JUSTC

Article

http://justc.ustc.edu.cn

Received: April 18, 2023; Accepted: September 12, 2023

An empirical study on the effect of user engagement on personalized free-content promotion based on a causal machine learning model

Shuang Wang, Hanbing Xue, and Lizheng Wang ✉

School of Management, University of Science and Technology of China, Hefei 230026, China

✉Correspondence: Lizheng Wang, E-mail：lzwang@mail.ustc.edu.cn

Graphical abstract

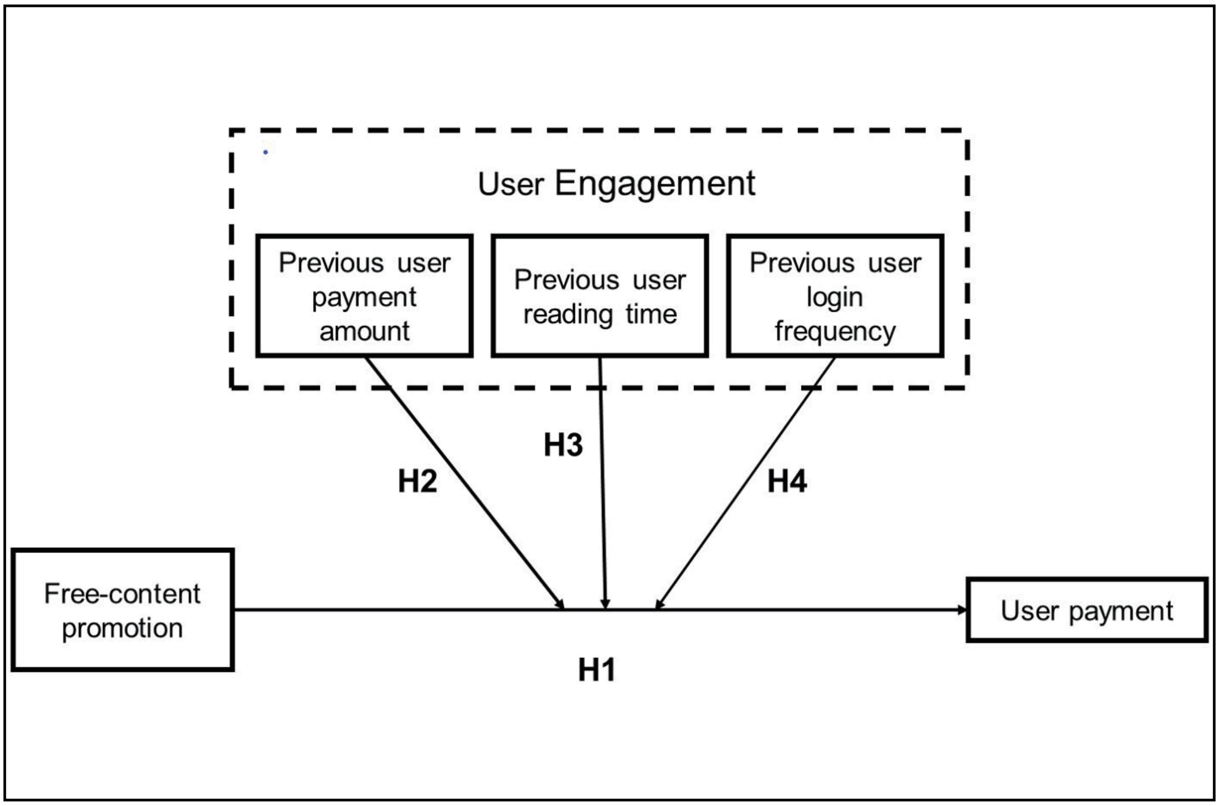

The hypothesis test result of this study.

Public summary

- This study demonstrates that, on average, free-content promotion can significantly increase users' payment amounts.
- However, these individual effects of promotion show considerable variation among users, with more-engaged users exhibiting a lower positive response to this promotion compared to less-engaged users.
- We observe that while the cannibalization effect of free-content promotion is minimal for less-engaged users, the expansion and acceleration effects are pronounced; conversely, this promotion has a notable cannibalization effect on more-engaged users.

Citation: Wang S, Xue H B, Wang L Z. An empirical study on the effect of user engagement on personalized free-content promotion based on a causal machine learning model. *JUSTC*, **2024**, 54(10): 1005. CSTR: 32290.14.JUSTC-2023-0063. DOI: 10.52396/JUSTC-2023-0063

JUSTC

Article

http://justc.ustc.edu.cn

Received: April 18, 2023; Accepted: September 12, 2023

An empirical study on the effect of user engagement on personalized free-content promotion based on a causal machine learning model

Shuang Wang, Hanbing Xue, and Lizheng Wang ✉

School of Management, University of Science and Technology of China, Hefei 230026, China

✉Correspondence: Lizheng Wang, E-mail：lzwang@mail.ustc.edu.cn

Cite This: *JUSTC*, **2024**, 54(10): 1005 (12pp)   Read Online

Abstract: Many digital platforms have employed free-content promotion strategies to deal with the high uncertainty levels regarding digital content products. However, the diversity of digital content products and user heterogeneity in content preference may blur the impact of platform promotions across users and products. Therefore, free-content promotion strategies should be adapted to allocate marketing resources optimally and increase revenue. This study develops personalized free-content promotion strategies based on individual-level heterogeneous treatment effects and explores the causes of their heterogeneity, focusing on the moderating effect of user engagement-related variables. To this end, we utilize random field experimental data provided by a top Chinese e-book platform. We employ a framework that combines machine learning with econometric causal inference methods to estimate individual treatment effects and analyze their potential mechanisms. The analysis shows that, on average, free-content promotions lead to a significant increase in consumer payments. However, the higher the level of user engagement, the lower the payment lift caused by promotions, as more-engaged users are more strongly affected by the cannibalization effect of free-content promotion. This study introduces a novel causal research design to help platforms improve their marketing strategies.

Keywords: free-content promotion; user engagement; random experiment; causal machine learning; individual-level treatment effect

CLC number: F274; TP181   **Document code:** A

1 Introduction

Digital content has become one of the fastest growing industries in the new economy. Many countries have introduced promotional policies to foster business opportunities. However, direct quality information about a digital content product can be obtained only after it has been viewed or played, and imperfect information about its content often reduces users' willingness to pay for it. Owing to the low cost of duplicating digital content products, many platforms have begun to adopt a free-sampling strategy to reduce user uncertainty about digital content products, applying it to products such as e-books, TV series, and music albums. For example, e-book platforms, such as Wattpad and iReader, provide free online previews of e-books, usually the first 5% of the total number of chapters in the book. Similarly, video content platforms, such as Youku, iQiyi, and Netflix, provide free partial viewing; users who want to watch the full video must pay extra or become a premium member. These free-sample strategies can enhance consumers' direct product experience and stimulate purchases[1–3]. However, the effects of free promotions are highly heterogeneous among digital content products and consumer groups. Some users may respond positively to free promotions, whereas others may perceive them as a derogation of product quality. Many companies have spent much of their marketing budgets with little gain. Therefore, companies should consider personalized promotional strategies to avoid or mitigate the negative impact of free promotions and target only those individuals who respond positively to free promotions, thereby saving marketing resources.

However, the empirical research on personalized promotions has utilized econometric models with interaction terms to study the heterogeneity of promotional effects, focusing on the development of personalized promotion strategies. Although this method is causally rigorous, it lacks personalization. Meanwhile, machine learning methods based on outcome prediction have achieved deep levels of personalization but lack causal rigor[4]. Therefore, a novel causal reasoning framework is required to address personalized promotions.

Many scholars have researched personalized promotional strategies for products other than digital content products. These studies suggested that consumers' past transactional behavior or user engagement plays a vital role in determining the target groups for various promotional interventions[5]. Furthermore, they all seemed to find that consumers with higher levels of engagement respond more strongly to promotional interventions[6–8].

However, the free promotion of digital content products differs significantly from the free promotion types studied in the literature, and the studies' conclusions may thus not apply. Consumers face high uncertainty levels about digital content products, and their willingness to pay for them is low[3, 9]. Moreover, digital content products such as e-books and TV episodes are serial in nature, and are offered using various payment forms, such as full subscription, per-unit subscription, and batch subscription. These characteristics lead to significant differences between digital content products and other products in terms of their promotion. Furthermore, the free promotion of an e-book involves offering part of the book for free; thus, free content is closely related to paid content but is only a part of the total content. For other products, however, the free sample in a promotion is very similar to the paid content. Few studies have examined the heterogeneous effects of free-content promotions for e-books, particularly concerning how users with varying levels of engagement respond to such promotions.

Therefore, this study addresses the following research questions: Which users are more sensitive to free-content promotion for digital content products, and to which users should promotional activities be targeted? What can account for the heterogeneity of free-content promotion effects, especially the role played by historical user engagement levels?

We explore our research questions through a random-field experiment for free content promotion using a popular mobile reading platform in China. We utilize a state-of-the-art causal inference algorithm for causal effect identification, which combines the causal forest (CF) technique and double/debiased machine learning (DML) model. Based on the individual treatment effects obtained by this algorithm, we can develop personalized promotion strategies and continue to examine the triggers responsible for the heterogeneity of treatment effects.

Our results reveal that, on average, free-content promotion leads to a significant increase in user payments. However, there is considerable heterogeneity in the treatment effects across users, wherein payment increases are lower for more-engaged users. We then utilize the acceleration-cannibalization-expansion (ACE) model proposed by Bawa and Shoemaker[10] to explain the underlying mechanism. We find that cannibalization effects offset a considerable portion of the positive effect of free-content promotions for more-engaged users. By contrast, the expansion and acceleration effects were dominant for less-engaged users.

This study has several important theoretical and practical implications. First, we utilize causal inference-based machine learning methods to obtain individual treatment effects (ITE), which enables personalized promotion strategies. Second, employing a framework that combines machine learning with econometric causal inference methods allows us not only to quantify how promotions affect users but also to analyze the causes of the heterogeneity of the promotion effect, particularly the moderating effect of user engagement-related variables. This method enriches the causal effect research. Third, our study extends the research on user engagement. User engagement cannot be measured meaningfully by any single indicator, and managers should define and explore it based on the specific business context to enable personalized strategies. Finally, we propose that managers should avoid or reduce the adverse effects of marketing interventions (such as the cannibalization effect in free promotions) by targeting only those who respond positively to the intervention, rather than basing their assessments exclusively on the magnitude of the post-intervention results.

2 Literature and hypotheses

In this section, we first survey the research on personalized promotion strategies. We then review the literature on the role of user engagement in promotional personalization. Finally, we review the literature on free promotion and develop the study's hypotheses from it.

2.1 Personalized promotion

Personalized promotion is the practice of tailoring marketing activities and promotional offers to customers based on their demographic characteristics, behavioral preferences, and real-time contexts. Empirical studies on personalized promotion have proposed promotion methods based on geographic location[11], demographic characteristics[12], and consumer behavior history[5]. For example, Ho et al.[11] analyzed how distance and local competition (i.e., the number of alternative options near consumers) affect mobile coupon redemptions. Ghose[13] modeled coupon redemption rates as a function of consumer characteristics and targeting strategies to analyze the impact of targeting strategies on coupon redemption. Bonfrer and Drèze[14] established bivariate hazard models that considered consumers' historical click behavior to personalize e-mail promotions. Venkatesan and Farris[15] explored how coupon campaigns tailored to consumer preferences affect consumer purchase behavior. Our personalized free-content promotion strategy is based on consumers' historical behavior and the characteristics of digital content products.

We review the literature on models and algorithms related to personalized promotional strategies. Most of the empirical literature mentioned above uses parametric methods based on econometric models such as linear models that consider interaction terms, logistic models, and hierarchical Bayesian models. Several researchers also used machine-learning tools to predict customer churn propensity and formulate marketing interventions for customers with high churn propensities, such as random forests[16], ensemble algorithms[17], and neural networks[18]. However, much of the recent literature suggested that outcome prediction-based machine learning methods are inferior to causal inference-based machine learning methods (e.g., uplift model)[19]. For example, Ascarza[4] used uplift models to estimate customer responses to retention programs and suggested that firms should target customers based on their sensitivity to customer marketing interventions rather than on their churn propensity. Customers with higher churn propensity do not necessarily respond positively to marketing interventions. Hitsch and Misra[20] combined a causal inference framework with the k-nearest neighbor (KNN) algorithm to develop a causal KNN regression algorithm by which to design a strategy for personalized e-mail promotions. Based on these insights, our work integrates

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

state-of-the-art causal machine learning with traditional econometric models. This integration aims to develop personalized promotion strategies and to analyze the moderating role of user engagement levels.

2.2 User engagement and personalized promotion

With the development of computer network technology, user engagement has become popular in mobile applications such as shopping websites, health applications, mobile reading platforms, and online gaming platforms. Human–computer interaction (HCI) scholars initially identified behavioral components as an essential aspect of user engagement[21]. Kim et al.[22] proposed that mobile users' engagement is a product of utilitarianism, hedonism, and social motivations. O'Brien and Toms[23] suggested that interactivity, user control, feedback, and novelty are important attributes for expressing user engagement. O'Brien and Cairns[24] further refined the theory of user engagement, arguing that user engagement is a quality of user experience. Researchers generally posited that user behavior plays an important role in determining user engagement levels[24]. Many scholars argued that user engagement is highly context-specific and should be measured appropriately based on the specific usage scenarios involved[25].

The relationship between enterprises and customers used to be limited to a simple exchange of goods or services. However, with the advent of Web 2.0, the concept of user engagement underwent a fundamental transformation. Users cease to be mere consumers and instead become creators and collaborators. Some platforms, such as social media and movie rating sites, lack transactional attributes. Consequently, new dimensions are required to measure user engagement. Van et al.[26] argued that customer engagement goes beyond transactional behavior and that the user interaction level is also a significant and distinct component for measuring user engagement. Researchers in the HCI field measured user engagement levels via app use time. Gu et al.[25] measured user engagement based on the duration of game sessions. However, categorizing users based on a single dimension is no longer adequate for dealing with increasingly complex business scenarios. In the context of online marketing, studies measured customer engagement by examining the incidence of specific activities, such as the number of user subscriptions and comments[27]. Lee et al.[28] defined engagement as the number of likes, comments, shares, and message clicks in their study on the impact of media advertising content on user engagement.

In addition, different user behaviors signify distinct levels of user engagement[27]. For instance, users who make more purchases on e-book platforms exhibit "higher engagement" than those who log in more frequently. Therefore, it is necessary to consider multiple dimensions of user engagement in a more comprehensive manner. Some researchers considered user payment behaviors and other behavioral patterns to define user engagement more comprehensively. Zhang et al.[9] proposed a hidden Markov model (HMM) for detecting user engagement levels, in which e-book platform users were clustered into different "hidden engagement" stage strata based on their reading time and subscription payments. Studies on the online retail industry have considered multiple dimensions of user engagement. Ellickson et al.[6] implemented a k-means clustering algorithm on variables related to historical customer behavior and classified customers into high and low engagement level groups. In their study, consumer engagement-related behaviors include opening e-mails, making purchases, and all behaviors associated with purchases or overall expenditures.

Our study defines user engagement in terms of three historical pre-promotion behavioral indicators—payment amount, reading time, and login frequency—rather than by using a single metric. This definition provides a comprehensive representation of user engagement and offers valuable insights helpful for creating targeted marketing strategies. Payment amount reflects users' transactional behavior, while reading time and login frequency represent the usage time and frequency of the software, respectively. Our definition aligns with O'Brien's theory of user engagement, the variables of which describe individuals' engagement levels during interactions[24]. These metrics play an important role in marketing interventions, similar to the recency, frequency, and monetary (RFM) metrics in online retailing.

2.3 Product uncertainty

Product uncertainty is an important concept in marketing. Consumers are unable to fully ascertain whether the quality, functionality, and applicability of a product aligns with their needs and expectations before purchasing it. This uncertainty may lead to consumer dissatisfaction after purchase, resulting in financial losses. Hong and Pavlou[29] divided product uncertainty into product quality and fit uncertainty. Product quality uncertainty refers to a product's failure to meet the promises made by manufacturers[30] or a supplier's failure to communicate product information to consumers effectively[31]. Product fit uncertainty refers to the degree to which consumers cannot evaluate whether a product's attributes align with their preferences[32].

Researchers classified product categories into experience and search goods based on search efficiency[33]. Products whose utility can be assessed through a simple search before purchase are referred to as "search goods". By contrast, products whose utility can be evaluated only via post-purchase experience are referred to as "experience goods." Consumers require personal experience to evaluate the utility of digital content products such as e-books, TV series, and music albums. Furthermore, the search process for digital content goods differs from that for physical goods because each digital product is unique and consumers must conduct a new search for each purchase[3]. Consequently, digital content products are considered experience goods. The quality of digital content products originates from the product experience, in which product fit plays a significant role. The quality of e-books can be assessed only after purchase, leading to a relatively high level of uncertainty among consumers[1].

Some studies analyzed various mechanisms for reducing consumer uncertainty about product attributes, including free sampling[34], reviews[34, 35], and word-of-mouth[36]. Chen et al.'s empirical analysis using product rating data from Amazon suggested that, while consumer reviews can alleviate uncertainty somewhat, their ability to address product fit

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

uncertainty is limited[35]. Consumers tend to rely on firsthand rather than secondhand experiences to obtain information about product quality when making purchase decisions. Free sampling can convey product quality to consumers efficiently, thereby reducing search costs and stimulating sales[3]. However, different individuals may have different evaluations and uncertainties regarding the same product, which requires companies to design strategies that reflect user heterogeneity[3]. However, research on this topic is limited.

2.4 Free promotion

Free promotion in this study refers specifically to the free-content promotion of e-books. Although this type of promotion is endowed with unique characteristics owing to the nature of digital content products, it is still closely related to conventional free-sample promotions. Early research on free-sample promotion was conducted in the context of physical products. Retailers have long offered free samples to incentivize consumers to try their products. This promotional practice conveys product quality to consumers effectively, reduces search costs, and encourages unplanned buying behavior[3]. Marks and Kamins[37] found that direct product experience (sampling) reduces uncertainty in product beliefs more effectively than indirect product experience (advertising).

Although digital content products incur high development costs, they can be replicated and distributed relatively cheaply[3]. Therefore, numerous digital content providers have begun offering partially free content to consumers to reduce search costs and stimulate purchases. In a study examining whether free promotions for digital TV services lead to more paying customers, Foubert and Gijsbrechts[38] found that free trials were more effective than advertising in reducing uncertainty and influencing consumers' quality beliefs. Choi et al.[1] demonstrated that free samples of digital content products provide users with a direct experience, thereby reducing the psychological distance between the user and the product and significantly increasing purchase likelihood. Hoang and Kauffman[3] also proved that free sampling of a household video-on-demand series positively affected consumers' likelihood of purchasing the series. Hence, by extending the e-book scenario, we propose the following hypothesis:

H1. Free content promotions have a positive effect on users' payment behavior.

However, free-content promotions may cannibalize regular purchases among users who intend to buy. In marketing science, "cannibalization" refers to the phenomenon in which a company's efforts to promote its products or services cause a decrease in market share for similar products or services. Researchers frequently utilized Bawa's ACE model to study the cannibalization effects of free-sample promotions. Bawa and Shoemaker[10] argued that free sampling produces an aggregate effect that decomposes into acceleration, cannibalization, and expansion effects. They proposed an ACE model to measure the impact of free-sample promotions on various consumer groups. The ACE model involves several core concepts. The first is the acceleration effect, wherein consumers who are likely to buy without free samples repeat the purchase of the sample product earlier after the promotion. The second is the cannibalization effect, wherein free samples reduce the number of repeat purchases made by consumers who would otherwise have purchased the product. The third is the expansion effect, whereby free samples increase the purchase volume of consumers who have no intention of buying. Halbheer et al.[39] argued that the cannibalization effect is a direct effect of free sampling, whereas the expansion effect is indirect. As with physical products, the free sampling of information products has expansion, acceleration, and cannibalization effects on sales[38]. However, because of the easy availability of the product, free sampling of information products is more severely affected by cannibalization. For example, Godinho de Matos and Ferreira[40] probed the impact of binge watching on video-on-demand subscriptions and found that it reduced customers' willingness to pay for them.

The e-books examined in this study are information products. Free promotions can cannibalize potential payments from highly engaged users willing to pay without them. Conversely, free-content promotions have little cannibalization effect on users with lower engagement levels and weaker willingness to pay, and mainly generate expansion and acceleration effects. Therefore, the moderating effect of user engagement on the effectiveness of free-content promotions can be hypothesized as follows:

H2. Free-content promotion for e-books has a less-positive impact on the payment behavior of users who have previously made more payments than on those who have made fewer payments.

H3. Free-content promotion for e-books has a less-positive impact on the purchasing behavior of users who have previously spent more time reading than on those who have spent less time reading.

H4. Free-content promotion for e-books has a less-positive impact on the purchasing behavior of users with a higher frequency of prior logins than those with a lower frequency.

3 Methodology

This section outlines our empirical methodology (see Fig. 1). First, we explain the implementation of a randomized field experiment regarding free-content promotion. Then, we provide a descriptive statistical analysis of the data. Next, we discuss our empirical research strategy, the CausalForestDML algorithm, which enables us to obtain individual-level treatment effects.

3.1 Experimental context and data description

We collaborate with a popular mobile reading platform in China to conduct a randomized field experiment on free content promotion. Established in September 2008, this platform is one of the world's leading digital reading platforms. It has over 140 million active monthly users and 500000 book copyright resources. The platform frequently employs promotional measures, such as discounts and free content, to promote its books and encourage users to make purchases.

The company has been considering free-content promotion as a means to stimulate consumption since 2018 due to the limited payment rate increase among users in the company's reading app business sector. Randomized experiments are the gold standard for causal inference and are considered the cleanest method of causal identification. Thus, on May 1,

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

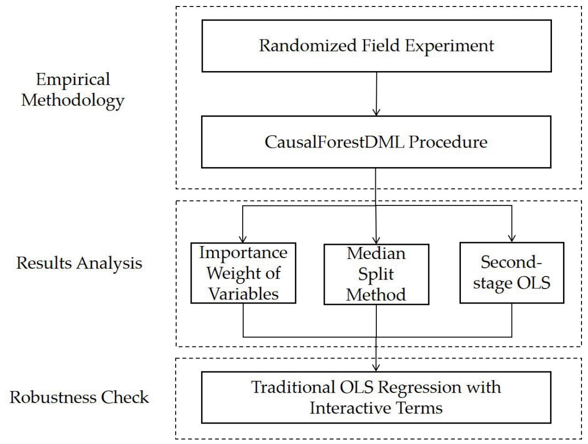

Fig. 1. Technology roadmap.

2018, the company launched a random field experiment to examine the effect of free-content promotion and provide a decision-making basis for implementing free-content promotion strategies. The core feature of this promotion was that the books pushed to users in the treatment group included coupons for 100 free chapters. Users could read the book during the promotion month; after the promotion ended, the extra free chapters reverted to paid status. By contrast, users in the control group received push notifications only for books that did not contain coupons. These users were unaware of the randomized experiment. Users in the treatment and control groups were randomly selected, each comprising 50% of the total sample, for a total of 7152 users. Notably, the pools of the pushed books in the treatment and control groups were the same. The book-push algorithm is based on user preferences, demographics, and other characteristics. The users who received push notifications chose whether to read the book and pay for the subsequent non-free chapters.

First, we obtain behavioral data for all participating users covering 30 days before the promotional campaign (i.e., before May 1), including purchase amounts, reading time, and login frequency. As these data were generated by users who participated in the experiment before the promotional campaign, they represented the users' level of historical engagement. In addition, we include the users' registration tenure data at the start of the promotional campaign. Because users may exhibit different payment behaviors toward different books, we also collect feature data for the pushed books, which served as control variables in the subsequent analyses. These data include the total number of chapters, the time the book spent on the shelf, the number of positive reviews, the number of negative reviews, and whether VIP users could purchase at a discount. Finally, at the end of May, we obtain the corresponding payment amounts for the recommended books from each participating user during that month. Table 1 presents a description of the relevant variables.

The summary statistics of the user and book characteristics are presented in Table 2. Although users logged an average of approximately 25 times before the promotion, their average payment amount and reading time remained low, at CNY 0.878 and 14.439 min, respectively. Therefore, a single metric is insufficient for reflecting user engagement levels.

The *t*-test results for these characteristics (see Table 3) show no significant difference in users between the control and treatment groups, providing strong evidence of balance between the two groups. The postpromotion user payment data refer to the user's payment for the books pushed and not the total payment in the promotion month. This practice avoids a misestimation of the effects of free-content

Table 1. Variable descriptions of users and book.

Variable	Description
Prepromotion user characteristics	
PrePay	Paid amount of the user in the month prior to the experiment (CNY)
PreRead	Reading time of the user in the month prior to the experiment (min)
PreLogin	Number of user logins in the month prior to the experiment
Tenure	Registration time of the user (d)
Pushed book characteristics	
ChapterCount	Number of chapters in the book
Updays	Time the book has been on shelf (d)
GoodComments	Number of positive reviews
BadComments	Number of negative reviews
Vip	A binary variable indicating whether it is a book that VIP users can purchase at a discount (Yes: 1, No: 0)
Postpromotion user payment	
Y	Amount paid by the user for the pushed book in the promotion month (CNY)

Table 2. Summary statistics of user and book characteristics.

Variable	Mean	Std	Min	Max
Prepromotion user characteristics				
PrePay	0.878	7.317	0	103.140
PreRead	14.439	99.000	0	4007.874
PreLogin	25.401	62.686	0	507
Tenure	108.596	321.556	10	3467
Pushed book characteristics				
ChapterCount	172.7480	48.766	128	465
Updays	654.128	551.447	4	3547
GoodComments	573.256	1733.386	0	25562
BadComments	60.267	143.505	0	3082
Vip	0.374	0.484	0	1
Postpromotion user payment				
Y (CNY)	6.344	8.752	0	154.620

Table 3. Randomization check.

	Control		Treatment		*p*-values
	Mean	Std	Mean	Std	
Number of participants	3576		3576		
PrePay	0.907	7.152	0.849	7.480	0.737
PreRead	16.352	89.089	12.526	107.985	0.102
PreLogin	25.812	68.653	24.990	56.094	0.580
Tenure	102.872	327.093	114.321	315.865	0.132

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

promotion.

3.2 Identification strategy

To infer the causal effects of different treatments, researchers typically estimate the average treatment effect of random interventions, which, in this study, is the average payment lift among users caused by a free promotion. However, there is significant heterogeneity in user responses to free promotions owing to the diversity of books and differences across user groups. Therefore, we are particularly interested in the heterogeneous effects of promotions. The empirical literature has used traditional regression models with interaction terms to analyze the heterogeneity of treatment effects (e.g., linear, squared, or cubic terms for each variable in the interaction terms and two- , three- , or higher-order interactions). These approaches require strong assumptions regarding possible linear, nonlinear, or interaction model specifications for the variables involved. Moreover, as there is an infinite number of such functions, whether such assumptions are correct cannot be guaranteed. To track recent advances in causal inference algorithms, we employ the CausalForestDML algorithm, which combines causal forests and double/debiased machine learning models to address this challenge[19, 41].

To describe the procedure in concrete terms, suppose we have data on n independent and identically distributed user-book pair samples $\{Y_i, X_i, W_i\}$, $i = 1, \cdots, n$, each composed of a vector of covariates or features $X_i \in \mathbb{R}$, a response outcome $Y_i \in \mathbb{R}$, and a binary treatment indicator $W_i \in \{0, 1\}$. In our research scenario, Y_i refers to the amount paid by the user for the pushed book during the promotion month, X_i represents the user characteristics before the promotion and the pushed book characteristics, and W_i indicates whether the user is from the promotion treatment group. Following the potential outcomes framework[42], the individual treatment effect (ITE) or conditional average treatment effect (CATE) of our interest can be expressed as $\tau(x) = \mathbb{E}[Y_i(1) - Y_i(0)|X_i = x]$, where $Y_i(1)$ and $Y_i(0)$ correspond, respectively, to what we would observe if we assign either treatment or control to user i.

In fact, we observe only one of the two potential outcomes: $Y_i(1)$ and $Y_i(0)$. Therefore, we must invoke the following assumptions to estimate CATE. The first is the unconfoundedness or ignorability assumption, namely, $W_i \perp \{Y(1), Y(0)\}|X$, which requires that potential outcomes and treatment assignments are independent of the observed covariates. The second is the overlap or positivity assumption, namely $0 < P(W = 1|X = x) < 1$, whereby, for any individual in the population, the probability of being assigned to the treatment and control groups is strictly greater than zero. It is reasonable to make these two assumptions because our data are taken from random experiments and the sample size is sufficiently large. Under these two assumptions, there are many candidate methods for estimating the heterogeneity treatment effects, including stratification methods, matching, regression adjustment, inverse propensity weighting, and various combinations thereof. The CausalForestDML was used in this study.

The causal forest (CF) is based on the classic random forest algorithm used for statistical learning. A random forest is a large collection of decision trees, each capable of predicting an outcome variable given a vector of covariates; the outcomes of all decision trees are then averaged. The difference between CF and classical random forests is that the split criteria for growing individual trees (also known as "causal trees") are specially designed to identify partitions in which the difference in the treatment effect is greatest.

First, each tree in CF recursively partitions the observations into subgroups along the chosen features, thereby clustering users with similar features into the same leaf. Once the tree partitions are obtained, the within-leaf treatment effect is estimated by calculating the difference between the average outcomes of the treated and control units on the same leaf:

$$\hat{\tau}_b(x) = \frac{1}{|\{i : W_i = 1, X_i \in L_b(x)\}|} \sum_{\{i:W_i=1,X_i\in L_b(x)\}} Y_i - \frac{1}{|\{i : W_i = 0, X_i \in L_b(x)\}|} \sum_{\{i:W_i=0,X_i\in L_b(x)\}} Y_i. \tag{1}$$

Wager et al.[19] proposed the concept of an "honest tree" to avoid overfitting and estimating sample noise as a heterogeneous treatment effect, as well as to ensure the consistency and asymptotic normality of the estimated effects. For each sample, the outcome variable Y can be used only to estimate the within-leaf treatment effect in Eq. (1) or to split the covariate space. The CF produces a collection of trees, each with an estimated value $\hat{\tau}_b(x)$. It then calculates a simple average of the estimates across B trees to obtain the CATE or ITE:

$$\hat{\tau}(x) = B^{-1} \sum_{b=1}^{B} \hat{\tau}_b(x). \tag{2}$$

This aggregation scheme also helps reduce variance and smooths out sharp decision boundaries.

However, Athey et al.[43] suggested that a centralized (or orthogonalized) CF can better handle confounding and heterogeneity, and also provides smaller mean squared errors, asymptotic normality, and the construction of confidence intervals. They had used the DML model to orthogonalize the causal forest. DML builds on and extends the method proposed by Robins et al.[44] The DML model makes the following structural equation assumptions regarding the data-generation process:

$$Y = \tau(x) \cdot W + g(X) + \varepsilon, \tag{3}$$

$$W = f(X) + \eta, \tag{4}$$

where $E[\varepsilon|X] = 0$, $E[\eta \mid X] = 0$.

Although the most direct method is to directly estimate $\tau(x)$ with Eq. (3), the estimated $\tau(x)$ is often biased, which comes partly from overfitting the sample and partly from the estimation of $g(X)$ deviation. Chernozhukov et al.[41] demonstrated that bias correction could be achieved using the following method. They fit Y and W using arbitrary machine learning (ML) models and then obtained the residuals (or centered outcomes), $\tilde{Y}$ and $\tilde{W}$, where $\tilde{Y} = Y - E[Y|x]$, $\tilde{W} = W - E[W|x]$. Then, they regressed the residual $\tilde{Y}$ with the residual $\tilde{W}$ (i.e., $\tilde{Y} = \tau(x) \cdot \tilde{W} + \varepsilon$). The regression coefficient $\tau(x)$ is the heterogeneous treatment effect we expect, which is

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

$$\tau(x) = \frac{\mathrm{Cov}\left[\left(\tilde{W}\right),\left(\tilde{Y}\right)\right]}{\mathrm{Var}\left(\tilde{W}\right)} = \frac{\mathrm{Cov}\left[(W - E[W|x]),(Y - E[Y|x])\right]}{\mathrm{Var}(W - E[W|x])}. \tag{5}$$

We implement the above process using the CF and aggregate the B-tree estimates into a weighted average to estimate the treatment effect for users with feature x. The data-adaptive weight is

$$\alpha_i(x) = \frac{1}{B}\sum_{b=1}^{B}\alpha_{ib}(x), \tag{6}$$

where $\alpha_{ib}(x) = \frac{\mathbb{I}(X_i \in L_b(x))}{|L_b(x)|}$. Here, $L_b(x)$ is the set of users that fall on the same leaf as that of a user whose covariate is x in the bth tree. The weight, $\alpha_i(x)$, represents how similar the ith user is to another user with feature $x(L_b(x))$, adjusted by the size of the leaf. These weights sum to one and are defined as forest-based adaptive neighborhoods x. Thus, we can modify Eq. (5) as follows:

$$\hat{\tau}(x) = \frac{\sum_{i=1}^{n}\alpha_i(x)(Y_i - \hat{m}^{(-i)}(X_i))(W_i - \hat{e}^{(-i)}(X_i))}{\sum_{i=1}^{n}\alpha_i(x)(W_i - \hat{e}^{(-i)}(X_i))^2}, \tag{7}$$

where $\hat{m}^{(-i)}(X_i)$ is the leave-one-out estimator for $m(x) = E[Y|x]$, and $\hat{e}^{(-i)}(X_i)$ is the leave-one-out estimator for $e(x) = E[W|x]$. This method reduces the estimation bias caused by overfitting and increases estimation robustness.

In marketing science, such causal inference algorithms (causal forests or DML) have been applied to analyses of customer retention[4], e-commerce cart targeting optimization[45], targeted e-mail promotions[6], the effect of information disclosure on industry payments to physicians[46], and the role of “Live” in live streaming markets[47]. We exploit the CausalForestDML procedure to probe the heterogeneous effects of e-book free-content promotion in the context of digital information goods.

4 Results analysis

First, we set the number of causal trees in the CausalForestDML program to 4000. Furthermore, given that the data utilized in this study are not overly complex, we train the estimator for the nuisance function $E[Y|x]$ using a LASSO regression model. We input three categories of data $\{Y_i, X_i, W_i\}$ into the corresponding CausalForestDML program. After training, we ultimately obtain the treatment effect estimate for each user $\hat{\tau}_{id}$, that is, the free-content promotion effect of the ith user for whom the dth book is pushed.

First, we discuss the average effect of promotions on user payments. Next, based on the technical roadmap shown in Fig. 1, we employ the median split method and second-stage OLS regression to explore the sources of heterogeneity in promotional effects across users[46]. Finally, we conduct robustness checks using a traditional OLS regression with interaction terms.

4.1 Treatment effect estimates

First, we report the average effect of free-content promotions on user payments. The results show that the promotion increases users’ payment by an average of CNY 9.4, with a 95% confidence interval of [6.8, 12.0], indicating that the payment lift caused by this promotion is statistically significant, thus supporting H1. However, this average treatment effect does not indicate how promotion effects vary among users with different levels of engagement. There may be differences in the treatment effects between less- and more-engaged users.

Therefore, we focus on the heterogeneous effects of free-content promotion. Recall that CausalForestDML is the bagging thousands of trees, providing point estimates and confidence intervals for each treated unit and allowing us to make statistical inferences about the promotion effect for each user. Therefore, we can observe the magnitude and significance of promotional effects at the individual level. Fig. 2 shows the distribution of free-content promotion effect estimates for each user. Although almost all users were affected positively, the payment lift caused by the promotion varied widely among them, with increases ranging from CNY 0 to 14. Approximately 15% of users who received the free-content promotion increased their payment by CNY 11 or more, approximately 18% increased their payment by only CNY 8 or less, and most increased their payment by CNY 8 to 11.

After estimating the individual treatment effects, the next step is to strategically assign free content coupons to users. A logical approach is to promote users whose payment lift is greater than the cost of the promotion (i.e., the opportunity cost of being free). The internal manager of this platform disclosed that each chapter of the e-book costs about CNY 0.08 and that providing 100 chapters for free costs about CNY 8. Therefore, approximately 18% of the users were not suitable for coupon distribution. However, companies often have relatively large restrictions and constraints on the number of coupons and requirements for minimum return on investment (ROI). We recommend sorting in descending order according to individual treatment effects and prioritizing users with larger individual treatment effects until the coupon limit is reached.

4.2 Sources of treatment effect heterogeneity

As understanding this heterogeneity is important for both theory and management practice, we examine its source. Similar to random forests, CausalForestDML provides an important

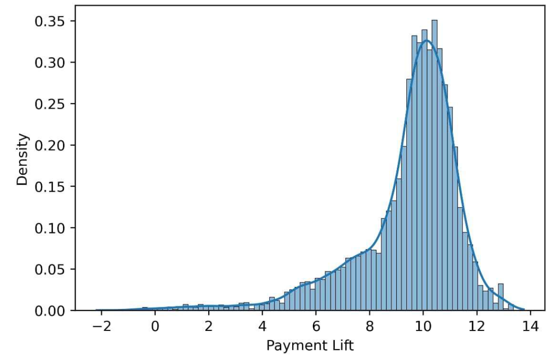

Fig. 2. Distribution of the treatment effect per user, $\hat{\tau}$.

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

measure of the covariates used in the estimation. This measure is computed using the frequency at which each covariate is split and weighted by the depth of each tree to reflect the relative weights of the covariates.

Table 4 presents the importance weights for each covariate and its rank. This result shows that the variables (PrePay, PreRead, and PreLogin) are closely related to user engagement account for approximately 63% of the total splits, indicating that engagement metrics are essential moderators of free-content promotion. This high percentage indicates the need for further study of these variables. Similar to the RFM metrics in online retailing, these variables play essential roles in determining the effectiveness of marketing intervention[6, 8]. In addition, the total number of chapters in the pushed books and shelf time were important variables that moderated the effect of free-content promotion (approximately 23%), indicating the need to control for book variables.

4.2.1 Median split method

Next, we explore the moderating role of covariates related to engagement in treatment effects and discuss how treatment effects vary according to user engagement level. First, we employ a grouping approach to compare user responses to free promotions at different levels of engagement. Specifically, for the three indicators of PrePay, PreRead, and PreLogin, we use the median split method (i.e., splitting by below or above the median) to divide users into two groups[46].

As shown in Table 5, the average payment lift for users with PrePay = 0 (approximately 83%) is CNY 9.647 (std, 1.585), while that of users with PrePay > 0 is CNY 4.465 (std, 3.168). Regarding prepromotion reading time, the average payment lift for users with PreRead = 0 is approximately CNY 2.521 more than that of users with PreRead > 0. Similarly, users with PreLogin ≤ 1 display an average payment lift of approximately CNY 1.957 more on average than users with PreLogin > 1.

4.2.2 Second-stage OLS regression

Although the simple median segmentation results show heterogeneous treatment effects across the three engagement-related dimensions, we also formally examined the existence

Table 4. Importance of covariates in heterogeneous treatment effects.

Variable	Importance	Rank
Prepromotion user data		
PrePay	8.973%	4
PreRead	10.837%	3
PreLogin	42.861%	1
Tenure	6.875%	6
Pushed book feature		
ChapterCount	14.530%	2
Updays	8.387%	5
GoodComments	3.730%	7
BadComments	3.408%	8
Vip	0.400%	9

Table 5. Payment lift by group.

Payment lift	N	Mean	Std	[95% conf. interval]	
$\hat{\tau}^{\text{PrePay}=0}$	6850	9.647	1.585	7.244	12.050
$\hat{\tau}^{\text{PrePay}>0}$	302	4.465	3.168	−1.015	9.945
$\hat{\tau}^{\text{PreRead}=0}$	5915	9.864	1.436	7.497	12.231
$\hat{\tau}^{\text{PreRead}>0}$	1237	7.343	2.744	3.805	10.881
$\hat{\tau}^{\text{PreLogin}\leq 1}$	3798	10.346	0.937	8.318	12.373
$\hat{\tau}^{\text{PreLogin}>1}$	3354	8.389	2.308	5.252	11.526
$\hat{\tau}$	7152	9.428	1.979	6.821	12.036

of this pattern after controlling for other user and pushed book characteristics. Following common practice in Ref. [46], we employ a second-stage OLS regression to investigate the heterogeneity of treatment effects across users. Although nonparametric methods, such as CausalForestDML, are not interpretable, we can perform post-hoc interpretability using the second model. Because these covariates related to user characteristics were all skewed, we log transformed them in the OLS regression, done in the $\log(*+1)$ form to keep all zero observations. Using the estimates of the treatment effect from CausalForestDML, we perform the following OLS regression on the estimated payment lift $\hat{\tau}$ for free-content promotion:

$$\begin{aligned}\hat{\tau}_{id} = & \alpha_0 + \alpha_1 \text{LogPrePay}_i + \alpha_2 \text{LogPreRead}_i + \\ & \alpha_3 \text{LogPreLogin}_i + \alpha_4 \text{LogTenure}_i + \\ & \beta_1 \text{ChapterCount}_d + \beta_2 \text{Updays}_d + \beta_3 \text{Comments}_d + \\ & \beta_4 \text{Vip}_d + \varepsilon_{id},\end{aligned} \tag{8}$$

where Comments_d represents the total number of comments on the dth book. The regression results are presented in Table 6.

Table 6. Relationship between payment lift and user/book characteristics.

Variable	$\hat{\tau}_{id}$	
	Coeff.	SE
Prepromotion user data		
LogPrePay	−0.722***	(0.031)
LogPreRead	−0.282***	(0.014)
LogPreLogin	−0.746***	(0.012)
LogTenure	−0.129***	(0.019)
Pushed book feature		
ChapterCount	0.003***	(0.000)
Updays	−0.0002***	(0.000)
GoodComments	0.000	(0.000)
BadComments	0.042	(0.033)
VIP	11.024***	(0.100)
R^2	0.587	
Observations	7152	

*** p<0.01, ** p<0.05, * p<0.1. Values in parentheses are standard errors.

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

The significant effects reported in Table 6 represent the relative difference in payment lift within each subgroup and not the relative difference in payments. The estimated coefficient of LogPrePay is negative and significant ($\alpha_1 = -0.722$, $p < 0.01$), suggesting that users with higher prepromotion payments will experience lower payment lift after the free-content promotion. A similar phenomenon occurs for users who spend more time reading before the promotion or log more before the promotion, which is consistent with our first analysis step. Therefore, when we control for other characteristics, the conclusion that less-engaged users are more responsive to free-content promotions than more-engaged users is supported. These findings thus support H2, H3, and H4.

We also draw interesting conclusions regarding the moderating effect of the pushed book characteristics. The coefficient for Updays is negative and significant ($p < 0.01$), suggesting that the longer the book spends on the shelf, the smaller the increase in the paid amount due to promotion. This finding shows that free content is an effective means of promoting new books. Additionally, the coefficient of ChapterCount is positive and significant ($p < 0.01$), indicating that the more chapters a book has, the greater the user's payment lift.

4.3 Robustness check for alternate specification

The CausalForestDML procedure we employed is a nonparametric method that makes us agnostic about the real data-generating process. In this regard, we employ traditional linear regression models to analyze the heterogeneity of treatment effects and test the main conclusions. Heterogeneity is reflected in the added interaction term in the traditional linear regression models. We use the median split method for three indicators: PrePay, PreRead, and PreLogin. We then construct the dummy variable vector Subclass $= \{\mathbb{I}(\text{PrePay} > 0), \mathbb{I}(\text{PreRead} > 0), \mathbb{I}(\text{PreLogin} > 1)\}$, where $\mathbb{I}$ is an indicator function. Note that the dependent variable Y_{id} is the amount paid by the ith user for the dth book in the promotional month, and we still control for the relevant feature Z_d of the pushed book. Therefore, we consider the following linear regression model:

$$Y_{id} = \gamma_0 + \gamma_1 \text{Treatment}_i + \gamma_1 \text{Subclass}_i + \gamma_3 \text{Treatment}_i \times \text{Subclass}_i + \phi Z_d + e_{id}. \tag{9}$$

Table 7 shows that the Treatment coefficient is positive and significant ($p < 0.01$), indicating that free-content promotion significantly impacts user payment behavior. Moreover, these results reveal that the coefficients in the Subclass vector are negative and significant ($p < 0.01$), indicating that the higher the level of user engagement, the smaller the effect of free-content promotion. These results prove that the conclusions obtained by our CausalForestDML are reliable.

Table 7. Results for the interaction between treatment and engagement.

Variable	Y_{id}	
	Coeff.	SE
Treatment	10.817***	(0.235)
Treatment interacted with Subclass		
PrePay > 0	−2.632***	(0.888)
PreRead > 0	−4.141***	(0.538)
PreLogin > 1	−0.945**	(0.365)
Intercept	0.855***	(0.370)
Control variables	Included	
R^2	0.387	
Observations	7152	

*** p<0.01, ** p<0.05, * p<0.1. Values in parentheses are standard errors.

5 Discussion

5.1 Potential mechanisms discussion for H1

We employ construal level theory (CLT) to account for the significant positive effect of free-content promotion on users' payment behavior. The CLT is a psychological representation theory that evolved from time construal theory. The CLT categorizes the processing of distant objects abstractly as high-level construals and the representation of close objects concretely as low-level construals[48]. In the marketing field, high-level psychological construals include reviews and brand-generated content[1], while low-level construals refer to the direct experience of a product. Although online reviews or brand-generated content may provide important information about product quality, such information is often inaccurate and biased, and cannot replace direct experience. Direct experience through free sampling reduces user uncertainty about a product more effectively than does indirect experience through advertising[1, 37].

In this study, users with free-content promotion coupons could obtain first-hand information and direct experiences. Additionally, they gain a deeper understanding of the book's quality, reducing their psychological distance from it. If readers are unable to learn how a story ends, it may cause dissatisfaction and regret. In general, coupon users are unlikely to abandon further reading because of mismatched preferences. In our setting, books were recommended based on user preferences. However, users without coupons generally have access to second-hand information and indirect experiences, resulting in weaker perceptions of e-book quality. Therefore, users with coupons were more willing to pay than those without.

5.2 Potential mechanisms discussion for H2–H4

Our explanation for hypotheses 2–4 is based on the ACE model proposed in Ref. [10], which decomposes the aggregate effect of free samples into three parts. The cannibalization effect is the direct negative effect of free-sample promotions, which is difficult to avoid. Digital content products are more severely affected by the cannibalization effect when free samples are offered[38, 40]. By contrast, the acceleration and expansion effects are indirect and positive effects expected from promotions. Following the practice of Reza et al.[49], in this study, the expansion effect refers to a situation in which $\mathrm{E}[Y(1) - Y(0) | Y_{\text{pre}} = 0] > 0$ for users with $\mathrm{E}[Y(0) = 0 | Y_{\text{pre}} = 0]$. The acceleration effect refers to a situation in which $\mathrm{E}[Y(1) - Y(0) | Y_{\text{pre}} = 0] > 0$ for users with $\mathrm{E}[Y(0) > 0 | Y_{\text{pre}} = 0]$.

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

Y_{pre} denotes the amount paid by the user for the pushed book before the promotion month. Since the user did not read the pushed book before the promotion month, $Y_{\mathrm{pre}}=0$. As users with a higher level of engagement display a lower payment lift from free-content promotions, can we conclude that cannibalization effects are more severe for more-engaged users?

To assess this explanation, we perform the following analyses. First, we obtain the counterfactual outcome for each user based on the factual outcome from the randomized experiments and the individual treatment effects estimated by CausalForestDML. Specifically, assuming that we observe a factual outcome of promotion intervention $Y_i(1)$ for a user, we can obtain the counterfactual outcome $Y_i(0)$ without a free-content coupon according to the individual treatment effect $Y_i(1)-Y_i(0)$. Similarly, given factual outcome $Y_i(0)$, we can obtain counterfactual outcome $Y_i(1)$.

We consider PreRead as an example. We find that, without free-content coupons, users whose PreRead equaled CNY 0 paid an average of approximately CNY 1, which was cannibalized by free promotions. Approximately 34% of users would not pay for the pushed books without coupons, but they would pay CNY 5.738 on average when they had coupons. Free-content promotions thus had an expansion effect on these users. The remaining users paid a small amount of money without coupons but they paid CNY 13.330 on average with coupons. Free-content promotion thus had accelerating effects on the remaining users. Therefore, the effect of free-content promotion on less-engaged users occurs mainly through expansion and acceleration effects, and the cannibalization effect is the weakest, similar to the findings of Ref. [49].

By contrast, users whose PreRead was more than CNY 0 paid an average of approximately CNY 15 when there were no free-content coupons. Among them, expansion users paid an average of approximately CNY 5.280 when they received coupons. For acceleration users with coupons, the average payment lift was only CNY 4.026. These findings show that, although free-content promotion had expansion and acceleration effects on more engaged users, the cannibalization effect offset a considerable part of this positive effect. Thus, the payment lift caused by free promotions is insufficient for more-engaged users. We find similar patterns when we analyze PrePay and PreLogin separately.

6 Conclusions

Although free samples, such as online previews, can alleviate consumer uncertainty about digital content products and stimulate purchases, the positive impact of a free sample strategy that targets all users is limited. Therefore, some companies have begun to consider personalized promotional strategies to increase user consumption and activity. In the context of digital content products, we study personalized free-content promotion strategies based on individual treatment effects and explore the sources of heterogeneity in promotion effects. We also explore how user engagement levels (historical payments, reading time, and number of logins) moderated the effects of free-content promotion interventions.

Utilizing random field experiment data provided by a popular e-book platform, we measure the impact of a free-content promotion on users' payment behavior. First, we obtain fine-grained estimates of promotion effects using a state-of-the-art causal inference algorithm (CausalForestDML) that combines causal forests with a DML model. Then, we analyze the source of the promotion effect heterogeneity using variable importance, median split, and second-stage OLS regression methods.

We find that, on average, free-content promotions can increase user payments significantly. However, our analysis of the individual treatment effects reveal that the effect of the promotion varied widely among users. Moreover, more-engaged users are less responsive to this promotion than low-engagement users are; that is, the payment lift is lower. Our conclusions are consistent across the three metrics closely related to engagement. We leverage the ACE model developed by Bawa and Shoemaker[10] to identify the potential mechanisms of users' heterogeneous responses to free-content promotion. Our analysis reveals that the cannibalization effect of free promotions on low-engagement users is weak but that the expansion and acceleration effects are significant. By contrast, this promotion has a significant cannibalization effect on more-engaged users. Moreover, more-engaged users already have a higher spending capacity, and providing additional free content may quickly diminish their interest in pushed books, thereby cannibalizing potential revenue.

6.1 Theoretical contributions

First, we employ causal inference-based machine learning methods to obtain individual treatment effects, which were then used to implement a personalized promotion. Moreover, we have expanded the scope of causal research by combining causal inference-based machine learning with traditional econometric methods. Second, this study extends the literature on user engagement. Prior research on user engagement has rarely considered the platforms on which digital content products are sold. Moreover, studies have represented user engagement through only a single metric[25], which is unreasonable. In addition to usage frequency, usage time, and transaction behavior are also important measures of user engagement levels on digital content platforms. Finally, we contribute to the theory of free promotion. Our analysis of the free promotion of digital content products differs significantly from previous promotion research on usage scenarios and methods. The e-books we study have various forms of payment, and their free promotional content is closely related to, but is not the same as, their paid content. We use the ACE model to explore the main reasons for the small increase in user fees from highly engaged users.

6.2 Managerial implications

Our findings have direct managerial implications for promotional strategies targeting digital content products. First, to save marketing resources, managers should focus their marketing activities on groups that are more sensitive to such activities. Traditional machine learning methods based on outcome prediction may inadvertently include users who would exhibit strong purchase intentions even without a promotion[4, 20]. Therefore, we recommend that managers use the

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

causal-inference-based machine learning approach proposed in this study to target users who are more sensitive to promotions. Second, managers can combine machine learning and causal research to formulate personalized promotional strategies based on individual treatment effects. Our causal research methodology and analysis conclusions can be generalized to other digital content products, such as comics, TV series, and music albums. Third, free-content promotion can significantly increase payments from less-engaged users, which shows that it is an effective means of stimulating new users or users who are about to churn. Therefore, this study can help companies target users and optimize the allocation of marketing resources when designing free promotions. Lastly, the cannibalization effect of promotional intervention is worth managers' attention and is not limited to the free promotions discussed in this paper. For example, discount coupons can cannibalize profits from loyal customers[50]. Therefore, managers must develop specialized marketing plans for users susceptible to promotion cannibalization[40].

6.3 Limitations and future research

Our study has several limitations that should be addressed in future research. First, we find that free-content promotion has a greater cannibalization effect on more-engaged users, but we did not analyze how to deal with this problem. For example, Godinho de Matos and Ferreira[40] showed that expanding recommendation reminders about consumer preferences for existing content can help mitigate the negative impact of cannibalization. Second, we include only two periods of user-behavior data (one month before the promotion and one month after it), which is the timeframe used by partner companies. Thus, we could not evaluate the long-term impact of free-content promotions on users. Third, we had no access to complete data on the platform, such as data on other dimensions related to engagement, certain demographic variables, and other book characteristics. Using such data would likely not only improve the accuracy of the individual treatment effect estimates but also enable us to perform a more fine-grained grouping of users. Finally, our experimental data collection lagged due to the company's data security and confidentiality concerns. However, this limitation does not affect the validity of our conclusions.

Acknowledgements

This work was supported by the Anhui Postdoctoral Scientific Research Program Foundation (2022B579).

Conflict of interest

The authors declare that they have no conflict of interest.

Biographies

Shuang Wang is a graduate student at the School of Management, University of Science and Technology of China. His research mainly focuses on big data marketing.

Lizheng Wang is a Special Associate Researcher at the University of Science and Technology of China (USTC). He received his Ph.D. degree from USTC in 2020. His research mainly focuses on big data marketing, econometrics, consumer dynamics analysis, and energy and environmental evaluation.

References

[1] Choi A A, Cho D, Yim D, et al. When seeing helps believing: The interactive effects of previews and reviews on e-book purchases. *Information Systems Research*, **2019**, *30* (4): 1164–1183.
[2] Ghose A, Han S P. Estimating demand for mobile applications in the new economy. *Management Science*, **2014**, *60* (6): 1470–1488.
[3] Hoang A-P, Kauffman R J. Content sampling, household informedness, and the consumption of digital information goods. *Journal of Management Information Systems*, **2018**, *35* (2): 575–609.
[4] Ascarza E. Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, **2018**, *55* (1): 80–98.
[5] Fudenberg D, Villas-Boas J M. Behavior-based price discrimination and customer recognition. In: Handbook on Economics and Information Systems, Volume 1. Amsterdam: Elsevier, **2006**: 377–436.
[6] Ellickson P B, Kar W, Reeder J C. Estimating marketing component effects: Double machine learning from targeted digital promotions. *Marketing Science*, **2022**, *42* (4): 704–728.
[7] Gopalakrishnan A, Park Y-H. The impact of coupons on the visit-to-purchase funnel. *Marketing Science*, **2021**, *40* (1): 48–61.
[8] Sahni N S, Zou D, Chintagunta P K. Do targeted discount offers serve as advertising? Evidence from 70 field experiments. *Management Science*, **2017**, *63* (8): 2688–2705.
[9] Zhang Y, Li B, Luo X, et al. Personalized mobile targeting with user engagement stages: Combining a structural hidden Markov model and field experiment. *Information Systems Research*, **2019**, *30* (3): 787–804.
[10] Bawa K, Shoemaker R. The effects of free sample promotions on incremental brand sales. *Marketing Science*, **2004**, *23* (3): 345–363.
[11] Ho Y-J, Dewan S, Ho Y-C. Distance and local competition in mobile geofencing. *Information Systems Research*, **2020**, *31* (4): 1421–1442.
[12] Sahni N S, Wheeler S C, Chintagunta P. Personalization in email marketing: The role of noninformative advertising content. *Marketing Science*, **2018**, *37* (2): 236–258.
[13] Ghose A, Li B, Liu S. Mobile targeting using customer trajectory patterns. *Management Science*, **2019**, *65* (11): 5027–5049.
[14] Bonfrer A, Drèze X. Real-time evaluation of e-mail campaign performance. *Marketing Science*, **2009**, *28* (2): 251–263.
[15] Venkatesan R, Farris P W. Measuring and managing returns from retailer-customized coupon campaigns. *Journal of Marketing*, **2012**, *76* (1): 76–94.
[16] Lemmens A, Croux C. Bagging and boosting classification trees to predict churn. *Journal of Marketing Research*, **2006**, *43* (2): 276–286.
[17] Risselada H, Verhoef P C, Bijmolt T H. Staying power of churn prediction models. *Journal of Interactive Marketing*, **2010**, *24* (3): 198–208.
[18] Hung S-Y, Yen D C, Wang H-Y. Applying data mining to telecom churn management. *Expert Systems with Applications*, **2006**, *31* (3): 515–524.
[19] Wager S, Athey S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, **2018**, *113* (523): 1228–1242.
[20] Hitsch G J, Misra S. Heterogeneous treatment effects and optimal targeting policy evaluation. SSRN 3111957, **2018**.
[21] Sutcliffe A. Designing for User Engagement: Aesthetic and Attractive User Interfaces. Cham, Switzerland: Springer, **2009**.
[22] Kim Y H, Kim D J, Wachter K. A study of mobile user engagement (MoEN): Engagement motivations, perceived value, satisfaction, and continued engagement intention. *Decision Support Systems*, **2013**,

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005

56: 361–370.
[23] O'Brien H L, Toms E G. What is user engagement? A conceptual framework for defining user engagement with technology. *Journal of the American Society for Information Science and Technology*, **2008**, *59* (6): 938–955.
[24] O'Brien H, Cairns P. Why Engagement Matters: Cross-Disciplinary Perspectives of User Engagement in Digital Media. Cham, Switzerland: Springer, **2016**.
[25] Gu Z, Bapna R, Chan J, et al. Measuring the impact of crowdsourcing features on mobile app user engagement and retention: A randomized field experiment. *Management Science*, **2022**, *68* (2): 1297–1329.
[26] Van Doorn J, Lemon K N, Mittal V, et al. Customer engagement behavior: Theoretical foundations and research directions. *Journal of Service Research*, **2010**, *13* (3): 253–266.
[27] Ahuja V, Medury Y. Corporate blogs as e-CRM tools — Building consumer engagement through content management. *Journal of Database Marketing & Customer Strategy Management*, **2010**, *17* (2): 91–105.
[28] Lee D, Hosanagar K, Nair H S. Advertising content and consumer engagement on social media: Evidence from Facebook. *Management Science*, **2018**, *64* (11): 5105–5131.
[29] Hong Y, Pavlou P A. Product fit uncertainty in online markets: Nature, effects, and antecedents. *Information Systems Research*, **2014**, *25* (2): 328–344.
[30] Pavlou P A, Liang H, Xue Y. Understanding and mitigating uncertainty in online exchange relationships: A principal-agent perspective. *MIS Quarterly*, **2007**, *31* (1): 105–136.
[31] Dimoka A, Hong Y, Pavlou P A. On product uncertainty in online markets: Theory and evidence. *MIS Quarterly*, **2012**, *36* (2): 395–426.
[32] Matt C, Hess T. Product fit uncertainty and its effects on vendor choice: An experimental study. *Electronic Markets*, **2016**, *26* (1): 83–93.
[33] Nelson P. Information and consumer behavior. *Journal of Political Economy*, **1970**, *78* (2): 311–329.
[34] Lin Z, Zhang Y, Tan Y. An empirical study of free product sampling and rating bias. *Information Systems Research*, **2019**, *30* (1): 260–275.
[35] Chen P, Hitt L M, Hong Y, et al. Measuring product type and purchase uncertainty with online product ratings: a theoretical model and empirical application. *Information Systems Research*, **2021**, *32* (4): 1470–1489.
[36] Godes D. Product policy in markets with word-of-mouth communication. *Management Science*, **2017**, *63* (1): 267–278.
[37] Marks L J, Kamins M A. The use of product sampling and advertising: Effects of sequence of exposure and degree of advertising claim exaggeration on consumers' belief strength, belief confidence, and attitudes. *Journal of Marketing Research*, **1988**, *25* (3): 266–281.
[38] Foubert B, Gijsbrechts E. Try it, you'll like it—or will you? The perils of early free-trial promotions for high-tech service adoption. *Marketing Science*, **2016**, *35* (5): 810–826.
[39] Halbheer D, Stahl F, Koenigsberg O, et al. Choosing a digital content strategy: How much should be free. *International Journal of Research in Marketing*, **2014**, *31* (2): 192–206.
[40] Godinho de Matos M, Ferreira P. The effect of binge-watching on the subscription of video on demand: Results from randomized experiments. *Information Systems Research*, **2020**, *31* (4): 1337–1360.
[41] Chernozhukov V, Chetverikov D, Demirer M, et al. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, **2018**, *21* (1): C1–C68.
[42] Rubin D B. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, **1974**, *66* (5): 688–701.
[43] Athey S, Tibshirani J, Wager S. Generalized random forests. *The Annals of Statistics*, **2019**, *47* (2): 1148–1178.
[44] Robins J M, Rotnitzky A, Zhao L P. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, **1994**, *89* (427): 846–866.
[45] Luo X, Lu X, Li J. When and how to leverage e-commerce cart targeting: The relative and moderated effects of scarcity and price incentives with a two-stage field experiment and causal forest optimization. *Information Systems Research*, **2019**, *30* (4): 1203–1227.
[46] Guo T, Sriram S, Manchanda P. The effect of information disclosure on industry payments to physicians. *Journal of Marketing Research*, **2021**, *58* (1): 115–140.
[47] Cong Z, Liu J, Manchanda P. The role of 'live' in livestreaming markets: Evidence using orthogonal random forest. SSRN 3878605, **2021**.
[48] Kruglanski A W, Higgins E T. Social Psychology: Handbook of Basic Principles. New York: The Guilford Press, **2013**.
[49] Reza S, Ho H, Ling R, et al. Experience effect in the impact of free trial promotions. *Management Science*, **2021**, *67* (3): 1648–1669.
[50] Reimers I, Xie C. Do coupons expand or cannibalize revenue? Evidence from an e-market. *Management Science*, **2019**, *65* (1): 286–300.

CSTR: 32290.14.JUSTC-2023-0063
DOI: 10.52396/JUSTC-2023-0063
JUSTC, **2024**, 54(10): 1005