

The Value of Descriptive Analytics: Evidence from Online Retailers

Ron Berman,^a Ayelet Israeli^{b,*}

^aMarketing, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104; ^bMarketing Unit, Harvard Business School, Boston, Massachusetts 02163

*Corresponding author

Contact: ronber@wharton.upenn.edu,  <https://orcid.org/0000-0002-8594-3627> (RB); aisraeli@hbs.edu,  <https://orcid.org/0000-0002-5996-3202> (AI)

Received: August 20, 2019

Revised: November 20, 2020;

December 1, 2021

Accepted: January 4, 2022

Published Online in Articles in Advance:
March 16, 2022

<https://doi.org/10.1287/mksc.2022.1352>

Copyright: © 2022 INFORMS

Abstract. Does the adoption of descriptive analytics impact online retailer performance, and if so, how? We use the synthetic difference-in-differences method to analyze the staggered adoption of a retail analytics dashboard by more than 1,500 e-commerce websites, and we find an increase of 4%–10% in average weekly revenues postadoption. 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 customer relationship management, personalization, and prospecting 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 repeat customers. In contrast, there is no change in basket size. These findings are consistent with a complementary effect of descriptive analytics that serve as a monitoring device that helps retailers control additional martech tools and amplify their value. Without using the descriptive dashboard, retailers are unable to reap the benefits associated with these technologies.

History: Avi Goldfarb served as the senior editor and associate editor for this article.

Funding: This work was supported by The Mack Institute for Innovation Management.

Supplemental Material: The web appendix and data are available at <https://doi.org/10.1287/mksc.2022.1352>.

Keywords: descriptive analytics • big data • difference-in-differences • synthetic control • e-commerce • online retail • martech

1. Introduction

Marketers are often encouraged to invest in analytics-driven decisions, and indeed, survey research by Germann et al. (2013, 2014) reports a positive association between deploying marketing analytics technology and firm performance. Recently, chief marketing officers (CMOs) have been spending 22%–29% of their budgets on marketing technologies (martech) with martech spending set to increase to \$122 billion in 2022.¹ This trend led to research that documents the benefits of adopting popular, yet specific, prescriptive technologies, such as retargeting and A/B testing (e.g., Lambrecht and Tucker 2013, Koning et al. 2022) and contributes to the literature that shows a general positive benefit for firm performance from deploying analytics (Brynjolfsson et al. 2011a, Brynjolfsson and McElheran 2016, Müller et al. 2018).

Analytics technologies are often classified into 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). Whereas academic research often focuses on the potential benefit

from predictive or prescriptive technology that uses sophisticated modeling (Pauwels et al. 2009, Hanssens and Pauwels 2016, Wedel and Kannan 2016, Bradlow et al. 2017), most adopting firms typically deploy descriptive analytics and simple key performance indicator (KPIs) dashboards in the sales and marketing departments (Bughin 2017, Delen and Ram 2018, Mintz et al. 2019). Despite the popularity of descriptive analytics dashboards, little is known about how to interpret and turn descriptive metrics into actionable insights, which raises questions about the value of descriptive analytics.

In this paper, we set out to measure and document the value of adopting a marketing analytics dashboard. We utilize detailed panel data from more than 1,500 online retailers in a variety of industries to test whether there are benefits from adopting the dashboard, and we analyze these benefits. For each retailer, we observe multiple performance outcomes and decisions before and after they adopted the dashboard. We focus on trying to provide a causal estimate of the change in revenue that the retailers experience when they adopt the dashboard, 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 can shed light on the additional capabilities that descriptive analytics enables for marketers. 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 into the mechanisms that drive the performance gains from these heuristics.

How might descriptive analytics benefit retailers? As mentioned, there is no clear evidence that descriptive analytics helps at all as most prior research does not have appropriate data or could not distinguish between the types of analytics used by companies. If descriptive analytics does provide benefits for retailers, there are a few potential mechanisms that can lead to increased revenue. Our research aims to distinguish among them. First, descriptive dashboards can aid retailers in adjusting traditional marketing levers, such as pricing or advertising in response to the KPIs they observe. For example, the common metric of customer acquisition cost (CAC) can be compared with the profit margin of the retailer and can yield an adjustment in advertising to reduce the CAC. We call this mechanism the “direct mechanism.”

Second, the descriptive dashboard may be used as a monitoring tool to assess the value of other operations of the retailer but does not drive decisions directly. For example, if a retailer invests in retargeting campaigns or website personalization for customers, it is difficult to measure the effects of these investments and optimize parameters without monitoring potentially affected KPIs. In this case descriptive analytics is used as a complement to other martech to assess and amplify the strategies that work best. We refer to this mechanism as a “complementary mechanism.”

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, it also takes 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. We refer to this mechanism as an “unrelated mechanism.”

Prior attempts to investigate the relationship between data analytics and firm performance find 3%–7% higher productivity for firms that adopt data-driven decision making or big data assets. However, these attempts are often limited by access to nongranular or nonprimary data² and do not allow researchers to gain insight into the different actions that firms have taken to achieve these outcomes, the resulting firm output,

or the changes in customer behavior. A further challenge is that the results may be correlational and potentially suffer from endogeneity. For example, one confounder is that more productive firms might have more funds available to invest in analytics. Indeed, even research with detailed primary data—such as Bajari et al. (2019), who find that increased data availability improves the accuracy of firms’ demand prediction (with diminishing returns)—suffers from a lack of random assignment of firms into adopting and using analytics, which makes causal interpretation difficult. One exception, although in a very different context, is Anderson et al. (2020), who use a field experiment with small, physical, mom-and-pop stores in Rwanda and also find benefits to descriptive analytics. In their case, the benefit is due to training in data literacy and interacting with technology products, whereas our paper examines e-commerce firms with existing data collection capabilities.

To overcome the data and identification challenges in prior work, we use detailed panel data from a descriptive marketing analytics dashboard about online retailers that adopted the dashboard. Our data are unique in that they also contain 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 retailers’ actions (such as price changes and technology integration) as well as transaction-level data of individual customers. For the main analysis, we adapt the synthetic difference-in-differences (SynthDiD) method of Arkhangelsky et al. (2021) to scenarios with varying (staggered) treatment times. We also carefully confirm the robustness of our results using multiple empirical methods.

The three mechanisms we describe (direct, complementary, and unrelated) provide us with testable empirical predictions to understand how analytics operate to help retailers. Our data, combined with the SynthDiD method, allow us to provide evidence to investigate these mechanisms. We note that our data do not allow us to perform a full causal mediation analysis, and we, therefore, cannot interpret all results as fully causal. Further, because our data are collected from firms that chose to adopt the dashboard, selection into treatment may limit the interpretation and generalization of the effects we estimate. 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.

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 the analytics service is associated with an increase of 4%–10% in average weekly revenues of the firm. The results are robust to using multiple methods including staggered difference-in-difference (SDiD), SynthDiD, and instrumental variables (IVs). We also find that the smallest retailers (in terms of revenues or transactions) benefit most from adoption compared with larger retailers.

Second, we are able to disentangle usage of the dashboard postadoption from the adoption itself 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 in customer relationship management (CRM), personalization, and prospecting technologies. Among those who invested in these technologies, we observe gains in revenues only for firms that actively used the analytics dashboard. Taken together, these results suggest that the complementary mechanism of descriptive analytics may be in operation and the value of descriptive analytics is that it allows the firm to better control additional martech tools, 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 if correctly optimized.

The third major finding is that the adoption of analytics resulted in an increase in the number of transactions, number of new website visitors, number of unique customers, revenue from repeat customers, and the diversity of products purchased, but it did not result in increased basket size per transaction or reduced CAC. Again, all of these changes occur only 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 also give more credence to the complementary 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, which are more likely to drive these changes. Consequently, we observe that, whereas retailers 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

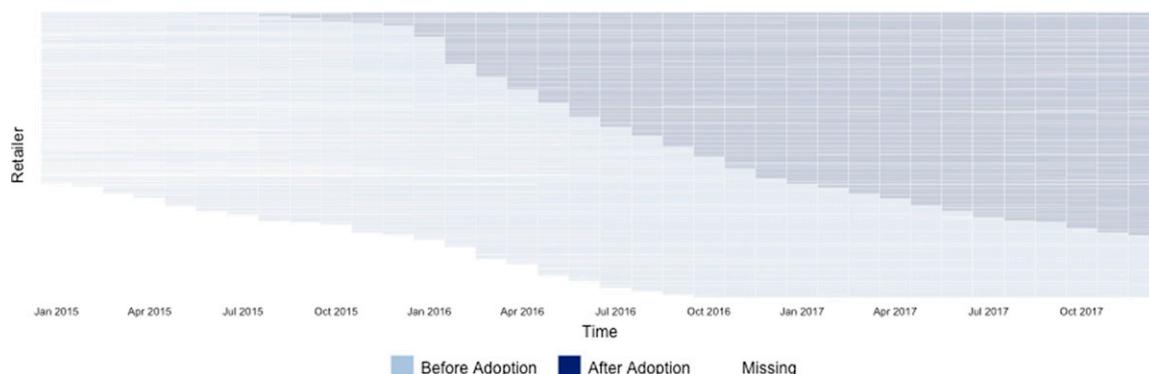
The analysis focuses on online retailers from a variety of industries that operate their own online stores.³ Many of the retailers in our data set manufacture and sell their own brands of products. For example, in the clothing and fashion categories, many of them produce and sell their individual designs and do not carry clothing from other brands.

Our data come 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 that sell through Amazon or use Paypal or Stripe as a checkout mechanism).⁴ 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 are collected daily at the transaction level, going backward 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.⁵

The dashboard uses the transaction-level data to generate metrics and visual reports on aggregate sales, average basket sizes, share of repeat customers' revenues, cost of new customer acquisition, 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 weekly and monthly levels. Compared with basic data reports from Google Analytics or e-commerce hosting providers, the benefit 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) an 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,

Figure 1. (Color online) Treatment Variation Plot



Notes. Each horizontal line corresponds to a retailer–channel combination. Light regions represent units before adoption of the analytics dashboard. Dark regions represent units after adoption of the analytics dashboard. White regions represent units with no data in the time period.

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 contain 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 between January 2016 and December 2018 and had historic annual revenue of at least US\$100,000. We use retailers that adopted the dashboard in the first 18 months as our focal treated group and retailers that adopted in the last 18 months as a control-only group. For the treated group, we use retailers that have at least 12 monthly observations, out of which there was at least one observation that occurs in 2015 (before signing up), one observation in the month of adoption, and one observation after adoption.

This yields a total of 1,001 retailers. Some of the retailers use multiple online channels to sell their products (e.g., their own website and Amazon), which yields 1,173 distinct retailer–channel combinations. The data are aggregated to a monthly level. Because retailers adopt the analytics service in a staggered manner, the data for each retailer in this time window can also serve as a control group during the period before their adoption. For the control-only group, in addition to the 12-month requirement, we include only retailers that had at least one observation in 2015 (or in 2016 for the 2018 adopters) to ensure sufficient overlap with the treated group. This yields 508 companies and 679 retailer–channel combinations out of a total of 1,631 companies that adopted the service during these 18 months. In our robustness tests, we leverage a larger portion of the retailers and find consistent results (see Section A.3.1 for details).

Figure 1 presents the treatment variation plot (Imai et al. 2022) that illustrates the dynamics of adoption of the analytics service in the data. Summary statistics for the data set are reported in Table 1. The industry categorization for the 1,509 retailers in our data are

Table 1. Summary Statistics

Variable	Mean	Standard deviation	Median	N
Panel A: Retailer level				
Number of observations per retailer	49.06	21.31	43.00	1,509
Average weekly revenue (US\$)	15,187.76	49,324.79	5,220.22	1,509
Average number of weekly transactions	188.00	517.85	58.66	1,509
Average number of unique customers	169.46	472.64	50.35	1,509
Average basket size (US\$)	179.61	375.60	87.41	1,509
Number of channels per retailer	1.26	0.52	1.00	1,509
Panel B: Observation level				
Average weekly revenue (US\$)	14,193.84	48,088.17	4,520.89	51,380
Average number of weekly transactions	181.74	695.95	49.00	51,380
Average number of unique customers	164.46	608.24	42.20	51,380
Average basket size (US\$)	173.42	407.64	81.65	51,380

Table 2. Distribution of Industries

Industry	Frequency	Percentage
Clothing and fashion	366	24.3
Health and beauty	185	12.3
Food and drink	110	7.3
Home and garden	107	7.1
Sports and recreation	93	6.2
Electronics	89	5.9
Jewelry and accessories	87	5.8
Toys and games	46	3.0
Other	426	28.2
Total	1,509	100

described in Table 2. Table 3 displays the distribution of the different hosting platforms for each of the 1,859 retailer–channel combinations in our data. Most of the retailers in the sample operate their own website, whereas roughly 7% of the retailers sell on Amazon either as a vendor or a seller. Shopify, BigCommerce, and Magento represent 61% of the retail channels in our sample. These platforms together held 36% market share in 2016 among the top eight e-commerce hosting platforms for independent sellers.⁶ Among the top 100,000 web stores, they held 37% in 2017.⁷

We augment the retailer data with data from four additional sources: (i) data about nonsubscribed 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 that the retailer installed at the online store, such as advertising 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 United States and the United Kingdom. Table 4 includes summary statistics of the main variables collected for each of the retailers.

4. Empirical Strategy and Results

This section details the empirical strategy we use to identify treatment effects and then provides estimates of the revenue gains retailers experience when they adopt the analytics dashboard. Section 5 then analyzes the mechanism behind these gains.

Table 4. Summary Statistics: Additional Data

Variable	Mean	Standard deviation	Median	N
Average monthly number of logins	0.52	0.92	0.20	42,925
Monthly indicator for any report views	0.08	0.27	0	42,925
Monthly number of technologies	55.46	23.42	52	40,754
Google average monthly ad costs (US\$)	2,351.36	12,275.70	0	18,464
Spyfu monthly advertising spending (US\$)	1,838.67	13,193.40	423.90	12,357
Spyfu monthly number of advertising keywords	27.43	44.50	8	7,024

Notes. The observations in this table are at the retailer-month level because 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.

Table 3. Distribution of Hosting Platforms

Hosting platform	Frequency	Percentage
Shopify	1,010	54.33
Paypal	484	26.04
Amazon Seller	129	6.92
Stripe	114	6.13
BigCommerce	94	5.06
Magento	24	1.29
Amazon Vendor	4	0.22
Total	1,859	100

4.1. Identification Challenges

The observational nature of the data poses several challenges for identification of the treatment effects: (i) every firm in our data is either treated or untreated in each time period, and we do not observe its counterfactual outcome in the unobserved condition; (ii) firms in our data adopt the dashboard in a staggered pattern with different adoption times, requiring careful computation and aggregation of treatment effects (Goodman-Bacon 2018, de Chaisemartin and d'Haultfoeuille 2020, Borusyak et al. 2021, Callaway and Sant'Anna 2021, Wooldridge 2021); (iii) the decision of when to adopt is endogenous and may be correlated with firm expectations of the benefit from the dashboard, resulting in a potential upward bias of our estimates; (iv) we only observe data from firms that eventually adopted the dashboard, and these firms might have been those expecting the highest benefit from the dashboard, leading to a potential upward bias of the estimates because of selection; and (v) a retailer could have implemented other unobserved policies concurrently with adopting the dashboard, and the observed effect cannot be separately attributed to the analytics dashboard.

To address these challenges, we utilize two empirical methods: SDiD based on Wooldridge (2021) and cohort-based SynthDiD based on Arkhangelsky et al. (2021), which borrows strengths from the SDiD method as well as the synthetic control method (Abadie et al. 2010, Abadie 2021). The technical details of these methods are provided in Section 4.2, but first, we provide an overview of how these methods address the identification challenges.

Challenge (i) of unobserved potential outcomes is addressed by constructing a control group to predict the counterfactual potential outcomes of an adopting firm as if it did not adopt analytics. To do so, we treat the cohorts of retailers that adopted the analytics service in our data window as our focal treatment group. The control group for each adopting cohort comprises data from firms that haven't adopted the dashboard yet or have adopted the dashboard for the first time outside the data window (Manchanda et al. 2015). Specifically, we designate retailers that adopted the analytics dashboard between January 2016 and June 2017 as the treated cohorts, whereas retailers that adopted between July 2017 and December 2018 serve as control-only cohorts.

This approach requires a conditional parallel-trends assumption on pretreatment outcomes to hold in the case of SDiD, which we address by controlling for linear cohort-level time trends. The SynthDiD method relaxes the parallel-trends assumption and allows for a more flexible preadoption pattern. We test the validity of these methods both visually and statistically.

Challenge (ii) of staggered adoption is addressed by using both the SDiD and SynthDiD methods to compute treatment effects for each cohort of adopting firms separately and then aggregating them into an overall average treatment effect on the treated (ATT), resolving issues such as negative weighting of treatment effects identified by, for example, Goodman-Bacon (2018), de Chaisemartin and d'Haultfoeuille (2020), and Callaway and Sant'Anna (2021).

We address challenge (iii) of endogenous adoption timing in two ways. First, SynthDiD is consistent even if there is unobserved correlation between treatment assignment and firm-level time trends, alleviating much of the concern. Second, we perform an instrumental variables analysis using three instruments for the timing of adoption that likely shift adoption timing but are plausibly unrelated to the retailer's revenue.

Regarding challenge (iv), because all of the retailers in our data adopt the analytics service eventually, they may differ from unobserved nonadopters, causing a selection bias when interpreting our estimates as an average treatment effect (ATE). Our estimates are, therefore, the effect on retailers that choose to adopt the service, that is, the ATT.⁸ Despite this potential limitation, we believe that the ATT 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.⁹ Accordingly, we measure the ATT and expect that similar retailers that are interested in adoption of descriptive analytics exhibit

similar outcomes and experience similar benefits as those in our sample. A second concern of selection is that firms with a higher benefit might choose to adopt with higher probability, and indeed, our estimates are representative for the set of firms in our data. In Section 4.4, we provide more details about treatment effect heterogeneity in our sample. A final concern is that firms with a higher expected benefit might select earlier into treatment, but this is handled by the solutions to challenge (iii).

Finally, to address challenge (v) of unobserved confounders, we use a unique aspect of our data set by which we observe firms' usage of the dashboard post-adoption. We compare the treatment effect for firms that made use of the dashboard (users) versus those that didn't (nonusers) to provide evidence that the benefit we observe can be attributed to the dashboard and not to other confounders.

4.2. Effect of Analytics Adoption on Retailer Revenue

4.2.1. SDiD. We start by analyzing the impact of analytics adoption on average weekly revenue using difference-in-differences. When adoption is not staggered, a common approach to estimating the ATT is to utilize a two-way fixed effects (TWFE) estimator by using ordinary least squares (OLS):

$$\mathbb{E}[\log(Y_{ijt} + 1)] = \alpha_{ij} + \gamma_t + \beta \text{AfterAdopt}_{ijt}, \quad (1)$$

where Y_{ijt} is average weekly revenue for retailer i in channel j in month t , α_{ij} are retailer-channel fixed effects, γ_t are time fixed effects, and AfterAdopt_{ijt} indicates whether retailer-channel ij adopted the dashboard at or before time t . The coefficient β in this specification estimates the ATT.

Wooldridge (2021) shows that, in the case of staggered adoption, the TWFE estimator is equivalent to a pooled OLS regression in which unit fixed effects are replaced with cohort fixed effects and the treatment effects are allowed to vary by cohort and time. We additionally allow for flexible linear time trends by cohort, which yields the following specification:

$$\begin{aligned} \mathbb{E}[\log(Y_{ijt} + 1)] &= \sum_{r=q}^T \lambda_r \text{Adopt}_{ijr} + \gamma_t \\ &+ \sum_{r=q}^T \tau_{rt} \text{Adopt}_{ijr} \mathbb{I}[t \geq r] \\ &+ \sum_{r=q}^T \phi_r \text{Adopt}_{ijr} \times t, \end{aligned} \quad (2)$$

where Adopt_{ijr} indicates whether retailer-channel ij first adopted the dashboard at time period (cohort) r , $\mathbb{I}(\cdot)$ is the indicator function, λ_r is a fixed effect for adoption cohort r , and τ_{rt} measures the treatment effect at time t

for retailer channels that adopted in time period r . The value q denotes the time period of the first adoption cohort in the data, whereas T denotes the last time period. The coefficients γ_t are time fixed effects as before, and ϕ_r captures any linear time trends in outcomes of retailer channels that adopted in cohort r .

The main difference between Specifications (1) and (2) is that the unit (retailer channel) fixed effects α_{ij} are replaced by cohort fixed effects λ_r and treatment effects are allowed to be heterogeneous across both cohorts and time. For example, τ_{23} measures the effect at time $t = 3$ for units that adopted in cohort $r = 2$, and it may be different from τ_{33} (effect at time $t = 3$ for cohort 3 units) and from τ_{22} (effect at time $t = 2$ for cohort 2 units). Specification (2) resolves the bias created by Specification (1) when staggered adoption with heterogeneous treatment effects is in play. If the true ATT was homogeneous across units and time or if adoption occurred in a single period, then both specifications would yield the same estimate.

We use two-way clustering of standard errors by retailer and month to address serial correlation (Bertrand et al. 2004). Most retailers sell using only one channel—typically their own website. However, we use retailer-channel observations (i.e., a retailer that sells on its own website and also through Amazon has two observations at each time period) because retailers may adopt the dashboard at different times for different online channels. Results are consistent when we aggregate the data to the retailer level.

The coefficient τ_{rt} measures the change in average weekly revenue after the adoption of analytics of cohort r at time t . It is identified by comparing the change in revenue of cohort r adopters from periods before r to period t with the change in revenues in the same time frame for retailers that adopt the dashboard after time t . Unlike Datta et al. (2018), we do not observe a group of nonadopters, and hence, we rely on future adopters as a control group. The identifying assumption in our SDID analysis is a conditional parallel-trends assumption: that there were no differential trends in revenues before adoption between retailers that adopted the service and those that did not after conditioning on unit- or time-invariant covariates. 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.

Given the cohort-level estimates $\widehat{\tau}_{rt}$, we can compute a dynamic ATT estimate, $\widehat{\tau}_\ell = \frac{\sum_{r=q}^{T-\ell} \widehat{\tau}_{r(r+\ell)}}{T-\ell-r+1}$, which depends on the length ℓ of exposure to treatment ($0 \leq \ell \leq T - q$) and measures the average effect ℓ periods after treatment for all treated units. We can also

aggregate all the treatment effects into an overall ATT estimate $\widehat{\tau}$:

$$\widehat{\tau} = \frac{\sum_{r=q}^T \sum_{t=r}^T \widehat{\tau}_{rt}}{(T - q + 1)(T - q - 2)/2}.$$

To ensure that there are sufficient control units in each time period, we truncate our data after June 2017, allowing all firms that adopted after this month to serve as controls only (whereas, in each time period, firms in cohorts that have not yet adopted serve as controls before their adoption, similar in spirit to Wang and Goldfarb (2017)). Because of this data truncation, we limit our SDID analysis to include only observations between July 2015 and June 2017 for retailers that adopted the dashboard in 2016–2018. This results in 36,433 observations of 18 retailer-channel cohorts that adopted the dashboard between January 2016 and June 2017, totaling 1,173 treated units and 679 control-only units that adopted after June 2017.

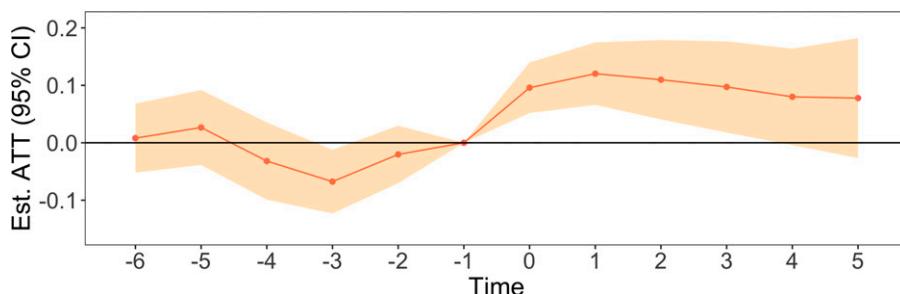
The validity of the estimates depends on the validity of the identifying assumption of conditional parallel trends, which we can test because we observe retailer channels over multiple time periods before adoption. Following Borusyak et al. (2021) we run an OLS analysis that uses only the pretreatment¹⁰ data and estimate the following specification:

$$\begin{aligned} \mathbb{E}[\log(Y_{ijt} + 1)] = & \sum_{r=q}^T \lambda_r Adopt_{ijr} + \gamma_t + \sum_{\ell=1-T}^{-2} \delta_\ell Adopt_{ij(t-\ell)} \\ & + \sum_{r=q}^T \phi_r Adopt_{ijr} \times t, \end{aligned} \quad (3)$$

where, as before, λ_r are adoption cohort fixed effects, γ_t are time fixed effects, and ϕ_r are adoption-cohort linear time trends. The indicator $Adopt_{ij(t-\ell)}$ denotes units that have adopted $|\ell|$ time periods after period t , and the coefficients δ_ℓ are lagged treatment effects that measure the difference in outcomes between treatment and control units ℓ periods before adoption. The effect at lag $\ell = -1$ is normalized to zero, and lag $\ell = -7$ is omitted to serve as the baseline level of outcome because the data truncation in June 2015 means that the earliest treated cohort of January 2016 has at most six periods of pretreatment data. We then use a Wald test to test the null hypothesis $\delta_\ell = 0$ for $-6 \leq \ell \leq -2$ that states that the differences in outcomes between treatment and control units in the six periods before adoption are zero.

Figure 2 presents the dynamic event study treatment effect estimates $\widehat{\tau}_\ell$ from Specifications (2) and (3) along with their 95% confidence intervals (CIs). We observe that most of the pretreatment estimates contain zero in their confidence intervals (with a violation

Figure 2. (Color online) SDiD Treatment Effects



Note. Coefficients and 95% confidence intervals obtained by estimating Equations (2) and (3).

in period -3), whereas the treatment effect estimates are positive and inch toward zero in later periods. The Wald test of the null hypothesis of parallel trends is unable to reject the null with a p -value of 0.1079, providing some evidence for the validity of the parallel-trends assumption. Column (1) of Table 5 presents the estimate of the ATT for the first quarter postadoption, which is valued at 0.108 (95% CI: 0.056, 0.161), and column (2) presents the six-month ATT, which is valued at 0.098 (95% CI: 0.034, 0.162). Because revenues are logged, the six-month ATT estimate suggests an increase in average weekly revenues of roughly 10.3% each month. The unconditional median average weekly revenue before adoption in our data are US\$4,364. Therefore, the median retailer experienced an increase of approximately US\$449.

Equations (2) and (3) allow for flexible cohort-level linear time trends, measured by the coefficients ϕ_r . Such a specification is uncommon in standard analysis that uses a specification such as (1) but is necessary in our setting to obtain parallel trends. When we reestimate Specifications (2) and (3) after omitting the linear time trend coefficients ϕ_r , the analysis yields nonzero estimates of $\hat{\tau}_\ell$ for most of the pretreatment periods $\ell = -2, \dots, -6$. Additionally, the Wald test of the null hypothesis $\delta_\ell = 0$ rejects the null with a p -value of 0.0011. These results raise concerns about the validity of the estimates from SDiD without controlling for time trends at the cohort level.

In the next section, we turn to identify the treatment effects using synthetic difference-in-differences, which does not rely on a parallel-trends assumption and also

resolves a few additional identification challenges mentioned in Section 4.1.

4.2.2. Aggregated SynthDiD. 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 preadoption outcomes of the adopting retailer, and thus, they capture any possible trends that might affect identification without requiring a parallel-trends assumption. The difference between the observed outcomes postadoption and the predicted outcomes are the estimated treatment effects from the method (Abadie et al. 2010, Abadie 2021).

Synthetic control analysis is often applied to a single or small number of treated units with a small number of control units and requires long and balanced preadoption panels. In our setting, the choice of preadoption panel length creates a trade-off between bias and variance of the estimates. If we require long preadoption panels, the risk of biased estimates is lowered, but the number of eligible retailers to include in the analysis is smaller, which increases the variance of the estimates. If we shorten the number of preadoption time periods, more retailers are eligible to be included in the analysis, but the risk of biased estimates increases. We address this challenge by employing synthetic difference-in-differences (Arkhangelsky et al. 2021), which borrows strengths from both the difference-in-differences (DiD) and SCM methods. SynthDiD finds optimal weights for control units and pretreatment periods to minimize the

Table 5. Effect of Analytics Adoption on Retailer Revenue

Methodology	Staggered DiD		SynthDiD		IV (third stage) (one month) (5)
	(three months) (1)	(six months) (2)	(three months) (3)	(six months) (4)	
After Adoption	0.108** (0.026)	0.098** (0.031)	0.056** (0.019)	0.038 ⁺ (0.020)	0.167** (0.05)

Notes. Point estimates for the treatment effects using SDiD in columns (1) and (2) (average effects for three and six months postadoption); SynthDiD in columns (3) and (4) (average effects for three and six months postadoption); and IV in column (5) (effect for the month of adoption only).

Significance level: 10% (+), 5% (*), 1% (**).

mean squared error of the target ATT to be estimated. Similarly to SCM, SynthDiD uses pretreatment data of control units to create a synthetic control for the average outcome of treated units and does not rely on a strong parallel-trends assumption for identification. Similarly to DiD, SynthDiD is invariant to unit-level shifts in outcomes and also allows for inference with large panels even when the pretreatment period is short.

The SynthDiD method is designed for a balanced panel of units in which the treatment timing is identical for all treated units. We adapt the method to estimate the ATT with staggered adoption by estimating a cohort-level ATT and then aggregating the estimates in a method similar to that used in Section 4.2.1 and also in applications of SCM with staggered adoption (Ben-Michael et al. 2022).

To perform the analysis, for each adoption cohort r , we construct a balanced panel in which the treatment group comprises retailer channels that adopted the dashboard in period r and have outcome data available $\ell_{\min} = -6$ periods before adoption and $\ell_{\max} + 1 = 6$ periods after and the control group comprises retailer channels that have data available for the same time frame but that adopted the dashboard after more than ℓ_{\max} periods after cohort r .

If we denote by N_r the set of units in the balanced panel of cohort r , by N_r^{co} the set of units in the control group, and by N_r^{tr} the set of units in the treatment group,¹¹ then for each cohort r , the SynthDiD estimation procedure solves

$$(\widehat{\tau}_r, \widehat{\alpha}_0, \widehat{\alpha}_{ij}, \widehat{\gamma}_t) = \arg \min_{\tau_r, \alpha_0, \alpha_{ij}, \gamma_t} \left\{ \sum_{ij \in N_r} \sum_{t=r+\ell_{\min}}^{r+\ell_{\max}} (\log(Y_{ijt} + 1) - \alpha_0 - \alpha_{ij} - \gamma_t - \text{AfterAdopt}_{ijt} \cdot \tau_r)^2 \widehat{\omega}_{ij} \widehat{\lambda}_t \right\}, \quad (4)$$

where τ_r is the average ATT of cohort r in the $\ell_{\max} + 1$ periods after adoption, α_0 is an intercept, α_{ij} and γ_t are retailer–channel and time fixed effects as before, and AfterAdopt_{ijt} indicates whether retailer–channel ij adopted the dashboard by time period t .

Equation (4) estimates a two-way-fixed-effect model identical to Equation (1) with the addition of unit- $\widehat{\omega}_{ij}$ and time-specific weights $\widehat{\lambda}_t$. The unit weights $\widehat{\omega}_{ij}$ are selected such that the pretreatment control outcomes weighted by $\widehat{\omega}_{ij}$ have a similar trend to that of the average outcomes of the treatment units; that is, for time periods $t < r$,¹²

$$\widehat{\omega}_0 + \sum_{i \in N_r^{co}} \widehat{\omega}_{ij} \log(Y_{ijt} + 1) \approx \frac{\sum_{i \in N_r^{tr}} \log(Y_{ijt} + 1)}{|N_r^{tr}|}.$$

The time weights $\widehat{\lambda}_t$ are designed so that the average posttreatment outcome for each of the control

units differs by a constant from the weighted average of the pretreatment outcomes of the same control units, that is,

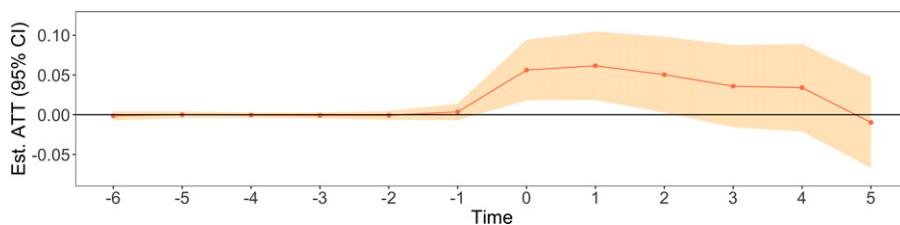
$$\widehat{\lambda}_0 + \sum_{t=r+\ell_{\min}}^{r-1} \widehat{\lambda}_t \log(Y_{ijt} + 1) \approx \frac{\sum_{t=r}^{r+\ell_{\max}} \log(Y_{ijt} + 1)}{\ell_{\max} + 1}.$$

The unit weights $\widehat{\omega}_{ij}$ serve the same role as in the standard SCM to align the preexposure trends in the outcomes of treated and control units. The time weights $\widehat{\lambda}_t$ balance preexposure time periods with postexposure ones; if a specific preexposure period is more predictive of postexposure outcomes, it receives a higher weight. Arkhangelsky et al. (2021) show that, under quite general conditions, including correlation between treatment assignment and unit-level time trends as well as heterogeneous treatment effects, $\widehat{\tau}_r$ is a consistent and asymptotically normal estimator of τ_r as long as the combination of the number of control units and preadoption periods is large enough compared with the combination of the number of treated units and postadoption periods.¹³ We believe that these assumptions hold in our setting because lowering the number of preadoption periods indeed increases the number of control units faster than the number of treated units.

Similarly to the analysis in Section 4.2.1, we truncate the data between July 2015 and December 2017, and obtain 18 cohort-level ATT estimates $\widehat{\tau}_r$ for cohorts that adopted between January 2016 and June 2017 using the synthdid package in the R programming language.¹⁴ We then compute the weighted average ATT as $\widehat{\tau} = \frac{\sum_r N_r^{tr} \cdot \widehat{\tau}_r}{\sum_r N_r^{tr}}$. Standard errors for each $\widehat{\tau}_r$ are estimated using the jackknife (algorithm 3 of Arkhangelsky et al. 2021) or placebo method (algorithm 4 of Arkhangelsky et al. 2021) if a cohort has only one treated unit. Standard errors for the overall ATT are computed as a weighted average of the cohort-level standard errors. Across the 18 cohorts, there are 982 treated retailers that have sufficient observations for the analysis. Table A.1 details the breakdown of the number of treated and control units used in each cohort. On average, there are 55 treated units and 485 control units in each cohort.¹⁵

Columns (3) and (4) of Table 5 present the estimated ATT from SynthDiD in the three and six months after adoption of the dashboard with the first quarter's ATT valued at 0.056 (95% CI: 0.019, 0.090), and the six-month ATT valued at 0.038 (95% CI: -0.001, 0.078). We also present an event study plot in Figure 3 that shows the pretreatment fit of SynthDiD between the treatment and control units as well as the evolution of the treatment effect over time.¹⁶ We observe that SynthDiD produces better pretreatment fit between the trends of the treated and control units

Figure 3. (Color online) SynthDiD Treatment Effects



Notes. SynthDiD ATTs of adopting analytics on revenues over time. Time 0 indicates month of adoption. Other times are relative to adoption.

compared with SDiD. The treatment effect estimates after adoption are smaller in magnitude than those achieved with SDiD. The effects are positive in the first few months after adoption and then attenuate in later periods.¹⁷ The six-month treatment effect of 0.038 is equivalent to an average increase of 3.9% in revenues or US\$170 for the median retailer in the data.

Table A.2 presents multiple analyses that demonstrate the robustness of our SynthDiD estimates to different numbers of preadoption lags (ℓ_{\min}) and postadoption leads (ℓ_{\max}) used to create the balanced panel, aggregation of the data to the retailer level instead of retailer–channel level, and using an additional sample in which 2017 adopters serve as treated and 2018 adopters serve as controls.

Because SynthDiD ensures an appropriate number of control units, accounts for potential preadoption trends, and allows for some level of endogenous treatment selection and timing, we use SynthDiD as the main analysis method in the remainder of the paper.

4.3. Addressing Potential Selection Bias

In this section, we assess the sensitivity of our findings to selection. We ask how the estimate that weekly revenues increase by an average of 3.9% in the six months after adoption of the dashboard might be affected by endogenous timing of the adoption or by selection into treatment. First, we use an instrumental variables analysis to alleviate concerns about endogenous adoption timing. Second, we present evidence that adopting the dashboard without using it 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 unrelated mechanism described in Section 1.

4.3.1. Instrumental Variables Analysis of Endogenous Adoption Timing.

As described in challenge (iii) in Section 4.1, firms may choose to adopt the dashboard when they anticipate the greatest benefit. We address this issue using IV analysis. Our IV strategy hinges on identifying times of increased awareness and attention to the analytics service among online retailers and on assuming that such increased awareness drives some of

the adoption exogenously. Specifically, we use three time-varying exogenous factors that plausibly impact the timing of analytics adoption but that are uncorrelated with the revenue of the retailer. The first IV utilizes the adoption rates among other retailers within each retailer's industry, and the other two IVs use data from the website of the *analytics service provider* (not the retailers) to measure interest in signing up to the dashboard from other websites in the retailer's geographic region and hosting channel. We view these three IVs as exogenous shifters of adoption timing that are not correlated with the performance of an adopting retailer. Web Appendix WA-1 provides details of the data used to construct the three instruments and discusses their relevance for shifting adoption and the validity of the exclusion restriction assumption.

Because retailer channels adopt the dashboard only once, the instruments cannot continue to be relevant adoption shifters postadoption. We, therefore, limit the scope of our analysis to include only observations up to the first month of adoption (including that month). Because we use an IV as the identification strategy, we consequently treat adoption timing as random and no longer rely on preadoption matching of trends between treatment and control. This allows us to use all adopting cohorts in the IV analysis.

Our IV strategy has two limitations. First, the instruments would ideally have exogenous variation at the retailer–channel level in each time period. However, our instruments only have variation at an industry–region–platform level, implying that two units from the same region, in the same industry, and using the same platform have identical instruments. Luckily, these three instruments together create enough subgroups so that the variation is almost at the retailer–channel level (generating 6,686 unique values for 12,115 observations in this analysis). Second, because the estimates limit the data to incorporating only one period after adoption, the IV estimates measure the ATT only during the first month of adoption. For these two reasons, we consider the IV estimates only as supporting evidence for the positive effect of adoption on performance outcomes that we found using SDiD and SynthDiD.

Because the endogenous adoption variable is binary, we utilize the three-step estimation procedure proposed in Deng et al. (2019) based on Wooldridge (2005, 2010, 2019). The first step is a probit model predicting the decision to adopt the dashboard. The resulting predicted probability is then used as an instrument in a two-step IV procedure. The estimation procedure and results are detailed in Web Appendix WA-2. The two-stage least squares instrument derived from the first stage passes the Stock–Yogo, Cragg–Donald, Anderson–Rubin, and Stock–Wright weak instrument tests.

Column (5) of Table 5 reports the IV analysis results that confirm our previous findings and show a revenue increase of 18.2% (95% CI: 7.1%, 30.5%). Because the IV analysis only estimates effects for the month of adoption, the estimates are substantially larger than those obtained in the SDiD analysis presented in Figure 2 and in the SynthDiD analysis presented in Section 4.2.2.¹⁸ Note that the IV analysis uses the same data used in the SynthDiD analysis. When limiting the data to the data used in the SDiD analysis, the IV estimate is 0.099 with a *p*-value of 0.057. That suggests a revenue increase of 10.4% (95% CI: –0.003%, 22.3%).

4.3.2. Disentangling Adoption and Usage. A second issue we examine is potential simultaneity bias in the estimate (challenge (v) in Section 4.1). Firms that adopted the dashboard could have made other concurrent changes, such as changing their management team or their website design, and these actions could have generated the change in performance outcomes we observe. If this concern is valid, we would expect firms that adopted the dashboard, but *did not* use it, to exhibit increased performance as well.

Our data are unique in that they include the dashboard login times and the reports accessed by each retailer, which allow us to separately estimate the effects of the dashboard's adoption by level of usage. We examine three dimensions of usage: any usage, intensity of dashboard logins, and report viewership. Because the access and login information are not available at the retailer–channel level, but rather at the retailer level, we aggregate observations to the retailer level and use the first time of adoption as the retailer's adoption time. We use the SynthDiD method to analyze the differential treatment effects based on each of the different aspects of usage.

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 nonusers. We then compare postadoption treatment effects across the groups of users and nonusers to examine whether retailers that use the service exhibit different outcomes than those that do not.¹⁹ Figure 4 reports the treatment effects of all time

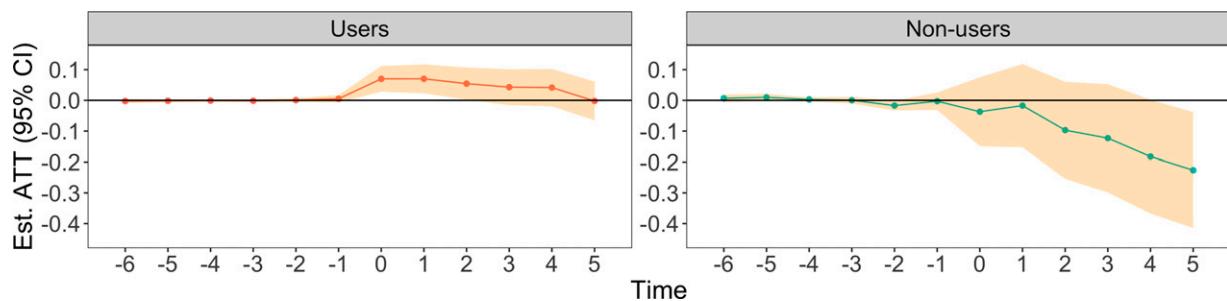
periods from six months before to six months after adoption. Before adoption, for both users and nonusers, most point estimates are statistically indistinguishable from zero (those that are different than zero are very close economically to zero). After adoption, only the point estimates of users are larger than zero. Additionally, the ATT for the six months postadoption are 0.046 (95% CI: 0.002, 0.090), and –0.11 (95% CI: –0.247, 0.019) for users and nonusers, respectively.²⁰ When we compute the confidence interval for the difference between the ATTs in the six months postadoption, we find a difference of 0.16 (95% CI: 0.090, 0.230), which is significantly different from zero with a *p*-value less than 0.001. In fact, we find that, for every time period in the six months after adoption, the point estimates for the group of users and nonusers are statistically different from each other with a *p*-value of less than 0.001 except for time period 1 in which the *p*-value for the differences equals 0.017.

Reassuringly, we find that our positive effects are driven by the 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.

Next, we examine the relationship between the intensity of dashboard usage and performance outcomes. For each retailer, we compute the average monthly number of dashboard events and split the sample of retailers to those above and below the median average postadoption usage.²¹ We use SynthDiD to compute the ATTs for each intensity subgroup. Column (B) of Table A.3 presents the results, demonstrating that intensity of analytics service usage matters. Overall, retailers with a larger average number of monthly events reap additional increases in revenues compared with those with less activity. (For example, the ATT for revenues six months after adoption is –0.018 (95% CI: –0.071, 0.035) for the below-median usage group versus 0.099 (95% CI: 0.031, 0.167) for the above-median group.)

Finally, we examine the usage of reports. The dashboard provides five main reports as detailed in Section 2. We split the data to 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 80% 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

Figure 4. (Color online) Treatment Effects of Users vs. Nonusers



Notes. ATTs (and 95% confidence intervals) of adopting analytics on revenues over time. Time 0 indicates the month of adoption. Other times are relative to adoption. Column (A) of Table A.3 presents the corresponding coefficients.

described, and then, for each retailer, we compute the average number of monthly reports they examined. In the first test, we compare SynthDiD estimates between retailers that viewed at least one report and 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 with those that did not, and retailers that accessed reports more frequently had larger gains compared with those that accessed reports less frequently. Columns (C) and (D) of Table A.3 present the results for these SynthDiD analyses. (For example, the ATT for revenues six months after adoption is -0.009 (95% CI: $-0.064, 0.028$) for the below-median report usage group versus 0.077 (95% CI: $0.012, 0.142$) for the above-median group.)

To summarize, we find that the intensity of dashboard usage 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 dashboard usage.

4.4. 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 SynthDiD to estimate heterogeneous effects by comparing the ATTs between subgroups of large and small firms.

Because we do not have an external measure of firm size (such as number of employees), we use revenues and number of transactions as proxies for size. To avoid using the same observations for both determining firm size as well as matching pretrends, we make sure there

is enough time difference between observations used to determine size and those used to match pretrends. To do so, we focus on firms that adopted in the first half of 2017, and we use two proxies for retailer size: (i) median split of the average monthly revenues in 2015 and (ii) median split of the average number of transactions in 2015. Table A.4 presents the results. For both revenue and transaction medians, retailer channels with below-median size exhibit a statistically significant increase in revenues after adoption, whereas the above-median retailers exhibit marginal and small decreases or no statistically significant changes in performance after adoption. Specifically, the ATT for revenues six months after adoption is 0.224 (95% CI: $0.064, 0.384$) for the below-median revenue-based size group versus -0.033 (95% CI: $-0.108, 0.042$) for the above-median group. Similarly, the ATT for revenues six months after adoption is 0.220 (95% CI: $0.069, 0.370$) for the below-median transaction-based size group versus -0.019 (95% CI: $-0.110, 0.073$) for the above-median group. Additionally, we compute the confidence intervals for the differences between the ATTs in the six months postadoption between the below- and above-median firms and find statistically significant differences. Therefore, we conclude that smaller retailers reap more of the benefits of dashboard adoption. This effect might be expected with descriptive analytics; larger and well-established retailers might have already optimized their performance and experience 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 translate to changes in customer behavior. These changes in customer behavior are 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 operate to drive the gains in revenue described in Section 1.

The three mechanisms we describe previously (direct, complementary, and unrelated) provide us with testable empirical predictions to understand how analytics operate to help retailers. For the direct mechanism, we expect to see changes in retailer decisions that are unrelated to the integration of additional technologies, and we 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 (nonusers).

For the complementary mechanism, we expect to see that the highest gain from adopting analytics appears for retailers that adopted additional technologies, but this gain only happens 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 is consistent with this 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 the unrelated 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 an unrelated process is the mechanism behind the effect, the effect should also disappear when we instrument for the timing of adoption.

We note that, ideally, we would perform a causal mediation analysis (Imai et al. 2010a, b; Pearl 2014) to disentangle between the different mechanisms. However, our data do 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, but more research with better data is needed to fully explore it.²²

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 CAC—retailers often know the value of a customer (in terms of the average margin), and if the CAC is higher than the desired margin of the retailer, it is obvious they should change their advertising spending. However, 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 in either size or 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 the adoption of analytics.

5.1. Decisions Driven by the Adoption of Analytics

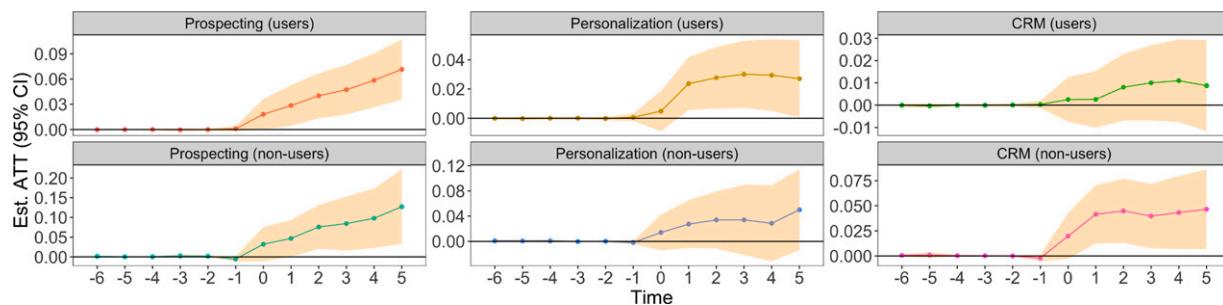
To analyze the changes in observable retailer actions after adoption, we use the SynthDiD approach. The application of the method is similar to that of Section 4.2.2, but in the current analysis, the dependent variables are the level of the actions taken (e.g., amount of discounts, price variation, or advertising allocation) during the observation window.

In a few cases, such as those of advertising and pricing, our data are limited to retailers that had this data available and made it accessible to the analytics service. This limits the analysis only to retailers that had connected their Google Analytics advertising or Shopify store data at the time of the data collection, resulting in 87% of the retailers using Google Analytics for the advertising analysis and 54.3% of the retailers using Shopify for the pricing analysis. We confirm that the previous main results of the effect of analytics hold for each subsample.

For pricing, we analyze the transaction data of the retailers and compute five measures as the dependent variables: (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 yields statistically significant differences after adoption of the analytics service (for users or nonusers of the dashboard).

For advertising, we use the retailers' own data collected through Google Analytics and estimate spending data 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 do not exhibit any statistical difference in spending after adopting analytics. This is the case for all retailers

Figure 5. (Color online) Adoption of Subtechnologies



Notes. ATTs (and 95% confidence intervals) of adopting analytics on personalization, prospecting, and CRM technologies over time. Time 0 indicates month of adoption. Other times are relative to adoption.

regardless of analytics usage.²³ The SpyFu data also did not reveal any meaningful changes in advertising behavior after adoption.²⁴

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.²⁵ Examples of web technologies include A/B testing, advertising retargeting, product recommendation, and personalization. The data covers 97% of the retailers in our data set. 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 and tracking (565 different technologies), advertising (532), e-commerce (263), email hosting (121), and payment (71). For the Builtwith.com data, we only observe the adoption of technology and not the usage. However, many of these technologies do not require active usage but only installation.

Table A.5 reports the SynthDiD coefficients for overall technologies as well as the three leading technologies for both users and nonusers.²⁶ Although the pretrend matching for nonusers is noisier, we still detect increases in adoption of the analytics and 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 and tracking, advertising, and e-commerce technologies.

To better understand retailers' decisions, we further break down the analytics and tracking, advertising, and e-commerce categories into subcategories of more specific technologies. We obtain six meaningful subcategories: CRM, website design and optimization, lead generation and prospecting, personalization, other advertising technologies, and other e-commerce technologies. Prospecting technologies are focused on attracting new users

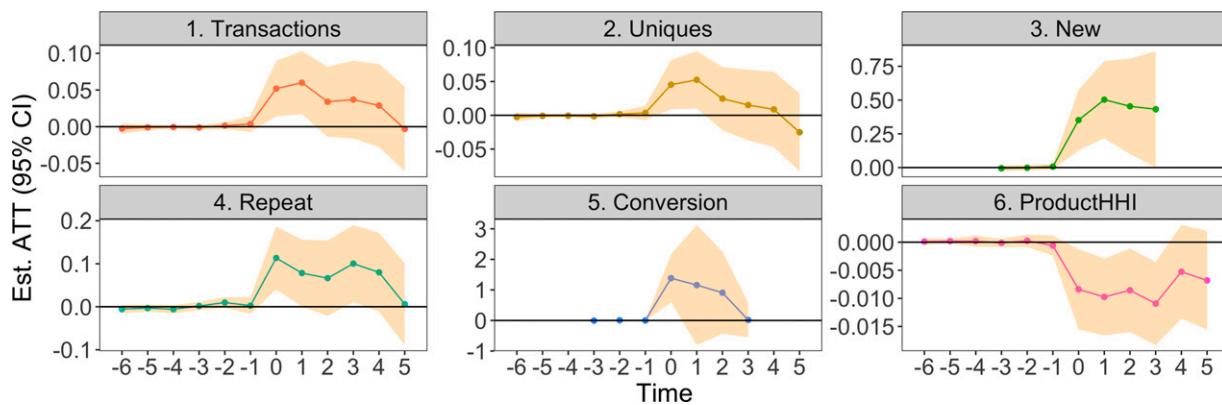
to websites, whereas personalization technologies are focused on optimizing the onsite experience of existing visitors and retargeting previous visitors.

Using the same approach as before, we analyze the number of technologies in each subcategory that retailers use as the dependent variable. We find an increase in prospecting, personalizing, and CRM technologies for both users and nonusers. The other subcategories do not exhibit significant changes after analytics adoption. Figure 5 displays the results for personalization, prospecting, and CRM technologies. Note that, whereas CRM for users and personalization for nonusers have a positive pattern postadoption, they are not statistically significant.

We conclude that the adoption of descriptive analytics was associated with the adoption of additional CRM, personalization, and prospecting technologies. There is no evidence of association with price changes, advertising allocation changes, changes in advertising spending, 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 (nonusers). One explanation is that retailers were encouraged to adopt both a descriptive dashboard and CRM, personalization, and prospecting technologies. However, the fact that we only observe an increase in revenue for adopters who are users of the dashboard 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 CRM, 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

Figure 6. (Color online) Customer Behavior Outcomes for Users

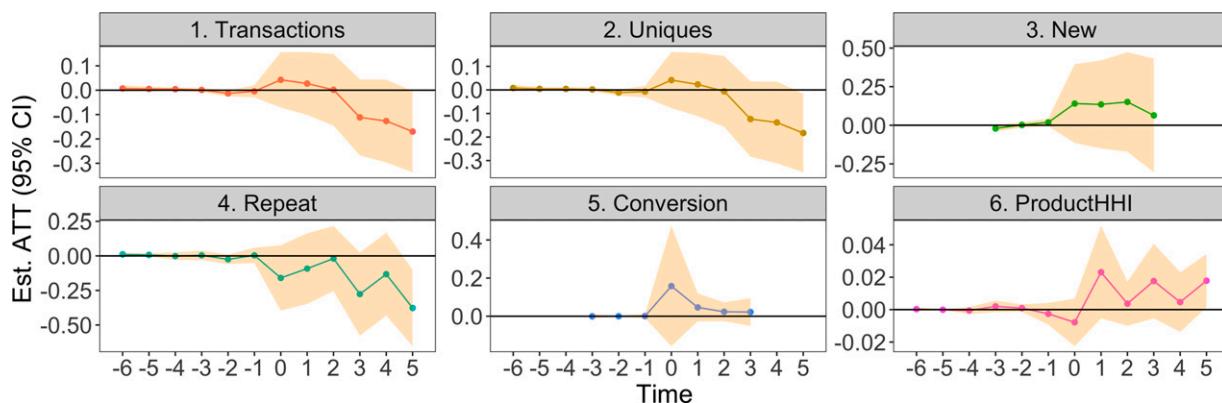
Notes. ATTs (and 95% confidence intervals) of adopting analytics on customer behavior-related outcomes over time. The analyses are limited to retailers that used the dashboard (or had high frequency of usage for the *New* variable). Time 0 indicates month of adoption. Other times are relative to adoption.

through recommender systems (Brynjolfsson et al. 2011b) as well as conversion rates from visits into purchases. Improved CRM technologies should increase repeat purchase rates. If pricing were to change, we would expect to see a change in basket size (the amount purchased) and potentially the number of products purchased. Advertising optimization would affect the number of new visitors to the store and would also show an improved CAC.

We use the following dependent variables to test for changes in customer behavior that 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 U.S. dollars (*Basket*), (v) average weekly amount of revenue from repeat customers (*Repeat*), (vi) CAC, (vii) conversion rates

from website visits to transactions²⁷ (*Conversion*),²⁸ (viii) unique number of products sold (*Products*), and (ix) product concentration measured using the Herfindahl index (HHI) (*ProductHHI*).²⁹

Figures 6 and 7 report the results of a SynthDiD analysis in which these variables were used as dependent variables. We only report those variables that exhibit statistically significant changes in outcomes after adoption. Figure 6 presents the results only for those retailers that used the analytics service. For users, we observe an increase in *Transactions*, *Uniques*, *New*, *Repeat*, and *Conversion* and a decrease in *ProductHHI*. The three-month ATTs are significantly different from zero for all variables, and the six-month ATTs are significantly different from zero for *Transactions*, *Repeat*, and *ProductHHI*. These changes are consistent with our predictions that CRM, personalization, and prospecting technologies contribute to a changing mix of customers in the store, but they do

Figure 7. (Color online) Customer Behavior Outcomes for Nonusers

Notes. ATTs (and 95% confidence intervals) of adopting analytics on customer behavior-related outcomes over time. The analyses are limited to retailers that did not use the dashboard (or had low frequency of usage for the *New* variable). Time 0 indicates month of adoption. Other times are relative to adoption.

not necessarily change how much customers spend. Additionally, the changes in *ProductHHI* suggest that the assortment of products purchased has changed after adoption. This change in product assortment can result from two sources we observe in the results: 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 CRM and personalization efforts increase the likelihood of existing customers to try new products. This interpretation 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).

Finally, we do not find evidence for changes in *Basket* or the CAC, which is consistent with the absence of changes in prices and advertising after the adoption of analytics reported in Section 5.1. 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 7 reports the results of the same analysis but for adopting retailers that did not use the dashboard.³⁰ As can be seen in the figures, whereas the preadoption period matching worked well for all outcome variables, retailers that did not use the service exhibit null effects for all of the variables. None of the three-month ATTs are different from zero, and among the six-month ATTs the only variable significantly different from zero is *Repeat* (estimated at -0.176 with *p*-value 0.071). This is particularly striking because this effect was positive for users of the dashboard and because nonusers invested in CRM technologies, which would presumably increase repeated revenues.

The results are consistent with our prediction on how different firm actions contribute to the observed gain in revenues. As before, we caution that these results should be interpreted carefully because it is difficult to disentangle the effect of two simultaneous 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 the 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 retailers to make. Given the nature of the descriptive dashboard, however, 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 or in customer behavior).

Overall, the evidence suggests that, whereas all retailers adopt the focal analytics service and are also likely to adopt additional CRM, 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 sold products, and a corresponding increase in performance outcomes such as revenues and transactions. Therefore, we conclude that the descriptive dashboard and the CRM, personalization, and prospecting technologies are complementary but they require usage to realize the benefits of the analytics service. Overall, it appears that a benefit of using a descriptive dashboard is that it allows retailers to evaluate how different technologies impact their outcomes and fine-tune them. Put together with the results on usage and the IV results, we conclude that the mechanism behind improvement in firm outcomes resulting from descriptive analytics adoption is the complementary mechanism.

In the context of our conclusions, there are two alternative interpretations that emerge and can potentially explain the results. First, nonusers of the dashboard may also be nonusers of the technologies, which may explain why they do not accrue the benefits of these technologies. Because we do not observe usage of technologies, we are unable to rule that out. In that case, the benefit for users may be simply a result of usage and not the dashboard adoption. However, our IV strategy that shifts timing of the dashboard but not the technologies adoption suggests that the improvement in firm outcomes is due, at least partially, to the adoption of analytics. In addition, many of the CRM, personalization, and prospecting technologies do not require human intervention after the initial setup, which also alleviates this concern (e.g., the Facebook Pixel technology). Second, it is possible that the decision maker in the firm who adopts the dashboard also controls other technology adoption but does not control pricing or advertising decisions. Although we do not observe the number of employees in each firm, most of the firms in the data are small and are likely to have one or two employees. Moreover, as demonstrated in Section 4.4, most of the benefit is accrued by the smaller firms in our sample. Therefore, we believe it is typically the same person who makes all of the aforementioned decisions, which may mitigate this concern.³¹

6. Conclusion

Although the interest in marketing 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 because of a lack of data and also the inability to explain what drives the benefits that are observed. In this paper, we describe the effect of the adoption of a descriptive analytics service by a wide variety of online retailers, and we make an effort to provide causal estimates of this effect. Our unique data set allowed us to not only provide estimates of the value created by descriptive analytics but also to explore 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 4%–10% in the six months after adoption. Although this range is wide and results from using multiple estimation approaches, we demonstrate 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 technologies.

The magnitude of our estimates is economically meaningful at 4%–10%, suggesting an increase of US\$170–US\$449 per week for a median retailer. Prior literature that investigates the relationship between data analytics and firm performance finds a positive relationship of 3%–7% greater productivity for firms that adopt data-driven decision making or big data assets (e.g., Brynjolfsson et al. 2011a, Brynjolfsson and McElheran 2016, Müller et al. 2018). However, their adoption and outcome measures were coarser. Recent papers that investigate the adoption of analytics technologies using more detailed data found larger effect sizes in line with our findings. Koning et al. (2022) find an increase of 10% in page views after adoption of A/B testing technologies by startup. Runge and Nair (2021) find an average increase of 42% in conversion rates among low-spending firms and 21.5% in high-spending firms after adopting randomized control trial tools on Facebook. Finally, Anderson et al. (2020) find an increase of 45% in sales and 36% in profits after adoption of descriptive analytics by microentrepreneurs in Rwanda.

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. Regarding this concern, we note that, 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, whereas the focus is on adoption of one focal analytics service, our data and analysis include 1,509 e-commerce firms from varying industries and countries. These are not small firms, but ones with annual revenues of \$100,000 or more, and nearly all of them have been using Google Analytics prior to adoption of the service.³² Therefore, we believe that our findings are generalizable to adoption of other descriptive dashboards by a variety of firms.

The second concern is about the endogenous timing of adoption, which may cause simultaneity and threaten the causal interpretation of our findings. Potentially, firms could have selected an ideal time to adopt in which they believe that the dashboard would be most effective, or 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 unrelated mechanism in Section 1). A unique result that we provide in our analysis is to show that the *usage* of analytics and not its adoption, *per se*, is what drives the improved firm performance. Further, the SynthDiD method we employ is robust to correlation between unobserved time-varying factors and a retailer's decision to adopt, and the IV strategy that shifts the timing of adoption also shows robust results, suggesting that the improvement in firm outcomes is due, at least partially, to the adoption of analytics.

A third concern is that our estimates might be biased because of selection of the retailers that chose to adopt the dashboard. Indeed, Section 4.4 shows that there is heterogeneity in the effect with smaller retailers experiencing a larger benefit from the dashboard compared with larger one. This result is not surprising given that larger retailers have probably already optimized much of their operation and have less to benefit from simpler analytics. The result also alludes to the potential threat of selection bias to our analysis—the results would generalize to firms that are similar to those that chose to adopt the dashboard in our data and might not generalize to very large firms (which are rare in our data) or those not interested in using the dashboard as the results for nonusers show. 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) prevents us from performing this analysis, which is left for future work. Additionally, there might be other mediating firm actions that we do not observe.

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, in which firms within the sample are difficult to compare. Focusing on online retailers provides 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. Because 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 the resulting assortment of purchases.³³ These changes may be due to the firm changing the inventory of products it sells, but they may also be due to the firm affecting the type of products to which consumers are exposed 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 CRM, personalization, and prospecting 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 is not an "install and forget" solution but, rather, one that requires continuous monitoring and from which the benefits may accrue over time with experience and also with additional investment. The analysis shows that both users and nonusers 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 users to generate their own insights from the data, they provide retailers with 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 to integrate new technologies. In turn, this research suggests future avenues for researchers to create better predictive and prescriptive solutions to solve these challenges for marketers.

Acknowledgments

The authors are grateful for constructive feedback by the Senior Editor and the review team, and for helpful comments from Eva Ascarza, Eli Ben-Michael, Kathy Li, Ilya Morozov, Simha Mummalaneni, Davide Proserpio, Dave Reibstein, Christophe Van den Bulte, Jeff Wooldridge, and Lynn Wu as well as seminar participants at Ben Gurion University, Brock University, Harvard Business School, Hebrew University of Jerusalem, Northeastern University, Northwestern University, University of Massachusetts Amherst, and Vrije Universiteit Amsterdam and participants at the Marketing Analytics Symposium Sydney, the Marketing Science Conference, the Northeastern Marketing Conference, and the Marketing in Israel Conference. They are indebted to the anonymous analytics service provider that provided the data. Shawn Zamechek and Hengyu Kuang provided excellent research assistance.

Appendix

A.1. Constructing the SynthDiD Event Study Plot

SynthDiD is designed to minimize the mean squared error of a target estimated ATT and not to separately measure effects in specific time periods. The event study plot we construct is, therefore, used to illustrate the effect although aggregating the standard errors that appear in this plot into an ATT requires taking serial correlation into account, which is done automatically by SynthDiD but cannot be done using the standard errors from the plot (or from the associated period-by-period standard errors we report in Appendix A.3).

To compute the treatment effects preadoption, we compute for each adoption cohort r and each time period t between $r + \ell_{\min}$ and $r - 1$,

$$\widehat{\tau}_{r(t-r)} = \left(\frac{\sum_{i \in N_r^{tr}} \log(Y_{ijt} + 1)}{|N_r^{tr}|} - \left(\widehat{\omega}_0 + \sum_{i \in N_r^{co}} \widehat{\omega}_{ij} \log(Y_{ijt} + 1) \right) \right) \cdot \widehat{\lambda}_t.$$

Because the weights $\widehat{\lambda}_t$ for time periods before adoption sum up to one, summing up $\widehat{\tau}_{r(t-r)}$ yields values that are approximately zero, which shows a good fit between treatment outcomes and the synthetic control preadoption. The standard errors for these values are computed using the jackknife method (or in case there is only one treated unit, the placebo method). The values for each cohort are then averaged, and the standard errors are aggregated appropriately.

Computing the cohort-level effects postadoption is done in a similar manner. For each adoption cohort r and each time period t between r and $r + \ell_{\max}$, we compute

$$\widehat{\tau}_{r(t-r)} = \left(\frac{\sum_{i \in N_r^{tr}} \log(Y_{ijt} + 1)}{|N_r^{tr}|} - \left(\widehat{\omega}_0 + \sum_{i \in N_r^{co}} \widehat{\omega}_{ij} \log(Y_{ijt} + 1) \right) \right).$$

In this case, we do not weight the estimated effect by $\widehat{\lambda}_t$ because, postadoption, the weight is just $\frac{1}{\ell_{\max+1}}$. Averaging

the effects within cohorts and the resulting averages across cohorts then produces the ATT reported by SynthDiD. The standard error for the aggregate ATT is computed using the jackknife method (or the placebo method as before) directly for the aggregated value, which takes into account potential serial correlation between effects across time.

A.2. Computing Subgroup Effects Using SynthDiD

There are three possible approaches to compute subgroup effects using SynthDiD. The first is to split the data into separate subsets for each subgroup and compute the SynthDiD effects for each subset. This approach, therefore, compares each treated unit to control units in the same subgroup. The disadvantage of this method is that each subset uses a different control group, resulting in a different synthetic control, and thus, the effects cannot be compared across subgroups.

The second approach is to use a control group comprising all nonadopters (regardless of subgroup membership) in each balanced panel cohort analysis and then switch the treatment units to the subgroup to be analyzed and perform a SynthDiD analysis. The advantage of this method is that the control weights $\hat{\omega}_i$ computed for each treated subgroup match the pretrends the best, but the disadvantage is that every control unit receives a different weight for each treatment subgroup being analyzed, which again makes comparing the results harder as the synthetic control used to estimate effects is different.

The third approach, which is the one we implement and report,³⁴ uses all the data to estimate the SynthDiD control weights $\hat{\omega}_i$, but it then computes the treatment effects $\hat{\tau}_{r(t-r)}$, which are defined in Section A.1 using *only* the treated subgroup units as N_r^{tr} . This has the advantage that all units have the same synthetic control, but it has the disadvantage that the pretrend match might not be as good because it was designed to match the average outcome of all treated units and not just a subgroup.

A.3. Additional Details and Analyses Using SynthDiD

Table A.1 shows the number of treated and control units in each cohort.

A.3.1. Main Effect: Robustness. We perform a series of robustness tests to our main effect results. These are available in Table A.2. In this table, column “Baseline” presents the results corresponding to Figure 3.

First, because of the importance of preperiod synthetic control matching, we show that our effects are robust to the choice of the number of lags and leads; we display overall robust results for additional 12-month windows with seven, eight, and nine preadoption lags (see columns “Lag-1,” “Lag-2,” and “Lag-3” in the appropriate tables).

Second, to verify that the effects we measure are not a result of the aggregation at the retailer-channel level, we aggregate the data to the retailer level, redefine “after adoption” to occur once the first channel data are added to the dashboard, and estimate the models at the retailer level. Corresponding results appear in each table as the “Company level” column.

Finally, to demonstrate the generalizability of our results, we collected additional data on 1,091 retailers that

Table A.1. Number of Units Used in SynthDiD Analysis

Cohort	Treatment	Control
Jan 2016	57	431
Feb 2016	131	446
Mar 2016	58	418
Apr 2016	72	406
May 2016	71	430
Jun 2016	67	417
Jul 2016	55	413
Aug 2016	49	435
Sep 2016	53	497
Oct 2016	49	494
Nov 2016	67	538
Dec 2016	56	524
Jan 2017	30	551
Feb 2017	18	557
Mar 2017	39	566
Apr 2017	33	576
May 2017	45	539
Jun 2017	32	494
Average	54.6	485

Table A.2. SynthDiD Robustness

Lag	Baseline	Lag-1	Lag-2	Lag-3	Company level	2017/2018 sample
-9					-0.0003 (0.002)	
-8			-0.001 (0.002)	-0.0005 (0.002)		
-7		-0.001 (0.002)	-0.0004 (0.002)	0.00002 (0.001)		
-6	-0.001 (0.003)	-0.001 (0.002)	-0.001 (0.001)	-0.0003 (0.001)	-0.001 (0.003)	-0.001 (0.003)
-5	0.00000 (0.002)	-0.00001 (0.002)	-0.00001 (0.001)	0.00003 (0.001)	0.00000 (0.002)	-0.0002 (0.001)
-4	-0.0004 (0.002)	-0.0001 (0.001)	0.0001 (0.001)	0.00004 (0.001)	-0.0003 (0.002)	-0.001 (0.001)
-3	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	0.0001 (0.002)
-2	-0.001 (0.003)	-0.001 (0.003)	0.0002 (0.002)	-0.0001 (0.003)	-0.001 (0.003)	-0.003 (0.003)
-1	0.004 (0.005)	0.004 (0.005)	0.002 (0.006)	0.002 (0.005)	0.004 (0.006)	0.005 (0.006)
0	0.056** (0.019)	0.053** (0.019)	0.048* (0.019)	0.036 ⁺ (0.019)	0.057** (0.020)	0.087** (0.022)
1	0.062** (0.022)	0.059** (0.022)	0.041 ⁺ (0.022)	0.017 (0.022)	0.060** (0.023)	0.075** (0.027)
2	0.050* (0.024)	0.037 (0.024)	0.021 (0.025)	-0.009 (0.025)	0.036 (0.026)	0.082** (0.030)
3	0.036 (0.026)	0.037 (0.026)	0.013 (0.027)		0.023 (0.029)	0.062* (0.031)
4	0.034 (0.028)	0.030 (0.029)			0.014 (0.030)	0.049 (0.031)
5	-0.010 (0.029)				-0.029 (0.031)	0.042 (0.032)
Three months	0.056** (0.019)	0.050** (0.019)	0.037* (0.019)	0.014 (0.018)	0.051* (0.020)	0.081** (0.023)
Six months	0.038 ⁺ (0.020)				0.027 (0.022)	0.066** (0.024)

Significance level: 10% (+), 5% (*), 1% (**).

Table A.3. Usage Results

Lag	A: Usage dummy		B: Usage intensity		C: Report usage		D: Report intensity	
	Users	Nonusers	Below	Above	Users	Nonusers	Below	Above
-6	-0.002 (0.003)	0.007 (0.006)	0.001 (0.004)	-0.004 (0.006)	-0.002 (0.004)	0.0005 (0.006)	0.001 (0.004)	-0.004 (0.005)
-5	-0.001 (0.002)	0.010 ⁺ (0.006)	0.001 (0.003)	-0.002 (0.004)	-0.001 (0.003)	0.006 (0.004)	0.003 (0.003)	-0.004 (0.004)
-4	-0.001 (0.002)	0.003 (0.003)	-0.003 (0.002)	0.003 (0.002)	0.0004 (0.002)	-0.003 (0.005)	-0.002 (0.002)	0.002 (0.002)
-3	-0.001 (0.002)	0.0004 (0.006)	-0.006* (0.003)	0.006* (0.003)	-0.002 (0.002)	0.001 (0.004)	-0.006* (0.003)	0.005 ⁺ (0.003)
-2	0.001 (0.003)	-0.017* (0.008)	-0.003 (0.004)	0.002 (0.004)	-0.001 (0.003)	-0.004 (0.007)	-0.004 (0.004)	0.002 (0.004)
-1	0.005 (0.006)	-0.003 (0.014)	0.010 (0.007)	-0.005 (0.009)	0.005 (0.006)	-0.001 (0.011)	0.008 (0.007)	-0.002 (0.009)
0	0.070** (0.021)	-0.037 (0.057)	0.039 (0.025)	0.086** (0.033)	0.069** (0.023)	0.008 (0.043)	0.049 ⁺ (0.025)	0.068* (0.032)
1	0.070** (0.024)	-0.017 (0.069)	0.028 (0.029)	0.111** (0.035)	0.070** (0.025)	0.018 (0.055)	0.042 (0.031)	0.084* (0.033)
2	0.054* (0.027)	-0.097 (0.080)	0.0002 (0.033)	0.094* (0.040)	0.046 ⁺ (0.028)	-0.006 (0.062)	0.003 (0.033)	0.082* (0.040)
3	0.043 (0.030)	-0.123 (0.090)	-0.029 (0.036)	0.106* (0.046)	0.037 (0.031)	-0.034 (0.068)	-0.032 (0.037)	0.099* (0.044)
4	0.042 (0.031)	-0.183 ⁺ (0.094)	-0.043 (0.037)	0.108* (0.050)	0.033 (0.033)	-0.059 (0.071)	-0.029 (0.037)	0.075 (0.048)
5	-0.002 (0.032)	-0.226* (0.096)	-0.103** (0.038)	0.090 ⁺ (0.051)	-0.0001 (0.033)	-0.147* (0.075)	-0.089* (0.039)	0.055 (0.049)
Three months	0.065** (0.021)	-0.050 (0.060)	0.022 (0.025)	0.097** (0.031)	0.062** (0.022)	0.007 (0.047)	0.031 (0.026)	0.078** (0.030)
Six months	0.046* (0.023)	-0.114 ⁺ (0.068)	-0.018 (0.027)	0.099** (0.035)	0.042 ⁺ (0.024)	-0.037 (0.052)	-0.009 (0.028)	0.077* (0.033)

Significance level: 10% (+), 5% (*), 1% (**).

Table A.4. Heterogeneity by Size

Lag	Sample	Revenue medians		Transaction medians	
		Below	Above	Below	Above
-6	-0.0001 (0.006)	-0.015 (0.009)	0.014 ⁺ (0.009)	-0.014 (0.010)	0.012 (0.008)
-5	0.0005 (0.004)	-0.002 (0.007)	0.003 (0.004)	-0.008 (0.005)	0.008 (0.005)
-4	-0.001 (0.002)	-0.005 (0.004)	0.003 (0.002)	-0.008 ⁺ (0.004)	0.005** (0.001)
-3	0.001 (0.003)	-0.001 (0.005)	0.003 (0.003)	0.004 (0.004)	-0.001 (0.003)
-2	-0.003 (0.005)	0.005 (0.007)	-0.010 (0.007)	0.003 (0.006)	-0.008 (0.007)
-1	0.003 (0.009)	0.018 (0.014)	-0.013 (0.012)	0.023 ⁺ (0.013)	-0.016 (0.012)
0	0.056 (0.034)	0.142* (0.061)	-0.014 (0.038)	0.227** (0.061)	-0.076* (0.036)
1	0.101* (0.041)	0.201* (0.079)	0.013 (0.038)	0.223** (0.079)	0.002 (0.038)
2	0.119* (0.055)	0.254* (0.099)	0.010 (0.058)	0.265** (0.096)	0.012 (0.061)
3	0.113* (0.055)	0.278** (0.098)	-0.034 (0.055)	0.252** (0.089)	-0.001 (0.069)
4	0.086	0.240 ⁺	-0.045	0.193 ⁺	0.004

Table A.4. (Continued)

Lag	Sample	Revenue medians		Transaction medians	
		Below	Above	Below	Above
5	(0.062)	(0.125)	(0.056)	(0.116)	(0.070)
	0.038	0.234*	-0.126*	0.159	-0.052
	(0.062)	(0.114)	(0.064)	(0.106)	(0.078)
Three months	0.092*	0.199**	0.002	0.238**	-0.021
	(0.038)	(0.073)	(0.033)	(0.071)	(0.034)
Six months	0.086*	0.224**	-0.033	0.220**	-0.019
	(0.043)	(0.082)	(0.038)	(0.077)	(0.047)

Notes. Column (1) reports the period-by-period estimates for the sample of firms that adopted the service between January 2017 and June 2017. The next columns use the same sample of firms and separate them to subgroups based on performance in 2015.

Significance level: 10% (+); 5% (*); 1% (**).

adopted the service in 2018. For this sample, we only observe a subset of the variables and, thus, use it only for robustness tests of our main effect. We use these firms to repeat the analysis without the 2016 adopters, using all 2017 retailers as the treatment group and the 2018 retailers as the control-only group. To ensure sufficient overlap, we restrict the sample to include retailers that had at least one observation in 2016. This yielded 765 firms out of the 958 firms that adopted the dashboard in 2017 as treated (833 retailer–channel combinations), and for 2018, we could use 757 firms out of 1,091 (886 retailer–channel combinations). Column “2017/2018 sample” reports the

results. The results are similar to the baseline albeit with slightly larger effect sizes postadoption. Additionally, the effect stays significantly positive for one additional time period.

As the table shows, the SynthDiD method matches the treatment and control units such that the difference between them is indistinguishable from zero for all of the different specifications.

A.3.2. Additional Tables. The tables in this section present period-by-period estimates and the three- and six-month ATTs. Table A.3 presents estimates of the effect of

Table A.5. Technology Adoption

Lag	Overall		Analytics and tracking		Advertising		E-commerce	
	Users	Nonusers	Users	Nonusers	Users	Nonusers	Users	Nonusers
-6	0.00003	-0.001	-0.0001	0.0003	-0.00005	0.001	-0.001	0.001 ⁺
	(0.0004)	(0.001)	(0.001)	(0.0004)	(0.0002)	(0.001)	(0.001)	(0.001)
-5	-0.0001	0.001*	0.00004	0.0001	-0.00002	-0.0001	0.0001	-0.002
	(0.0002)	(0.0003)	(0.0004)	(0.0004)	(0.0003)	(0.0001)	(0.001)	(0.002)
-4	0.00000	0.00000	-0.0002	0.0005	-0.00001	-0.0002	0.0001	-0.0001
	(0.00001)	(0.00000)	(0.001)	(0.001)	(0.0003)	(0.0003)	(0.0002)	(0.001)
-3	-0.0001	0.0004	-0.0001	0.00000	0.0002	-0.0002	-0.0001	-0.0002
	(0.0002)	(0.0005)	(0.0002)	(0.0003)	(0.001)	(0.001)	(0.0003)	(0.001)
-2	-0.0001	-0.0001	-0.0002	-0.001	0.001	-0.002	-0.0001	0.0002
	(0.0003)	(0.0005)	(0.0004)	(0.001)	(0.001)	(0.003)	(0.0003)	(0.001)
-1	0.0002	-0.0001	0.001	0.0001	-0.001	0.002	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.001)	(0.004)
0	0.004	0.017 ⁺	0.013*	0.026	0.019 ⁺	0.009	0.007	0.006
	(0.004)	(0.009)	(0.006)	(0.017)	(0.010)	(0.024)	(0.005)	(0.012)
1	0.010+	0.021*	0.022*	0.053*	0.048**	0.021	0.011	0.005
	(0.006)	(0.010)	(0.009)	(0.021)	(0.013)	(0.023)	(0.007)	(0.013)
2	0.012 ⁺	0.022 ⁺	0.029**	0.075**	0.047**	0.026	0.019*	0.026
	(0.006)	(0.012)	(0.010)	(0.023)	(0.015)	(0.027)	(0.008)	(0.020)
3	0.008	0.010	0.025*	0.067*	0.052**	0.019	0.021*	0.038
	(0.007)	(0.014)	(0.011)	(0.026)	(0.017)	(0.033)	(0.009)	(0.025)
4	0.006	0.001	0.025*	0.067*	0.048**	0.005	0.022*	0.055*
	(0.008)	(0.015)	(0.011)	(0.029)	(0.017)	(0.034)	(0.009)	(0.026)
5	0.002	-0.0004	0.030*	0.058 ⁺	0.060**	-0.007	0.019 ⁺	0.041
	(0.009)	(0.016)	(0.013)	(0.032)	(0.019)	(0.037)	(0.010)	(0.027)
Three months	0.009 ⁺	0.020*	0.021**	0.051**	0.038**	0.018	0.012*	0.012
	(0.005)	(0.009)	(0.008)	(0.019)	(0.011)	(0.022)	(0.006)	(0.013)
Six months	0.007	0.012	0.024**	0.058**	0.045*	0.012	0.016*	0.028 ⁺
	(0.006)	(0.011)	(0.009)	(0.022)	(0.013)	(0.026)	(0.007)	(0.017)

Significance level: 10% (+); 5% (*); 1% (**).

adoption on revenues, conditional on usage intensity, and Table A.4 presents estimates of the effect conditional on firm size. Table A.5 presents the analysis of subtechnologies for the mechanism section.

Endnotes

¹ This is according to the 2019 Gartner CMO Spend Survey and the 2017–2019 Forester Marketing Technology Services Outlook report, published in April 2018.

² For example, Brynjolfsson et al. (2011a) and Brynjolfsson and McElheran (2016) use survey data, Müller et al. (2018) use general measures of big data assets, and Brynjolfsson et al. (2011a) and Müller et al. (2018) use public firms' financial performance

³ Fewer than 7% of these retailers also sell on Amazon.com as third-party sellers.

⁴ For confidentiality reasons, we cannot provide identifying information about the analytics service provider.

⁵ We only observe a one-time snapshot of the subscription data from late 2018; in that time, there are no significant differences between users and nonusers of the service in subscription rates.

⁶ See <https://www.engadget.com/2016-11-03-ecommerce-platform-market-share-looking-at-the-companies-that-d.html>, accessed November 1, 2020.

⁷ See <https://aheadworks.com/blog/ecommerce-market-2017/>, accessed November 1, 2020.

⁸ In this setting, because the dashboard was available to all retailers and installation requires awareness, the ATT is equivalent to the local average treatment effect and not the ATE in the population.

⁹ A randomized controlled trial would not resolve this issue completely because of compliance and statistical power challenges. For example, Anderson et al. (2020) analyze the impact of basic analytical training on small, mom-and-pop stores in Rwanda. Despite using randomized controlled field experiments, Anderson et al. (2020) do not estimate the ATE in the general population because they recruited a specific group of entrepreneurs based on their growth potential. In our setting, adopters operate technically sophisticated e-commerce stores and are expected to have these very basic analytical skills.

¹⁰ Borusyak et al. (2021) show that using both pretreatment and posttreatment data to test for parallel trends conflates multiple identifying assumptions and tries to simultaneously test the parallel-trends assumption and estimate the treatment effects under this assumption. Instead, they recommend using only untreated units (never treated or not yet treated).

¹¹ That is, $N_r = N_r^{co} \cup N_r^{tr}$.

¹² We omit the full details of the estimation, which can be found in section 2 of Arkhangelsky et al. (2021). Briefly, the method uses regularized ridge regression to regress the average of treatment unit outcomes on control unit outcomes while constraining the resulting weights to sum to one.

¹³ The exact conditions appear in assumption 2 of Arkhangelsky et al. (2021).

¹⁴ The data window is longer in this analysis to allow a balanced panel of the same length for all cohorts. Using the same truncation rule as in Section 4.2.1 doesn't affect the results qualitatively.

¹⁵ The control group is larger than the treated group because it comprises all later adopters from multiple future cohorts and not just units from one cohort.

¹⁶ Appendix A.1 describes the methodology used to construct this plot.

¹⁷ Column (1) of Table A.2 reports the period-by-period estimates that correspond to Figure 3.

¹⁸ To further evaluate these results, Web Appendix Table W-3 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 standard OLS effect size is less than half of the IV's third-stage estimate.

¹⁹ Appendix A.2 details the method of computing subgroup effects. Whereas the number of nonusers is relatively small, SynthDiD is designed 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.

²⁰ For preadoption matching, SynthDiD targets the ATT of the entire sample, and thus, estimates for individual time periods might be inaccurate, whereas their average is an unbiased estimate. The method also does not optimize the preadoption fit for each subgroup as can be seen in Figure 4, and each subgroup is compared with a mix of units from both groups as the control. We should, therefore, focus on the difference between the subgroups and not the absolute values of the period-level effects.

²¹ We use median split rather than quartiles or deciles to ensure preadoption matching for the treated subgroup. Results are similar but noisier if we split the data into more quantiles.

²² A randomized controlled trial in which analytics capabilities are allocated randomly to retailers will not suffice for this analysis. The mediators themselves need to be randomized as well.

²³ We caution that the Google Analytics advertising data are sparser, so they include lower coverage of time periods.

²⁴ We were able to identify the monthly advertising budget for 874 of the retailers (58%) and the number of terms and display ads for 486 of the retailers.

²⁵ This analysis excludes the focal analytics service, which is not detected by Builtwith.com.

²⁶ The remaining two categories do not show significant effects.

²⁷ We thank an anonymous reviewer for suggesting this analysis.

²⁸ For *New* and *Conversion*, we could only compute three lag and four lead periods because of data constraints. Whereas these measures are not perfect, the results are consistent with other metrics.

²⁹ Product-level data are 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 zero and one for each retailer-month.

³⁰ For *New*, we use low frequency of usage because there were not enough retailers that did not use the service to perform an analysis for this variable.

³¹ We thank an anonymous referee for pointing out this alternative interpretation.

³² For comparison, in our data, retailers had an average basket size of \$180 with a median of \$87 and annual revenues with an average of \$735,000 and a median of \$250,000. Shopify's retailers in 2019 had an average basket size of \$67–\$101 depending on region and an average annual revenue of \$74,000 (Shopify 2019).

³³ Jin and Sun (2019) also do not find price changes when analytics are adopted. Their setting is competitors on one specific online shopping platform.

³⁴ The first approach yielded similar but noisier results.

References

Abadie A (2021) Using synthetic controls: Feasibility, data requirements, and methodological aspects. *J. Econom. Literature* 59(2):391–425.

Abadie A, Diamond A, Hainmueller J (2010) Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *J. Amer. Statist. Assoc.* 105(490):493–505.

Akter S, Wamba SF, Gunasekaran A, Dubey R, Childe SJ (2016) How to improve firm performance using big data analytics

capability and business strategy alignment? *Internat. J. Production Econom.* 182:113–131.

Anderson SJ, Chintagunta P, Kaul R, Vilcassim NJ (2020) Impact of marketing analytics 1.0 on entrepreneur and firm performance: Field experiment evidence from Rwanda. Working paper, University of Texas, Austin, Austin, TX.

Arkhangelsky D, Athey S, Hirshberg DA, Imbens GW, Wager S (2021) Synthetic difference-in-differences. *Amer. Econom. Rev.* 111(12):4088–4118.

Bajari P, Chernozhukov V, Hortaçsu A, Suzuki J (2019) The impact of big data on firm performance: An empirical investigation. *AEA Papers Proc.* 109:33–37.

Ben-Michael E, Feller A, Rothstein J (2022) Synthetic controls with staggered adoption. *J. Royal Statist. Soc.: Series B (Statistical Methodology)*. Forthcoming.

Bertrand M, Duflo E, Mullainathan S (2004) How much should we trust differences-in-differences estimates? *Quart. J. Econom.* 119(1):249–275.

Blum BS, Goldfarb A (2006) Does the internet defy the law of gravity? *J. Internat. Econom.* 70(2):384–405.

Borusyak K, Jaravel X, Spiess J (2021) Revisiting event study designs: Robust and efficient estimation. Preprint, submitted August 27, <https://arxiv.org/abs/2108.12419>.

Bradlow ET, Gangwar M, Kopalle P, Voleti S (2017) The role of big data and predictive analytics in retailing. *J. Retailing* 93(1):79–95.

Brynjolfsson E, McElheran K (2016) Data in action: Data-driven decision making in U.S. manufacturing. Center for Economic Studies Working Paper 16-06.

Brynjolfsson E, Hitt L, Kim HH (2011a) Strength in numbers: How does data-driven decision making affect firm performance? Working paper, Stanford University, Stanford, CA.

Brynjolfsson E, Hu Y, Simester D (2011b) Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Sci.* 57(8):1373–1386.

Bughin J (2017) Ten big lessons learned from big data analytics. *Appl. Marketing Analytics* 2(4):286–295.

Callaway B, Sant'Anna PHC (2021) Difference-in-differences with multiple time periods. *J. Econometrics* 225(2):200–230.

Datta H, Knox G, Bronnenberg BJ (2018) Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Sci.* 37(1):5–21.

de Chaisemartin C, d'Haultfoeuille X (2020) Two-way fixed effects estimators with heterogeneous treatment effects. *Amer. Econom. Rev.* 110(9):2964–2996.

Delen D, Ram S (2018) Research challenges and opportunities in business analytics. *J. Bus. Analytics* 1(1):2–12.

Deng Y, Lambrecht A, Liu Y (2019) Spillover effects and freemium strategy in mobile app market. Working paper, University College London, London.

Germann F, Lilien GL, Rangaswamy A (2013) Performance implications of deploying marketing analytics. *Internat. J. Res. Marketing* 30(2):114–128.

Germann F, Lilien GL, Fiedler L, Kraus M (2014) Do retailers benefit from deploying customer analytics? *J. Retailing* 90(4):587–593.

Goodman-Bacon A (2018) Difference-in-differences with variation in treatment timing. Technical report, National Bureau of Economic Research, Cambridge, MA.

Han C (2016) Efficiency comparison of random effects two stage least squares estimators. *Econom. Lett.* 148:59–62.

Hanssens DM, Pauwels KH (2016) Demonstrating the value of marketing. *J. Marketing* 80(6):173–190.

Imai K, Keele L, Tingley D (2010a) A general approach to causal mediation analysis. *Psych. Methods* 15(4):309–334.

Imai K, Keele L, Yamamoto T (2010b) Identification, inference and sensitivity analysis for causal mediation effects. *Statist. Sci.* 25(1):51–71.

Imai K, Kim IS, Wang E (2022) Matching methods for causal inference with time-series cross-section data. *Amer. J. Political Sci.* Forthcoming.

Jin Y, Sun Z (2019) Information acquisition and the return to data. Working paper, University of Toronto, Toronto.

Koning R, Hasan S, Chatterji A (2022) Experimentation and startup performance: Evidence from a/b testing. *Management Sci.* Forthcoming.

Lambrecht A, Tucker C (2013) When does retargeting work? Information specificity in online advertising. *J. Marketing Res.* 50(5): 561–576.

LaValle S, Lesser E, Shockley R, Hopkins MS, Kruschwitz N (2011) Big data, analytics and the path from insights to value. *MIT Sloan Management Rev.* 52(2):21–31.

Lismont J, Vanthienen J, Baesens B, Lemahieu W (2017) Defining analytics maturity indicators: A survey approach. *Internat. J. Inform. Management* 37(3):114–124.

Manchanda P, Packard G, Pattabhiramaiah A (2015) Social dollars: The economic impact of customer participation in a firm-sponsored online customer community. *Marketing Sci.* 34(3): 367–387.

Mintz O, Bart Y, Lenk P, Reibstein D (2019) Drowning in metrics: How managers select and trade-off metrics for making marketing budgetary decisions. Preprint, submitted January 1, 2020, <https://dx.doi.org/10.2139/ssrn.3502600>.

Müller O, Fay M, vom Brocke J (2018) The effect of big data and analytics on firm performance: An econometric analysis considering industry characteristics. *J. Management Inform. Systems* 25(2):488–509.

Oestreicher-Singer G, Sundararajan A (2012) Recommendation networks and the long tail of electronic commerce. *Management Inform. Systems Quart.* 36(1):65–83.

Pauwels K, Ambler T, Clark BH, LaPointe P, Reibstein D, Skiera B, Wierenga B, Wiesel T (2009) Why, what, how, and what research is needed? *J. Service Res.* 12(2):175–189.

Pearl J (2014) Interpretation and identification of causal mediation. *Psych. Methods* 19(4):459–481.

Runge J, Nair H (2021) The role of randomized control trials in online demand generation: Exploratory evidence from Facebook. Preprint, submitted November 16, <https://dx.doi.org/10.2139/ssrn.3794028>.

Seddon PB, Constantinidis D, Tamm T, Dod H (2017) How does business analytics contribute to business value? *Inform. Systems J.* 27(3):237–269.

Shopify (2019) Shopify's state of commerce report, June 2019. Accessed November 1, 2020, <https://news.shopify.com/shopify-unveils-first-state-of-commerce-report>.

Wamba SF, Akter S, Edwards A, Chopin G, Gnanzou D (2015) How "big data" can make big impact: Findings from a systematic review and a longitudinal case study. *Internat. J. Production Econom.* 165:234–246.

Wang K, Goldfarb A (2017) Can offline stores drive online sales? *J. Marketing Res.* 54(5):706–719.

Wedel M, Kannan P (2016) Marketing analytics for data-rich environments. *J. Marketing* 80(6):97–121.

Wooldridge J (2005) Fixed-effects and related estimators for correlated random-coefficient and treatment-effect panel data models. *Rev. Econom. Statist.* 87(2):385–390.

Wooldridge J (2010) *Econometric Analysis of Cross Section and Panel Data (Chapter 21) Estimating Average Treatment Effects* (The MIT Press, Cambridge, MA), 937–940.

Wooldridge J (2019) Correlated random effects models with unbalanced panels. *J. Econometrics* 211(1):137–150.

Wooldridge J (2021) Two-way fixed effects, the two-way Mundlak regression, and difference-in-differences estimators. Preprint, submitted August 18, <https://dx.doi.org/10.2139/ssrn.3906345>.