Journal of Economic Theory and Econometrics, Vol. 27, No. 1, Mar. 2016, 16–33

Spillover Effects of Temporary Price Cuts: Evidence from U.S. Scanner Data

Hyunchul Kim Kyoo il Kim*

Abstract Temporary price cut is one of the most popular marketing tactics in the supermarket industry. We examine the effects of price promotions on storewide profitability. Using a treatment effect setting we estimate the short-term and long-term promotional effects that spill over into other product categories using supermarket scanner data. In an effort to remove the selection bias due to differences in the subpopulations of the treatment and the control groups, we use several nonparametric imputation methods. The detailed information on household level purchases and store visits allows us to improve matching quality in the estimation. We find that the effects of price promotion on storewide sales deteriorate over time and that the spillover effects are substantial. This suggests that pricing strategies of multi-product retailers maximizing storewide profits should take into account the spillover effects of short-term pricing.

Keywords Price promotion, Spillover effects, Multi-product retailers, Supermarket industry

JEL Classification D1, L1, L8, M3

*We are grateful to Pat Bajari for his advice and support. We would also like to thank the seminar participants at University of Minnesota for many useful comments. Correspondence: Hyunchul Kim, Department of Economics, Sungkyunkwan University, email: hchkim@skku.edu; Kyoo il Kim, Department of Economics, Michigan State University, email: kyookim@msu.edu.

Received February 4, 2016, Revised March 4, 2016, Accepted March 5, 2016

1. INTRODUCTION

Price promotion is one of the most heavily used tactics in the supermarket industry and there have been scores of studies to establish consumer responses to temporary price cuts. Particularly, price changes affect store choice decisions for upcoming shopping events, and storewide sales may change consequently. In this paper, we attempt to measure the spillover effects of price promotions on consumer choices in multi-product retail stores and examine the dynamics (i.e. persistency) of these effects.

Consumer response to price cut is either instantaneous or persistent. Consumers may impulsively decide to purchase unplanned products or increase the quantity purchased of planned products when they encounter price promotions at stores. These instantaneous promotion effects only occur within the category for which price cuts are offered. On the other hand, price cuts may also affect store choices in subsequent shopping plans by altering price beliefs of consumers about store prices. When consumers have incomplete price information for each store, the update in price expectation that arises from the experience of price promotions can have systematic effects on their store choices.[1] Since shoppers typically purchase multiple categories in a single store trip due to considerable transportation and search costs, the promotion effects on store choices can spread to other categories which have not been promoted. In such a situation, the externality of price cuts is an important factor in pricing strategies of the retail stores who maximize storewide profitability.

Previous work in the marketing and economics literature restricted the attention to the within-category effects of price promotions (e.g. Blattberg et al., 1995; Kalyanaram and Winer, 1995; Erdem, 1996; Mela et al., 1997; Ailawadi and Neslin, 1998; Van Heerde et al., 2000; Hendel and Nevo, 2003, among others). The within-category or within-brand analysis, however, does not allow for the storewide effects of price changes and thus provides store managers with limited guidance in setting optimal pricing strategies. There is a relatively small literature that examined the cross-category effects of price promotions between different but related categories. Walters (1991) investigates the within- and inter-store effects of retail price cuts on promoted and non-promoted products across complementary categories and Niraj et al. (2008) show that the cross-category effects of price cuts on purchase incidences and quantities are asymmetric between two complementary categories. The spillover effects we study in this pa-

[1]Consumers rely on their price expectation in store choices since they do not know (or have limited information on) store prices before they visit a store. See the evidences reported in Bell et al. (1998) and Kim and Kim (2015).

per, in contrast, examine the externality of price cuts across different categories, which are not necessarily related to each other.

When consumers rely on their price beliefs which change over time according to the store prices they observe or pay in previous store trips, the storewide effects of price promotions persist for a while depending on shopping patterns of each consumer. For this reason, the spillover effects of price promotions can be analyzed in a dynamic framework. In this paper, using methods of estimating average treatment effects, we evaluate the short-term and long-term effects of price promotions allowing for different time spans for which those effects last.

In our framework the treatment status is about purchasing promoted items at the household level and we measure the average effects of the price promotions on storewide dollar spending of consumers. The fundamental empirical challenge is we do observe any potential outcomes (spillovers in the future purchasing behaviors) in the absence of the treatment (purchasing promoted items). Moreover, the purchasing decision is endogenous because it is the choice of households e.g. maximizing their utilities. In an effort to deal with or at least mitigate this selection and the missing data problem we resort to a few different nonparametric imputation methods (see e.g. Rosenbaum and Rubin, 1983; Abadie and Imbens, 2002; Hahn, 1998; Imbens, 2004). One method is a regression approach, which requires estimating a regression function of the predicted outcomes under no treatment for the treated group and the other is a matching method that imputes the missing outcomes for the treated group using the outcomes of their nearest neighbors in the control group and vice versa, with the distance as defined in terms of the covariates. These treatment effects estimators address the issues under the arguable assumption that the systematic difference in the outcomes between treated and control groups with similar characteristics be attributed to the treatment, often referred to as unconfoundedness, selection on observables, and conditional independence in the treatment effects literature.

For our analysis we use the supermarket scanner data that contain household level grocery purchases for 31 different categories. The availability of extremely detailed information on the historical purchases and demographics of about 6,000 households allows the data-driven covariates to sufficiently account for the differences between treated and control groups, which plays a key role in consistently estimating the average treatment effects.

We find that the spillover effects of price cuts are substantial. The estimated spillover effects are particularly larger in some categories than in others. More importantly, the within-category effects of price promotions for these categories are much smaller than the spillover effects. In other words, price promotions on products in a certain category give rise to an increase in dollar sales in other categories to a larger extent than that occurring within the same category. We also find that the spillover effects deteriorate over time but persist for a long time. The evidence of the long-lasting promotion effects implies that consumers not only respond to price cuts by immediately increase purchases of the products but they also update price expectation based on these price changes in their upcoming store choice decisions.[2]

This paper, to the best of our knowledge, is the first attempt to evaluate the effects of price cuts that spill over into other categories. Most of the existing literature focuses on the within-category or within-brand effects of price promotions. Given the complexity of consumer choices in response to price changes in the supermarket industry, one of the challenging problems retail stores face is how they choose the loss-leaders or set the frequency and depth of price cuts maximizing storewide profits. This paper provides a useful guideline for these problems.

While the short-term effects of price cuts have been extensively examined in the literature, the long-lasting promotional effects have received relatively small attention by researchers.[3] The notable exceptions are Erdem et al. (2003), Hendel and Nevo (2006), and Pesendorfer (2002). Erdem et al. (2003) and Hendel and Nevo (2006) study the demand of storable goods subject to frequent price promotions in a dynamic framework, and Pesendorfer (2002) examines the pricing strategies of grocery stores taking into account the inter-temporal demand effects of temporary price cuts. We contribute to this literature of the long-term promotional effects.

The paper is organized as follows. Section 2 provides the description of data. Section 3 outlines our identification strategies and estimation methods. The results are presented in Section 4. Section 5 concludes.

2. DATA

We use two data sets. The first is household level panel data of weekly grocery purchases. The second data set is store level data containing weekly store sales. These two data sets were collected by IRI with scanning devices for 31 different categories from January 2001 to December 2007.

[2]Kim and Kim (2015) develop and estimate a structural store choice model that takes into account the role of price expectation in store choices, and find that the updates in price expectation affect upcoming store choice decisions.

[3]Blattberg et al. (1995) point out that the long-term promotional effect is "the most debated issue in the promotional literature and one for which the jury is still out."

Table 1: Summary Statistics of Household-Level Data

	Mean	Median	Std	Min	Max
Number of trips per month	6.4	5	3.6	1	46
Number of visited stores	3.1	3	1.2	1	6
Weekly spending (31 categories)	16.18	14.26	9.35	2.28	104.87
Weekly spending (all categories)	61.58	48.58	49.72	.49	415.65
Number of categories per trip	3.2	3	1.2	1	11
Demographics					
Income (thousand)	40.2	40.0	23.6	8	87.5
Family size	2.4	2	1.3	1	6
Home owner (percentage)	0.80	-	-	0	1

The household level data contain information on weekly purchases of the sample households at the product level, including price and quantity for each purchase. Each product is identified by the UPC (universal product code), which is a unique code assigned to each product. The data cover seven retail stores of four supermarket chains in a small city in Wisconsin. The store level data have weekly store sales at the UPC level. In addition to price and quantity, this data set includes information on promotional activities such as price discount, advertising, and shelf display. In particular, information on price promotion is a dummy variable for 5% or higher price discount. We merged this variable into the household level data.[4] These two data sets include 31 grocery goods, which consist of 17 food and beverage goods (e.g. soft drink, salted snack, peanut butter) and 14 non-food household goods (e.g. laundry detergent, shampoo, facial tissue). We also have a supplementary data set of individual store visits for each household, which includes weekly total expenditure for each store visit covering all goods purchased at the stores. The household level data also include demographic characteristics of each household.

Table 1 provides the summary statistics of the household level data and the demographics of the sample households. Households visit grocery stores about 1.5 times a week on average and they visit about three different stores in a year. Based on the household level trip data, the weekly dollar spending on the 31 categories included in the household level data is about 25 percent of the total

[4]The price variable in the store level data is the weekly average of the prices actually paid by customers and the prices in the household level data are drawn from the store level data.

Table 2: Price Promotions by Chain

	Number of stores	Frequency of promotion	Depth of price cuts (%)	Assortment size
Chain A	1	9.45%	18.5	13,522
Chain B	1	21.3%	17.6	14,064
Chain C	1	12.8%	17.4	14,872
Chain D	2	10.6%	17.6	10,985

Note: Frequency of promotion reports the percentage frequency of price promotions with more than 5% discounts. Assortment size is the number of different products of the 31 categories.

weekly spending on all categories at stores. Among the 31 included categories, households buy three different categories for each store trip on average. This bundle purchase behavior provides the motivation for examining spillover effects of price discounts.

Table 2 presents the descriptive statistics of price promotions offered by each chain supermarket. According to the nondislosure agreement with IRI, the name of chain supermarkets are masked. Among the seven stores included in the data, we exclude two stores that entered or exited the market during the data period. The frequency of price promotion is the percentage frequency of promotion weeks for each product with price discounts higher than 5%. The percentage price discounts is based on regular prices calculated by taking the mode of unpromoted prices over the year. Although the percentage price discounts are similar among supermarket chains, there is a substantial variation in the frequency of price promotions across chain stores and the amount of such variations also varies across different categories. Assortment size is the number of different products of the 31 categories and the four chain stores differ in the assortment size.

3. ECONOMETRIC FRAMEWORK AND ESTIMATION

The object of our study is to measure the impact of price promotions on storewide spending of households over different time frames. To this end, we use semiparametric estimation strategies for average treatment effects using a few different nonparametric imputation methods. In our framework the treatment group is defined as the households who have purchased *promoted* products within a specific period and the control group includes those who have not. Decisions of purchasing promoted products are the *optimal* choices made by the households, so we have the potential selection issue. Our identifying assumption of the treatment effects rely on the arguable assumption (e.g. unconfoundedness) that receiving a treatment is independent of potential outcomes given observables of consumers. Underlying key content of this assumption is that either those observables are all variables that enter the purchasing decision process of the households or if the decision is driven also by differences in other unobserved factors, then those unobserved factors should be unrelated to the potential outcomes. In other words the assumption asserts that controlling for only observable differences in consumers is sufficient to remove the selection bias. Our identification strategy is then that the average differences in spending between the treated and control groups, given the same values for covariates, are attributable to the treatment.

To formalize our framework we adopt the standard notations in the literature of the average treatment effects (e.g. Imbens, 2004; Rosenbaum and Rubin, 1983; and Rosenbaum and Rubin, 1984). The treatment group is a random sample of N_{1t} households who purchase promoted items at time t and the control group is a random sample of size N_{0t} drawn from the control population, where $N_t = N_{0t} + N_{1t}$ is the total sample size. Household i spend y_{it} at month t. Let $Y_{it,m} = \sum_{s=t+1}^{t+m} y_{is}$ denote the cumulative total spending of household i for the next m months after one month since the time t. Household i has an outcome $Y_{it,m}(1)$ if he or she receives the treatment of price promotions at time t and an outcome $Y_{it,m}(0)$ if not treated. Let W_{it} indicate whether household i is actually treated ($W_{it} = 1$) or not ($W_{it} = 0$). Note that $Y_{it,m}(0)$ is *not* observed in the data if $W_{it} = 1$ since it is the outcome for the treated household had she not been treated, and *vice versa* if $W_{it} = 0$.

Here note that for the outcome variables we transform weekly data into monthly level. We do this for a couple of reasons. First, the underlying intuition behind the spillover effects is the update in consumers' price expectation, and thus we define outcome variable with a longer time period than a week to more sufficiently capture the effects of price updates on the upcoming store choices or purchases. Furthermore, the within-category effects of price promotions can be estimated in a meaningful way with monthly data–instead of weekly data–because purchase incidence of most categories is not as frequent as weekly. Second, to define the outcome variable by month allows a reasonable size of both the treatment and the control groups.

Next we assume a vector of characteristics or pretreatment variables, denoted by X_{it}, are available for the household i at time t. These variables are what we use to control for the selection problem. They may include pretreatment variables such as dollar spending, purchase frequency, and price discounts for category purchases in the previous period and include demographic variables such as income, age, home ownership, family size, number of children, and indicators for marriage and pet ownership. Using predetermined outcome variables, which are realized before evaluation periods, as the conditioning covariates is well justified in the literature. See e.g. Imbens (2004).

The following two assumptions play an important role in identifying the average treatment effects.

Assumption 1. *(Unconfoundedness) For almost all x,* $(Y_{it,m}(1), Y_{it,m}(0))$ *is independent of* W_{it} *conditional on* $X_{it} = x$*, that is,*

$$(Y_{it,m}(1), Y_{it,m}(0)) \perp W_{it} | X_{it} = x.$$

Assumption 2. *(Overlap) For some* $c > 0$ *and almost all x,*

$$c \leq Pr(W_{it} = 1 | X_{it} = x) \leq 1 - c.$$

The *unconfoundedness* (or *conditional independence*) means that receiving a treatment is independent of the potential outcomes with and without the treatment, being conditioned on the observable covariates. This assumption is analogous to the standard exogeneity assumption that W_{it} is exogenous given the covariates. The *overlap* assumption requires that both treated and nontreated households with any given $X = x$ are present in the data so that there exist one or more untreated households that are use to match with each treated household and *vice versa*. If the second assumption does not hold, then for households who has a particular value of $X = x$ we cannot impute the potential missing outcomes, so the strategy of estimating the average treatment effects by averaging over the conditional treatment effects for a subpopulation with covariates $X = x$ does not work.

For our study we are interested in the average effects for the treated who purchased promotion items at time t (the population-average treatment effect for the treated: PATT) written as

$$\tau_t^T = \mathbb{E}[Y_{it,m}(1) - Y_{it,m}(0)|W_{it} = 1]. \tag{1}$$

3.1. INTERPRETATION OF THE TREATMENT EFFECTS

Note that in our framework the treatment status indicator W_{it} depends on both i and t because the households who responded to the promotion at one period may not respond in other periods. Also our treatment effects can vary by time t depending on the promotion strengths and also how the treated and the control groups vary by time. Complications that may arise in our framework are in two-fold: first (a) both the treated and the control groups we consider at time t can be further exposed to the upcoming promotions that arrive between $t+1$ and $t+m$ during the evaluation periods and second (b) the two groups defined at time t can make persistently different responses to those ongoing promotions during the evaluation periods.

Because of the issue (a) the treatment effects we identify should be interpreted as the accumulated effects of promotions during the evaluation periods due to (persistently) different responses of the treated and the control groups. Therefore the issue of (a) is only a matter of how we interpret the treatment effects. In dealing with the issue (b), our assumption to ensure validity of the treatment effects in (1) based on the unconfoundedness is that controlling for the observed differences in covariates of the two groups at time t should be sufficient to adjust for the systematic differences in their responses to the ongoing promotions during the evaluation periods.

However, there might remain a potential endogeneity in the treatment we define (i.e. purchasing promoted items in a given period). To avoid this potential endogeneity (even after being conditioned on the observables), we define the outcome variables as the future spendings of households after one month since receiving the treatment. One could alternatively define the outcome variables as the future spendings after several months since the treatment instead of one month.

Finally we use only a binary indicator of the treatment in our framework but one could develop a multi-valued treatment effect framework to evaluate the effects of promotions depending on the levels of responses to the promotions e.g., W_{it} is not a binary indicator but is an integer value of how many times the households purchased the promoted items. This multi-valued treatment framework can also explicitly handle the issue of (b) but this route is beyond our scope in this paper.

3.2. TREATMENT EFFECTS ESTIMATORS

Below we briefly review a few different methods of estimating these treatments effects, which we use for our analysis.

3.2.1 Regression Approach

We begin with a regression approach as a benchmark strategy of estimating the average treatment effects. It is based on the consistent estimation of the conditional regression function $\mu_\omega(x)$ for $\omega \in \{0,1\}$, defined as (suppressing the subscript m)

$$\mu_\omega(x) \equiv \mathbb{E}[Y_{it}(\omega)|X_{it} = x]. \tag{2}$$

Given the estimate $\hat{\mu}_\omega(x)$ of the regression function, the average treatment effect for the treated group is estimated by averaging the difference between the actual outcomes for the treated group and the estimates of their outcomes without the treatment. Formally, the regression estimator of the average treatment effect for the treated group is

$$\hat{\tau}^T_{t,reg} = \frac{1}{N_{1t}} \sum_{i=1}^{N_t} W_{it} \cdot [Y_{it} - \hat{\mu}_0(X_{it})]. \tag{3}$$

The estimated regression function is the predicted outcomes under no treatment for the treated group. The regression function for the control group, $\mu_0(x)$, is thus used to predict the missing outcomes for the treated group. Hence, if the first step estimation of $\mu_\omega(x)$ relies on parametric regressions, the results may be sensitive to the distributional differences in the covariates between the treated and control group since the estimators depend heavily on extrapolation. Hahn (1998) and Imbens et al. (2003) proposed nonparametric methods for estimating $\mu_\omega(x)$.

3.2.2 Matching

The second approach is matching estimation. We follow Abadie and Imbens (2002) and Abadie and Imbens (2006)'s methodology. The matching estimators impute the missing outcomes for the treated group using the outcomes of their nearest neighbors in the control group and vice versa, with the distance as defined in terms of the covariates. For household i in the treated group, define d_{it} as the distance measured by the Euclidean metric between the vectors of the covariates for i and the nearest match in the control group.

$$d_{it} = \min_{j=1,\ldots,N||W_{jt}=0} \| X_{it} - X_{jt} \| . \tag{4}$$

Then, let

$$\mathscr{T}_t(i) = \{ j \in \{1,2,\ldots,N\} : W_{jt} = 0, \| X_{it} - X_{jt} \| = d_{it} \}$$

be the set of the closest matches for the treated household i. Note that $\mathscr{T}_t(i)$ contains a single match if X_{it} is continuously distributed. Now, the estimate of the missing outcome $Y_{it}(0)$ of the treated household i is the average outcomes of the matched households in the control group:

$$\widehat{Y}_{it}(0) = \frac{1}{\#\mathscr{T}_t(i)} \sum_{j \in \mathscr{T}_t(i)} Y_{jt}, \tag{5}$$

where $\#\mathscr{T}_t(i)$ is the number of elements in $\mathscr{T}_t(i)$. Similarly we can obtain $\widehat{Y}_{it}(1)$ of the control household as the average outcomes of the matched households in the treated group.

Then, the matching estimator of PATT is written as

$$\hat{\tau}^T_{t,match} = \tfrac{1}{N_{1t}} \sum_{i:W_{it}=1} (Y_{it} - \widehat{Y}_{it}(0)). \tag{6}$$

We note that this matching estimator can potentially suffer from finite sample bias. Abadie and Imbens (2002) show that the bias term due to the matching discrepancy that arises in (4) is the order of $N_0^{(-1/k)}$, where N_0 is the size of the control group and k is the number of covariates. So, when k is large, the resulting treatment effect estimator can potentially perform poorly in the finite sample. However, they also show that this bias can be ignored if the number of controls is quite large relative to the number of treated as in our application.

4. ESTIMATION RESULTS

We estimate the average treatment effects of price promotions in four product categories. We choose these categories because the purchase frequency is relatively large and allow a larger size of treatment group. We focus on temporary price cuts in one month of the data period, July 2003. For a sensitivity analysis, we also do the same estimations for five other (randomly-picked) months,

Table 3: Expenditure and Store Visits of Treated Group

	Category spending ($)	Store visits	Spending on other categories	Number of households
Soft-drinks	13.37	1.4	49.54	1,931
Cereal	8.27	0.9	58.90	713
Laundry detergent	2.53	0.2	70.83	232
Salted snack	6.34	1.2	58.00	1,540

Note: Spending and store visits are monthly average of the households in the treatment group and they are calculated based on the data from July 2003 and January 2004 because the prediction of future spending in the estimation covers this time period.

and the qualitative results remain the same. The data used for the estimation include all purchases with or without price promotion in the five grocery stores, which belong to four different chain supermarkets. The treatment group is a set of households who purchased products with price promotion in each category in July 2003. The control group includes those who did not buy any promoted product of the category during the same period. We also include the households who did not visit the store during the period to the control group.[5]

Table 3 shows the expenditure and store visits of the households in the treatment group for each category. The spending and the trip frequency in the table are monthly average and computed based on the data from July 2003 and January 2004, which cover the time period of the predicted future spending in the estimation. The households who purchased products on sale in the treatment period visit stores about once in a month for purchasing soft-drink, cereal, and salted snack, whereas they visit stores every five months to buy laundry detergent. Particularly, the average monthly spending on laundry detergent reflects this infrequency of store visits.

Table 4 provides a list of covariates used for predicting the future spending of each household. They include dollar spending, purchase frequency, and price discounts for category purchases in the previous period. Percentage discounts in the current period, which exclude those for the category of interest, are used for

[5]If shoppers who did not purchase promoted products observed the price discounts during their store trips and this affects their future spending at the store, adding the customers who visit the store during the treatment period to the control group may underestimate the effects of price promotions. However, given the variety of products and the limited amount of time for each shopping trip, it is reasonable to assume that customers perceive price promotions only when they purchase the products.

Table 4: Definition of Explanatory Variables

Variable	Definition
PREDOL_*category*	Spending on each category during the previous period.
PREDOL_OTHER	Spending on other categories during the previous period.
PREFREQ_*category*	Frequency of previous store visits for the purchase of each category.
PREPROM_*category*	Previous spending on promoted products of each category.
PREPROMDISC_*category*	Average percentage price discounts of previous purchases for each category.
PROMO_*category*	Dummy for price promotion in the current period.
PROMDISC_*category*	Average percentage price discounts of current purchases for each category.
FIRST_VISIT	Dummy for first visit to the store during the previous period.
LOYALTY_LENGTH	Length of time (by months) since the first visit to the store.
Demographic variables	Income, age, marriage, home ownership, family size, number of children, pet (dog or cat) ownership.

Note: Previous spending on other categories is only used in estimating the spillover effects. The dummy variables and percentage discounts of current promotions do not include the promotion of the category of interest. Income variables include dummies for 12 income groups.

controlling for the effects of promotions in other categories. We also control for store loyalty measured by the time period of previous store visits and the indicator of first visit during the previous period.[6] Demographic variables include income, age, home ownership, family size, number of children, and indicators for marriage and pet ownership.

Table 5 and 6 report the effects of price promotions. In the regression approach, we use a parametric regression function to estimate the missing outcomes of the treated group. R^2 of the regressions based on the control group is fairly high for most of the cases. For the matching estimation strategy, we use the bias-corrected matching estimator suggested by Abadie and Imbens (2002).[7]

[6]For some categories, the length of time since the first visit to a store does not vary across households because all of them started to visit the store even before the data period. In this case, the dummy for first visit during the previous time more effectively captures store loyalty.

[7]They show that, with k continuous covariates, the simple matching estimator has a bias term caused by the matching discrepancy between the treatment group and their matches in the control group.

Table 5: Within-Category Effects of Price Promotions on Expenditure

	Regression estimator	Matching estimator	Number of treated	Number of controls
Soft-drinks				
1 month	3.94	3.68	1,931	2,279
	(0.28)	(0.42)		
3 months	3.73	3.71	1,931	2,471
	(0.24)	(0.33)		
6 months	2.95	2.93	1,931	2,938
	(0.18)	(0.24)		
Cereal				
1 month	2.65	2.62	713	2,874
	(0.27)	(0.34)		
3 months	1.75	1.94	713	2,943
	(0.18)	(0.21)		
6 months	1.32	1.44	713	3,334
	(0.15)	(0.18)		
Laundry detergent				
1 month	0.89	1.12	232	3,082
	(0.24)	(0.29)		
3 months	0.73	0.64	232	3,111
	(0.15)	(0.16)		
6 months	0.50	0.59	232	3,474
	(0.12)	(0.14)		
Salted snack				
1 month	2.22	3.72	1,540	2,499
	(0.15)	(0.21)		
3 months	1.92	3.32	1,540	2,650
	(0.10)	(0.15)		
6 months	1.71	1.55	1,540	3,093
	(0.09)	(0.11)		

The estimates in the tables represent the average differences in monthly spending on each category between the treated and the control groups. They also report short term and long term effects of price discounts for different time frames. Table 5 presents the promotion effects within the same category. The estimates evaluate the extent to which price promotions in a category increase future spending for the same category. The estimates from the matching estimator are quite similar to those from the regression approach and all statistically significant. Compared to the average monthly spending on each category presented in Table 3, the increase in monthly sales for the next one month induced by price promotion is about 23% on average for soft-drink, laundry detergent, and salted

Table 6: Spillover Effects of Price Promotions on Expenditure

	Regression estimator	Matching estimator	Number of treated	Number of controls
Soft-drinks				
1 month	5.21 (0.66)	3.76 (0.94)	1,931	2,279
3 months	2.56 (0.51)	1.90 (0.69)	1,931	2,471
6 months	3.97 (0.42)	3.62 (0.63)	1,931	2,938
Cereal				
1 month	6.72 (1.25)	7.71 (1.51)	713	2,874
3 months	2.89 (0.93)	3.32 (1.18)	713	2,943
6 months	3.24 (0.78)	2.79 (0.91)	713	3,334
Laundry detergent				
1 month	6.51 (2.80)	7.33 (3.45)	232	3,082
3 months	5.71 (2.05)	6.34 (2.61)	232	3,111
6 months	4.24 (1.71)	4.66 (1.95)	232	3,474
Salted snack				
1 month	7.41 (0.88)	5.92 (1.18)	1,540	2,499
3 months	6.65 (0.72)	6.19 (0.90)	1,540	2,650
6 months	5.89 (0.63)	5.79 (0.83)	1,540	3,093

snack, and 17% for cereal based on the regression estimator. Promotion effects decrease but do not vanish over time for all of the four categories.

Table 6 reports the estimates of the spillover effects of price promotion. The estimates represent the average of future spending in other categories that are triggered by price discounts in each category. All of the estimates are statistically significant. For all of the four categories, the effects of price promotion that spill over into other categories are greater than the within-category effects presented in Table 5. Although price cuts drive a smaller increase in future spending within the same product category, they raise store sales in other categories to a large extent. This is particularly true for the laundry detergent category Spillover effect over the next one month after price promotion on these categories is on average 10% increase in storewide sales for the 31 categories included in the data.

5. CONCLUSION

Using the treatment effects setting we examine the short-term and long-term promotional effects that spill over into other product categories using supermarket scanner data. The detailed information on individual households' purchases and store visits allows us to improve the quality of imputation for the missing potential outcomes in evaluating the treatment effects. We find that the spillover effects across different categories are substantial compared to the within-category effects. This suggests that pricing strategies of multi-product retailers should take into account the spillover effects of short-term pricing as well as the direct promotion effects within the promoted category. In our current framework we consider only a binary treatment status whether a household has purchased a promoted item or not. In a future work we can also extend our analysis to a multi-valued treatment effect framework where we evaluate the effects of promotions, which depend on levels of households' responses to the promotions given a time period.

REFERENCES

Abadie, A., and G. W. Imbens (2002). Simple and Bias-Corrected Matching Estimators for Average Treatment Effects, NBER working paper no. 283.

Abadie, A., and G. W. Imbens (2006). Large Sample Properties of Matching Estimators for Average Treatment Effects, *Econometrica* 74, 235–267.

Ailawadi, K. L., and S. A. Neslin (1998). The Effect of Promotion on Consumption: Buying More and Consuming It Faster, *Journal of Marketing Journal* 35, 390–398.

Bell, D. R., T. H. Ho, and C. S. Tang (1998). Determining Where to Shop: Fixed and Variable Costs of Shopping, *Journal of Marketing Research* 35, 352–369.

Blattberg, R. C., R. Briesch, and E. J. Fox (1995). How Promotions Work, *Marketing Science* 14, G122–G132.

Erdem, T. (1996). A Dynamic Analysis of Market Structure Based on Panel Data, *Marketing Science* 15, 359–378.

Erdem, T., S. Imai, and M. P. Keane (2003). Brand and Quantity Choice Dynamics Under Price Uncertainty, *Quantitative Marketing and Economics* 1, 5–64.

Hahn, J. (1998). On the Role of the Propensity Score in Efficient Semiparametric Estimation of Average Treatment Effects, *Econometrica* 66, 315–331.

Hendel, I., and A. Nevo (2003). The Post-Promotion Dip Puzzle: What do the Data Have to Say?, *Quantitative Marketing and Economics* 1, 409–424.

Hendel, I., and A. Nevo (2006). Measuring the Implications of Sales and Consumer Inventory Behavior, *Econometrica* 74, 1637–1673.

Imbens, G. W. (2004). Nonparametric Estimation of Average Treatment Effects under Exogeneity: A Review, *Review of Economics and Statistics* 86, 4–29.

Imbens, G. W., W. Newey, and G. Ridder (2003). Mean-Squared-Error Calculations for Average Treatment Effects, Working paper.

Kalyanaram, G., and R. S. Winer (1995). Empirical Generalizations from Reference Price Research, *Marketing Science* 14, G161–G169.

Kim, H., and K. Kim (2015). Estimating Store Choices with Endogenous Shopping Bundles and Price Uncertainty, Working paper.

Mela, C. F., S. Gupta, and D. R. Lehmann (1997). The Long-Term Impact of Promotion and Advertising on Consumer Brand Choice, *Journal of Marketing Research* 34, 248–261.

Niraj, R., V. Padmanabhan, and P. B. Seetharaman (2008). A Cross-Category Model of Households' Incidence and Quantity Decisions, *Marketing Science* 27, 225–235.

Pesendorfer, M. (2002). Retail Sales: A Study of Pricing Behavior in Supermarkets, *Journal of Business* 75, 33–66.

Rosenbaum, P. R., and D. B. Rubin (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects, *Biometrika* 70, 41–55.

Rosenbaum, P. R., and D. B. Rubin (1984). Reducing Bias in Observational Studies Using Subclassification on the Propensity Score, *Journal of the American Statistical Association* 79, 516–524.

Van Heerde, H. J., P. S. H. Leeflang, and D. R. Wittink (2000). The Estimation of Pre- and Post-promotion Dips with Store-Level Scanner Data, *Journal of Marketing Research* 37, 383–395.

Walters, R. G. (1991). Assessing the Impacts of Retail Price Promotions on Product Substitution, Complementary Purchase, and Interstore Sales Displacement, *Journal of Marketing* 55, 17–28.