

The Value of Descriptive Analytics: Evidence from Online Retailers

Ron Berman
Ayelet Israeli

Working Paper 21-067



The Value of Descriptive Analytics: Evidence from Online Retailers

Ron Berman

The Wharton School, University of Pennsylvania

Ayelet Israeli

Harvard Business School

Working Paper 21-067

Copyright © 2020 by Ron Berman and Ayelet Israeli.

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Funding for this research was provided in part by Harvard Business School and Mack Institute for Innovation Management.

The Value of Descriptive Analytics: Evidence from Online Retailers*

Ron Berman[†]

Ayelet Israeli[‡]

November 2020

*We are grateful for helpful comments from Eva Ascarza, Kathy Li, Simha Mummаланeni, Davide Proserpio, Dave Reibstein, 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.

[†]The Wharton School, ronber@wharton.upenn.edu.

[‡]Harvard Business School, aisraeli@hbs.edu.

The Value of Data Analytics: Evidence from Online Retailers

Abstract

Does the adoption of descriptive analytics impact online retailer performance, and if so, how? We use the synthetic control method to analyze the staggered adoption of a retail analytics dashboard by more than 1,000 e-commerce websites, and find an increase of 13–20% in average weekly revenues post-adoption. We demonstrate that only retailers that adopt and use the dashboard reap these benefits. The increase in revenue is not explained by price changes or advertising optimization. Instead, it is consistent with the addition of prospecting and personalization technologies to retailer websites.

The adoption and usage of descriptive analytics also increases the diversity of products sold, the number of transactions, the numbers of website visitors and unique customers, and the revenue from repeated customers. In contrast, there is no change in basket size.

Put together, these findings are consistent with an indirect effect of descriptive analytics, and they suggest that the adoption of analytics serves as a monitoring device that allows retailers to assess the value of new technologies. Without using the descriptive dashboard, retailers are unable to reap the benefits associated with these technologies.

Keywords: descriptive analytics, big data, synthetic control, e-commerce, online retail.

JEL Classification: C55, L25, M31.

1 Introduction

A common advice for marketers is to base decision making on careful data analysis to generate better marketing outcomes. When surveying marketers, previous research by Germann et al. (2013) and Germann et al. (2014) has shown a positive association between deploying marketing analytics and performance. However, despite such evidence and other research that points to the benefit of deploying analytics solutions (Brynjolfsson et al. 2011a, Brynjolfsson and McElheran 2016, Müller et al. 2018), little is known about how these benefits are generated, how marketers use the information from analytics dashboards, and how customers change their behavior after a marketer deploys an analytics solution.

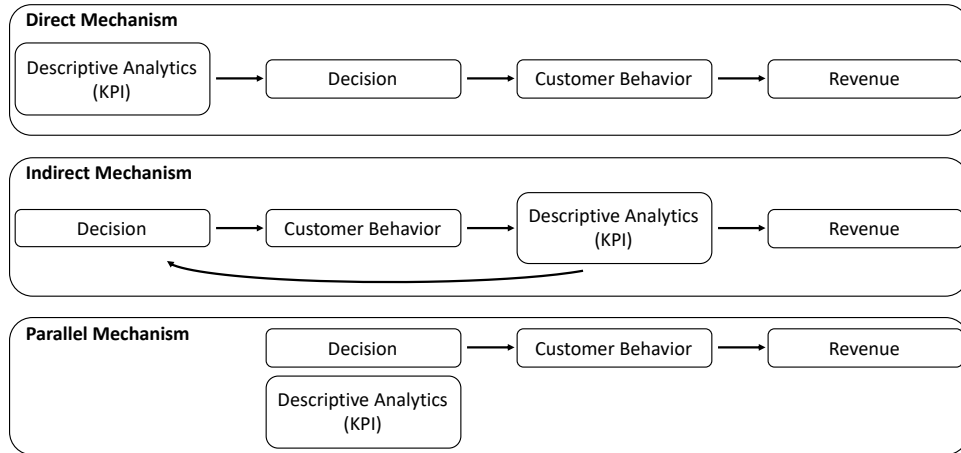
Analytics solutions are often categorized into one of four categories (Lismont et al. 2017): (i) descriptive (what happened), (ii) diagnostic (why did it happen), (iii) predictive (what will happen next), and (iv) prescriptive (what should be done about it). Most firms that adopt these solutions often deploy them first in the sales and marketing functions of the organization (Bughin 2017), and among marketing practitioners, descriptive analytics and simple key performance indicators (KPI) dashboards are the most popular choice (Delen and Ram 2018, Mintz et al. 2019). This preference for descriptive analytics stands in contrast to academic research, such as Pauwels et al. (2009), Hanssens and Pauwels (2016), Wedel and Kannan (2016), and Bradlow et al. (2017), that emphasizes the value of sophisticated modeling to derive insights from analytics.

In this paper we test whether there is a benefit from adopting descriptive analytics for online retailers and analyze these benefits. We focus on a causal estimate of the change in revenue that retailers experience when they adopt descriptive analytics, followed by an analysis of the decisions that retailers make as a result of the adoption and how these decisions drive changes in customer behavior. Our results should shed light on the additional capabilities that descriptive analytics enable for marketers. In practice, although descriptive analytics solutions are very popular among marketers and retailers, there is little clarity on how to interpret descriptive metrics and turn them into actionable insights. This raises questions about the value of this type of analytics. Answering these questions is useful for both practitioners and researchers. For practitioners, we provide a benchmark on the benefits they might expect when using a descriptive dashboard, and we illustrate how to potentially best extract these benefits. For researchers, our results shed light on the importance of simple heuristics and provide insights on the mechanisms that drive the

performance gains from these heuristics.

How might descriptive analytics benefit retailers? As mentioned above, there is no clear evidence that descriptive analytics helps at all, as most prior research did not have appropriate data or could not distinguish between the types of analytics used by companies. If descriptive analytics do provide a gain for retailers, there are a few potential mechanisms that can lead to an increase in revenue. We illustrate these in Figure 1 and aim to disentangle between these different mechanisms in our analyses. The first is a direct effect, where the descriptive KPIs are used to directly drive decisions. A direct effect is possible when it is clear what actions should be taken after observing some metric or outcome on the dashboard and when this action can be easily taken by the retailer. As an example, one common KPI retailers use is the cost of customer acquisition (CAC) or a similar measure of advertising cost per sale. If this metric is high compared to the profit margin of the retailer, it is clear to the retailer that they should optimize their advertising to reduce their CAC. Another example of a direct effect occurs when a retailer observes a low conversion rate of visitors into buyers in their online store. If they lower their prices, which they can do directly with ease, the conversion rate will increase.

Figure 1: Potential mechanisms for the effect of descriptive analytics



The second route through which descriptive analytics can operate is an indirect mechanism. In this case, the descriptive dashboard is used as a monitoring tool to assess the value of other operations of the retailer, but it does not drive decisions directly. For example, if the retailer invests in a retargeting campaign, or in website personalization for customers, it would be difficult

to measure the effects of these investments without monitoring KPIs that may change due to these actions. That is, descriptive analytics are useful to help the retailer identify which actions are most valuable. In this case, the dashboard is used to assess and determine which strategies work best and are complementary to other decisions taken by the retailer. Without it, the retailer is unable to properly optimize these decisions and reap the full benefit potential.

Finally, descriptive analytics might not operate at all to generate any value for the retailer but may be correlated with other actions the retailer takes. If, for example, when the retailer integrates a descriptive dashboard they also take simultaneous, unrelated actions (such as changing the store design or hiring new managers), then we might observe a correlation between increased performance and analytics adoption. If this is the case, we will observe that descriptive analytics are associated with improved performance, but we will not be able explain what decisions the retailers take or how customer behavior changes as a result. We refer to this mechanism as a “parallel mechanism.”

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, one confounder is that more productive firms might have more funds available to invest in 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 the following: firm adoption and usage of analytics, the resulting firm output and customer behavior, and 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 yields diminishing returns (Bajari et al. 2019). Despite this initial positive evidence, even with primary data an additional challenge is the lack of random assignment of firms into adopting and using analytics, which makes causal interpretation difficult. One exception is Anderson et al. (2020) that uses a field experiment in the context of small physical mom-and-pop stores in Rwanda and also finds benefits to descriptive analytics.

To overcome these challenges, we use detailed panel data from over 1,000 online retailers in

a variety of industries that adopted a descriptive marketing analytics dashboard. For the main analysis, we develop a novel adaptation of the Synthetic Control Method (SCM) for scenarios with multiple treated units and varying (staggered) treatment times. We also carefully confirm the robustness of our results using multiple empirical methods. Our data is unique in that it also contains detailed observational data on retailers’ behavior and outcomes both before and after the adoption of the analytics service. In particular we observe metrics of the retailer’s actions (such as price changes and technology integration), as well as transaction level data of individual customers.

The three mechanisms we described (direct, indirect, and unrelated mechanisms) provide us with testable empirical predictions to understand how analytics operate to help retailers. Our data along with the adapted SCM allow us to provide evidence to investigate these mechanisms. We note that our data does not allow us to perform a full causal mediation analysis, and we therefore cannot interpret all results as fully causal. Throughout the paper we indicate which assumptions might be needed for a causal interpretation of the evidence, and which results are consistent with the theory but might require further research.

For the direct mechanism, we would expect to see changes in retailer decisions that are unrelated to the integration of additional technologies, and we would expect to see these changes only for retailers that make use of the analytics dashboard (users) versus those that adopt analytics but do not login to look at the reports (non-users).

For the indirect mechanism, we would expect to see that the highest gain from adopting analytics appears for retailers that adopted additional technologies, but this gain will only happen if they also made use of the descriptive analytics dashboard. That is, if descriptive analytics is not complementary to the other technologies, we would not expect to see a difference in the benefit generated among users and non-users of the analytics dashboard. Other evidence that would be consistent with an indirect mechanism are changes in customer behavior that are predicted by the addition of specific technologies but that are unrelated to other actions of the retailer.

Finally, for a parallel mechanism we would observe that adopting analytics presents a gain in performance regardless of usage of the dashboard, and regardless of what other actions the retailer takes as a result of adopting analytics. If a parallel process is the mechanism behind the effect, the effect will also disappear when we instrument for the timing of adoption.

Three major findings emerge from our analysis regarding (i) the overall effect of analytics

adoption, (ii) the mechanisms through which the benefits are gained, and (iii) the changes in customer behavior once retailers make decisions based on analytics. First, we find causal evidence for a positive main effect of adopting descriptive analytics technologies on retailer revenue. In our sample, adoption of an analytics service is associated with an increase of 13% to 20% in average weekly revenues of the firm. The results are robust to using multiple methods that include a staggered difference-in-difference (SDD), adapted Synthetic Control (SCM), and instrumental variables (IV). We also find that the smallest retailers (in terms of revenues or transactions) benefit from adoption, while larger retailers do not.

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. We do not observe that firms made direct changes to pricing or advertising, but they did invest in adopting new technologies, and particularly prospecting and personalization technologies. Among those who invested in these technologies, we only observe gains in revenues for firms that actively used the analytics dashboard. Taken together these results suggest that the indirect mechanism of descriptive analytics may be in operation, and that the value of descriptive analytics is as a device that allows firms to assess the value of other technologies, but it does not necessarily drive direct decisions by retailers. This finding is quite interesting, because it shows that marketing actions beyond pricing or advertising changes also have a large potential value.

The third major finding is that the adoption of analytics results in an increase in the number of transactions, number of new website visitors, number of unique customers, revenue from repeated customers, and the number of unique products purchased, but did not result in increased basket size per transaction or reduced CAC. Again, all of these changes only occur for those retailers that used the analytics service, and not for those that adopted but didn't use it. As we explain in detail in Section 5, these changes in customer behavior give more credence to the indirect mechanism through which analytics affects retailer revenue. These results suggest that retailers that adopt descriptive analytics mostly gain from monitoring of complementary investments in additional technologies. Particularly, for our sample, we observe that descriptive analytics enables retailers to gain benefits from personalization and prospecting technologies.

The theory suggests that 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). Additionally, personalization efforts increase the likelihood of existing customers to try new products, thereby also resulting in more diverse purchases. We note that while retailers improve their ability to attract more customers to their website, and increase repurchase rates among existing customers, the basket size does not change. In this sense, descriptive analytics with complementary technologies affects the extensive margin of customer revenue, but not the intensive margin of profit per customer.

2 Institutional Background

The analysis focuses on online retailers from a variety of industries that operate their own online stores.¹ Many of the retailers in our dataset manufacture and sell their own brands 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.

Our data comes from an analytics service provider that offers a popular analytics dashboard for online retailers. The software collects and analyzes data from the retailer’s online store, as well as from other payment and fulfillment channels such as Amazon, Paypal, and Stripe (for sellers who sell through Amazon, or use Paypal or Stripe as a checkout mechanism).² The analytics service was launched 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 dashboard was ranked as one of the top analytics dashboards for online stores, giving it substantial visibility among retailers. Once installed, the software collects the retailer’s data after installation, as well as historical data. The data is collected daily, at the transaction level, going backwards for up to 24 months before the installation of the dashboard. 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.

The dashboard uses the transaction level data to generate metrics and visual reports on ag-

¹Fewer than 5% of these retailers also sell on Amazon.com as third-party sellers.

²For confidentiality reasons, we cannot provide identifying information about the analytics service provider.

gregate sales, average basket sizes, share of repeated customers’ revenues, cost of new customer acquisition, the average customer lifetime value, and many other metrics. Depending on data availability (i.e., whether the retailer also uses Google Analytics or Facebook advertising) more than 20 descriptive metrics are calculated and displayed on a weekly or a monthly level. Compared to basic data reports from Google Analytics or e-commerce hosting providers, the benefits of the analytics service is that it integrates and aggregates data from the retailer’s different data sources. However, the analysis is primarily descriptive, providing retailers a view into their performance. Nearly all of the retailers in our data have been using Google Analytics prior to adoption of the analytics dashboard.

The dashboard presents information and metrics in five main reports: (i) customer acquisition costs, (ii) revenue by hosting platform report, (iii) benchmark reports that compare the retailer’s metrics to a set of benchmark retailers selected by industry and revenues, (iv) executive report that summarizes all reports, and (v) an insights report that uses an algorithm to provide individualized recommendations. Although the insights report has a prescriptive flavor, it was in early development during the time window of the data, and the recommendations were not very useful. For example, a common recommendation was “reduce customer acquisition costs by x%” or “increase the number of visitors to the website by y” without additional details. Additionally, our data also contains detailed information about each retailer’s access to each report over time, and the insights report was the least frequently viewed report.

3 Data and Sample Construction

We use data from all retailers that adopted (signed-up for) the analytics service in 2016 and had at least one year of annual revenue of USD 100,000 or more in the data. We ensure that each retailer has at least 12 monthly observations, out of which at least one observation occurs before signing-up, one observation at the month of adoption, and one observation after adoption. This yields a total of 836 retailers who serve as a focal treated group. Some of the retailers use multiple online channels to sell their products (e.g., their own website and Amazon), which yields 964 distinct retailer-channel combinations. The data is aggregated to a monthly level. Because retailers adopt the analytics service in a staggered manner, the data of retailers that adopted the service in 2016 can 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 include only retailers that had historical data that included at least one observation in 2015. This yields 328 companies and 425 retailer-channel combinations out of a total of 958 companies that adopted the service during 2017.

Figure 2 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 1,164 retailers in our data is described in Table 2. Table 3 displays the distribution of the different hosting platforms for each of the 1,389 retailer-channel combinations in our data. Most of the retailers in the sample operate their own website, while roughly 5% of the retailers sell on Amazon either as a vendor or a seller. Shopify, BigCommerce, and Magento represent 63% of the retail-channels in our sample. These platforms held together 36% market share in 2016 among the top 8 e-commerce hosting platforms for independent sellers.³ Among the top 100k web-stores, they held 37% in 2017.⁴

Figure 2: Treatment variation plot




Figure 2 Treatment variation plot (Imai et al. 2018). Each horizontal line corresponds to a retailer—channel combination. Blue regions represent units before adoption of the analytics dashboard. Dark gray regions represent units after adoption of the analytics dashboard. Light gray regions represent units with no data in the time period.

We augment the retailer data with data from four additional sources: (i) data about non-

³<https://www.engadget.com/2016-11-03-ecommerce-platform-market-share-looking-at-the-companies-that-d.html>

⁴<https://aheadworks.com/blog/ecommerce-market-2017/>

Table 1: Summary statistics

Variable	Mean	Std. Dev.	Median	N
Panel A: Retailer level				
No. obs. per retailer	51.63	21.89	45.00	1,164
Avg. weekly revenue (USD)	12,545.25	31,359.89	4,741.05	1,164
Avg. no. weekly transactions	157.84	492.04	53.68	1,164
Avg. no. unique customers	143.59	456.21	46.05	1,164
Avg. basket size (USD)	182.61	372.58	87.38	1,164
No. channels per retailer	1.23	.50	1.00	1,164
Panel B: Observation level				
Avg. weekly revenue (USD)	12,732.18	37,006.22	4,268.11	38,074
Avg. no. weekly transactions	167.13	732.86	45.75	38,074
Avg. no. unique customers	152.42	637.28	40.00	38,074
Avg. basket size (USD)	174.75	397.47	82.17	38,074

subscribed retailers that visited the analytics service website; (ii) data about usage (login times, reports viewed) of the analytics service by retailers; (iii) data on additional technologies the retailer installed on their online store, such as advertising tracking and email tracking (collected via Builtwith.com); and (iv) data on historical keyword advertising collected via Spyfu.com, which includes periodic 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 provides evidence for gains in retailer revenues due to analytics adoption. Once we demonstrate that analytics adoption yields a positive gain in revenues, in Section 5 we analyze the mechanism behind this gain. First, we look at actions that firms take, from changing prices to adding technologies to their websites. Then, we analyze changes within and across the customer base of the retailers and describe their contribution to the observed impact on revenues.

There are three main challenges to the identification of the effects that stem from the fact that the analytics service was not randomly assigned to retailers. The retailers chose whether to adopt analytics, which may introduce selection bias to the analysis. Once choosing to adopt analytics,

Table 2: Distribution of industries

Industry	Frequency	Percent
Clothing & Fashion	288	24.7%
Health & Beauty	134	11.5%
Food & Drink	86	7.4%
Home & Garden	81	7.0%
Jewelry & Accessories	77	6.6%
Electronics	71	6.1%
Sports & Recreation	62	5.3%
Toys & Games	36	3.1%
Other	329	28.3%
Total	1,164	100%

Table 3: Distribution of hosting platforms

Hosting platform	Frequency	Percent
Shopify	813	58.5%
Paypal	359	25.8%
Stripe	97	7.0%
Amazon Seller	59	4.2%
BigCommerce	48	3.5%
Magento	9	0.6%
Amazon Vendor	4	0.3%
Total	1,389	100%

the retailer could choose the best timing for adoption, which might upward-bias our estimates due to endogeneity of adoption timing. Further, the retailer could also implement other policies concurrently with adopting analytics, creating an omitted variable bias that would also upward-bias our estimates.

We employ a fourfold approach to identify the main effects of interest. First, as described in section 3, we construct a control group that is used to predict the potential outcomes of the adopting firm as if it did not adopt analytics. 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 a staggered difference-in-difference (SDD) analysis with time and retailer fixed effects to net out the effects of time trends and retailer-specific growth trends unrelated to the adoption of analytics. Third, we employ a staggered synthetic control method (SCM) that ensures appropriate pre-treatment trends between treatment and control units. Fourth, we perform additional analyses to alleviate concerns about selection of adoption timing and omitted confounders. For endogenous adoption timing, we construct two instrumental variables that are likely shifters of adoption timing but are plausibly unrelated to the retailer’s revenue. For omitted confounders, we examine the effect on revenue for retailers that have used the service after adoption versus those that did adopt the service but did not use it.

Because all of the retailers in our data adopt the analytics service eventually, they may differ from unobserved non-adopters. Our estimates are therefore the effect on retailers that choose to

Table 4: Summary statistics—Additional data

Variable	Mean	Std. Dev.	Median	N
Avg. monthly no. logins	.34	.58	.15	32,321
Monthly indicator for any report views	.10	.3	0	32,321
Monthly no. technologies	56.3	23.1	53	31,110
Google avg. monthly ad costs (USD)	2,502.5	12,316	0	17,959
Spyfu monthly advertising spending (USD)	1,818.5	13,061.3	397.7	9,164
Spyfu monthly no. advertising keywords	29.1	46.7	9	5,151

The observations in this table are retailer-month level, since these data are collected at the retailer level. Some data sources do not include all of the retailers in our sample. This leads to a smaller number of observations for the respective variables.

adopt the service, i.e., the local average treatment effect (LATE) and not the average treatment effect (ATE) in the population. Despite this potential limitation, we believe that the LATE is a relevant and appropriate measure to focus on because gaining value from analytics requires engaging with the reports and making data driven decisions, which are all endogenous decisions. Even if retailers are exogenously assigned to adopt a descriptive analytics solution, we do not expect to see any benefit for those that do not use it.⁵ Accordingly, we measure LATE and expect that similar retailers that are interested in adoption of descriptive analytics will exhibit similar outcomes to those in our sample.

4.1 Effect of analytics adoption on retailer revenue

4.1.1 Staggered difference-in-difference (SDD)

We start by analyzing the impact of analytics adoption on average weekly revenue using classic staggered difference-in-difference. We estimate the following SDD model using OLS:

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

where Y_{ijt} is average weekly revenue for retailer i and channel j in month t . **AfterAdopt** indicates whether the retailer has adopted analytics at or before time t . We control for retailer-channel fixed effects f_{ij} and month fixed effects g_t . We use two-way clustering of standard errors by retailer

⁵For example, even research such as Anderson et al. (2020) that uses randomized controlled field experiments recruits entrepreneurs based on their growth potential, and thus estimates a LATE rather than an ATE.

and month to address serial correlation (Bertrand et al. 2004). Most retailers sell only using one channel—typically their own website. However, we use retailer-channel observations (i.e., a retailer that sells on their own website and through Amazon has two observations at each time period) because retailers may adopt the service at different times for different online channels. Results do not substantially differ when we aggregate the data to the retailer level.

The coefficient β measures the change in average weekly revenue after the adoption of analytics. It is identified by comparing the change in revenue before and after adoption for adopting retailers with the change in revenues in the same time periods for retailers that haven’t yet adopted the service, similar in spirit to Wang and Goldfarb (2017). 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. The identifying assumption in our SDD analysis is that there were no differential trends in revenues before adoption between retailers that adopted the service and those that did not. We also note that it is unlikely that those retailers that haven’t yet adopted the service are indirectly affected by those that adopted the service because adoption is not observable by other firms and there is little competition between the retailers in the sample.

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. Due to this data truncation, we limit our SDD analysis to include only observations between July 2015 and up to 18 months after adoption for each retailer.

We first examine the identifying assumption of parallel trends between treatment and control units prior to adoption. Figure 3 presents average revenue for future adopters from the 2016 and 2017 cohorts over time and demonstrates visual evidence for similar trends before adoption. We also estimate the following OLS model to statistically test the identifying assumption:

$$\log(Y_{ijt} + 1) = f_{ij} + g_t + \sum_{k \neq -K, -1} \gamma_k D_{ijt} + \epsilon_{ijt} \quad (2)$$

where f_{ij} and g_t are, respectively, retailer-channel and time fixed effects, as before. γ_k captures the effect of analytics adoption in time k relative to adoption. For example, γ_1 captures the effect one month after adoption, and γ_{-2} captures the effect two months prior to adoption. D_{ijt} indicates whether the time t equals k periods relative to adoption for retailer-channel ij . Following Boryusak and Xavier (2018), we exclude the minimum lag $-K$ and the lag of $k = -1$ from our analysis, for identification purposes. This exclusion causes the baseline of the comparison to be one month

before adoption, and the value at $k = -1$ is set to zero. We use up to 35 months of data per retailer, i.e., $K = 17$.

Figure 3: Average weekly revenues for the 2016 and 2017 cohorts before adoption




Figure 3 plots average weekly revenues of adopters in 2016 (solid red) vs. 2017 (dashed blue) before adoption over time. Each line displays average revenues of retailers that have not adopted yet in the specific month.

Figure 4 plots the coefficient estimates $\hat{\gamma}_k$ and the 95% confidence intervals for the monthly indicators six periods before and after adoption. Time 0 in the figure is the month of adoption. As can be seen in the figure, the coefficients are statistically indistinguishable from zero in the two periods before adoption. However, there seems to be a slight increasing trend in the coefficients in the pre-adoption period. Specifically, three months before adoption and earlier the treatment group had lower revenue than the control group, and this trend reverses after adoption. The figure also reveals that the main effects decay over time (although, as indicated by Boryusak and Xavier (2018), this regression does not estimate the treatment effects efficiently, and equation (1) should be used instead).

The pre-trends analysis raises two concerns about the validity of the SDD analysis. First, the parallel trends assumption required for identification of equation (1) might not hold. Second, recent research about interpretation of SDD designs has shown that the resulting ATE estimates are a weighted average of the ATEs of different cohorts (with weights potentially being negative), which may cause a bias and are difficult to interpret (Goodman-Bacon 2018, de Chaisemartin and d’Haultfoeuille 2019, Callaway and Sant’Anna 2020). We therefore proceed by reporting the SDD estimates for completeness, and we then implement a synthetic control approach for the analysis of the effect of adoption, which provides a solution to these issues.

Figure 4: Revenue before and after adoption




Figure 4 plots coefficients and 95% confidence intervals obtained by estimating equation (2).

Column 1 of Table 5 presents the results of the estimation of equation (1). The results suggest that after adopting the analytics service, retailers exhibit an average increase in weekly revenues of roughly 18% [95% CI: 10%, 28%]. The unconditional median average weekly revenue before adoption in our data is USD 3,704. Therefore, the median retailer experienced an increase of approximately USD 684 per week.

4.1.2 Aggregated staggered synthetic control

The synthetic control method (SCM) uses a weighted average of outcomes from control retailers to predict the outcomes of the adopting retailer “as-if” they did not adopt the analytics service. The weights are chosen to optimally match the outcomes of the adopting retailer *pre-adoption*, and thus capture any possible trends that might affect identification. The difference between the observed outcomes post-adoption and the predicted outcomes are the estimated treatment effects from the method (Abadie et al. 2010, Abadie 2020).

Synthetic control analysis is often applied to a single or small number of treated units with a small number of control units and long balanced panels. Analysis with multiple treated units and staggered adoption is challenging, as described in Ben-Michael et al. (2019), and currently no method exists for unbalanced panels with staggered adoption and multiple treated units as in our case. Part of our contribution is to extend the single treated unit method described in Doudchenko and Imbens (2016) to allow for multiple units with staggered adoption and unbalanced panels.

The analysis is performed in cohorts of adoption month and then aggregated to provide a cumulative ATE for time from adoption. In our data, we focus on 12 cohorts of adopting firms,

from January 2106 (cohort 1) to December 2016 (cohort 12). The data window for each retailer in the analysis consists of ℓ lagged time periods pre-adoption and f leading time periods post-adoption. Each retailer thus contributes $\ell + f + 1$ observations to this analysis. Control retailers are those that have $\ell + f + 1$ observations during the analysis window and adopted analytics outside of the analysis window (i.e., at least $f + 1$ periods after the cohort’s adoption time). This means that every cohort of adopting firms potentially has different retailers contributing to its synthetic control. For example, for cohort adoption time 1 with $\ell = 8$ and $f = 5$, control units are those that adopted at least six months after period 1, and have 14 continuous observations, from eight periods before the adoption to five periods after. Treated units are all units that adopted analytics in the cohort’s adoption time and have ℓ lags and f leads.

We denote the time period of adoption as $t = 0$, the time periods pre-adoption (lags) as $t = -\ell, \dots, -1$, and the time periods post-adoption (leads) as $t = 1, \dots, f$. We assume the control group for the specific adoption cohort is indexed by $i = 1, \dots, n$.

To generate the synthetic control, we use the following procedure for each treated retailer in each cohort. For every retailer j in the treated group, we use only the ℓ observations in the time periods pre-adoption, and we estimate:

$$\log(Y_{jt} + 1) = \alpha + \sum_{i=1}^n \omega_{ij} \log(Y_{it} + 1) + \epsilon_{jt} \quad (3)$$

We do not restrict the weights ω_{ij} to be positive or sum to 1; rather we use a regularized elastic-net regression (Zou and Hastie 2005). This allows the method to control for retailer fixed effects and level differences in outcomes between the treatment and controls. The tuning parameters for regularization are selected using cross validation over the control units as described in Doudchenko and Imbens (2016).

The estimated treatment effect $\hat{\tau}_{jt}$ of each treated unit j at time period t is then:

$$\hat{\tau}_{jt} = \log(Y_{jt} + 1) - \left(\hat{\alpha} + \sum_{i=1}^n \hat{\omega}_{ij} \log(Y_{it} + 1) \right) \quad (4)$$

In this estimation, $\hat{\alpha} + \sum_{i=1}^n \hat{\omega}_{ij} \log(Y_{it} + 1)$ is the predicted counterfactual outcomes of treated unit j “as if” it didn’t adopt analytics at times $t \geq 0$. If the synthetic control is able to match trends in the outcomes of the treated unit, the estimated treatment effects $\hat{\tau}_{jt}$ pre-treatment ($t < 0$) should be close to zero. If the adoption of analytics has (a positive) impact on the firm’s performance, we

expect that additionally, at $t \geq 0$, the estimated treatment effects $\hat{\tau}_{jt}$ would be positive. Positive effects post-treatment show that the predicted placebo outcome without adoption and the true observed outcomes are different.

To estimate the standard errors of $\hat{\tau}_{jt}$, we apply the placebo method from Doudchenko and Imbens (2016). For every control unit i , we use it as a treated unit in the analysis above, while the rest of the control units are used to create a synthetic control for it. Because the analysis is done on control units that did not adopt analytics in the time window analyzed, any estimated effect is attributed to noise, and the variance of this noise across control units estimates the standard error of a treatment effect under the null hypothesis of no effect.

The outcome of this analysis is a treatment effect estimate $\hat{\tau}_{jt}$ and a standard error estimate \hat{se}_{jt} for every retailer j in relative time period t between $-\ell$ and f . To aggregate the retailer level estimates into an estimate for each time period, we use random effects meta-analysis (Borenstein et al. 2011), where for each time period t , we calculate a weighted mean of the individual retailer treatment effects, weighted by the inverse of their variances. We elected to use random effects meta-analysis over fixed-effects meta-analysis because the estimates were practically identical, but the random effects method produced a more conservative (wider) confidence interval. For robustness we also computed estimates using unweighted means and the jackknife for confidence intervals, which generated similar results.

We apply this aggregate staggered SCM to analyze the change in revenue post-adoption, using eight-month lags and five-month leads (total of 14 months). Because this methodology ensures that each retailer-channel has observations in each of the 14 time periods, we have fewer treated units than in the SDD analysis. However, because SCM was developed initially for the case of only one treated unit, this methodology is appropriate for our sample. Across the cohorts that adopted in 2016, there are 733 retailers that have sufficient observations for the analysis. Table W-1 in the Web Appendix details the breakdown of the number of treated and control units used in each cohort. On average there are 50–60 treated units and 300 control units in each cohort.⁶

Figure 5 plots the estimated SCM ATEs of adopting analytics over time using both the meta-analysis method⁷ and the jackknife method, and present similar results (although the matching

⁶The control group is larger than the treated group because it is comprised of all later adopters from multiple cohorts (both in 2016 and 2017) and not just units from one cohort in 2016.

⁷Column 1 of Table W-2 in the Web Appendix reports the corresponding coefficients.

is a bit noisy in the jackknife method). Overall, the SCM procedure matches the pre-treatment revenues well, as the pre-treatment ATEs are indistinguishable from zero. Compared to Figure 4, we do not observe pre-adoption differential trends between treatment and control units, which should alleviate concerns about the validity of the identifying assumptions underlying the interpretation of the SCM results. The results show that there is a significant increase in revenues post-adoption. Six months after adoption, the cumulative increase in weekly revenues is roughly 15.1% [95% CI: 8.3%, 22.3%] (on average 13.4% per period). The treatment effect for each month estimates the total effect accumulated up to that month. Note that at period 5 the treatment effect is lower than in period 4, suggesting a decaying effect for revenues. However, all ATE are statistically significant and positive six months after adoption.

Figure 5: Staggered SCM ATE




Figure 5 plots the staggered SCM ATEs of adopting analytics on revenues over time. Time 0 indicates month of adoption. Other times are relative to adoption.

Section B of the Web Appendix presents multiple analyses that demonstrate robustness of our SCM estimates to different number of lags and leads used for SCM, aggregation of the data to the retailer level instead of retailer-channel level, and inclusion of early 2017 adopters as treated.

The estimates of the SCM analysis are slightly smaller in magnitude compared to the SDD analysis presented in the previous section. We report the six months cumulative ATE results of the SCM estimates in Column 2 of Table 5. Because SCM ensures appropriate controls and accounts for potential pre-adoption trends, we use SCM as the main analysis method in the remainder of the paper.

4.2 Addressing selection issues

We have presented evidence for an average increase of approximately 15% in weekly revenues once a retailer adopted the analytics dashboard. A limitation of the analysis includes potential selection-related threats to the validity of our estimates. In this section we use instrumental variables to alleviate concerns regarding endogenous timing of adoption, and we present evidence that adopting analytics without using the dashboard does not produce an observed increase in performance, which should alleviate concerns about unobserved confounding factors. The results in this section allow us to rule out the parallel mechanism described in Figure 1.

4.2.1 Instrumental variables to control for endogenous adoption timing

Because firms choose when to adopt analytics, they may choose to adopt at an opportune time. For example, a retailer may predict the benefit from adopting analytics and will choose to adopt it when the benefit is the highest. We address this issue using instrumental variables (IV). Specifically, we use data from the website of the *analytics service provider* to identify two time-varying exogenous factors that plausibly impact the timing of analytics adoption but that are uncorrelated with the revenue of the retailer. We view these IVs as time shifters of adoption.

Our IV strategy hinges on identifying times of increased awareness and attention to the analytics service among online retailers, and assuming that such increased awareness drives some of the adoption exogenously. We leverage traffic to the analytics provider’s website and not to the individual retailer websites. New visitors to the analytics provider website are likely online retailers that are potential clients of the service and not end consumers who buy items at the retailers’ website. Figure A-1 in the Appendix illustrates the traffic we leverage. The number of unique new visitors to the analytics provider’s website measures how many retailers are interested in the analytics service in that specific month. We expect increased attention to the analytics service to exogenously shift the timing of adoption for adopting retailers to be earlier than they would have adopted without increased attention. The variation in the number of new visitors to the analytics provider’s website stems from online mentions, articles, and advertising about the provider’s capabilities, and should be uncorrelated with the performance of a specific retailer. Hence, we also expect the exclusion restriction to hold for this instrument. In our measures, we use unique *new visitors* to the service provider website to identify traffic that arrived at the website to learn about the service or to sign

up.

First, we use the number of unique new visitors that arrived at the service provider’s website from a particular hosting platform (“referralVisits”) as an IV. This variable captures the hosting platform promotion or increase in awareness of the analytics dashboard. The analytics dashboard prominently appeared as a top popular service for retailers on some of the hosting platforms for a period of time. 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 via adoption. We did not include traffic from channels with targeted advertising because they could violate the exclusion restriction 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. Figure A-2 in the Appendix illustrates the variation in this instrument.

Second, using a similar logic, we construct an IV using the number of unique new visitors from a particular geographic region to the analytics provider’s website. In this case, this variable captures the effect of local media and “buzz” about the analytics service, 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 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.⁸ We use the state for U.S. retailers, and country information elsewhere. Roughly half of the companies are based in the U.S. with the majority in California. Table A-1 in the Appendix displays the distribution of regions of retailers in our sample, and Figure A-3 in the Appendix illustrates the variation in this instrument.

A potential threat to the validity of the exclusion restriction for the regional IV is country-specific performance waves that cause companies in regions that do well financially to both visit the analytics provider’s website and to sign-up. 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 knowledge. Additionally, our results are robust to adding region fixed effects.⁹ We also use the referral IV which is less likely

⁸We were unable to identify the location of 102 of the retailers. Thus we define this variable for 1,059 of the retailers, which represent 1,267 of the retailer-channel combinations.

⁹We do not include the 65 region fixed effects in our main regression specification due to efficiency concerns and

to suffer from hosting platform performance waves.

Our instruments are relevant to the endogenous independent variable, because traffic of new users from a region or a platform is expected to be correlated with adoption. Therefore, we are confident that these are plausible adoption time shifters. However, this restriction does not 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). The IV estimates therefore measure the treatment effect during the first month of adoption and not over a longer period of time. Note that, because the IV apparatus is used as the identification strategy, we no longer rely on pre-adoption matching of trends between treatment and control. This allows us to use all adopting cohorts in the IV analysis.

Our IVs are admittedly limited. Ideally, we would have retailer-channel-specific time-varying exogenous variation we could exploit. Instead, we use time-varying instruments (hosting platform traffic and country traffic) that apply to a subset of retailers equally to estimate retailer-channel-specific changes in performance due to adoption. We note, however, that “referralVisits” and “regionVisits” together create 200 distinct groups, with a median group including two retailers, and the largest group “California-Shopify” comprising roughly 12% of the observations. 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 the SDD and SCM methods.

Because the endogenous adoption variable is binary, we utilize 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-step procedure. The estimation procedure and results are detailed in the Appendix. The 2SLS instrument derived from the first stage passed the Stock-Yogo, Cragg-Donald, Anderson-Rubin, and Stock-Wright weak instrument tests.

Column 3 of Table 5 reports the IV results that confirm our findings so far and show a revenue increase of 19.6% [95% CI: 6.9%, 33.8%]. Note that since the IV approach only estimates effects for

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.

the month of adoption, the effect sizes for revenues are substantially larger than those obtained in the SDD analysis (displayed in Figure 4) and those computed in SCM analysis in section 4.1.2.¹⁰

Table 5: Effect of analytics adoption on retailer revenue

Methodology	Staggered DiD	Staggered SCM	IV (Third stage)
	(1)	(2)	(3)
After	.169**	.141**	.179**
Adoption	(.037)	(.031)	(.057)

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

4.2.2 Disentangling adoption and usage

A second issue we examine is a simultaneity bias in the estimate due to omitted variables, which we call the parallel mechanism. It is possible that firms that adopted the analytics service made other changes in parallel to adoption, such as changing their management team or their website design, and that 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.

We examine this issue using the unique access we received to data that includes the dashboard login times and the reports accessed by each retailer. We examine three different aspects of usage: any usage, intensity of usage, and report viewership. The access and login information is available at the *retailer level*, and not the retailer-channel level. Therefore, we aggregate observations to the retailer level and use the first time of adoption as the retailer’s adoption time. We use the SCM approach to analyze the differential treatment effects based on each of the different aspects of usage. Accordingly, we consider these tests as heterogeneity tests of the effect of adoption.

First, we identify “users” as those retailers that ever logged-on to the analytics service. Roughly 10% of retailers that adopted the service never used it, and we call them “non-users”. We then compare post-adoption treatment effects across the groups of users and non-users to examine whether

¹⁰To further evaluate these results, Table A-3 in the Appendix provides a direct comparison and reports the results of an OLS regression using the same observations and specifications as the IV regressions, without instrumenting for the Adoption indicator. The OLS effect size is less than half of the IV third stage estimate.

retailers that use the service exhibit different outcomes than those that do not.¹¹ Figure 6 reports the treatment effects of all time periods from eight months before adoption to six months after adoption. Before adoption, for both users and non-users most point estimates are statistically indistinguishable from zero (those that are different than zero are very close economically to zero). After adoption, only the point estimates of users are larger than zero.

Reassuringly, we find that our effects are driven exclusively by those retailers that use the analytics dashboard. Although we are not able to completely rule out simultaneity of actions by the retailer, a correlation between usage and performance is a necessary condition for analytics to have a causal effect. Because 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.

Figure 6: Users versus non-users




Figure 6 plots ATEs (and 95% confidence intervals) of adopting analytics on revenues over time, computed using random effects meta-analysis. Time 0 indicates month of adoption. Other times are relative to adoption. Column A of Table W-3 in the Web Appendix presets the corresponding coefficients.

Next, we examine whether the intensity of usage has a differential relationship with performance outcomes. For each retailer, we compute the average monthly number of events logged when using the dashboard. We then split the sample of all retailers by the median of the average post-adoption usage.¹² We use the SCM to analyze each half separately. Therefore, in this case, 2016 cohorts are matched to future cohorts within the same usage group. Column B of Table W-3 in the Web

¹¹While the number of non-users is relatively small, SCM methodologies are constructed to compute treatment effects even if only one unit is treated, as long as there are enough control units. Therefore, the method works well in this setting.

¹²We use median split rather than tertiles or quartiles to ensure a sufficiently large control group. Results are similar but noisier if we split the data into more quantiles.

Appendix presets the results, demonstrating that intensity of usage matters. Overall, retailers with a larger average number of monthly events reap additional increases in revenues compared to those with less activity. (For example, the effect size for revenues six months after adoption is .12 for the below-median usage group vs. .26 for the above-median group.) The divergence starts in later periods and might suggest that there is learning due to usage.

Finally, we examine the usage of reports. The service provides five main reports, as detailed in section 2. We split the data by those that use reports and those that do not. 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, and then for each retailer we compute the average number of monthly reports they examined. In the first test we compare SCM estimates between retailers that viewed at least one report to retailers that did not view any reports. In the second test, we again compare the revenue of firms with below- and above-median average numbers of reports viewed.

We find that those retailers that accessed reports exhibit higher increases in revenues compared to those that did not, and retailers that accessed reports more frequently had larger gains compared to those that accessed reports less frequently. Columns C and D of Table W-3 in the Web Appendix present the results for these SCM analyses.

To summarize, we find that the intensity of usage of the dashboard is associated with increased revenue. If retailers adopted other methods that increased their performance simultaneously to the adoption of the dashboard, we would not have found that the increase in revenues is associated with the dashboard usage.

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. We use the SCM approach to estimate heterogeneous effects by running the analysis separately for each size group.

To analyze heterogeneity of the effect based on firm size, we use two proxies: i) median split

of size based on the average monthly revenues in the six months before adoption of each retailer-channel; ii) median split of size based on the average number of transactions in the six months before adoption for each retailer-channel. Table W-4 in the Web Appendix presents the results. For both revenue and transaction medians, retailer-channels with below-median size exhibit a statistically significant increase in revenues after adoption, while the above-median retailers exhibit marginal and small increases or no statistically significant changes in performance after adoption. Therefore, we conclude that smaller retailers reap more of the benefits due to adoption. This effect might be expected with descriptive analytics—larger and well-established retailers might have already optimized their performance and have lower returns from this type of analytics.

5 Mechanism

Because we observe evidence of a positive effect of adopting descriptive analytics, we expect the adoption of analytics to drive better retailer decisions which would translate to changes in customer behavior. These changes in customer behavior would be associated with generating higher revenues. In this section we first examine which decisions retailers make after they adopt descriptive analytics, and we further examine whether the observed changes in customer behavior are consistent with the predicted changes from the actions taken. The analysis allows us to provide evidence for the potential mechanisms through which descriptive analytics operates to drive gains in revenue described in Figure 1. This analysis of decision making by retailers and customer behavior is part of the unique contribution of our paper.

We note that, optimally, we would like to perform a causal mediation analysis (Imai et al. 2010a;b, Pearl 2014) on the model described in the figure. However, our data does not provide enough power and exogenous variation to properly test these claims in a causal framework. Our results are therefore indications for consistency between our theory and observed behavior, yet more research with better data will be needed to fully explore it.

Descriptive analytics does not provide retailers with specific guidance on which actions to take. Even the availability of simple benchmarks, which may tell a retailer whether they are underperforming on some metric, does not prescribe how to improve that metric. One possible exception is customer acquisition cost (CAC)—if this value is higher than the desired threshold of the retailer, it is obvious they should change their advertising spending. It is less clear how this change should

be made (e.g., decreased budget or change in allocation). We therefore analyze the main observable actions that retailers can take to affect customer behavior: (i) changing prices; (ii) reallocating and optimizing advertising; (iii) in-store personalization; and (iv) driving traffic to the store.

If prices in the store change, or if personalization is improved, we expect the purchased baskets to change, either in size or in product composition. If advertising allocation or optimization changes, and if technologies that drive customers (new or repeating) to the websites change, we expect to see a change in the composition of customers that make purchases after adopting analytics.

5.1 Decisions driven by the adoption of analytics

To analyze the changes in observable retailer actions after adoption, we use the SCM approach. The application of the method is similar to that of Section 4.1.2, but in the current analysis the dependent variable is the level of the action taken (e.g., amount of discounts, price variation, or advertising allocation) during the observation window.

In a few cases, such as that of advertising and pricing, our data is limited to retailers that had this data available and made it accessible to the analytics service. This limits the analysis only to retailers who had connected their Google Analytics advertising data or Shopify at the time our data was collected, resulting in 87% of the retailers using Google Analytics for the advertising analysis and 58.5% of the retailers using Shopify for the pricing analysis. We have confirmed that the previous main results of the effect of analytics hold for each subsample.

For pricing, we analyzed the transaction data of the retailers and computed five measures as the dependent variable: i) the average price for all the products sold by a retailer in a particular month; ii) the average monthly product price 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; and (v) the average discount rate. None of these variables yielded statistically significant differences after adoption of the analytics service (for neither users nor non-users of the dashboard).

For advertising, we used the retailers' own data collected through Google Analytics, and estimated spending data which we collected through SpyFu.com. We look at (i) the total monthly advertising spend of a retailer reported by each of these sources; (ii) the monthly number of search terms each retailer targeted; and (iii) the number of new display ad copies each month. The

Google Analytics advertising spending data does not exhibit any statistical difference in spending after adopting analytics. This is the case for all retailers regardless of analytics usage.¹³ In contrast, the SpyFu estimates do exhibit a decrease in budget towards the end of our observation period (see point estimates in Table W-5 in the Web Appendix). Note that the SCM did not perfectly match the pre-adoption period budgets for users, which is likely due to the smaller coverage of retailers in the SpyFu data.¹⁴ We interpret these results as inconclusive evidence for changes in advertising spending. Potentially there is a decrease in the advertising budget in the long run, but it might be minor. For the two other measures of advertising allocation (number of search terms and new display copies) we do not find any statistical difference following the adoption of analytics.

Turning to analyze the changes to the retailer’s online store, we collected longitudinal data from Builtwith.com on the adoption of different web technologies by the retailers. Examples of web technologies include A/B testing, advertising retargeting, product recommendation, and personalization. The data covers 95% of the retailers in our dataset. We first examine whether the overall number of different technologies installed on a website increases with the adoption of analytics, as well as how the number varies by category, for technologies related to retail management: Analytics & Tracking (565 different technologies), Advertising (532), E-commerce (263), E-mail hosting (121), and Payment (71).

Table W-6 in the Web Appendix reports the SCM coefficients for overall technologies as well as the three leading technologies for both users and non-users.¹⁵ Although the non-users matching is noisier, we still detect increases in adoption of the Analytics & Tracking and E-commerce technologies also among that group. The results suggest that with the adoption of descriptive analytics, retailers increase the overall number of technologies installed on their online store, and in particular adopt more analytics & tracking, advertising, and E-commerce technologies.

To better understand retailers’ decisions, we further break down the analytics & tracking, advertising, and E-commerce categories into sub-categories of more specific technologies. We obtain six meaningful sub-categories: customer relationship management (CRM), website design and optimization, lead generation and prospecting, personalization, other advertising technologies, and

¹³We caution that the Google Analytics advertising data is sparser so it includes lower coverage of time periods.

¹⁴We were able to identify the monthly advertising budget for 644 of the retailers (55%) and the number of terms and display ads for 324 of the retailers.

¹⁵The remaining two categories do not show significant effects and are unreported in the interest of space.

other E-commerce technologies. Prospecting technologies are focused on attracting new users to websites, while personalization technologies are focused on optimizing the on-site experience of existing visitors and retargeting previous visitors.

Using the same approach as before, we analyze the number of technologies in each sub-category that retailers use as the dependent variable. We find an increase in prospecting technologies and in personalizing technologies, for both users and non-users. The other sub-categories do not exhibit significant changes after analytics adoption. Figure 7 displays the results for prospecting and personalizing technologies.

Figure 7: Sub-technologies Adoption

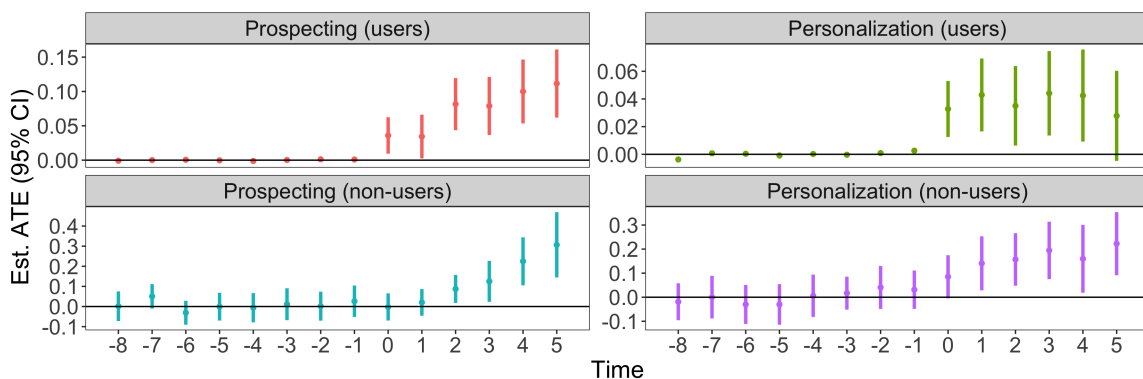


Figure 7 plots ATEs (and 95% confidence intervals) of adopting analytics on personalization and prospecting technologies over time, computed using random effects meta-analysis. Time 0 indicates month of adoption. Other times are relative to adoption.

We conclude that the adoption of descriptive analytics was associated with adoption of additional prospecting and personalization technologies. It was potentially associated with a decrease in advertising spending. There is no evidence of association with price changes, advertising allocation changes, or adoption of other technologies. Notably, the increase in the specific technology adoption exists for both adopters who used the dashboard (users) and adopters who did not use it (non-users). One explanation is that retailers were encouraged to adopt both a descriptive dashboard and personalization and prospecting technologies. However, the fact that we only observe an increase in revenue for adopters who are users implies that adding these technologies by themselves would not be a main contributor to a gain in revenue.

5.2 Changes to customer behavior

The observed increased usage of personalization and prospecting technologies is predicted to contribute to changes in customer behavior. Specifically, enhanced prospecting should increase the number of new visitors to the store and new customers (those visitors who convert) at the store. Improved personalization is predicted to increase repeat visits through retargeting, and the diversity of products purchased through recommender systems (Brynjolfsson et al. 2011b). If pricing were to change, we would have expected to see a change in basket size (the amount purchased) and potentially the number of products purchased. Advertising optimization would have affected the number of new visitors to the store and would also have shown an improved CAC.

We use the following dependent variables to test for changes in customer behavior which are predicted by the observed changes in retailer decisions. We focus on monthly measures of: (i) Average number of weekly transactions (*Transactions*); (ii) Average number of weekly unique customers (*Uniques*); (iii) Number of new visitors as counted by Google Analytics (*New*); (iv) average basket size in USD (*Basket*); (v) average weekly amount of revenue from repeat customers (*Repeat*); (vi) the CAC; (vii) the unique number of products sold (*Products*); and (viii) the product concentration measured using the Herfindahl Index (HHI) (*ProductHHI*).¹⁶

Figures 8 and 9 report the results of an SCM analysis where the above variables were used as dependent variables. We only report those variables that exhibit statistically significant changes in outcomes after adoption. In the interest of space, we report only five lags to demonstrate pre-adoption matching. Figure 8 presents the results only for those retailers that used the analytics service. For users, we observe an increase in *Transactions*, *Uniques*, *New*, *Repeat*, *Products* (for four periods), and a decrease in *ProductHHI*. These changes are consistent with our predictions that personalization and prospecting technologies contribute to a changing mix of customers in the store, but they do not necessarily change how much customers spend. Additionally, the changes in *Products* and *ProductHHI* suggest that the assortment of products purchased has changed after adoption. We do not find evidence for changes in *Basket* or the CAC, which is consistent with the absence of changes in prices and advertising post analytics adoption reported in section 5.1.

¹⁶Product level data is limited to Shopify transactions. *ProductHHI* measures what fraction of sales each SKU generates, squares that figure, and sums those up to create a measure of concentration between 0 and 1 for each retailer-month.

The fact that we observe an increase in revenue without a change in basket size suggests that the increases are in the extensive margin of customer revenue, but not the intensive margin of profit per customer. Figure 9 reports these results for the retailers that adopted the service but did not use it.¹⁷ As can be seen in the figures, while the pre-adoption period matching worked relatively well for most outcome variables, retailers that did not use the service exhibit null effects for all of the variables.

Figure 8: Customer behavior outcomes for users




Figure 8 plots ATEs (and 95% confidence intervals) of adopting analytics on customer behavior-related outcome variables over time using the staggered SCM with random effects meta-analysis. The analyses are limited to retailers that used the service. Time 0 indicates month of adoption. Other times are relative to adoption.

The results are consistent with our prediction on how different firm actions contribute to the observed gain in revenue. As before, we would caution that these results should be interpreted carefully because it is difficult to disentangle the effect of two actions (e.g., pricing changes and recommender systems) that may contribute to the same outcome (e.g., changes in basket size).

5.3 Discussion

One surprising result from our mechanism analysis is that the gain in revenue after adopting descriptive analytics is not achieved through changes in pricing or advertising, which are often the easiest changes for the retailer to make. Given the nature of the descriptive dashboard, however,

¹⁷for *New*, *Products*, and *ProductHHI*, we use low frequency of usage because there were not enough retailers that did not use the service to perform SCM for these variables. For the *New* variable, we could only compute four lags and four leads due to data constraints.

Figure 9: Customer behavior outcomes for non-users




Figure 9 plots ATEs (and 95% confidence intervals) of adopting analytics on customer behavior-related outcome variables over time using the staggered SCM with random effects meta-analysis. The analyses are limited to retailers that did not use the service (or had low frequency of usage for the *New*, *Products*, and *ProductHHI* variables). Time 0 indicates month of adoption. Other times are relative to adoption.

it would indeed be hard to come up with a specific price or advertising allocation using simple KPIs. Potentially, once retailers realize that the descriptive dashboard doesn't provide specific recommendations, they might also install other technologies that automate price and advertising changes. However, we don't observe any changes that are consistent with such actions (not in pricing and advertising decisions, nor in customer behavior).

Overall, this section suggests that while all retailers adopt the focal analytics service and are also likely to adopt additional personalization and prospecting technologies, only those retailers that *use* the service also exhibit an increase in the number of new visitors, a reduction in the concentration of new products, and a corresponding increase in their performance outcomes, such as revenues and transactions. Therefore, we conclude that the descriptive dashboard and the personalization and prospecting technologies are complementary but that they require usage to realize the benefits of the analytics service. It appears that a benefit of using a descriptive dashboard is that it allows retailers to evaluate how different technologies impact their outcomes. Put together with the results on usage and the IV results, we conclude that the mechanism behind improvement in firm outcomes due to descriptive analytics adoption is the indirect mechanism.

6 Conclusion

Although 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 inability to explain what drives the benefits that are observed. In this paper we have aimed to identify the causal effect of the adoption of a descriptive analytics service by a wide variety of online retailers. Our unique dataset allowed us to not only provide estimates of the value created by descriptive analytics but also to provide evidence for the mechanism behind that value creation.

The results of our analysis show that the adoption of the analytics service in our sample of firms increases weekly revenue by an average of 13–20%. Although this range is wide and results from using multiple estimation approaches, we demonstrated that the positive effect is substantial and robust using multiple methods and alternative analyses. In addition, we provide evidence consistent with the interpretation that the dashboard benefits are accrued indirectly, likely by using the dashboard as a monitoring tool to assess the impact of other decisions.

There are a few potential concerns regarding the interpretation of our results. The first concern regards the generalizability of our results because we analyze the adoption of a single analytics service. First, because the analytics service provides a descriptive dashboard and does not incorporate algorithmic recommendations or predictions, we believe it is representative and not very different from other descriptive solutions. Second, this specific analytics dashboard was featured as a top selection by Shopify, one of the most widely used e-commerce hosting platforms, which indicates that it is not a small analytics provider. Third, while the focus is on adoption of one focal analytics service, our data and analysis include 1,164 e-commerce firms from varying industries and countries. These are not small firms, but ones with annual revenues of \$100K or more, and nearly all of them have been using Google Analytics prior to adoption of the service.¹⁸ Therefore, we believe that our findings are generalizable to adoption of other descriptive dashboards by a variety of firms. Moreover, the model presented in Figure 1 is applicable to any descriptive solution.

The second concern is about simultaneity and a causal interpretation of our findings. Potentially,

¹⁸For comparison, in our data retailers had an average basket size of \$182 with a median of \$87, and annual revenues with an average of \$652k and a median of \$246k. Shopify’s retailers in 2019 had an average basket size of \$67–\$101 depending on region, and an average annual revenue of \$74k (Shopify 2019)

firms made many changes (e.g., hired more skilled employees) and one of these changes was adding an analytics solution, with analytics having no direct impact on firm actions or performance (the parallel mechanism in the Figure 1). A unique result that we provide in our analysis is to show that the *usage* of analytics, and not their adoption per se, is what drives the improved firm performance. Further, we also use an IV strategy that shifts timing of adoption and shows that the results still hold, suggesting that the improvement in firm outcomes is due, at least partially, to the adoption of analytics. To further alleviate concerns about the causal relationship between retailer decisions and outputs, one could use causal mediation analysis (Imai et al. 2010a;b, Pearl 2014). However, the endogeneity of the mediators in our sample (firm actions) have prevented us from performing this analysis, which is left for future work. Additionally, there might be other mediating firm actions that we do not observe.

A third 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 likely 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 a LATE—among firms that choose to invest in analytics, what do we expect the impact to be?

Building on the main effect that we identify, we focus on disentangling different potential avenues 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 it is also due to using 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.¹⁹ These changes may be due

¹⁹Jin and Sun (2019) also do not find price changes when analytics are adopted. Their setting is of competitors on one specific online shopping platform.

to the firm changing the inventory of products it sells, but they may also be 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 and personalization technologies. These technological changes, coupled with changes in assortment and new visitors, are consistent with the finding that the customer’s basket size does not change, but the number of consumers as well as repeat revenues both increase.

One conclusion from our results is that retailers should not expect to generate actionable insights from descriptive dashboards easily. That is, descriptive analytics are 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 both users and non-users of the dashboard adopt additional technologies, but only the users of the dashboard experience benefits that these technologies are likely to provide. That is, only users saw increases in new visitors, in the number of unique products sold, in revenue from repeat customers, and overall in the number of transactions and revenue.

Why are descriptive analytics solutions so popular then? Although they rarely provide recommendations and leave the user to generate their own insights from the data, they allow retailers a simple way to monitor and assess the performance of different decisions, thus enabling marketers to extend the range of actions they can take and integrate new technologies. In turn, this suggests future avenues for researchers to create better predictive and prescriptive solutions to solve these challenges for marketers.

References

- A. Abadie. Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 2020.
- A. Abadie, A. Diamond, and J. Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program. *Journal of the American Statistical Association*, 105(490):493–505, 2010.
- 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.
- S. J. Anderson, P. Chintagunta, R. Kaul, and N. J. Vilcassim. Impact of marketing analytics 1.0 on entrepreneur and firm performance: Field experiment evidence from rwanda. Working paper, 2020.
- 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.
- E. Ben-Michael, A. Feller, and J. Rothstein. Synthetic controls and weighted event studies with staggered adoption. *arXiv preprint arXiv:1912.03290*, 2019.
- 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.
- M. Borenstein, L. V. Hedges, J. P. Higgins, and H. R. Rothstein. *Introduction to Meta-Analysis*. John Wiley & Sons, 2011.
- K. Boryusak and J. Xavier. Revisiting event study designs. Working paper, 2018.
- E. T. Bradlow, M. Gangwar, P. Kopalle, and S. Voleti. The role of big data and predictive analytics in retailing. *Journal of Retailing*, 93(1):79–95, 2017.
- 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.
- B. Callaway and P. H. C. Sant’Anna. Difference-in-differences with multiple time periods. *SSRN Electronic Journal* 3148250, 2020.
- 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.
- C. de Chaisemartin and X. d’Haultfoeuille. Two-way fixed effects estimators with heterogeneous treatment effects. Technical report, National Bureau of Economic Research, 2019.
- D. Delen and S. Ram. Research challenges and opportunities in business analytics. *Journal of Business Analytics*, 1(1):2–12, 2018.
- Y. Deng, A. Lambrecht, and Y. Liu. Spillover effects and freemium strategy in mobile app market. Working paper, 2019.
- N. Doudchenko and G. W. Imbens. Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. Technical report, National Bureau of Economic Research, 2016.
- F. Germann, G. L. Lilien, and A. Rangaswamy. Performance implications of deploying marketing analytics. *International Journal of Research in Marketing*, 30(2):114–128, 2013.
- F. Germann, G. L. Lilien, L. Fiedler, and M. Kraus. Do retailers benefit from deploying customer analytics? *Journal of Retailing*, 90(4):587–593, 2014.
- A. Goodman-Bacon. Difference-in-differences with variation in treatment timing. Technical report, National Bureau of Economic Research, 2018.
- C. Han. Efficiency comparison of random effects two stage least squares estimators. *Economic Letters*, 148:59–62, 2016.
- D. M. Hanssens and K. H. Pauwels. Demonstrating the value of marketing. *Journal of Marketing*, 80(6):173–190, 2016.
- 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.
- Y. Jin and Z. Sun. Information acquisition and the return to data. Working Paper, 2019.
- 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.
- J. Lismont, J. Vanthienen, B. Baesens, and W. Lemahieu. Defining analytics maturity indicators: A survey approach. *International Journal of Information Management*, 37(3):114–124, 2017.
- 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. Mintz, Y. Bart, P. Lenk, and D. Reibstein. Drowning in metrics: How managers select and trade-off metrics for making marketing budgetary decisions. *SSRN Electronic Journal 3502600*, 2019.
- 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*, 36(1):65–83, 2012.
- K. Pauwels, T. Ambler, B. H. Clark, P. LaPointe, D. Reibstein, B. Skiera, B. Wierenga, and T. Wiesel. Why, what, how, and what research is needed? *Journal of service research*, 12(2):175–189, 2009.
- 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.
- Shopify. Shopify’s state of commerce report, June 2019. <https://news.shopify.com/shopify-unveils-first-state-of-commerce-report>.
- 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.
- M. Wedel and P. Kannan. Marketing analytics for data-rich environments. *Journal of Marketing*, 80(6):97–121, 2016.
- J. Wooldridge. Fixed-effects and related estimators for correlated random-coefficient and treatment-effect panel data models. *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.
- H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: series B (statistical methodology)*, 67(2):301–320, 2005.

A Appendix

A.1 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:

$$Pr(Adopt_{ijt} = 1) = \Phi(\alpha + \beta X_{ijt-1} + \gamma Z_{ijt} + \delta \bar{X}_{ij} + \theta \bar{Z}_{ij} + g_t + \mu_{ij}) \quad (5)$$

where $Adopt_{ijt}$ indicates whether retailer i adopted the service in channel j in month t . Observations after adoption time t are not used to estimate equation 5 since the adoption decision is made once. X_{ijt-1} are retailer-channel-time control variables that include eight industry dummies, six channel dummies, the number of other channels of retailer i (excluding channel 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 hosting platform (referralVisits) 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-channel 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 μ_{ij} as retailer-channel 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 second and third stages of our estimation procedure, and is specified as:

$$\log(Y_{ijt} + 1) = \alpha + \beta \mathbf{AfterTreatment}_{ijt} + \delta X_{ijt-1} + \theta \bar{X}_{ij} + g_t + \mu_{ij} + \epsilon_{ijt} \quad (6)$$

where the variables are defined as in the previous equations for each retailer i with channel 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 (5) and (6). Note that this approach allows us to incorporate month-year fixed effects because there are adoptions and controls in each month and year (up to December 2017). These specifications increase the number of adoptions in our data (because all retailers adopt) compared to the limitation we had in our sample in section 4.1.1. For

each retailer-channel, the controls are every retailer-channel that have not yet adopted, and the effects are identified off individual retailer-channel variation.

Figure A-1: Traffic to the service provider’s website

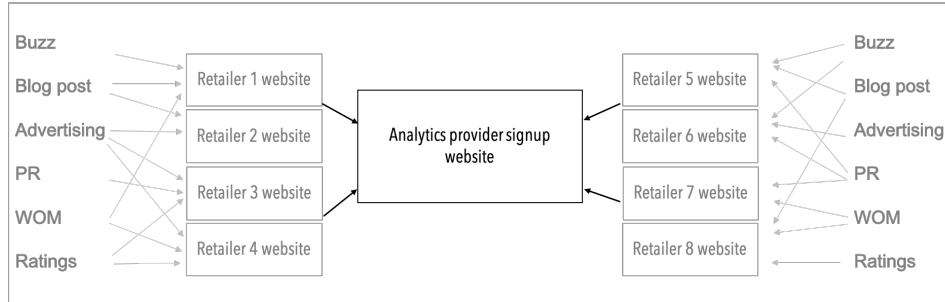


Figure A-1 illustrates the traffic we leverage for our IVs. We use the traffic to the focal service provider sign-up page (illustrated with the black thick arrows) as a proxy for various “buzz” channels about the analytics service provider (gray thin arrows).

Table A-1 presents the distribution of retailers’ regions. Figures A-2 and A-3 present the variation in the instruments over time.

A.2 IV Results

Panel A of Table A-2 presents the first-stage results and shows that our IVs are positively correlated with adoption as expected. Prior to estimating the two-step IV, we preform a Durbin-Wu-Hausman endogeneity test using the constructed IV, \widehat{Adopt}_{ijt} . The test produces p-values consistent with evidence of endogeneity (p-value=0.0003). We also confirm that our instruments are not weak using Stock-Yogo, Anderson-Rubin, and Stock-Wright tests.

Panel B of Table A-2 reports the second 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 is significantly different than zero. Additionally, the IV variable \widehat{Adopt}_{ijt} is positively associated with adoption. Column 1 of Table A-3 reports the third-stage IV results, and Column 2 reports comparable OLS estimates.

Table A-1: Distribution of retailers' regions

Region	Frequency	Percent	Region	Frequency	Percent
California	216	20.4%	New Jersey	14	1.3%
United Kingdom	99	9.3%	Washington	14	1.3%
Australia	85	8.0%	Oregon	13	1.2%
Canada	71	6.7%	Colorado	12	1.1%
New York	55	5.2%	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%
India	14	1.3%	Other	256	24.2%
Total			1,059	100.0%	

Table A-1 presents the distribution of location. In the US, we report states; otherwise countries are reported. In this table, locations with fewer than 10 retailers were consolidated into “other” for confidentiality reasons and in the interest of space. In our estimation we use the actual locations and not the “other” category.

Figure A-2: Hosting platform-based instrument over time




Figure A-2 illustrates the variation in traffic we leverage for our hosting platform-based IV. For each channel, we plot the number of new users that arrived from that channel, scaled such that the maximum number of users is 100 for each channel.

Figure A-3: Location-based instrument over time




Figure A-3 illustrates the variation in traffic we leverage for our location-based IV. For the top 20 locations in terms of number of companies in the data, we plot the number of new users that arrived from that location, scaled such that the maximum number of users is 100 for each location.

Table A-2: IV first and second steps

Panel A: First-Stage Probit (DV: Adoption dummy)	
Region visits	.0004** (.00013)
Referral visits	.000074** (.00001)
Log (Lag revenue)	.055** (.021)
# other channels	-.13* (.057)
Panel B: Second Stage (DV: Adoption dummy)	
\widehat{Adopt}_{ijt}	.129** (.0034)
F-Stat	366.89

Table A-3: IV and OLS estimates

Specification	IV (third stage)	OLS
After adoption	.18** (.057)	.076** (.019)
Log (Lag revenue)	.53** (.017)	.54** (.017)
# other channels	.0046 (.012)	.004 (.013)
Observations	12,115	12,115

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

SE in parentheses are clustered by retailer.

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

Web Appendix

A Additional Tables

Table W-1: Number of units used in SCM analysis

	Full sample		Repeated sample	
Cohort	Treatment	Control	Treatment	Control
Jan	55	356	50	315
Feb	122	363	109	323
Mar	56	350	53	311
Apr	66	358	60	314
May	68	333	54	288
Jun	60	313	54	268
Jul	50	336	41	276
Aug	47	345	42	279
Sep	50	331	41	265
Oct	48	296	40	231
Nov	62	268	56	206
Dec	49	221	38	166

B Main effect: Robustness

We perform a series of robustness tests to our main effect results. These are available in Table W-2. In this table, column “Baseline” presents the results corresponding to Figure 5 in the paper.

First, due to the importance of pre-period synthetic control matching, we show that our effects are robust to the choice of the number of lags and leads; we display robust results for additional 14-month windows with 9, 10, and 11 per-adoption lags (see columns “Lag-1,” “Lag-2,” “Lag-3,” in the appropriate tables).

Second, 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 hosting platform data is added to the dashboard , and estimate the models at the retailer level. Corresponding results appear in each table as the “Company-level” column.

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 the first six cohorts of 2017 adopters in our analysis as well. Column “2017 treated” presents the corresponding results.

As can be seen in the table, the SCM method matches the treatment and control units such that the difference between them is indistinguishable from zero for all of the different specifications.

Table W-2: Robustness for Revenues

Lag	Baseline	Lag-1	Lag-2	Lag-3	Company-level	2017 treated
-11				-0.002** (0.0003)		
-10			-0.0005 (0.001)	0.001** (0.0004)		
-9		-0.002* (0.001)	-0.002 (0.001)	-0.001* (0.0004)		
-8	-0.003* (0.001)	-0.001 (0.001)	-0.001 (0.001)	0.00005 (0.001)	-0.002 (0.001)	-0.003** (0.001)
-7	-0.003+ (0.002)	-0.004** (0.002)	-0.002+ (0.001)	-0.001 (0.001)	-0.004* (0.002)	-0.003 (0.002)
-6	-0.002 (0.002)	-0.0001 (0.002)	-0.001 (0.001)	0.0001 (0.001)	-0.001 (0.002)	-0.001 (0.002)
-5	0.003+ (0.002)	0.002 (0.002)	0.0005 (0.001)	0.0003 (0.001)	0.002 (0.002)	0.003+ (0.002)
-4	0.001 (0.002)	0.001 (0.002)	0.0002 (0.001)	-0.0001 (0.0004)	-0.0002 (0.002)	0.002 (0.002)
-3	0.001 (0.002)	0.002 (0.002)	0.002+ (0.001)	0.0005 (0.001)	0.001 (0.002)	0.0002 (0.002)
-2	0.003 (0.002)	0.001 (0.002)	0.002 (0.001)	0.001 (0.001)	0.004 (0.002)	0.002 (0.002)
-1	0.001 (0.002)	0.0002 (0.001)	-0.001 (0.001)	0.00005 (0.001)	0.002 (0.002)	0.002 (0.002)
0	0.088** (0.021)	0.075** (0.021)	0.078** (0.021)	0.082** (0.021)	0.097** (0.021)	0.093** (0.018)
1	0.102** (0.024)	0.093** (0.023)	0.101** (0.023)	0.093** (0.024)	0.109** (0.024)	0.124** (0.021)
2	0.108** (0.026)	0.111** (0.025)	0.100** (0.026)	0.079** (0.027)	0.114** (0.027)	0.138** (0.024)
3	0.143** (0.028)	0.138** (0.028)	0.117** (0.029)		0.152** (0.030)	0.162** (0.026)
4	0.172** (0.031)	0.146** (0.029)			0.180** (0.032)	0.184** (0.028)
5	0.141** (0.031)				0.146** (0.033)	0.146** (0.028)

SE in parentheses are computed using the meta-analysis method.

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

C Usage Results

Table W-3: Usage results

Lag	A: Usage dummy		B: Usage Intensity		C: Report Usage		D: Report Intensity	
	Users	Non-users	Below	Above	Users	Non-users	Below	Above
-8	-0.002 (0.002)	-0.013+ (0.007)	-0.001 (0.002)	-0.016** (0.005)	-0.004+ (0.002)	-0.004 (0.004)	-0.004 (0.003)	-0.006* (0.002)
-7	-0.004+ (0.002)	0.001 (0.011)	-0.010+ (0.005)	-0.007 (0.004)	-0.006* (0.003)	-0.002 (0.006)	-0.007 (0.005)	-0.004 (0.004)
-6	-0.001 (0.002)	-0.026* (0.013)	-0.004 (0.006)	-0.012* (0.005)	-0.001 (0.003)	-0.014+ (0.008)	-0.004 (0.005)	-0.002 (0.004)
-5	0.002 (0.002)	0.006 (0.017)	-0.0003 (0.006)	0.001 (0.005)	0.004 (0.003)	0.004 (0.009)	0.001 (0.006)	0.0002 (0.004)
-4	-0.002 (0.002)	0.025+ (0.013)	0.003 (0.005)	-0.003 (0.007)	-0.002 (0.003)	0.011 (0.007)	0.003 (0.005)	0.002 (0.004)
-3	0.002 (0.003)	-0.006 (0.012)	0.001 (0.006)	0.008 (0.006)	0.001 (0.003)	0.0005 (0.008)	0.004 (0.006)	0.005 (0.005)
-2	0.006* (0.003)	-0.009 (0.015)	0.004 (0.006)	0.022** (0.007)	0.007* (0.003)	-0.003 (0.007)	-0.002 (0.005)	0.010* (0.004)
-1	0.002 (0.002)	0.017 (0.013)	0.014** (0.005)	0.014* (0.007)	0.003 (0.003)	0.007 (0.007)	0.008+ (0.005)	0.0005 (0.004)
0	0.122** (0.023)	-0.036 (0.054)	0.112** (0.029)	0.120+ (0.061)	0.118** (0.024)	0.024 (0.042)	0.115** (0.029)	0.116** (0.033)
1	0.127** (0.026)	0.049 (0.064)	0.117** (0.035)	0.220** (0.063)	0.130** (0.028)	0.041 (0.050)	0.129** (0.036)	0.117** (0.035)
2	0.136** (0.029)	0.009 (0.079)	0.119** (0.038)	0.189** (0.061)	0.134** (0.030)	0.057 (0.060)	0.126** (0.038)	0.142** (0.041)
3	0.170** (0.033)	0.088 (0.086)	0.152** (0.042)	0.371** (0.073)	0.170** (0.034)	0.117+ (0.064)	0.153** (0.043)	0.210** (0.046)
4	0.199** (0.034)	0.082 (0.090)	0.192** (0.045)	0.392** (0.066)	0.194** (0.036)	0.142* (0.069)	0.190** (0.043)	0.223** (0.049)
5	0.165** (0.035)	0.029 (0.095)	0.120** (0.046)	0.262** (0.070)	0.168** (0.037)	0.052 (0.073)	0.119** (0.044)	0.231** (0.052)

SE in parentheses are computed using the meta-analysis method.

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

D Heterogeneity results

Table W-4: Heterogeneity by size

Lag	Revenue medians		Transaction medians	
	Below	Above	Below	Above
-8	-0.016** (0.005)	-0.003 (0.003)	-0.009* (0.004)	-0.006* (0.003)
-7	-0.013* (0.006)	-0.003 (0.005)	-0.025** (0.006)	0.001 (0.005)
-6	-0.003 (0.005)	0.006 (0.005)	-0.008 (0.006)	-0.002 (0.005)
-5	0.009+ (0.005)	-0.003 (0.006)	0.012+ (0.007)	0.003 (0.005)
-4	0.003 (0.005)	-0.001 (0.005)	-0.002 (0.006)	-0.002 (0.005)
-3	-0.0003 (0.005)	0.006 (0.006)	0.006 (0.008)	0.005 (0.005)
-2	0.014** (0.005)	-0.001 (0.006)	0.022** (0.006)	-0.004 (0.005)
-1	0.011* (0.005)	0.00002 (0.005)	0.014* (0.006)	0.005 (0.005)
0	0.156** (0.031)	0.033 (0.031)	0.206** (0.034)	-0.003 (0.025)
1	0.187** (0.037)	0.043 (0.031)	0.185** (0.039)	0.032 (0.027)
2	0.199** (0.040)	0.016 (0.034)	0.215** (0.042)	0.004 (0.031)
3	0.246** (0.043)	0.060 (0.038)	0.242** (0.046)	0.050 (0.034)
4	0.287** (0.046)	0.055 (0.041)	0.321** (0.048)	0.015 (0.036)
5	0.237** (0.049)	0.075* (0.038)	0.277** (0.048)	0.015 (0.038)

SE in parentheses are computed using the meta-analysis method.

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

E Mechanism: Additional Results

Table W-5: Advertising Budget on SpyFu

Lag	All retailers	Only users
-8	-0.020 (0.019)	-0.045 (0.034)
-7	0.010 (0.028)	-0.014 (0.041)
-6	0.055* (0.026)	0.092** (0.031)
-5	-0.020 (0.029)	-0.010 (0.037)
-4	0.020 (0.029)	0.044 (0.036)
-3	0.012 (0.026)	0.014 (0.031)
-2	-0.017 (0.027)	0.003 (0.040)
-1	-0.015 (0.016)	-0.100** (0.034)
0	-0.043 (0.108)	-0.153* (0.069)
1	-0.076 (0.132)	-0.160 (0.147)
2	-0.172 (0.113)	-0.337* (0.131)
3	-0.180 (0.139)	-0.316* (0.155)
4	-0.264+ (0.150)	-0.354* (0.157)
5	-0.406** (0.142)	-0.413** (0.149)

SE in parentheses are computed using the meta-analysis method.

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

Table W-6: Technology Adoption

Lag	Overall		Analytics&Tracking		Advertising		E-commerce	
	Users	Non-users	Users	Non-users	Users	Non-users	Users	Non-users
-8	-0.0001+ (0.0001)	0.001 (0.007)	-0.001** (0.0003)	-0.015 (0.015)	-0.0003 (0.0003)	-0.111** (0.037)	-0.0003 (0.0003)	-0.003 (0.003)
-7	0.0001 (0.0001)	-0.006 (0.012)	-0.0002 (0.0004)	0.009 (0.027)	0.0002 (0.0004)	0.073** (0.022)	-0.0003 (0.0005)	-0.026+ (0.015)
-6	-0.00002 (0.0001)	-0.017+ (0.010)	0.001+ (0.0005)	-0.019 (0.021)	-0.0003 (0.0005)	0.005 (0.027)	0.001 (0.001)	0.018* (0.008)
-5	-0.0003+ (0.0001)	-0.001 (0.010)	0.0001 (0.0004)	0.003 (0.016)	-0.001 (0.0003)	-0.069** (0.026)	-0.001+ (0.001)	0.009 (0.027)
-4	-0.0001 (0.0001)	-0.004 (0.009)	0.0001 (0.0003)	-0.003 (0.020)	-0.00000 (0.0002)	-0.018 (0.026)	0.002* (0.001)	-0.008 (0.011)
-3	-0.0002 (0.0002)	0.002 (0.008)	-0.0003 (0.0004)	0.002 (0.019)	0.0003+ (0.0002)	0.059+ (0.033)	0.0003 (0.0004)	-0.004 (0.010)
-2	0.0003 (0.0002)	-0.005 (0.010)	0.0001 (0.001)	0.009 (0.016)	0.00000 (0.0002)	0.052 (0.035)	-0.001 (0.0003)	-0.005 (0.006)
-1	0.0003* (0.0001)	0.005 (0.009)	0.001+ (0.0004)	0.026+ (0.014)	0.0005+ (0.0003)	0.079+ (0.046)	0.001** (0.0002)	0.007 (0.004)
0	0.011** (0.004)	0.032* (0.013)	0.018** (0.007)	0.061* (0.025)	0.043** (0.012)	0.034 (0.076)	0.018** (0.004)	0.008+ (0.004)
1	0.027** (0.005)	0.085** (0.017)	0.035** (0.011)	0.058** (0.020)	0.091** (0.018)	-0.006 (0.108)	0.022** (0.008)	0.016* (0.006)
2	0.038** (0.006)	0.129** (0.019)	0.056** (0.012)	0.076** (0.023)	0.095** (0.019)	0.011 (0.118)	0.031** (0.009)	0.006 (0.005)
3	0.051** (0.008)	0.137** (0.020)	0.065** (0.012)	0.113** (0.031)	0.121** (0.022)	-0.003 (0.104)	0.043** (0.009)	0.021+ (0.011)
4	0.069** (0.009)	0.145** (0.023)	0.068** (0.013)	0.183** (0.050)	0.120** (0.022)	0.017 (0.099)	0.050** (0.010)	0.040** (0.011)
5	0.072** (0.010)	0.155** (0.026)	0.073** (0.014)	0.205** (0.054)	0.132** (0.024)	0.043 (0.088)	0.048** (0.012)	0.042** (0.011)

SE in parentheses are computed using the meta-analysis method.

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