnenberg (2017), respectively.

Raghuram Iyengar is Miers-Busch W'1885 Professor and Professor of Marketing, Department of Marketing, The Wharton School, University of Pennsylvania, USA (email: riyengar@wharton.upenn.edu). Young-Hoon Park is Sung-Whan Suh Professor of Management and Professor of Marketing, Samuel Curtis Johnson Graduate School of Management, Cornell University, USA (email: yp34@cornell.edu). Qi Yu is Assistant Professor of Marketing, Lee Kong Chian School of Business, Singapore Management University, Singapore (email: qiyu@smu.edu.sg).

A subscription program may have two types of effects on customer purchases. The first effect, which we term the economic effect of the program, captures the changes in purchase behavior attributable to the tangible benefits that reduce the effective price of the products (e.g., exclusive member discounts). Accordingly, the noneconomic effect of a program includes any remaining effects on the demand that are likely due to psychological drivers. For instance, a subscription program may lead customers to feel enhanced status or to form a new consumption habit. The economic and noneconomic effects of the subscription program have different implications on customer acquisition and retention, as well as the firm's bottom line. The economic effect can be enhanced by altering the tangible benefits of the program. However, the economic effect may not endure once the benefits are removed or matched by competing firms. In contrast, the noneconomic effect is likely to be applicable in the absence of the tangible benefits and can be managed by altering the noneconomic aspects of the program. Therefore, identifying two types of effects on customer purchases is substantively important in evaluating the profitability of subscription programs and improving their design. As an extreme scenario, if additional sales are generated only by reducing (effective) prices, the program might even negatively affect a firm's performance in the long term.[2]

The purpose of this article is to take a first step toward assessing the causal impact of customers' joining a subscription program on their purchase behavior. We hereinafter refer to the impact as the total treatment effect of subscription. We also aim to decompose the total treatment effect into the economic effect due to program benefits and the noneconomic effect that cannot be explained by the tangible benefits of the program. Specifically, we are interested in addressing the following questions: Does a subscription program generate value for a firm? Is the subscription program effective in inducing customers to change their behavior because of the economic benefits and/or the psychological drivers? How do these effects vary over time and across customers? What are the underlying drivers of any documented effects? We address these questions in close collaboration with a company that launched a subscription program. The program offers members a few exclusive benefits for an up-front fee. Our data contain individual-level transactions before and after the launch of the program and other information on various components of the program, thus allowing us to examine its effect on purchase behavior.

A key concern in the estimation of the impact of a subscription program on purchase behavior arises from the lack of random assignment. We exploit the panel structure of our data and rich information on customer characteristics and rely on a quasi-experimental design to control for self-selection and identify the effect at the individual level. Specifically, our baseline model uses a difference-in-differences (DD) specification that controls for unobserved individual and time fixed effects (Angrist and Pischke 2008). In addition, to enhance the comparability between members and nonmembers, we create a weighted set of neighboring observations for each member based on a large set of observed characteristics following a generalized random forests (GRF) procedure (Athey, Tibshirani, and Wager 2019) and estimate the DD model using the weighted sample to obtain the treatment effect on customer behavior. The combination of the DD approach and the GRF procedure is robust to selection bias based on observed as well as time-invariant unobserved characteristics. It also provides individual-level estimates of the treatment effect. Within this framework, we also quantify the noneconomic effect by evaluating the remaining effect of the program after controlling for the economic benefits members were exposed to. The individual-level estimates of economic and noneconomic effects allow us to get a richer understanding of the impact of the subscription program and its underlying drivers.

We find that subscription is effective in lifting sales. On average, members increased their purchases by about $27 per month over a 12-month period after subscribing, which is more than double the customer monthly purchase amount prior to subscription. The treatment effect of subscription is economically significant, persistent over time, and heterogeneous across customers. The subscription keeps members more engaged in terms of frequency and variety of their purchases. Interestingly, only one-third of the treatment effect on purchase amount is due to the economic benefits of the program, and the remaining two-thirds is attributed to the noneconomic effect. Our main findings are robust to potential confounding effects of self-selection and unobservables, different treated groups, and different outcomes of purchase behavior.

We investigate the possible explanations for our findings. We find that, in addition to the psychological underpinnings documented in loyalty programs (e.g., affect, habit, status), a unique feature of subscription programs helps lift sales: as customers pay a fee up front in exchange for future benefits, they experience sunk cost fallacy (e.g., Arkes and Blumer 1985; Thaler 1980). Although the up-front fee has already been incurred and should no longer be relevant to future decisions, it poses a psychological cost for the members so that they are incentivized to change their behavior to amortize the psychological cost. We provide evidence supporting this mechanism. Further, we illustrate how firms can calculate the profitability of the program.

Our article is related to several streams of research. It makes both substantive and theoretical contributions to the literature on subscription programs as well as the closely related literature on loyalty programs. While the empirical research on subscription programs is limited, there is extensive research on loyalty programs, which have long been used across a wide array of industries to reward repeat purchases. The key difference between these two types of programs lies in the up-front

[2] For instance, MoviePass, a subscription service that offered its members one free movie each day for $9.95 per month, managed to attract more than 2 million subscribers but failed to build a deeper relationship with customers. The company reported a loss of $266 million in 2018 and ended the service in September 2019.

payment. Some studies find that loyalty programs can increase customer lifetime value and share of wallet (e.g., Gopalakrishnan et al. 2021; Kopalle et al. 2012; Lal and Bell 2003; Liu 2007). Others find no or weak evidence that loyalty programs are effective (e.g., Hartmann and Viard 2008). Several researchers have documented that loyalty programs can induce positive affect (e.g., Leenheer et al. 2007), lead to the development of habitual consumption (e.g., Wood and Neal 2009), and enhance members' perception of status (e.g., Drèze and Nunes 2009). To the best of our knowledge, our research is the first to study the causal impact of subscription programs on customer purchases and their underlying drivers. Substantively, we provide empirical evidence of the causal effect of subscription programs on customer behavior using quasi-experimental data. We demonstrate a persistent increase in both purchase amount and purchase variety. Theoretically, we demonstrate that the effect of a subscription program can go beyond the economic benefits it offers. We document a novel mechanism that subscription programs may create a sunk cost and encourage purchases through the up-front fee. As the up-front fee is a feature common to subscription programs, we believe our findings have broad implications for the design of such programs.

Our findings suggest that customers behave in a boundedly rational manner, which adds to empirical evidence for such behavior found in the lab and in other field settings. The sunk cost fallacy has implications in a variety of contexts, and extensive evidence of it has been found in lab studies (Thaler 1980). Relatively few studies, however, provide evidence for the sunk cost fallacy in the field. Arkes and Blumer (1985) conduct a field experiment in theaters and find that the attendance rate is positively correlated with the ticket price. Ho, Png, and Reza (2017) find evidence for the sunk cost fallacy in the Singapore automobile market by demonstrating that driving time is positively correlated with the price of the license. More recently, Goli, Chintagunta, and Sriram (2021) find a sunk cost effect in massive open online courses (MOOCs) such that paying users demonstrate a higher level of motivation to complete the courses. We add to this literature by showing that subscription programs also induce the sunk cost fallacy and contribute to our understanding of consumer behavior in the field using observational data.

The remainder of the article is organized as follows. We first describe the empirical context and the data. Next, we present the empirical strategy and analyses, followed by the findings and possible explanations. We then offer several robustness checks and discuss the profitability of the subscription program. We conclude with a discussion of implications for the design of subscription programs and areas for future research.

Empirical Context and Data

We obtained the data for our empirical analysis from a retailer in Asia. The retailer sells a wide range of products in the personal care category (e.g., skin care). It has significant brick-and-mortar and online presence. In December 2015, a subscription program was launched on its e-commerce website. The launch of the program and its benefits were communicated through mass emails and on the website, and no targeting was involved. The program provided both unlimited use of a service and access to exclusive member offers on the e-commerce website for a subscription fee of $50 per year.[3] Specifically, upon joining the program, members were provided with a $50 gift card that could be used with no restrictions. Members also received a $3 gift card per month for purchases during the month. Several products were occasionally coupled with exclusive member discounts. Free samples were offered to members monthly with a purchase. Finally, members also had access to unlimited free shipping with no minimum purchase requirements. Similar to other subscription programs, these benefits were offered beyond the first year of subscription rollout. Other than the program benefits mentioned here, there was no difference in the communication between members and nonmembers during the data period.

Our data include 10,811 customers who joined the subscription program between December 2015 and February 2017. On average, 720 customers joined the program per month. The monthly number of members has a moderate level of variation, ranging from 342 to 1,062, with a standard deviation of 210. For the purpose of comparison, we also obtained a random sample of 13,768 customers who had yet to subscribe the program as of July 2017. The data consist of two parts: transaction data and program usage data. The transaction data contain detailed information on each purchase made by a customer, including when a customer purchased a product and how much the customer paid for it. The program usage data contain information on how a member benefited from the program, including amounts spent with gift cards and free samples received. Our data also contain sociodemographic characteristics of customers, including age, gender, and address.

We conduct the analysis at the customer-month level. Using transaction data, we define a set of outcome measures associated with customer purchases. Because the program was offered through the online channel only, the measures are based on online purchases unless specified otherwise. As our main interest is to assess how effective the program is in lifting sales, our primary measure is the amount spent by a customer per month. In addition, we consider two other (monthly) measures of customer purchases: number of purchases made (purchase frequency) and basket size (in dollars) conditional on purchase. Although these two measures positively correlate with the amount spent, they may have differing implications for the firm in terms of engagement and costs. We also characterize the variety in purchase behavior with a few metrics. We classify a product (and its category) that a customer purchased as a new

[3] All transactions were recorded in the currency of the country in which the company's headquarters was located. We converted purchase amounts to U.S. dollars using the average exchange rate over the data period.

versus known product (and category) on the basis of whether the customer had purchased it in the presubscription period. The first set of metrics relates to the variety at the product level: amount spent for new versus known products. The second set of metrics relates to the variety at the category level: amount spent for new versus known categories.[4] As a proxy for engagement with the firm, these measures are useful to investigate how customers change their behavior after beginning a subscription.

Method

In this section, we first give an overview of the empirical strategy. We then discuss the samples to establish the causal effect of subscription program and the implementation details of our empirical model.

Empirical Strategy

A key challenge in identifying the impact of subscription on purchase behavior is due to self-selection: members may differ from nonmembers even before they subscribe to the program, and this difference may lead to biased results if we estimate the effect by directly comparing purchases between members and nonmembers. We also aim to identify the economic and noneconomic effects of subscription on customer behavior separately. The former capture the changes in purchase behavior attributable to the tangible benefits of the program (e.g., reduced prices due to exclusive member discounts), and the latter include any remaining effect on demand. Finally, we are interested in examining the heterogeneity in the treatment effect across customers.

To control for self-selection and identify the effect of subscription on purchase behavior at the individual level, we rely on a quasi-experimental design. Our baseline model uses a DD approach (Angrist and Pischke 2008) and controls for selection based on time-invariant unobservables. Within the regression framework, we separately identify the noneconomic effect on purchase behavior by evaluating the remaining effect of the program after controlling for the economic benefits members were exposed to. Conceptually, our framework is similar to previous work that has examined how the different components of a pricing scheme may have an impact on demand over and above its economic effects (e.g., Bertini and Wathieu 2008; Iyengar et al. 2011). We complete our modeling framework by embedding the DD specification within a GRF procedure (Athey, Tibshirani, and Wager 2019). Briefly, we estimate the DD specification for each member using a subsample of comparable customers defined by a random forest in a high-dimensional covariate space. In doing so, we account for selection based on observables and heterogeneity across customers in a nonparametric manner and obtain individual-level treatment effects.

Treated and Control Groups

We focus on a cohort of members who joined the program around the same time in our main analysis. Such a cohort-level analysis is common when analyzing customer value (e.g., McCarthy, Fader, and Hardie 2017). Focusing on a cohort of members is conducive to examining the effect of the subscription program, as it gives well-defined pretreatment and post-treatment periods for the analysis. Our main findings consider the cohort of 721 members who joined the program in April 2016, four months after the launch of the program.[5] As reported in the discussion of our robustness checks, our findings are robust in that the subscription program has largely similar effects across several cohorts.

Before we establish the causal effect of the program on customer purchases, we examine the purchase amount for members and nonmembers over a 24-month period from April 2015 to March 2017. The first 12 months (April 2015 to March 2016) are before the members' subscription. Figure 1 offers model-free evidence that purchase behavior differed considerably between members and nonmembers, and the difference persisted over time. On average, members spent \$43.16 per month during their subscription, whereas nonmembers spent only \$3.93.

To assess whether nonmembers were similar to members before joining the program, we compare their purchases during the 12-month period before subscription and their individual characteristics. Table 1 shows that members and nonmembers differed significantly on both their purchases and their demographics. On average, members spent more per month than nonmembers (diff. = 5.61, $p < .001$), which is consistent with the intuition that customers who spent more were more likely to join the program as they could benefit more from the program. Members were older than nonmembers (diff. = 3.09, $p < .001$). Clearly, the results would be biased if we estimated the effect of subscription by comparing customer purchases between the two groups.

Difference-in-Differences

The descriptive analysis reveals that members purchased considerably more than nonmembers during their subscription. However, this analysis may suffer from self-selection. In this section, we employ a DD approach that controls for time-invariant unobserved variables to quantify the treatment effect

[4] From conversations with the retail partner, we decided to use five product categories in our empirical analysis to correspond with the key metrics that the firm monitors regarding customer purchases: skin care, makeup, hair care, bath and body care, and other purchases, in which we aggregated and grouped fragrances and the rest of the categories (e.g., tools, brushes, accessories).

[5] To mitigate concerns of selection and unmeasured confounders, we deliberately excluded early subscribers because they may systematically differ from other customers (e.g., Rogers 2003).

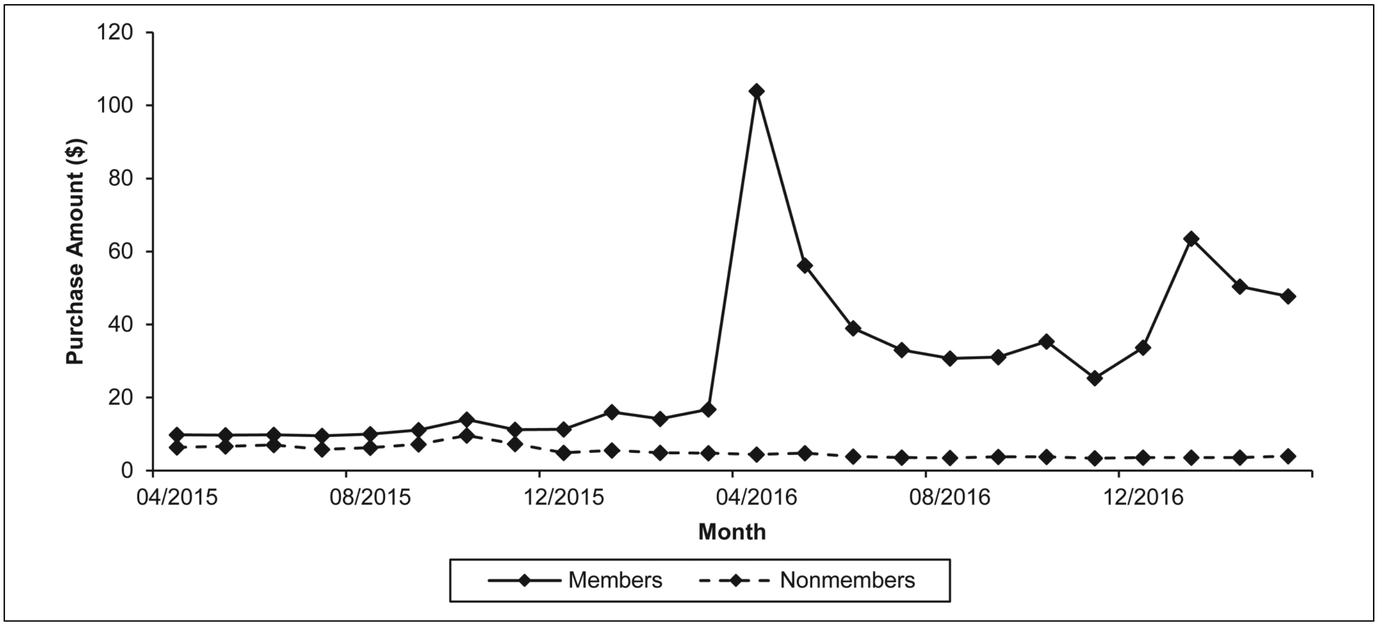

Figure 1. Customer purchases of members versus nonmembers.

of the subscription program. We estimate the following DD model:

$$Y_{it} = \tau Member_{it} + \alpha_i + \gamma_t + \varepsilon_{it}, \quad (1)$$

where Y_{it} is the outcome measure of customer i in month t and $Member_{it}$ is an indicator variable for the treatment that equals 1 if customer i was a member in month t and 0 otherwise. The parameters α_i and γ_t are customer and month fixed effects, respectively, and ϵ_{it} is the error term. By including the two-way fixed effects, we control for time-invariant customer characteristics as well as common time trends and month-to-month fluctuations. The parameter of interest is τ, which captures the total treatment effect (sum of economic and noneconomic effects) on purchase behavior. Our panel consists of data over a 24-month period, where the first 12 months are prior to subscription. We excluded customer purchases in the month of adoption (April 2016) to avoid any potential simultaneity bias with the adoption itself.

The total treatment effect of the subscription program can be attributed to the economic benefits and the psychological drivers. In our context, members received unlimited free shipping service and exclusive member offers. We find that free shipping and product samples had limited impact.[6] However, members benefited from price discounts and gift cards: on average, members had additional discounts of about 6% as compared with nonmembers, and about 20% of monthly gift cards were redeemed. To disentangle the economic and noneconomic effects, we extend Equation 1 to account for the price discounts and monthly gift cards members were exposed to. We conceptualize that customers decide on their purchase behavior given the price and their monthly budget, which may be altered by price discounts and monthly gift cards, respectively. We specify a log-linear model in which the customer's purchase is linear in the logarithm of the price and the logarithm of the customer's monthly budget. The log-linear model has been used widely in the marketing literature to characterize consumption patterns (e.g., Dubé, Hitsch, and Rossi 2018) because it is consistent with utility maximization (e.g., Sato 1972).[7] Specifically,

$$\begin{aligned} Y_{it} = {} & \tau Member_{it} + \beta_1 \log(Price_{it}) \\ & + \beta_2 \log(Baseline_i + Giftcard_{it}) + \alpha_i + \gamma_t + \varepsilon_{it}, \end{aligned} \quad (2)$$

where $Price_{it}$ is the store-level price for customer i in month t, $Baseline_i$ is the baseline budget of customer i, and $Giftcard_{it}$ is the gift card amount that customer i received in month t. The parameters β_1 and β_2 are the semi-elasticities of price and gift card amounts, respectively, and the parameter τ captures the noneconomic effect.

We compute the store-level price for customer i in month t ($Price_{it}$) as the weighted average price for a basket of products (e.g., Dubé, Hitsch, and Rossi 2018). Given that we do not observe the prices of all products that a specific customer was exposed to, we use prices at the group level (members and nonmembers) as a proxy for customer-level prices. This assumption is reasonable, as there was no other source of price discrimination between members and nonmembers beyond the discounts from

[6] The firm offered customers, regardless of subscription, free shipping on orders above a certain threshold, which was satisfied by most orders. In the "Noneconomic Effect" section, we provide a detailed discussion of the effect of free shipping. We also examined whether a customer purchased a product after receiving a free sample of that product, and we find that less than 1% of purchases were induced by free samples. Nevertheless, we excluded the purchases induced by free samples in estimating the noneconomic effect of the program.

[7] Web Appendix A discusses conditions under which our demand model is consistent with utility maximization.

Table 1. Summary Statistics of Members Versus Nonmembers.

Variable	Members	Nonmembers	Difference	*p*-Value
Purchase amount ($)				
Apr. 2015	9.86	6.36	3.50	.000
May 2015	9.77	6.68	3.09	.000
Jun. 2015	9.81	7.03	2.78	.000
Jul. 2015	9.54	5.84	3.70	.000
Aug. 2015	10.05	6.33	3.72	.000
Sep. 2015	11.17	7.25	3.92	.000
Oct. 2015	14.02	9.63	4.40	.000
Nov. 2015	11.18	7.27	3.90	.000
Dec. 2015	11.33	4.90	6.43	.000
Jan. 2016	16.08	5.50	10.57	.000
Feb. 2016	14.21	4.93	9.27	.000
Mar. 2016	16.81	4.77	12.05	.000
Demographics				
Age	35.94	32.84	3.09	.000
Gender (1 if female)	.93	.94	−.01	.162
Observations	721	13,768		

subscription. Specifically, we operationalize the monthly store-level price as the weighted average price of a basket consisting of all products purchased by both members and nonmembers in that month.[8] In addition, we compute the baseline budget of a customer (without gift cards) as the customer's maximum monthly spending one year prior to subscription.[9]

The identification of Equation 2 exploits the panel structure of our data and heterogeneity in the baseline budget. We use the month-to-month variation in store-level price to identify the price coefficient (β_1). Note that all members received a gift card worth $3 each month. The parameter β_2 is identified by the difference of the change in purchases upon subscription across members who had different baseline budgets. For example, a positive β_2 indicates that the marginal effect of the gift card is lower for members with larger budgets. The parameter τ is identified by the remaining change in member purchases upon subscription and provides an estimate for the noneconomic effect of the program.[10]

Generalized Random Forests

While the DD model controls for time-invariant heterogeneity and estimates the average treatment effect, two issues are worth discussing. First, Figure 1 and Table 1 show that the difference between the two groups' purchase amounts widens over time, suggesting that compared with an average nonmember's monthly spending, the monthly spending by an average member increased over time even before subscription. Thus, the parallel time trend assumption may not hold, and the validity of the DD estimator is questionable (e.g., Bertrand, Duflo, and Mullainathan 2004).[11] Second, a typical way to accommodate heterogeneity in the treatment effect is to interact the treatment dummy with individual characteristics. This approach can become cumbersome as the covariate space increases and moderates the treatment effect in a nonlinear manner.

Recent developments in the machine learning literature allow us to address both issues in a principled manner. We employ the GRF method (Athey, Tibshirani, and Wager 2019). Similar to other methods for causal inference using observational data, such as kernel matching (e.g., Hastie, Tibshirani, and Friedman 2009), propensity score matching (e.g., Hirano, Imbens, and Ridder 2003), and synthetic control (e.g., Abadie, Diamond, and Hainmueller 2010), the key idea of GRF is to define for each member a weighted set of neighbors with similar covariates and then fit the model of interest using these neighbors. As an improvement over the traditional methods where weights are chosen by deterministic kernel functions (kernel matching), parametric models (propensity score matching), or trend matching (synthetic control), Athey, Tibshirani, and Wager (2019) propose learning the weights using a revised random forest algorithm that is designed to minimize the estimation error.

Given the forest, we can define for each member with covariates x a weighted set of neighboring customers by locating which customers fall into the leaves that contain the same covariates, weighted by the associated frequency. The treatment effect for this member is estimated by fitting the DD specification on the weighted set. The identifying assumption is that the treatment status is independent of the unobservables conditional on all the observables (Rosenbaum and Rubin 1983). As compared with other commonly used methods for matching, GRF is

[8] Web Appendix B describes how store-level price is operationalized for the main analysis and discusses the robustness of our results under alternative operationalizations.

[9] Our results are robust under alternative operationalizations of the baseline budget. See Web Appendix C for a discussion.

[10] Web Appendix D shows how different patterns in purchase data could help identify model parameters.

[11] Web Appendix E presents the results of empirical tests for the parallel time trend assumption.

nonparametric and robust to model misspecification. The tree structure and the ensemble of many trees naturally account for complex interactions among covariates. The adaptive nature of trees can substantially increase the accuracy of the weighting function with a large space of covariates. Another advantage of GRF is that it uncovers the point estimates and confidence interval of the treatment effect at the individual level with formal asymptotic guarantees. These estimates allow us to explore heterogeneous treatment effects in a systematic manner and can sharpen our understanding of underlying drivers of the program's success.

Implementation

For the outcome variables, we analyze transaction data over a 24-month period (April 2015 to March 2017) because we are interested in examining the long-term effect of subscription. Of these, the first 12 months are prior to subscription. As noted previously, we exclude customer purchases in the month of subscription adoption (April 2016).[12]

To control for potential confoundedness, we include three sets of covariates that describe members and nonmembers in the pretreatment period. The first set of covariates relates to the customer–firm relationship, which would be associated with the adoption of a service, namely, tenure, breadth, and depth (e.g., Bolton, Lemon, and Verhoef 2004; Prins and Verhoef 2007). We calculate tenure on the basis of elapsed time since the customer created an account on the website. We measure breadth by the number of unique categories purchased and depth by the number of transactions made. We also include the average basket size. In addition, we include monthly purchase amount per category during the 12-month presubscription period, instead of the total amount across product categories, because it could help find clusters of customers with similar purchase patterns across categories. We also include the standard deviation of monthly purchase amount to capture the tendency to place small orders and thus customer response to the unlimited free shipping service.

The second set of covariates relates to customers' purchase characteristics (e.g., Baumgartner 2002). We create summary statistics to capture different aspects of purchase behavior: exploratory (purchases of new products), repetitive (repeat purchases of products), and promotional (price discount rate received for purchases).[13]

The third set of covariates relates to sociodemographic characteristics of customers. We include age and gender. We also include the coordinates of a customer's mailing address because it can help control for other unobserved sociodemographics that affect subscription, such as education, income, and lifestyle. Table 2 summarizes the covariates and describes how the variables are operationalized. Altogether, we use 72 covariates to build the trees for the random forest.

We briefly describe the procedure for building the forest.[14] We grow each decision tree in the forest by iteratively partitioning the data into subgroups (leaves). In each iteration, the algorithm chooses a covariate and the cutoff to find partitions where treatment effects most differ. As a single tree is likely to overfit the data, an ensemble of trees is generated. Each tree is constructed using a random subsample, and a random set of the covariates can be potentially used to form the splits. The number of trees, minimum number of treatment and control observations in each leaf, subsample size, and size of the set of covariates used to build each tree are the hyperparameters of a forest, which we choose by cross-validation. In the context of panel data, we also account for clusters at the individual level in the sampling and estimation processes. To ensure that the final estimates are credible and robust, we restrict the sample to the range of common support of propensity scores (Imbens and Rubin 2015). We excluded 50 members and 2,023 nonmembers on the basis of propensity scores estimated by a regression forest. Our findings are based on a final sample of 671 members and 11,745 nonmembers.

The forest performs well in balancing members and nonmembers. Following Imbens and Rubin (2015), we use the normalized absolute mean difference to assess the degree of balance of the observables. Figure 2 shows the normalized mean difference of the variables before and after the adjustments by GRF. After the adjustment, the normalized absolute mean differences of the variables between the members and the matched nonmembers are mostly below .1, and the members and nonmembers are indistinguishable in terms of their observed characteristics.[15]

Finally, given the forest, we define for each member a weighted set of neighboring customers, which may include both nonmembers and other members. For a member with covariates x, the neighbors are the customers appearing in the leaves that contain x, weighted by the frequency of the appearance.[16] The procedure then fits Equations 1 and 2 with the weighted set of observations for each member. To obtain the treatment effect across members, we construct a doubly robust average treatment-effect estimator by augmenting the naive plug-in estimator with a residual-based correction. The doubly robust estimator combines results from GRF and a regression-based prediction such that it is robust to misspecification of either the matching model or the regression model (Chernozhukov et al. 2018).[17]

[12] The estimate we report provides a conservative estimate of the effect of subscription on customer purchases. Detailed results that include the month of subscription adoption are available from the authors on request.

[13] The results are largely similar with alternative operationalization of the variables, such as the share of repeat product purchases per month and category, and with alternative thresholds to construct the variables.

[14] Web Appendix F provides the implementation details.

[15] Appendix A shows that members and matched nonmembers have statistically comparable pretreatment time trends.

[16] We used out-of-bag predictions to avoid overfitting; that is, only trees that do not include the member during tree building were used to produce the weights.

[17] Web Appendix G describes how the individual-level estimates are identified.

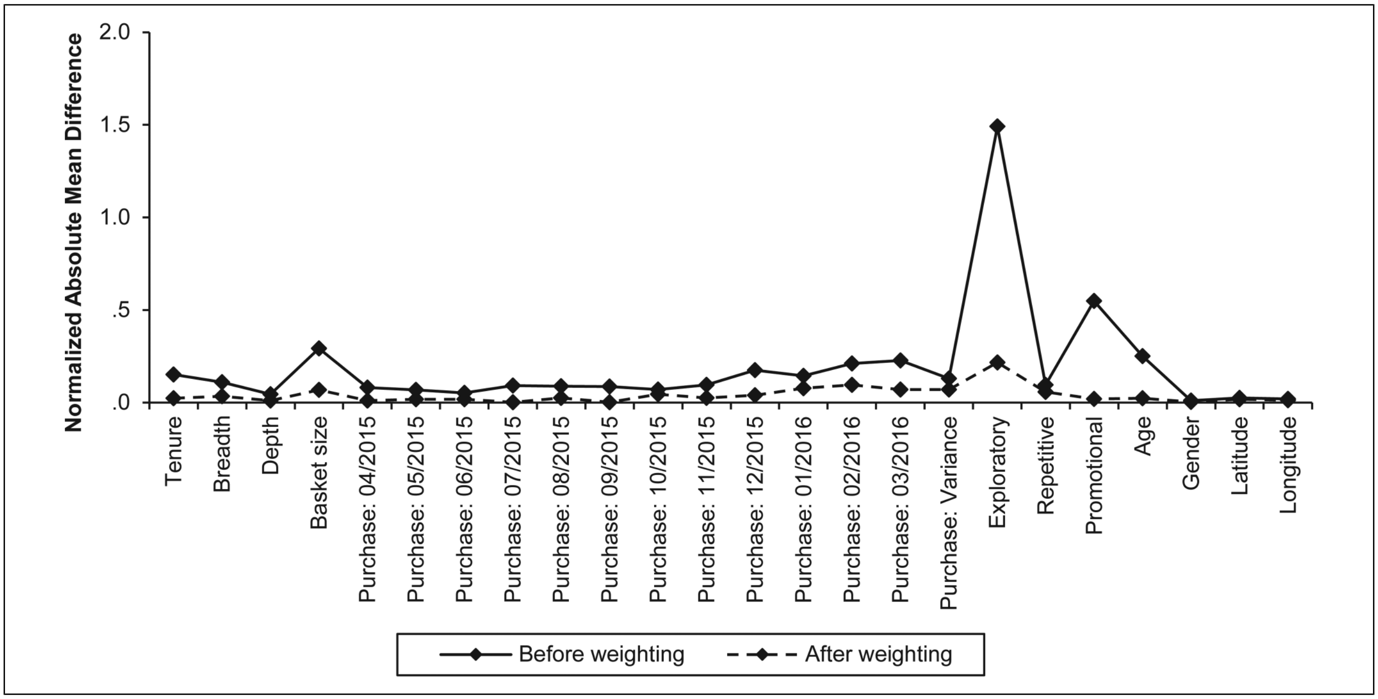

Figure 2. Covariate balance before and after weighting.
Notes: For the purpose of illustration, we plot the normalized absolute mean difference of the monthly purchase amount along with other observed characteristics. In the matching procedure, however, monthly purchase amounts from all product categories were used.

Findings

In this section, we discuss the main findings on the total treatment and noneconomic effects of subscription on customer purchases. We also discuss how these effects vary over time and across customers. Finally, we explore some possible explanations underlying the effect.

Total Treatment Effects

Table 3 reports the average treatment effects of subscription on customer purchases.[18] Because our objective is to identify the effect of subscription for the long term, we first discuss the treatment effect of the program over the whole subscription period, the 12 months after the subscription began. The first four columns in Table 3 report the average total treatment effect of subscription on customer purchases. As shown in the first column, on average, the purchase amount per month among members increased by $27.45. The effect is economically significant, as the purchase amount per month was about $12 prior to subscription.

The treatment effect on the purchase amount is striking. As shown in Figure 1, data patterns suggest that our finding is not an artifact. Although the empirical research on subscription programs is limited, the literature on loyalty programs documents that their impacts have substantial variation, ranging from no effect to an increase as high as 150% in customer response across various empirical settings.[19] For instance, using an application of a "buy ten get one free" program offered by a golf course, Hartmann and Viard (2008) find no changes in customer response. Similarly, Lewis (2004) evaluates a loyalty program from an online merchant and finds about a 2% increase in customer revenue. Using data from a men's hair-salon chain, Gopalakrishnan et al. (2021) find that the introduction of a nontiered loyalty program increases customer value by 19% over a five-year horizon. Kopalle et al. (2012) use data from the loyalty program of a hotel chain and find about a 30% increase in customer spending due to the program. Using data from a convenience store chain's loyalty program, Liu (2007) finds that consumers whose initial patronage levels were low or moderate considerably increase their spending by around 150% under the loyalty program.

Whereas Hartmann and Viard (2008), Kopalle et al. (2012), and Gopalakrishnan et al. (2021) study customer response to loyalty programs in a single category (e.g., hotel), Liu (2007) examines the impact of a loyalty program on customer behavior in a firm offering multiple categories. Our empirical context is close to Liu's because the focal firm offers a wide assortment of brands and products across categories. The effect size in our

[18] Appendix B reports the results in which outcome measures are transformed with natural log in the estimation. The results are qualitatively similar.

[19] Lal and Bell (2003) examine the impact of frequent shopper programs in grocery retailing and find increases of $98, $141, and $150 across three different segments. As they do not report the baseline prior to the frequent shopper program in their study, however, we are not able to compute the relative impact of the program.

Table 2. Covariates for GRF.

Covariates	Operationalization	Mean	SD
Customer–firm relationship			
Tenure	Elapsed time (years) since the customer created an online account	5.55	3.19
Breadth	Number of unique categories purchased	1.97	1.23
Depth	Number of transactions made	1.75	1.27
Basket size	Average basket size ($)	33.51	31.26
Monthly purchase amount: category level	Monthly spending ($) in each product category	1.33	5.91
Monthly purchase amount: SD	Standard deviation of monthly purchase amount ($)	15.07	20.77
Purchase characteristics			
Exploratory	Inverse of average time (years) taken for the purchase among three new products purchased since the launch of the product	3.12	13.39
Repetitive	1 if a customer repeatedly purchased a product more than 10 times, 0 otherwise	.03	.18
Promotional	Average price discount rate received for purchases	.31	.23
Sociodemographic characteristics			
Age	Age	33.96	8.96
Gender	1 if female, 0 otherwise	.94	.23
Address	Coordinates of home address		

Table 3. Total Treatment and Noneconomic Effects.

	Total Treatment Effect				Noneconomic Effect			
Variable	All Months	Month 2	Months 3–4	Months 5+	All Months	Month 2	Months 3–4	Months 5+
Purchase amount ($)	27.45***	40.57***	28.95***	25.53***	17.91***	37.28***	22.26***	12.31***
	(1.75)	(4.70)	(2.33)	(2.03)	(1.96)	(9.51)	(2.70)	(1.99)
Purchase frequency	1.15***	1.37***	1.29***	1.10***	.97***	1.20***	1.24***	.80***
	(.05)	(.02)	(.03)	(.01)	(.04)	(.11)	(.07)	(.05)
Basket size ($)	−5.14**	−9.44*	−5.01***	−5.92***	−6.19	−2.88	−4.63*	−6.67
	(2.47)	(5.23)	(1.76)	(2.02)	(12.77)	(8.89)	(2.80)	(10.47)
Variety ($)								
Purchase amount in known categories	26.19***	40.62***	27.22***	23.61***	17.44***	35.74***	21.77***	12.36***
	(1.77)	(4.83)	(2.30)	(2.07)	(1.96)	(7.22)	(3.00)	(2.46)
Purchase amount in new categories	1.99***	1.00***	2.64***	1.77***	1.40***	1.06	1.44**	1.27***
	(.29)	(.32)	(.57)	(.27)	(.23)	(1.24)	(.58)	(.16)
Purchase amount of known products	7.13***	6.02**	7.65***	7.18***	3.26**	7.20**	1.06	3.99***
	(.68)	(3.05)	(.84)	(.73)	(1.29)	(3.29)	(1.93)	(1.31)
Purchase amount of new products	20.89***	35.74***	22.22***	17.86***	14.23***	28.59***	20.47***	10.56***
	(1.67)	(3.26)	(2.19)	(1.93)	(1.42)	(6.02)	(2.19)	(2.02)

*$p < .1$.
**$p < .05$.
***$p < .01$.
Notes: Robust standard errors appear in parentheses. Customer purchases in the month of subscription adoption (Month 1) were excluded from the analysis.

study is considerably larger than those documented in the literature on (free) loyalty programs. Our finding is an important addition to the literature because our study examines customer response in a subscription program where a one-time purchase of a subscription can lead to recurring sales. Subscribers are more likely to make repeat purchases during their subscription. In contrast, existing literature examines customer response in noncontractual loyalty programs.

The increase in customer purchases could be driven by the increase in purchase frequency and/or basket size. As shown in Table 3, members made about one additional purchase per month (1.15) after beginning the subscription. Interestingly, we find a small but significant decrease in basket size (decrease of $5.14, or 15% of the average basket size prior to subscription). While this reduction in basket size may appear to be due to free shipping benefits, our results on the noneconomic effects of subscription rule out this explanation. However, it is possible that members might make their baskets into smaller ones to utilize, for example, recurring monthly gift cards.

We next examine the effect of subscription on the variety in purchase behavior. Recall that we classify the products (categories) that a customer purchased into new versus known products (categories) on the basis of prior purchase behavior. We find a significant increase across all variety measures.[20] At the category level, more than 95% of the increase in purchase amount ($26.19 out of $27.45) came from known categories. At the product level, approximately 75% of the increase in purchase amount ($20.89 out of $27.45) was from new products that a customer had never purchased. Taken together, our evidence supports that subscription makes members purchase more frequently, with a greater variety of products and categories, leading to increased customer loyalty and share of wallet.

We also examine the temporal variation in the total treatment effect on purchase behavior. For instance, the program could create an initial excitement among members, leading to increased purchasing. If the novelty effect of the program were the only underlying reason for the behavioral change, the effect would likely fade over time, and the program would have limited impact on the firm's long-term revenue (e.g., Galak and Redden 2018). To examine the temporal effect of subscription on purchase behavior, we utilize the purchase patterns shown in Figure 1 and distinguish between the treatment effects within the first two months (excluding the first month of adoption), the next two months, and the remaining months in the subscription period. We estimate the temporal effects by applying GRF on data in the corresponding time periods relative to the 12-month pretreatment period. We find that the effect of subscription on customer purchases is the largest (an increase of $40.57) within the first two months and persisted (an increase of $25.53 per month) after four months of subscription.

In summary, we find a causal impact of subscription on customer purchases. The effect is economically and managerially significant and is persistent over time. The program keeps customers more engaged both in terms of frequency and variety in their purchases.

Noneconomic Effect

We next discuss the noneconomic effect of subscription, controlling for the impact of the tangible benefits of the program. The last four columns of Table 3 report the average noneconomic effect of the program. Interestingly, we find that about two-thirds of the treatment effect on purchase amount ($17.91 out of $27.45) is due to the noneconomic effect of the subscription program. By calculating the relative contribution of economic benefits (price discounts and gift cards), we find that price discounts accounted for about $4.20 of the increase in purchases and that gift cards explained about $5.30 of the increase in purchases.[21]

Looking at the temporal patterns of the noneconomic effect, we find that the increase in purchase amount is largest (an increase of $37.28) within the first two months of subscription. Although it gradually faded over time, the effect is persistent and managerially important. Across most of the other metrics as well, interestingly, a significant part of the increase in customer purchases is not accounted for by the economic benefits of the program. The noneconomic effect on basket size is not significant, indicating that after accounting for the economic benefits, we find no further impact on the basket size. In summary, even after controlling for the economic benefits of the program, we can attribute a significant part of the effect on customer purchases to the noneconomic effect of the subscription program.

Before discussing the heterogeneity in the treatment effect, we investigate whether the increase in purchasing can be attributed to the free shipping benefit associated with the subscription program. Previous research has found that free shipping can be effective in increasing order incidence and purchase amounts, with the effect size varying by the savings in the shipping fee (e.g., Lewis 2006; Lewis, Singh, and Fay 2006). In our context, the firm offered free shipping to all customers (regardless of subscription) on orders above a certain threshold, which was satisfied by a large majority of orders. Thus, we expect that the free shipping benefit associated with the subscription program likely did not play an important role in the increase in purchasing.

We present a few pieces of evidence corroborating our conjecture. First, although members' average basket size decreased slightly, this change becomes insignificant after we account for other economic benefits of the program.[22] Second, we estimate a lower bound of the noneconomic effect sans the effect of free shipping and compare it with the estimates in Table 3. In particular, the effect of free shipping on the purchase amount should be bounded above by the change in purchases from the orders below the free shipping threshold.[23] We can therefore obtain

[20] The treatment effect on purchase amount is slightly different from the sum of the treatment effects on purchase amount in the known and new products and categories. Because our GRF procedure minimizes the mean squared error, matches depend on both covariates and the outcome variable. Therefore, matches can differ slightly when the effects on different outcome variables are evaluated. However, unconfoundedness guarantees that all the estimates are unbiased.

[21] We checked the face validity of the results by comparing our price and coupon elasticities to those documented in previous research. We find that the price elasticity (percentage change in purchase amount in response to a 1% chang,e in price) is −1.75 and the marginal effect of a coupon (dollar change in purchase amount in response to a $1 change in coupon) is 1.76. As a point of comparison, in two meta-analyses, Tellis (1988) and Bijmolt, Van Heerde, and Pieters (2005) report average price elasticities of −1.76 and −2.62, respectively. Venkatesan and Farris (2012) report the marginal effect of coupons to be greater than 2, as the mere exposure to coupons can help increase sales.

[22] We also replicated this finding when we examined the changes in the number of orders below the free shipping threshold. We find members had significantly more orders below the free shipping threshold after subscription. However, after taking price discounts and gift cards into account, the effect becomes insignificant.

[23] We considered the two ways in which free shipping may affect the purchase amount. First, members would likely place more orders below the free shipping threshold, whereas these orders would have been abandoned prior to subscription. Second, members would no longer make the effort to increase their basket size to meet the minimum threshold. With these two factors taken together, the effect of free shipping on the purchase amount should be bounded above by the increase in purchasing from the orders below the free shipping threshold.

a more conservative estimate of the noneconomic effect by excluding all purchases from orders below the threshold and revisit the analysis described in the "Method" section. We find that after excluding the orders below the threshold, members' purchase amount increased by \$16.29 after controlling for other economic benefits. This estimate is close to \$17.91, which is the magnitude of the noneconomic effect reported in Table 3. In summary, the free shipping benefit offered through the subscription program is not likely the driver of the observed effects in our context.

Heterogeneous Treatment Effects

Panels A and B of Figure 3 show the distribution of the total treatment and noneconomic effects of the subscription program on purchase amount, respectively. Both panels of the figure show that the impact of subscription on customer purchases is heterogeneous across customers. These results illustrate the benefits of employing GRF, in that we can obtain individual-level treatment effects.

Panel A of Figure 3 shows significant variation in the total treatment effect across members, ranging from less than \$10 to around \$80. About 20% of the members increased their purchases by \$15 or less, and approximately 27% increased their purchases by \$40 or more. Panel B of Figure 3 illustrates large heterogeneity in the noneconomic effect as well. About 40% of the members increased their purchases by \$15 or less, and around 14% increased by \$40 or more. The proportion of the treatment effect explained by the noneconomic effect also has substantial heterogeneity across members. On average, 65% of the increase in purchase amount was driven by the noneconomic benefit of the program. Among 15% of members, almost 90% of the increase in purchases could be attributed to the economic benefits of the program. In contrast, 40% of members would change their behavior even without economic benefits, in that more than 80% of the increase in purchases was due to the noneconomic effect of the program. These results show that paid subscription programs can have an impact on purchase behavior even after accounting for their economic benefits (e.g., Bolton, Kannan, and Bramlett 2000).

Possible Explanations

We discuss the likely underlying mechanism(s) for the increase in purchase. We consider four psychological drivers motivated by the institutional details of subscription programs and the closely related literature on loyalty programs. In what follows, we first examine each mechanism in isolation and then confirm these insights in a unified framework.

Sunk cost fallacy. In typical subscription programs, members pay an up-front subscription fee. This unique feature distinguishes subscription programs from free loyalty programs. While rational customers should not take the up-front fee into account when making subsequent purchase decisions, studies suggest that people exhibit sunk cost fallacy, a tendency to hold the initial (sunk) cost in a mental account and change their behavior to amortize the psychological burden of the cost (e.g., Thaler 1980). In the context of subscription programs, we posit that the up-front fee of the program creates a sunk cost and induces customers to increase their purchases so that they can utilize the program benefits and recover the initial payment. Following this logic, we hypothesize that as customers benefit more from the program, the sunk cost is amortized and its effect wears off.

To test this hypothesis, we first obtain the estimate of the noneconomic effect of subscription ($\hat{\tau}_{im}$) at the individual-month level. We specify the following DD model that allows for lags of the treatment effect (Angrist and Pischke 2008) and embed it within GRF:

$$Y_{it} = \sum_{m=1}^{11} \tau_m \text{Member}_{i,t-m} + \beta_1 \log(\text{Price}_{it}) + \beta_2 \log(\text{Baseline}_i + \text{Giftcard}_{it}) + \alpha_i + \gamma_t + \varepsilon_{it}, \quad (3)$$

where $\text{Member}_{i,t-m}$ equals 1 if customer i joined the subscription program in month $t-m$ and 0 otherwise. That is, $\text{Member}_{i,t-m}$ is an indicator variable for m months since subscription. For each member i, we estimate Equation 3 on the weighted sample defined by GRF to obtain $\hat{\tau}_{im}$, which captures the noneconomic effect in month m.

To test whether the noneconomic effect is moderated by the usage of program benefits, we estimate the following fixed-effect model:

$$\hat{\tau}_{im} = \gamma \times X_{im} + \delta_i + \xi_{im}, \quad (4)$$

where X_{im} is the cumulative program benefits utilized by member i prior to month m. We also include individual-level fixed effects to account for time-invariant heterogeneity across members. As the dependent variable is estimated with error, we calculate heteroskedastic-robust standard errors clustered at the individual level to account for the variance of the dependent variables (e.g., Hanushek 1974; Liang and Zeger 1986). As reported in Column 1 of Table 4, we find that, consistent with the sunk cost fallacy, as customers utilized more program benefits and recovered their sunk cost, the noneconomic effect of the program declined.[24]

By construction, the dependent and independent variables in Equation 4 may display systematic time trends; the noneconomic effect may likely decline while the cumulative benefits monotonically increase over time. These time trends may result in a spurious correlation between the noneconomic effect and the cumulative benefits. We detrend the dependent and independent variables through first differencing (e.g., Hamilton 2020) and estimate the following

[24] We also investigated the presence of sunk cost fallacy by examining whether members decelerated their purchases after using the \$50 gift cards they received upon subscription. We do not find a decrease in customer purchases. Because most customers used their \$50 gift cards in the first month after joining, there was limited variation to detect an effect.

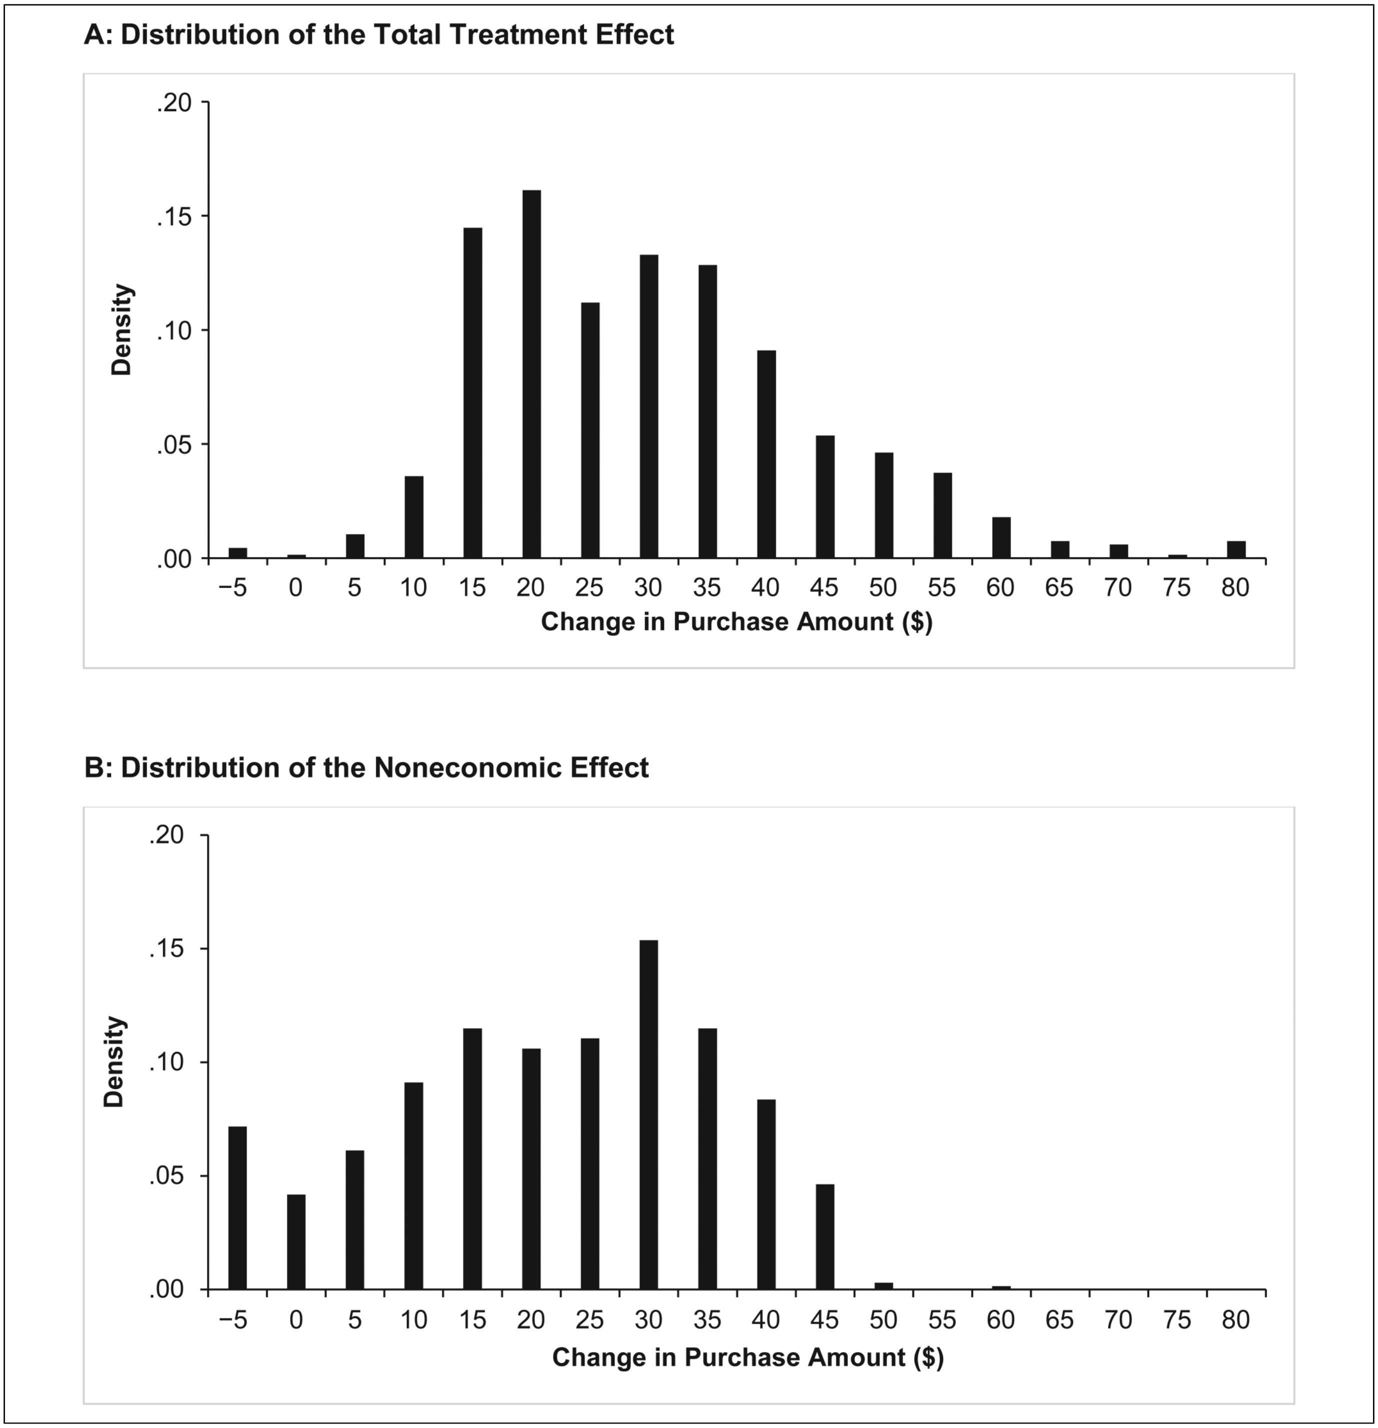

Figure 3. Distribution of the total treatment and noneconomic effects.

regression:

$$\hat{\tau}_{im} - \hat{\tau}_{i,m-1} = \gamma \times (X_{im} - X_{i,m-1}) + e_{im}. \quad (5)$$

In essence, the first difference model tests whether the use of more benefits in month $m - 1$ leads to a larger decline in the noneconomic effect in month m. As shown in the second column in Table 4, we find that consistent with the sunk cost fallacy, the increase in purchases is smaller when customers utilized more program benefits.

Affect. The literature on loyalty programs suggests that a membership program can induce a positive affect toward the firm and lead to increased purchases (e.g., Leenheer et al. 2007). Our prediction based on positive affect is that customers act

Table 4. Results of Possible Explanations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variable	**Linear**	**First Diff.**	**Linear**	**First Diff.**	**Linear**	**First Diff.**	**First Diff.**
Cumulative benefits ($)	−.13***	−.23***					−.18*
	(.03)	(.02)					(.10)
Cumulative purchases ($)			−.02***	−.03***			−.009
			(.004)	(.003)			(.01)
Number of members					−.003***	−.0002	.0004**
					(.0001)	(.0001)	(.0002)
Observations	6,710	6,710	6,710	6,710	6,710	6,710	6,710
R-squared	.08	.08	.09	.08	.07	.008	.08

*$p < .1$.
**$p < .05$.
***$p < .01$.
Notes: Robust standard errors clustered at the individual level appear in parentheses.

favorably toward the program shortly after they become a member but experience hedonic decline as they continue to purchase from the firm (e.g., Galak and Redden 2018). Thus, we would expect the impact of the program to decrease with past purchase amount. To test this prediction, we estimate Equations 4 and 5 using the cumulative purchase amount as the explanatory variable. Similar to our previous analyses, we estimate the fixed-effect model and account for the common time trends of the noneconomic effect and the cumulative variable through first differencing. Both results, as reported in Columns 3 and 4 of Table 4, show that the noneconomic effect is negatively correlated with past purchases, suggesting the presence of hedonic decline with past purchases.

Habit formation. A subscription program may encourage customers to purchase upon joining, and this increase in the short run may lead to a habitual increase in the long run (e.g., Wood and Neal 2009). This mechanism would predict that a high level of customer purchases in the long term is a result of state dependence (i.e., habits formed) based on the increase in purchases in the short term. However, as we find that the noneconomic effect is negatively correlated with past purchases, we conclude that habit formation is unlikely to be the driver of the purchase patterns.

Status. Research on loyalty programs shows that members may feel superior to other customers when they have access to exclusive offers, and their enhanced status can encourage purchases as well (e.g., Drèze and Nunes 2009). This literature suggests that the value of status created by membership is associated with its distinctiveness (e.g., Grier and Deshpandé 2001). If members derive status from the program, the noneconomic effect of the program should be smaller as the number of members increases. To test this hypothesis, we again estimate Equations 4 and 5, using the total number of members as the explanatory variable. Columns 5 and 6 of Table 4 report the results from these models. While results from the linear model indicate that the noneconomic effect of the program decreases with the number of members, the correlation becomes insignificant when first differencing is done. We conclude that status effect is weak in the context of subscription programs.[25]

So far, we have examined the four mechanisms in isolation. We now explore possible co-existence of multiple mechanisms using a unified regression framework because we have estimated the noneconomic effect at the customer-month level. The last column of Table 4 reports the results from the first-difference regression model where all three moderators are included in the vector X_{im}. We find that the insights from the previous analyses continue to hold. The purchase pattern is consistent with the hypothesis that members increased their purchases because of positive affect toward the firm. After accounting for all the other patterns, we also find evidence that customers exhibited a sunk cost fallacy.[26] This finding is in line with previous research that documents consumer response to sunk costs in various contexts. For instance, Arkes and Blumer (1985) show that people who pay full price for theater season tickets attend more shows than those who receive discounts. Ho, Png, and Reza (2017) find that the price of license plates in the Singapore automobile market has an influence on driving time. Goli, Chintagunta, and Sriram (2021) find in MOOCs that paying users are more likely to complete a course than free users are.

Robustness Checks

In this section, we analyze the robustness of our findings by addressing potential confounding effects of self-selection and unobservables, different treated groups, and different outcomes of purchase behavior.[27]

[25] To tease out the status effect, we also examine the noneconomic effects for different cohorts. We do not find a systematic pattern. The noneconomic effects across cohorts are reported in Table 5.

[26] Web Appendix H shows that there is enough variation between program benefits and purchases to identify sunk cost fallacy.

[27] The total treatment and noneconomic effects on all outcome measures show similar temporal patterns as the ones in the main analysis. As we are primarily interested in the long-term effect on purchase amount, we only report the effect on purchase amount for all months after subscription. Other results are available from the authors on request.

Table 5. Total Treatment and Noneconomic Effects by Cohort.

	Total Treatment Effects				Noneconomic Effects			
Cohort	**All Months**	**Month 2**	**Months 3–4**	**Months 5+**	**All Months**	**Month 2**	**Months 3–4**	**Months 5+**
Feb. 2016	26.53***	35.36***	31.31***	25.49***	20.29***	33.93***	29.22***	17.22***
	(2.02)	(2.46)	(5.52)	(2.14)	(2.05)	(1.29)	(4.05)	(1.59)
Mar. 2016	25.68***	43.05***	26.58***	28.29***	17.67***	32.87***	25.69***	15.03***
	(2.42)	(4.47)	(2.88)	(2.51)	(2.14)	(1.38)	(2.91)	(1.75)
Apr. 2016	27.45***	40.57***	28.95***	25.53***	17.91***	37.28***	22.26***	12.31***
	(1.75)	(4.70)	(2.33)	(2.03)	(1.96)	(9.51)	(2.70)	(1.99)
May 2016	26.52***	35.47***	25.60***	26.04***	17.23***	32.24***	15.02***	15.28***
	(1.68)	(2.76)	(1.76)	(1.69)	(1.91)	(1.12)	(2.42)	(1.64)
Jun. 2016	33.52***	37.89***	35.05***	32.84***	23.50***	36.94***	30.64***	19.02***
	(2.00)	(3.15)	(3.13)	(2.03)	(2.14)	(3.21)	(2.71)	(1.86)
Jul. 2016	26.82***	33.03***	27.74***	26.39***	19.51***	29.03***	21.87***	16.44***
	(1.39)	(2.34)	(1.91)	(1.36)	(2.03)	(4.18)	(2.60)	(1.80)
Aug. 2016	35.80***	37.60***	32.34***	33.97***	24.99***	34.22***	28.77***	22.02***
	(3.95)	(4.72)	(4.03)	(4.33)	(2.93)	(1.35)	(3.10)	(2.69)
Average	29.35***	38.77***	29.61***	28.29***	20.09***	34.63***	24.67***	17.19***
	(2.24)	(3.64)	(3.05)	(2.87)	(2.35)	(1.26)	(2.98)	(2.21)

*$p < .1$.
**$p < .05$.
***$p < .01$.
Notes: Robust standard errors appear in parentheses. Customer purchases in the month of adoption (Month 1) were excluded from the analysis.

Selection on Unobservables

The GRF framework has advantages over other causal inference methods, such as propensity score matching and nearest neighbor matching, in that it matches members and nonmembers in a nonparametric and robust manner. Because the treatment is not assigned randomly, the validity of the method still hinges on the assumption of unconfoundedness; that is, the treatment status is not correlated with unobservables.

While the unconfoundedness assumption is usually not directly testable, we present an additional piece of evidence to alleviate concern for this assumption. Specifically, we use late adopters, rather than nonmembers, as controls for early adopters. The late adopters could display a closer resemblance to the early adopters than nonmembers if their adoption time is close enough (e.g., Datta, Knox, and Bronnenberg 2017; Goldfarb and Tucker 2011; Manchanda, Packard, and Pattabhiramaiah 2015; Narang and Shankar 2019). We choose customers who joined the program between August 2016 and November 2016 as the control group, allowing us to have enough customers for matching and enough time periods to estimate the effect. We find qualitatively similar results. On average, customers increased their monthly purchase amount by \$31.38 (SE = 2.18). After accounting for the economic benefits, their monthly purchase amount increased by \$18.56 (SE = 1.92).

Alternate Treated Groups

In our main analysis, we used a single cohort of the members who joined the subscription program in April 2016. As a robustness check, in Table 5, we replicate the analysis for members who joined the program during other months. We also measure the average total treatment and noneconomic effects across cohorts by estimating the DD model in Equation 1 on a matched sample combining all the (weighted) samples from the cohort-specific analyses. The last row of Table 5 reports the average treatment and noneconomic effects of all cohorts. Our results suggest that the effects of subscription on purchase amount across several cohorts are largely similar and that our findings are indeed robust.

Alternate Outcomes

Because the subscription program we study is an online-only program, our main analysis focused on customer purchases only on the website. The firm we partnered with has both brick-and-mortar and online presence and is able to link customer purchases between online and offline channels at the individual level through its reward program. We thus investigate whether the increase in online purchases through subscription was due to channel-switching behavior to online from offline (e.g., Forman, Ghose, and Goldfarb 2009; Wang and Goldfarb 2017).

We perform the analysis described previously, replacing online purchases with combined online and offline purchases and retaining the operationalization of all covariates in Table 2. We find that the treatment and noneconomic effects on total purchase amount are \$26.72 (SE = 1.68) and \$19.41 (SE = 1.75), respectively, suggesting that in our context, the online and offline channels are only weak

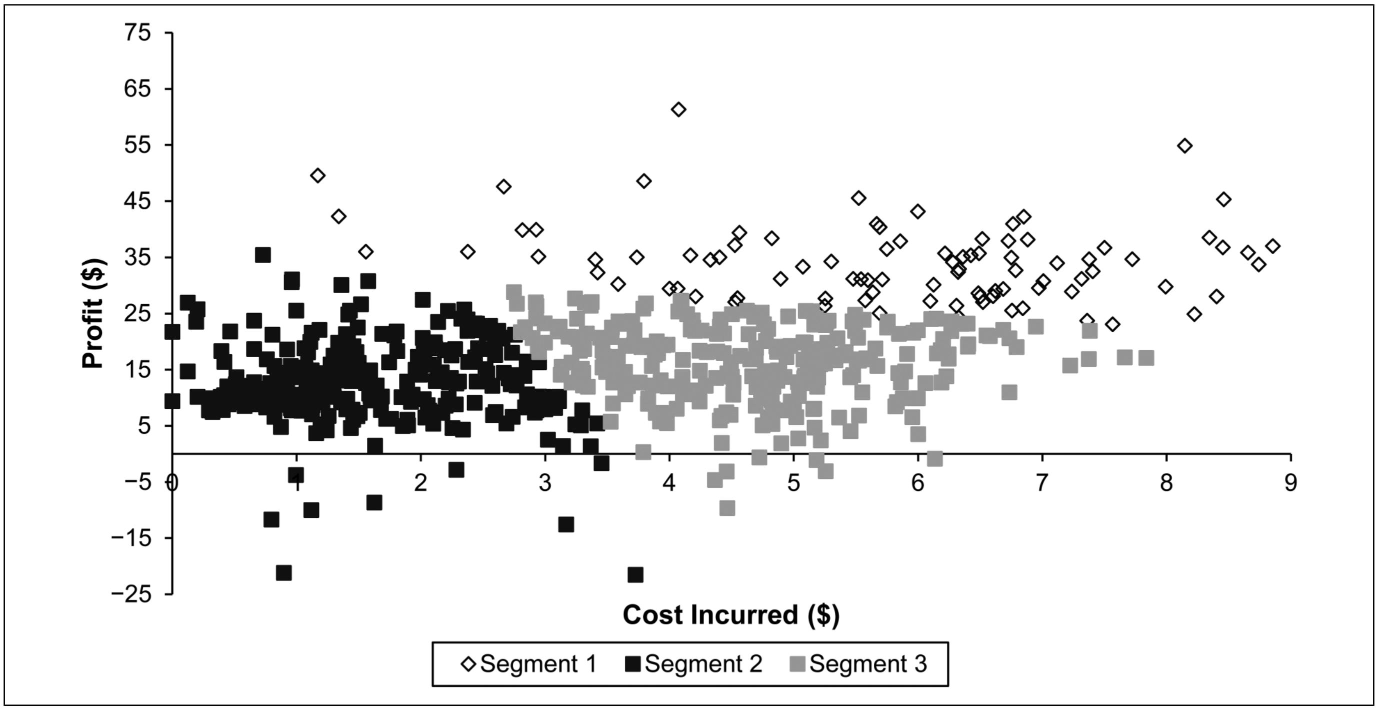

Figure 4. Scatterplot of profit and cost.

substitutes.[28] In summary, the program is effective in lifting overall revenue for the firm in our context.

Profitability of the Program

A successful subscription program can lead to an increase in sales but may induce additional costs for the firm. Because the profitability of the program will depend on whether the increase in revenue outweighs the additional cost, we illustrate how firms can calculate the profitability of the subscription program. We also investigate the characteristics of members who were most profitable. We leverage information on how members utilized program benefits and construct a cost measure to calculate the profitability. We note that some of the cost information for the profit calculation is not readily available (e.g., profit margins at the product level). Therefore, the cost measure we construct in this section is a proxy for the actual costs to the firm. Our intention is to illustrate the key trade-offs a firm may face when launching a subscription program or expanding its scale.

We construct the gross profit as the increase in revenue less the increase in the cost of goods sold (COGS).[29] We also consider three types of costs associated with the program: the free shipping service, the $3 gift card per month, and the free samples with purchase. According to our results, members ordered more frequently than nonmembers, resulting in slightly higher shipping costs for the firm. The program incurred costs from gift cards and free samples as well. Using data on how members utilized the benefits of the program, we construct the cost measure as the increase in shipping cost (number of orders made multiplied by the shipping cost per order) and add the costs incurred through redemptions of gift cards and free samples.[30]

We find that the program generated an average monthly net profit of $13.71 per member. Although information on customer purchases beyond the first year of the subscription program is not available, we calculate the lifetime profit based on our estimates of the one-year treatment effects and the information on member retention. The expected lifetime profit per member is $366.[31] We find significant variation in profit across members. Figure 4 plots each member

[28] We also performed the analysis controlling for additional covariates using offline purchases, such as the monthly purchase amount at the category level. The results are similar.

[29] We observe that the COGS typically ranges from 20% to 30% of revenue in the beauty industry. In our analysis, we assumed that the COGS is 25% of the list price.

[30] We cannot reveal shipping costs that the firm paid to delivery service companies because of the nondisclosure agreement, and we note that shipping costs in Asia were considerably lower than those in the United States. The cost of a sample was derived from the price of the full-sized regular product assuming that cost is proportional to size.

[31] We observe that 55% of the members renewed their subscription after the first year. Assuming that the program's attrition rate stayed constant over time, we calculated the expected lifetime profit per member as follows: (Profit per year)/(1 − Renewal rate) = ($13.71 × 12)/(1 − .55) = $366.

Table 6. Profitability and Characteristics by Segment.

Variable	Segment 1	Segment 2	Segment 3
Profit ($)	35.21	12.53	16.16
Cost ($)	5.84	1.52	4.70
Segment size (% of all members)	14%	46%	40%
Customer–firm relationship			
Tenure	6.46	6.02	6.37
Breadth	1.50	1.38	2.66
Depth	1.57	1.29	2.45
Basket size ($)	13.92	26.94	29.54
Monthly purchase amount ($): category level	1.56	2.45	3.53
Monthly purchase amount ($): SD	14.22	18.79	25.71
Purchase characteristics			
Exploratory	53.09	4.03	23.41
Repetitive	.35	.02	.15
Promotional	.63	.58	.36
Demographics			
Age	36.02	35.37	35.80
Gender (1 if female)	.98	.93	.95

according to the cost incurred through subscription and (gross) profit.

Taking the analysis a step further, we apply the K-means clustering algorithm to segment the members based on these two measures. We find that members can be grouped into three segments: one small segment (Segment 1, accounting for 14% of members) and two large segments (Segments 2 and 3, accounting for 46% and 40% of members, respectively). Segment 1 (represented by diamonds in Figure 4) includes the members who contributed the highest profit but incurred the largest cost for the firm. This segment generated about 30% of the total profit. The largest segment, Segment 2 (black squares in Figure 4), represents the members who generated a sizable increase in profit but did not incur as much cost related to subscription. Segment 3 (gray squares in Figure 4) is similar to Segment 1 in terms of the incurred cost but differs significantly in terms of the profit.

We also find that the three segments differed in terms of their observed characteristics. Table 6 reports the summary statistics of observed characteristics for the three segments. Interestingly, the members who contributed highest profit for the firm upon subscription (Segment 1) are not the ones who purchased the most prior to subscription. Rather, they were relatively less active customers in the presubscription period but were willing to explore new products (exploratory), had repeated purchases (repetitive), and were responsive to promotions (promotional) so that they reengaged with the firm upon subscription. The high-value customers (Segment 3) based on past purchases, while enjoying program benefits, increased their purchases only moderately, which is likely due to a ceiling effect. These results can assist managers in scoring and targeting future customers for the subscription program.

Conclusions

In this article, we examine the causal effect of a subscription program on customer behavior. We combine the DD approach with the GRF procedure and obtain estimates of the treatment effect at the individual level. We find that subscription is effective in lifting sales and keeps members more engaged in terms of frequency and variety in their purchases. The effect is economically significant, persistent over time, and heterogeneous across customers. Our main findings are robust to potential confounding effects of self-selection and unobservables, different treated groups, and different outcomes of purchase behavior.

Interestingly, we find that only one-third of the effect on customer purchases is due to the economic benefits of the subscription program, whereas the remaining two-thirds is attributed to the noneconomic effect. To uncover the underlying mechanism that leads to the behavioral changes, we leverage the treatment effects at the individual-month level in our estimation. In addition to multiple drivers found in the context of loyalty programs (e.g., affect, habit, status), notably we document evidence that members experience a sunk cost fallacy due to the up-front payment that subscription programs entail. Further, we discuss the profitability of the subscription program.

Our findings shed light on the practices of subscription-based businesses for customer retention and development. We find that subscription programs are broadly effective in lifting sales and enhancing customer engagement (e.g., breadth and depth of purchases in terms of products and categories). Our findings also have direct implications for the design of subscription programs. Although we study one subscription program launched by a single firm, our finding that customers experience a sunk cost fallacy due to the up-front payment is applicable to a wide range of contexts, as the up-front fee is a feature common to subscription programs. Companies may benefit from making the up-front fee more salient to customers after they become members to incentivize their purchases. This suggestion is also in line with prior work showing that making prices salient can make members consume a service on a more consistent basis (e.g., Gourville and Soman 2002). In addition, we find that members' enhanced status encourages purchases, suggesting that managers can also emphasize the exclusiveness of the program to boost sales. Our findings contribute to the ongoing debate concerning the amplified frictions in the retail industry. As firms are able to lock in customers with subscription programs by creating a sunk cost fallacy, the result may be insufficient switching and increased market concentration due to the rise of subscription programs (e.g., Amazon Prime).

As our research is the first attempt to identify the effect of subscriptions in the retail area on customer behavior, naturally there are limitations that should be acknowledged and addressed in future research. First, as our study focused on the effect of a subscription program in a given firm, it is likely that some of our findings could reflect the customer base and product categories of our partner firm. The subscription period is also reasonably long (one year) and early termination was not allowed, so the sunk cost is prominent. The subscription program examined in this research is an online-only program, whereas members in other subscription programs could benefit at both online and offline channels (e.g., Amazon, Barnes & Noble). With these limitations in mind, we hope our approach provides a framework for further studies on subscription programs with various pricing schemes or framing (e.g., subscription programs with varying time windows) in other product categories and across channels. Second, while we examined customer behavior in the first year of subscription rollout, it could be interesting to measure the long-term (e.g., three-year) effect of subscriptions on customer behavior beyond the first year of subscription rollout. Finally, with the relevance and popularity of subscription-based businesses, it is possible that many more companies will have their own subscription programs. With competition in play, the effect of a subscription program on customer engagement and purchase remains unclear. We hope that our work will inspire further studies to deepen the understanding of this nascent and important area of research.

Table A1. Placebo Tests.

Variable	Placebo Effect
Purchase amount ($)	1.50
	(1.49)
Purchase frequency	.03
	(.03)
Basket size ($)	1.20
	(2.93)

*$p < .1$.
**$p < .05$.
***$p < .01$.
Notes: Robust standard errors appear in parentheses. Placebo tests on the variety metrics are omitted. By design, the purchase amount of known products (and categories) equals the monthly purchase amount before subscription, and the purchase amount of new products (and categories) remains unchanged (at zero) before subscription.

Appendix A: Placebo Tests

We defined a placebo treatment on members six months prior to the actual subscription date. We estimated the effect of such a placebo treatment on customer purchases using a DD model on all weighted presubscription data. If the parallel trend assumption holds, we should expect null effects. As shown in Table A1, we find that members and matched nonmembers have statistically comparable pretreatment time trends.

Appendix B: Average Treatment Effects with log(Y_{it})

Table B1 reports the average treatment effects of subscription on customer purchases with log-transformed outcome measures.

Table B1. Total Treatment and Noneconomic Effects with log(Y_{it}).

	Total Treatment Effect				Noneconomic Effects			
Variable	All Months	Month 2	Months 3–4	Months 5+	All Months	Month 2	Months 3–4	Months 5+
Purchase amount ($)	1.57***	2.05***	1.80***	1.50***	1.34***	1.63***	1.49***	1.11***
	(.05)	(.13)	(.08)	(.27)	(.05)	(.11)	(.08)	(.05)
Purchase frequency	1.47***	1.87***	1.57***	1.67***	.47***	.54***	.54***	.40***
	(.11)	(.13)	(.12)	(.15)	(.01)	(.04)	(.02)	(.01)
Basket size ($)	−.34***	−.17*	−.24***	−.33***	−.24	−.16*	−.24***	−.29***
	(.14)	(.09)	(.05)	(.11)	(.17)	(.09)	(.04)	(.12)
Variety ($)								
Purchase amount in known categories	1.82***	2.08***	1.86***	1.63***	1.32***	1.58***	1.46***	1.13***
	(.05)	(.11)	(.08)	(.06)	(.05)	(.11)	(.07)	(.06)
Purchase amount in new categories	.19***	.15***	.24***	.18***	.15***	.12**	.21***	.17***
	(.02)	(.03)	(.03)	(.02)	(.03)	(.05)	(.02)	(.04)
Purchase amount of known products	.08**	.06	.05	.05	.07**	.11***	.04	.05
	(.04)	(.04)	(.04)	(.04)	(.03)	(.04)	(.03)	(.05)
Purchase amount of new products	1.92***	2.26***	2.13***	1.77***	1.42***	1.64***	1.51***	1.27***
	(.04)	(.09)	(.06)	(.04)	(.04)	(.08)	(.05)	(.05)

*$p < .1$.
**$p < .05$.
***$p < .01$.
Notes: Robust standard errors appear in parentheses. Customer purchases in the month of adoption (Month 1) were excluded from the analysis.

Acknowledgments

The authors thank several managers at the participating company for their efforts in making this research possible. They thank Ryan Dew, Pinar Yildirim, Christophe Van den Bulte, and seminar participants at the 2017, 2018, and 2020 ISMS (INFORMS Society for Marketing Science) Marketing Science Conferences, the 2019 Choice Symposium, the University of Chicago, the University of Michigan at Ann Arbor, the University of Pittsburgh, and the University of Texas at Austin for their comments.

Author Contributions

All authors contributed equally and are listed in alphabetical order.

Associate Editor

S. Sriram

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship and/or publication of this article.

ORCID iD

Qi Yu https://orcid.org/0000-0001-7889-0806

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010), "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 105 (490), 493–505.

Angrist, Joshua D. and Jorn-Steffen Pischke (2008), *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.

Arkes, Hal R. and Catherine Blumer (1985), "The Psychology of Sunk Cost," *Organizational Behavior and Human Decision Processes*, 35 (1), 124–40.

Athey, Susan, Julie Tibshirani, and Stefan Wager (2019), "Generalized Random Forests," *Annals of Statistics*, 47 (2), 1148–78.

Baumgartner, Hans (2002), "Toward a Personology of the Consumer," *Journal of Consumer Research*, 29 (2), 286–92.

Bertini, Marco and Luc Wathieu (2008), "Research Note—Attention Arousal Through Price Partitioning," *Marketing Science*, 27 (2), 236–46.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004), "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119 (1), 249–75.

Bijmolt, Tammo H.A., Harald J. van Heerde, and Rik G.M. Pieters (2005), "New Empirical Generalizations on the Determinants of Price Elasticity," *Journal of Marketing Research*, 42 (2), 141–56.

Bolton, Ruth N., P.K. Kannan, and Matthew D. Bramlett (2000), "Implications of Loyalty Program Membership and Service Experiences for Customer Retention and Value," *Journal of the Academy of Marketing Science*, 28 (1), 95–108.

Bolton, Ruth N., Katherine N. Lemon, and Peter C. Verhoef (2004), "The Theoretical Underpinnings of Customer Asset Management: A Framework and Propositions for Future Research," *Journal of the Academy of Marketing Science*, 32 (3), 271–92.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen,, Whitney Newey, et al. (2018), "Double/Debiased Machine Learning for Treatment and Structural Parameters," *Econometrics Journal*, 21 (1), C1–C68.

Datta, Hannes, George Knox, and Bart J. Bronnenberg (2017), "Changing Their Tune: How Consumers' Adoption of Online Streaming Affects Music Consumption and Discovery," *Marketing Science*, 37 (1), 5–21.

Drèze, Xavier and Joseph C. Nunes (2009), "Feeling Superior: The Impact of Loyalty Program Structure on Consumers' Perceptions of Status," *Journal of Consumer Research*, 35 (6), 890–905.

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi (2018), "Income and Wealth Effects on Private-Label Demand: Evidence from the Great Recession," *Marketing Science*, 37 (1), 22–53.

Forman, Chris, Anindya Ghose, and Avi Goldfarb (2009), "Competition Between Local and Electronic Markets: How the Benefit of Buying Online Depends on Where You Live," *Management Science*, 55 (1), 47–57.

Galak, Jeff and Joseph P. Redden (2018), "The Properties and Antecedents of Hedonic Decline," *Annual Review of Psychology*, 69, 1–25.

Goldfarb, Avi and Catherine E. Tucker (2011), "Privacy Regulation and Online Advertising," *Management Science*, 57 (1), 57–71.

Goli, Ali, Pradeep Chintagunta, and S. Sriram (2021), "Effect of Payment on User Engagement in Online Courses," *Journal of Marketing Research*, 59 (1), 11–34.

Gopalakrishnan, Arun, Zhenling Jiang, Yulia Nevskaya, and Raphael Thomadsen (2021), "Can Non-tiered Customer Loyalty Programs Be Profitable?" *Marketing Science*, 40 (3), 508–26.

Gourville, John and Dilip Soman (2002), "Pricing and the Psychology of Consumption," *Harvard Business Review*, 80 (9), 90–96.

Grier, Sonya A. and Rohit Deshpandé (2001), "Social Dimensions of Consumer Distinctiveness: The Influence of Social Status on Group Identity and Advertising Persuasion," *Journal of Marketing Research*, 38 (2), 216–24.

Hamilton, James Douglas (2020), *Time Series Analysis*. Princeton, NJ: Princeton University Press.

Hanushek, Eric A. (1974), "Efficient Estimators for Regressing Regression Coefficients," *The American Statistician*, 28 (2), 66–67.

Hartmann, Wesley R. and V. Brian Viard (2008), "Do Frequency Reward Programs Create Switching Costs? A Dynamic Structural Analysis of Demand in a Reward Program," *Quantitative Marketing and Economics*, 6 (2), 109–37.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York: Springer.

Hirano, Keisuke, Guido W. Imbens, and Geert Ridder (2003), "Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score," *Econometrica*, 71 (4), 1161–89.

Ho, Teck-Hua, I.P.L. Png, and Sadat Reza (2017), "Sunk Cost Fallacy in Driving the World's Costliest Cars," *Management Science*, 64 (4), 1761–78.

Imbens, Guido W. and Donald B. Rubin (2015), *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge, UK: Cambridge University Press.

Iyengar, Raghuram, Kamel Jedidi, Skander Essegaier, and Peter J. Danaher (2011), "The Impact of Tariff Structure on Customer Retention, Usage, and Profitability of Access Services," *Marketing Science*, 30 (5), 820–36.

Kopalle, Praveen K., Yacheng Sun, Scott A. Neslin, Baohong Sun, and Vanitha Swaminathan (2012), "The Joint Sales Impact of Frequency Reward and Customer Tier Components of Loyalty Programs," *Marketing Science*, 31 (2), 216–35.

Lal, Rajiv and David E. Bell (2003), "The Impact of Frequent Shopper Programs in Grocery Retailing," *Quantitative Marketing and Economics*, 1 (2), 179–202.

Leenheer, Jorna, Harald J. van Heerde, Tammo H.A. Bijmolt, and Ale Smidts (2007), "Do Loyalty Programs Really Enhance Behavioral Loyalty? An Empirical Analysis Accounting for Self-Selecting Members," *International Journal of Research in Marketing*, 24 (1), 31–47.

Lewis, Michael (2004), "The Influence of Loyalty Programs and Short-Term Promotions on Customer Retention," *Journal of Marketing Research*, 41 (3), 281–92.

Lewis, Michael (2006), "The Effect of Shipping Fees on Customer Acquisition, Customer Retention, and Purchase Quantities," *Journal of Retailing*, 82 (1), 13–23.

Lewis, Michael, Vishal Singh, and Scott Fay (2006), "An Empirical Study of the Impact of Nonlinear Shipping and Handling Fees on Purchase Incidence and Expenditure Decisions," *Marketing Science*, 25 (1), 51–64.

Liang, Kung-Yee and Scott L. Zeger (1986), "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika*, 73 (1), 13–22.

Liu, Yuping (2007), "The Long-Term Impact of Loyalty Programs on Consumer Purchase Behavior and Loyalty," *Journal of Marketing*, 71 (4), 19–35.

Manchanda, Puneet, Grant Packard, and Adithya Pattabhiramaiah (2015), "Social Dollars: The Economic Impact of Customer Participation in a Firm-Sponsored Online Customer Community," *Marketing Science*, 34 (3), 367–87.

McCarthy, Daniel M., Peter S. Fader, and Bruce G.S. Hardie (2017), "Valuing Subscription-Based Businesses Using Publicly Disclosed Customer Data," *Journal of Marketing*, 81 (1), 17–35.

Narang, Unnati and Venkatesh Shankar (2019), "Mobile App Introduction and Online and Offline Purchases and Product Returns," *Marketing Science*, 38 (5), 756–72.

Prins, Remco and Peter C. Verhoef (2007), "Marketing Communication Drivers of Adoption Timing of a New E-Service among Existing Customers," *Journal of Marketing*, 71 (2), 169–83.

Reisinger,, Don (2017), "Here's How Much Amazon Prime Customers Spend Per Year," *Fortune* (October 18), https://fortune.com/2017/10/18/amazon-prime-customer-spending/.

Rogers, Everett M. (2003), *Diffusion of Innovations*, 5th ed. New York: Simon and Schuster.

Rosenbaum, Paul R. and Donald B. Rubin (1983), "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika*, 70 (1), 41–55.

Sato, Kazuo (1972), "Additive Utility Functions with Double-Log Consumer Demand Functions," *Journal of Political Economy*, 80 (1), 102–24.

Tellis, Gerard J. (1988), "The Price Elasticity of Selective Demand: A Meta-Analysis of Econometric Models of Sales," *Journal of Marketing Research*, 25 (4), 331–41.

Thaler, Richard (1980), "Toward a Positive Theory of Consumer Choice," *Journal of Economic Behavior and Organization*, 1 (1), 39–60.

Venkatesan, Rajkumar and Paul W. Farris (2012), "Measuring and Managing Returns from Retailer-Customized Coupon Campaigns," *Journal of Marketing*, 76 (1), 76–94.

Wang, Kitty and Avi Goldfarb (2017), "Can Offline Stores Drive Online Sales?" *Journal of Marketing Research*, 54 (5), 706–19.

Wood, Wendy and David T. Neal (2009), "The Habitual Consumer," *Journal of Consumer Psychology*, 19 (4), 579–92.