

# The Added Value of Data Analytics: Evidence from Online Retailers\*

Ron Berman<sup>†</sup>

Ayelet Israeli<sup>‡</sup>

August 2019

---

\*We are grateful for helpful comments from Eva Ascarza, Simha Mummalaneni, Davide Proserpio, Christophe Van den Bulte, Lynn Wu, and participants at the Marketing Science 2019 Conference. We are indebted to the anonymous analytics service provider that provided us with the data. Shawn Zamechek and Hengyu Kuang provided excellent research assistance. Research support from The Mack Institute for Innovation Management is acknowledged.

<sup>†</sup>The Wharton School, ronber@wharton.upenn.edu

<sup>‡</sup>Harvard Business School, aisraeli@hbs.edu

# The Added Value of Data Analytics: Evidence from Online Retailers

## Abstract

Does the adoption of data analytics impact retailers' performance, and if so, how? Exploiting the staggered adoption of a retail analytics service by more than 1,000 e-commerce websites, we find an average increase of 8-29% in monthly revenues post adoption. Further analysis shows an increase in the number of transactions, the number of unique customers, and revenue from repeated customers. In contrast, we do not find a change in basket size. These results are robust to different identification strategies and specifications. Importantly, we demonstrate that only retailers that adopt and use the service reap these benefits.

We explore subsequent marketing decisions to uncover the mechanism underlying retailers' improved performance. While retailers do not seem to change their advertising spending or pricing strategies after adoption, they install more prospecting technologies that allow websites to better attract and serve new customers. Further, retailers' exhibit an increase in the number of paid visitors to their sites, an increase in the diversity of products sold. Put together, these findings suggest that the adoption of the data-analytics service added value to retailers, primarily through improved prospecting activities that drive new customers to the retailer's website with more heterogeneous preferences.

# 1 Introduction

Businesses have invested substantially in developing analytics capabilities in recent years. These capabilities span a wide variety of resources, including the ability to collect large amounts of fine grained data, the computational resources to analyze this data, and shifting the organization towards data-driven decision making by enhancing the skill sets of employees (Henke and Kaka 2018). In 2018, the market for big data and analytics solutions included both established companies and startups with hefty funding (Whiting 2019), totalling nearly \$169 billion USD in revenues, with roughly a third in software solutions revenue (Shirer and Goepfert 2019).

The usage of analytics tools may benefit businesses through many organizational improvements (such as in supply chain and operations, strategic competitiveness, internal human resources management, etc.), but studies such as Bughin (2017) have shown that many companies that adopted big data and analytics technologies report that the fastest and widest adoption is in the marketing and sales functions, with the majority of the return of investment coming from this adoption. However, when delving deeper into how these companies utilize these tools to drive decisions, researchers conclude that “analytics can create new opportunities and disrupt entire industries. But few leaders can say how,” and that “surprisingly few companies know where and how analytics can create value” (Henke and Kaka 2018).

Prior literature that investigated the relationship between data analytics and firm performance found a positive relationship of 3%-6% greater productivity for firms that adopt data driven decision making or big data assets. However, most of these results are correlational and suffer from endogeneity issues. For example, more productive firms might have more funds available to invest in data analytics. Part of the difficulty in pinpointing the value of data analytics and gaining insight into how to best utilize these tools is the lack of primary data about firm adoption and usage of analytics, and about the resulting firm output as well as the different actions that firms have taken to achieve these outcomes. For example, past research is based on survey data (e.g. Brynjolfsson et al. (2011a) and Brynjolfsson and McElheran (2016)), general measures of big data assets (Müller et al. 2018), or public firms’ financial performance (e.g. Brynjolfsson et al. (2011a) and Müller et al. (2018)). Related research on the value of big data that utilized detailed primary data has shown that an increase in the amount of data available to firms improves prediction accuracy of demand, but has diminishing returns (Bajari et al. 2019). Despite this initial positive evidence, an additional

challenge is the lack of random assignment of firms into adopting and using data analytics, which makes causal interpretation difficult. One paper that attempted causal identification using Bartik (1991)-type instruments for adoption (Müller et al. 2018) found that these instruments eliminate evidence of a positive effect of investments in big data assets on firm performance.

Our paper aims to overcome these data and identification challenges, and provide deeper insights about the value of analytics tools and how they are used in the context of online retailing. We focus on two research questions. First, we attempt to provide a causal estimate of the added value of data analytics on online retailer performance using a unique dataset that allows us to observe a panel of performance data before and after retailers adopted an analytics solution. We also ask whether all firms experience a similar benefit and quantify the amount of heterogeneity in the impact of retail analytics using detailed firm level data. Our dataset comprises data from hundreds of online stores, in a wide variety of retail categories. Focusing on online retailers allows us to compare similar firms, which results in a more accurate estimate compared to previous work, and the ability to provide a causal and nuanced interpretation of the effect.

Second, there is no theoretical prediction regarding the route through which analytics impacts the performance of retailers. For example, firms can hire more employees skilled in interpreting analytics reports and making decisions, or undertake operational or financial improvements etc. We ask what observable changes in marketing actions firms take as a result of adopting an analytics service. In other domains (e.g., Wu et al. 2017) data analytics has been shown to increase diversity of innovation which in turn increased firm productivity through uncovering potential complementarities between existing technologies. In online retailing (Oestreicher-Singer and Sundararajan 2012), peer-based recommendations have been shown to increase diversity of products sold. We expect data analytics to provide insights and uncover potential improvements that will be observed in the marketing actions that operate to impact the retailer’s bottom line. Specifically we examine the pricing, advertising, product assortment, and technological investment decisions of the online retailer. This allows us to provide evidence about the potential mechanisms through which the adoption of data analytics creates benefits.

The dataset we analyze comprises of monthly data for 1,169 online retailers that have adopted the focal analytics service in the years 2016 and 2017. The analytics service integrates different sources of data and displays reports to its clients. The retailers in our sample adopted the service

in different times, which creates a staggered adoption pattern. Moreover, when a retailer decided to install the analytics service, the service provider collected historical data going back up to 24 months prior. The staggered adoption pattern and availability of panel data before and after the adoption of the analytics service allow us to causally identify the impact on different measures of firm output. A second unique aspect of our data is that it allows us to observe transaction level data by the retailers as well as other changes the retailers have made to their online stores during that period (such as investment in advertising or in technological changes). Using this augmented dataset we create a set of indicators for firm actions (such as changing advertising spend, or changing prices) which we use to analyze the paths through which data analytics may operate to create value for the online retailer.

Three major findings emerge from our analysis regarding (i) the overall effect of analytics adoption, (ii) who reaps the most benefits and (iii) the mechanisms through which the benefits are gained. First, we find evidence of a positive overall main effect of adopting analytics technologies on firm performance, consistent with prior literature. Our estimate, however, is causal. In our sample, adoption of an analytics service is associated with an increase of 8% to 29% in monthly revenues of the firm, as well as increases in the number of monthly transactions, number of unique customers, and revenue from repeated customers. To alleviate concerns about a bias in the estimates due to endogeneity of the service adoption, we use a matching estimator as well as instrumental variables that shift the timing of adoption, and find consistent results. We also find that the smallest retailers (in terms of revenues or transactions) do not benefit from adoption, while larger retailers do.

Second, we are able to disentangle usage from adoption and show that among adopters of the analytics service, only *users* of the service improve their performance, and these improvements increase with usage of the analytics reports. Moreover, higher level of usage is associated with greater increases in outcomes. These findings provide more insight about who is expected to gain from adopting the service, and strengthens the causal claim that the benefits are due to the analytics service and not due to simultaneity of adoption and other actions. Further analysis shows that usage of a few specific reports is associated with the largest gains the retailers experience. Specifically the reports that analyze revenue and those that benchmark the relative performance of the store are associated with the largest gain. Unlike prior work that observed investment in analytics, but not actual usage, this analysis not only contributes by strengthening causal claims, but also by

providing further insight into the mechanism through which analytics affects firm outcomes.

The third major finding is that the adoption of the analytics service resulted in the firms investing more in “prospecting” technologies, an increase in visitors that arrived through sponsored channels, and an increase in the diversity of products purchased by customers. In contrast, the adoption of analytics did not seem to affect the pricing and advertising decisions by retailers. Overall, we conjecture that these effects are related: website prospecting technologies allow the retailer to better target new customers without changing advertising budgets, which in turn increases the heterogeneity of customer tastes for products and result in more diverse purchases. This explanation is consistent with previous research findings on the impact of technology adoption on product diversity in online retailing (Brynjolfsson et al. 2011b, Oestreicher-Singer and Sundararajan 2012).

Overall, we find that adoption of the retail analytics service has a positive impact on retailer performance, and that this impact is borne primarily through changes in prospecting new customers, and not through different pricing strategies or advertising spending. Customer prospecting often requires automated and dynamic software, and cannot be replaced by manual operations of store managers. This is in contrast to making changes to pricing and advertising that are often discrete choices that can be made (although potentially less efficiently) by people rather than by software. An implication of these findings is that investing in analytics technologies alone may not make it easier for firms to gather insights and improve decision making. To reap the benefits of analytics tools that collect data and allow analysis, firms may also need to invest in additional software that automate actions and decision making.

## 2 Institutional Background

Our analysis focuses on online retailers from a variety of industries that operate their own online stores.<sup>1</sup> Many of the retailers in our dataset manufacture and sell their own brand of products. For example, in the Clothing and Fashion categories, many of them produce and sell their individual designs and do not carry clothing from other brands.

The retailers operate their stores using E-Commerce platforms such as Shopify, Magento or BigCommerce. The E-Commerce platforms provide the retailer with the software needed to easily create and operate an online store, including web templates, inventory and fulfillment software,

---

<sup>1</sup>A small number of these retailers also sell on Amazon.com as 3rd party sellers.

and checkout and payment processing. These platforms also allow adding plug-ins (or apps) to enhance the basic E-commerce services provided to the retailers. Plug-ins can include product recommendation algorithms, advertising tracking, email targeting and hosting, and many more services. Among these plug-ins, an analytics plug-in allows retailers to access a dashboard that allows them to monitor and analyze the performance of their online store.

Our data was provided by an analytics service provider that offers an analytics dashboard software with ability to integrate, collect and analyze data from the three E-commerce platforms mentioned above, as well as from Amazon and from Paypal and Stripe (for sellers that use Paypal or Stripe as a checkout mechanism).<sup>2</sup>

The analytics service launched its dashboard in 2015 and reached a substantial volume of subscribed retailers in 2016. During the majority of the period for which we collected our data, the analytics service was ranked as one of the top analytics plug-ins on these E-commerce platforms, giving it substantial visibility among retailers. Once installed, the software collects historical data of the retailer from the E-commerce platforms as well as other data sources the retailer can connect to the dashboard. Examples would be data from Google Analytics and Facebook’s advertising platform. The data is collected daily, at the transaction level, going backwards for up to 24 months before the installation of the dashboard software. The service has a basic free version, but most retailers pay for the full range of services. Fees depend on the retailer’s annual revenue, but are generally lower than 1% of that revenue.

Using the transaction level data, the dashboard provides visual reports that include metrics such as aggregate sales, average basket sizes, share of repeated customers’ revenues, cost of new customer acquisition, the average customer life time value and many other metrics. Depending on data availability (i.e., whether Google Analytics or Facebook advertising access was provided to the service) more than 20 metrics are calculated and displayed as a weekly or a monthly report. In other words, the service does not only display already existing data that the retailer had, but integrates and aggregates data from the retailer’s different data sources.

The information is presented to the retailer in the form of reports that focus on one aspect of the retailer’s operation. Examples include a revenue focused report, a customer acquisition report and a benchmark report that compares the retailer to other similar retailers in the same industry

---

<sup>2</sup>For reasons of confidentiality and data privacy, we cannot provide identifying information about the analytics service provider.

with similar revenue profiles.<sup>3</sup> Among our data, we have detailed information about each retailer’s access to each report over time.

During the period for which we collected our data, the analytics service ran multiple advertising and marketing campaigns to try and encourage sign-up to its service. Although we do not have specific data about every type of marketing outreach, we have detailed information about visits to the analytics service’s website by new (non-registered) visitors, which we later utilize as part of the identification strategy.

### 3 Data

We use data from all retailers that adopted the analytics service in 2016 and had at least one year of annual revenue of USD 100,000 or more. We ensure that each retailer has at least 12 monthly observations, out of which at least one observation occurs before the adoption and one observation after. This yields a total of 840 retailers who serve as a focal treated group. Some of the retailers sell through multiple platforms (such as Shopify, Amazon, Paypal) which yields 969 distinct retailer-platform combinations. The data is aggregated to a monthly level. Note that since retailers adopt the service in a staggered manner, retailers that adopted the services in 2016 serve as a control group before their adoption.

We collect additional data on retailers that adopted the analytics service in 2017 and had historic annual revenue of at least USD 100,000 to serve as a control only group. To ensure sufficient overlap with the treated group, we selected retailers that had historical data that included at least one observation in 2015. This yielded 329 companies, and 426 retailer-platform combinations out of a total of 958 companies that adopted the service during 2017.

Figure 1 presents the treatment variation plot (Imai et al. 2018) that illustrates the dynamics of adoption of the analytics service in the data. Summary statistics for the dataset are reported in Table 1. The industry categorization for the 1,169 retailers in our data is described in Table 2. Table 3 displays the distribution of the different platforms for each of the 1,395 retailer-platform combinations in our data.

We complement the retailer data with data from five additional sources: (i) Data about visitors to the analytics provider’s website collected from Google Analytics. (ii) Data about usage (login

---

<sup>3</sup>Details on the different reports appear in later sections.



Figure 1: Treatment variation plot




Figure 1 presents the treatment variation plot (Imai et al. 2018). Each horizontal line corresponds to a retailer-platform combination. Blue regions represent units before adoption of the analytics service. Red regions represent units after adoption of the analytics service. White regions represent units with no data in the time period.

times, reports viewed) of the analytics service by retailers through MixPanel. (iii) Historical reports presented to the retailers by the analytics service provider. (iv) Data on additional technologies the retailer installed on their online website, such as advertising tracking and email tracking, collected via BuiltWith.com. (v) Data on historical keyword advertising collected via Spyfu.com, which includes periodical keywords used and ad spend for the U.S. and the UK. Table 4 includes summary statistics of the main variables collected for each of the retailers.

## 4 Empirical Strategy and Results

This section provides details of our empirical strategy, and then examines the results of the analysis of changes in retailer performance outcomes due to adoption of the analytics service. The main

Table 1: Summary statistics

Variable	Mean	Std. Dev.	Median	N
Panel A: Retailer level				
No. obs. per retailer	51.71	21.92	45	1,169
No. obs. per retailer-platform	27.52	3.84	30	1,169
Avg. monthly revenue (USD)	53,867.3	134,827.2	20,097.7	1,169
Total revenue (USD)	1,798,629.5	4,362,981	642,447.1	1,169
Avg. no. monthly transactions	679.75	2,127.12	231.3	1,169
Total no. transactions	23,641.3	82,419.14	7,089	1,169
Avg. monthly unique customers	618.24	1,970.56	198.7	1,169
Avg. basket size (USD)	183.03	372.33	87.42	1,169
No. platforms per retailer	1.23	0.5	1	1,169
Panel B: Observation level				
Avg. monthly revenue (USD)	54,938.3	160,210.7	18153.6	38,272
Monthly no. transactions	722.1	3,202.8	193	38,272
Monthly no. unique customers	658.3	2,776.2	170	38,272
Basket size (USD)	174.77	397.1	82.21	38,184

challenge to the identification of the effects is that retailers select *whether* and *when* to adopt the analytics service. In other words, there is no random assignment of adoption. The unique features of our data, that include staggered adoption of the service as well as a rich set of historical data for each adopter allow us to address this challenge.

We employ a threefold approach to identify the main effects of interest. First, as described in section 3, we construct a control group. To do so, we treat the cohorts of retailers that adopted the analytics service in 2016 as our focal treatment group, and use later adopters of the service as a control group (Manchanda et al. 2015). Second, we employ two types of difference-in-difference estimators: (i) a staggered difference-in-difference (SDD) analysis with time and retailer-platform fixed effects to net out the effects of time trends and retailer specific growth trends unrelated to the adoption of the service; (ii) A dynamic, Arellano-Bond type difference-in-difference analysis (Arellano and Bond 1991) that allows for dynamics in the growth of retailers. Third, we perform three additional analyses to alleviate concerns about selection of adoption timing and simultaneity of the effect with unobserved confounding factors: (i) propensity score matching to correct for potential bias due to adoption timing based on observables; (ii) construction of two instrumental

Table 2: Distribution of retailers' industries Table 3: Distribution of retailers' platforms

<b>Industry</b>	Frequency	Percent	<b>Platform</b>	Frequency	Percent
Clothing & Fashion	289	24.7%	Shopify	815	58.4%
Health & Beauty	136	11.6%	Paypal	359	25.7%
Food & Drink	86	7.4%	Stripe	99	7.1%
Home & Garden	82	7.0%	Amazon Seller	60	4.3%
Jewelry & Accessories	77	6.6%	BigCommerce	48	3.4%
Electronics	71	6.1%	Magento	9	0.6%
Sports & Recreation	63	5.4%	Amazon Vendor	5	0.4%
Toys & Games	36	3.1%	<b>Total</b>	1,395	100%
Other	329	28.1%			
<b>Total</b>	1,169	100%			

Table 4: Summary statistics - Additional data

<b>Variable</b>	<b>Mean</b>	<b>Std. Dev.</b>	<b>Median</b>	<b>N</b>
Monthly no. logins	.22	.95	0	32,412
Monthly indicator for any report views	.98	0.3	0	32,412
Monthly no. technologies	56.3	23.1	53	31,185
Monthly advertising spending (USD)	1,844.1	13,075.1	396.3	9,184
Monthly no. advertising keywords	29.4	47.3	9	5,183

The observations in this table are retailer-month level, since these data are collected at the retailer level. Since BuiltWith and Spyfu do not include all of the retailers in our sample, the number of observations is smaller for variables constructed from these websites.

variables that are likely shifters of adoption timing, but are plausibly unrelated to a retailer's performance outcomes; (iii) examining the effect on outcomes for retailers that have used the service after adoption vs. those who did adopt the service, but did not use it.

Since all of the retailers in the data adopt the analytics service at some point, and since they may differ from unobserved non-adopters, our estimates are essentially of the effect on retailers that choose to adopt the service, i.e., the local average treatment effect (LATE) and not the average treatment effect (ATE) in the population. Due to the endogenous adoption of analytics, LATE is an appropriate measure of the effect of adoption. We explore the mechanism by which the adopting retailers change their performance outcomes in Section 5.

## 4.1 Effect of adopting the analytics service on retailer performance

We examine five outcome variables (all at the monthly level): revenue, number of transactions, number of unique customers, basket size (defined as the average transaction size in USD), and repeated customer revenue (for retailers selling through Shopify or Paypal platforms only).<sup>4</sup>

### 4.1.1 Staggered difference-in-difference

First, we employ a staggered difference-in-difference (SDD) strategy to estimate the main effect of analytics adoption. We identify this effect by comparing the change in outcomes before and after adoption for adopting retailers with the change in outcomes in the same time periods for retailers that haven't yet adopted the service. That is, in each focal time period the treated group consists of the retailers that have adopted by the focal time period and the control group are the retailers that haven't yet adopted the service (similar in spirit to Wang and Goldfarb (2017)). Note, that unlike Datta et al. (2018), we do not observe a group of non-adopters and hence have to rely on future adopters as a control group.

Since retailers may provide a connection to the data of their sales platforms (e.g., Shopify, Amazon etc.) in different points in time, we use retailer-platform couples as the unit of observation.<sup>5</sup> We estimate the following SDD model using OLS:

$$\log(Y_{ijt} + 1) = \alpha + \beta \mathbf{AfterTreatment}_{ijt} + f_{ij} + g_t + \epsilon_{ijt} \quad (1)$$

$Y_{ijt}$  is the outcome of interest for retailer  $i$  and platform  $j$  in month  $t$ . **AfterTreatment** is an indicator variable which equals one starting at month  $t$  in which retailer  $i$  adopted the service using platform  $j$ , and zero otherwise. We control for retailer-platform fixed effects  $f_{ij}$  and month fixed effects  $g_t$ . We use two-way clustering of standard errors by retailer and month to address serial correlation (Bertrand et al. 2004).

To ensure that there are sufficient control units in each time period, we remove observations after adoption for the retailers that adopted the service in 2017.<sup>6</sup> Due to this data truncation, we

---

<sup>4</sup>Transaction level data from Shopify and Paypal identify repeated transactions, while data from other platforms does not.

<sup>5</sup>The results are robust to aggregating every retailer's outcomes from all platforms into one unit of observation. The treatment period is then defined using the earliest date of platform adoption. Table W-1 in the web appendix presents the results.

<sup>6</sup>The results are robust to including all of the 2017 observations (Table W-2 in the web appendix), and to including

limit our main analysis to include only observations between July 2015 and up to 6 months after adoption for each retailer-platform. An illustration of the remaining observations in the analysis appears in Figure A-1 in the Appendix.

The identifying assumption is that there were no differential trends in outcome variables before adoption between retailers that adopted the service and those that have not. The results of the regression based on equation (1) are first presented in Table 5, followed by evidence consistent with the identifying assumption. After adopting the analytics service, the results suggest that retailers exhibit an increase in revenues of roughly 29%. The unconditional median monthly revenues before adoption in our data is USD 15,548, therefore the median retailer experienced an increase of approximately USD 4,567. Similarly, the number of transactions increased by 21% after adoption, which translates to an increase of about 35 transactions for the median retailer. The number of unique customers increased by 20% after adoption, yielding 29 more customers for the median retailer. Next, while the coefficient is positive, there is no statistically significant change in basket size after adoption. Finally, there is an increase of 50% in repeated revenue among sellers on Shopify and Paypal after adoption, corresponding to an increase of USD 1,742 for the median retailer. These results imply that the increase in revenues is associated with an increase in the number of buyers *and* in the number of transactions, both by new customers and by repeated customers. However, we do not find compelling evidence for a change in basket size.

Table 5: Staggered difference-in-difference

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.258** (.043)	.194** (.038)	.183** (.038)	.0185 (.013)	.407** (.06)
Observations	24,124	24,124	24,124	24,043	20,316

All models include retailer-platform and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

We now turn to examine the identifying assumption of “parallel trends” over time between subsets of these observations, for example, including adopters before July 2017 in the treatment group (Table W-3 in the web appendix).

treatment and control units. Figure 2 presents average outcomes for future adopters from the 2016 and 2017 cohorts over time and demonstrates visual evidence for the similar trends before adoption of the average revenues and number of transactions. Furthermore, the staggered pattern of adoption allows us to examine the identifying assumption using regression analysis. We estimate regressions similar to (1) where the **AfterTreatment<sub>ijt</sub>** indicator is replaced with indicators for each month starting six months before adoption to five months after adoption (totaling a year). In this analysis, the baseline for comparison is more than six months before adoption. Figure A-2 in the Appendix plots the coefficients associated with each of these monthly indicators and illustrates that: i) the outcome variables were not significantly different between the treatment and control before the time of adoption; and ii) the effects measured occur after adoption. For each outcome variable, we plot the point estimate and 95% confidence intervals.<sup>7</sup> Time 0 in the figure is the month of adoption. As can be seen in the figure, the coefficients are statistically indistinguishable from zero before adoption. For some of the variables, there seems to be a slight increase one month before adoption. However, it is not statistically significantly different from the previous month. The figure also reveals that the main effects decay over time.

We consider the SDD results above as our baseline results and next conduct a series of robustness tests and additional analyses to confirm our findings. We summarize the results of the different approaches in Table 6 and provide detailed tables to support the analysis in the web appendix. We omit the basket size variable from Table 6 for the sake of space, as the parallel trends assumption does not hold, and there is typically no effect on basket size.


#### 4.1.2 Robustness of SDD estimates

We conduct a series of robustness tests to validate the results of the SDD estimates of the main effect. Specific details for these tests are detailed in section A of the web appendix. First, we verify that the effects we measure are not due to the aggregation at the retailer-platform level, by re-aggregating our data to the retailer level. Table W-1 in the web appendix presents the consistent results. Second, to verify that the effect we measure is not spurious, we perform a placebo test using the pre-adoption data (Table W-4 in the web appendix). Third, to ensure that the results


---

<sup>7</sup>In the interest of space, we omit basket size from this figure, as the main effect from basket size is indistinguishable from zero. Basket size coefficients are increasing starting one month before adoption but not all the coefficients for the months before the adoption are statistically different from zero.

Figure 2: Average outcomes for the 2016 and 2017 cohorts before adoption



(a) Revenues



(b) Transactions

The figures plot the average revenues and transactions for the 2016 cohort (those retailer-platforms that adopted the service in 2016) and for the 2017 cohort (those retailer-platforms that adopted the service in 2017) in the period before adoption (between July 2015 and November 2016). In 2015, the figure plots the averages of all the retailers that adopt in 2016 and in 2017. Once retailers adopt the service they are dropped from the plot. Therefore, over time there are fewer 2016 adopters in the 2016 cohort line, such that in November 2016 only those retailers that adopted in December 2016 are plotted for the 2016 cohort. Other outcome variables exhibit similar patterns (except for basket size) and appear in Figure W-1 the web appendix.

Table 6: Effect of adoption using various methods

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(repeated revenue+1)
Basic SDD (Table 5)	.258** (.043)	.194** (.038)	.183** (.038)	.407** (.06)
Retailer level SDD (Table W-1)	.275** (.044)	.208** (.039)	.201** (.039)	.42** (.064)
Arellano-Bond (Table W-5)	.0806** (.029)	.0617** (.02)	.0603** (.02)	.179** (.04)
Matched SDD (Table W-6)	.257** (.055)	.193** (.048)	.182** (.048)	.407** (.078)
IV (Table A-2)	.17* (.069)	.19** (.059)	.15** (.058)	-.15 (.11)
OLS on IV sample (Table W-8)	.12** (.024)	.071** (.019)	.065** (.019)	.17** (.04)

Coefficients reported in this table are all for the **AfterTreatment** variable in the respective regressions

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

are not driven by the 2016 adopters, we include sub samples of 2017 adopters in our specifications as well (Table W-2 and Table W-3 in the web appendix).

#### 4.1.3 Dynamic difference-in-difference

Turning to explore the dynamic nature of our panel data, it is possible that some of the effect we capture in our analysis accrues to firms with faster growth rates that adopt the analytics service. This may lead such differential growth to be incorrectly captured as a positive effect of using the analytics service. In particular, we look at each of the outcomes of interest and regress their value also on their lagged value, to control for possible auto-correlation between past and future performance that will confound with the effect of adopting analytics.

In particular, we estimate the following specification:

$$\log(Y_{ijt} + 1) = \alpha + \beta \mathbf{AfterTreatment}_{ijt} + \gamma \log(Y_{ijt-1} + 1) + f_{ij} + g_t + \epsilon_{ijt} \quad (2)$$

Since equation 2 includes retailer-platform fixed effects and the lagged DV, we use the Arellano-



Bond dynamic panel-data estimator (Arellano and Bond 1991). In this method the data is first-differenced to remove the fixed effects, and the lags of the dependent variable are used as instruments for differenced lags of the dependent variable. Additionally, first differences of the remaining regressors are used as instruments.

Table W-5 in the web appendix reports the results of the GMM estimates. The results are consistent with those of Table 5 however, the effect sizes are substantially smaller.<sup>8</sup> In particular, the increase in revenues after adoption of the analytics service is 8% (USD 1,306 for the median retailer-platform), the increase in the number of transactions is 6% (10 additional transactions for the median retailer-platform), the increase in the number of unique customers is 6% (9 additional customers for the median retailer-platform), and the increase in repeated revenue is 20% (USD 680 for the median retailer-platform).

## 4.2 Addressing selection issues

In this section we focus on potential selection-related threats to the validity of our estimates. We aim to alleviate concerns about selection of adoption timing and simultaneity of the effect with unobserved confounding factors using propensity score matching, instrumental variables, and disentangling adoption and usage.

### 4.2.1 Matching to generate common support on observables

A possible issue with selecting to adopt the service is that the firms who adopt the service early are somehow different than those who adopt later. To address this issue, we combine matching and SDD (Heckman et al. 1997, Datta et al. 2018), to ensure common support on observables. We utilize cohort-based propensity score matching to ensure common support on observables between the adopters (the treated retailers) and the later adopters (the control retailers). We then use the matches created by this algorithm to estimate equation (1). The matching generated control group is similar to the treatment group in all observables except adoption. The generated treatment group omits retailer-platform observations that do not have corresponding matched controls.

---

<sup>8</sup>This is due to inclusion of lags and the instruments, and not due to the first-differencing estimation that identifies the effect of adoption based on one period after adoption. In fact, if we run the prior specification only one month after adoption, the effects sizes are larger than those obtained by using the entire sample.

Due to the staggered adoption, the matching procedure sequentially matches each retailer-platform cohort to observations from adopters in later cohorts. Once a retailer-platform is matched as a control, their observations starting that month are excluded from subsequent analysis. We use logistic propensity score matching to estimate equation (3) for each adoption month  $t$ , starting January 2016:

$$Pr(Adopted_{ijt} = 1) = h(\alpha + \beta X_{ijt}) \quad (3)$$

where  $Adopted_{ijt}$  indicates whether retailer  $i$  has adopted the analytics service using platform  $j$  by month  $t$ ,  $X_{ijt}$  are control variables that include: three monthly lags of log revenue, three monthly lags of log number of transactions, three monthly lags of log number of payers, country dummies, industry dummies, platform dummies, and the number of other platforms of retailer  $i$  (excluding platform  $j$ ).  $h(\cdot)$  is the logistic function. We report the distribution of propensity scores before and after matching in Figure A-3 of the Appendix.<sup>9</sup> After matching, we are left with 927 retailer-platform combinations for 807 retailers. Out of these, only 557 are retailers that adopted the service in 2016.<sup>10</sup>

We re-estimate the main model using the matched sample. Retailer-platform observations that are matched to more than one control are weighted accordingly in the standard error computation. Table W-6 in the web appendix reports the results from estimating equation (1) using the matched sample. The estimates obtained using the matched sample are extremely similar to those obtained using the unmatched sample, despite the larger standard errors and the smaller sample size.

#### 4.2.2 Instrumental variables to control for endogenous adoption timing

A second concern is that firms select the time of adoption endogenously using unobservables that are correlated with adoption and the outcome measure. For example, a retailer may anticipate the value of the analytics service and will choose to adopt it when this value is the highest. We address this issue using instrumental variables (IV). Specifically, we use the analytics service provider's Google Analytics data to identify two time-varying exogenous factors that plausibly impact the timing of analytics adoption, but are uncorrelated with the performance outcomes. We view these

---

<sup>9</sup>Logistic regression results for each of the twelve regressions are available from the authors upon request.

<sup>10</sup>The smaller number of retailers is due to the fact that i) our matching procedure allows for matches within any future cohort and then removes them from future groups, and ii) we exclude observations not on the common support of observables for both treatment and control

IVs as time shifters of adoption.

First, we use the number of unique new visitors that arrived to the service provider’s website from a particular platform (“platformVisits”) as an IV. This variable captures platform-based promotion or increase in awareness of the provider’s service. For example, the provider’s app was listed as the top app in the analytics category on the Shopify application store for a period of time. During that time period there were spikes in traffic arriving from Shopify. Such promotion would have increased awareness and attention to the analytics service, and thus was likely to influence the timing of adoption for an individual retailer without correlating with performance except for via adoption. We did not include traffic from channels with targeted advertising because they could violate the exclusion restriction assumption by attracting retailers that were already interested in adopting analytics. We categorize platform specific sources into three categories - traffic from Shopify, traffic from BigCommerce, and all other traffic.

Second, using a similar logic, we construct an IV using the number of unique new visitors from a particular geographic region. In this case, this variable captures the effect of local media and “buzz” about the service provider, which in turn is more likely to be consumed by retailers of closer geographic proximity (Blum and Goldfarb 2006). Such information should impact the timing of adoption through increased awareness of local retailers, but is plausibly uncorrelated with performance except for via adoption. The “regionVisits” variable measures location specific time-varying traffic for each retailer. The retailer’s region was defined at adoption as the location of the retailer’s headquarters. For the U.S., if state information is available we use the state as the region of the retailer, otherwise, we use the country as the region. Table 7 displays the distribution of regions of retailers in our sample.<sup>11</sup> We were unable to identify the location of 106 of the retailers, thus we define this variable for 1,063 of the retailers, which represent 1,280 of the retailer-platform combinations. As can be seen from these tables, roughly half of the companies are based in the U.S. with the majority in California.

One potential threat to the validity of the exclusion restriction for the regional IV is country-specific performances waves in which companies in countries that are doing well financially are also more likely to visit the service provider website and adopt the service. This concern is mitigated thanks to the short time-span of the data which did not exhibit any such waves to the best of our

---

<sup>11</sup>Table W-7 in the web appendix displays countries only. Locations with less than 10 retailers were consolidated into “other” for confidentiality reasons, and in the interest of space.

Table 7: Distribution of retailers' regions

<b>Region</b>	Frequency	Percent	<b>Region</b>	Frequency	Percent
California	218	20.5%	India	14	1.3%
United Kingdom	99	9.3%	New Jersey	14	1.3%
Australia	85	8.0%	Oregon	13	1.2%
Canada	71	6.7%	Colorado	12	1.1%
New York	56	5.3%	Georgia	12	1.1%
Texas	28	2.6%	Illinois	12	1.1%
Virginia	28	2.6%	Israel	11	1.0%
Florida	25	2.4%	Utah	11	1.0%
New Zealand	21	2.0%	Michigan	10	0.9%
Germany	17	1.6%	North Carolina	10	0.9%
Pakistan	15	1.4%	Pennsylvania	10	0.9%
Washington	15	1.4%	Other	256	24.08%
			<b>Total</b>	1,063	100.0%

knowledge. Additionally, our results are robust to adding region fixed effects.<sup>12</sup> We also use the platform-specific IV that is less likely to suffer from platform-specific performance waves.

Our instruments are relevant to the endogenous independent variable, since traffic of new users from a region or a platform is plausibly correlated with adoption. Therefore, we are confident that these are plausible adoption time shifters. However, this restriction doesn't hold for post-adoption. Therefore, we limit the scope of our analysis to include only observations up to the first month of adoption (including that month).

Due to the binary nature of the endogenous variable, our analysis utilizes the three-step estimation procedure proposed in Deng et al. (2019) based on Wooldridge (2010), Wooldridge (2005), and Wooldridge (2019). The first stage is a probit model for the decision to adopt the analytics service, the resulting predicted probability of adoption is then used as an instrument in a two step IV, hence leading to a three steps procedure. The estimation procedure is detailed in Appendix A.2.

Table A-1 in the Appendix presents the first stage results and shows that our IVs are positively

<sup>12</sup>We do not include the 65 region fixed effects in our main regression specification due to efficiency concerns and due to the incidental parameter problem in probit analysis. However, the main effect results using region fixed effects are virtually identical to the reported results.

correlated with adoption as expected. Prior to estimating the two-step IV, we preform a series of Durbin-Wu-Hausman endogeneity tests using the constructed IV,  $\widehat{Adopt}_{ijt}$ , for each of our outcome variables. Aside from the *BasketSize* outcome variable, all the tests produce p-values consistent with evidence of endogeneity (p-values for the tests are: 0.0001, 0.0086, 0.0441, 0.1796, 0.0002 for revenues, transactions, unique customers, basket size and repeat revenue respectively). We also confirm that our instruments are not weak using Stock-Yogo tests.

Table A-2 in the Appendix reports the second and third stage results of our IV estimation using the Baltagi error component 2SLS estimators (EC2SLS) (Han 2016). The F-statistics on the instrument in the second stage are all significantly different than zero. Additionally, the IV variable  $\widehat{Adopt}_{ijt}$  is positively associated with adoption. Overall, the IV results confirm most of our findings so far — revenues, transactions, and number of customers increase following adoption. However, the increase in revenue is significant only at the 5% level (p-value=0.018), and the coefficient for repeated revenue is insignificant but no longer positive.

To evaluate these results, Table W-8 in the web appendix provides a direct comparison and reports the results of an OLS regression using the same set of observations and a similar specification using retailer-platform random effects and the various fixed effects, without instrumenting for the Adoption indicator. These OLS results are quite similar to those obtained using the Arellano-Bond GMM estimates in Table W-5. As can be seen in the 2SLS third stage results, for revenues, transactions, and number of customers, the effect sizes are larger than the OLS results that include lagged performance.

The interpretation of our IV estimates is limited. Ideally, we would have retailer-platform specific time-varying exogenous variation we could exploit. Instead, we use time varying instruments (platform traffic / country traffic) that apply to a subset of retailers equally to estimate retailer-platform specific changes in performance due to adoption.<sup>13</sup> Moreover, because our IVs shift adoption which is a one time decision, the estimates limit the data to incorporating only one period after adoption. Therefore, we consider the IV estimates as supportive evidence for the positive effect of adoption on performance outcomes that we found using SDD.

---

<sup>13</sup>In particular, for “platformVisits” Shopify represents roughly 60% of the retailers, BigCommerce represents 4% and the rest 36%. For “regionVisits” there is less concentration with California representing 20% of retailers, and the five largest regions comprising 50% of retailers.

### 4.2.3 Disentangling adoption and usage

A third issue we examine is that of simultaneity. It is possible that firms that adopted the analytics service made other changes in parallel to adoption, such as changing their management team, and these actions generate the increase in performance outcomes we observe. If this concern is valid, we would expect firms who adopted the analytics service, but did not use it, to exhibit increased performance as well.

To examine this issue, we utilize the unique access we were given to MixPanel data that includes the login times and list of analytics reports accessed by each retailer. We examine three different aspects of usage: any usage, intensity of usage, and reports viewed. Note that the MixPanel information is available on the retailer level, and not the retailer-platform level. Therefore, we aggregate observations to the retailer level and use the first time of adoption as the retailer’s adoption time.

First, we construct a simple indicator variable that equals one if the retailer ever logged-on to the analytics service post-adoption and zero otherwise (“everUsed”). Roughly 10% of retailers that adopted the service have never used it. We interact this variable with our post adoption indicator to examine whether retailers that use the service exhibit different outcomes than those that do not.<sup>14</sup> Accordingly, we consider these tests as heterogeneity tests of the effect of adoption. Table 8 reports these results. Reassuringly, we find that our effects are driven exclusively by those retailers that use the analytics service. Although we are not able to completely rule out simultaneity of actions by the retailer, seeing a correlation between usage and performance is a necessary condition for analytics to have a causal effect. Since our prior estimates were an overall average for all adopters, we find that the effects are larger for those that adopted the service and became users of the service. For convenience, we report the overall effect on users in the bottom row of the table.<sup>15</sup> Note that this heterogeneity test also reveals a marginal positive effect of adoption on basket size for those retailers that used the service.

Next, we examine whether the intensity of usage has a differential relationship with performance outcomes. We construct two different measures. First, we count the number of events logged on

---

<sup>14</sup>This variable is only defined for adopting retailers after adoption. Hence the regression we estimate does not include “everUsed” as a main effect.

<sup>15</sup>These estimates are computed using a linear combination of the estimates for “After adoption” and “After adoption  $\times$  Ever used”.

Table 8: Usage after adoption

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	-.0871 (.095)	-.0199 (.089)	-.0499 (.086)	-.0995+ (.054)	-.0724 (.13)
After adoption X Ever used	.418** (.11)	.263* (.096)	.289** (.093)	.161** (.053)	.57** (.15)
Observations	20,459	20,459	20,459	20,459	17,674
Effect size For Users	.331** (.047)	.243** (.041)	.239** (.041)	.0614** (.018)	.498** (.067)

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

Effect size for users is computed as a linear combination of the estimates for “After adoption” and “After adoptionXUser”.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

MixPanel for each retailer in a particular month and interact that variable with the post adoption indicator. Second, we generate a variable indicating the total number of months of activity for each retailer, and interact that variable with the post adoption indicator. Both of these interaction variables yield positive and significant results for most of the DVs. Tables W-9 and W-10 in the web appendix present the results for each of these regressions, demonstrating that intensity of usage matters. Since the monthly events variable changes over time, we also estimate this regression using lagged DV and Arellano-Bond procedure, in which case the estimates for intensity of usage are identified off the variation in the changes in usage over time. These regressions yield consistent results (see Table W-11 in the web appendix). Overall, retailers with larger number of monthly events reap additional increases in outcomes (for example, one additional monthly event is associated with 1.2%-3.6% increase in revenues). To explore the intensity results further, we add a squared term to the number of events regression in the SDD specification and find some evidence of diminishing returns for the revenue and number of transactions variables indicated by negative coefficients on these terms (p-value=0.045 and 0.065 respectively, see Table W-12 in the web appendix).

Finally, we examine the usage of reports. The service provides five main reports — (i) customer

acquisition costs, (ii) revenue by platform report, (iii) benchmark reports that include a variety of performance metrics for the retailer and its selected benchmarks which are defined based on industry and revenues, (iv) executive report which is a summary of all reports, and (v) an insights report which the service generates based on their algorithms that gives specific recommendations to its retailer clients. We conduct two different tests using the reports data. First, we generate an indicator variable for whether the retailer examined any of the reports during a particular month. Nearly 78% of retailers viewed a report at least once, but there is substantial variation in report usage over time. Second, we create monthly indicator variables for whether a retailer examined a particular type of the five reports described above. These measures take into account the time-varying usage of reports after adoption. In each test, we interact each of these variables with the post adoption indicator.

We find that those retailers that accessed reports exhibit higher performance outcomes compared to those that did not. Web appendix Tables W-13 and Table W-14 present the SDD and Arellano-Bond results for these regressions.

With regards to the specific reports, a consistent pattern across our estimates is that those retailers that examined revenue reports or benchmark reports exhibited higher increases in all outcome variables, whereas those that examined the insight reports exhibit lower increases. Since the choice of reports to examine is endogenous, we are unable to provide further interpretation of this finding. Both of the regression results are reported in Table W-15 and Table W-16 in section D of the web appendix.

### **4.3 Heterogeneity of the impact of the analytics service**

Given the evidence for the effect of adopting the analytics service, we turn to ask whether the benefits are distributed uniformly across firms. In particular, we look at the heterogeneity of the effect based on firm size, platform through which the retailers operate, and the industry of the retailer. We use the SDD apparatus to measure these heterogeneous effects.

To analyze heterogeneity of the effect based on firm size, we use two proxies: i) size quartiles based on the average monthly revenues in the six months before adoption of each retailer-platform; ii) size quartiles based on the average number of transactions in the six months before adoption for each retailer-platform. We interact each of the quartile indicators with the post adoption indicator



in a series of regressions corresponding to our main regressions.

For revenue quartiles, retailer-platforms from the lowest quartile of revenue do not exhibit a statistically significant increase in outcomes after adoption (except for the repeated revenue measure, with  $p\text{-value}=0.081$ ), while the three other quartiles all exhibit similar positive increases (see Table W-17 in the web appendix). For transaction quartiles, the lowest quartile in terms of transactions exhibits increases in the revenue and repeated revenue outcome variables ( $p\text{-values}$  of 0.061 and 0.01, respectively), but not in the number of transactions or number of customers (see Table W-18 in the web appendix). All other transaction-based quartiles exhibit increases, with no robust significant differences between the different quartiles. Overall, while there are no robust significant differences between most of the size quartiles, there is an indication that smaller firms (average monthly pre-adoption revenues smaller than \$10,400 USD or those with fewer than 107 monthly transactions on average before adoption) exhibit less of the upsides from adoption of the service.

Next, we examine whether there are heterogeneous effects for different sales platforms the retailers utilize. We do so by interacting an indicator for each platform with the post adoption indicator, and run the appropriate series of regressions. The results reveal that our effects are driven mainly by changes on Shopify and Stripe, and not other platforms (see Table W-19 in the web appendix). Note that together these two platforms account for 65.5% of our retailer-platform combinations. Additionally, connecting the analytics service to either platforms reveals detailed transaction-item information including the specific items that were purchased and the price for each, potentially yielding more useful insights compared to other platforms. Note that Amazon vendors exhibit a negative effect, however we only have five such retailers in the data.

Finally, we examine differential effects by industry by interacting industry dummies with the post adoption indicator. We find that the results (see Table W-20 in the web appendix) are driven by the largest industry categories in our sample: other, clothing & fashion, health & beauty, and food & drink, which together account for 72% of the industries in the sample. While it is relatively easy to alter pricing and advertising for any retailer, it is typically easier to alter assortment in the apparel, personal care, and food and drink industries compared to the other industries. Such differential ability to change assortment is consistent with the evidence we find for potential mechanisms in the next section, and might explain why these industries experience a larger benefit.

## 5 Mechanism

Our second research question examines the mechanisms through which retailers generate the gains we identified in Section 4.1 from adopting the analytics service. In particular, we explore the actions retailers take to understand what drives better outcomes. The two primary avenues we examine first are the pricing and advertising decisions of retailers. We focus on these actions for two reasons. First, pricing and advertising decisions are dynamic and short term decisions retailers can make without too much investment and pre-planning, as they are performed during the normal operation of the business. Second, the analytics service provides a customer acquisition report that presents detailed information about customer acquisition cost (CAC), including comparing those values over time and to the retailer’s peers, as well as a revenue report that among other indicators compares the average transaction value of a retailer to its peers. These reports could indicate to retailers that their advertising and pricing are not optimized, leading to a change in advertising and pricing after adopting the analytics service.

Following this analysis, we analyze the web-technologies that retailers installed after adopting the analytics service. If there is a substantial increase in a specific type of technology use (for example, product recommendation engine, or advertising tracking), this may allow us to explain how retailers reap the gains of the analytics service. In this section, we use the SDD apparatus to measure the effects.

### 5.1 Does adoption affect advertising or pricing?

We focus on three measures of advertising decisions: spending, number of unique keywords used for Google advertising (AdWords) bidding, and changes to the ad text being displayed. We use two different sources of data to create the dependent variables - Google Analytics for those retailers that connected their Google Analytics data to the service, and SpyFu.com data for all retailers that were available on SpyFu.com.

First, we used the Google Analytics data variable “ad costs” as the monthly advertising spend for the 983 of our retailers (roughly 84%) which connected their Google Analytics data to the service. Since this variable passed our parallel trends tests, we used it as a dependent variable. However, we find no statistically significant difference in Google advertising spending after adoption of the service. We also examine the number of paid visitors, as a proxy for changes in advertising

strategy or customer acquisition costs. We find that both the number of paid visitors and the number of paid visits increased by about 21% (p-value=0.030). This suggests a potential change in advertising strategy which does not operate via advertising spending.

To further explore these findings, we collected additional data about retailers' advertising using SpyFu.com. SpyFu collects advertising data for websites in the U.S. and the UK. In particular, we identified each retailer's website and used Spyfu to collect: i) the estimated monthly advertising budget for each website, ii) the number of terms used for AdWords for each website, and iii) the number of new display ad copies. We were able to identify the monthly advertising budget for 520 of the retailers (roughly 44%) and the number of terms for 268 of the retailers. While all of these variables pass the parallel trends tests, we find no statistically significant changes in any of them following the adoption of the service. We caution that given the relatively small coverage of the SpyFu data we may not have the appropriate statistical power to identify any changes.

Next, we examine whether retailers change their pricing due to the adoption of the service. Despite the fact the service doesn't provide detailed reports on specific SKUs, the customer acquisition report and the revenue report may lead retailers to examine their pricing in order to improve these outcomes (in addition to the advertising spend, which we have explored above).

To investigate this question we would ideally have observed each retailer's entire assortment and pricing for each day of the month for the entire observation period. Instead, we rely on transaction level data that reflect information on products and prices only for products purchased. In particular, we use a sub-sample of our retailers that use Shopify as a platform and have detailed transaction data. This sample contains 801 of the Shopify sellers (roughly 98%).

We generate five pricing measures: i) the average price for all the products sold by a retailer in a particular month; ii) the average price (i) weighted by total quantity of sales for each product; iii) the average variance in prices for each product, to reflect changes in prices; iv) the average variance in prices for each product weighted by total quantity of sales for each product; (v) the average discount rate. While all of these variables passed our parallel trends tests (and hence we used them as a dependent variable in an SDD setup), none of them yielded statistically significant differences.

To conclude, we do not find compelling evidence for a change in advertising spend, advertising allocation, or pricing strategy following the adoption of the analytics service that can explain the

benefit we observe from adopting analytics. However, more visitors to retailers’ website arrive from paid advertising after adoption, suggesting potential (unobserved to the researchers) changes in advertising strategy, or an alternative method that optimized advertising spending.

## 5.2 Adoption of additional technologies and the potential impact on product assortment

Since we have not found robust evidence for the influence of analytics on pricing and advertising decisions, we now focus on another observable action the retailer can take – the adoption of additional web technologies. In particular, because we observe an increase in number of paid visitors to retailers’ websites, perhaps they arrive due to an underlying change in web technologies.

To perform this analysis, we identify each retailer’s website and find the list of web technologies installed on their website over time using BuiltWith.com. A web technology is a piece of software that can be added to a website and is often provided by a third party as a modular plug-in. Builtwith identifies over 32,000 web technologies that retailers can install on their websites, including web hosting technologies, widgets, advertising technologies, analytics plug-ins, payment processing software, etc. Among the 15 most common technologies installed on retailers’ websites in our data are jQuery (a simplified javascript library for websites), Google Analytics, the advertising technology DoubleClick.Net, Facebook for websites, and Shopify. We have complete website and technology data for 1,142 of the retailers in our data.<sup>16</sup>

We first look at the changes in the total technology count after adoption, and then look at the change in the type of technologies installed. We focus on five main types of technologies that are most common in terms of number of technologies and are related to retail and consumer management: Analytics and Tracking (565 different technologies), Advertising (532), E-commerce (263), E-mail hosting (121), and Payment (71). Since we are interested in how the number of technologies changes after adoption, we analyze it as a dependent variable. We use the proportion of each of the five types of technologies installed on the website as the dependent variables. These variables pass the parallel trends tests.

The results suggest that after adoption of the data analytics service, retailers decrease the share of email hosting technologies, and instead increase the share of analytics and tracking technologies

---

<sup>16</sup>This leaves us with 95% of our sample, our main results from previous analyses hold for this smaller sample.

and advertising technologies. Table W-21 in the web appendix reports the coefficients on the adoption indicator for each of the five technology categories.

To further understand the role of the increase in analytics and tracking as well as advertising technologies, we collect additional data on the sub-categories for these categories. We generate five meaningful sub-categories — consumer relationship management (CRM), website design and optimization, lead generation and prospecting, personalization, and other advertising technologies. Uncategorized technologies are identified as such. Prospecting technologies are focused on attracting new users to websites, while personalization technologies are focused on optimizing the on-site experience of existing visitors.

Using the same approach as before, we generate independent variables that measure the proportion of each of the technologies out of the total sum of these sub-categories.<sup>17</sup> We find a decrease in the proportion of website design and optimization, and uncategorized technologies. Conversely, we find an increase in prospecting technologies. The rest of the technologies exhibit no significant changes. Table W-22 in the web appendix displays the coefficients on the adoption indicator for each of the portions. The increase in prospecting technologies is consistent with the finding that more paid visitors arrive to retailers’ websites.

We conclude that the adoption of the analytics service encouraged the integration of more prospecting technologies into the retailer’s website. A possible driver of this increase may be the information on new visitors and returning visitors that the analytics service reported to retailers.

The increase in the proportion of prospecting technologies suggests that retailers should experience an increase in the number of customers (as observed in the previous subsection), and a shift in the composition of the client base towards bigger heterogeneity in their preferences (unobserved to the researchers). Such increase in the number and heterogeneity of customers is predicted (Brynjolfsson et al. 2011b) to manifest itself through an increase in the variety of products sold by the retailer. Hence, we turn to investigate changes in assortment that are associated with the adoption of the analytics service.

We generate two measures: i) number of unique products sold each month; ii) a Herfindahl Index (HHI) like score that measures concentration of sales of each retailer’s SKUs.<sup>18</sup> Both of these

---

<sup>17</sup>As before, there is an increase in the number of technologies for each of the sub categories, and several do not pass the parallel trends test. Instead, we are focusing on the proportions.

<sup>18</sup>In particular, we measure what fraction of sales each SKU generates, square that figure, and sum those up to

variables pass the parallel trends tests, and we preform our analysis on both. For the number of unique products, we find an increase of 15.4% (the post adoption coefficient is .143, with p-value of 0.03). For HHI, we find a 2% decrease in SKU concentration (coefficient=-.021, p-value=.0055). We interpret both of these results as evidence supporting increased heterogeneity in the composition of distinct products sold.

Admittedly, there could be other reasons why the concentration of unique items sold decreased. For example these effects may be due to changes in advertising or site design or different product recommendations. While we find no evidence of changes in advertising spending or an increase in personalization technologies, this increase is consistent with the increase in prospecting technologies and number of paid visitors that are expected to increase customer heterogeneity. However, the data we observe does not allow us to robustly examine this link.

## 6 Conclusion

While the interest in analytics technologies and their impact has been tremendous in the past few years, causal evidence for the efficacy of these technologies is surprisingly rare. This is partially due to lack of data, but also due to the lack of the ability to explain what drives the benefits that are observed in the output of the firms. In this paper we have aimed to address that by identifying the causal effect of the adoption of an analytics service by a wide variety of online retailers. Our unique dataset has allowed us to not only provide estimates, but also provide evidence that, as one would expect, just installing an analytics solution is useless if one does not use it.

The results of the analysis show that the adoption of the analytics service in our sample of firms increases monthly revenue by an average of 8 - 29%. Although this range is wide and results from using multiple estimation approaches, we have performed multiple robustness checks and alternative analyses to demonstrate that the effect is positive and substantial.

There are a few concerns with regard to the interpretation of our results that may arise from the analysis. The first concern is about simultaneity. Potentially, firms made many changes (e.g., hired more skilled employees) and one of them was adding an analytics solution, but this solution had no direct impact on firm actions or performance. A unique result that we provide in our analysis is to show that the usage of analytics, and not the adoption per-se, is what drives the improved

---

create a measure of concentration between 0 and 1 for each retailer-month.

firm performance. It is possible, of course, that the analytics service is used just for monitoring purposes and not used to drive decisions, but even then, this would imply that the monitoring aspects of analytics services provide firm value. An interesting avenue for future work would be to disentangle the direction of causality between analytics and other firm actions - if analytics is used to monitor other actions, then the value is in the data collection and analysis. If analytics is used to change actions, then the value is in the interpretation of the results from the analysis.

A second concern that arises from the analysis is that in our sample, all firms eventually adopted the service, raising a concern about selection. As we noted in the analysis, we do not claim that randomly assigning firms to adopting an analytics service will increase their revenues. On the contrary, random assignment would probably not affect usage enough to generate benefits for firms. Our results are applicable to those firms that choose to adopt analytics, and can be thought of as an LATE - among firms that choose to invest in analytics, what should they expect the impact to be?

Building on the main effect that we observed, we focused on identifying different potential mechanisms through which analytics may benefit retailers. The research on big data analytics (e.g., Brynjolfsson et al. 2011a, LaValle et al. 2011, Wamba et al. 2015, Akter et al. 2016, Brynjolfsson and McElheran 2016, Seddon et al. 2017) does not provide details beyond strategic and organizational considerations on how firms derive their observed benefits. Partially this is due to lack of detailed firm data, but more so this might be because past research has often aggregated data from many industries, where firms within the sample are difficult to compare.

Focusing on online retailers provided a better ability to inspect these companies and their actions. Specifically, some of the major decisions for retailers are their advertising, pricing and assortment choices. As there isn't a clear theoretical argument for how firms should best capitalize on their investments in analytics, our findings may provide some guidance. We do not find any changes in firms' actions with regard to pricing strategies or advertising spending, but we do find changes in resulting assortment of purchases. This change may be due to the firm changing the inventory of products it sells, but also due to the firm affecting the type of products consumers are exposed to or the type of consumers the firm attracts. Although the data cannot rule out (or support) the former explanation, the analysis of web technologies on the site provides further evidence that supports the latter - the adoption of analytics increases retailer integration of prospecting

technologies. These technological changes, coupled with changes in assortment and paid visitors, are consistent with the finding that the consumer’s basket size does not change, but the number of consumers as well as repeat revenues both increase.

These results, where we analyze changes in firm actions and connect them with firm outputs, might be driven by other causal paths than direct effects. This concern may potentially be alleviated by performing causal mediation analysis (Imai et al. 2010a;b, Pearl 2014), but the endogeneity of the mediators in our sample (firm actions) have prevented us from doing so. We believe that this approach, left for future work, would provide stronger evidence and clear predictions about how the usage of analytics operates through firm actions.

One interpretation of our results relates to what firms may expect when they adopt an analytics solution. Analytics solutions should not be expected to create actionable insights for firms easily. That is, analytics is not an “install and forget” solution, but rather one that requires continuous monitoring, and from which the benefits may accrue over time with experience, but also with additional investment. The analysis shows that among the actions that firms could take (and are observable to us), the firms that adopted an analytics solution adopted more automatic technologies (prospecting) instead of changing prices or reallocating advertising, which can be done manually. This may stem from the difficulty a business owner may have when interpreting analytics reports. Although the reports provide well crafted statistics and allow a deep dive into the data, they rarely provide recommendations and leave the user to generate insights from the data.

One potential implication from this observation is that firms should invest in additional technologies that build upon data analytics to take automatic actions for business owners. Without such investments, analytics services, on their own, might not live up to the expectations. A second potential implication is that as retailers adopt more advanced analytics solutions, firm competition might gradually be controlled more and more by automatic algorithms, providing a fascinating direction for researchers to understand these effects.

## References

S. Akter, S. F. Wamba, A. Gunasekaran, R. Dubey, and S. J. Childe. How to improve firm performance using big data analytics capability and business strategy alignment? *International Journal of Production Economics*, 182:113–131, 2016.



- M. Arellano and S. Bond. Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2):277–297, 1991.
- P. Bajari, V. Chernozhukov, A. Hortaçsu, and J. Suzuki. The impact of big data on firm performance: An empirical investigation. *AEA Papers and Proceedings*, 109:33–37, 2019.
- T. J. Bartik. Who benefits from state and local economic development policies? W.E. Upjohn Institute., 1991.
- M. Bertrand, E. Duflo, and S. Mullainathan. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1):249–275, 2004.
- B. S. Blum and A. Goldfarb. Does the internet defy the law of gravity? *Journal of international economics*, 70(2):384–405, 2006.
- E. Brynjolfsson and K. McElheran. Data in action: Data-driven decision making in u.s. manufacturing. Center for Economic Studies Working Paper 16-06, 2016.
- E. Brynjolfsson, L. Hitt, and H. H. Kim. Strength in numbers: How does data-driven decisionmaking affect firm performance? Working paper, 2011a.
- E. Brynjolfsson, Y. Hu, and D. Simester. Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science*, 57(8):1373–1386, 2011b.
- J. Bughin. Ten big lessons learned from big data analytics. *Applied Marketing Analytics*, 2(4):286–295, 2017.
- H. Datta, G. Knox, and B. J. Bronnenberg. Changing their tune: How consumers’ adoption of online streaming affects music consumption and discovery. *Marketing Science*, 37(1):5–21, 2018.
- Y. Deng, A. Lambrecht, and Y. Liu. Spillover effects and freemium strategy in mobile app market. Working paper, 2019.
- C. Han. Efficiency comparison of random effects two stage least squares estimators. *Economic Letters*, 148:59–62, 2016.


- J. J. Heckman, H. Ichimura, and P. E. Todd. Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *Review of Economic Studies*, 64(4):605–654, 1997.
- N. Henke and N. Kaka. Analytics comes of age. *McKinsey Analytics*, 2018. <https://www.mckinsey.com/~media/McKinsey/Business%20Functions/McKinsey%20Analytics/Our%20Insights/Analytics%20comes%20of%20age/Analytics-comes-of-age.ashx>, Accessed May 31, 2019.
- K. Imai, L. Keele, and D. Tingley. A general approach to causal mediation analysis. *Psychological methods*, 15(4):309, 2010a.
- K. Imai, L. Keele, T. Yamamoto, et al. Identification, inference and sensitivity analysis for causal mediation effects. *Statistical science*, 25(1):51–71, 2010b.
- K. Imai, I. S. Kim, and E. Wang. Matching methods for causal inference with time-series cross-section data. Working paper, 2018.
- S. LaValle, E. Lesser, R. Shockley, M. S. Hopkins, and N. Kruschwitz. Big data, analytics and the path from insights to value. *MIT sloan management review*, 52(2):21, 2011.
- P. Manchanda, G. Packard, and A. Pattabhiramaiah. Social dollars: The economic impact of customer participation in a firm-sponsored online customer community. *Marketing Science*, 34(3):367–387, 2015.
- O. Müller, M. Fay, and J. vom Brocke. The effect of big data and analytics on firm performance: An econometric analysis considering industry characteristics. *Journal of Management Information Systems*, 25(2):488–509, 2018.
- G. Oestreicher-Singer and A. Sundararajan. Recommendation networks and the long tail of electronic commerce. *MIS Quarterly*, pages 65–83, 2012.
- J. Pearl. Interpretation and identification of causal mediation. *Psychological methods*, 19(4):459, 2014.
- P. B. Seddon, D. Constantinidis, T. Tamm, and H. Dod. How does business analytics contribute to business value? *Information Systems Journal*, 27(3):237–269, 2017.

- M. Shirer and J. Goepfert. Idc forecasts revenues for big data and business analytics solutions will reach \$189.1 billion this year with double-digit annual growth through 2022. *IDC.com*, 2019. <https://www.idc.com/getdoc.jsp?containerId=prUS44998419>, Accessed August 14, 2019.
- S. F. Wamba, S. Akter, A. Edwards, G. Chopin, and D. Gnanzou. How ‘big data’ can make big impact: Findings from a systematic review and a longitudinal case study. *International Journal of Production Economics*, 165:234–246, 2015.
- K. Wang and A. Goldfarb. Can offline stores drive online sales? *Journal of Marketing Research*, 54(5):706–719, 2017.
- R. Whiting. 10 data analytics companies to watch in 2019. *CRN.com*, 2019. <https://www.crn.com/slide-shows/applications-os/10-data-analytics-companies-to-watch-in-2019>, Accessed August 14, 2019.
- J. Wooldridge. Fixed-effects and related estimators for correlated random-coefficient and treatment-effect panel data models. *The Review of Economics and Statistics*, 87(2):385–390, 2005.
- J. Wooldridge. *Econometric analysis of cross section and panel data*, pages 937–940. The MIT press, 2010.
- J. Wooldridge. Correlated random effects models with unbalanced panels. *Journal of Econometrics*, 211(1):137–150, 2019.
- L. Wu, L. M. Hitt, and B. Lou. Data analytics skills, innovation and firm productivity. The Wharton School Research Paper, 2017.

# A Appendix

## A.1 Figures

Figure A-1: Treatment variation plot for the main analysis



Treatment variation plot (Imai et al. 2018) of the subset of data presented in Figure 1 that we use in our main analysis. Each horizontal line corresponds to a retailer-platform combination. Blue regions represent units before adoption of the analytics service. Red regions represent units after adoption of the analytics service. White regions represents units with no data in that time period.

Figure A-2: Outcome variables before and after adoption

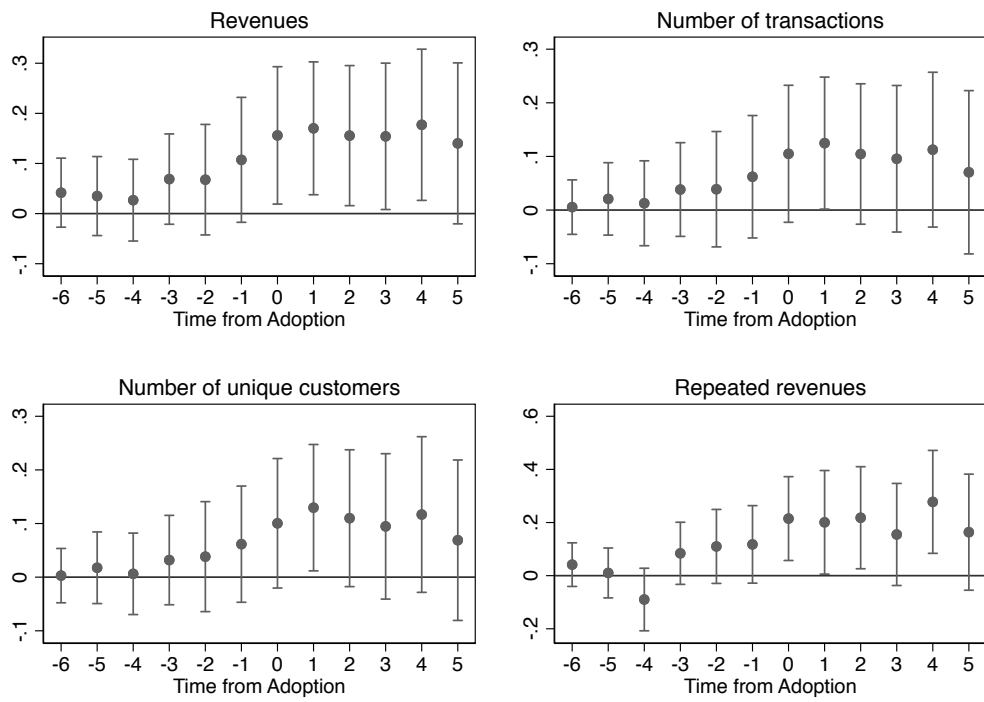
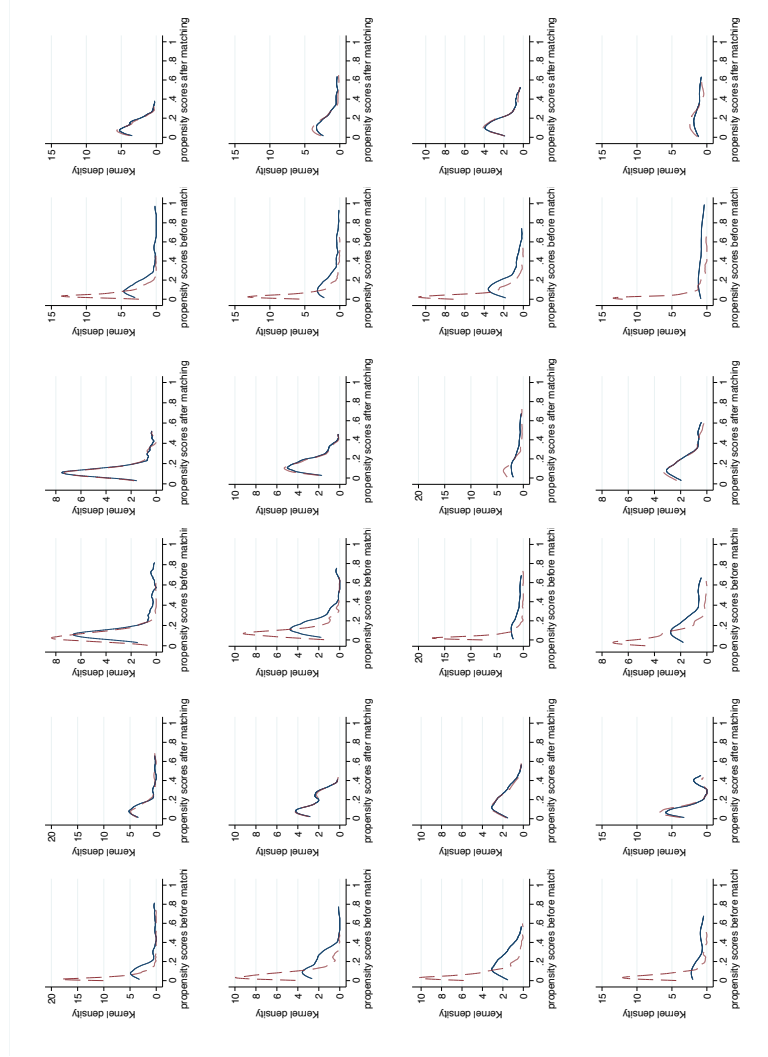


Figure A-3: Distribution of propensity scores before and after matching for each adopting cohort



The distribution of propensity scores before and after matching for each cohort, starting from January 2016 on the top right. The treated units are plotted using a solid line, and the control units are plotted using a dashed line.

## A.2 IV Procedure

In the first stage, we estimate a random effects probit model for the decision to adopt the analytics service. We estimate the following model:

$$Pr(Adopt_{ijt} = 1) = \Phi(\alpha + \beta X_{ijt-1} + \gamma Z_{ijt} + \delta \bar{X}_{ij} + \theta \bar{Z}_{ij} + g_t + \mu_{ij}) \quad (4)$$

Where  $Adopt_{ijt}$  indicates whether retailer  $i$  adopted the service in platform  $j$  in month  $t$ . Observations after adoption time  $t$  are not used to estimate equation 4 since the adoption decision is made once.  $X_{ijt-1}$  are retailer-platform-time control variables that include 8 industry dummies, 6 platform dummies, the number of other platforms of retailer  $i$  (excluding platform  $j$ ), and the lag DV of the relevant DV from the next stage for each regression.  $Z_{ijt-1}$  are the instrumental variables, which indicate the number of new visitors to the service provider's website from the platform (platformVisits) and from the region (regionVisits). Following Deng et al. (2019), and as suggested by Wooldridge (2005) and Wooldridge (2019) we include  $\bar{X}_{ij}$  and  $\bar{Z}_{ij}$ , the mean value of variables per retailer-platform for time-varying variables.  $g_t$  are month-year fixed effects. Due to the large number of retailers, and to address the incidental parameter problem with fixed effects in probit models, we include  $\mu_{ij}$  as retailer-platform random effects instead of fixed effects.

Following the first stage, the predicted probability of adoption,  $\widehat{Adopt}_{ijt}$ , is used as an instrument for the endogenous adoption indicator  $AfterTreatment_{ijt}$  in a two stage least squares regression. This regression comprises the the second and third stages of our estimation procedure, and is specified as:

$$\log(Y_{ijt} + 1) = \alpha + \beta \mathbf{AfterTreatment}_{ijt} + \delta X_{ijt-1} + \theta \bar{X}_{ij} + g_t + \mu_{ij} + \epsilon_{ijt} \quad (5)$$

Where the variables are defined as in the previous equations for each retailer  $i$  with platform  $j$  in month  $t$ . When estimating this equation we use all observation up to the first observation after adoption. Because we use an IV to alleviate endogeneity concerns, we use all adoptions that occur in 2016 or 2017 to estimate equations (4) and (5). Note that this approach allows us to incorporate month-year fixed effects, since there are adoptions and controls in each month and year (up to December 2017). These specifications increase the number of adoptions in our data (since all retailers adopt) compared to the limitation we had in our sample in section 4.1.1. For each retailer-platform, the controls are every retailer-platform that have not yet adopted, and the effects are identified off individual retailer-platform variation.

Table A-1: Adoption likelihood (First-stage probit)

Dependent Variable	Adopt log(revenue+1)	Adopt log(no. transactions +1)	Adopt log(no. unique customers+1)	Adopt log(basket size+1)	Adopt log(repeated revenue+1)
	(1)	(2)	(3)	(4)	(5)
Region visits	.0004** (.00013)	.0004** (.00013)	.0004** (.00013)	.0004** (.00013)	.0005** (.00014)
Platform visits	.000074** (.00001)	.000075** (.00001)	.000074** (.00001)	.000073** (.00001)	.00008** (.00001)
Log (Lag DV)	.055** (.021)	.079** (.025)	.077** (.025)	-.072 (.052)	.03* (.013)
No. other platforms	-.13* (.057)	-.13* (.056)	-.13* (.056)	-.12* (.051)	-.086 (.054)
Observations	12,351	12,351	12,351	12,310	10,206

Each column represents the first stage for a different logged DV: revenues, transactions, unique customers, basket size, repeated revenue (respectively). Lag(DV) in each column is the appropriate lagged dependent variable.

All models include (i) retailer-platform random effects, (ii) include industry platform and month fixed effects, and (iii) mean of retailer-platform controls.

SE in parentheses are clustered by retailer.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.



Table A-2: IV regression

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
Panel A: Second Stage (DV: After adoption)					
$\widehat{Adopt}_{ijt}$	.130** (.0231)	.130** (.0230)	.131** (.0230)	.130** (.0231)	.148** (.0240)
F-Stat	291.99	245.17	243.47	143.89	187.524
Panel B: Third Stage					
After adoption	.17* (.069)	.19** (.059)	.15** (.058)	-.035 (.026)	-.15 (.11)
Log (Lag DV)	.47** (.023)	.58** (.017)	.6** (.016)	.29** (.033)	.47** (.021)
No. other platforms	-.0081 (.014)	.0029 (.011)	-.00021 (.011)	-.0042 (.0055)	-.014 (.019)
Observations	12,351	12,351	12,351	12,310	10,206

All models include (i) retailer-platform random effects, (ii) include industry platform and month fixed effects, and (iii) mean of retailer-platform controls.

SE in parentheses are clustered by retailer.

Lag(DV) in each column is the appropriate lagged dependent variable.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

## Web Appendix

### A Main effect: Robustness

First, to verify that the effects we measure are not due to the aggregation at the retailer-platform level, we aggregate the data to the retailer level, redefine “after adoption” to occur once the first platform adopted the service, remove duplicate retailer-month level observations, and estimate the models at the retailer level. Accordingly, we use retailer fixed effects instead of retailer-platform fixed effects. Table W-1 presents the results. Although the point estimates are slightly higher, the results are consistent with our main analysis. In addition, this aggregation results in a statistically significant increase of basket size at the 5% level, and suggests an average increase of 4.6% in basket size after adoption.

Second, to verify that the effect we measure is not spurious, we perform a placebo test using the pre-adoption data. To account for seasonality in sales, we define “placebo adoption” a year before the true adoption period. We then estimate an SDD model for each of the outcome variables. Table W-4 presents the results and shows that the main effects at the placebo timing are either not distinguishable from zero or negative (at the 5% level), suggesting that our main findings are indeed due to the adoption.

Third, to ensure that our results are not driven by the 2016 adopters (that are early adopters and therefore may exhibit different effects than later adopters) we include 2017 adopters in our regressions as well. Table W-2 includes all of the 2017 adopters as treated, and Table W-3 includes only adopters up to June 2017 as treated. The results are similar to the main results.

Table W-1: Robustness to unit level

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.275** (.044)	.208** (.039)	.201** (.039)	.04* (.019)	.42** (.064)
Observations	20,459	20,459	20,459	20,459	17,674

In Table W-1, we aggregate retail-platform combination to the platform level. “After adoption” equals one when the first platform of that retailer adopted the service. We estimate equation 1, but the fixed effects are at the retailer level instead of the retailer-platform level.

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-2: Including 2017 adopters as treated

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.251** (.043)	.204** (.039)	.186** (.038)	.0113 (.0091)	.366** (.059)
Observations	42,491	42,491	42,491	42,368	36,248

In Table W-2, we treat all 2017 adopters as treated, and do not omit observations once they adopted. Data includes all months in years 2015-2017. We estimate equation 1.

All models include retailer-platform and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-3: Including adopters up to June 2017 as treated

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.344** (.044)	.29** (.039)	.266** (.039)	.0108 (.0096)	.446** (.06)
Observations	33,659	33,659	33,659	33,575	29,082

In Table W-3, we treat adopters up to June 2017 as treated, and do not omit observations once they adopted. Retailer that adopted between July 2017 and December 2017 are treated as a control group. Data includes months starting July 2015 through 2017. We estimate equation 1.

All models include retailer-platform and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-4: Placebo test

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
Placebo after adoption	-.0721* (.034)	-.0665* (.028)	-.0631* (.027)	-.00891 (.012)	-.0518 (.062)
Observations	18,098	18,098	18,098	18,018	15,408

In Table W-4, we use pre-adoption data only. “Placebo after adoption” equals one a year before adoption. We estimate equation 1 using this new “post” variable.

All models include retailer-platform and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

## B Main effect: Additional Analyses

Table W-5 reports the Arellano-Bond results, and Table W-6 reports the matching results.

Table W-5: Arellano-Bond GMM estimates

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.0806** (.029)	.0617** (.02)	.0603** (.02)	.00239 (.013)	.179** (.04)
Lag(DV)	.379** (.023)	.479** (.018)	.478** (.018)	.179** (.028)	.284** (.022)
Observations	22,773	22,773	22,773	22,668	19,251

Lag(DV) is the lagged dependent variable.

All models include retailer-platform and month-year fixed effects. Retailer-platform FE are differenced out. Asymptotic SE robust to general cross-section and time series heteroskedasticity are reported in parentheses.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-6: Matched staggered difference-in-difference

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.257** (.055)	.193** (.048)	.182** (.048)	.0266 (.017)	.407** (.078)
Observations	18,414	18,414	18,414	18,368	15,981

All models include retailer-platform and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

## C Main effect: IV Specification

Table W-7: Distribution of retailers' countries

<b>Country</b>	Frequency	Percent
United States	570	53.6%
United Kingdom	99	9.3%
Australia	85	8.0%
Canada	71	6.75%
New Zealand	21	2.0%
Germany	17	1.6%
Pakistan	15	1.4%
India	14	1.3%
Israel	11	1.0%
Other	160	15.05%
<b>Total</b>	1,063	100.0%

Table W-8: Comparable OLS estimates

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.12** (.024)	.071** (.019)	.065** (.019)	.013 (.0087)	.17** (.04)
Lag(DV)	.48** (.023)	.57** (.017)	.57** (.018)	.27** (.034)	.4** (.022)
# Other platforms	-.008 (.016)	.001 (.013)	-.00002 (.014)	-.003 (.0066)	-.02 (.025)
Observations	12,351	12,351	12,351	12,310	10,206

All models include (i) retailer-platform random effects, (ii) include industry platform and month fixed effects, and (iii) mean of retailer-platform controls.

SE in parentheses are clustered by retailer.

Lag(DV) in each column is the appropriate lagged dependent variable.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

## D Usage Results

Table W-9: SDD with Monthly Number of Events

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.25** (.047)	.186** (.04)	.178** (.04)	.0349 (.02)	.393** (.069)
After adoption X Activity Level	.0356** (.012)	.0313** (.0098)	.0322** (.0096)	.00727 (.0057)	.0396* (.018)
Observations	20,459	20,459	20,459	20,459	17,674

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-10: SDD with Number of Months Active

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.199** (.059)	.151** (.05)	.139* (.05)	.0138 (.025)	.314** (.082)
After adoption X No. Months Active	.0201* (.0087)	.0149+ (.0076)	.0163* (.0077)	.00691* (.0031)	.0289* (.013)
Observations	20,459	20,459	20,459	20,459	17,674

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.



Table W-11: Arellano-Bond with Monthly Number of Events

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.073* (.03)	.0651** (.021)	.0642** (.021)	-.00832 (.022)	.29** (.048)
After adoption X Activity Level	.0121+ (.0067)	.0129* (.0058)	.0117* (.0058)	.00355 (.0038)	-.00604 (.014)
Lag(DV)	.394** (.026)	.503** (.02)	.501** (.02)	.177** (.033)	.271** (.024)
Observations	20,459	20,459	20,459	20,459	17,674

All models include retailer and month-year fixed effects. Retailer FE are differenced out. Asymptotic SE robust to general cross-section and time series heteroskedasticity are reported in parentheses.

Lag(DV) is the lagged dependent variable.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-12: SDD with Squared Monthly Number of Events

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.241** (.048)	.179** (.041)	.172** (.041)	.0332 (.021)	.38** (.071)
After adoption X Activity Level	.0671** (.021)	.0564** (.018)	.0554** (.018)	.0131 (.01)	.0872* (.039)
After adoption X Activity Level Squared	-.00492* (.0023)	-.00391+ (.002)	-.00364 (.0022)	-.000909 (.0011)	-.00771 (.0047)
Observations	20,459	20,459	20,459	20,459	17,674

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-13: SDD with Report usage after adoption

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.227** (.048)	.167** (.041)	.161** (.041)	.0312 (.021)	.36** (.071)
After adoption X Any report	.176** (.047)	.148** (.037)	.145** (.037)	.032 (.02)	.223** (.064)
Observations	20,459	20,459	20,459	20,459	17,674
Effect size for report users	.403** (.053)	.315** (.045)	.306** (.045)	.0632** (.019)	.583** (.07)

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

Effect size for report viewing is computed as a linear combination of the estimates for “After adoption” and “After adoption X Any Report”.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-14: Arellano-Bond with report usage after adoption

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.0672* (.03)	.0641** (.021)	.0652** (.021)	-.011 (.022)	.273** (.05)
After adoption X Any report	.0513* (.026)	.0391+ (.021)	.0294 (.02)	.0183 (.015)	.0306 (.048)
Lag(DV)	.394** (.026)	.502** (.02)	.501** (.02)	.177** (.033)	.271** (.024)
Observations	19,425	19,425	19,425	19,425	16,815
Effect size for report users	.119** (.032)	.103** (.025)	.0946** (.025)	.00736 (.021)	.304** (.056)

All models include retailer and month-year fixed effects. Retailer FE are differenced out. Asymptotic SE robust to general cross-section and time series heteroskedasticity are reported in parentheses.

Effect size for report viewing is computed as a linear combination of the estimates for “After adoption” and “After adoption X Any Report”.

Lag(DV) is the lagged dependent variable.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-15: Specific report usage after adoption - Staggered difference-in-difference

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.231** (.048)	.17** (.041)	.163** (.041)	.0319 (.021)	.371** (.07)
After adoption X	.0331	-.00384	.00257	.0343	-.0807
Acquisition report	(.061)	(.056)	(.055)	(.02)	(.11)
After adoption X	.0692	.0815	.0947	-.0136	.187*
Benchmark report	(.063)	(.058)	(.056)	(.021)	(.079)
After adoption X	.00683	.0405	.0475	-.0134	.0317
Executive report	(.049)	(.043)	(.042)	(.022)	(.089)
After adoption X	-.0843+	-.0683+	-.0884*	-.011	-.162+
Insights report	(.048)	(.039)	(.04)	(.019)	(.093)
After adoption X	.179**	.132*	.126*	.0358	.258*
Revenue report	(.064)	(.053)	(.053)	(.023)	(.11)
Observations	20,459	20,459	20,459	20,459	17,674

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-16: Specific report usage after adoption - Arellano-Bond

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption	.0719* (.03)	.0666** (.021)	.0669** (.021)	-.00975 (.022)	.285** (.05)
After adoption X	.0232 (.044)	.00274 (.04)	.02 (.039)	.0173 (.024)	-.115 (.093)
Acquisition report					
After adoption X	.104* (.046)	.0608 (.038)	.0627+ (.036)	.0366+ (.022)	.129 (.079)
Benchmark report					
After adoption X	-.0473 (.045)	.0252 (.042)	.0257 (.039)	-.0518* (.025)	-.0183 (.095)
Executive report					
After adoption X	-.122** (.042)	-.111** (.036)	-.141** (.036)	-.00162 (.023)	-.15+ (.091)
Insights report					
After adoption X	.0977* (.047)	.0625 (.042)	.0631 (.039)	.0255 (.023)	.146 (.1)
Revenue report					
Lag(DV)	.393** (.026)	.502** (.02)	.501** (.02)	.177** (.033)	.27** (.023)
Observations	19,425	19,425	19,425	19,425	16,815

All models include retailer and month-year fixed effects. Retailer FE are differenced out. Asymptotic SE robust to general cross-section and time series heteroskedasticity are reported in parentheses.

Lag(DV) is the lagged dependent variable.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

## E Heterogeneity results

Table W-17: Heterogeneity by revenue quartiles

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption X	.0553	.0208	.00894	.00666	.184+
Quartile 1	(.07)	(.066)	(.065)	(.02)	(.1)
After adoption X	.329**	.266**	.256**	.00275	.616**
Quartile 2	(.08)	(.066)	(.066)	(.026)	(.12)
After adoption X	.255**	.184*	.17*	.0323	.328**
Quartile 3	(.074)	(.07)	(.067)	(.023)	(.1)
After adoption X	.417**	.325**	.317**	.0358+	.532**
Quartile 4	(.095)	(.082)	(.082)	(.019)	(.13)
Observations	24,124	24,124	24,124	24,043	20,316

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-18: Heterogeneity by transaction quartiles

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption X	.142+	.0276	.0236	.0499*	.434**
Quartile 1	(.072)	(.056)	(.056)	(.022)	(.11)
After adoption X	.168*	.131*	.122+	-.00593	.276*
Quartile 2	(.072)	(.062)	(.063)	(.027)	(.1)
After adoption X	.303**	.228**	.219**	.0408*	.455**
Quartile 3	(.077)	(.066)	(.065)	(.018)	(.11)
After adoption X	.453**	.433**	.408**	-.0168	.477**
Quartile 4	(.097)	(.09)	(.09)	(.023)	(.13)
Observations	24,124	24,124	24,124	24,043	20,316

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-19: Heterogeneity by platform

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption X Shopify	.361** (.056)	.256** (.047)	.247** (.046)	.0299* (.014)	.564** (.072)
After adoption X Paypal	.0489 (.062)	.0531 (.062)	.0633 (.062)	-.0124 (.029)	.0765 (.085)
After adoption X Stripe	.368* (.15)	.335* (.13)	.268+ (.13)	.0266 (.039)	
After adoption X Amazon Seller	.111 (.2)	.0805 (.19)	.0789 (.19)	.0284 (.036)	
After adoption X BigCommerce	.0969 (.15)	.0168 (.13)	-.14 (.11)	.0744 (.049)	
After adoption X Magento	1.83 (1.7)	1.45 (1.4)	1.27 (1.3)	.249 (.17)	
After adoption X Amazon Vendor	-1.55* (.67)	-.942** (.29)	-.918** (.29)	-.249 (.28)	
Observations	24,124	24,124	24,124	24,043	20,316

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-20: Heterogeneity by platform

Dependent Variable	log(revenue+1)	log(no. transactions +1)	log(no. unique customers+1)	log(basket size+1)	log(repeated revenue+1)
After adoption X	.224*	.222**	.201*	-.024	.275*
Clothing & Fashion	(.08)	(.076)	(.077)	(.028)	(.12)
After adoption X	.496**	.412*	.402*	.0215	.924**
Health & Beauty	(.18)	(.15)	(.15)	(.029)	(.24)
After adoption X	.35*	.298*	.304**	.0076	.325
Food & Drink	(.13)	(.11)	(.1)	(.041)	(.19)
After adoption X	.331*	.172	.139	.0423	.497*
Home & Garden	(.14)	(.12)	(.11)	(.045)	(.22)
After adoption X	.247	.193	.185	.0317	.312
Jewelry & Accessories	(.15)	(.13)	(.14)	(.037)	(.2)
After adoption X	-.0107	-.0375	-.0452	-.00379	.297
Electronics	(.13)	(.13)	(.13)	(.049)	(.18)
After adoption X	.0549	.00272	.00138	.0185	.204
Sports & Recreation	(.13)	(.11)	(.11)	(.04)	(.18)
After adoption X	.127	.0661	.0624	.0651	.0942
Toys & Games	(.16)	(.14)	(.14)	(.052)	(.22)
After adoption X	.281**	.182**	.173**	.0442*	.478**
Other	(.072)	(.061)	(.061)	(.019)	(.099)
Observations	24,124	24,124	24,124	24,043	20,316

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

## F Mechanism: Additional Results

Table W-21: Proportion of technology adoption

Dependent variable	Advertising proportion	Analytics and Tracking proportion	Payment proportion	E-commerce proportion	E-mail hosting proportion
After adoption	.0041* (.0017)	.0035* (.0014)	.00047 (.00087)	-.00027 (.00095)	-.00404** (.0014)
Observations	19,425	19,425	19,425	19,425	19,425

All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.

\* significant at 5%; \*\* significant at 1%; + significant at 10% level.

Table W-22: Proportion of advertising and analytics and targeting technology adoption

Dependent variable	CRM proportion	Website proportion	Prospecting proportion	Personalization proportion	Other Adv. proportion	Uncategorized proportion
After adoption	-.0016 (.0025)	-.014** (.004)	.013** (.004)	.004 (.003)	.004 (.003)	-.005+ (.002)
Observations	19,425	19,425	19,425	19,425	19,425	


All models include retailer and month-year fixed effects.

SE in parentheses are two-way clustered by retailer and month.


\* significant at 5%; \*\* significant at 1%; + significant at 10% level.



Figure W-1: Additional average outcomes for the 2016 and 2017 cohorts before adoption



(a) Customers



(b) Basket size



(c) Repeated revenue

The figures plot average outcome for the 2016 cohort (those retailer-platforms that adopted the service in 2016) and for the 2017 cohort (those retailer-platforms that adopted the service in 2017) in the period before adoption (between July 2015 and November 2016). In 2015, the figure plots the averages of all the retailers that adopt in 2016 and in 2017. Once retailers adopt the service they are dropped from the plot.