



Marketing Science Institute Working Paper Series 2021

Report No. 21-102

Fewer Clicks, More Purchases

Murat Unal and Young-Hoon Park

“Fewer Clicks, More Purchases” © 2021

Murat Unal and Young-Hoon Park

MSI Working Papers are Distributed for the benefit of MSI corporate and academic members and the general public. Reports are not to be reproduced or published in any form or by any means, electronic or mechanical, without written permission.

Fewer Clicks, More Purchases

Murat Unal and Young-Hoon Park*

November 2020

*Murat Unal is a Ph.D. student in Marketing at the Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853; email: mu96@cornell.edu. Young-Hoon Park is the Sung-Whan Suh Professor of Management and Professor of Marketing at the Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853; email: yp34@cornell.edu. We thank several managers at the participating company for their efforts in making this research possible.

Fewer Clicks, More Purchases

Abstract

E-commerce retailers are increasingly faced with the challenges of finding ways to provide a seamless shopping experience to customers. The checkout process and its related touch points are especially critical in shaping customer experience, which would lead to purchases at an online store. We examine how enrollment in an option that reduces the number of steps required to place a purchase order, referred to as one-click buying, affects subsequent customer behavior online. Using quasi-experimental data over a period of 32 months from an online retailer before and after the launch of one-click buying, we find it is effective in lifting sales and does so by making treated customers purchase more often with more items per order. The impact of one-click buying on customer purchases post adoption is economically significant, persistent over time and heterogeneous across customers. Analyzing clickstream data of customer activity online and purchases across product categories, we provide evidence that the increase in customer purchases is driven by richer engagement through both more visits to the website and more searches upon visit as well as the expansion of purchases across categories. We discuss the implications of our findings for customer experience and targeting.

Keywords: Customer experience, One-click buying, Retailing, E-commerce, Causal inference, Machine learning, Generalized random forests

1 Introduction

E-commerce has become one of the most popular shopping channels worldwide and continues to increase its share in retailing. In the US, e-commerce sales account for about 14.5% of total retail sales in 2020 (eMarketer 2020) and are expected to reach \$970 billion by 2023, up from \$450 billion in 2017 (eMarketer 2019). Because e-commerce will continue to grow in the foreseeable future, increasing purchase conversion as well as customer loyalty is of great importance for e-commerce companies.

To achieve higher sales and repeat purchases, retailers have been working on finding ways to enhance customer experience. In addition to adopting policies such as free shipping and returns and buy-online-pick-up-in-store, retailers have transformed a myriad of touch points along the customer decision journey (e.g., Verhoef et al. 2009; Lemon and Verhoef 2016). For instance, virtual assistants enabled through automated bots help customers find products much faster and easier than ever before. Augmented reality tools provide customers the try-before-you-buy experience while shopping for apparel as well as furniture and beauty products. Thanks to voice-ordering devices, such as Amazon's Alexa, a new and convenient form of placing orders has also been introduced in e-commerce. Additionally, retailers try to enrich customer experience by increasing engagement via mobile apps, social media and online communities post purchase.

Despite many of these investments and efforts toward improving customer experience, most retailers still struggle to increase conversion rates, and shopping-cart abandonment remains a major challenge. According to a study, about 70% of e-commerce visitors abandon their shopping cart during the checkout process (Baymard Institute 2017). A long and complicated checkout process is consistently ranked as a top reason for not having made the purchase. Researchers have long suggested retailers create a seamless checkout process to ensure more customers will complete their transaction. Nevertheless, e-commerce websites continue to vary widely in the number of steps that customers have to go through from login to payment before confirming their order (e.g., Baymard Institute 2019).

One globally known online retailer that has been successful in minimizing frictions during

the checkout process is Amazon. Since 1997, Amazon has been using its 1-Click ordering technology to allow its customers to make their purchases in a single click. Amazon patented the feature in 1999 and prevented others from adopting the technology in their checkouts. It has been speculated that this feature alone increased Amazon's sales significantly and the patent was estimated to be valued at \$2.4 billion annually (e.g., Digiday 2017; Rejoinder 2017).

As of September 12, 2017, Amazon lost its exclusive right to the one-click buying technology because its 1-Click patent expired. Other retailers have gained the opportunity to adopt the feature without being forced to pay license fees. Indeed, several firms have responded to this opportunity. Magento, for example, adopted one-click buying under "Instant Purchase" and began offering it to its retail clients starting in December 2017 (Magento 2017). With the hurdles being removed, many more retailers and their partners are expected to adopt the one-click checkout technology and create a seamless checkout process (e.g., CNET 2017).

However, several aspects of one-click buying continues to remain a source of debate for e-commerce companies in deciding whether to implement this feature on their websites. Apart from anecdotal estimates, the economic value of ordering online in a single step is a secret known only by Amazon (e.g., Wagner and Jeitschko 2017). To date, no research has examined how one-click buying affects subsequent customer behavior. Even if we can estimate the causal effect of one-click buying on customer purchases post adoption, a general set of questions remains unanswered. Does it lift purchase amount due to the shift in order frequency and/or the change in order size? Does the impact last for the long term? What type of customers change their behavior most after registering for one-click buying? Does ordering in a single step serve as an effective tool only for repeat customers who purchase regularly? Or can it equally affect customers who like to explore and try new products (Lindner 2017)?

The objective of this paper is to shed light on the above issues by seeking answers to the following questions: Does the one-click checkout technology generate value for an online retailer? Specifically, is the feature effective in inducing customers to change their subsequent behavior? How does the effect vary over time and across customers? What are the underlying drivers? We

address these questions in close collaboration with a retailer in Asia that launched the one-click buying feature on its website in January 2017.¹ Using quasi-experimental data over a period of 32 months before and after the launch of one-click buying, we apply propensity score matching and combine it with the generalized random forests (GRF) procedure (Athey et al. 2019) to obtain treatment-effect estimates at the individual level.

Because adopting one-click buying is a voluntary decision on the part of customers, we measure the treatment effect by comparing the changes in purchase behavior of the adopters with those of the non-adopters after accounting for self selection into the treatment of one-click buying. We also apply the staggered-adoption framework (e.g., Bertrand and Mullainathan 2003; Goldfarb and Tucker 2014) because adoption dates of one-click buying vary among treated customers. Furthermore, GRF combine machine-learning tools with econometric methods to generate doubly robust individual-level treatment effects, which allow us to examine the heterogeneity in the treatment effect.

We find one-click buying is effective in lifting sales. On average, treated customers increased their purchase amount by 17.4% over a period of 12 months after adopting one-click buying. The feature increased the purchase amount by making treated customers purchase more often (an increase of 5.1%) with more items per order (an increase of 6.2%). We find the effect of one-click buying on customer purchases is economically significant and persistent over time. Our findings are robust to potential confounding effects of self-selection and unobservables, different treated and control groups and different outcomes of purchase behavior.

We also find there is a large variation in the treatment effect across customers. Specifically, the magnitude of the increase in purchase amount varies considerably, ranging from 0.01% to 30% post treatment. One-click buying had a limited impact on high-value customers (based on past purchases) but had a larger impact on those customers who previously purchased less in terms of amount and frequency. Compared to customers who like to explore and try new products, those customers who made repeat purchases increased their purchases more post adoption. We also find

¹Because Amazon's patent was only applicable in the US, our partner firm in Asia was free to introduce the feature on its online channel without any restrictions.

younger people were more responsive to the feature than older people. These results demonstrate the benefits of obtaining individual-level treatment-effect estimates by applying the GRF procedure. Our findings on heterogeneous treatment effects can assist marketers in scoring and targeting customers and also in allocating resources across individual customers when designing marketing activity associated with one-click buying.

To investigate the potential mechanisms that might explain our findings, we leverage a variety of data including offline purchases, product returns, customer activity online and purchases across product categories. Using these data, we explore channel-switching behavior, impulse-buying behavior and customer engagement as possible explanations. We find no evidence for channel substitution and limited evidence for impulse buying. We suggest one-click ordering provides an avenue for the online retailer to improve customer engagement, which leads to changes in purchase behavior. We provide evidence consistent with this explanation based on data of customer activity online and purchases across product categories. Specifically, treated customers visited the website significantly more often (an increase of 18.5%), viewed more pages (an increase of 7.3%) and spent more time (an increase of 9.4%) upon visit, and deepened their relationship with the firm by expanding their purchases across categories (an increase of 43.3%). Taken together, the improvement in the checkout process through one-click buying feature offers customers an enhanced customer experience, which increases their engagement at the e-commerce website and thereby leads to higher conversions.

Our paper is related to a few streams of research. Substantively, our study complements and extends the literature on customer experience (e.g., Verhoef et al. 2009; Lemon and Verhoef 2016). Previous research in this realm has reported how certain online features, such as decision aids and personalized shopping lists, improve customer experience and thereby lead to increased satisfaction and loyalty to the focal retailer (e.g., Häubl and Trifts 2000; Palmer 2002; Parboteeah et al. 2009; Shi and Zhang 2014). This paper contributes to this literature by using quasi-experimental data to measure the effect of one-click buying on customer behavior and shedding light on several aspects of the mechanism of the effect which the previous literature has not considered. Broadly,

our study fits in the literature of the impact of technology on customer behavior. Extant studies have analyzed how various applications of the latest technology reshape the retail landscape and impact customer behavior (e.g., Xu et al. 2016; Gallino and Moreno 2018; Shankar 2018). We add to this literature by documenting how one-click checkout technology affects customer behavior.

Additionally, we contribute to the nascent and growing literature on the effects of various marketing interventions on customer behavior using quasi-experimental data (e.g., Manchanda et al. 2015; Datta et al. 2017; Narang and Shankar 2019; Bell et al. 2020). Only a few papers in this area have examined heterogeneous treatment effects using machine-learning methods (e.g., Ascarza 2018; Fong et al. 2019). Our paper adds to this stream of research by offering an application of the combination of state-of-the-art machine learning methods with well-established econometric methods to a marketing related problem.

The remainder of the paper is organized as follows. §2 gives an overview of one-click buying and describes our research setting and data. §3 discusses our empirical methodology. §4 and §5 present the findings and discuss possible explanations for the effect. §6 presents our robustness checks. We conclude in §7.

2 Research Setting and Data

In this section, we describe one-click buying and the data that we have obtained to study how one-click buying affects customer behavior.

2.1 One-click Buying

One-click buying is a feature that allows customers to make purchases online in a single step. The hallmark of the feature is that it removes some of the frictions customers face in the checkout process. Specifically, it allows customers to buy items with just one click without having to re-enter billing, payment, or shipping information at each purchase. Also, it allows customers to make purchases without adding items to a shopping cart. To have the option to buy products through one-click ordering, customers are required to register at the e-commerce website offering the feature and agree it can store their billing and shipping information in the firm's database.

In 1997, Amazon launched its 1-Click checkout technology, which at the time was hailed as an innovation due to its ability to allow visitors to make purchases on the website with a single click. Indeed, by making the checkout fast, easy and simple, the feature revolutionized consumer shopping online and created an important advantage for Amazon in the online marketplace. In September 1999, Amazon secured the patent for 1-Click ordering and protected the invention, after suing Barnes & Noble for implementing a similar ordering technology called “Express Lane.” The patent allowed Amazon to protect the technology from other online retailers and platforms in the US, furthering its lead in the online marketplace. One year later, in September 2000, Apple licensed 1-Click ordering and later incorporated it into iTunes, iPhoto and Apple’s App Store. Since then, Amazon and Apple have been the prominent two companies that offer their customers the ability to make purchases with one click.

Amazon’s U.S. patent for 1-Click buying expired in September 2017. Practically, online retailers and platforms, which have either had to not use one-click buying or pay Amazon licensing fees to do so, have gained the opportunity to incorporate the feature into their businesses. Magento and Shopify are among the first few firms that have taken this opportunity by introducing their versions of one-click checkout and making it available for their retail partners. More recently, one-click ordering has also made inroads into social media platforms (e.g., Instagram).

2.2 Data

We obtained the data for our empirical analysis from a retailer in Asia that prefers to remain anonymous. The retailer we collaborated with specializes on certain product categories, such as shoes and clothes, and sells a wide range of consumer goods at both brick-and-mortar and online channels. The retailer launched a single-step checkout technology, which we refer to as one-click buying, on its e-commerce website on January 1, 2017. The introduction of one-click buying was communicated to online customers through mass emails and on the website, and no specific targeting was involved. The one-click buying feature in our study is the same 1-Click feature that Amazon employs on its website. After joining one-click buying, customers have the option to order products in a single step without adding them into the shopping cart and going through the

checkout process. Importantly, the retailer did not alter the configuration since the feature was introduced and maintained it throughout our data period.

Our data span a period of 32 months, starting from January 2016 to August 2018. It includes a random sample of 977 customers who registered for one-click buying between January 2017 and September 2017 on a voluntary basis without any economic incentives.² These 977 customers maintained their registration throughout the data period and constitute the treatment group in our analysis. For the purpose of comparison, we also obtained a random sample of 17,229 customers who made at least one purchase online over the 12-month period prior to the launch of one-click buying (January 2016 to December 2016) and had yet to adopt the feature as of August 2018.

The data consist of three parts: transaction data of customer purchase (and return) behavior, clickstream data of customer activity and socio-demographic data. The transaction data contain detailed information on each order made by a customer, that is, when a customer purchased a product and how much she paid for it. These data also include information on products returned and the categories of products purchased and returned. Using purchase data, we define a set of outcome measures associated with customer purchases. Because one-click buying applies to the online channel only, unless specified otherwise, these measures are based on online purchases and are constructed at the customer-month level, which is the unit of analysis in this research.

Because our main interest is in assessing how effective one-click buying is in lifting sales, our primary measure is the amount spent by a customer per month.³ In addition, we consider two other (monthly) measures of customer purchases, that is, number of orders made (order frequency) and number of items per order (order size), because the change in purchase amount through one-click buying can arise in multiple ways. For example, one-click buying could lift purchase amount due to the increase in order frequency and/or order size. Order frequency and order size could also change in opposite directions, but the overall change in purchase amount might still be positive.

The second type of data we obtained is clickstream data. They contain detailed individual-

²Because of the non-disclosure agreement we have with the collaborating firm, we are unable to disclose the total number of adopters at the firm.

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

level information on each visit to the website and search activity upon visit, that is, when a customer visited the website, and which pages (and how long) she viewed and searched on the website. Using these micro-level data, we present a set of measures associated with customer activity online, for example, website visits and page views upon visit, to explore possible explanations of our findings. Finally, our data also contain socio-demographic characteristics of customers, for example, age, gender and address, which we utilize to further control for customer heterogeneity.

Figure 1 plots the average purchase amount spent by customers in the treated and control groups over the 32-month observation period. The vertical dotted line in the figure indicates the launch of one-click buying on the retailer’s website (January 2017). The figure offers model-free evidence about the substantial difference in customer purchases between the two groups as well as its persistence over time.

Insert Figure 1 about here

3 Empirical Framework

In this section, we first give an overview of our identification strategy and describe our treated and control groups. We next discuss the details of our econometric approach and the GRF procedure, which we employ to estimate the treatment effect.

3.1 Identification Strategy

Our primary objective is to identify the impact of adopting one-click checkout on subsequent customer behavior and examine heterogeneous treatment effects. We face two major challenges. The first challenge is due to the fact that treatment assignments are not random. An ideal evaluation strategy would start by randomly assigning customers to the one-click buying feature, i.e., intention-to-treat, allowing adopters to make purchases using the feature and tracking subsequent customer behavior (e.g., 12 months). While random assignment would make customers in the intention-to-treat and control groups comparable, it would not be enough for attributing any difference in customer behavior between the two groups to the adoption of one-click buying. Econometric methods would still be necessary for correcting the issue of self-selection into treatment.

In addition to being insufficient for overcoming self-selection, designing and implementing a large-scale field experiment in the course of commercial operations is also challenging because it is subject to the constraints of a partnering firm's requirements and preferences. Moreover, carrying out a long-term experiment for, say 12 months, is difficult because it requires to control for other marketing efforts over the data period. Maintaining the feature of one-click checkout only for adopters in the treated group becomes increasingly difficult over time, unless a huge effort is put into maintaining the integrity of the website. Therefore, we approach this causal-inference problem using observational data and applying a quasi-experimental matching procedure in which we match adopters with similar non-adopters based on propensity scores constructed from observable variables prior to adopting one-click buying.

The second challenge we face is the fact that adoption dates of one-click buying are not fixed among treated customers. Although the launch of one-click buying was an exogenous event that happened on a specific date (i.e., January 1, 2017), customers had the option to register for the feature at any date post launch. Because the date at which a customer was first exposed to one-click buying varied by customer, we apply the staggered-adoption framework (e.g., Bertrand and Mullainathan 2003; Goldfarb and Tucker 2014). Under this framework, treatment adoption by customers occurs at different points in time, and the pre-treatment and post-treatment periods of one-click buying are defined based on the times treated customers registered for one-click buying.

3.2 Treated and Control Groups

Before we establish the effect of one-click buying on customer purchases, as a first step, we assess whether customers in the control group have the same purchase patterns as those in the treatment group prior to adoption. Because we are mainly interested in examining the effect for the long term, we analyze data over a 24-month period and define the first 12 months prior to adoption as the pre-treatment period and the second 12 months as the post-treatment period. If a customer joined one-click buying in September 2017, for example, we analyze her behavior over the 24-month period between September 2016 and August 2018, with the first 12 months (September 2016 to August 2017) as the pre-treatment period and the next 12 months (September 2017 to August

2018) as the post-treatment period. For the non-treated customers, we examine their behavior over the 24-month period between January 2016 and December 2017 because one-click buying was introduced in January 2017. For this control group, we define the 12 months in 2016 as the pre-treatment period and the 12 months in 2017 as the post-treatment period.

Panel A in Table 1 shows that, on average, treated customers spent \$22.19 per month post treatment, whereas those in the control group spent only \$5.29 (diff. = 16.90, p -value < 0.001). Looking at the purchase amount in the pre-treatment period, we see that treated and non-treated customers also had different purchase patterns. On average, treated customers spent \$14.98 per month, whereas non-treated customers spent \$6.21 (diff. = 8.76, p -value < 0.001). Put together, the monthly purchase amount differed considerably between the two groups, and treated customers spent much more than non-treated customers both in the pre-treatment and post-treatment periods. Importantly, customers who spent more were more likely to adopt one-click buying, arguably because they could benefit more from it.

Insert Table 1 about here

We observe similar patterns in two other purchase measures: order frequency and order size. The naïve comparison of the changes in customer purchases suggests one-click buying had a significant impact on purchase behavior among treated customers. However, because customers with different pre-treatment characteristics likely self-selected into the treatment of one-click buying, we would have biased results if we estimated the effect of one-click buying by merely comparing customer purchases between the two groups.

3.3 Propensity Score Matching

The selection on observables approach attempts to mitigate the effect of selection bias and mimic an experimental research design using observational data. The key identifying assumption under this approach is that conditional on observables, the treatment and control groups would differ only on their treatment status and comparing their outcomes of interest would result in unbiased estimates of the treatment effect (Imbens and Rubin 2015).

Matching based on the estimated propensity score is an effective method for applying the selection on observables strategy. We implement matching by first estimating the propensity score, defined as a customer's likelihood of adopting one-click buying using logistic regression and three sets of covariates. The first set of covariates relates to customer-firm relationship which would be associated with the adoption of a new technology (e.g., Bolton et al. 2004; Prins and Verhoef 2007). We use tenure, recency, frequency, monetary (RFM), average amount of discounts received and number of products returned in the pre-treatment period. We also include the pre-treatment 12-month averages of online activity measures: number of visits to the website, number of pages viewed per visit and time duration of each visit.

The second set of covariates relates to psychographic measures that reflect customer preferences or interests that might affect joining one-click buying (e.g. Baumgartner 2002). We operationalize these measures by capturing variety seeking and repeat (or replenishment) behavior in customer purchases, because customers might have different preferences toward or interest in one-click buying depending on their purchase patterns being variety seeking and/or repetitive. For each measure, in particular, we include covariates to capture overall purchase patterns across product categories as well as to reflect purchase patterns within each product category.⁴

The third set of covariates relates to socio-demographics of customers. We include age, gender and address for which we use customers' zip codes and classify them into five regions by income per capita in 2016. Because customers' lifestyles and purchase patterns might differ across regions, these covariates can help control for other unobserved socio-demographics that might affect joining one-click buying, for example, education, income, life style and so on. Table 2 shows the summary statistics of the covariates and describes how the variables are operationalized.

Insert Table 2 about here

To accommodate a flexible specification, we implement several functions of the covariates for inclusion in the logistic regression. We generate interactions and higher-order terms, both

⁴Based on conversations with the retail partner, we decided to have covariates across five product categories which corresponds the way the e-commerce website is organized and measures the business performance.

within and across the three sets of covariates. Following the approach proposed in Imbens and Rubin (2015), we apply an iterative procedure to determine the second-order covariates and the interactions of the main covariates for estimating the propensity score. This approach leads to 28 first-order and 44 second-order covariates and interactions, which altogether constitute 72 variables for estimating the propensity score.

Letting $e(x; \beta)$ denote the model for the propensity score parameterized by β , we obtain pairs of matches by matching on the linearized (estimated) propensity score, i.e., log-odds ratio:

$$l(x; \beta) = \ln \left(\frac{e(x; \beta)}{1 - e(x; \beta)} \right).$$

This transformation linearizes values on the unit interval and can improve estimation (Imbens and Rubin 2015). Consider customer i who joined one-click buying, i.e., received treatment. We ask the question whether, for this customer, there is customer i' in the control group such that the difference (in absolute value) in linearized propensity scores, $l(x_i; \beta) - l(x_{i'}; \beta)$, is less than or equal to a threshold u . In our application, we focus on a threshold of $u = 0.15$ (Stuart 2010), meaning that the difference in propensity scores is approximately less than 15%. The matching algorithm results in 801 unique pairs that are closest to each other in the propensity scores.

The quality of the matching algorithm can be assessed by examining whether the distribution of the propensity scores and the distribution of the covariates are similar across the treated and control groups after matching. Figure 2 shows the density of the estimated propensity scores by treatment status, before and after matching on the propensity scores. Before matching the densities share overlap but vary significantly over the range. Matching balances the densities across treatment status to the extent that little bias seems to remain in the difference of the propensity scores between treated and control groups. We also assess whether matching reduces the mean differences in covariates. Following Austin (2009) and Imbens and Rubin (2015), we examine the standardized differences in covariate means between treated and control groups. Figure 3 presents the standardized differences for each variable used in estimating the propensity scores. Matching

leads to a substantial improvement in balance. After matching, none of the normalized differences exceeds 0.1, a degree of balance comparable to what one might expect in a completely randomized experiment (Stuart 2010; Imbens and Rubin 2015).

Insert Figures 2 and 3 about here

As a further check on the goodness of our propensity score matches, we leverage the panel structure of our data and compare the purchase trends of the treated and control groups before and after matching by estimating models of the following type:

$$\log Y_{it} = \theta_i + \sum_{m=1}^{24} \lambda_t \cdot \mathbb{1}(t = m) + \sum_{m=1}^{24} \beta_m \cdot W_i \times \mathbb{1}(t = m) + \epsilon_{it}, \quad (1)$$

where $\log Y_{it}$ is the dependent variable of customer i at month t with log transformation, W_i indicates whether customer i belongs to the treatment or control group, and the indicator variables $\mathbb{1}(t = m)$ are 1 if month t is m . θ_i is a customer-level fixed effect, λ_t is a month-level fixed effect, and ϵ_{it} is the error term. The two-way fixed-effects specification controls for time-invariant customer characteristics as well as common time trends and month-to-month fluctuations. Our primary interest is the parameter β_m . We normalize the first month (i.e., $m = 1$) in the pre-treatment period as the baseline of 0, and distinguish between the treatment effects over a 24-month period. This way the parameter β_m captures the average difference in purchase measures between treated and non-treated relative to the baseline. To the extent that purchase trends are common prior to one-click buying, estimates of β_m should result in a null effect in the months prior to the treatment, that is, $m \leq 12$. We use robust standard errors clustered at the customer level to account for any serial correlation (Bertrand et al. 2004).

Figure 4 shows the estimates for β_m over the data period before and after matching. The vertical dotted line in this figure indicates the first month post treatment (i.e., month 13). Before matching the estimates are positive and statistically significant throughout the post-treatment period (months 13 to 24). Importantly, the estimates for β_m in the pre-treatment period (months 1 to 12) are also statistically significant except for a few months. As shown in the bottom panel of Fig-

ure 4, matching renders the two groups indistinguishable in terms of their pre-treatment purchase behavior. None of the coefficient estimates in the pre-treatment period are statistically different from zero, whereas all estimates are positive and statistically significant in the post-treatment period. This is further evidence that our matching resulted in a valid control group for the treated.

Insert Figure 4 about here

Panel B in Table 1 shows the summary statistics for the matched sample. The difference-in-differences between the treated and control are substantially lower for the matched sample compared to the unmatched sample, which is shown in Panel A in Table 1. Taken together, our selection on observables strategy through propensity score matching ameliorates the self-selection concern.

3.4 Generalized Random Forests

As discussed earlier, we are interested not only in estimating the average treatment effect but also in analyzing the heterogeneity of the estimates. We employ a recently developed machine-learning based procedure, called GRF (Athey et al. 2019). The procedure is a nonparametric statistical estimation method for causal inference in observational studies and provides a general framework to obtain unbiased estimates of the average treatment effects and also to capture heterogeneity in parameters of interest. GRF is an extension of the causal forest method (Wager and Athey 2018), which is based on the classic random forests algorithm used for statistical learning (Breiman 2001). As a forest-based method for treatment-effect estimation, causal forests apply the same general training and prediction framework used in building random forests, such as resampling, recursive partitioning and averaging across many trees.

What distinguishes GRF from classic random forests is that the splitting criteria for growing individual trees are specifically designed to find partitions where treatment effects most differ. That way, GRF provide a data-driven procedure for selecting the dimensions that are most important for capturing heterogeneity in the treatment effects. Compared to conventional model specifications with interaction terms to capture nonlinear relationships, the method flexibly accommodates complex interplay in data and estimates the parameters of interest without making assumptions on the

functional form.

Built on the same general framework as causal forests, GRF critically rely on sample splitting, which is referred to as the honesty condition, and uses different subsamples of the data for growing trees and making predictions at the leaves of the trees. However, the method has an important additional feature that is designed to improve the performance of causal forests. Specifically, instead of obtaining treatment-effect estimates at the tree level for each test example, it creates a list of neighboring training examples and records the frequency by which the test example and each training example share the same leaf in the trees built during training. Based on this information, it assigns (similarity) weights to each neighboring training example and together with their treatment status and outcomes uses them to make predictions for the test example.

To make the procedure concrete, suppose we have data on n independent and identically distributed samples $(Y_i, X_i, W_i), i = 1, \dots, n$, where Y_i is the observed outcome of interest for customer i , X_i denotes the set of covariates that were measured before the treatment, and W_i is the binary treatment indicator. Using the potential-outcomes framework, we are interested in estimating the conditional average treatment effect (CATE):

$$\tau(x) = \mathbb{E}[Y_i(1) - Y_i(0)|X_i = x], \quad (2)$$

where $Y_i(1)$ and $Y_i(0)$ correspond to the outcome we would have observed had we assigned treatment or control to customer i . CATE is identified under the assumption that potential outcomes and treatment assignment are independent conditional on the observed covariates, $\{Y_i(1), Y_i(0)\} \perp\!\!\!\perp W_i|X_i$, and the propensity scores of the treated and non-treated customers overlap at every point in the covariate space, $0 < P(W = 1|X = x) < 1$.

The forest-based algorithm as implemented in GRF estimates $\tau(x)$ by first creating similarity weights $\alpha_i(x)$ that measure the relevance of the i -th training example to fitting $\tau(\cdot)$ at x : $\alpha_i(x) = \frac{1}{K} \sum_{k=1}^K \alpha_{ik}(x)$, where $\alpha_{ik}(x) = \frac{\mathbf{1}(X_i \in L_k(x))}{|L_k(x)|}$, $L_k(x)$ is the set of training examples that fall in the same leaf as the test example x in the k -th tree and K is the number of trees. The similarity

weights add up to 1 and capture the forest-based adaptive neighborhood of x . They are applied to fit $\tau(x)$ using the empirical version of the moment condition:

$$E\left[\alpha_i(x) \cdot \left((Y_i - W_i \cdot \tau(x) - \mu(x))(1 - W_i^T)^T\right) | X_i = x\right] = 0,$$

where the intercept $\mu(x)$ is a nuisance parameter. The resulting estimator yields personalized treatment-effect estimates that are shown to be consistent for the true effect and are asymptotically Gaussian (Athey et al. 2019):

$$\hat{\tau}(x) = \frac{\sum_{i=1}^n \alpha_i(x) \cdot (Y_i - \bar{Y}_\alpha) \cdot (W_i - \bar{W}_\alpha)}{\sum_{i=1}^n \alpha_i(x) \cdot (W_i - \bar{W}_\alpha) \cdot (W_i - \bar{W}_\alpha)^T}, \quad (3)$$

where $\bar{W}_\alpha = \sum_{i=1}^n \alpha_i(x) W_i$ and $\bar{Y}_\alpha = \sum_{i=1}^n \alpha_i(x) Y_i$. The procedure then uses the personalized estimates and implements the augmented inverse propensity weighting scheme to obtain doubly-robust scores (Robins et al. 1994) and yields doubly robust average treatment effect estimates by averaging the estimated scores.

In this research, we apply the GRF procedure to the change in the outcomes between the post-treatment and pre-treatment periods, instead of using the outcome itself. Hence, we effectively combine GRF with the difference-in-differences estimation on the propensity score matched sample. This combination of estimation methods has the additional benefit of accounting for time-invariant unobservables that might still persist even after matching (Smith and Todd 2005) and thus enhances the robustness of our inference. We define the outcome (Y_i) as the log-transformed change in purchase behavior for customer i as follows:

$$Y_i = \frac{1}{|T_A|} \sum_{t \in T_A} \log Y_{it} - \frac{1}{|T_B|} \sum_{t \in T_B} \log Y_{it}, \quad (4)$$

where Y_{it} denotes the outcome variable for customer i at month t , and T_B and T_A denote the pre-treatment and post-treatment periods, respectively.

We next describe implementation details. For the outcome variables, we analyze transaction

data over a 24-month period under the staggered-adoption strategy. Of these, the first 12 months are prior to the treatment. The covariates that we are interested in finding the heterogeneity of the treatment effect are a subset of the covariates we used in estimating the propensity scores. We include tenure, recency, frequency, monetary and discount from the customer-firm relationship set. In addition, we include the number of visits to the website during the last month in the pre-treatment period. From the set of psychographic variables we include variety and repeat purchase measures across five product categories. Finally, we include demographics to understand to what extent treatment effects vary by age and gender.⁵

4 Findings

Using the GRF procedure, we trained causal forests by setting the number of trees (K) to 10,000 and obtained estimates of personalized treatment effects as well as estimates of average treatment effects on the treated. We first discuss the main findings of the treatment effect on customer purchases. We next describe the heterogeneity in the effect across customers.

4.1 Average Treatment Effects

We report the average treatment effect on customer purchases in Table 3. Because our primary objective is to identify the causal effect of one-click buying on customer purchases for the long term, we first discuss the estimates over a 12-month period. Note that we used the natural log transformations of the outcomes in our estimations.⁶ Column 1 shows that, on average, one-click buying resulted in a 17.4% ($= \exp(0.16) - 1$) increase in purchase amount among treated customers, with the estimate being statistically significant (p -value < 0.005).

Insert Table 3 about here

To investigate how one-click buying affected order frequency and order size, we estimate the GRF model with the difference in each of these two measures as an outcome measure using the same time window. As shown in columns 2 and 3 in Table 3, we find the number of orders made

⁵We obtained estimates by standardizing all continuous covariates.

⁶To deal with 0 in the original scale, we added a constant of 1 before taking the natural log transformation.

by treated customers and the number of items per order increased by an average of 5.1% and 6.2%, respectively, post treatment. Both estimates are statistically significant (p -value < 0.005).

We are also interested in examining how the treatment effect on customer purchases changes over time because the impact of one-click buying could be short lived or long lasting. To that end, we investigate how customer purchases changed within the first three and six months post treatment relative to the 12-month pre-treatment period by applying the GRF procedure. Our findings for both time windows suggest a novelty effect of one-click buying insofar as the impact is larger in the initial months post treatment. Specifically, the average increase in purchase amount is 29.7% ($= \exp(0.26) - 1$) and 23.4% ($= \exp(0.21) - 1$) within the first three and six months post treatment, respectively. The novelty effect also holds true with respect to order frequency and order size. The former increased by 8.3% and 5.1% within the first three and six months post treatment, respectively, while latter increased by 9.4% and 7.3%, respectively. All estimates are statistically significant (p -value < 0.005).

It is also useful to take a closer look and examine whether the increase in purchase amount, as a result of one-click buying, is uniform or differential across product categories. To that end, we divide the purchase amount into spending across product categories and perform our analysis separately with outcomes defined as the changes in purchase amount in each of the product categories. We find there is a large variation in the treatment effect across product categories, ranging from 3% in category 3 (not significant) to 23.4% in category 1, primarily because customers did not spend uniformly across categories.

Taken together, our findings of the average treatment effects provide evidence that one-click buying is effective in lifting sales and does so by making treated customers purchase more often with more items per order.⁷ Furthermore, the effect on customer purchases is economically significant and persistent over a time window of 12 months in the post-treatment period.

⁷As our primary interest in this study is to identify the causal effect of adopting one-click feature on purchase behavior, rather than to investigate how treated customers used one-click feature post adoption, we chose to measure it with customer purchases constructed at the customer-month level. We find that after joining one-click buying, treated customers used the feature in 51% of their subsequent orders, showing that adopters used one-click often post adoption. Single-item orders made up 56% and 53% of the orders placed with one-click and without one-click, respectively. This suggests that purchase behavior among adopters did not differ based on the usage of one-click buying.

4.2 Heterogeneous Treatment Effects

Using the personalized treatment effects obtained through the GRF procedure, we examine the heterogeneity of the treatment effects using the 12-month window post treatment.

Figure 5 shows the distribution of the treatment-effect estimates in purchase amount. We find there is significant variation in the changes in purchase amount across customers. Although the effect is positive for almost all (791 out of 801) customers, the magnitude of the increase in purchase amount varies considerably, ranging from a mere 0.01% to 30% post treatment. About 27% of the customers increased their purchase amount by 20% or less and approximately 8% increased their purchase amount by 25% or more. A majority of the treated increased their purchases between 20% and 25%. This result illustrates the benefits of obtaining individual-specific treatment-effect estimates by applying the GRF procedure.

Insert Figure 5 about here

Since we observe large variation in the treatment-effect estimates, we examine the source of the heterogeneity because understanding such variation is important both theoretically and managerially. A useful feature that is embedded in GRF is the importance measure of covariates used in the estimation. Because GRF build trees during training by splitting on covariates where treatment effects most differ, the importance measure of covariates reflects the relative weight of each covariate in generating the splits. The measure is computed by taking the frequency for which each covariate is used for splitting the nodes and weighting them by the depth of each tree.

Table 4 shows the importance weight of each covariate as well as its rank among all covariates used in our estimation. The causal forest spent about 75% of its splits using the covariates that reflect the customer-firm relationship. In particular, RFM measures, which summarize a customer's past transaction history, were among the most important variables and together accounted for about 40% of the splits. This result is in line with the findings in the literature that RFM measures could be moderators of various marketing activities (e.g., Rossi et al. 1996; Kumar and Shah 2004). Interestingly, website visit, which is operationalized as the number of visits to the website

in the last month before treatment, was important and accounted for about 11% of the splits. This variable is similar to RFM measures because it reflects customer engagement with the retailer.

Insert Table 4 about here

The causal forest also spent over 15% of its splits based on the variables that measure the extent of variety and repeat behavior in customer purchases. These variables reflect customer preferences or interests and are less precise and more nuanced because they are not directly observable (e.g., Baumgartner 2002). Nevertheless, variety and repeat measures in customer purchases contributed to the identification of the heterogeneity in our context. Socio-demographic variables accounted for less than 9% of the splits in our estimation.

Now that we have identified the source of the heterogeneity in the treatment effect, we take a more granular look and seek to identify subgroups of customers for whom one-click buying might have a stronger impact on subsequent purchases. To that end, we relate heterogeneous treatment effects to observed covariates and divide the customers into two equal-sized groups based on the personalized treatment-effect estimates (i.e., below and above the median). We then compute the means of the covariates across the two groups.

Figure 6 shows how the group with high treatment effect differs from the group with low treatment effect in terms of the covariates used in the analysis. The horizontal axis in the figure refers to the mean of the standardized value of each covariate. The figure presents a clear pattern in the changes in purchase behavior with respect to the RFM measures: Customers who purchased more recently, more frequently and spent more in the pre-treatment period had a smaller increase in purchases after adopting one-click buying. Similarly, customers who had more discounts and visited the online store more often in the last month prior to joining one-click had a smaller treatment effect. On the other hand, customers who had a longer tenure with the firm had a larger treatment effect. In terms of customer-firm relationship, this suggests that joining one-click buying had a stronger effect on those customers who were familiar with the firm but were not necessarily loyalists. Put differently, the new feature had a limited impact on customers who already had a strong relationship with the firm, but had a larger impact on those customers who previously purchased

less in terms of amount and frequency. This is in line with previous work that report a diminished impact of an experience-centric channel on customers who had more experiences with the firm prior to receiving the treatment (e.g., Bell et al. 2020).

Insert Figure 6 about here

With respect to the psychographic measures, we find overall repeat purchase behavior is similar between groups, whereas overall variety-seeking behavior is higher in the group of low treatment effect. Together these findings suggest that compared to customers who like to explore and try new products, those customers who purchased repeatedly increased their purchases more due to one-click buying. We also find slight variation in customer socio-demographic characteristics between the high and low treatment effect groups. As expected, younger customers were more responsive to one-click buying than older customers. Among all of the covariates used, gender had the least variation in the treatment-effect estimates between the two groups.

To summarize, we find the impact of one-click buying on customer purchases is heterogeneous across customers. Our findings on heterogeneous treatment effects can assist marketers in scoring and targeting customers and allocating resources across individual customers when designing marketing activity associated with one-click buying.

5 Potential Mechanisms

Now that we have established the impact of one-click buying on customer behavior, in this section, we explore some possible explanations underlying the effects.

5.1 Channel Switching

Our first explanation is based on the organizational structure of the partner retailer. The firm we partnered with has both a 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 therefore investigate whether the increase in online purchases through one-click buying was due to the channel-switching behavior to online from offline.

To examine whether online and offline channels are complements or substitutes (e.g., Forman et al. 2009; Wang and Goldfarb 2017), we perform the GRF analysis described in §3.4 by replacing online purchases with offline purchases and analyze the 12-month post-treatment period. We retain the operationalization of all covariates from our main analysis of online purchase measures. Table 5 presents the treatment effects on offline purchases. As shown in columns 1-3, none of the estimates in all outcome measures are statistically significant. These results suggest the increase in online purchases through one-click buying did not come at the cost of offline purchases. Therefore, the one-click checkout technology was effective in lifting overall sales for the firm in our context.

Insert Table 5 about here

5.2 Impulse Buying

Another explanation that might exist between one-click buying and purchase outcomes is impulse buying, a purchase that is unplanned, the result of an exposure to a stimulus and decided on the spot (e.g., Piron 1991). In his seminal work, Stern (1962) argues impulse buying is related to ease of buying to the extent that a purchase involves less of one's resources (e.g., money, time, effort). Impulse buying can be triggered by a variety of stimuli ranging from the product itself and its attributes to the environment in which the purchasing event takes place (e.g., Hui et al. 2013). Studying impulse buying in online settings, Parboteeah et al. (2009), for example, find a web interface with high-quality task-relevant characteristics, which make shopping efficient and effective, increases the likelihood of customers to buy impulsively. Given that one-click buying adds to an online store's task-relevant features, it might lead to increased impulse purchases.

An additional way in which one-click buying could lead to impulse buying is by hampering customers' mental accounting (Thaler 1985). Customers exert more time and effort for the act of buying when multiple steps are involved in finalizing a purchase. While making buying more burdensome, multiple-step ordering as opposed to one-click buying could provide ample opportunities for customers to deliberate their choice and register the cost of the transaction in their mind. However, when the ordering process is reduced to a single step, the convenience of shopping can

override the cognitive processes that keep track of spending, and the losses involved in the purchase might not be apparent, thereby increasing impulse buying. Dutta et al. (2003) show the number of steps involved in the payment process has a significant impact on the subjects' recall of past expenses. Compared to subjects in the multiple-step payment condition, those under the one-step payment condition have lower recall of past expenses, which in turn leads to impulse buying.

We examine whether impulse-buying behavior contributed to increased purchases after adopting one-click buying. To evaluate this explanation, we utilize data on product returns at the online channel and use product returns as a proxy for impulse purchases, because impulse buying could lead to more product returns. In fact, Ridgway et al. (2008) study impulse buying as a dimension of compulsive buying and report that it causes behaviors such as hiding and making frequent returns of purchased items due to feelings of remorse or guilt. Therefore, we construct two measures of impulse-buying (number of products returned and proportion of products returned among products purchased) to assess impulse buying-behavior as an explanation in our data. We examine both the number and proportion of products returned to provide a more comprehensive picture of the results.

As a test for impulse-buying behavior, we conduct GRF analysis, in which we use each measure of product returns as the dependent variable and retain all covariates from our main analysis. Columns 4-5 in Table 5 show the results. When analyzing the 12-month post-treatment period, we observe one-click buying had a marginal impact on customer behavior associated with product returns. In particular, one-click buying increased the number of products returned by 1.4% and the proportion of products returned by 0.8%. As the increases in return measures are very small relative to the increases in purchase measures, we interpret these results as somewhat limited evidence for the impulse-buying explanation for the increase in purchases in our context.

5.3 Customer Engagement

As another potential explanation, we propose to examine the impact of one-click buying on customer experience that materializes as customers interact with the retailer. Each of the customer-retailer interactions could evoke customer responses across multiple dimensions (e.g., cognitive,

emotional), which then influence customer satisfaction and engagement (e.g., Verhoef et al. 2009).

Among the three stages in the customer's journey with the firm, namely, pre-purchase, purchase and post-purchase, the purchase stage is key in shaping customer experience because it is related to a variety of behaviors, such as choice, ordering and payment (e.g., Lemon and Verhoef 2016). Online retailers, equipped with useful features such as online decision aids, can facilitate and enable customers to attain their shopping goals. Customers exert less effort and yet make better decisions when online decision aids are available (Häubl and Trifts 2000). Moreover, usage of online decision aids is associated with increased satisfaction and the likelihood of repeat use of the website (e.g., Palmer 2002; Parboteeah et al. 2009). Similar to online decision aids, one-click buying is an online feature that enhances customer experience by reducing the time and effort during checkout and enhances the usability of an online store (e.g., Dutta et al. 2003; Parboteeah et al. 2009). This increased convenience could connect to a better shopping experience online, which increases customer engagement for the retailer.

Richer engagement could manifest itself through customers' tendency to make more frequent visits to the retailer's online store and spend more time upon visits. Providing evidence for this proposed explanation of customer engagement requires clickstream data of customer activity online from our retail partner for all customers in our study over the same period of 32 months (between January 2016 and August 2018), which we obtain and complement with our purchase data. Using these micro-level data, we construct three measures of online activity, that is, visit instances to the website, number of page views per visit and duration (minutes) per visit on the retailer's website. We also construct a measure of category expansion by computing the number of product categories a customer made purchases from, because customer engagement could also be reflected in customers' tendency to expand their purchases across product categories. To test our suggested explanation, we again employ the GRF procedure on the matched sample. We use the change in each of the measures on online activity and category expansion as the dependent variable. In each estimation, we retain all covariates used in our main analysis.

Table 6 shows the results. We find treated customers increased website visits by 18.5%

(= $\exp(0.17) - 1$), viewed more pages (an increase of 7.3%) and spent more time (an increase of 9.4%) upon visit in the 12-month period post adoption. We also find treated customers deepened their relationship with the firm by expanding their purchases across categories (an increase of 43.3% (= $\exp(0.36) - 1$)) in the 12-month post-treatment period.

Insert Table 6 about here

In summary, our results on customer engagement are largely consistent with earlier findings that customers become more engaged with the focal retailer as their interactions with purchase touch points (e.g., decision aids) improve (e.g., Shi and Zhang 2014). We provide evidence supporting this mechanism by demonstrating that the increase in purchases was to a large extent due to improved customer experience which enhanced customer engagement with the retailer.

6 Robustness Checks

In this section we evaluate the robustness of our findings. In §6.1 and §6.2 we address potential effects of unobservables. In §6.3 we employ alternate outcome measures. Finally, in §6.4 we show the results from an alternate specification.

6.1 Selection on Unobservables

In our matching procedure, we used a rich set of covariates based on individual-level data to control for purchase and other behavior with the firm prior to one-click buying. Furthermore, we employed our estimation procedure on the changes of outcome measures, which effectively eliminates individual-specific time-invariant unobservables. However, if the treated customers were systematically different from the non-treated customers on various time-dependent aspects that are unobservable to us but are relevant for our outcome measures, our estimated treatment effects may be biased. Therefore, we revise our identification strategy against this possibility and present an additional piece of evidence to alleviate the concern for the unconfoundedness assumption.

We obtain the treatment-effect estimates using the treated only. Specifically, we divide treated customers based on their adoption dates and use late adopters, rather than non-adopters, as controls for early adopters in our estimations (e.g., Goldfarb and Tucker 2011; Manchanda et al.

2015; Datta et al. 2017; Narang and Shankar 2019). We define early adopters as treated customers who joined one-click buying in the first two quarters of 2017 and late adopters as those who joined in the third quarter of 2017.⁸

We find the results are consistent with those in our main analysis. Specifically, treated customers who joined in the first two quarters of 2017 increased their purchases by 28.4% (p -value < 0.005) in the three months post treatment, as compared to a 29.7% increase in our main analysis. Similarly, we find both order frequency and size increased by an average of 8.3% in the three months post treatment. Both estimates are very similar from our main analysis and statistically significant (p -value < 0.005). As these results obtained from a within-treated analysis are similar to those from our main analysis, they provide further support that our matching procedure created a valid control group for the treated and does not unduly suffer from selection on unobservables.

6.2 Rosenbaum Bounds

We assess the sensitivity of positive and significant treatment effects to potential unobservable effects by implementing Rosenbaum bounds test (Rosenbaum 2002). The underlying idea behind this test is to evaluate how much the change in odds ratio of joining one-click buying, due to unobservables, would be required to nullify the treatment effect identified by propensity score method. Define Γ as the level of hidden bias caused by unobservables in the odds ratio of differential treatment assignment. If $\Gamma=1$, customers matched on observables are indeed the same and they have equal chances of treatment. If $\Gamma > 1$, however, unobservables affect treatment assignment and the propensity score for treated customers to select into one-click buying is higher than for the control customers conditional on observables.

We use different values of Γ to examine how inferences on the Hodges-Lehmann's estimates (Rosenbaum 1993) of one-click buying might change if hidden bias were present. Table 7

⁸We excluded the June and July adopters from the treated and controls. By limiting the adopters from January to May and the non-adopters from August to September, we can ensure the 3-month post-treatment period among adopters does not overlap with the time during which the non-adopters were actually treated. However, we cannot obtain the treatment-effect estimates for the 12-month period because late adopters in this case were exposed to the treatment in the third quarter of 2017. Nevertheless, the analysis using 494 treated and 269 control customers, is useful in gauging the robustness of our results with regards to unobservables.

reports the p -values from Wilcoxon signed-rank tests for the treatment effect on the treated while setting the level of hidden bias to a certain value of Γ . It also reports both lower and upper bounds of the Hodges-Lehmann point estimates. We find the estimated treatment effect on customer purchases is positive and significant for $\Gamma \leq 1.7$.⁹ It means that for a positive estimated treatment effect to be nullified, the potential unobserved factors affecting treatment assignment would have to be large enough to increase odds ratio of joining one-click buying by at least 80%. Our GRF estimate is within the range of the Rosenbaum bounds in the values of Γ . Hence, our conclusion that one-click buying leads to increased purchases is unlikely due to unobserved selection effects.

Insert Table 7 about here

6.3 Alternate Outcomes

To ease interpretation and reduce skewness, we used the log transformations of the outcome measures in our estimations. We assess the robustness of our findings by performing our analysis without employing this transformation to the dependent variables. Table 8 shows that all estimates have the same sign and similar statistical significance as the estimates from our main analysis, suggesting that our findings are robust in terms of functional forms of the outcome measures.

Insert Table 8 about here

As another robustness check for alternate outcomes, we perform our analysis by excluding the month in which customers joined one-click buying. Customers could adopt one-click buying with our retail partner at any time during their visits to the website. If customers adopted it during the checkout process, they would have additional purchases even if their purchase behavior did not change significantly post adoption. We thus investigate to what extent our findings are sensitive to this possibility. We find the treatment-effect estimates are 17%, 5.1% and 6.2% for purchase amount, order frequency and order size, respectively, which are consistent with our main findings.

⁹The estimated treatment effect for the number of orders made and the number of items per order is positive and significant for $\Gamma \leq 1.6$. The Γ we find is on the same order of magnitude as the Rosenbaum bound results reported by previous work, e.g., DiPrete and Gangl (2004), Sun and Zhu (2013), Manchanda et al. (2015).

This suggest that the treatment effect is not an artifact of the purchases made during the adoption stage, but rather a change in subsequent customer behavior.

6.4 Alternate Specification

Employing a non-parametric statistical method, such as GRF, for estimating the treatment effects allowed us to remain agnostic about the true data-generating process. We report additional results using panel regressions that rely on the assumption that treatment effects are linear, constant and additive. Although these are strong constraints on how customers responded to treatment, the estimates serve as a useful baseline. We modify Equation (1) and obtain treatment-effect estimates based on the following model for each of the outcome measures:

$$\log Y_{it} = \theta_i + \lambda_t + \beta_{3m} \cdot W_i \times \mathbb{1}(t \leq 3) + \beta_{6m} \cdot W_i \times \mathbb{1}(4 \leq t \leq 6) + \beta_{12m} \cdot W_i \times \mathbb{1}(7 \leq t \leq 12) + \epsilon_{it}, \quad (5)$$

where the indicator variable $\mathbb{1}(\cdot)$ is 1 if it satisfies the condition (\cdot). This way we obtain treatment effects in the short term (within the first three months of adoption, β_{3m}), medium term (between four and six months after adoption, β_{6m}), and long term (between seven and twelve months after adoption, β_{12m}). We use robust standard errors clustered at the customer level to account for any serial correlation. The results, as shown in Table 9, demonstrate the robustness of positive and significant treatment effects of one-click buying across all of outcome measures.

7 Conclusions

The end of Amazon's hold on the one-click checkout technology offers opportunities to retailers, and naturally leads to the question about the economic value of one-click buying for an online retailer. Utilizing quasi-experimental data over a 32-month period at a company that launched the one-click buying feature on its website, we measure the causal effect of one-click buying on customer behavior post adoption. Our identification is based on selection on observables whereby we create matched pairs and obtain treatment-effect estimates at the individual level via the recently developed GRF procedure.

We find one-click buying is effective in lifting sales and does so by making treated customers purchase more often with more items per order. The impact of one-click buying on customer purchases post adoption is economically significant, persistent over time and heterogeneous across customers. Our findings are robust to potential confounding effects of self selection and unobservables, different treated and control groups and different outcomes of purchase behavior.

To uncover the underlying mechanisms that lead to the behavioral changes, we leverage a variety of data including offline purchases, product returns, customer activity online and purchases across categories. Using these data, we explore channel-switching behavior, impulse-buying behavior and customer engagement as possible explanations. Although multiple drivers can be at work, we suggest one-click ordering provides an avenue for the online retailer to improve customer engagement, which eventually leads to the positive changes in purchase behavior. We provide evidence consistent with this explanation based on micro-level data on customer activity online and purchase behavior across categories. We find no evidence for channel substitution and weak evidence for impulse buying.

Our findings offer three important managerial implications. First, this research presents evidence on the economic value of the one-click buying feature for e-commerce companies. Because an increase in purchases as well as customer activity is expected for an e-commerce company after launching the technology, having management solutions in place to effectively and efficiently manage traffic to the website and fulfill orders so that customers have positive experience is essential. Along these lines, it is also critical to improve product assortments (both breadth and depth) for cross-selling opportunities because one-click buying motivates customers to expand categories for their purchases. Second, because companies are increasingly concerned about customer engagement, our findings illustrate the importance of investments and efforts toward improving customer experience through advanced features in e-commerce (e.g., seamless checkout process). Third, this study provides evidence on the heterogeneous treatment effects of one-click buying that can assist managers in scoring and targeting customers and also in allocating resources across individual customers when designing marketing activity associated with one-click buying.

Because this research is the first attempt to quantify the effect of one-click buying at a retailer across a variety of behavioral measures, a number of limitations should be acknowledged and perhaps addressed in future research. First, given that our context lacks random assignment into treatment and control groups, our identification strategy hinges on the unconfoundedness assumption that treatment is determined only by observed characteristics. We believe that unconfoundedness is a reasonable assumption in our context because we have included a rich set of observables. However, we are unable to rule out with certainty any bias that might result from self selection. Second, our study focused on a single retailer in a specific industry. Therefore, replication across other firms and industries would be needed to build empirical generalizations on this topic. With that in mind, we hope our approach provides a framework for further studies. Third, customers who adopt one-click buying feature could change their behavior related with, for example, shopping cart (e.g., add items to cart, abandon cart). Studying how treated customers use one-click feature post adoption would be interesting. Finally, we are unable to study how customer behavior at the focal retailer may change when competitors introduce their versions of one-click checkout technology. With competition in play, the effect of one-click buying on customer behavior remains unclear. We hope our work will generate further interest in expanding our understanding of the impact of technology in a growing e-commerce landscape.

References

- Ascarza, Eva. 2018. Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research* 55(1):80–98.
- Athey, Susan, Julie Tibshirani, Stefan Wager. 2019. Generalized random forests. *Annals of Statistics* 47(2):1148–1178.
- Austin, Peter C. 2009. Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples. *Statistics in Medicine* 28(25):3083–3107.
- Baumgartner, Hans. 2002. Toward a personology of the consumer. *Journal of Consumer Research* 29(2):286–292.
- Baymard Institute. 2017. E-commerce checkout usability report. <http://baymard.com/checkout-usability>.
- Baymard Institute. 2019. The average checkout flow has 14.88 form fields – twice as many as necessary. <http://baymard.com/blog/checkout-flow-average-form-fields>.
- Bell, David R, Santiago Gallino, Antonio Moreno. 2020. Customer supercharging in experience-centric channels. *Management Science* 66(9):4096–4107.
- Bertrand, Marianne, Esther Duflo, Sendhil Mullainathan. 2004. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119(1):249–275.
- Bertrand, Marianne, Sendhil Mullainathan. 2003. Enjoying the quiet life? Corporate governance and managerial preferences. *Journal of Political Economy* 111(5):1043–1075.
- Bolton, Ruth N, Katherine N Lemon, 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):271–292.
- Breiman, Leo. 2001. Random forests. *Machine Learning* 45(1):5–32.
- CNET. 2017. The death of amazon’s 1-click patent could save shoppers. <https://www.cnet.com/news/the-death-of-amazons-1-click-patent-could-save-shoppers/>.
- Datta, Hannes, George Knox, Bart J Bronnenberg. 2017. Changing their tune: How consumers’ adoption of online streaming affects music consumption and discovery. *Marketing Science* 37(1):5–21.
- Digiday. 2017. End of an era: Amazon’s 1-click buying patent finally expires. <https://digiday.com/marketing/end-era-amazons-one-click-buying-patent-finally-expires/>.
- DiPrete, Thomas A, Markus Gangl. 2004. 7. assessing bias in the estimation of causal effects: Rosenbaum bounds on matching estimators and instrumental variables estimation with imperfect instruments. *Sociological Methodology* 34(1):271–310.

- Dutta, Ranjan, Sirkka Jarvenpaa, Kerem Tomak. 2003. Impact of feedback and usability of online payment processes on consumer decision making. *ICIS 2003 Proceedings 2*.
- eMarketer. 2019. US ecommerce 2019. <http://www.emarketer.com/content/us-ecommerce-2019>.
- eMarketer. 2020. US ecommerce 2020. <http://www.emarketer.com/content/us-ecommerce-2020>.
- Fong, Nathan, Yuchi Zhang, Xueming Luo, Xiaoyi Wang. 2019. Targeted promotions on an e-book platform: Crowding out, heterogeneity, and opportunity costs. *Journal of Marketing Research* 56(2):310–323.
- Forman, Chris, Anindya Ghose, 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.
- Gallino, Santiago, Antonio Moreno. 2018. The value of fit information in online retail: Evidence from a randomized field experiment. *Manufacturing & Service Operations Management* 20(4):767–787.
- Goldfarb, Avi, Catherine E Tucker. 2011. Privacy regulation and online advertising. *Management Science* 57(1):57–71.
- Goldfarb, Avi, Catherine E Tucker. 2014. Conducting research with quasi-experiments: A guide for marketers. Working paper, University of Toronto.
- Häubl, Gerald, Valerie Trifts. 2000. Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science* 19(1):4–21.
- Hui, Sam K, J Jeffrey Inman, Yanliu Huang, Jacob Suher. 2013. The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies. *Journal of Marketing* 77(2):1–16.
- Imbens, Guido W, Donald B Rubin. 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, UK.
- Kumar, Viswanathan, Denish Shah. 2004. Building and sustaining profitable customer loyalty for the 21st century. *Journal of Retailing* 80(4):317–329.
- Lemon, Katherine N, Peter C Verhoef. 2016. Understanding customer experience throughout the customer journey. *Journal of Marketing* 80(6):69–96.
- Lindner, Matt. 2017. E-commerce platforms rush to develop their own one-click ordering tools. *Digital Commerce 360*, <https://www.digitalcommerce360.com/2017/09/22/e-commerce-platforms-rush-develop-one-click-ordering-tools-amazons-patent-expires/>.
- Magento. 2017. New instant purchase checkout boosts sales. <https://magento.com/blog/magento-news/new-instant-purchase-checkout-boosts-sales>.

- Manchanda, Puneet, Grant Packard, Adithya Pattabhiramaiah. 2015. Social dollars: The economic impact of customer participation in a firm-sponsored online customer community. *Marketing Science* 34(3):367–387.
- Narang, Unnati, Venkatesh Shankar. 2019. Mobile app introduction and online and offline purchases and product returns. *Marketing Science* 38(5):756–772.
- Palmer, Jonathan W. 2002. Web site usability, design, and performance metrics. *Information Systems Research* 13(2):151–167.
- Parboteeah, D Veena, Joseph S Valacich, John D Wells. 2009. The influence of website characteristics on a consumer's urge to buy impulsively. *Information Systems Research* 20(1):60–78.
- Piron, Francis. 1991. Defining impulse purchasing. *Advances in Consumer Research* 18:509–514.
- Prins, Remco, Peter C. Verhoef. 2007. Marketing communication drivers of adoption timing of a new e-service among existing customers. *Journal of Marketing* 71(2):169–183.
- Rejoinder. 2017. How valuable is amazon's 1-click patent? It's worth billions. <http://rejoinder.com/resources/amazon-1clickpatent>.
- Ridgway, Nancy M, Monika Kukar-Kinney, Kent B Monroe. 2008. An expanded conceptualization and a new measure of compulsive buying. *Journal of Consumer Research* 35(4):622–639.
- Robins, James M, Andrea Rotnitzky, Lue Ping Zhao. 1994. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* 89(427):846–866.
- Rosenbaum, Paul R. 1993. Hodges-lehmann point estimates of treatment effect in observational studies. *Journal of the American Statistical Association* 88(424):1250–1253.
- Rosenbaum, Paul R. 2002. Sensitivity to hidden bias. *Observational studies*. Springer, 105–170.
- Rossi, Peter E, Robert E McCulloch, Greg M Allenby. 1996. The value of purchase history data in target marketing. *Marketing Science* 15(4):321–340.
- Shankar, Venkatesh. 2018. How artificial intelligence (ai) is reshaping retailing. *Journal of Retailing* 94(4):6–11.
- Shi, Savannah Wei, Jie Zhang. 2014. Usage experience with decision aids and evolution of online purchase behavior. *Marketing Science* 33(6):871–882.
- Smith, Jeffrey A, Petra E Todd. 2005. Does matching overcome lalonde's critique of nonexperimental estimators? *Journal of Econometrics* 125(1–2):305–353.
- Stern, Hawkins. 1962. The significance of impulse buying today. *Journal of Marketing* 26(2):59–62.
- Stuart, Elizabeth A. 2010. Matching methods for causal inference: A review and a look forward. *Statistical Science* 25(1):1–21.

- Sun, Monic, Feng Zhu. 2013. Ad revenue and content commercialization: Evidence from blogs. *Management Science* 59(10):2314–2331.
- Thaler, Richard. 1985. Mental accounting and consumer choice. *Marketing Science* 4(3):199–214.
- Verhoef, Peter C, Katherine N Lemon, Ananthanarayanan Parasuraman, Anne Roggeveen, Michael Tsiros, Leonard A Schlesinger. 2009. Customer experience creation: Determinants, dynamics and management strategies. *Journal of Retailing* 85(1):31–41.
- Wager, Stefan, Susan Athey. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523):1228–1242.
- Wagner, R Polk, Thomas Jeitschko. 2017. Why amazon’s ‘1-click’ ordering was a game changer. *Knowledge@Wharton*, <http://knowledge.wharton.upenn.edu/article/amazons-1-click-goes-off-patent/>.
- Wang, Kitty, Avi Goldfarb. 2017. Can offline stores drive online sales? *Journal of Marketing Research* 54(5):706–719.
- Xu, Kaiquan, Jason Chan, Anindya Ghose, Sang Pil Han. 2016. Battle of the channels: The impact of tablets on digital commerce. *Management Science* 63(5):1469–1492.

Table 1: Summary Statistics of Treated and Control Groups

	Treated	Control	Treated – Control
Panel A. Before Matching			
Pre-treatment period			
Purchase amount (\$)	14.98	6.21	
Order frequency	0.39	0.20	
Order size	0.59	0.44	
Post-treatment period			
Purchase amount (\$)	22.19	5.29	
Order frequency	0.61	0.15	
Order size	0.79	0.29	
Difference			
Purchase amount (\$)	7.21	-0.92	8.13
Order frequency	0.22	-0.05	0.27
Order size	0.20	-0.15	0.35
Observations	977	17,229	
Panel B. After Matching			
Pre-treatment period			
Purchase amount (\$)	11.61	11.17	
Order frequency	0.35	0.35	
Order size	0.55	0.69	
Post-treatment period			
Purchase amount (\$)	17.32	12.50	
Order frequency	0.55	0.38	
Order size	0.75	0.63	
Difference			
Purchase amount (\$)	5.71	1.33	4.38
Order frequency	0.20	0.03	0.17
Order size	0.20	-0.06	0.26
Observations	801	801	

Table 2: Covariates for Propensity Score Estimation

Variable	Operationalization	Mean	Std. Dev.
Customer-firm relationship			
Tenure	Number of months since having online account	54.55	42.50
Recency	Elapsed time (days) since last purchase	158.80	115.00
Frequency	Number of purchases made	5.28	10.95
Monetary	Purchase amount (\$) spent	80.19	235.60
Discount	Average amount (\$) of discount received	6.02	6.14
Return	Number of products returned	1.57	12.22
Visit	Average number of website visits	2.50	7.55
Page view	Average number of pages viewed per visit	11.23	36.44
Duration	Average time (minutes) spent on the website per visit	14.86	56.38
Variety in purchase			
Overall	Number of unique categories purchased	1.71	1.11
Category 1	Number of unique products purchased in category 1	1.81	3.24
Category 2	Number of unique products purchased in category 2	0.93	1.63
Category 3	Number of unique products purchased in category 3	0.31	1.11
Category 4	Number of unique products purchased in category 4	0.38	1.31
Others	Number of unique products purchased in other categories	0.59	1.52
Repeat in purchase			
Overall	Share of unique categories purchased more than once	0.12	0.25
Category 1	Share of unique products purchased more than once in category 1	0.08	0.23
Category 2	Share of unique products purchased more than once in category 2	0.07	0.23
Category 3	Share of unique products purchased more than once in category 3	0.02	0.12
Category 4	Share of unique products purchased more than once in category 4	0.01	0.10
Others	Share of unique products purchased more than once in other categories	0.03	0.13
Socio-demographics			
Age		34.90	9.50
Gender	1 if female, 0 if male	0.94	0.25
Region 1	Zip-code based classification 1	0.39	0.50
Region 2	Zip-code based classification 2	0.25	0.43
Region 3	Zip-code based classification 3	0.17	0.38
Region 4	Zip-code based classification 4	0.10	0.29
Region 5	Zip-code based classification 5	0.10	0.30

Table 3: Treatment Effects

	Purchase Amount (\$)	Order Frequency	Order Size
	(1)	(2)	(3)
Post Treatment			
12 months	0.16*** (0.05)	0.05*** (0.02)	0.06*** (0.02)
6 months	0.21*** (0.05)	0.05*** (0.02)	0.07*** (0.02)
3 months	0.26*** (0.06)	0.08*** (0.02)	0.09*** (0.03)
Observations	1,602	1,602	1,602

Notes: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$. All outcomes are in natural log terms. Bootstrapped standard errors are shown in parentheses.

Table 4: Importance of Covariates in Heterogeneous Treatment Effects

Variable	Importance (%)	Importance Rank
Customer-firm relationship		
Tenure	10.41	6
Recency	12.58	3
Frequency	12.53	4
Monetary	16.56	1
Discount	13.82	2
Visit	10.81	5
Variety in purchase: Overall	5.85	9
Repeat in purchase: Overall	9.31	7
Socio-demographics		
Age	6.98	8
Gender	1.16	10

Table 5: Treatment Effects on Offline Purchases and Product Returns

	Purchase Amount (\$)	Order Frequency	Order Size	Returns	Prop. of Returns
	(1)	(2)	(3)	(4)	(5)
Post Treatment					
12 months	0.02 (0.03)	0.01 (0.02)	0.00 (0.02)	0.014** (0.006)	0.008*** (0.002)
Observations	1,602	1,602	1,602	1,602	1,602

Notes: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$. All outcomes are in natural log terms. Bootstrapped standard errors are shown in parentheses.

Table 6: Treatment Effects on Online Activity and Category Purchases

	Website Visits	Page Views per Visit	Duration per Visit (min)	Category Expansion
	(1)	(2)	(3)	(4)
Post Treatment				
12 months	0.17*** (0.04)	0.07* (0.03)	0.09*** (0.03)	0.36*** (0.03)
Observations	1,602	1,602	1,602	1,602

Notes: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$. All outcomes are in natural log terms. Bootstrapped standard errors are shown in parentheses.

Table 7: Rosenbaum Bounds

Γ	p -value	Lower Bound	Upper Bound
1.4	0.000	0.155	0.442
1.5	0.002	0.126	0.473
1.6	0.013	0.099	0.502
1.7	0.049	0.074	0.530
1.8	0.129	0.051	0.557
1.9	0.262	0.028	0.582

Table 8: Treatment Effects without Log Transformation

	Purchase Amount (\$)	Order Frequency	Order Size
	(1)	(2)	(3)
Post Treatment			
12 months	3.48*	0.14***	0.18***
	(1.67)	(0.04)	(0.06)
6 months	4.04*	0.12***	0.20***
	(1.67)	(0.04)	(0.06)
3 months	4.73*	0.15***	0.21*
	(1.89)	(0.04)	(0.09)
Observations	1,602	1,602	1,602

Notes: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$. Bootstrapped standard errors are shown in parentheses.

Table 9: Treatment Effects Using Panel Regression

	Purchase Amount (\$)	Order Frequency	Order Size
	(1)	(2)	(3)
Post Treatment			
7-12 months	0.28***	0.09***	0.09***
	(0.05)	(0.02)	(0.02)
4-6 months	0.26***	0.07***	0.10***
	(0.05)	(0.02)	(0.02)
≤ 3 months	0.36***	0.09***	0.11***
	(0.05)	(0.01)	(0.02)
Observations	38,448	38,448	38,448

Notes: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$. All outcomes are in natural log terms. Individual and month fixed effects are included in all estimations. Robust standard errors are clustered at the customer level and shown in parentheses.

Figure 1: Purchase Amount (\$) of Treated and Control Customers




Figure 2: Distribution of the Propensity Score

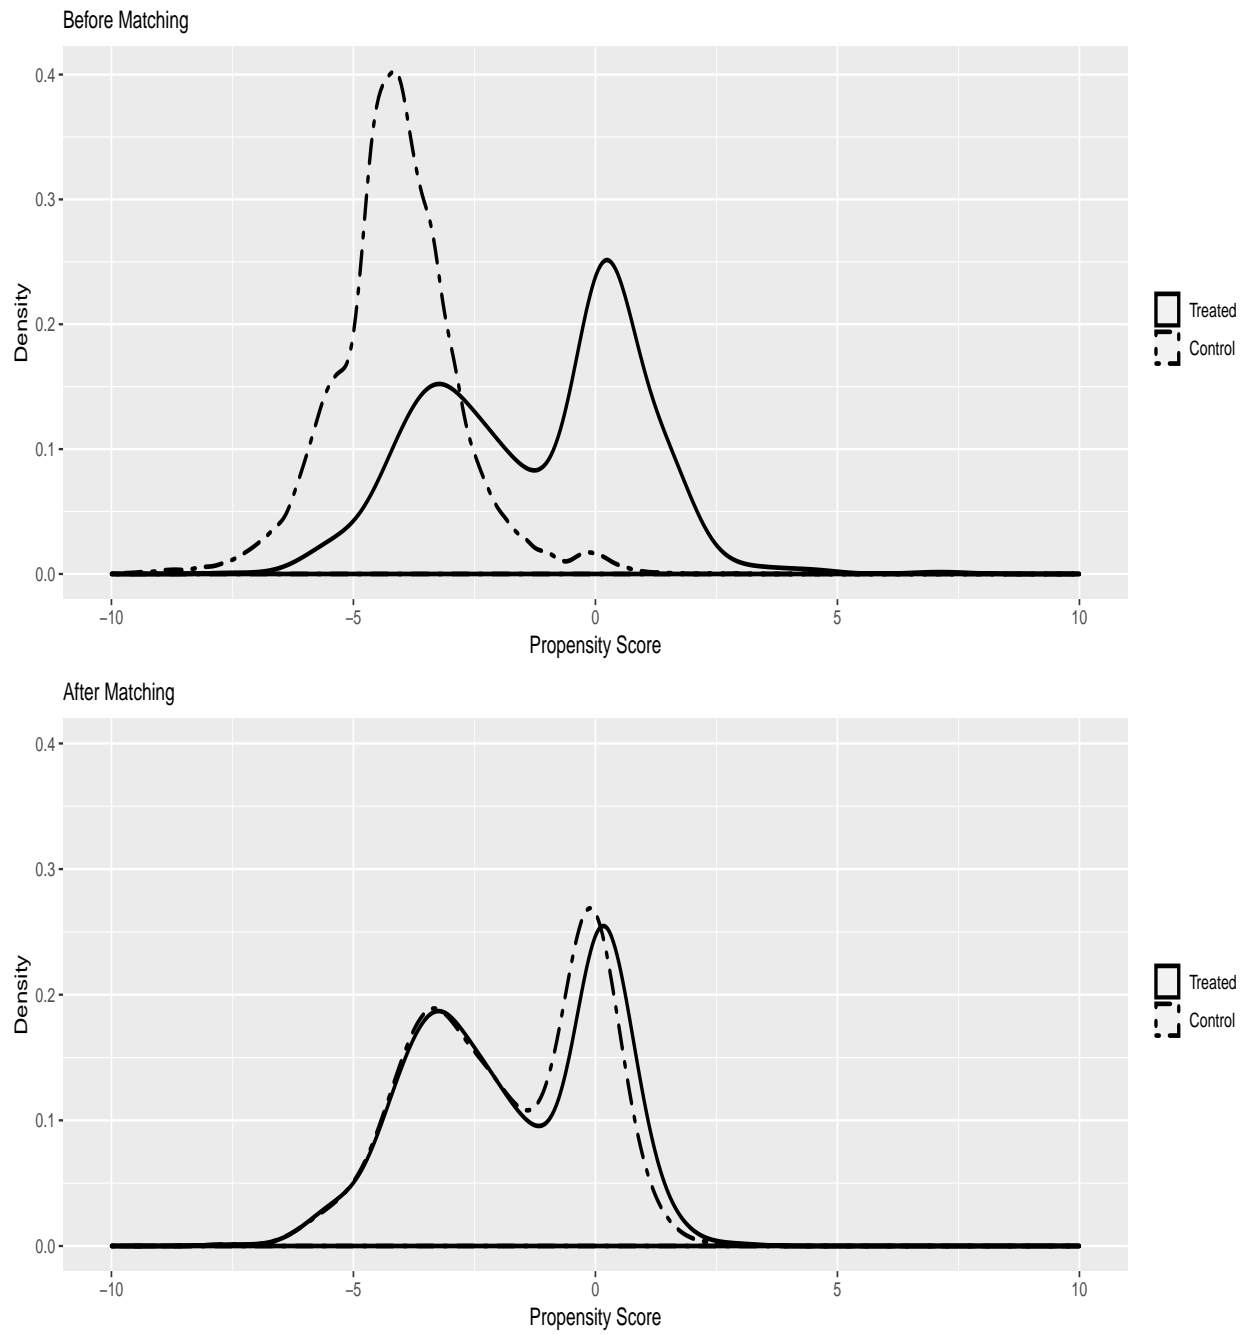


Figure 3: Covariate Balance





Figure 4: Comparison of Purchase Trends Between Treated and Control Groups



Note: The solid lines show the coefficient estimates from Equation (1) which capture the average difference in purchases between the treated and control groups over time relative to the baseline of 12 months prior to the treatment. The pre-treatment and post-treatment months are from 1 to 12 and 13 to 24, respectively. The error bars represent the 95% confidence intervals of the estimates. Standard errors are heteroskedasticity robust and clustered at the customer level. All outcomes are in natural log terms.

Figure 5: Distribution of the Treatment Effects




Figure 6: Standardized Covariate Means between High and Low Treatment Effects

