ADDING BRICKS TO CLICKS:
THE EFFECTS OF STORE OPENINGS ON SALES THROUGH DIRECT CHANNELS

Jill Avery, Simmons School of Management, jill.avery@simmons.edu*
Thomas J. Steenburgh, Harvard Business School, tsteenburgh@hbs.edu
John Deighton, Harvard Business School, jdeighton@hbs.edu
Mary Caravella, University of Connecticut, mary.caravella@business.uconn.edu
(*corresponding author)

*Simmons School of Management
409 Commonwealth Avenue
Boston, Massachusetts 02115
617.521.3853
Fax 617.521.3880

February, 2008

ADDING BRICKS TO CLICKS:
THE EFFECTS OF STORE OPENINGS ON SALES THROUGH DIRECT CHANNELS

ABSTRACT

We assess the effects of opening physical retail stores on direct-to-consumer channel sales. Our data come from a leading U.S. retailer which opened four new stores over several years in different retail trading areas. We hypothesize two effects, cannibalization and complementarity, and conjecture that the magnitude of these effects may change over time and may differ between the catalog and online channels. We find that opening retail stores cannibalizes sales in the catalog channel in the short term, but produces complementary effects in both the catalog and the online channels in the long term; the complementary effects, which are magnified in the online channel, more than overcome the initial losses in the catalog channel. Customer analysis suggests that opening retail stores paves the way for higher rates of customer acquisition and higher rates of repeat purchasing among existing customers in the direct channels in the long term. Our results are based on intervention analysis with a treatment/control group design. We achieve greater balance between the groups by matching zip codes in the treatment and control regions; these procedures have been developed by scholars in other fields to approximate datasets that would have resulted from random experimentation.

KEYWORDS: Multichannel Retailing, Channels of Distribution, Direct Marketing,

E-commerce, Intervention Analysis, ARIMA, Time Series Models

ADDING BRICKS TO CLICKS:
THE EFFECTS OF STORE OPENINGS ON SALES THROUGH DIRECT CHANNELS

1. Introduction

Multichannel marketing is of increasing interest as new technologies open new paths to market for brands. This paper investigates what happens to sales through existing channels when new channels are added. Aggregate sales tend to increase, all other things being equal, when a brand adds a channel because the new channel provides access to new customers. However, sales through any of the existing channels may increase or decrease because the new channel has two effects. It advertises the brand, giving buyers more opportunity to become aware of it and experience it, but it also cannibalizes sales through existing channels, and the aggregate increase in sales may not be so great as to offset the cannibalization. Thus new channels complement, but also compete with, existing channels. Whether the net effect of new on existing channels is complementary or cannibalizing depends on a number of contingencies, among them the type of channel, the type of customer, the passing of time, and the maturity of the brand and, therefore, its responsiveness to the advertising effect.

Understanding the effects of these contingencies in field contexts must be pieced together as field data sets become available. This article reports on the analysis of a data set recording sales of a national retailer with a catalog and on-line retail presence as well as brick-and-mortar stores. It analyzes, for four new store openings in a three year period, whether each store cannibalizes sales from direct-to-consumer channels in the new store's trading area, or has a complementary effect. The data set also classifies customers according to whether they are first-time or repeat customers of the store and the direct channels, allowing us to analyze whether cannibalization is due to customers switching from the old channels to the new or reducing their order size, and whether complementarity is due to improved customer acquisition or better retention.

This analysis contributes to the cumulative understanding of the contingencies governing new channel effects on existing channel sales in several ways. First, it is one of the few studies to be able to explore the contingency of channel type, and to say whether online and catalog channels respond differently to expansion of channels. Second, because the replications varied with respect to whether the brand had previously existed in the trading areas, it investigates the brand maturity contingency. Third, our data allow us to study how complementarity and cannibalization vary contingent on the passing of time, which earlier studies have not done.

We organize the paper as follows. We begin by describing existing research on the contingencies affecting sales response to channel expansion. Then, we describe our dataset and the matching method we use to account for the nonrandom store location selection decision. Next, we describe the modeling approach we use to understand the impact of a retail store opening on direct sales over time. We use intervention time series analysis (Hanssens, Parsons and Schultz 2001) in an autoregressive integrated moving average framework (ARIMA). Our first series of analyses models the effect of store openings on both catalog and online sales series over time and uses an innovative test/control group experimental design to control for extraneous variables affecting the sales time series. Our second series of analyses models the effect of store openings on the number of new and existing customers using direct channels. We conclude with a discussion of the implications of our work for research on multichannel retailing and the managerial implications for channel management strategy.

2. Conceptual Background

Past research suggests two consequences from the addition of a new retail store to an existing pattern of channels – cannibalization and complementarity. In this section, we examine the theory and empirical evidence for each effect. We explore theorized substitution relationships among retail store, catalog, and online channels and conjecture that asymmetries among channel response may occur.

2.1 The Cannibalization Hypothesis.

Cannibalization views channels as competing against one other for customers. This perspective dominates channel management research, perhaps due to a historical emphasis on studying channel expansion as a between-firm construct rather than as a within-firm construct as we study it here. For example, researchers have measured the cannibalization effect of mass merchandisers on grocery stores (Fox, Montgomery and Lodish 2002) and game theorists have designed models which capture the detrimental impact the practice of manufacturers selling direct to consumers has on their reseller channel partners, and models which capture competitive dynamics between retail firms and their direct marketer competitors (Balasubramanian 1998; Zettelmeyer 2000; Chiang, Chhajed and Hess 2003).

Past work has suggested some conditions which favor cannibalization. First, cannibalization occurs when two channels too closely duplicate each other and do not provide adequate product and/or service differentiation between the channels (Deleersnyder, Geyskens, Gielens and Dekimpe 2002). Second, cannibalization occurs more frequently for products where delivery costs are low, such as information products (Shapiro and Varian 1999). Third, cannibalization occurs when the channels target the same consumers (Deleersnyder et al. 2002). Finally, when handling the product is important, a retail store may cannibalize direct sales because customers using direct channels cannot experience the product first hand prior to purchase (Coughlan, Anderson, Stern and El-Ansary 2001).

2.2 Complementary Hypotheses.

Other research has suggested conditions that favor channel complementarity. The first is heterogeneity across consumers. Alba et al.(1997) argue that consumers have heterogeneous purchase preferences which affect their choice of channel; for example, a consumer who is housebound will value the ability to order from an online or catalog channel, while a consumer who needs to use a product right away will prefer to buy it in a retail store rather than wait for shipment. Hence, a new channel can complement existing channels if it serves customers who were previously not served (Moriarity and Moran 1990).

The second is heterogeneity within a customer but across purchase occasions. If a customer wants a store for some types of purchases and a catalog for others, channels will be complementary. Studies have shown that retailers who offer multiple channels to consumers command higher levels of loyalty among their customers (Shankar, Smith and Rangaswamy 2003; Wallace, Giese and Johnson 2004). These multichannel shoppers (Reda 2002) "are combining various channels and approaches, searching online to buy offline, searching offline to buy online," (Wind and Mahajan 2002, p. 65), increasing their interactions with the retailer and obtaining better service. In business-to-business settings, multichannel shoppers have been shown to be more loyal and more profitable than single channel shoppers (Kumar and Venkatesan 2005).

Third, there is an advertising effect. Adding a new channel may increase brand awareness as consumers are exposed to more local advertising and promotion communications. This increased brand awareness may lead to increased sales in existing channels even if the new channel is not utilized. Two studies of multichannel shoppers (Ansari, Mela and Neslin 2005; Kumar and Venkatesan 2005) show that they are exposed to more marketing communications from the firm than single channel shoppers. Additionally, valuable brand associations attributed to the distribution channel may transfer to the brand (Jacoby and Mazursky 1984; Keller 1993). For instance, consumers perceive that retailers with an online presence offer lower prices (Cotlier 2001) and that retailers with a local retail store presence are more trustworthy on the online than retailers who sell only over the online, perhaps making website or catalog sales seem less risky and more reliable (Tang and Xing 2001).

2.3 Relative Magnitudes of Cannibalistic and Complementary Effects.

A few empirical studies have directly tested the effects of channel expansion on sales and customer acquisition and retention. In a study of the impact on stock prices of the decision of European newspaper companies to add an online edition to a paper edition, both demand expansion (via market expansion, brand switching, and relationship deepening), and demand reduction (via channel shift without lift, loss of impulse purchases, and lack of service support) were identified (Geyskens, Gielens and Dekimpe 2002). Although, on average, the net effect of adding the online channel was positive, more than 30% of the time, expanding distribution to an online channel resulted in negative stock returns. A follow-up study (Deleersnyder et al. 2002) found some evidence of cannibalization in a small number of cases (7% of newspapers experienced drops in circulation and 6% experienced drops in advertising revenues), but, on average, publishers did not experience cannibalization of their traditional distribution channels by the online channel. However, complementary effects were not robustly observed either, with only 14% of firms showing small, but positive, circulation growth rate changes following the launch of the online channel and no firms showing any change in the growth of advertising revenues. An empirical analysis of Tower Records' incursion into online distribution yielded similar results (Biyalogorsky and Naik 2003). Negligible cannibalization effects on the retail store business were found. However, this study featured a one-period design, with cannibalization measured as a change in retail store sales within the same period as the online sale, so longer term or lagged cannibalization effects were not explored. The time series was twelve months, leaving the longer term sales response unknown.

In a study of the sales impact of opening physical retail stores on a firm's existing catalog and online direct-to-consumer sales, Pauwels and Neslin (2006) uncover cannibalization of the catalog channel, but find that the online channel was unaffected by the opening of a retail store, so that cannibalization effects were not equal across different channels. The empirical data used in this study came from a durables and apparel retailer whose sales came mainly from direct-to-consumer channels. Following the opening of the retail store, catalog order sizes decreased by 14% and existing catalog customers purchased less frequently. The authors partially explain the cannibalization of catalog sales as accelerated by the firm's actions in the catalog channel following the opening of the retail store: when the sales of catalog customers in the store region initially dipped following the store's opening, the company's RFM models, which calculate whether to send a customer a catalog based on their purchase recency, frequency, and monetary value to the firm, began to send less catalogs to these customers, furthering the cannibalization effect. The authors conclude with the implication that retail stores are closer substitutes for catalog channels rather than online channels due to the fact that consumers with similar demographics shop in stores and in catalogs.

Online consumer survey data were used by Ward and Morganosky (2000) to explore cannibalization between traditional retail store channels and online channels across a wide array of product categories, including computer hardware and software, home electronics, food, investment services, music, books, travel, and apparel. They found that consumers who searched for product information online increased their purchases in retail stores. However, the reverse did not hold; searching for product information in offline channels did not increase purchases online, indicating a potential asymmetry in complementarity. This finding is supported by Verhoef, Neslin, and Vroomen (2005) who calculate differential "lock-in" rates for various channels, a channel's ability to capture a sale following product search in that channel, and show that retail stores have higher lock-in rates than online websites, where many consumers search for information and then leave to purchase in an offline channel. Ward (2001) found that direct channels, such as catalogs and websites, are close substitutes for each other and tend to cannibalize each other, while direct channels are less likely to cannibalize sales from retail stores, a finding seemingly at odds with the results of the Pauwels and Neslin (2006) study.

2.4 Conjectures and Contributions

Nothing in past research, conceptually or empirically, rules out the possibility that new channels might both compete with and complement existing ones. In this paper, we explore four contingencies which help explain when cannibalization or complementarity effects dominate. We hypothesize that store openings create both types of effects and that the emergence of and strength of each effects varies over time and across type of channel, type of customer, and type of market retail conditions. Our specific conjectures relate to these contingencies that govern the relative importance of cannibalization and complementarity and are outlined below.

Time Effects

Although some studies have found evidence that new channels compete with existing ones and others have found that new channels complement them, previous research has not examined how inter-channel relationships play out over time. We hypothesize that cannibalization and complementary effects will vary over time because some changes to consumer behavior are likely to occur immediately while others are likely to take time to manifest. More specifically, we hypothesize that cannibalization will be the dominant short-term effect as some existing customers try the new channel by shifting either all or some of their purchases to the retail store to ascertain if it better suits their needs. Studies have shown that consumers will shift their purchases when a new channel offers them service or convenience features not previously available to them in existing channels (Coughlan et al. 2001). In our retail setting, experience goods make up a significant portion of the retailer's product offerings; hence, existing direct channel customers are likely to shift some of their purchases to the retail store in order to physically experience the products prior to purchase.

While the short-term effects are likely to be harmful to the direct channels, we hypothesize that the store opening will complement the direct channels over time. First, the store acts as an advertisement for the retailer's brand. This advertising creates and strengthens brand awareness and associations in the retail trading area surrounding the store, attracting new customers to both the store and to the direct channels, and reminding existing direct channel customers to purchase from the retailer via the catalog and online channels. Second, the physical presence of a store in a market reassures customers previously hesitant to patronize direct channels or to use direct channels for certain types of purchases. Third, the opening of the store creates additional points of contact with existing customers, which has been shown to increase customer satisfaction and loyalty as cited above. Branding and loyalty effects take time to build; hence, we propose that the positive impact of these effects on direct channel sales will be more strongly felt over the long term than in the short term.

Thus, we hypothesize that cannibalistic effects will be felt immediately and will dominate in the short term, while complementary effects will be felt gradually and cumulatively over time and will dominate in the longer term.

Channel Effects

Second, we hypothesize that the magnitude of cannibalizing and complementary effects will differ across channels. Previous research suggests that retail stores will have an asymmetric cannibalizing effect on the catalog and Internet channels, where the opening of a retail store will cannibalize catalog sales, but not Internet sales, due to shared demographics among store and catalog shoppers and different demographics among store and Internet shoppers. Following this, we hypothesize that cannibalization effects will be larger in the catalog channel than in the Internet channel following the opening of our retail stores.

However, we also hypothesize that the store channel will asymmetrically complement the two direct channels, and, specifically, that the opening of a retail store will complement the online channel to a greater degree than it will the catalog channel, for reasons that relate to the way demand is generated in each channel. The catalog sales channel may be characterized as a discriminating channel in the sense that catalogs are mailed only to addresses that meet particular criteria. The online sales channel is a non-discriminating channel, accessible by any consumer who knows of its existence. Traffic on an online channel is more broadly and immediately responsive to traditional brand advertising or search engine advertising than traffic on the catalog channel, which is narrower (because it is limited to addresses receiving the catalog) and lagged (because catalogs are mailed in waves.) We argue that the pattern of demand in the online channel is more similar to that of the retail store channel, another non-discriminating channel, and, therefore, the cross-channel interaction between the two may be intensified, a theory in contrast to previous theories in which the online and catalog are viewed as substitutes for each other (Ward 2001). The opening of a retail store will act as a brand-building advertisement for both the catalog and online channels, driving more consumers to seek out the retailer on more purchase occasions; however this branding effect more quickly, strongly, and directly impacts sales through the online channel, given that online channels are always open for business and open to everyone, whereas catalog channels are only open to select consumers only once the retailer has identified them as a suitable prospect and only after the retailer has mailed them a catalog.

Thus, we hypothesize that a retail store will cannibalize catalog sales more than online sales <u>and</u> that it will complement online sales more quickly and to a greater extent than catalog sales; hence, online sales will show shallower decreases and faster and greater recoveries following store openings than catalog sales.

Customer Effects

Third, we conjecture that first time and repeat customers will be differentially impacted by the opening of a retail store. While the branding effects which contribute to complementary effects outlined above will impact both first time and repeat customers, by enticing first time customers to try the direct channels and by reminding repeat customers to purchase via the catalog and Internet, short term cannibalization is likely to occur only with existing customers of direct channels and not with potential first time customers (i.e. prospects).

Hence, we hypothesize that the direct channel will initially lose some of its repeat customers to the store, but that after the store opening the remaining repeat customers will increase their purchasing through the direct channel over time due to branding and multichannel loyalty effects. Over the longer term, we hypothesize that some portion of existing direct channel customers who initially defect to the store channel are likely to return to the direct channels for some of their purchasing after trying the store. We also hypothesize that first time customers will be acquired at a faster rate due to branding effects, hence purchases from first time users of the direct channels will increase faster in the trading area of a new store relative to the rate of increase in markets served by direct channels that do not have a store opening.

Market Effects

Lastly, past research has not explored whether the relative magnitude of cannibalizing and complementary effects in a market depends on the strength of the brand in the market. In particular, if a store is opened in a retail trading area where other stores of the same brand already exist, we hypothesize that both the cannibalizing and the complementary effects will be weaker than if the new store opens in a retail trading area not previously served by the brand. The short term cannibalization effect will be weaker because direct channel customers will have already had the opportunity to visit the store. The longer term complementary effects due to store advertising will be weaker because familiarity with the brand will already exist in the retail trading area due to the presence of the other existing retail store and its associated branding efforts.

3. Data Description

We exploit a natural experiment, the opening of four new retail stores in a single U.S. state, to observe and analyze the effect of retail store openings on direct sales. We use proprietary, longitudinal transaction data from a multichannel retailer of apparel, accessories, and home furnishings. The retailer operates stores in shopping malls and also sells to consumers through a direct channel that combines a well established direct mail catalog operation with newer online website operations, which began during the period of our study. The retail store channel and the direct channels both draw on the same merchandise selection and use the same price points for regular ticket pricing; however specific merchandise promotion, price promotion, advertising and communication decisions are made at the operating unit level, and day-to-day operations are largely separate across the two channels. Overall, sales from retail stores have been significantly higher than sales from direct channels, but growth in direct channels has been dramatic over the last decade, as it has been for many retailers.

Our data consists of monthly direct sales, aggregated by zip code, for transactions in a single U.S. state. Using data aggregated at the zip code level has three benefits for our analysis. First, it preserves the privacy of the firm's individual customers and allows us access to customer data that was not previously available to researchers. Second, because few individual consumers purchase from this retailer at least once in every month, aggregating to the zip code level generates a continuous sales variable over time which allows us to use time-series based longitudinal modeling approaches. Time series analyses have generated considerable insight into persistent, long term consumer response to strategic marketing changes (Dekimpe and Hanssens 1995; Pauwels and Srinivasan 2004). Finally, because demographic information is widely available from third party sources at the zip code level, we can account for consumer heterogeneity across units and over time (Steenburgh, Ainslie and Engebretson 2003).

Our outcomes of interest are net catalog sales and net online sales from zip codes in a retail trading area. These are defined as sales, net of returns, generated by all direct-to-consumer media (i.e. catalog mailings, email campaigns) and transacted by telephone and online on the firm's website.

3.1 Treatment Events

The retailer opened four stores over the time period of our study which serve as the treatment events. Two of the stores opened in retail trading areas which were previously served by the retailer only through direct channels, while the other two stores opened in retail trading areas previously served by the retailer through both direct channels and retail stores. The existing stores in these retail trading areas had been open for more than five years prior to our observation period. The retail trading areas receiving the new retail store opening treatment are identified in the table below:

Retail Trading Area	Year Opened	Existing Stores
Store A	Fall 2000	Yes
Store B	Fall 2001	No
Store C	Fall 2002	Yes
Store D	Fall 2002	No

We use 36 months of sales data before each of the four store opening events (34 months in Store A due to lack of data) and we bring each post-store opening time series through December 2005, resulting in a time series of 96 months for Store A, 88 months for Store B, and 75 months for Stores C and D. Hence, depending on the store region, our data covers the time period from 1998-2005.

4. Model Development

4.1 Intervention Analysis of the Direct Sales Time Series

Given that we were interested in observing the effect of the store opening events on direct sales, we used intervention time series analysis (Box and Tiao 1975) in an autoregressive integrated moving average framework (ARIMA) to model the effect of store openings on both catalog and online sales over time. This allows us to test whether changes in the time series occur and then to specify the nature, timing, and magnitude of the changes.

It is common in intervention analysis to predict future observations from a pre-intervention trend. A downside of this approach is that it cannot control for events which coincide with the occurrence of the intervention and which may interfere with its effects. For example, an economic recession occurring at the time of a store opening could attenuate the store opening's effects. A quasi-experimental design, however, circumvents this problem by introducing a control group to control for such unrelated, yet temporally correlated, events (Cook and Campbell 1979).

With this in mind, we build upon a methodology developed by Krishnamurthi, Narayan, and Raj (1986) which combines traditional intervention analysis with a quasi-experimental design. In this methodological design, a control time series is used to approximate the path the sales time series would have taken if the intervention had not occurred, rather than estimating it from a pre-intervention trend. Hence, we compare net direct sales in store retail trading areas to net direct sales in control retail trading areas where the retailer had not yet opened stores. This type of comparative approach strengthens the internal validity of our study by ruling out alternative explanations and establishes a baseline of sales so that economic effect sizes resulting from the intervention may be measured.

The internal validity of quasi-experiments improve as treatment and control groups become more comparable (Cook and Campbell 1979). Thus, in a perfect world, we would randomly assign store openings to different retail trading areas which would ensure that the effect of the store opening was independent of other factors influencing direct sales. This is not possible, of course, so our analysis needed to address the potential for selection bias due to nonrandom treatment (store opening location) assignment. Scholars in other fields, such as sociology, finance, political science, economics, and epidemiology, have addressed this problem through matching (Jaffe, Trajtenberg and Henderson 1993; Meyer 1995; Heckman, Ichimura, Smith and Todd 1998; Winship and Morgan 1999; Lee and Wahal 2004; Ho, Imai, King and Stuart 2005).

In this paper, we introduce to marketing several matching procedures (Ho, Imai, King and Stuart 2004; 2007) that are used in other fields to create datasets which approximate ones that would result from random experiments. These procedures, described in greater detail in the next section, are a preprocessing step taken to control for variables other than the intervention of interest that may differentially impact treatment and control groups. The addition of this step enhances the internal validity of our intervention analysis and reduces model dependence.

The final step in the analysis was to identify a form for the response to the intervention, which has to be specified in advance by the researcher (Hanssens et al. 2001). Generally, the form is identified based on theory or prior empirical results about similar phenomena. The existing literature on channel expansion points to several different potential forms, from an immediate and long-term drop suggested by the cannibalization literature, to a gradual and long-term increase suggested by the complementary literature. Given that prior research has found evidence of both cannibalizing and complementary effects and given our hypothesis that the emergence and dominance of the effects may be time dependent, we chose to model the response to capture a short term intercept change in level and a longer term change in the trend and then inspect the fit of the model versus the true series for evidence of a more nuanced response form.

Given these considerations, our model was specified as:

$$(1)\quad treatment_t = \alpha + \beta_1 pre.open.months_t + \beta_2 store.open_t + \beta_3 post.open.months_t + \beta_4 control_t + N_t$$

where *treatment* represents the aggregate monthly sales across zip codes in the retail trading area in which the store opened and *control* represents the weighted aggregate monthly sales across the zip codes in the matched control retail area. The *store.open* variable is the step function indicator for the store-opening intervention (taking a value of zero prior to the store opening and one starting in the month that the store opens), *pre.open.months* measures the number of months from the observed month to the store opening month (-36 to 0), *post.open.months* measures the number of months from the store opening month to the observed month (store opening month through December 2005). N_t is assumed to follow an ARIMA(p, d, q) model. Separate models were run for catalog channel sales and for online channel sales.

In these models, the coefficients of *store.open* and *post.open.months* identify the general nature of the response over time. The *store.open* coefficient β_2 measures the immediate shift in the sales mean that occurs when the store opens; a significant β_2 coefficient supports a store opening effect that causes at least a short-term change in level and estimates the short-term impact of the sales response. The *post.open.months* coefficient β_3 measures changes in the underlying sales trend after the store opens; a significant coefficient supports a response to the store opening which changes over time. The *pre.open.months* coefficient β_1 measures an underlying time trend in the store series prior to the store opening; a significant coefficient suggests that the retail trading areas surrounding the new stores differ from the control retail trading areas on our outcome variable of interest prior to the opening of the store and hence, that our control series is not accounting for an important unobserved component.

4.2 Constructing the treatment and control retail trading areas.

We construct our treatment and control retail trading areas by using a preprocessing matching procedure to select control zip codes that match zip codes in our treatment retail trading areas on demographic and geographic characteristics. Matching procedures allow us to construct control groups that closely match our treatment groups on key variables that we believe impact the choice of a retail store location and direct sales levels. Matching protocols are designed to address selection bias in observational studies like this one in which the treatment variable is observed, rather than manipulated by the researcher.

We generated datasets for each of the four stores by matching zip codes in the retail trading area surrounding the new store location (the treatment group) with zip codes from the rest of the state (the control group). To select the control zip codes, we first identified three metropolitan areas in the state which were broadly similar in population and income to the metropolitan areas where the new stores were opened, but where the firm had not yet opened stores. In each of these areas, we identified a control shopping mall that was most similar to the malls in which the new stores opened, looking specifically for direct competitors' stores.

To assign zip codes to a treatment or control retail trading area, we generated a drive-time variable from each zip code to each shopping mall. We assigned each zip code to either a treatment or control retail trading area based on the drive-time to the nearest store, using a maximum drive-time cutoff of sixty minutes for each retail trading area. This cutoff represents a reasonable maximum drive-time from which a shopping mall would draw. For Stores A and C, where there were also existing stores within a 60 minute drive, we considered only those zip codes where the new store was the closer store. We assigned all zip codes within sixty minutes of any of the control shopping malls to a single control retail trading area. This resulted in a pool of 743 zip codes: 551 in the four store retail trading areas and 192 in the control region. (These numbers reflect the elimination of seven zip codes, two in the treatment regions and five in the control region, for which demographic data was not available to assist with the matching procedure.)

Following the nonparametric preprocessing procedures outlined in Ho, Imai, King and Stuart (2007), we matched treatment zip codes to control zip codes for each of the four store regions. We matched treatment and control zip codes based on their joint similarity on multiple factors hypothesized to affect sales: drive-time to the closest store, a five year population level average based on annual population levels for 2000 to 2005, a five year median income level based on annual median income for 2000 to 2005, the compound annual growth rates for population and median income from 2000 to 2005, a two year median age based on median age for 2000 and 2002, the percentage of households with online access in 2002, the percentage of adults who purchased goods online in the past year in 2002, and the percentage of adults who purchased goods via a catalog in the past year in 2002. The latter two variables measured online and catalog purchasing in general and not specific to the retailer of this study. We evaluated five different matching algorithms with the goal of constructing a balanced sample: to generate common distributions across the treatment and control zip codes for the nine demographic and geographic variables that described the zip codes.

The five matching algorithms we evaluated were derived from the MatchIt software program (Ho et al. 2004). The "sub-classification" algorithm forms subclasses such the distribution of covariates for the treated and control groups are as similar as possible in each. The "nearest neighbor" algorithm selects the best control zip code matches for each zip code in the treatment group using a distance measure. A control zip code is matched to each treatment zip code one at a time, and, at each step in the matching process, the control zip code that has not yet been matched and is closest to the treatment zip code is chosen. Unlike the "nearest neighbor" algorithm, which chooses the closest control zip code for each treatment zip code one at a time, the "optimal matching" algorithm finds the matched treatment and control samples with the smallest average absolute distance across all of the matched pairs. The "full" matching algorithm (Rosenbaum 2002; Hansen 2004) delivers a fully matched sample in which matched sets (consisting of one treatment zip code and one or more control zip codes) minimize a weighted average of the estimated distance measure between each treatment zip code and each control zip code. Finally, the "genetic" matching algorithm (Abadie and Imbens 2004; Diamond and Sekhon 2005) automates the process of achieving an optimal balance between the treatment and control zip codes by matching with replacement, searching for a set of weights for each covariate such that optimal balance is achieved after matching.

We used several numerical and graphical methods to assess the balance of our matches derived from the five algorithms. First, within each matched sample, we evaluated the standardized bias (the difference in the means of the treatment and the control groups divided by the treatment group's standard deviation) of each geo-demographic variable for each of the five match sets. Second, we constructed quantile-quantile plots for visual inspection to assess the distributions of each geo-demographic variable for each of the five match sets. The quantile-quantile plots allow us to identify deviations between the treatment and the control group in any part of their distributions by plotting the quantiles of the treatment group and the quantiles of the control group for a particular variable in a square plot.

Following our assessment of the five matching algorithms, we determined that the "genetic" algorithm produced the best balance between the treatment and control groups. The results of the matching tests for the genetic match are contained in Table 1. Both the standardized bias summary statistics as well as a visual inspection of the quantile-quantile plots show the greatest improvement in balance between the treatment and control groups. As a rule of thumb, the literature suggests "good" matches generally produce standardized biases less than 0.25, indicating that the means of the treatment and the control are less than a quarter of a standard deviation apart (Ho et al. 2004; 2007). Across our nine geo-demographic matching variables in each of the four store/control groups, all standardized bias statistics are less than 0.25, with the exception of the age variable in Store C. Furthermore, even in this one case, visual inspection of the quantile-quantile plot for the age variable in Store C, contained in Figure 1, shows that our genetic matching procedure greatly improves the balance between our treatment and control groups, leaving us satisfied with the genetic matching results. The standardized bias summary statistics for the other matching procedures are reported in Table 2. While these procedures often produced good matches, the standardized biases were greater than 0.25 more frequently and were larger on the whole under these procedures than they were under the genetic matching procedure. Hence, we went forward using the genetic matching results.

5. Direct Sales Time Series Model Estimation

We followed the Box-Jenkins three-step modeling procedure to specify the ARIMA processes (Box and Tiao 1975; Box and Jenkins 1976). First, we tested for evolution versus stationarity of the store and control series. We used the Phillips-Perron unit root test (Phillips and Perron 1988) to test for stationarity of the store and the control series in each of the samples. This test is more robust to heteroskedasticity in the error term than the more well-known augmented Dickey Fuller test (Dickey and Fuller 1979) and has been used previously in marketing applications (Pauwels and Srinivasan 2004). The unit root tests showed all time series to be stationary at a 0.01 significance level except for the time series for the treatment group's online sales for Stores A, B, C, and D (all were not stationary at a .05 significance level; however, Store B was directionally stationary at a .08 significance level). We then ran Phillips-Ouliaris cointegration tests (Phillips and Ouliaris 1990) of each store series and its corresponding control series and found them to be cointegrated at a 0.01 significance level; hence, the linear combination of them contained in our model makes their difference stationary, allowing us to substitute a value of 0 for the trend component (d term) in the ARIMA models. Results of the stationarity and cointegration analyses are contained in Table 3.

We then examined the autocorrelation (ACF) and partial autocorrelation function plots (PACF) to identify patterns of autocorrelation and moving averages. We used the ACF and PACF plots to estimate values for the autoregressive (p term) and moving average components (q term) of the ARIMA models. Based on our observations of the plots, we estimated a series of ARIMA models which varied the values of p and q from 0 to 2 and then selected the best fitting model for each analysis, assessing goodness of fit using log likelihood and choosing the model with the lowest Akaike Information Criterion (AIC) value (1974; 1981). We then ran the ARIMA models and conducted Ljung-Box tests (Ljung and Box 1978) on the residuals of the models to determine if our specified ARIMA models left any systematic variation remaining. Systematic variation was not observed in the residuals of any of our models ($p > .10$).

We first ran an analysis combining the sales from all four stores in order to increase the power of our test, given our limited time series. For this "All Store" analysis, we stacked the data from each of the four retail trading areas and used dummy variables for each of the stores to capture store fixed effects. Then, we ran separate sales analyses for each store, where the trend variables measure the number of months from the observed month to the store opening month, centered on the store opening date and running through December 2005 (from -34 months to +61 months in Store A, from -36 months to +51 months in Store B, and from -36 to +38 months in Stores C and D).

5.1"All Stores" sales model time series results.

Catalog Sales. Table 4 shows the results of our "All Stores" catalog sales model aggregated across the four stores. First, the *pre.open.months* coefficient β_1 was not significant in the model, providing support for the adequacy of the matches after accounting for the effects of the store openings (β_1 = 15, p = .9455). This non-significant coefficient supports our assumption that the retail stores were not placed into retail trading areas that had higher (or lower) catalog sales than control retail trading areas. The *store.open* coefficient β_2, indicating a short term change in sales level following a store opening, showed a significant short term drop in catalog sales in the store retail trading areas (β_2 = -12,947, p = .0193). The *post.open.months* coefficient β_3, identifying differences between the store series and the control series sales trends following the store openings, showed a significant increase in the catalog sales trend in store retail trading areas relative to the control retail trading areas following the opening of the stores (β_3 = 395, p = .0013). These coefficients equate to a significant short-term drop in catalog sales of approximately 12.1% followed by a return at approximately a 0.4% per month growth rate to the previous catalog sales levels and beyond; hence, it takes, on average, 33 months for catalog sales to recover short term cannibalization and begin to exhibit incremental complementary effects which lead to net increases in catalog sales following the introduction of a retail store.

Online Sales. Table 4 shows the results of our aggregate online sales model across the four stores. First, the *pre.open.months* coefficient β_1 was not significant in the full model, providing support for the adequacy of the matches after accounting for the effects of the store openings (β_1 = 252, p = .7419). The *store.open* coefficient β_2 was not significant (β_2 = -1,610, p = .8147). The *post.open.months* coefficient β_3 identified significant increases in the online sales trend in store retail zip codes relative to the control zip codes following the opening of the stores (β_3 = 2,360, p < .0001). These coefficients equate to an insignificant short-term drop in online sales followed by a significant increase in the sales growth trend at approximately a 34% per month growth rate. Hence, online sales show much larger complementary effects than catalog sales and do not suffer from significant levels of cannibalization as the catalog channel sales do, indicating an asymmetrical response between the two direct channels to the opening of a retail store. Even if we were to ignore the statistical insignificance of the online sales short-term drop for a moment, one can see that the sales recovery in the online channel is much faster than in the catalog channel, bringing online sales back to pre-store levels within one month of the store opening (compared to 33 months in the catalog channel) and then quickly moving sales levels beyond the control baseline for the months following.

5.2 "Store by Store" sales time series results.

Catalog Sales. Table 5 shows the results for catalog sales in each of the four retail trading areas. First, the *pre.open.months* coefficient β_1 was not significant in three of the four store models, providing support for the adequacy of the matches after accounting for the effects of the store opening. In Store D, a negative *pre.open.months* coefficient was evident (β_1 = -196, p = .0759), but this coefficient suggests that sales in the Store D retail trading area were growing at a slower rate than the control group sales prior to the store opening, which should make it harder for us to find a post-store opening complementary effect. (This also was the only significant *pre.open.months* coefficient in the study.) The *store.open* coefficient β_2 showed a significant drop in catalog sales in one of the four stores' retail trading areas; the other three stores exhibited no significant short term drop in catalog sales (Store A β_2 = -14,280, p = .1320, Store B β_2 = -25,684, p = .0033; Store C β_2 = -6,165, p = .4438; and Store D β_2 = -5,640, p = .0806). The *post.open.months* coefficient β_3 identified significant increases in the catalog sales trends in two of the four stores' retail trading areas relative to the control retail trading areas (Store A: β_3 = 450, p = .0046); Store B: β_3 = 517, p = .0062; Store C: β_3 = 100, p = .6908; and Store D: β_3 = 129, p = .1928.)

Online Sales. Table 5 shows the results for online sales in each of the four retail trading areas. First, the *pre.open.months* coefficient β_1 was not significant in any of the four store models, providing support for the adequacy of the matches after accounting for the effects of the store opening. The *store.open* coefficient β_2 showed a directional short term drop in online sales in only one of the four stores' retail trading area; all other store retail trading areas showed no significant short-term drop in online sales (Store A β_2 = -6,657, p = .4956; Store B β_2 = -5,775, p = .0771; Store C β_2 = -11,660, p = .4396; and Store D β_2 = 832, p = .8915). The *post.open.months* coefficient β_3 identified significant increases in the sales trends in all of the four store retail trading areas relative to the control retail trading areas (Store A: β_3 = 977, p < .0001; Store B: β_3 = 794, p < .0001; Store C: β_3 = 5,853, p = .0115; and Store D: β_3 = 1,455, p = .0117). Small, mostly insignificant drops in sales followed by immediate build ups of sales beyond pre-store opening levels over time indicate the complementary effects of a retail store opening for the online channel. Comparisons of the online results with the catalog results suggest that complementary effects exert themselves much stronger and more quickly in the online channel than in the catalog channel.

5.3 Discussion.

Adding a new retail store channel to existing direct sales channels increases firm sales in the long run, as sales from new retail stores are incremental to sales from direct channels, which show no long-term damage and significant increased growth rates from intra-firm channel competition. Our results illustrate that the addition of a new retail store channel to an established direct sales channel has a short term cannibalistic effect on catalog sales, but not online sales, in the retail trading area surrounding the retail store. However, this cannibalistic effect is short-lived and does not have lingering temporal effects over the longer term as complementary effects begin to take effect to bring catalog sales levels higher than pre-store opening levels. Catalog sales in a retail trading area fall immediately following a store opening, but slowly recover back towards and beyond the pre-store opening sales levels over the next 33 months, on average, as complementary effects begin to manifest themselves. Online sales exhibit no significant short term cannibalizing effects from the opening of a retail store and show longer term complementary effects which are larger in magnitude than those seen in the catalog channel. Hence, the retail store serves as a substitute for the catalog channel in the short, but not the long, term and as a complement for both the catalog and online channels in the long term, albeit to a greater extent for the online channel.

Two competing theoretical explanations can describe this direct sales response pattern. First, the opening of the retail store may induce existing catalog customers to try shopping at the store. During this trial period, these shoppers reduce their expenditures in the catalog channel and shift their sales to the retail store channel. However, over time, these shoppers eventually work their way back to the catalog channel from which they came. This explanation would indicate a short term cannibalistic effect for catalog sales which decays over time. The second explanation is that cannibalistic and complementary effects are both in play in the catalog channel following the opening of a retail store and that the sales response we model is a combination of the two effects working together. For example, the store opening causes an immediate cannibalistic effect, which endures over time. The opening of the retail store may induce existing catalog customers to try shopping at the store and their expenditures previously spent in the catalog channel shift to the retail store channel for the remainder of their duration as a customer. Hence, these customers are lost to the catalog channel. However, simultaneously, a complementary effect is occurring with a different set of customers due to the branding effect a new retail store brings which both increases brand awareness and improves brand associations of the previously direct-to-consumer retailer. The opening of the retail store attracts new customers to the direct channels who did not previously purchase via direct before due to a new awareness or appreciation for the retailer. Over the long term, this new influx of shoppers contributes incremental sales to both direct channels which compensate for the loss of sales from existing catalog customers who migrate to the retail store channel. In order to test which of these two competing theories best explains our results, we conducted an empirical analysis of customer development in the direct channels at the household level. The model developed for this analysis and its results are discussed below.

6. The Customer Development Model

6.1 Intervention analysis of the customer count time series.

Following the same intervention analysis with a control series modeling approach we used with the sales data, we ran an intervention analysis model on customer data, substituting the number of customer households purchasing during each month through online and catalog channels for channel sales in the previous model. Again, we chose to model the response to capture a short term intercept change in level and a longer term change in the trend and then inspect the fit of the model versus the true series for evidence of a more nuanced response form.

Following the same procedure as before, we first used the Phillips-Perron unit root test and the Phillips-Ouliaris cointegration tests to assess the stationarity of the store and control customer count series and their cointegration levels. The unit root tests showed all time series to be stationary at a 0.01 significance level except for the time series for the treatment group's repeat customers for Stores C and D ($p = .23$). We then ran Phillips-Ouliaris cointegration tests of each store series and its corresponding control series and found Store C to be cointegrated at a 0.01 significance level and Store D at a .02 significance level; hence, the linear combination of them contained in our model makes their difference stationary, allowing us to substitute a value of 0 for the trend component (d term) in the customer count ARIMA models. Results of the stationarity and cointegration analyses are contained in Table 6.

We then examined the autocorrelation (ACF) and partial autocorrelation function plots (PACF) to estimate a series of ARIMA models which varied the values of p and q from 0 to 2 and then selected the best fitting model for each analysis, assessing goodness of fit using log likelihood and choosing the model with the lowest Akaike Information Criterion (AIC) value. We then ran the ARIMA models and conducted Ljung-Box tests on the residuals of the models to determine if our specified ARIMA models left any systematic variation remaining. Systematic variation was not observed in the residuals of any of our models ($p > .10$).

In the customer count time series model, *treatment* represents the aggregate number of customer households purchasing in the direct channels during each month across all zip codes in the retail trading area in which the store opened and *control* represents the weighted aggregate number of customer households purchasing during each month across the zip codes in the matched control retail area. Separate models were run for first-time customer households and for repeat customer households to determine if the cannibalization and complementary effects were being driven by the purchasing behaviors of new or existing customers. The number of first-time customer households acquired each month represents customer households who have not previously purchased from the direct channels in the past who make a purchase in a direct channel during that month. First-time customer households are only identified as "new" in the first month in which they make a purchase; their subsequent purchases in the time series appear in the repeat customer households data. Repeat customer households represent customer households who have previously purchased from direct channels in the past who make a purchase in a direct channels during that month. The customer count analysis compared customer households for the direct channels in aggregate as catalog and online channel breakdowns were not available.

In these models, the coefficients of the *store.open* and *post.open.months* variables identify the general nature of the response over time. The *store.open* coefficient β_2 measures the immediate shift in the mean of direct sales channel new customer household acquisition rates or repeat customer household retention rates that occur when the store opens. A significant β_2 coefficient supports a store opening effect that causes at least a short-term change in level, and estimates the short-term impact of the new customer household acquisition and repeat customer household retention response. The *post.open.months* coefficient β_3 measures changes in the underlying new customer household acquisition or repeat customer household trends in the direct channels after the store opens; a significant coefficient supports a response to the store opening which changes over time. The *pre.open.months* coefficient β_1 measures an underlying time trend in the store series prior to the store opening; a significant coefficient suggests that the retail trading areas surrounding the new stores differ from the control retail trading areas on our outcome variables of interest prior to the opening of the store and hence, that our control series is not accounting for an important unobserved component.

Again, we ran an aggregate "All Stores" customer count model first to maximize statistical power and then ran separate customer count models for each of the four stores.

6.2 "All Stores" customer count model time series results.

First-Time Customers. Table 7 shows the results of our "All Stores" customer count model. First, the *pre.open.months* coefficient β_1 was not significant in the full model, providing support for the adequacy of the matches after accounting for the effects of the store opening (β_1 = -0.44, p = .5327). The *store.open* coefficient β_2, indicating a short term change in the first-time customer count level following store openings, showed no significant change (β_2 = 1.32, p = .9179). This implies that new customers did not delay shopping through direct channels when the store opened. The *post.open.months* coefficient β_3, identifying differences in the monthly trends of the number of first-time customer households between the store series and the control series sales trends following the store openings, identified significant increases in the new customer acquisition trends in the store retail trading areas relative to the control retail trading areas (β_3 = 1.74, p = .0023). These coefficients equate to significant long term growth trends of approximately 1.3% per month in the number of first-time direct customer households acquired (relative to the control series) after the opening of retail stores. Hence, the opening of a store helps direct channels acquire new customers at a faster rate, perhaps due to increased brand awareness and positive brand associations the store brings to consumers in the retail trading area.

Repeat Customers. We used data on repeat customers to ascertain the store opening effect on the direct channels' existing customers. As explained above, the repeat customer count for each month after the store's opening includes 1.) customers who purchased that month who previously bought in direct channels prior to the store opening, as well as 2.) new customers from all of the post-store opening months prior to the month in question (i.e. a first-time customer in month +1 becomes a repeat customer in months +2, +3, +4…).

Table 7 shows the results for our aggregate model for the number of repeat customers purchasing in the direct channels. First, the *pre.open.months* coefficient β_1 was not significant in the full model, providing support for the adequacy of the matches after accounting for the effects of the store opening (β_1 = .2218, p = .8556). The *store.open* coefficient β_2, indicating a short term change in the repeat customer count level following the store openings, showed no significant change (β_2 = -22.15, p = .3339). The *post.open.months* coefficient β_3 identified significant increases in the repeat customer count trends in the store retail trading areas relative to the control retail trading areas (β_3 = 3.16 , p = .0012), indicating that more repeat customer households were purchasing in direct channels following the opening of a store.
These coefficients equate to significant long term growth trends of approximately 0.7% per month in the number of repeat customer households purchasing each month after the opening of retail stores in the retail trading areas into which stores were launched versus control retail trading areas.

6.3 Store by store customer count results.

First-Time Customers: Table 8 shows the results for our separate models for the four store retail trading areas for the number of first-time customer households acquired. First, the *pre.open.months* coefficient β_1 was not significant in any of the four store models, providing support for the adequacy of the matches after accounting for the effects of the store opening. The *store.open* coefficient β_2 did not show a significant drop in any of the four store retail trading areas (Store A $\beta_2 = 2.75$, p = .8335; Store B $\beta_2 = -10.09$, p = .3822; Store C $\beta_2 = 8.20$, p = .8138; and Store D $\beta_2 = -8.16$, p = .2679). The *post.open.months* coefficient β_3 identified significant (or directional) increases in the new customer acquisition trends in three of the four store retail trading areas relative to the control retail trading areas (Store A: $\beta_3 = .19$, p = .4115; Store B: $\beta_3 = 1.06$, p = .0001; Store C: $\beta_3 = 5.79$, p = .0667; and Store D: $\beta_3 = 1.56$, $p < .0001$). Hence, the opening of a new store does not appear to impact the acquisition of first-time customers in the month in which the store opens, but contributes to higher first-time customer household acquisition growth rates over the long term.

Repeat Customers: Table 8 shows the results for our model for the separate model for each of the four store retail trading areas for the number of repeat customers purchasing in the direct channels. First, the *pre.open.months* coefficient β_1 was not significant in any of the four store models, providing support for the adequacy of the matches after accounting for the effects of the store opening. The *store.open* coefficient β_2 showed a significant drop in only one of the four store retail trading areas (Store A $\beta_2 = -31.21$, p = .1688; Store B $\beta_2 = -72.96$, p = .0063; Store C $\beta_2 = 8.68$, p = .8582; and Store D $\beta_2 = -32.73$, p = .2438). The *post.open.months* coefficient identified significant increases in the repeat customer count trends in the store retail trading areas relative to the control retail trading areas in three out of the four stores (Store A: $\beta_3 = .92$, p = .0194; Store B: $\beta_3 = 1.59$, p = .0114; Store C: $\beta_3 = 9.00$, p = .0131; and Store D: $\beta_3 = 2.70$, p = .1720). Hence, the number of existing customer households purchasing from direct channels grows at a faster rate following the opening of a retail store.

6.4 Discussion.

The empirical analysis of customer development in direct channels at the household level helps us disentangle the two competing theoretical explanations outlined above for the direct sales response we found in retail trading areas surrounding new stores. First, the sales response exhibits a short term cannibalistic effect, which does not appear to be a result of a reduction in the purchasing frequency of existing households following the new store opening as the number of repeat customer households purchasing in that month does not significantly change. Nor is the short term drop in sales due to a decrease in the acquisition rate of new customers which also demonstrates no significant change following the opening of the store. Hence, it must be due to a reduction in the order size of existing and/or new customers in the direct channels (i.e. the same number of households are purchasing less). The customer analysis also suggests that the complementary sales effects identified above are partially caused by incremental faster acquisition of new customer households, as well as increased purchase frequency of repeat customer households. Hence, it offers support that both cannibalizing and complementary effects operate in tandem and that the store opening has a branding effect which induces new customers to try the direct channels who did not shop there previously, and induces existing direct channel households to shop in the catalog and online channels more frequently.

7. General Discussion

This research set out to test how adding a physical store would affect direct channel sales. We pursued four hypotheses which proposed that both cannibalizing and complementary effects occur in direct-to-consumer channels in the trading area of the new store and which specified the conditions under which each effect would dominate: 1.) Direct channel sales will be cannibalized immediately by store openings, and complementarity effects will be slower to be felt, 2.) Catalog sales will be more negatively impacted than online sales, and online sales will be more positively impacted than catalog sales when a retail store is opened, 3.) Direct sales from first-time and repeat direct channel customers would be complemented by the opening of a retail store, but only sales from existing customers of the direct channels will be cannibalized, and 4.) The pre-existence of a retail store in a trading area in which a new store is opened moderates both the cannibalizing and complementary effects of the new store on direct sales.

Our results show that a new retail store both cannibalizes and complements existing direct-to-consumer channels in its retail trading area and that the passage of time largely determines when each effect arrives and when each effect dominates the other. In the short term, a retail store is cannibalistic to the catalog channel, reducing sales in the catalog channel by 12% as existing customers of the catalog channel reduce their purchasing through the catalog and presumably begin shopping in the retail store. This suggests that heterogeneity within customers across purchase occasions occurs, as customers switch some of their demand for the retailer's goods from the direct channels to the retail store. However, over the longer term, a retail store is complementary to both catalog and online channels and allows them to grow sales, new customer household acquisition, and repeat customer household purchasing frequency at a greater than expected rate which more than makes up for the short term sales cannibalization. Hence, the data empirically supports prior researchers' suppositions that multichannel customers will be better customers than customers shopping in a single channel; in our study, existing customers begin purchasing more in the direct channels than they would have without a retail store presence.

Our results show that catalog and online sales exhibit similar patterns of complementary effects; however, only the catalog channel is cannibalized by the opening of a retail store in the short term. Over the long run, both direct channels are helped by the introduction of a retail store. However, again in contrast to previous empirical studies, the online channel's complementary response is magnified versus the response of the catalog channel, with online sales exhibiting a faster growth rate (34% vs. 0.4% on average) due to stronger complementary effects in the longer term. This asymmetry is due to an important difference between catalog and online sales channels which offers insight into the origins of demand for online and catalog retailing.

Finally, the existence of a retail store presence in the retail trading area prior to the opening of a new store does not seem to affect cross-channel complementarity, but does affect cross-channel cannibalization. Stores C and A, which open 34 miles and 52 miles respectively from existing stores, are still able to produce similar complementary effects in the catalog and online channels that Stores B and D, which open in virgin retail trading areas previously unserved by the retailer with retail stores, produce. However, Stores C and A do not exhibit significant cannibalization effects in the catalog channel as Stores B and D do, suggesting that catalog cannibalization may have already occurred back when the earlier stores were first opened in those retail trading areas. Hence, while the opening of a new store fails to cannibalize catalog sales in markets where the retailer has a previous store presence, it still reinforces and enhances pre-existing brand awareness and positive brand associations in markets in which the retailers has a store presence, as well as build new brand awareness and positive brand associations in markets in which the retailer is less known.

Care should be taken in extrapolating these results to other retailers as our study involved only stores opened by a single retailer with a well established and respected brand. Additional empirical studies of other retailers could add robustness to our findings and combined with the other studies outlined in the literature review move us closer as a field to empirical generalizations about cross-channel effects. Direct retailers with less established brands may exhibit direct sales responses which differ from the one we uncovered here. Specifically, we hypothesize that direct retailers with less established brands may benefit more from branding effects that the opening of a new store brings, such that their direct business may recover from the initial drop in sales from cannibalization and grow faster towards and beyond recovery through complementary effects than this retailer. Structurally, our retailer maintained the retail store operating unit and the direct channel operating unit as separate entities, perhaps limiting complementary effects from occurring across channels due to a dearth of cross-channel promotion and marketing coordination. However, this separation of the two units may also have artificially preserved relationships with direct customers who were not made aware of the store opening.

Additionally, the timing of the opening of a retail store in relation to the growth of the online as a viable direct-to-consumer channel may also affect cross-channel sales interactions. In the early days of the Internet as studied here, online sales were not cannibalized by the opening of a retail store; however, as online penetration grows and as shopping via the online becomes more predominant, the opening of retail stores may begin to cannibalize online sales. This may reflect the fact that early users of the online and those that were the early adopters of online shopping were demographically different (perhaps younger) from the shoppers in retail stores, so the opening of a store had little impact on them. However, as online shopping diffuses into the mainstream, the demographic differences between online shoppers and retail store shoppers are beginning to disappear, making it more likely that the opening of a retail store will both cannibalize online sales in the short term and complement online sales in the longer term.

Methodologically, this study illustrates the potential of zip code level data as a useful intermediate aggregation level for analysis of multichannel retailing research questions, one which is analogous to using store-level data in scanner models. In addition to the benefits described earlier in this paper, from a managerial perspective, it is also easier and less invasive to capture store purchases at the zip code level rather than at a household level which may trigger privacy concerns. The use of a less invasive identifier improves the likelihood of consumer-enabled identification, improving the likelihood that data samples more completely capture all transactions both within single households and across households.

It is also our hope that the matching process used in our study will be useful for researchers investigating natural experiments in marketing. Although becoming well established in other social sciences, matching has not been utilized in marketing research and serves an important role in improving our empirical methods. Managerially, matching techniques may also improve the precision of planned experimental methods such as A/B testing which are increasingly being used in interactive marketing contexts. In our study, matching allowed us to increase the internal validity of our findings and better isolate the effect of opening a retail store on direct sales by parceling out variance due to alternative explanations, such as income or population growth, which could potentially contribute to changes in direct sales.

REFERENCES

Abadie, A. and G.W. Imbens. 2004. Large sample properties of matching estimators for average treatment effects. *Working paper, http://elsa.berkeley.edu/~imbens/wp.shtml*.

Akaike, H. 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*. **19**(6), 716-723.

---. 1981. Likelihood of a model and information criteria. *Journal of Econometrics*. **16**, 3-14.

Alba, J., J. Lynch, B. Weitz, C. Janiszewski, R. Lutz, A. Sawyer and S. Wood. 1997. Interactive home shopping: Consumer, retailer, and manufacturer incentives to participate in electronic markets. *Journal of Marketing*. **61**(3), 38-53.

Ansari, A., C.F. Mela and S.A. Neslin. 2005. Customer channel migration. *Working paper, Teradata Center, Duke University*. Durham, NC.

Balasubramanian, S. 1998. Mail versus mall: A strategic analysis of competition between direct marketers and conventional retailers. *Marketing Science*. **17**(3), 181-195.

Biyalogorsky, E. and P. Naik. 2003. Clicks and mortar: The effect of on-line activities on off-line sales. *Marketing Letters*. **14**(1), 21-32.

Box, G.E.P. and G.M. Jenkins. 1976. *Time series analysis: Forecasting and control*. San Francisco, CA: Holden-Day.

Box, G.E.P. and G.C. Tiao. 1975. Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association*. **70**(349), 70-79.

Chiang, W.-y.K., D. Chhajed and J. Hess. 2003. Direct marketing, indirect profits: A strategic analysis of dual-channel supply-chain design. *Management Science*. **49**(1), 1-20.

Cook, T.D. and D.T. Campbell. 1979. *Quasi-experimentation*. Boston, MA: Houghton-Mifflin.

Cotlier, M. 2001. Allocating the sale. *Catalog Age*. **18**(3), 58-59.

Coughlan, A.T., E. Anderson, L.W. Stern and A.I. El-Ansary. 2001. *Marketing channels (Sixth Edition ed.)*. Upper Saddle River, NJ: Prentice Hall.

Dekimpe, M.G. and D.M. Hanssens. 1995. The persistence of marketing effects on sales. *Marketing Science*. **14**(1), 1-21.

Deleersnyder, B., I. Geyskens, K. Gielens and M.G. Dekimpe. 2002. How cannibalistic is the Internet channel: A study of the newspaper industry in the United Kingdom and the Netherlands. *International Journal of Research in Marketing*. **19**, 337-348.

Diamond, A. and J. Sekhon. 2005. Genetic matching for estimating causal effects: A new method of achieving balance in observational studies. http://jsekhon.fas.harvard.edu.

Dickey, D.A. and W.A. Fuller. 1979. Distribution of the estimators for autoregressive time series with a unit root. *Econometrica*. **49**, 1057-1072.

Fox, E.J., A.L. Montgomery and L.M. Lodish. 2002. Consumer shopping and spending across retail formats. *The Journal of Business*. **77**(2), S25-S60.

Geyskens, I., K. Gielens and M.G. Dekimpe. 2002. The market valuation of Internet channel additions. *Journal of Marketing*. **66**(2), 102-119.

Hansen, B.B. 2004. Full matching in an observational study of coaching for the SAT. *Journal of the American Statistical Association*. **99**, 609-618.

Hanssens, D.M., L.J. Parsons and R.L. Schultz. 2001. *Market response models: Econometric and time series analysis (2nd Edition ed.)*. Norwell, MA: Kluwer Academic Publishers.

Heckman, J., H. Ichimura, J. Smith and P. Todd. 1998. Characterizing selection bias using experimental data. *Econometrica*. **66**(5), 1017-1098.

Ho, D.E., K. Imai, G. King and E. Stuart. 2004. Matchit: Matching as nonparametric preprocesing for parametric causal inference.

---. 2005. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Working paper, Harvard University Institute for Quantitative Social Science*. Cambridge, MA.

---. 2007. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis*. **15**, 199-236.

Jacoby, J. and D. Mazursky. 1984. Linking brand and retailer images -- do the potential risks outweigh the potential benefits? *Journal of Retailing*. **60**(2), 105-122.

Jaffe, A.B., M. Trajtenberg and R. Henderson. 1993. Geographic location of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*. **108**(3), 577-598.

Keller, K.L. 1993. Conceptualizing, measuring, and managing customer-based brand equity. *Journal of Marketing*. **57**(1), 1-31.

Krishnamurthi, L., J. Narayan and S.P. Raj. 1986. Intervention analysis of a field experiment to assess the buildup effect of advertising. *Journal of Marketing Research*. **23**(November), 337-345.

Kumar, V. and R. Venkatesan. 2005. Who are the multichannel shoppers and how do they perform?: Correlates of multichannel shopping behavior. *Journal of Interactive Marketing*. **19**(2), 44-62.

Lee, P.M. and S. Wahal. 2004. Grandstanding, certification and the underpricing of venture capital backed IPOs. *Journal of Financial Economics*. **73**, 375-407.

Ljung, G.M. and G.E.P. Box. 1978. On a measure of a lack of fit in time series models. *Biometrika*. **65**, 297-303.

Meyer, B.D. 1995. Natural and quasi-experiments in economics. *Journal of Business and Economic Statistics*. **13**(2), 151-161.

Moriarity, R.T. and U. Moran. 1990. Managing hybrid marketing systems. *Harvard Business Review*. **90**(6), 146-155.

Pauwels, K. and S.A. Neslin. 2006. Building with bricks: The long-term revenue impact of opening physical stores. *Working paper, Dartmouth College*. Hanover, NH.

Pauwels, K. and S. Srinivasan. 2004. Who benefits from store brand entry? *Marketing Science*. **23**(3), 364.

Phillips, P.C.B. and S. Ouliaris. 1990. Asymptotic properties of residual based tests for cointegration. *Econometrica*. **58**, 165-193.

Phillips, P.C.B. and P. Perron. 1988. Testing for a unit root in time series regression. *Biometrika*. **75**, 335-346.

Reda, S. 2002. Active multi-channel shoppers may be a liability. *Stores*. **84**, 78-82.

Rosenbaum, P.R. 2002. *Observational studies (2nd Edition ed.)*. New York: Springer Verlag.

Shankar, V., A. Smith and A. Rangaswamy. 2003. Customer satisfaction and loyalty in online and offline environments. *International Journal of Research in Marketing*. **20**, 153-175.

Shapiro, C. and H.R. Varian. 1999. *Information rules*. Boston: Harvard Business School Press.

Steenburgh, T.J., A. Ainslie and P.H. Engebretson. 2003. Massively categorical variables: Revealing the information in zip codes. *Marketing Science*(Winter), 40-57.

Tang, F.-F. and X. Xing. 2001. Will the growth of multi-channel retailing diminish the pricing efficiency of the web? *Journal of Retailing*. **77**, 319-333.

Verhoef, P.C., S.A. Neslin and B. Vroomen. 2005. Browsing versus buying: Determinants of customer search and buy decisions in a multi-channel environment. *Working paper, University of Groningen*. Groningen, Netherlands.

Wallace, D.W., J.L. Giese and J.L. Johnson. 2004. Customer retailer loyalty in the context of multiple channel strategies. *Journal of Retailing*. **80**, 249-263.

Ward, M.R. 2001. Will online shopping compete more with traditional retailing or catalog shopping? *Netnomics*. **3**, 103-117.

Ward, M.R. and M. Morganosky. 2000. Online consumer search and purchase in a multiple channel environment. *Working paper, Department of Agricultural and Consumer Economics, University of Illinois*. Urbana Champaign, IL.

Wind, Y. and V. Mahajan. 2002. Convergence marketing. *Journal of Interactive Marketing*. **16**, 64-74.

Winship, C. and S.L. Morgan. 1999. The estimation of causal effects from observational data. *Annual Review of Sociology*. **25**, 659-707.

Zettelmeyer, F. 2000. Expanding to the Internet: Pricing and communications strategies when firms compete on multiple channels. *Journal of Marketing Research*. **37**(3), 292-308.

FIGURE 1: QUANTILE-QUANTILE PLOT

STORE C AVERAGE AGE MATCHING VARIABLE

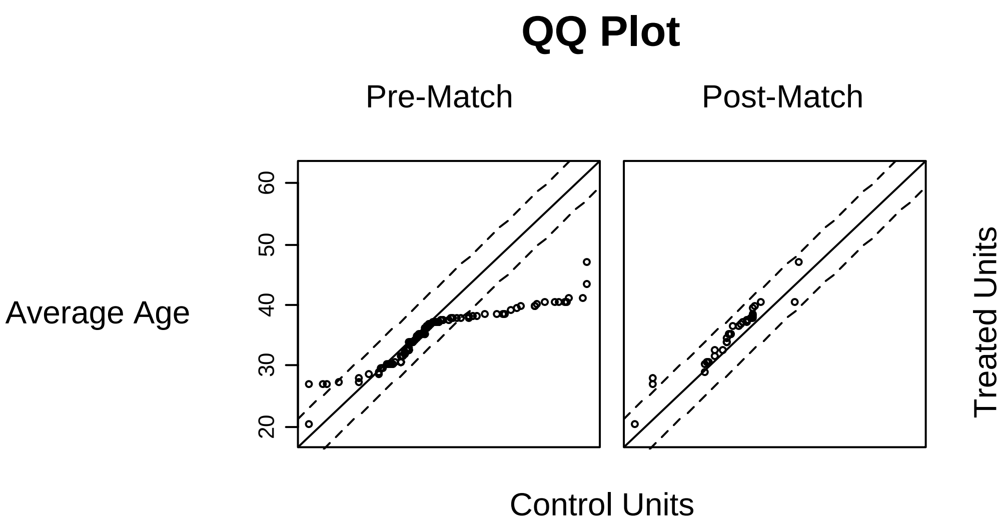

TABLE 1:
MATCHING ASSESSMENT FROM GENETIC PROCEDURE

Standardized Mean Difference between Treatment and Control

	Prior to Matching				Genetic Matching			
	Store A	Store B	Store C	Store D	Store A	Store B	Store C	Store D
Drive-time	-0.138	0.285	-0.860	0.406	-0.142	-0.031	-0.174	-0.074
Average Population	0.602	0.341	0.670	0.294	0.138	0.035	0.124	-0.014
Population CAGR	-0.026	-0.329	-0.100	0.075	-0.005	-0.017	-0.205	0.034
Average Income	0.550	-0.138	0.015	-0.021	0.068	0.033	0.104	0.049
Income CAGR	-0.874	0.190	-0.965	0.038	0.076	0.071	-0.154	-0.051
Average Age	0.554	0.113	-0.895	-0.315	0.060	0.049	0.336	-0.144
Adults Buying via Catalog	-0.153	0.147	-0.423	-0.118	-0.108	0.014	-0.044	-0.066
Adults Buying via Internet	-0.230	-0.065	-0.263	-0.019	-0.100	0.016	0.185	0.120
HH w/ Internet Access	0.571	-0.207	0.000	-0.052	0.051	-0.018	0.029	-0.032

TABLE 2:
MATCHING ASSESSMENT FROM OTHER PROCEDURES

Standardized Mean Difference between Treatment and Control

	Sub-classification Matching				Nearest Neighbor Matching			
	Store A*	Store B	Store C*	Store D	Store A	Store B	Store C	Store D
Drive-time	0.276	0.122	0.292	0.113	0.119	0.342	-0.135	0.358
Average Population	0.406	0.113	0.414	0.167	-0.022	0.428	0.327	0.285
Population CAGR	0.318	0.098	0.144	0.090	-0.014	-0.422	-0.071	0.037
Average Income	0.300	0.110	0.083	0.129	0.078	-0.192	-0.126	-0.029
Income CAGR	0.233	0.102	0.354	0.129	-0.245	0.207	-0.545	0.039
Average Age	0.235	0.127	0.215	0.121	0.249	0.178	-0.333	-0.241
Adults Buying via Catalog	0.219	0.174	0.131	0.136	-0.194	0.237	-0.389	-0.097
Adults Buying via Internet	0.083	0.163	0.180	0.181	-0.224	-0.056	-0.467	-0.014
HH w/ Internet Access	0.325	0.066	0.104	0.176	0.040	-0.281	-0.055	-0.047

	Optimal Matching				Full Matching			
	Store A	Store B*	Store C	Store D	Store A	Store B	Store C	Store D
Drive-time	0.121	0.356	-0.135	0.358	0.438	0.119	-0.051	0.017
Average Population	-0.009	0.278	0.327	0.285	0.295	-0.067	0.417	0.032
Population CAGR	-0.016	0.064	-0.071	0.037	-0.099	-0.057	-0.124	0.019
Average Income	0.066	-0.049	-0.126	-0.029	-0.067	0.078	-0.422	-0.156
Income CAGR	-0.234	0.032	-0.545	0.039	0.069	-0.024	-0.366	0.002
Average Age	0.280	-0.260	-0.333	-0.241	0.132	0.154	-0.151	0.007
Adults Buying via Catalog	-0.209	-0.094	-0.389	-0.097	0.385	0.285	0.284	-0.108
Adults Buying via Internet	-0.247	-0.012	-0.467	-0.014	0.497	0.249	0.160	0.033
HH w/ Internet Access	0.028	-0.055	-0.055	-0.047	0.057	0.056	-0.462	-0.126

* The number of subclasses was restricted to five or fewer in Store A and to three or fewer in store C to allow the sub-classification procedure to converge.

** Several treatment zip codes were excluded in Store B to allow the Optimal matching procedure to converge.

TABLE 3: DIRECT SALES TIME SERIES TESTS

FOR STATIONARITY AND COINTEGRATION

	PHILLIPS-PERRON UNIT ROOT TESTS		**PHILLIPS-OULIARIS COINTEGRATION TESTS**	
	Dickey-Fuller Z (alpha)	p-value	Phillips Ouliaris demeaned	p-value
Store A				
Catalog Sales Treatment	-50.96	.01	-78.41	.01
Catalog Sales Control	-52.77	.01		
Online Sales Treatment	-12.61	.38	-99.03	.01
Online Sales Control	-53.5	.01		
Store B				
Catalog Sales Treatment	-40.96	.01	-58.32	.01
Catalog Sales Control	-42.65	.01		
Online Sales Treatment	-18.61	.08	-84.1	.01
Online Sales Control	-23.38	.02		
Store C				
Catalog Sales Treatment	-38.7	.01	-65.43	.01
Catalog Sales Control	-35.65	.01		
Online Sales Treatment	1.601	.99	-31.51	.01
Online Sales Control	-35.86	.01		
Store D				
Catalog Sales Treatment	-37.45	.01	-73.02	.01
Catalog Sales Control	-32.57	.01		
Online Sales Treatment	-12.1	.40	-33.11	.01
Online Sales Control	-27.49	.01		

TABLE 4: "ALL STORE" DIRECT SALES MODEL RESULTS

	CATALOG SALES	**ONLINE SALES**
Intercept	22,803 (7897)	-35,271 (47,731)
Pre.open.months β_1	15 (224)	252 (765)
Store.open β_2	-12,947* (5503)	-1,610 (6861)
Post.open.months β_3	395** (121)	2,360*** (491)
Control β_4	1.23*** (.08)	1.33*** (.06)
Store Fixed Effects	Yes	Yes
	ARIMA (0,0,1) AIC = 7,101	ARIMA (1,0,1) AIC = 6,526

*** p < .001, ** p < .01, * p < .05, † p < .10

TABLE 5: STORE BY STORE DIRECT SALES MODEL RESULTS

	CATALOG SALES				ONLINE SALES			
	Store A	**Store B**	**Store C**	**Store D**	**Store A**	**Store B**	**Store C**	**Store D**
Intercept	29,707 (11,068)	42,898 (9365)	58,030 (7758)	14,017 (4562)	-582 (9180)	6,592 (2651)	22,924 (3,391,100)	544 (531,544)
Pre.open.months β_1	323 (384)	-224 (312)	-166 (288)	-196† (109)	-120 (1232)	290 (210)	584 (2366)	19 (583)
Store.open β_2	-14,280 (9397)	-25,684** (8495)	-6,165 (8004)	-5,640† (3181)	-6,657 (9718)	-5,775† (3218)	-11,660 (15000)	832 (6080)
Post.open.months β_3	450** (155)	517** (184)	100 (250)	129 (98)	977*** (155)	794*** (108)	5,853* (2254)	1,455* (562)
Control β_4	2.80*** (.22)	1.07*** (.09)	2.67*** (.32)	1.042** (.08)	1.60*** (.16)	1.3*** (.05)	3.09*** (0)	1.21*** (0.1)
	ARIMA (0,0,0) AIC = 2,197	ARIMA (0,0,0) AIC = 1,994	ARIMA (0,0,0) AIC = 1,683	ARIMA (1,0,2) AIC = 1,579	ARIMA (1,0,2) AIC = 1,547	ARIMA (1,0,2) AIC = 1,539	ARIMA (1,0,2) AIC = 1,572	ARIMA (1,0,2) AIC = 1,480

*** p < .001, ** p < .01, * p < .05, † p < .10

TABLE 6: CUSTOMER COUNT TIME SERIES TESTS FOR STATIONARITY AND COINTEGRATION

	PHILLIPS-PERRON UNIT ROOT TESTS		**PHILLIPS-OULIARIS COINTEGRATION TESTS**	
	Dickey-Fuller Z (alpha)	p-value	Phillips Ouliaris demeaned	p-value
Store A				
First-Time Customers Treatment	-41.33	.01	-102.5	.01
First-Time Customers Control	-46.54	.01		
Repeat Customers Treatment	-44.29	.01	-92.48	.01
Repeat Customers Control	-41.78	.01		
Store B				
First-Time Customers Treatment	-35.87	.01	-75.7	.01
First-Time Customers Control	-43.36	.01		
Repeat Customers Treatment	-30.78	.01	-52.48	.01
Repeat Customers Control	-36.17	.01		
Store C				
First-Time Customers Treatment	-29.69	.01	-24	.01
First-Time Customers Control	-44.17	.01		
Repeat Customers Treatment	-14.90	.23	-26.76	.02
Repeat Customers Control	-29.67	.01		
Store D				
First-Time Customers Treatment	-28.75	.01	-62.34	.01
First-Time Customers Control	-38.62	.01		
Repeat Customers Treatment	-23.47	.23	-37.03	.01
Repeat Customers Control	-28.79	.01		

TABLE 7: "ALL STORE" CUSTOMER COUNT MODEL RESULTS

	FIRST-TIME CUSTOMERS	REPEAT CUSTOMERS
Intercept	8.12	-21.24
Pre.open.months β_1	-.44	.22
Store.open β_2	1.32	-22.15
Post.open.months β_3	1.74**	3.16**
Control β_4	1.66***	1.55***
Store Fixed Effects	Yes	Yes
	ARIMA (2,0,1) AIC = 3,336	ARIMA (1,0,2) AIC = 3,733

*** p < .001, ** p < .01, * p < .05, † p < .10

TABLE 8: STORE BY STORE CUSTOMER COUNT MODEL RESULTS

	First-Time Customers				Repeat Customers			
	Store A	Store B	Store C	Store D	Store A	Store B	Store C	Store D
Intercept	22.08	3.34	56.17	-1.84	104.61	55.80	166.50	11.78
Pre.open.months β_1	-.05	-.57	.48	-.35	.90	-1.28	1.88	-1.39
Store.open β_2	2.75	-10.09	8.20	-8.16	-31.21	-72.96**	8.68	-32.73
Post.open.months β_3	.19	1.06***	5.79†	1.56***	.92*	1.59*	9.00*	2.70
Control β_4	2.53***	1.59***	3.53***	1.42***	2.50***	1.40***	2.75***	1.22***
	ARIMA (0,0,0) AIC = 934	ARIMA (0,0,0) AIC = 832	ARIMA (1,0,2) AIC = 747	ARIMA (1,0,0) AIC = 661	ARIMA (1,0,2) AIC = 1,034	ARIMA (1,0,2) AIC = 938	ARIMA (1,0,2) AIC = 811	ARIMA (1,0,2) AIC = 747

*** p < .001, ** p < .01, * p < .05, † p < .10