The Surprising Breadth of Harbingers of Failure

The Surprising Breadth of Harbingers of Failure

Duncan Simester, Catherine Tucker, Clair Yang

MIT

June 2019

Previous research has shown that there exist "harbinger customers," who systematically purchase new products that fail and that are discontinued by retailers. We extend this finding in two ways. First, we show that there are also some zip codes that are "harbinger zip codes." If households in these zip codes adopt a new product, this is a signal that the new product will fail. Second, we show that households in these zip codes make choices that are systematically different from other households across a surprising array of decisions.

We identify harbinger zip codes using purchases of new products from one retailer. We then analyze purchases from these zip codes at a different retailer. We show that households in the harbinger zip codes identified at the first retailer purchase products from the second retailer that households in other zip codes are less likely to purchase. We then investigate whether harbinger zip codes are predictive outside the retail domain. Households in harbinger zip codes donate to different congressional election candidates than households in neighboring zip codes, and they donate to candidates that are less likely to win. We also compare changes in house prices between 2002 and 2015. House prices in harbinger zip codes increased at slower rates than other zip codes.

By studying households that change zip codes, we are able to investigate whether harbinger zip codes result from households with harbinger "tendencies" choosing to cluster together, or households learning these tendencies by observing their neighbors. The evidence strongly suggests that harbinger tendencies are a sticky trait, and the harbinger zip code effect is more due to where customers choose to live, rather than households influencing the tendencies of their neighbors.

Keywords: preference heterogeneity, new product development, real estate prices, campaign donations

1. Introduction

Presumably, someone liked Diet Crystal Pepsi. Unfortunately for Pepsi, there were not enough of these people, so the product flopped. What is surprising is that the customers who purchased Diet Crystal Pepsi may also have liked other new products that flopped, like Colgate Kitchen Entrees (a range of frozen meals). Anderson, Lin, Simester and Tucker (2015) documented the existence of these customers, whom they labelled "harbingers of failure." Harbingers are more likely to purchase products that other customers do not buy.

We extend this finding in two ways. First, we show that there are not just customers who are harbingers; there are also harbinger zip codes. Purchases of new products by households in these zip codes is a signal that the new product will fail. More precisely, holding the number of purchases by non-harbingers fixed, the number of purchases by harbinger zip codes is higher on products that fail than on products that succeed.

Second, we show that harbingers make decisions that are different from neighboring zip codes across a wide range of decision contexts. In particular, we first identify harbinger zip codes using new product purchases from a mass merchandise retailer. We show that households in these zip codes purchase new products that tend to fail, and also purchase niche existing products that neighboring zip codes are less likely to purchase. The same pattern also extends to purchasing decisions at another retailer selling private label apparel. Zip codes identified as harbingers at the mass merchandise store are also harbinger zip codes at the apparel retailer.

We then ask whether our results extend beyond the retail domain, starting with a comparison of contributions to congressional election campaigns. We identify the top two candidates and compare donations in each 5-digit zip code with donations in neighboring zip codes. Zip codes identified as harbingers at the mass merchandise store are systematically less likely to donate to the most popular candidates. They are also more likely to donate to candidates that lose their elections. We also compare changes in house prices between 2002 and 2015. Zip codes identified as harbingers at the mass merchandise store had systematically smaller increases in their house prices during periods that prices were increasing for neighboring zip codes.

Data from the apparel retailer also allows us to observe households moving between zip codes. We use these events to ask whether the harbinger zip code effect results from households with harbinger tendencies clustering together, or whether households learn these tendencies by observing their neighbors. We first ask whether households who move from harbinger zip codes tend to move to other harbinger zip codes. The evidence strongly supports this pattern. Households who leave harbinger zip codes tend to go to other harbinger zip codes. The reverse is true for households that start in non-harbinger zip codes; they tend to move to other non-harbinger zip codes. We then ask whether households' tendencies change when they move, which might suggest that households' tendencies are in part learned from their new neighbors. We find no support for this; there is little evidence that harbinger households' tendencies become more like their neighbors' after they move. We conclude that harbinger tendencies are a sticky trait, and the clustering of harbingers within harbinger zip codes appears to be caused more by the tendencies households bring than the tendencies that they learn.

Collectively, the results reveal that harbinger tendencies are correlated across a wide range of decisions, including purchasing decisions, political decisions, and housing decisions. We do not know what characteristics of customers' decision-making contribute to these tendencies. For purchasing decisions, the harbinger effect could result from correlations in product preferences or responsiveness to marketing variables. We also recognize that the effect could reflect customer tendencies that are unrelated to product features. For example, harbingers may have a preference for variety (variety-seekers), they may have more (or less) traditional values, they may be contrarians, or they may have a greater (or lesser) willingness to bear risk. The evidence that the effect extends beyond purchasing decisions increases the likelihood that the harbinger effect reflects these more generic customer tendencies.

In this paper we do not seek to evaluate why the harbinger effect occurs. Instead, the focus of the paper is empirical. We generalize the harbinger effect from an individual level effect to a zip code level effect – opening up the use of the harbingers phenomenon to retailers who do not collect individual-level data. We also show that the findings extend from a single retailer, to other retailers and other (non-purchasing) decision contexts.

Implications

The findings in this paper provide a conceptual foundation for collaborative filtering. Collaborative filtering is widely used in product recommender systems. It relies on the assumption that customers who agreed in the past will agree in the future, and their similarities will extend across different decision contexts. The findings in this paper provide direct support for this assumption.[1]

The findings also have implications for managers engaged in new product development. For a firm introducing a new product in a new product category, the findings indicate that the firm can more accurately predict product success by asking: which customers are early adopters of the new product? If the new product is purchased by a high proportion of customers who are harbingers in the firm's existing categories, then the new product is less likely to succeed. We show that this is true, even if the new product is in a different category (including perhaps a very different category).

The evidence that the findings in one category extend to other categories is also relevant for market research firms that support product development. A market research firm can identify harbinger zip codes in one product category and use this to improve new product forecasting for a different firm in a different product category (or even different industry). Identifying the harbinger effect at the zip code level facilitates the portability of this knowledge. Firms do not need to match specific customers across different firms or markets. Instead, our findings show that it is sufficient to match zip codes. This alleviates privacy concerns that may arise when using individual data (and may potentially help to reduce the risk of data breaches because the analysis does not rely on personally identifiable data).

The Harbinger Effect

Anderson et al. (2015) documented the existence of customers who systematically purchase new products that fail. Their findings were based upon purchases of new products at a chain of convenience stores that sells consumer packaged goods. They divided the new products into

[1] The findings also provide a conceptual foundation for the application of transfer learning to marketing problems. Transfer learning provides a mechanism for transporting knowledge from one decision context to another. This is only valuable if the information in the first decision context is relevant to the second decision context. The findings in this paper provide evidence that customer tendencies in one decision context are relevant in other decision contexts.

two subsets and used the outcomes in the first subset to classify customers as harbingers (or not). They then showed that customers who purchased flops in the first subset of products (the harbingers) were also systematically more likely to purchase flops in the second subset of products. Adoption of a new product by these harbingers was a strong signal that a new product would fail. They argue that in the setting they study, it is unlikely that the finding is due to observational learning or information spillovers between customers. Instead, they attribute the finding to harbingers having unusual tendencies that are not representative of other customers. Adoption by the harbingers is a signal that other customers are less likely to adopt the product, which leads to product failure.

The original Anderson et al. (2015) study has since been replicated and extended in two recent working papers. Anderson, Chen and McShane (2018) replicate the findings using an IRI panel dataset including over 100,000 consumers and 400 retailers. Anderson, Chen, Liu and Simester (2017) demonstrate the existence of "harbinger products." While purchases of harbinger customers signal that a new product will fail, harbinger products represent the inverse. Purchases of harbinger products on a customer's first shopping visit signals that the new customer will fail (not return).

Perhaps the most surprising aspect of the harbinger customer effect is that the signal extends across consumer packaged goods (CPG) categories. Customers who purchase new oral care products that flop also tend to purchase new haircare products that flop. Anderson et al. (2015) interpret their findings as evidence that customers who have unusual preferences in one product category also tend to have unusual preferences in other categories. In other words, the customers who liked Diet Crystal Pepsi also tended to like Colgate Kitchen Entrees (which also flopped). If the most surprising aspect of the original study is that the harbinger effect extends across product categories, then the results in this paper magnify that surprise. Anderson et al. (2015) show that the effect extends across CPG categories within a single retailer. This paper show that the effect extends across different retailers, beyond consumer packaged goods, and beyond purchasing decisions. Customers who purchase products that fail at a mass merchandise retailer, also purchase less popular items at an apparel retailer, support less popular congressional election candidates, and live in zip codes that have smaller house

price increases. We view this as a form of purchase-based segmentation that adds incremental insight over traditional socio-demographic segmentation variables.

Other Related Literature

The harbinger effect is related to, though distinct from, the preference minority and lead user literatures. The preference minority literature studies customers with unusual preferences. These customers have been used to explain the growth of Internet sales in some product categories. If bricks and mortar stores allocate shelf space according to the preferences that are most typical in their markets, then preference minorities may not find products that fit their preferences (Anderson 1979; Waldfogel 2003). This helps to explain why preference minorities are more likely to purchase from the Internet, and are less price-sensitive when doing so (see Choi and Bell 2011 and Brynjolfsson, Hu and Rahman 2009). A related explanation has also been used to explain why we see a long tail of purchases on niche items on the Internet (see Brynjolfsson, Hu and Smith 2003; Brynjolfsson, Hu and Simester 2011).

Lead users are customers whose preferences are more likely to identify "breakthrough" ideas (von Hippel 1986). Adoption of new products by these lead users is a signal that the product is more likely to succeed. This is the opposite signal that we should infer from the adoption of new products by harbingers. The lead user literature highlights that firms need to attract the right kind of adoption for their products. This contrasts with a broader assumption in the product diffusion literature that greater adoption by any customer is an indication that a new product will succeed. However, the adoption signal associated with harbingers is reversed. It is this aspect of the harbinger effect that makes it counter-intuitive.

The remainder of the paper is organized as follows. Section 2 describes the data we use to identify harbinger zip codes together with the other datasets that we match to this. Section 3 identifies harbinger zip codes using data from a mass merchandise store. In Section 4 we use this identification from the mass merchandise store to compare purchases of apparel from a private label retailer. These findings demonstrate that the harbinger effect persists across retailers. Section 5 presents evidence that the effect also extends to changes in house prices. In Section 6 we investigate the extent to which harbinger zip codes explain variation in donations to political campaigns. In Section 7, we invert the analysis. Instead of using new product success at one retailer to explain variation in house price increases, political donations and

purchases at another retailer, we use the decisions in those other contexts to explain variation in new product success at the original retailer. The paper concludes with Section 8.

2. Data and Initial Results

This paper uses multiple datasets. We start by describing the data that we use to identify harbinger zip codes and then briefly describe the other datasets used in the paper. To identify harbinger zip codes, we use data provided by a mass merchandise store. For confidentiality reasons we cannot identify which store, but for ease of exposition we will refer to it as "MassStore." The retailer sells a broad range of products including perishables, sundries, and durables. When customers purchase at this retailer they must identify themselves with their membership card, which allows us to link each transaction to each customer.

Our analysis uses three datasets provided by this retailer. The first dataset is a complete record of individual customer transactions (for every customer) between January 2013 and July 2016. Each line item in the data identifies a purchase of a unique item by a customer on a unique purchase occasion. The data identifies the customer, store, date and time of purchase, number of units purchased and price paid. The second dataset is a complete record of coupons used in the customer transactions included in the transaction data. The coupon data identifies how many coupons were used in each transaction and the amount saved with each coupon. However, we cannot match coupons with specific items. As a result, we will use the coupons to measure customers' deal sensitivity, but we cannot use them to measure how often a product was purchased at a discount.

MassStore also provided a third dataset containing demographic data for each customer, together with the 5-digit zip code from each customer's residential address (this address is provided by customers when registering for a membership). We augment this dataset by obtaining census data that describes the population density, ethnic composition, and average education attainment levels by zip code.[2] Educational attainment measures the highest degree of education an individual has completed. We measure the proportion of the zip code whose highest attainment was graduating high school, and the proportion whose highest attainment

[2] This data was obtained from: https://factfinder.census.gov

was graduating college with a four-year bachelor's degree. We use this augmented data both to compare the observable characteristics of harbinger zip codes with other zip codes, and to control for variation in these observable characteristics in our analysis. Definitions for all of these variables, together with summary statistics, are provided in the Appendix.

The classification of zip codes using the MassStore zip code data is then used to compare purchases from a private label apparel retailer. A confidentiality agreement again prevents us from identifying which retailer provided this data. For ease of exposition we will refer to it as "ApparelCo". The retailer sells through its own dedicated retailer stores, its own catalog stores and its own Internet site. It rarely sells products that do not carry its own brand, and its products are not sold by other retailers (with very few exceptions). The data provided by ApparelCo is a complete record of all purchases made through the retailer's catalog and Internet channels. We exclude store transactions as it is not always possible to identify the residential zip code of the customer making the purchase.[3] Like MassStore's transaction data, each line item in the data identifies a purchase of a unique item by a customer on a unique purchase occasion. The data identifies the customer, ordering channel, date and time of purchase, number of units purchased and price paid. We replicate our analysis using transactions from 2010 and 2011. Because the data sample ends on 9 December 2011, the actual data periods are:

2010: 10 December 2009 through 9 December 2010

2011: 10 December 2010 through 9 December 2011.

When analyzing the ApparelCo data, we identify households using the zip code of the purchaser. In some cases this differs from the zip code of the recipient. Reassuringly, our results replicate when we use the zip code of the recipient instead of the purchaser.

Our investigation of contributions to congressional election candidates is based on a dataset constructed from two sources. The individual campaign contribution dataset was provided by a third-party watchdog organization, the Center for Responsive Politics (CRP). CRP digitizes

[3] The retailer does attempt to match store transactions with Internet and catalog customers. However, this matching process is not complete. When we include the store transactions that are matched, we obtain a very similar pattern of findings.

reports published by the Federal Election Commission and organizes them into standardized data sets.[4] The data records any individual contribution made to a congressional candidate, including variables such as the individual's name, address, zip code, donation amount, recipient candidate and date of contribution. We supplement this data with congressional election outcomes from the Clerk of the US House of Representatives. The contribution data span 2000 to 2014, while the election outcomes data extend from 2000 to 2010.

To study the change in house prices, we obtained house price data from Zillow.com.[5] This dataset reports median home values by zip code (the Zillow Home Value Index) and includes single-family residences, condominiums and co-op homes.

Throughout the paper we indicate significance (statistically significantly different from zero) using: ** $p < 0.01$, * $p < 0.05$ and † $p < 0.10$. In the next section, we use data from MassStore to investigate whether there are zip codes that systematically purchase new products that fail.

3. Harbinger Zip Codes

Recall that the transaction data provided by MassStore extends from January 2013 through July 2016. We identify new items introduced between July 2013 and June 2014.[6] We then use the period from July 2014 to July 2016 to evaluate whether these new items survive 18 months. We label a new product as a failure if its last transaction was within 18 months of the product's introduction. We also replicate the results using 12, 15 and 21 month survival windows. We exclude items that survive less than 90 days or that are explicitly introduced with a short-term purpose (as labeled by the retailer). This helps avoid short-term or seasonal items, such as St Patrick's Day products. Furthermore, we require that in these 90 days the item is sold in at least a quarter of the retailer's stores. This restriction excludes products that were being tested in a few stores and items with limited geographic appeal (such as a Boston Red Sox hat). As an

[4] Federal campaign finance laws in the US require candidate committees, party committees, and PACs to file periodic reports disclosing the money they raise and spend. Federal candidates and committees must provide the names, occupations, employers and addresses of all individuals who give them more than $200 in an election cycle. The Federal Election Commission maintains this database and publishes the information about campaigns and donors on its web site.
[5] This data was obtained from: http://files.zillowstatic.com/research/public/Zip/Zip_Zhvi_AllHomes.csv.
[6] These items had no transactions between January 2013 and July 2013.

additional robustness check, we repeat the analysis when using items sold in at least 50%, 75%, 95% and 100% of the retailer's stores.

We randomly divide the new items into "classification" items and "validation" items. There are 2,324 items in the classification category and 2,388 items in the validation category. As we would expect, the 18-month survival rates are similar for this data split. The classification set had a 64.07% survival rate, and the validation set had 65.49%. This success rate is higher than in Anderson et al (2015). In their sample of 8,809 new products, 3,508 (40%) survived for 3 years (12 quarters). The higher success rate reported in this paper may reflect the use of an 18-month window rather than a three-year window. It could also reflect the more mainstream product focus of MassStore. The new products have all survived the retailer's initial market tests and are broadly introduced across the retailer's stores. If we were able to observe the full sample of new products proposed by manufacturers or subjected to initial market tests, the success rate would be considerably lower.

We classify five-digit zip codes into groups by focusing on purchases of the 2,324 new items in the classification category in the first 90 days after each new item is introduced. We calculate the proportion of new items purchased by customers in each zip code that succeeded (had sales after 18 months). To construct this zip code level average, we weight by the number of transactions for each new item. We then rank the zip codes based on this average and organize them into groups based on quartiles of the average success rate.[7]

Group 1	over 85.9% average success rate
Group 2	84.1% to 85.9% average success rate
Group 3	82.1% to 84.1% average success rate
Group 4	under 82.1% average success rate

We then estimate very similar OLS models to those used by Anderson et al. (2015):

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \quad (1)$$

[7] These are the inter-quartile cutoffs, not the average success rate per quartile. The average success rates by quartile are: 87.71%, 84.99%, 83.14%, and 79.98%. Notice that these average success rates are higher than the average new product survival rates reported above. Recall that the average success rates are weighted by purchase frequency and, as we might expect, the products that survive are purchased more frequently.

$$\text{Success}_j = \alpha + \beta_1 \log \text{Group 1 Orders}_j + \beta_2 \log \text{Group 2 Orders}_j + \beta_3 \log \text{Group 3 Orders}_j + \beta_4 \log \text{Group 4 Orders}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \quad (2)$$

The unit of analysis in both models is a new product (j) and the estimation sample is the 2,388 items in the validation sample. *Total Orders* describes the total number of purchases of new product j calculated using all of the households. For Equation 2, *Group 1 Orders* measures the total number of purchases of new product j by households from zip codes in Group 1 (with analogous definitions for the variables corresponding to the other groups). We use a $\log_{10}$ transformation for all of these total order variables.[8] Recall that the grouping of zip codes is based upon a different sample of new products (the classification items) than the items used for estimation (the validation items).

The $\boldsymbol{\beta}\mathbf{X}_j$ term represents nineteen control variables ($\mathbf{X}_j$) and their associated coefficients ($\boldsymbol{\beta}$). As we noted in the previous section, definitions for all of these variables (together with summary statistics) are provided in the Appendix. Notice that the unit of analysis in Equations 1 and 2 is a new product, while these control variables are all identified at the zip code level. In order to include these zip code level variables in the three models, we need to construct a product level version of each control variable. To accomplish this we recognize that each new product transaction (within the first 90 days of a new product's introduction) is made by a household associated with a zip code. Therefore, for each new product, we can use these transactions to calculate a transaction weighted average of each control variable ($\bar{x}_j$):

$$\bar{x}_j = \frac{\sum_z w_{jz} x_z}{\sum_z w_{jz}}$$

In this expression, w_{jz} measures the number of purchases (within the first 90 days) of product j by customers in zip code z, and x_z represents the value of control variable x in zip code z. These transaction-weighted averages are constructed separately for each of the nineteen control variables and included in the vector $\mathbf{X}_j$.

We also estimate an alternative model:

[8] In this respect the model is different from the model estimated by Anderson et al. (2015). The volume of purchases is much larger in the MassStore data than in the dataset they analyzed.

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \beta_2 \,\%\, \text{Group } 2_j + \beta_3 \,\%\, \text{Group } 3_j + \beta_4 \,\%\, \text{Group } 4_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \quad (3)$$

The *% Group 2*$_j$ variable measures the proportion of orders contributed by households in zip codes in Group 2. The *% Group 3*$_j$ and *% Group 4*$_j$ variables are defined similarly. The proportions for the four groups sum to 100% and so we omit the *% Group 1*$_j$ variable. The coefficient for the *% Group 2*$_j$ variable can be interpreted as the change in the probability of success when the proportion of sales contributed by households in Group 2 increases (and there is a corresponding decrease in the proportion contributed by households in Group 1).

In Table 1 we report the coefficients of interest (β_1 through β_4) obtained when estimating Equations 1, 2 and 3 using the 2,388 new products in the validation set. For ease of presentation, the coefficients for the nineteen control variables are reported separately in the Appendix (for Equation 2). Column (1) confirms that new products that sell more in their first 90 days are more likely to survive for 18 months. Columns 2 and 3 reveal that the likelihood a new product will survive depends not just upon total purchases; it also depends upon which households are making those purchases. An increase in purchases by households in Groups 1 and 2 is associated with a higher likelihood of success. However, additional purchases from Groups 3 and 4 are a signal that the new product will fail. An order of magnitude (10-fold) increase in orders from Group 4 is associated with a 109.09% reduction in the probability that the new item will survive. Similarly, if the proportion of purchases from zip codes in Group 4 increases by 10% (with 10% less from Group 1), there is a 66.12% reduction in the probability that the item will succeed.

Table 1 about here

While items with more *Total Orders* in the first 90 days are more likely to succeed (the *Total Orders* coefficient is positive), it is not just the number of orders that is important. It is also important to know which households are placing those orders. If the new product is purchased by households in harbinger zip codes, then more orders signals the reverse outcome - a lower probability of success. The significance of the coefficients, together with the increase in the R^2 values (in Equations 2 and 3 compared to Equation 1), indicate that the harbinger effect clearly explains additional variation in the success of the products that cannot be explained by the observable control variables alone.

We caution that the findings do not show that harbingers buy more failed products than non-harbingers. Instead, they show that holding the number of purchases by non-harbingers fixed, the number of purchases by harbingers is higher for products that fail than for products that succeed. We can illustrate this by restricting attention to new products for which the number of orders from Group 1 customers is within half a standard deviation of the mean number of orders from Group 1 customers. These 927 products include 361 products that failed and 566 products that succeeded. The mean number of purchases from Group 2, Group 3 and Group 4 customers for the products that succeeded and failed is summarized in the Appendix. The findings confirm that holding the number of purchases by non-harbingers approximately fixed (within half a standard deviation of the mean), the number of purchases by harbingers is higher for products that fail than for products that succeed.

The findings also suggest that there is a mirror image of harbingers - zip codes that consistently purchase new products that succeed. However, the existence of these non-harbinger zip codes is less surprising than the existence of a harbinger segment. A strong strawman argument is that more purchases are more likely to lead to new product success. The harbinger effect has to overcome this effect in order for more purchases by harbingers to signal a lower probability of new product success.

Robustness Checks

The results survive an extensive range of robustness checks. In particular, we replicate the analysis using the following modifications:

- Using an alternative approach to control for income differences across zip codes.
- Including additional variables to control for the average price charged for each product, together with the number of this retailer's stores that the product is sold in.
- Using 12 month, 15 month and 21 month survival windows (instead of 18 months).
- Varying the restriction on the number of items that the stores were sold in.
- Re-estimating the three models using logistic regression.
- Constructing the *Average Success Rate* by weighting the new item purchases using the number of households that purchased each item (instead of the number of orders).
- Using decile buckets instead of quartile buckets for the *Average Success Rate*.

We relegate these robustness checks to the Appendix, where we describe the analyses and report detailed findings. The pattern of findings is robust to all of these modifications. We next investigate whether harbinger zip codes can help managers *predict* new product success.

Predicting New Product Success

To develop a predictive model, we randomly divide the 2,388 new items in the validation set into a calibration sample (1,194 new items) and a holdout sample (1,194 new items). Using the calibration sample we estimate the following equations using the complete set of control variables ($\mathbf{X}_j$).

$$\text{Success}_j = \alpha + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(Controls Only)}$$

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(Total Orders)}$$

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \beta_2 \,\%\, \text{Group 1}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(\% Group 1)}$$

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \beta_2 \,\%\, \text{Group 2}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(\% Group 2)}$$

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \beta_2 \,\%\, \text{Group 3}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(\% Group 3)}$$

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \beta_2 \,\%\, \text{Group 4}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(\% Group 4)}$$

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \beta_2 \,\%\, \text{Group 3 and 4}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(\% Groups 3 and 4)}$$

$$\text{Success}_j = \alpha + \beta_1 \log \text{Total Orders}_j + \beta_2 \,\%\, \text{Group 2}_j + \beta_3 \,\%\, \text{Group 3}_j + \beta_4 \,\%\, \text{Group 4}_j + \boldsymbol{\beta}\mathbf{X}_j + \varepsilon_j \qquad \text{(Full Model)}$$

All of these models except the *Controls Only* and the *Total Orders* models make use of the harbinger classification. We retain the coefficients from each model and use these coefficients to predict new product success in the holdout sample. The accuracy is reported in Table 2.[9]

A naïve baseline prediction is that all of the new products in the holdout sample will succeed. This baseline prediction is correct 65.91% of the time (787 new products succeed from a total of 1,194 new products in the holdout sample). Including the control variables and *log Total*

[9] If the predicted probability that a new product will succeed is larger than 50%, we interpret this as a prediction of success. Predicted probabilities smaller than 50% are interpreted as predictions of product failure.

Orders increases predictive accuracy to 71.36%. Adding the *% Group 1* measure further lifts accuracy to 73.03%, and adding the *% Group 4* measure raises accuracy to 73.87%.

If we interpret the Group 3 and Group 4 zip codes as "harbinger zip codes", we can use the *% Groups 3 and 4* model to measure the increase in predictive power from knowing what proportion of new product purchases are made by harbinger zip codes. Predictive accuracy increases to 74.04% in this model.[10] The most accurate model is the *Full Model*, for which predictive accuracy is 75.21%. We conclude that knowing which zip codes are early adopters of a new product can help to predict whether the new product will succeed. As Anderson et al. (2015) discuss in the original harbingers paper, improvements in forecasting new product success can improve the profitability of the new product development process for both manufacturers and retailers. In our next analysis we use the nineteen control variables to highlight the distinguishing characteristics of harbinger zip codes.

Which Zip Codes are Harbinger Zip Codes?

We used the nineteen control variables together with *Total Orders* to establish a benchmark for predicting new product success. We then showed that the harbinger variables can predict variation in new product success that exceeds this benchmark. We can also use the nineteen control variables to identify characteristics that distinguish harbinger zip codes from other zip codes. In Table 3 we report the pair-wise correlation between each of these control variables and the *Average Success Rate*. The unit of analysis is a zip code, and so we use the raw values of each of the control variables (rather than the transaction weighted variables). The *Average Success Rate* measures the average success rate of new product purchases by households in each zip code. It is calculated using the full sample of 4,712 new products, including both the classification and validation sets.

Tables 2 and 3 about here

Recall that harbinger zip codes have low *Average Success Rates*, and so a positive (negative) correlation in Table 3 indicates that harbinger zip codes have relatively low (high) values on that variable. The findings reveal that harbinger zip codes are less urban than other zip codes.

[10] We can also allow *% Group 3* and *% Group 4* to enter the model individually. This yields a predictive accuracy of 74.79% (an 8.88% improvement over the baseline).

None of MassStore's stores are in rural locations, and so the non-urban locations can generally be interpreted as suburban locations. Perhaps consistent with suburban locations, the zip codes tend to have lower household incomes and home values, older heads of households, and a higher proportion of single family homes. They are also located further away from both MassStore's stores and its competitors' stores. Households in harbinger zip codes are relatively less educated (they are less likely to have graduated with bachelor's degrees). They also tend to have proportionately larger white populations, with fewer African Americans, Asians or Hispanics. We also see evidence that they are more deal sensitive, purchasing a higher proportion of items with coupons, using coupons that tend to offer larger discounts, and buying items that have a higher average unit price.

We also investigated how much of the variation in the *Average Success Rate* is explained by the control variables. These findings are reported in the Appendix. The findings reveal that collectively the 19 control variables explain 45.2% of the variation in the *Average Success Rate*. While this confirms that the control variables are associated with changes in the *Average Success Rate*, they also confirm that there is additional variation that these control variables do not capture. For completeness, in the Appendix we also report the coefficients when separately regressing *Average Success Rate* on the control variables within each of the four harbinger zip code groups. Before extending the analysis to other settings, we further investigate the differences between harbinger zip codes and non-harbinger zip codes at MassStore by comparing purchases of existing items (that were not newly introduced).

Low Volume or Niche Items

In this analysis we first construct a measure that identifies which items are purchased infrequently within a specific three-digit zip code. We then compare whether harbinger zip codes tend to purchase more of these items than their neighboring zip codes. We use the number of orders within a 3-digit zip code to define items that are "low volume" or "niche". In particular, we rank the items within each 3-digit zip code according to how frequently they were purchased in the 2013 calendar year. As an indication of the size of a 3-digit zip code, the

greater Boston metropolitan area includes eight 3-digit zip codes.[11] Within each 3-digit zip code we identify the least frequently ordered items using different cumulative purchase thresholds. We define *Low Volume* items as the lowest volume items that collectively contribute 50% of total orders (within a 3-digit zip code). We similarly define *Niche* items as the lowest volume items that collectively contribute 10% of total orders.

These labels identify items that are purchased infrequently within neighboring regional zip codes. We can then ask whether these items are purchased more frequently by harbinger zip codes than other zip codes. For each (5-digit) zip code we calculate a weighted average of the proportion of 2013 orders that are low volume (*% Low* Volume) or niche (*% Niche*), where we weight by the number of orders for each item in that zip code.

To reliably calculate the *Average Survival Rate*, we restrict attention to 5-digit zip codes with at least 200 orders of new products. To reliably calculate the *% Low Volume* and *% Niche* we use three-digit zip codes with at least 100,000 item purchases, and 5-digit zip codes with at least 1,000 item purchases. Our goal is to investigate whether harbinger zip codes that purchase new products that tend to fail also purchase existing products that neighboring zip codes are less likely to purchase. In particular, we examine whether there is an association between the *Average Success Rate* of new product purchases within a zip code and the *% Low Volume* and *% Niche* in that zip code. We use a multivariate approach, which controls for the nineteen zip code level characteristics used in our analysis of new product success. In particular, we estimate the following model:

$$Y_z = \alpha + \beta_1 \, Average\ Success\ Rate_z + \boldsymbol{\beta}\ \mathbf{3\text{-}Digit} + \boldsymbol{\beta}\mathbf{X}_z + \varepsilon_z \qquad (4)$$

The unit of analysis is a zip code and the **β 3-Digit** term represents a complete set of 3-digit zip code fixed effects. We estimate Equation 4 separately using either the *% of Low Volume* items or *% of Niche* items as the dependent variable. The $\boldsymbol{\beta}\mathbf{X}_z$ term describes the nineteen zip-code level control variables. Unlike the previous analysis, where the unit of analysis was a new product, these are the raw zip code level measures, not the transaction-weighted measures. We include the full set of 19 control variables. The coefficient of interest is the *Average Success*

[11] Three-digit zip codes comprise the first three digits of each five-digit zip code. For example, the three-digit zip code for 5-digit zip code 12345 is 123. More information about three-digit zip codes can be found at: http://pe.usps.com/Archive/PDF/DMMArchive20030810/L002.pdf

Rate coefficient (β_1), which identifies how the outcome variable varies with the *Average Success Rate*. Recall that this variable measures the success rate of new products purchased from MassStore by customers living in that zip code. This measure is calculated using all of the new products (including new products in both the classification and validation sets). In particular, we use the quartile grouping of zip codes to construct binary indicators revealing which group each zip code is in. We then estimate Equation 5:

$$Y_z = \alpha + \beta_1 \textit{ Group 2}_z + \beta_2 \textit{ Group 3}_z + \beta_3 \textit{Group 4}_z + \boldsymbol{\beta}\ \textbf{3-Digit} + \boldsymbol{\beta}\mathbf{X}_z + \varepsilon_z \qquad (5)$$

Because we omit the Group 1 binary indicator from this model, the outcome in Group 1 serves as the baseline. The coefficients estimated for the *Group 2*, *Group 3* and *Group 4* indicator variables compare the change in the dependent variable in groups 2 through 4 (compared to this baseline). The model again includes fixed effects identifying each 3-digit zip code, together with the full set of control variables.

The coefficients of interest are reported in Table 4. They confirm that harbinger zip codes not only purchase new products that have a lower *Average Success Rate*, they are also more likely to purchase low volume and niche existing products. Recall that low volume and niche products are defined within a 3-digit zip code and so harbinger zip codes purchase existing products that customers in *neighboring zip codes* are less likely to purchase. This is consistent with our interpretation that households in these zip codes have tendencies that are not representative of other customers. These results also serve to introduce Equations 4 and 5. We will use these equations repeatedly throughout the paper to investigate how the MassStore *Average Success Rate* (calculated using new item purchases) explains variation in decisions across zip codes in different decision contexts. We next investigate whether the harbinger effect may result from differences across zip codes in firm marketing activities.

Table 4 about here

Differences in Marketing Activities

Our interpretation of harbinger tendencies embraces systematic differences in customers' sensitivities to the firm's marketing actions. An alternative explanation is that there is a difference in the MassStore's marketing activities across zip codes. Before investigating this explanation, we first note that in order for the findings in this paper to be explained by

systematic differences in firms' marketing actions across zip codes, it would need to be the case that MassStore, ApparelCo, and congressional election candidates all adopt similar systematic differences in their marketing actions across zip codes. This does not seem plausible, unless MassStore, ApparelCo, and congressional election candidates are all responding to differences in customers' sensitivities to marketing actions. This reverses the causality: harbinger tendencies may create systematic differences in firms' marketing actions, rather than the firms' marketing actions creating harbinger tendencies.

To evaluate the variation in marketing activities across zip codes we investigated whether the firm engaged in target marketing during our data window. This information is described in the Appendix. We learned that at that time there was little variation in this retailer's marketing activities across zip codes. Although not conclusive, it suggests that the effects we report are unlikely to be attributable to the firm's marketing activities. To supplement this evidence, we also conducted additional analysis using newly acquired customers. There are two findings of interest. First, the results replicate the harbinger effect: new customers in harbinger zip codes are more likely to purchase *Low Volume* items, compared to new customers in non-harbinger zip codes. This replication is reassuring, as both the customers and the time period are different than the results reported in the main text. Second, within a single promotion condition there is strong evidence of a harbinger effect. Notably, within a single promotion condition, the new customers were all exposed to the same marketing actions, and so the effects cannot be attributed to differences in marketing actions. Together, these findings allow us to conclude that the harbinger effect is robust, and is not solely due to variation in marketing activities.

Summary

Using data from MassStore, we have shown that harbinger zip codes exist. Households in these zip codes are consistently more likely to purchase new products that fail than households in other zip codes. They are also more likely to purchase niche products that are rarely purchased by neighboring zip codes. These harbinger zip codes tend to be located in suburban rather than urban areas. They also tend to be located further away from MassStore's (and its competitors') stores, they contain fewer households, and they have lower average home values, more single-family homes, and older heads of households.

The findings extend the results reported by Anderson et al. (2015) by identifying clusters of harbingers at the zip code level. Recall that Anderson et al. (2015) used household level data to identify harbinger customers - households whose purchases of a new product are an indication that the product will fail. Anderson et al. (2015) acknowledged that an important challenge limiting the application of their finding is that detailed household level data is not always available to identify harbingers. By extending the result to the zip code level, we help overcome this challenge and expand the potential applications for their result.

The evidence that we can identify harbingers at the zip code level will be particularly helpful if the same zip codes exhibit harbinger tendencies across a broad range of decision contexts. We investigate this in the reminder of this paper, where we use the classification of harbinger zip codes in the MassStore data to explain variation in purchasing decisions at another retailer (Section 4), donations to congressional election campaigns (Section 5), and changes in house prices (Section 6).

4. Harbinger Zip Codes at MassStore and Purchases from ApparelCo.

Unlike MassStore, where successful new products are continued and unsuccessful new products are discontinued, ApparelCo often updates its products between seasons irrespective of whether the products are successful. This means that observing that a new product was discontinued does not imply that the product flopped. ApparelCo does not have an obvious metric for evaluating new product success or failure, and so we are unable to repeat the same new product analysis that we conducted with the MassStore data.

Instead, we focus our analysis of the ApparelCo data on three questions. First, we investigate whether harbinger zip codes were more likely to purchase "niche" items that households in neighboring zip codes were less likely to purchase. Second, we investigate whether harbinger zip codes are more likely to purchase products that other customers return. Finally, we identify a sample of households that changed zip codes. We use this sample to ask whether the clustering of harbingers in harbinger zip codes occurs because of where households choose to live, or because households change their tendencies when they move into a harbinger zip code. An important difference in this analysis is that we classify the zip codes into harbinger and non-harbinger zip codes using new product purchases from one company (MassStore) and use

the classifications to compare purchases from a different company (ApparelCo). We start with the analysis of niche purchases.

Do Harbingers Purchase Niche Items from ApparelCo?

We use the same approach that we used to study niche purchases at MassStore. In particular, we use the number of purchases within a three-digit zip code to define items that are "low volume" and "niche" (see the earlier definitions). We then calculate a weighted average of apparel orders in a zip code that are low volume or niche. We construct the data sample using the same approach as our earlier analysis. We use three-digit zip codes with at least 100,000 purchases of any item from ApparelCo and zip codes with at least 1,000 new product purchases. We also restrict attention to zip codes with at least 200 orders of new products from MassStore. When calculating the *Average Success Rate* at MassStore, we use all 4,712 new MassStore products, including both the classification and validation samples. In Table 5 we report the coefficients of interest when estimating Equations 4 and 5 using *% Low Volume* or *% Niche* as dependent variables. Recall that these models include fixed effects at the 3-digit zip code level, together with the full set of 19 control variables.

Table 5 about here

The findings confirm that harbinger zip codes at MassStore purchase a higher percentage of low volume and niche items than their neighboring zip codes (that share the same 3-digit zip code). The Equation 5 coefficients indicate that the harbinger zip codes (Group 4), which purchase the highest proportion of MassStore new products that fail, purchase a 1.29% higher proportion of *low volume* ApparelCo items, than Group 1 zip codes. As a benchmark, the average proportion of low volume items purchased is 50% (by construction). They also purchase 0.42% more *niche* items (where the benchmark proportion of niche items is 10%). These is an extremely robust pattern of results. For example, the findings reported in Table 5 use ApparelCo purchases from 2011, and if we repeat the analysis using purchases from 2010, the pattern of results is unchanged.

In additional analysis reported in the Appendix, we also use the ApparelCo data to compare whether harbinger zip codes are more or less likely to purchase items that other customers return. If harbingers have unusual preferences, then we might expect that products purchased

by harbingers would tend to have relatively high return rates from other customers. The findings confirm that harbinger zip codes are more likely than neighboring zip codes to purchase items that other customers return. A 10% decrease in the *Average Success Rate* in a zip code is associated with an increase of more than 2% in the average return rate (by customers in other zip codes).

The analysis of return reasons reveals that the tendency of harbingers to buy items that other customers return is concentrated in returns due to (a) size and (b) because customers did not like the item. Preferences for the color, material or styling, together with the size and fit of an item are idiosyncratic, reflecting customers' body shapes and personal tastes. The evidence that harbingers are more likely to purchase items that other customers return for these idiosyncratic reasons suggests that harbinger tendencies may be associated with differences in some product preferences. In contrast, quality defects are not idiosyncratic and we would expect returns due to defects to affect all customers in approximately the same way. It is therefore notable that the effects are much smaller for returns due to defects. This is consistent with the interpretation that harbinger tendencies are associated with idiosyncratic product preferences.

In our final ApparelCo analysis we shift focus and ask whether harbinger zip codes result from households with harbinger tendencies choosing to cluster together, or whether households learn these tendencies by observing their neighbors. We do so by studying ApparelCo customers that changed zip codes.

Do Households Bring Their Tendencies or Learn Them?

We identify households' zip codes using the addresses to which they shipped orders.[12] All of the households placed at least one order in both 2007 and 2011. Their 2007 orders were all shipped to one zip code and their 2011 orders were all shipped to another zip code (we exclude households that shipped to multiple zip codes within either calendar year).[13] We first ask

[12] In the other analysis in this section we identify households' zip codes using their registered addresses. However, we only have registered addresses for a subset of the households in 2007. The registered address is generally the billing address for their credit cards (we only received data describing the 5-digit zip code, not the complete street address).

[13] Note that because the data period ends on 9 December 2011 the years are defined as ending on 9 December of that year and so include 22 days from the previous calendar year. There is a small risk of error because households sometimes ship items as gifts. Although, this is unlikely to provide an alternative explanation for our findings, we

whether households that started in harbinger zip codes in 2007 moved to other harbinger zip codes in 2011. There were 28,476 households that changed zip codes from 2007 to 2011. We grouped zip codes into quartiles using the MassStore new product survival rates, and mapped the transfers of these 28,476 households between each group of zip codes. This mapping is summarized in Table 6.

Table 6 about here

We see clear evidence that when households in harbinger zip codes move, they move to another harbinger zip code. Of the 6,375 households that started in a Group 4 zip code, 2,889 (45.3%) moved to another Group 4 zip code, but just 740 (11.6%) moved to a Group 1 zip code. In contrast, of the 8,950 that started in a Group 1 zip code, just 1,080 (12.1%) moved to a Group 4 zip code, while 4,097 (45.8%) moved to another Group 1 zip code. Across all 28,476 households that changed zip codes, the pair-wise correlation between the *Average Success Rate* in 2007 and 2011 is 0.63. These findings reveal systematic differences in where households choose to live. They suggest that one reason we see harbinger zip codes is that harbingers choose to move into zip codes occupied by other harbingers, while non-harbingers choose to do the opposite. For example, if non-harbingers like ocean views but harbingers do not, then we would expect to see clusters of harbingers in zip codes without ocean views.

We conducted two robustness checks. First, we grouped the households according to whether they moved within a 3-digit zip code or to a separate 3-digit zip code. The pattern of findings survives in either case. Second, instead of identifying households from the zip codes to which they mailed purchases, we identified them from the address they have registered with the retailer. In particular, we compared the zip codes in each household's registered address as of 19 October 2007 and 9 December 2011. We only know the registered addresses in 2007 for a randomly selected subsample of the households, which reduced the sample size. The pattern of findings was again unchanged.

We can also use households that changed zip codes to address a second question; do households change their purchasing decisions when they move into a harbinger zip code? In particular, we investigate whether households' purchases of niche items changed after they

will also replicate the results for a subset of the customers for which we have data describing their registered addresses in both 2007 and 2011.

moved zip codes. This analysis is reported in the Appendix. We see no evidence that movement from a more harbinger zip code to a less harbinger zip code results in a reduction in the proportion of niche items purchased (or vice versa). We conclude that households do not appear to acquire their harbinger tendencies from their new neighbors.

Summary

The findings reported in this section represent the first evidence that the identification of harbinger zip codes is stable across decision contexts. Identifying harbinger zip codes from purchases at one company helps to explain variation in purchasing at another company. We also used a sample of ApparelCo customers that changed zip codes to investigate whether the harbinger zip code effect is due to where harbingers choose to live, or due to changes in tendencies when households move into a harbinger zip code. There is strong support for the first interpretation; households leaving a harbinger zip code tend to move to another harbinger zip code, while the reverse is true of those leaving non-harbinger zip codes. However, there is no evidence that households learn harbinger tendencies when they move to a harbinger zip code. This suggests that harbinger tendencies are a sticky trait, and the harbinger zip code effect is more due to where customers choose to live, rather than households learning the tendencies of their neighbors.

5. Change in House Prices

In Section 4 we used ApparelCo customers that changed zip codes to investigate whether harbinger zip codes result from harbinger households choosing to cluster together, or from households changing their tendencies when they move into harbinger zip codes. The findings strongly supported the first explanation. Households that move from a harbinger zip code tend to move to another harbinger zip code, while non-harbinger zip codes do the reverse. We illustrated how this could contribute to the formation of harbinger zip codes using an example of ocean views. If non-harbingers like ocean views but harbingers do not, then we would expect to see clusters of harbingers in zip codes without ocean views. By choosing where they live, households also determine (in part) how much their home values will increase. In this section we investigate whether harbinger zip codes experience larger or smaller house price

increases. We measure the relationship between the *Average Success Rate* of new products purchased at MassStore and the average change in house prices.

The decision to live in a harbinger zip code may reflect more than just a preference for the location. It may also reflect a preference for the housing stock within that location, and this could also contribute to the rate at which house prices change. For example, a recent study by Trulia (reported in Bloomberg 2016) discovered that customers who buy unusual houses suffer smaller house price increases. Not only are "McMansions" unattractive to look at, but their prices increases at slower rates than other houses. If unusual tendencies lead customers to purchase unusual houses, then these tendencies may also affect home values.

The findings in this section are different in several respects to the findings in the other sections. First, we focus on the housing market rather than product choices or political decisions. Second, while the datasets used in the previous sections document explicit customer decisions, the connection between changes in house prices and household decision-making is less direct. Changes in house prices are measured using transactions, and so changes in house prices are related to the pricing decisions of some households.[14] However, the decision that is perhaps most relevant to this study is the choice of which zip code to live in. Households make this decision when they first enter the zip code (unless they inherit the property), and at least implicitly again when they decide to continue to live there.

A third difference between the findings in this section and the results in the previous section is that it is difficult to argue that the results are solely attributable to unusual customer preferences. We do not have any data describing features of the location or housing stock that reflect unusual preferences. This means that there are a range of intervening factors that could contribute to the findings, including for example the earlier evidence that harbinger zip codes tend to be more suburban. Rather than trying to establish the mechanism that drives our results, we focus instead on documenting a surprising relationship between the adoption of successful (or unsuccessful) new products at MassStore and changes in house prices. We also conduct extensive robustness checks to determine the limits of this relationship. We begin by

[14] Obviously, many households in a zip code do not participate in the transactions that determine the change in house prices. However, it is also true that many of the households in each zip code do not purchase from MassStore and ApparelCo, and do not donate to congressional election campaigns.

illustrating the general trend in house prices using Zillow's US median price index.[15] We use the Zillow data to calculate the year-on-year *Price Change* in zip code *z* in year *t*:

$$Price\ Change_{zt} = \frac{Price\ End_{zt} - Price\ Start_{zt}}{0.5 * Price\ End_{zt} + 0.5 * Price\ Start_{zt}}$$

Using the average of the two prices in the denominator ensures that price increases and decreases are treated symmetrically. A positive (negative) value of *Price Change* indicates a price increase (decrease). We then calculate the average of this *Price Change* for each year by averaging across the zip codes. The averages are calculated using a common sample of 4,291 zip codes for which both Zillow data and MassStore data are available. The average year-on-year *Price Change* is reported in Figure 1.

Figure 1 about here

We see that 2002 and 2006 was a period during which house prices increased on average, before the fall in prices from 2007 to 2011. House prices then recovered between 2012 and 2015. We might expect that the zip codes in which house prices increased the most between 2002 and 2006 are also the zip codes in which house prices decreased the most between 2007 and 2011. Therefore, we will investigate changes in house prices across three separate periods: January 2002 to December 2006, January 2007 to December 2011, and January 2012 to December 2015. For each zip code we construct $Price\ Change_{zt}$ using prices at the start and end of these multi-year periods and use them as dependent variables to estimate Equation 4 (which we re-state below):

$$Y_z = \alpha + \beta_1\ Average\ Success\ Rate_z + \boldsymbol{\beta}\ \textbf{3-Digit} + \boldsymbol{\beta}\mathbf{X}_z + \varepsilon_z \qquad (4)$$

Recall that the unit of analysis is a zip code and $Average\ Success\ Rate_z$ is the average success rate of new products purchased from MassStore by households in zip code *z*. The **β 3-Digit** term denotes fixed effects identifying each 3-digit zip code and the $\mathbf{X}_z$ term represents the control variables. We re-estimate Equation 4 separately using the *Price Change* measures calculated for each of the three multi-year periods. Because each model focuses on a single multi-year period, the coefficient of interest (β_1) is only identified by cross-sectional variation

[15] Available at: http://files.zillowstatic.com/research/public/State/State_Zhvi_Summary_AllHomes.csv

in the *Average Success Rate* (i.e. these are not time series models). All three models are estimated using the same sample of 4,291 zip codes. The coefficient of interest is β_1, which measures how changes in house prices vary within a 3-digit zip code according to the *Average Success Rate* of new product purchases at MassStore.[16]

The coefficients of interest are summarized in Table 7. They indicate that during periods in which house prices increased on average, the harbinger zip codes tended to have smaller house price increases. In particular, a 10% smaller *Average Success Rate* is associated with a 1.3% smaller house price increase between 2002 and 2006, and a 2.0% smaller increase between 2012 and 2015. Not surprisingly, these zip codes also experienced a smaller house price decrease when average house prices fall from 2007 to 2011. This is consistent with these zip codes experiencing a smaller run up in house prices between 2002 and 2006. We will focus our attention on price changes during the two periods that average house prices increased.

Table 7 about here

As a robustness check we controlled for the initial price at the start of each period prices as a control variable in the regressions. We also replaced the fixed effects identifying each 3-digit zip code with fixed effects identifying each 4-digit zip code. This reduces the price variation to even smaller regions, with at most ten (5 digit) zip codes within each 4-digit zip code. The pattern of findings was unchanged under both of these modifications (see the Appendix).

We caution that the relationship between new product failure rates at MassStore and changes in house prices is unlikely to be causal. Smaller house price increases are unlikely to have caused these households to purchase new product flops at MassStore. Instead, there is almost certainly an unobserved intervening variable(s) that explains which zip codes are harbinger clusters, and which zip codes had the smallest house price increases. As we discussed, in the previous sections of this paper we have interpreted this unobserved variable as "unusual tendencies." With these house price data, it could represent a range of different underlying

[16] The inclusion of the fixed effects again ensures that β_1 is only identified by variation in house prices within each 3-digit zip code (and not by regional variation across 3-digit zip codes).

explanations. Recall, however, that we included the full set of 19 control variables, making it unlikely the effect is solely attributable to income, age, or any of the other controls.

Summary

In this section, we showed that during periods in which house prices were generally increasing, harbinger zip codes experienced smaller house price increases than other zip codes. This is a robust finding that survives controlling for starting prices, controlling for a wide range of demographic variables, and limiting identification to variation within 3-digit and 4-digit zip codes.

In the next section, we further explore the extent to which harbinger tendencies are correlated across choices beyond purchasing. We focus on donations to congressional election candidates and investigate whether MassStore's harbinger zip codes contribute to the same candidates as households in neighboring zip codes. We also investigate whether the candidates supported by harbinger zip codes are more or less likely to win their elections.

6. Contributions to Federal Election Campaigns

In the previous sections, we have shown that zip codes identified from new product purchases at MassStore make decisions that are different from households in other zip codes. What is most surprising about these results is that they extend across product categories and across retailers. In this section, we investigate whether the effect also extends to political tendencies.

We start by comparing whether harbinger zip codes make contributions to election candidates in proportions that are systematically different than their neighboring zip codes (other zip codes in the same 3-digit zip). We use data describing the total $ amount contributed to congressional campaigns between 2000 and 2010. We first describe the variables that we constructed for this analysis, and then introduce the analysis itself.

Political Contributions

For each 3-digit zip code, we identified the two candidates that received the largest dollar contributions. We then calculated the proportion of dollars that each of these candidates received:

$$zip3\ amount = \frac{\$\ for\ candidate\ with\ most\ \$}{total\ \$\ for\ top\ 2\ candidates}$$

We also calculated similar measures at the 5-digit zip level, where the top 2 candidates were defined by the total dollar contributions at the zip3 level:

$$zip5\ amount = \frac{\$\ for\ candidate\ with\ most\ \$}{total\ \$\ for\ top\ 2\ candidates}$$

Our outcome measure is a binary indicator *More $ for Top Candidate*, which is equal to 1 if *zip5 nbr* is greater than *zip3 nbr*, and zero otherwise. This binary indicator helps to reduce noise introduced if zip codes have relatively few donors. Given these definitions, a value of 1 for *More $ for Top Candidate* indicates that this 5-digit zip code contributed a higher proportion of dollars to the most popular candidate than the 3-digit zip code's average.

To analyze how this outcome measure with the *Average Success Rate* of new product purchases at MassStore we estimated Equations 4 and 5 using *More $ for Top Candidate* as the dependent measure. The unit of analysis is a (5-digit) zip code x election cycle x congressional district.[17] The findings are reported in Table 8. The findings indicate that harbinger zip codes are less likely than other zip codes to donate to the candidates that receive the most donations in their 3-digit zip code. In particular, households in the zip codes with the lowest *Average Success Rate* of new products purchased at *MassStore* are approximately 2% less likely to contribute more dollars to the most popular candidate compared to zip codes with the highest *Average Success Rate*.

Table 8 about here

As a robustness check we repeated the analysis using the 'number of donors' instead of the 'dollar amount donated'. The findings were very similar. As an additional robustness check we also re-estimated the results using a more aggregate unit of analysis. The findings in Table 8 use a zip code x election cycle x congressional district as the unit of analysis. Reassuringly, when aggregating to the zip code x election cycle level, the pattern of findings is unchanged.

[17] For some congressional elections the total contribution was identical for the top two candidates in a 3-digit zip code. We omitted these observations. We also restrict attention to zip codes with at least one donation to the top 2 candidates. Finally, we also exclude a small number of observations in which a 3-digit zip code only had contributions from a single 5-digit zip code for that election cycle (by construction, the proportion contributed by the 5-digit zip code is the same as the 3-digit zip code in these observations).

An alternative approach to analyzing the political contributions data is to ask: are harbinger zip codes less likely to contribute to the winning candidate? We answer this question in additional analysis reported in the Appendix. The findings confirm that harbinger zip codes are systematically more likely than their neighboring zip codes to contribute to congressional election candidates that lose their elections. The effect is particularly strong in the zip codes with the lowest 5% of *Average Success Rates* in new product purchases at MassStore (the most harbinger-like zip codes).

Summary

We have shown that harbinger zip codes identified from new product purchases at MassStore tend to support congressional candidates that neighboring zip codes are less likely to support. They also support congressional candidates that are less likely to be elected than other zip codes. This confirms that harbinger zip codes can explain variation in choices that extend beyond retail settings. In our final set of analyses, we investigate reversing the relationship we have documented in this and the previous sections. We have shown that the *Average Success Rate* of new product purchases can help to explain variation in purchasing decisions at ApparelCo, variation in house price increases, and variation in contributions to federal election campaigns. In the next section we investigate whether the ApparelCo, house price and campaign contribution data can help to explain variation in new product success at MassStore.

7. Reversing the Relationship

Our analysis of whether decisions from other contexts can help explain variation in new product success at MassStore is analogous to the way we introduced control variables in Equation 1. We construct nine zip code level variables measuring: niche purchases at ApparelCo, price changes at Zillow, and contributions to federal election campaigns (we will label them collectively the *Zip Code Measures*). Each new product transaction (within the first 90 days of new product entry) is associated with a zip code. Therefore, for each new product we calculate a transaction-weighted average for each *Zip Code Measure*. We modify Equation 1 to include these transaction-weighted averages:

$$Success_j = \alpha + \beta_1 \, log\ Total\ Orders_j + \beta_2 \, Transaction\ Weighted\ Zip\ Code\ Measure_j + \mathbf{\beta X} + \varepsilon_j \quad (6)$$

The unit of analysis is a new product at MassStore and *Total Orders* measures the total number of purchases of item *j* in the first 90 days after it is introduced. The **βX** term describes the transaction weighted averages for all nineteen control variables. The *Zip Code Measures* are introduced to the model separately (we estimate separate models using each variable). These nine measures include:

ApparelCo Purchases	**Zillow House Prices**
% of Niche Purchases 2010	Price Change 2002 to 2006
% of Niche Purchases 2011	Price Change 2007 to 2011
% of Low Volume purchases 2010	Price Change 2012 to 2015
% of Low Volume purchases 2011	

Political Data (aggregated across election cycles and congressional districts)
More $ for Top Candidate
% of Winning Candidates

Our estimation sample includes all 4,712 new products at MassStore (pooling the classification and validation samples). The coefficient of interest is β_2, which measures how variation in each transaction weighted *Zip Code Measure* is associated with new product success at MassStore. We report these coefficients in Table 9.

Table 9 about here

Our earlier results suggest that new product success at MassStore should be associated with:

- a lower proportion of niche purchases at ApparelCo,
- a larger Zillow house price increase from 2002 to 2006 and 2012 to 2015,
- a lower Zillow house price decrease from 2007 to 2011,
- more $ for the top candidate, and
- a larger percentage of winning candidates.

The coefficients in Table 9 confirm that many of these relationships are significant. Moreover, the significant relationships are all in the direction that we expected. We conclude that decisions in other contexts can explain variation in new product success at MassStore. More generally, the evidence that the effect is robust to reversing the relationship is consistent with our interpretation that the harbinger effect reflects underlying differences in customer tendencies. These tendencies are revealed by the different decisions. What is remarkable is that

decisions in one context reveal tendencies that explain variation in decisions in very different contexts. These findings not only demonstrate the robustness of our findings, but also demonstrate the practical advantages of identifying harbingers at the zip code level instead of the individual level. There is a rich array of data available at the zip code level, in a variety of decision contexts. Even though the decision contexts are different, this zip code level data can be used to evaluate who the early adopters of a new product are (rather than just the total number of new adopters).

8. Conclusions

Using data from multiple sources, we have shown that the phenomenon of harbingers is surprisingly widespread. We begin by showing that harbinger zip codes exist. Households in these zip codes are more likely than households in other zip codes to purchase new products that fail. Their adoption of a new product is a signal that the new product will fail. We interpret this finding as evidence that households in these zip codes have tendencies that are not representative of households in other zip codes. We then show that the evidence of unusual tendencies extends across retail product categories and across retailers.

What makes these results particularly surprising is that while we measure the average outcome for a zip code, relatively few households in each zip code participate in each decision. Not every household purchases from the retailers that we study, and relatively few households contribute to congressional election candidates. Moreover, the households that participate will often be different households for each decision. It is unlikely that the households who purchase from one retailer are all the same households that purchase from the other retailer. They are also unlikely to all make donations to congressional election campaigns. Despite this, we observe similarities in zip code level decisions across these different purchasing contexts.

We explore two explanations for why households with unusual tendencies cluster together. This analysis uses a sample of households that changed zip codes. The analysis reveals that households that moved from a harbinger zip code tended to move to another harbinger zip code. Similarly, households that started in a non-harbinger zip code generally moved to another non-harbinger zip code. This suggests that harbinger zip codes arise at least in part from customers choosing to cluster with other households that have similar tendencies. We did

not find any support for the alternative explanation that customers learn their tendencies when they move into a harbinger zip code. It appears instead that harbinger tendencies are relatively sticky and that harbinger households bring their tendencies when they change zip codes, rather than learning them when they get there.

We have shown the predictive power of these harbinger zip codes using data spread over several years. Future research could investigate how the predictive power of this identification diminishes (or increases) over time as households move and environments change.

9. References

Anderson, Eric, Chaoqun Chen, and Blake McShane (2018), "Harbingers of Failure … and Success," working paper, Kellogg School of Management.

Anderson, Eric, Chaoqun Chen, Caiyun Liu and Duncan Simester (2017), "Do Harbinger Products Signal Which New Customers Will Stop Purchasing?" working paper, MIT Sloan School of Management.

Anderson, Eric, Song Lin, Duncan Simester, and Catherine Tucker (2015), "Harbingers of Failure," *Journal of Marketing Research*, 52(5), 580-592.

Anderson, Evan E. (1979), "An Analysis of Retail Display Space: Theory and Methods," *Journal of Business*, 52(1), 103-118.

Bass, Frank M. (1969), "A New Product Growth Model for Product Diffusion," *Management Science*, 15(5), 215-227.

Bloomberg (2016), "McMansions Define Ugly in a New Way: They're a Bad Investment," August 23, 2016, http://finance.yahoo.com/news/mcmansions-define-ugly-way-bad-142534896.html.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Mohammad S. Rahman (2009), "Battle of the Retail Channels: How Product Selection and Geography Drive Cross-Channel Competition," *Management Science*, 55(11), 1755–1765.

Brynjolfsson, Erik, Yu (Jeffrey) Hu and Duncan Simester (2011), "Goodbye Pareto Principle, Hello Long Tail: The Effect of Internet Commerce on the Concentration of Product Sales," *Management Science*, 57(8), 1373-1387.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith (2003), "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 49(11), 1580–1596.

Choi, Jeonghye and David R. Bell (2011), "Preference Minorities and the Internet," *Journal of Marketing Research*, XLVIII, August, 670-682.

Galak, Jeff, Kurt Gray, Igor Elbert and Nina Strohminger (2016), "Trickle-down Preferences: Preferential Conformity to High Status Peers in Fashion Choices," *PLOS ONE*, 11(5), May 4 2016, 1-11.

Mahajan, Vijay, Eitan Muller and Frank M. Bass (1990), "New Product Diffusion Models in Marketing: A Review and Directions for Research," *Journal of Marketing*, 54(1), 1-26.

Waldfogel, Joel (2003), "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated Product Markets," *RAND Journal of Economics*, 34(3), 557-568.

Table 1. Harbinger Zip Codes at MassStore

	Equation 1	**Equation 2**	**Equation 3**
Log Total Orders	13.06%** (1.49%)		4.30%** (1.60%)
Log Orders Group 1		168.38%** (15.66%)	
Log Orders Group 2		-16.59% (19.02%)	
Log Orders Group 3		-36.83%* (19.70%)	
Log Orders Group 4		-109.09%** (14.79%)	
% Group 2			-165.09%** (50.70%)
% Group 3			-287.58%** (48.91%)
% Group 4			-661.21%** (54.01%)
R^2	0.1692	0.2188	0.2246

The table reports the coefficients of interest from estimating Equations 1, 2 and 3. Coefficients for the control variables are reported in the Appendix. In both models the unit of analysis is a new product and the sample size in all models is 2,388. Standard errors are in parentheses.

Table 2. Predictive Model: Holdout Accuracy

Model	% Correct	Improvement Over Baseline
Naïve Baseline	65.91%	
Controls Only	71.36%	5.44%** (1.35%)
Total Orders	70.35%	4.44%** (1.42%)
% Group 1	73.03%	7.12%** (1.44%)
% Group 2	70.02%	4.10%** (1.44%)
% Group 3	71.44%	5.53%** (1.44%)
% Group 4	73.87%	7.96%** (1.45%)
% Group 3 and 4	74.04%	8.12%** (1.43%)
Full Model	75.21%	9.30%** (1.45%)

The table reports the % of predictions that are correct. The unit of analysis is a new product and the sample sizes for the estimation and the holdout sample are each 1,194. Standard errors are in parentheses.

Table 3. Demographic Characteristics

Age	-0.2971	High School	-0.1710
Home Value	0.2865	Bachelors	0.1197
Income	0.0646	White	-0.3919
Single Family	-0.4047	African American	0.2232
Multi-family	0.4132	Asian	0.3426
Distance	-0.3740	Hispanic	0.4854
Comp. Distance	-0.4156	Coupon Discount	-0.2333
Nbr Households	0.3160	Coupon Frequency	-0.3741
Urban	0.4190	Unit Price Paid	-0.0523
Urban Clusters	-0.1603		

The table reports the reports the pair-wise correlation between each (raw) control variable and the *Average Success Rate* of new product purchases in that zip code. The control variables are defined in the Appendix, where we also provide summary statistics. The unit of analysis is a zip code. The sample size for the pair-wise correlations is 4,689 zip codes. All of the pair-wise correlations are significantly different from zero at $p < 0.01$.

Table 4. Purchases of Niche Items at MassStore

	% Niche (Bottom 10%)		% Low Volume (Bottom 50%)	
	Equation 4	**Equation 4**	**Equation 4**	**Equation 5**
Average Success Rate	-14.31%** (0.92%)		-20.35%** (1.82%)	
Group 2		0.26%** (0.05%)		0.48%** (0.10%)
Group 3		0.52%** (0.06%)		0.88%** (0.11%)
Group 4		0.84%** (0.07%)		1.26%** (0.13%)
R^2	0.5847	0.5788	0.5983	0.5962

The table reports the coefficients of interest from estimating Equations 4 and 5 with different dependent variables. Fixed effects at the 3-digit zip code together with coefficients for the control variables are estimated but omitted from the table. The unit of analysis is a zip code and the sample size in both models is 4,679. Standard errors are in parentheses.

Table 5. Purchases of Niche Items at ApparelCo

	% Niche (Bottom 10%)		% Low Volume (Bottom 50%)	
	Equation 4	**Equation 5**	**Equation 4**	**Equation 5**
Average Success Rate	-5.12%* (2.39%)		-18.77%** (7.27%)	
Group 2		0.13% (0.11%)		0.42% (0.33%)
Group 3		0.36%** (0.12%)		0.91%* (0.36%)
Group 4		0.42%** (0.13%)		1.29%** (0.41%)
R^2	0.2336	0.2386	0.2115	0.2141

The table reports the coefficients of interest from estimating Equations 4 and 5 with different dependent variables. Fixed effects at the 3-digit zip code together with coefficients for the control variables are estimated but omitted from the table. The unit of analysis is a zip code and the sample size in both models is 1,416. Standard errors are in parentheses.

Table 6. Do Households in Harbinger Zip Codes Move to Other Harbinger Zip Codes?

	Ended in Group 1	Ended in Group 2	Ended in Group 3	Ended in Group 4	Number of Households
Started in Group 1	45.8%	25.7%	16.5%	12.1%	8,950
Started in Group 2	23.4%	29.7%	27.8%	19.1%	6,871
Started in Group 3	18.4%	25.7%	28.5%	27.4%	6,280
Started in Group 4	11.6%	17.0%	26.1%	45.3%	6,375

The table reports the movements of 28,476 households that changed zip codes from 2007 to 2011. The rows indicate the grouping of zip codes that households moved from in 2007, and the columns indicate the grouping of zip codes that the households lived in by 2011. The % entries measure the percentage of households that started in the corresponding group in 2007 and ended up in each group in 2009. The percentages in each row sum to 100%. The zip codes are grouped according to the *Average Survival Rate* for new product purchases at MassStore.

Table 7. Relationship Between Zillow House Price Changes and *Average Success Rate* at MassStore

	2002 to 2006	2007 to 2011	2012 to 2015
Average Success Rate	13.07%* (6.92%)	-16.13%† (8.27%)	20.22%** (6.39%)
R^2	0.8924	0.8851	0.8353

The table reports the *Average Success Rate* coefficients and R^2 values when estimating Equation 4 using different time periods. The unit of analysis is a zip code and the sample size for all of the models is 3,837. Standard errors are in parentheses.

Table 8. Donations to Political Candidates

	Equation 4	Equation 5
Average Success Rate	28.75%** (7.64%)	
Group 2		-1.04%** (0.36%)
Group 3		-2.05%** (0.43%)
Group 4		-1.82%** (0.47%)
R^2	0.0119	0.0120

The table reports the coefficients of interest from estimating Equations 4 and 5. The dependent variable in both models is *More $ for Top Candidate*. Fixed effects at the 3-digit zip code together with coefficients for the control variables are estimated but omitted from the table. The unit of analysis is a zip code x election cycle x congressional district and the sample size in both models is 186,149. Standard errors are in parentheses.

Table 9. Reversing the Analysis: New Product Success at MassStore

		Coefficients	R^2
ApparelCo Purchases	*% of Niche Purchases 2010*	-1532.83%* (714.55%)	0.1736
	% of Niche Purchases 2011	-1479.11%* (737.30%)	0.1735
	% of Low Volume Purchases 2010	-701.84%** (250.61%)	0.1742
	% Of Low Volume Purchases 2011	-514.88%* (248.79%)	0.1698
Zillow House Prices	*Price Change 2002 to 2006*	88.31% (60.16%)	0.1732
	Price Change 2007 to 2011	-184.59%** (55.64%)	0.1748
	Price Change 2012 to 2015	282.01%** (72.65%)	0.1755
Political Donations	*More $ for Top Candidate*	222.03%* (111.90%)	0.1735
	% of Winning Candidates	-109.83% (104.09%)	0.1730

The table reports the coefficients of interest and the R^2 values from estimating Equation 6 using each transaction-weighted *Zip Code Measure*. The sample size in all models is 4,712. Standard errors are in parentheses.

Figure 1. Year-on-Year Change in Zillow Median House Price Index

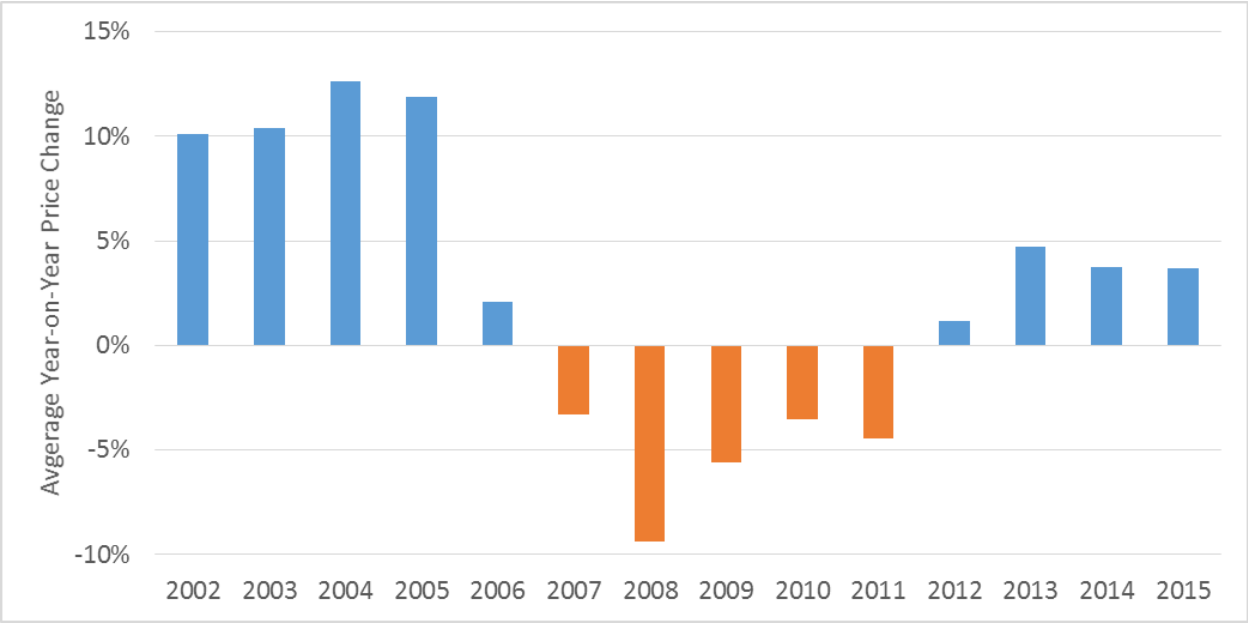

The figure reports the year-on-year Price Change averaged across zip codes. The averages are calculated using a common sample of 4,291 zip codes for which both Zillow data was available every year and the MassStore data provided an average success measure.