# Customer Supercharging in Experience-Centric Channels
# Online Appendix.

David R. Bell
Idea Farm Ventures, New York, NY 10010, david@ideafarmventures.com

Santiago Gallino
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104, sgallino@wharton.upenn.edu

Antonio Moreno
Harvard Business School, Boston, MA 02163, amoreno@hbs.edu

*History*: May 31, 2019.

## 1. Introduction

This document provides further details on the data, analysis, and findings given in the main paper. The paper is entirely self-contained and written for ease of exposition; hence, this Appendix is a repository of key supporting description and analysis that is essential to the overall robustness of the research, but not essential to the main narrative and insights. In particular, here we provide considerably more detail on the robustness of our econometric approach and the alternative analyses undertaken to confirm that our findings are valid.

We begin with some additional details on the data (Section 2). We continue with additional details of the analyses we conducted to assess the robustness of our results to the matching approach (Section 3). Section 4 presents additional results that are not included in the paper for ease of exposition. Finally, Section 5 provides a cost-benefit analysis to inform the decision on whether a retailer should operate conventional stores, ZIS, or maintain purely online operations.

## 2. Data Description and Summary Statistics

Our data were provided by a digitally native male apparel brand (the "company"). The data themselves cover all transactions (whether finalized at the website or at the firm's offline format) completed over a period of almost 10 years. More details on the firm and its products, the offline format, and summary properties of the data are provided below. Our individual order-level data include: a unique customer identification code, the description of each item (e.g., category and sub-category), and the price paid for each item in the order. In addition, we know returns information at the individual item level and whether the order was placed in connection to a ZIS visit or not. For "ZIS transactions" we know the specific ZIS location, and for each order, we know the ZIP code to which it was shipped.

We observe the data from the inception of the company in October 2007 through October 2016, i.e., there is no left-censoring and we can also determine the exact sequence of orders for each customer. Similarly, right censoring is not a concern either, e.g., for product returns, the management shared an additional three months of data through January 2017 (a three-month interval is used by management to classify whether an order "converted" into a sale). In total, we have 2,016,554 transactions corresponding to 596,740 unique customers (Table 1 provides a summary).

**Table 1 Data Description**

| Variable | Description | Units Count |
|---|---|---:|
| Items | Total Number of Items | 4,541,486 |
| Orders | Total number of Orders | 2,016,554 |
| Web Orders | Total number of Orders placed online | 1,751,612 |
| ZIS Orders | Total number of Orders in connection to a ZIS visit | 264,942 |
| Customers | Number of Unique Customers | 596,740 |
| ZIS Customers | Number of Unique Customers that visit a ZIS | 138,374 |
| Categories | Total Number of Categories | 19 |
| Subcategories | Total Number of Subcategories | 138 |
| Date Range | Range between First and Last Day of Observation | 3,276 |

Around 23 percent (138,374 / 596,740) of all customers visit a ZIS. Even though only 13 percent of all orders are placed in connection to a ZIS visit, the data suggest that a significant fraction of the remaining 87 percent of orders, while placed online, were likely to have been heavily influenced (in terms of value and velocity) by ZIS visits. Specifically, that a ZIS visit serves to "supercharge" the relationship between a firm and its customers (more details shortly).

To construct variables for use in estimation, we aggregate the data at the order level (across potentially more than one SKU in the basket). Table 2 presents order-level summary statistics deconstructed by dollar value, size, and so on, at the order and customer level. "Web Orders" are those placed directly online and "ZIS Orders" were placed in connection to a ZIS visit. "Web Customers" are those who during the observation period never visit a ZIS and "ZIS Customers" are those who visited a ZIS at least once.

ZIS Orders are larger (in dollar terms, number of items, and number of categories) than Web Orders. The (conditional) average return value for ZIS Orders is also higher in dollar terms, but

**Table 2 Main Summary Statistics**

| Variable | Description | Mean | Std. Dev. | Percentile 5th | 50th | 95th | N |
|---|---|---|---|---|---|---|---|
| $ Order | Order size in dollars | 191.0 | 292.4 | 40.6 | 128.0 | 540.6 | 2,016,554 |
| $ Order Web | Order size in dollars for Web Orders | 177.1 | 229.0 | 38.4 | 121.0 | 484.0 | 1,751,612 |
| $ Order ZIS | Order size in dollars for ZIS Orders[1] | 282.9 | 542.6 | 59.5 | 188.0 | 816.0 | 264,942 |
| $R Order | Return dollars per Order | 52.4 | 242.6 | 0.0 | 0.0 | 267.2 | 2,016,554 |
| $R Order Web | Return dollars per Order for Web Orders | 50.7 | 189.5 | 0.0 | 0.0 | 259.2 | 1,751,612 |
| $R Order ZIS | Return dollars per Order for ZIS Orders | 63.3 | 478.5 | 0.0 | 0.0 | 329.0 | 264,942 |
| # Order | Number of Items per Order | 2.3 | 3.2 | 1.0 | 2.0 | 6.0 | 2,016,554 |
| # Order Web | Number of Items per Order for Web Orders | 2.2 | 2.0 | 1.0 | 2.0 | 5.0 | 1,751,612 |
| # Order ZIS | Number of Items per Order for ZIS Orders | 2.9 | 7.2 | 1.0 | 2.0 | 8.0 | 264,942 |
| #Cat Cust. | Number of Unique Categories bought per Cust. | 1.5 | 0.8 | 1.0 | 1.0 | 3.0 | 596,740 |
| #Cat Web Cust. | Number of Unique Categories bought per Web Cust.[2] | 1.4 | 0.7 | 1.0 | 1.0 | 3.0 | 458,366 |
| #Cat ZIS Cust. | Number of Unique Categories bought per ZIS Cust.[3] | 1.7 | 1.0 | 1.0 | 1.0 | 4.0 | 138,374 |
| Dist. ZIS | Min. Distance (miles) from Customer to a ZIS | 80.9 | 184.7 | 0.0 | 10.3 | 329.3 | 596,740 |
| Dist. ZIS Web Cust. | Min. Distance (miles) from Web Cust. to a ZIS | 96.3 | 197.5 | 0.5 | 16.8 | 352.3 | 458,366 |
| Dist. ZIS ZIS Cust. | Min. Distance (miles) from ZIS. Cust. to a ZIS | 30.0 | 120.6 | 0.0 | 2.7 | 205.4 | 138,374 |

1. ZIS Orders are those orders place in connection to a ZIS visit.
2. Web Customers are those customers that never visited a ZIS.
3. ZIS Customers are those customers that visited a ZIS at least once.

considerably lower as a percentage of the average basket value (22.1% versus 28.6%). The average distance for a random customer to reach a ZIS is about 81 miles, whereas customers who shop at least once in a ZIS are only about 30 miles away from one. Conversely, customers who only ever shop on the website are even further away from a ZIS (96 miles) than a random customer is.

## 3. Robustness

In considering the impact of model implementation decisions, we carried out a number of different approaches to construct the matches, and our results are qualitatively robust to using different specifications in the matching methodology, such as the matching method, the specific covariates we use, the number of matches per treated unit, the distance metric, and caliper (the maximum distance at which two observations are a potential match). Our findings were qualitatively unchanged and also remained so under an alternative approach that retains the maximum number of pairs subject to a balance constraint (i.e., cardinality matching, as described in Zubizarreta et al. 2014). In the main results we report in Section 3 of the paper (which are conservative), we consider a nearest neighbor matching using the Mahalanobis distance of the five covariates described above, with 1:1 matching between treated and control and with a caliper of 0.3. Further detail on additional variations tested are available from the authors upon request. For more details on the definition of the caliper and the implementation of the Mahalanobis distance metric, we refer the interested reader to chapter 8 of Rosenbaum (2010).

The main concern for the validity of our results is the potential existence of significant self-selection. We assess the robustness of the overall effect on sales just described in Section 3 of the paper by following several complementary directions. First, we assess the sensitivity of our finding to the potential presence of unobservables that may affect the decision to visit a ZIS and acquire offline experience with the brand. Second, we assess the balance in covariates that were not used in the matching, which would be likely imbalanced if there was significant self-selection. Third, we exploit individual-level variation in the intensity of the treatment that different customers experience. Finally, we complement our analysis with a geographic-level approach that is not subject to potential self-selection at the individual level.

### 3.1. Rosenbaum's sensitivity analysis

Following Rosenbaum's sensitivity analysis (Rosenbaum 2010) we explored how large the unmeasured covariate would have to be in order to affect the conclusions of the study. To understand the mechanics of this analysis, suppose that two units (treated and untreated customers in our case) have the same covariates ($\mathbf{x}_j = \mathbf{x}_k$) but have a different probability of treatment ($\pi_j \neq \pi_k$). The odds that each unit receives the treatment are $\pi_j/(1-\pi_j)$ and $\pi_k/(1-\pi_k)$, respectively, and the odds ratio is

$$\frac{\pi_j/(1-\pi_j)}{\pi_k/(1-\pi_k)}$$

The sensitivity analysis assumes that this ratio is at most some number $\Gamma \geq 1$. When $\Gamma = 1$, that means that for equal observed covariates the odds of receiving the treatment are the same, i.e., there is no hidden bias. As $\Gamma$ increases, the impact of the unobserved covariates on the odds of receiving treatment becomes larger (i.e., there is more hidden bias). The idea of the sensitivity analysis is to evaluate how the estimated effect evolves as $\Gamma$ increases. In particular, we want to understand at what point, i.e., for what value of $\Gamma$, the inferences regarding the sign of the coefficient of interest would change. If this happens for values that are close to $\Gamma = 1$, the results are very sensitive to the potential presence of hidden bias. Table 3 reports the confidence intervals and Hodges-Lehmann point estimates at different values of $\Gamma$. We have to increase $\Gamma$ to at least 2.25 in order for our estimated to lose significance.

In other words, for our results to change, the hidden bias should be strong enough that for two customers who are identical in their observable covariates, the odds of receiving treatment are at least twice as large for one of the units. Thus, we can safely conclude that our results are not sensitive to the potential existence of unobserved characteristics that can potentially generate hidden bias.

### 3.2. Balance in additional socio-demographic variables

One of the main underlying assumptions of matching techniques is that by matching in observable factors we are able to eliminate, or mitigate, differences in other variables not observed in the data. There are some important covariates we are not able to observe for all customers (and therefore are unable use in the matching for the entire sample). To assess whether those covariates become balanced, we obtained additional socio-demographic variables on a subsample of customers. To do this we considered the 255,792 customers that form our matched sample, and focused on those

**Table 3 Sensitivity Analysis**

| Gamma | sig+ | sig- | t-hat+ | t-hat- | CI+ | CI- |
|---|---|---|---|---|---|---|
| 1.00 | 0.00 | 0.00 | 202.60 | 202.60 | 198.99 | 206.20 |
| 1.25 | 0.00 | 0.00 | 141.60 | 270.00 | 138.30 | 274.05 |
| 1.50 | 0.00 | 0.00 | 95.85 | 330.80 | 92.70 | 335.35 |
| 1.75 | 0.00 | 0.00 | 59.45 | 386.84 | 56.40 | 391.80 |
| 2.00 | 0.00 | 0.00 | 29.15 | 438.90 | 26.10 | 444.30 |
| 2.25 | 0.04 | 0.00 | 2.70 | 487.70 | -0.39 | 493.60 |
| 2.50 | 1.00 | 0.00 | -21.35 | 533.80 | -24.60 | 540.1 |

Gamma: log odds of differential assignment
due to unobserved factors
sig+: upper bound significance level
sig–: lower bound significance level
t-hat+: upper bound Hodges-Lehmann point estimate
t-hat–: lower bound Hodges-Lehmann point estimate
CI+: upper bound confidence interval (a= 0.95)
CI–: lower bound confidence interval (a= 0.95)

customers for which we obtained the additional covariates. We were able to obtain at least one additional covariate for 26,082 customers of the 255,792 customers included in the matching analysis. The table below shows the total number of customers for which we were able to obtain each of the additional covariates. In addition, the table shows the number of these customers that correspond to the treatment and the control group. If the assumption that the matching approach balance these groups on unobserved factors is valid, we should observe that the covariates for these two groups are not statistically different from each other.

**Table 4 Covariate Balance in Additional Socio-demographic variables**

| | Observations | | | Mean | |
|---|---|---|---|---|---|
| Covariates | Total | Treatment | Control | Treatment | Control |
| Age | 26,082 | 12,855 | 13,227 | 39.81 | 39.64 |
| Household Income | 25,576 | 12,643 | 12,933 | 116.21 | 115.33 |
| Education | 24,681 | 12,160 | 12,521 | - | - |
| Parental Status | 24,409 | 12,067 | 12,342 | 0.90 | 0.90 |
| Home Owner Status | 24,026 | 11,857 | 12,169 | 0.63 | 0.64 |
| Marital Status | 23,802 | 11,746 | 12,056 | 0.34 | 0.35 |
| Profession | 9,076 | 4,339 | 4,737 | - | - |

As illustrated in the last two columns of the table, the balance of these covariates is remarkably good. Thus, this statistical analysis shows that our assumption of balance between treated and control customers is justified. The comparison for these covariates indicates that there is not any statistically significant difference at the 5% level, for any of the factors. To further explore the balance of the covariates and their distributions, we generate histograms for each covariate and their different levels. The histograms resulting from these analyses are presented below.

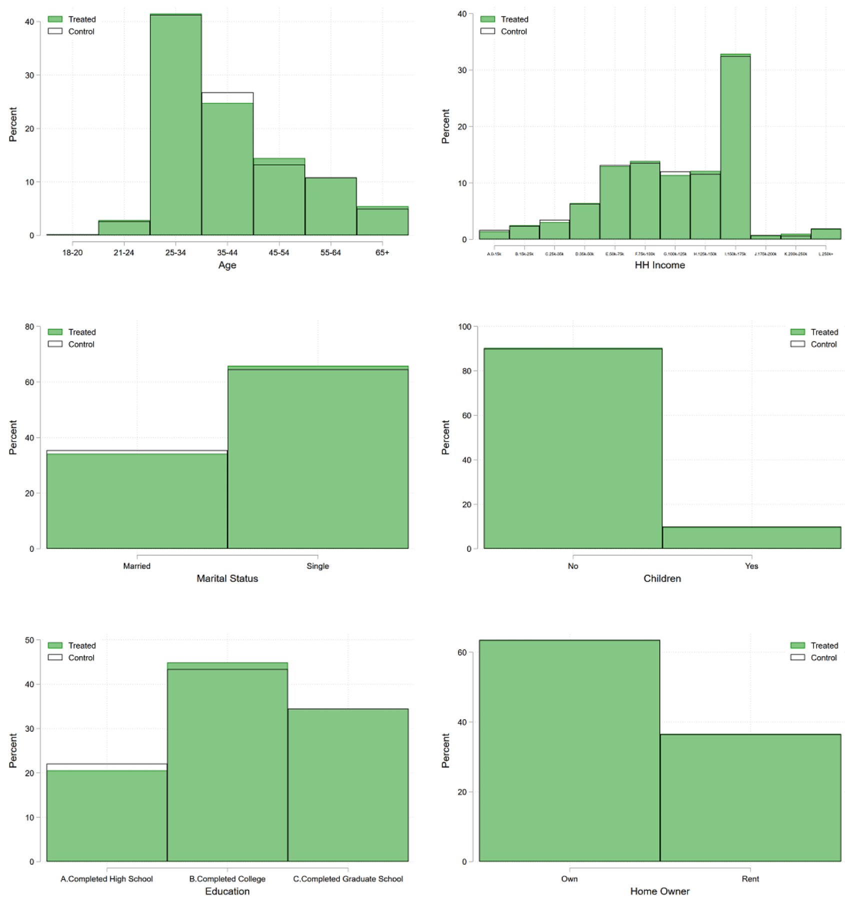


**Figure 1 Covariate Balance in Additional Socio-demographic variables**

The histograms show that not only the means, but also the distributions, are very well balanced. In general, we do not observe statistical differences across groups for the different levels. These are covariates with multiple levels and most of the levels for these covariates are well balanced.

This analysis and the data obtained allowed us to implement a careful validation of the balance on relevant covariates and the outcome of this analysis gives us reassurance on our results. We expect that if self-selection were an issue in our original matching i.e., if the treated and control individuals differed substantially in unobservables, it would result in numerous significant differences in these demographic covariates. As this is not the case, we are therefore reassured that the extent of self-selection, if present at all, does not seem to be significant.

### 3.3. Intensity of Treatment and the Role of ZIS Congestion

We explore the variation in the intensity of the treatment that different customers experience. Our underlying hypothesis is that an initial interaction at the ZIS "supercharges" a customer by providing an intimate experience and exposure to the offering of the brand. Our treated customers differ in the conditions in which they visit the ZIS, and therefore in the potential intensity of the "supercharging". Some customers visit the ZIS for the first time in a very uncongested day, when their first experience can be intimate, and some visit the ZIS for the first time in a very congested day. We hypothesize that customers who encounter a more congested ZIS when they visit it for the first time will experience lower supercharging than customers who visit the ZIS for the first time in a less congested period. We construct a daily metric of ZIS congestion by computing the number of customers served in the ZIS in that day.

We find that as the congestion in the customer's first visit increases, the effects on sales decrease, consistent with our hypothesis about the intensity of the treatment.

**Table 5** **Sales of Treated vs Control Customers After First ZIS Visit**

| Variable | Sales |
|---|---|
| $ZIS\,Treatment$ | 493.03*** |
| | (5.06) |
| $ZIS\,Treatment \times Congestion$ | -15.37*** |
| | (0.30) |
| Constant | 496.42*** |
| | (3.20) |
| Observations | 255,792 |

Robust Standard errors in parentheses
*** $p < 0.001$

### 3.4. Geographic-level Analysis

In addition to the aforementioned customer-level analysis, we conduct a complementary analysis of the effects of the ZIS channel at the geographic level. The new analysis uses a different unit of observation (zipcodes instead of customers), and takes advantage of the fact that ZIS opened in different locations at different times.

We use a propensity score approach at the zipcode level. For each zipcode in the US, we collect extensive demographic information and market information related to the fashion and apparel category. This information comes from ESRI.[1] We model the probability of a zipcode having access to a ZIS using a logit model with these variables. We then use the model to calculate the probability of each zipcode having access to a ZIS (i.e., its propensity score) and we match zipcodes that actually have access to a ZIS at some point with those that have a very similar propensity score but never have access to a ZIS. The analysis we report here uses 14 socio-demographic variables at the zipcode level to calculate the propensity score, but our results are consistent if we expand or reduce the

[1] See http://downloads.esri.com/esri_content_doc/dbl/us/2018_Updated_Demographics_Variable_List.pdf and http://downloads.esri.com/esri_content_doc/dbl/us/2018_Consumer_Spending_Variable_List.pdf

number of variables considered. Our analysis retains 1755 treated zipcodes (which had access to a ZIS at some point) and 1755 control zipcodes (which never had access to a ZIS).

During our period of observations, some of the zipcodes enter the catchment area of a ZIS (20 miles in our default model, but we conduct sensitivity analysis to different distances without qualitative changes in our results) due to ZIS openings. There is variation in the times when different ZIS opened. We conduct a difference-in-differences analysis using zipcode-week as our unit of observation. Our coefficient of interest is the interaction between an indicator of "treated" zipcode with and a variable that indicates whether treatment is active (which happens when the zipcode is in the catchment area of a ZIS). This allows us to use the variation in time of opening of the different ZIS for the identification of the effects of the ZIS. We consider sales at the zipcode and week as the main dependent variable of interest. As shown in Table 6, the effect of a ZIS opening on weekly sales in zipcodes in the catchment area is positive and significant, consistent with the patterns we document in our customer-level analysis. Table 6 also explores the effects of the opening of a ZIS on returns. Returns are two percentage points lower when ZIP codes enter the catchment area of a ZIS. Even if we focus our attention to transactions that do not originate in the ZIS (Column 3), returns are 1.4 percentage points lower when a zipcode enters the catchment area of a ZIS.

Overall, these patterns confirm that the findings obtained with our customer-level analysis are not an artifact of customer self-selection into visiting the ZIS.

**Table 6 Regression Results**

| | *Dependent variable:* | | |
|---|---|---|---|
| | Sales | Return (%) | Return online (%) |
| | (1) | (2) | (3) |
| ZIS_Open | $281.441^{***}$ | $-0.020^{***}$ | $-0.014^{***}$ |
| | (7.788) | (0.002) | (0.002) |
| Zip code fixed effects | YES | YES | YES |
| Week-year fixed effects | YES | YES | YES |
| Observations | 385,175 | 385,175 | 367,130 |
| $R^2$ | 0.589 | 0.040 | 0.042 |
| *Note:* | | | $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 |

### 3.5. Summary

The amplification effect, aka "supercharging" created by the ZIS is observed in different yet complementary analyses that are based on different sets of assumptions and levels of aggregation. Beginning with model-free evidence in the raw data, we are then able to detect strong effects in individual-level panel data analysis. Our focal matching-based analysis reproduces the key result, and the matching itself is shown to be robust under the Rosenbaum test and cross-validation using a rich set of demographic variables. Moreover, the supercharging mechanism responds as one would expect when store congestion is used as a proxy for the quality of the service experience. Finally, ZIP code-level analysis yields qualitatively identical findings, providing us with strong assurance that the overall effect and the approach used to identify it are sound.

## 4. Additional Results

Section 4 in the paper covers three critical dimensions of consumer shopping behavior that are impacted by experience in a ZIS: shopping velocity, basket composition (in terms of prices paid, items, and categories), and product returns. In addition, we show that the effect of the ZIS is greater when customers are in close proximity to it, and when their extent of prior brand experience is more limited.

One additional aspect of customer behavior not addressed in the paper but which is potentially of interest to the retailer, is the extent to which customers avail themselves of discounts. Using the standard specification (equation (1) in the paper), we model the ratio of dollars of discounts to dollars of sales after a customer has has a ZIS experience. Table 7 shows that post the ZIS experience, customers use 4% fewer discounts than before, which, all else equal, is positive for the brand.

**Table 7 Discount as a Percentage of Sales**

| | Sales |
|---|---|
| $ZIS\,Treatment$ | -0.02*** |
| | (0.00) |
| Constant | 0.16*** |
| | (0.00) |
| Observations | 255,792 |

Robust Standard errors in parentheses
*** $p < 0.001$

## 5. Cost-Benefit Analysis: ZIS plus E-Commerce Versus Brick and Mortar Stores

With the data at hand, we cannot offer a direct empirical comparison of a traditional brick and mortar (B&M) store that carries inventory, versus a ZIS; nevertheless, we are able to coalesce some key distinctions, that, in a digitally-enabled world, lead to important advantages to the ZIS format.

Specifically, the ZIS alleviates the logistics of inventory fulfillment, reduces the assortment burdens of retailing, and mitigates stock-outs, all while requiring a much smaller physical footprint, so there are important aspects to like. Moreover, employees, working in a smaller and more intimate space and freed from the burdens of managing inventory and assortment, can be more focused on value-add personal selling and service. On the downside, the ZIS does not and cannot cater to the need for immediate fulfillment sought by certain types of customers or on certain shopping occasions.

In order to offer a more quantitatively grounded comparison of the two formats, we developed a simulation. In the simulation we compare the profitability of ZIS versus B&M stores, under different coefficients of variation of demand, and demand correlation across offline locations. Our base case, consistent with what is observed in practice, assumes that the retailer operates with a 70% gross margin and an in-stock target of 95% for traditional B&M stores. In addition, we assume that the retailer has 300 stores and the demand correlation across stores is 0.5 and the coefficient of variation of the demand is 0.2.

The simulation works by focusing on B&M stores as a base case, and then computing the range of parameter values over which a ZIS would be at least as profitable as the store. A critical ingredient to the simulation is a valid estimate of the percentage of store sales that a ZIS would reasonably retain; as noted above, for customers and shopping occasions for which immediate fulfillment was required, sales would be "lost" by the ZIS. In the spirit of Decision Calculus (Little, 1970), we elicited plausible values from experienced senior retail managers. During retail industry workshop with more than one hundred executives present, we proposed the following scenarios:

**Scenario 1:** What would be the effect on sales if J. Crew converted stores into showrooms (with no inventory) and offered Free Shipping with delivery in 4 business days?

**Scenario 2:** What would be the effect on sales if J. Crew converted stores into showrooms (with no inventory) and offered Free Shipping with next day delivery?

As might be expected, more sales should be retained in Scenario 2, where the ZIS is better able to deliver something closer to "immediate fulfillment". Responses from the executives, who represent a wide cross-section of the retail sector, are summarized below.

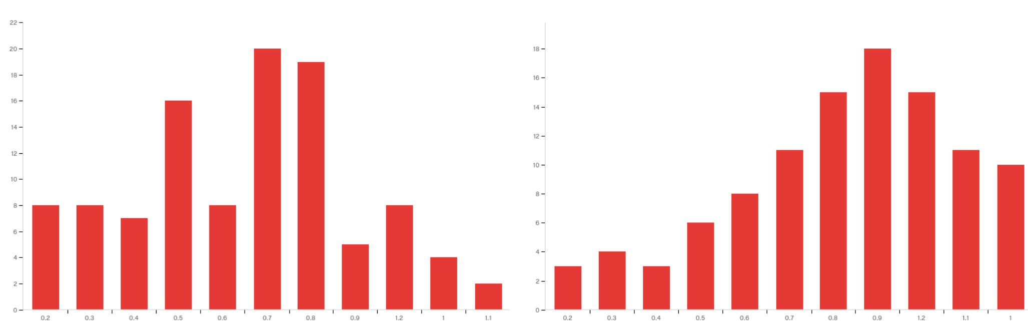


(a) Scenario 1: 105 Responses

(b) Scenario 2: 104 Responses

**Figure 2** **Managers' Survey Results**

The data shows a wide range of estimates ranging from the ZIS retaining 20 percent of the store sales to the ZIS capturing 10 percent more than the sales of a B&M store. For Scenario 1, the mean and standard deviation are 0.66 and 0.27, respectively, meaning that a when a B&M store is converted to a ZIS, it will retain about 66% of the original sales. Naturally, for Scenario 2 the mean expected percentage of sales retained by the ZIS increased to 0.83 (the standard deviation remained as 0.27). For our sensitivity analysis we established our base case scenario by considering that the ZIS would capture 83 percent of the sales of B&M stores, the average estimate of the managers under next day delivery. This survey provides a grounded approach to the simulation exercise we implemented and is consistent with the "decision calculus" view of simulations first introduced to the management science literature by Little (1970).

To complete the simulation also need estimates for operational efficiency parameters. Here we assume that the return rate is 15% for both formats, and the cost of handling returns is 10% of the sales value for the return items. The cost of the store operation in our base case is assumed to be higher for the B&M format (15% of sales) and lower for the ZIS (10% of sales), while customer e-commerce shipping costs under the ZIS format are 5% of sales. With the 70% profit margin for the traditional B&M store as a baseline and a ZIS profit margin that is 7.7% higher, we evaluate what happens when we change the coefficient of variation of demand and the correlation across locations (these are the two factors that affect the inventory levels).

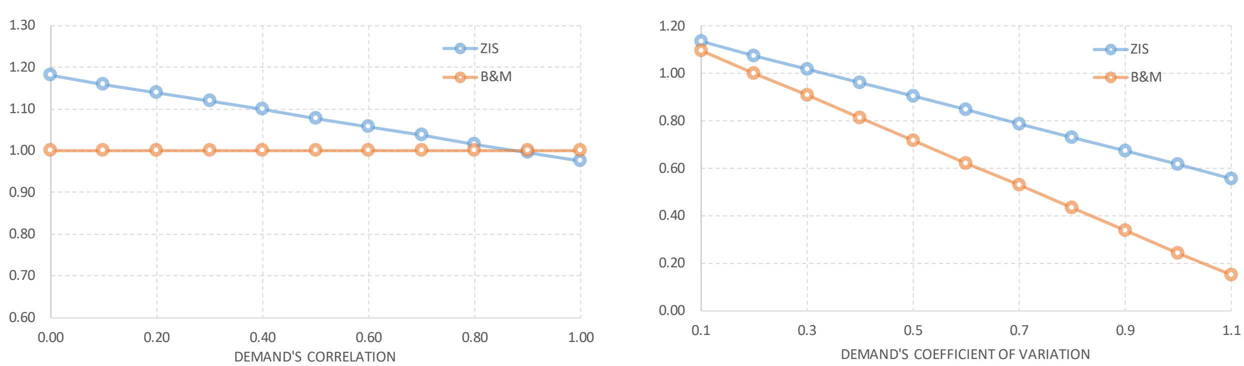


(a) Demand Correlation (b) Coefficient of Variation Sensitivity

**Figure 3 Cost Benefit: Sensitivity Analysis**

The sensitivity analysis reveals that the ZIS dominates B&M for most demand correlation scenarios. The benefits coming from the economies of scale of the showrooms, which pool demand uncertainty, are lost as the correlation of demand across locations gets closer to one. Increasing the coefficient of variation of demand also has a negative impact on profit margin for both the ZIS and the B&M format; it is, however, very interesting to see that the ZIS is less impacted.

Finally, we compare the profit (in absolute terms) that the two different formats generate. To do this we implemented a sensitivity analysis where we vary the coefficient of variation for the demand under the two formats and compute the lost revenue that the ZIS can suffer, compared to the traditional B&M format, to obtain the same level of profit.

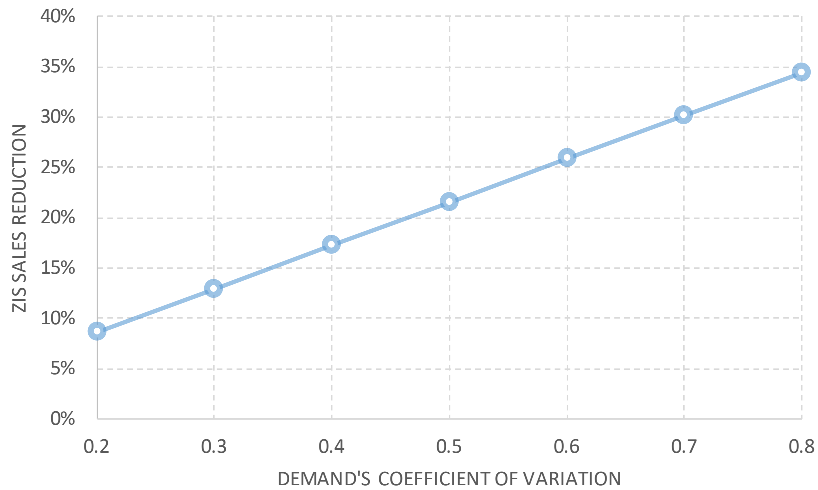


**Figure 4 Cost Benefit: Additional Demand for B&M Format**

The figure clearly shows that as the coefficient of variation of demand ($x$-axis) increases, the ZIS format can lose an increasing percentage of sales ($y$-axis), relative to the B&M format, and still match its profitability. Again, to get a sense of whether this relationship is practically meaningful for the apparel category and under our assumptions, we returned to the responses in our retail executive survey. First, the coefficient of variation common in the fashion apparel industry is around 0.5 (or higher). At this level of coefficient of variation, the ZIS could give up 22% of demand, relative to the B&M format, and yet still be equally profitable.

The executives in our sample estimate that with one-day shipping, the ZIS would, on average, lose only 17% of demand relative to the B&M (alternatively, 83% of demand would be retained in this instance). Hence, at the appropriate coefficient of variation in demand, the amount of demand that can be given up for the ZIS to match profitability with the B&M format is *larger* than the average and median loss estimated by the managers. Thus, our simulation, conducted with data inputs from real managers and under plausible data inputs, implies that the ZIS format holds considerable promise and that a switch from B&M stores to experience-centric ZIS formats is worth exploring.

## References

Little JD (1970) Models and managers: The concept of a decision calculus. *Management Science* 16(8):B–466.

Rosenbaum PR (2010) Design of observational studies. *Design of Observational Studies* (Springer).

Zubizarreta JR, Paredes RD, Rosenbaum PR (2014) Matching for balance, pairing for heterogeneity in an observational study of the effectiveness of for-profit and not-for-profit high schools in chile. *The Annals of Applied Statistics* 204–231.