Appendix: Generalizable and Robust TV Advertising Effects*

Bradley T. Shapiro
University of Chicago – Booth

Günter J. Hitsch
University of Chicago – Booth

Anna E. Tuchman
Northwestern University – Kellogg

August 5, 2020

A Advertising elasticities

To illustrate the possible interpretations of β, we drop the store and market indices and focus on one specific advertising component, a_t, with corresponding coefficient β. The elasticity of demand in period t with respect to advertising in period $\tau \in \{t-L,\ldots,t\}$ is given by

$$\frac{\partial Q_t}{\partial a_\tau}\frac{a_\tau}{Q_t} = \beta\delta^{t-\tau}\frac{a_\tau}{1+A_t}.$$

Furthermore, the advertising stock elasticity is equivalent to the total sum of the advertising elasticities:

$$\frac{\partial Q_t}{\partial A_t}\frac{A_t}{Q_t} = \beta\frac{A_t}{1+A_t} = \sum_{\tau=t-L}^{t}\frac{\partial Q_t}{\partial a_\tau}\frac{a_\tau}{Q_t}.$$

To further clarify the difference between the short-run and long-run effect of advertising, suppose that advertising is constant at the level $a_t \equiv a$, such that $A_t = \rho a$ in all periods t, where $\rho = (1-\delta)^{-1}(1-\delta^{L+1})$. Then the elasticity of per-period demand with respect to the constant

*All three authors contributed equally although not listed in alphabetical order. We acknowledge the superb research assistance of Jihong Song and Ningyin Xu. We thank Liran Einav, Paul Ellickson, Jeremy Fox, Wes Hartmann, Carl Mela, Matt Shum, and Sha Yang for helpful comments. We also benefited from the comments of seminar participants at Amazon, Bates White, Columbia, CUHK, HKUST, Johns Hopkins, NUS, Rice, UNC, UCSD, Yale, Marketing Science, the MSI Young Scholars Conference, the Wash U. Junior Faculty Development Forum, the NYC Media Seminar, the $(IO)^2$ Zoom Seminar and the 12th Workshop on the Economics of Advertising and Marketing. Calculated (or derived) based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

advertising flow a is

$$\frac{dQ_t}{da}\frac{a}{Q_t} = \beta\frac{\rho a}{1+\rho a}. \tag{1}$$

This elasticity measures the effect of a permanent percentage increase in advertising. Similarly, assuming again that $a_t = a$ in all periods t, and also that all other factors affecting demand (prices, etc.) are constant, we can derive the effect of a current increase in advertising at time t on total or long-run demand in periods $t, \ldots, t+L$:

$$\left(\frac{\partial}{\partial a_t}\sum_{\tau=t}^{t+L} Q_\tau\right)\frac{a_t}{Q_t} = \beta\frac{\rho a}{1+\rho a}. \tag{2}$$

The effect of permanent percentage increase in advertising (1) is equivalent to the cumulative, long-run increase in demand (2). Both effects are bounded above by β and will be approximately equal to β if the advertising stock, ρa, is large. For example, if $\delta = 0.9$, $L = 52$, and advertising $a = 20$ GRPs, then $\rho a/(1+\rho a) = 0.995$, and the long-run demand effect is well approximated by β.

The short-run advertising elasticity is

$$\frac{\partial Q_t}{\partial a_t}\frac{a_t}{Q_t} = \beta\frac{a_t}{1+A_t}.$$

If $a_t = a$ in all periods t and if the advertising stock is large, then

$$\frac{\partial Q_t}{\partial a_t}\frac{a_t}{Q_t} = \beta\frac{a}{1+\rho a} \approx \beta\frac{a}{\rho a}.$$

Hence, the ratio of the long-run effect to the short-run effect of advertising is ρ, which is approximately equal to $1/(1-\delta)$ if δ^L is small.

B Data construction

The objective of this project is to estimate the effect of TV advertising on retail sales for a wide range of brands. To do that, we need the following data for each brand:

- Weekly volume, price, promotion, and feature/display at store or market level.
- Weekly advertising (GRP, duration, or spending) at television market (DMA) level.

We create the data we want in the following steps:

1. Build Ad Intel Data
    (a) The ad occurrences and viewerships are separate in the raw Ad Intel data. We need to merge them in order to find the GRP for each advertisement.

   (b) There are some discrepancies between the national and local records of Network TV ads. We need to resolve those discrepancies.

2. Create brand map between Ad Intel and RMS data sets.

   (a) Ad Intel and RMS use different brand definitions, so for each RMS brand, we need to find all the corresponding Ad Intel brands.

3. Aggregate Data

   (a) RMS data come in UPC-Store-Week level, so we need to aggregate it to Brand-Store-Week level.

   (b) Ad Intel data come in {Ad Intel Brand}-Market-Channel-Second level, so we need to aggregate it to {RMS Brand}-Market-Week level.

4. Identify RMS stores to be used in estimation

5. Identify products to be used in estimation

Each of these steps is described in more detail below.

B.1 Build Ad Intel Data

B.1.1 General Concepts

Media Types Ad Intel covers 4 TV media types: Cable, Network, Syndicated, and Spot.

- For Cable TV, ads are purchased at a national level.
- For Network and Syndicated TV, ads are purchased at a national level. The programs are broadcast at local TV stations.
  - The local TV stations are typically affiliated to a national network. For example, WBZ is the Boston affiliate of CBS.
- For Spot TV, ads are purchased at the DMA level. The programs are also broadcast at local TV stations.

Since Network and Syndicated TV ads are purchased nationally but broadcast locally, the Ad Intel record them in two ways:

- The Network TV and Syndicated TV occurrence files record them at national level.
  - i.e., the date and time each ad is supposed to be broadcast at every local station
- The Network Clearance Spot TV and Syndicated Clearance Spot TV occurrence files record them at local channel level.

  - i.e., the date and time each ad is actually broadcast at every local station

- The local channels have some authority to replace or move nationally scheduled ads, and the Nielsen data are also not perfect. Hence there are discrepancies between those national and local files.

Occurrence Data The occurrence data provide detailed information for each advertisement, including:

- Date [AdDate]
- Time [AdTime]
  - Note that Ad Intel does not capture any local ads between 2AM and 5AM.
- Media Type [MediaTypeID]
- Channel [DistributorCode, DistributorID]
- Market (can be national) [MarketCode]
- Primary, Secondary, and Tertiary Brands [PrimBrandCode, ScndBrandCode, TerBrandCode]
- Duration [Duration]
- The associated TV program [NielsenProgramCode, TelecastNumber]
- Other time-related info [TVDayPartCode, DayOfWeek, TimeIntervalNumber]

Impression (Viewership) Data For the national media types (Cable, Network, and Syndicated), Ad Intel provides the estimated impression for each TV program--defined as a pair of NielsenProgramCode and TelecastNumber.

For the local media types (Network Clearance, Syndicated Clearance, and Spot), Ad Intel provides the estimated impression at {Local Station}-Month-{Day of Week}-{5 Minute Time Interval} level.

Note: There are only 25 markets (the "Local People Meter" markets) for which the local impressions are available in all months. For the rest of the markets, local impressions data are only available in four "sweeps months": February, May, July, and November. Therefore, we need to impute the impressions for the non-sweeps months in non-LPM markets. Now we use an average between the two closest available months, weighted by the time difference. For example, for June we use 1/2 May + 1/2 July, and for March we use 2/3 February + 1/3 May.

Universe Estimates Ad Intel also provides the estimated total number of TV audience at national and market level. Those universe estimates are updated yearly.

B.1.2 Build the Regular Parts

The logic of the regular build is very simple. For each media type in each month, we need to do the following:

1. Merge occurrences with impressions
   (a) For national data, merge on NielsenProgramCode and TelecastNumber
   (b) For local data, merge on DistributorID, DayOfWeek, and TimeIntervalNumber
   (c) Remember to do the imputation for non-LPM markets in non-sweep months.
2. Merge the result with universe estimates
3. Calculate the GRP as 100 * Impression / Universe for each ad occurrence

B.1.3 Resolve the "Missing Network" Discrepancy

The objective of this part is also simple: we need to find the national Network TV ads that are not recorded in the Network Clearance data, and if the missing cannot be reasonably explained, we believe that the local data are wrong, and we add those "unexpectedly missing" occurrences into the local records. We say a national ad is "expectedly missing" if it is replaced by another local ad, or if it's scheduled air-time is between 2AM and 5AM. In practice, this procedure is quite complicated to implement. We take the following steps:

1. Find the information for each local station, including:
   (a) The market (MarketCode) and network (Affiliation) for each local station (DistributorCode).
   (b) The DistributorID for each DistributorCode.
      i. This is in fact a one-to-one relationship, but we have to record that because the "Station Affiliation" data only has DistributorCode, while the impressions data only have DistributorID.
2. For each network and each local station, stack all the monthly data.
   (a) We cannot use the raw monthly data because the national and local files have different dates.
   (b) Stacking also prevents errors at month boundaries. For example, a national ad at 2012/05/31 23:30:00 may be distributed locally at 2012/06/01 00:30:00. This will not be captured if we process the data month-by-month.
3. For each local station, find the "unexpectedly missing" occurrences. In short, we categorize all the national ads as following:

(a) A national ad is directly matched to the local data if its closest local occurrence has the same primary brand code.

(b) A national ad is indirectly matched to the local data if there's a local occurrence that is aired within some time limit before or after the scheduled air-time. This step accounts for the ads that are moved around. The time limit is 3 hours for ETZ/CTZ, 6 hours for MTZ, and 7 hours for PTZ.

(c) A national ad is replaced by another ad if another spot / network clearance / syndicated clearance ad runs into its scheduled time slot.

(d) A national ad is not captured locally if its scheduled air-time is between 2AM and 5AM.

(e) We mark all remaining national ads as unexpectedly missing at this local station.

4. We get all the "unexpectedly missing" occurrences at each station, and we reorganize them into monthly files. We then merge those monthly files with the monthly local impressions data.

Note: The "broadcast delay" for mountain and pacific time zones causes trouble.

- A nationally scheduled program or ad can be broadcast with a delay of 0/1/2/3 hours in pacific-time markets or 0/1 hours in mountain-time markets. This delay can be pretty arbitrary.

- In step 3, we say a national ad is "unexpectedly missing" only if it is "unexpectedly missing" under all the possible delays, i.e. 0/1 hour in MTZ and 0/1/2/3 hours in PTZ.

- In step 4, for PTZ/MTZ markets we average the impressions at the airtime and 3/1 hours after the airtime.

B.2 Create Brand Map between RMS and Ad Intel

We create a map between the brands in the RMS and Ad Intel data sets using string matching. We classify the matches in 4 "tiers," which are described below. In theory, tier-1 and tier-2 advertising should have a positive effect on sales, while the effect of tier-3 and tier-4 ads can be either positive or negative. Based on this logic, we construct our measure of own advertising by grouping tier-1 and tier-2 advertising together and our measure of “affiliated brands” advertising by grouping tier-3 and tier-4 advertising together.

Own Advertising

- Tier 1: RMS and Ad Intel brand names are exact matches.

- Tier 2: Ad Intel brand is more specific than the RMS brand.

Table 1: Fraction of Own Advertising GRPs from Different Matches

	Median	Mean	Percentiles			
			10%	25%	75%	90%
Exact Matches	43.3214	47.8728	0	0.4212	97.2967	100
Inexact Matches	56.6786	52.1272	0	2.7033	99.5788	100

  - Example: Ad Intel brand LAYS POTATO CHIPS CHICKEN AND WAFFLE is a tier-2 match to RMS brand LAY'S.

- For the median brand in our data, tier-1 matches make up 43% of own advertising GRPs, while tier-2 matches make up the remaining 57% of own advertising GRPs. The distribution of GRPs coming from identical matches is shown in Table 1.

Affiliated Brands Advertising

- Tier 3: Ad Intel brand is more general than the RMS brand.
  - Example: Ad Intel brand COCA-COLA SOFT DRINKS is a tier-3 match to RMS brand COCA-COLA R.
- Tier 4: Ad Intel brand is an "associate" to the RMS brand.
  - Example: Ad Intel brand COCA-COLA ZERO DT is a tier-4 match to RMS brand COCA-COLA R.

We also carry out some module aggregation, which amounts to aggregating some very specific RMS modules together. For example, the RMS modules NUTS-BAGS, NUTS-CANS, NUTS-JARS, and NUTS-UNSHELLED are essentially the same thing, and advertisements never distinguish between them.

Finally, we do some aggregation across flavors and sub-brands. For example, the brand "Lean Cuisine Frozen Entree" has 50 sub-brands in RMS (e.g., LEAN CUISINE ONE DISH FAVORITE or LEAN CUISINE SPA COLLECTION). Aggregating them together makes the matching easier and creates more tier-2 matches and fewer tiers-3/4 matches.

B.3 Aggregate Data

Ad Intel The Ad Intel data build comes at the {Ad Intel Brand}-Channel-Time level, and in the end we want to aggregate it to the {RMS Brand}-Market-Week level.

First, we aggregate the ad data to the {Ad Intel Brand}-{Media Type}-Market-Week level. The aggregation here only involves adding up Duration and GRP.

- Some ad occurrences come with 2/3 brands, but those brands are mostly the same product (e.g., Snapple Black Tea and Snapple Green Tea). To avoid double-counting the ads, we use the following trick: if an occurrence has two/three brands, treat it as two/three occurrences with half/one-third of the Duration and GRP.

RMS The RMS data build comes at UPC-Store-Week level, and we want to aggregate it to Brand-Store-Week level.

- One RMS brand may contain hundreds of UPCs with different sizes (size1_amount, say 12 OZ or 24 OZ) and different multi-pack status (multi, say 6-pack or 12-pack).
  - Therefore, instead of using the units field in the RMS data, we need to calculate the volume in equivalency units: volume: = units*multi*size1_amount. We adjust price accordingly.
- For each store-week, the brand-level variables are calculated as follows:
  - Volume: sum of UPC-level volumes
  - Price: weighted average of UPC-level prices. The weight for a UPC is its average weekly revenue in this store.
  - Promotion: weighted average of UPC-level promotion indicators (price / base_price < 0.95).
  - Feature/Display: weighted average of UPC-level feature/display indicators (remove missing values).

B.4 Store and Border Selection

We remove the stores that switch between different counties and stores that are not continuously tracked by Nielsen between 2010–2014. We then rank the stores by the total 2010–2014 revenue (across all products), and find the stores that constitute 90% of total revenue. We use those stores for all of our analyses.

Nielsen provides a mapping between counties and DMAs. From this, we constructed a data set that flags the counties that lie on a border between DMAs. However, some counties change DMAs over time, since the borders are re-drawn periodically. Therefore, we removed all the counties that did not stay in a single DMA, and we removed the borders that were re-drawn.

B.5 Product Selection

We began our analysis with the top 500 national brands in the RMS data based on sales revenue between 2010–2014. The above flavor and module aggregation steps reduce the count of unique brands somewhat. We are able to match 358 of these aggregated RMS brands to brands in the Ad Intel data.

Table 2: Frequency of Departments and Revenue Share

Department	No. of brands	Homescan revenue share
DRY GROCERY	127	52.19
NON-FOOD GROCERY	50	13.47
HEALTH & BEAUTY CARE	33	4.39
FROZEN FOODS	23	10.75
DAIRY	21	9.90
ALCOHOLIC BEVERAGES	19	3.49
PACKAGED MEAT	11	3.83
DELI	5	2.24
FRESH PRODUCE	1	0.14
GENERAL MERCHANDISE	1	0.20

Note: Three brands in our sample have products in two departments.

Screening Based on Own Advertising For each of the 358 RMS brands in our universe, we calculate the fraction of market-weeks with positive own advertising GRPs, and the mean own advertising GRPs conditional on it being positive. We drop 70 brands that have positive GRPs in less than 5% of observations, or whose "positive mean" is below 10 GRPs.

C Affiliated brand and competitor advertising elasticities

In the main text we reported own-advertising elasticity estimates. All model specifications also control for “affiliated brand” advertising and top competitor advertising.[1] We now discuss the corresponding affiliated brand and competitor advertising effect estimates. While theory predicts that own-advertising effects should typically be positive, the direction of the affiliated brand and competitive advertising effects are both ambiguous. For affiliated brand products, the ad is relevant both to the focal product and other products that are potentially substitutes. If the partial ad effect on the substitutes is of equal or greater magnitude than the partial ad effect on the focal product, the net ad effect on the focal product could be negative. For example, a Coke Zero ad could reinforce the general Coca-Cola brand and lead to an increase in sales of Diet Coke, which would reflect a positive ad effect. But Coke Zero ads could also lead some consumers to buy Coke Zero instead of Diet Coke, which would appear as a negative ad effect. With regard to competitor ad effects, the previous literature has similarly found mixed results. Some papers have shown positive spillovers of advertising (e.g., Sahni 2016, Shapiro 2018, and Lewis and Nguyen 2015), while others have shown negative, business stealing effects (Sinkinson and Starc 2019). Advertising for a direct substitute may steal sales from the focal brand. However, a

[1] The top competitor is the competitor brand with the largest market share in the same product module.

Table 3: Frequency of Categories

Category	No. of brands	Category	No. of brands
PAPER PRODUCTS	16	VEGETABLES-FROZEN	3
SNACKS	13	CHEESE	3
CARBONATED BEVERAGES	13	LAUNDRY SUPPLIES	3
BEER	11	SANITARY PROTECTION	3
DETERGENTS	11	WRAPPING MATERIALS AND BAGS	3
CANDY	11	DEODORANT	3
JUICE, DRINKS - CANNED, BOTTLED	10	NUTS	3
PACKAGED MEATS-DELI	10	BABY FOOD	2
SOFT DRINKS-NON-CARBONATED	9	PREPARED FOOD-DRY MIXES	2
CEREAL	9	COOKIES	2
PREPARED FOODS-FROZEN	7	UNPREP MEAT/POULTRY/SEAFOOD-FRZN	2
SALAD DRESSINGS, MAYO, TOPPINGS	6	COT CHEESE, SOUR CREAM, TOPPINGS	2
PET FOOD	6	PACKAGED MILK AND MODIFIERS	2
BREAKFAST FOOD	6	WINE	2
LIQUOR	6	HOUSEHOLD SUPPLIES	2
VITAMINS	6	PET CARE	2
MEDICATIONS/REMEDIES/HEALTH AIDS	6	SKIN CARE PREPARATIONS	2
DISPOSABLE DIAPERS	6	SEAFOOD - CANNED	1
CONDIMENTS, GRAVIES, AND SAUCES	5	PREPARED FOOD-READY-TO-SERVE	1
CRACKERS	5	JAMS, JELLIES, SPREADS	1
COFFEE	5	DESSERTS, GELATINS, SYRUP	1
PIZZA/SNACKS/HORS DOEURVES-FRZN	5	TEA	1
DRESSINGS/SALADS/PREP FOODS-DELI	5	SPICES, SEASONING, EXTRACTS	1
YOGURT	5	FRESH MEAT	1
COUGH AND COLD REMEDIES	4	PUDDING, DESSERTS-DAIRY	1
ICE CREAM, NOVELTIES	4	EGGS	1
BUTTER AND MARGARINE	4	FRESH PRODUCE	1
MILK	4	PERSONAL SOAP AND BATH ADDITIVES	1
ORAL HYGIENE	4	CHARCOAL, LOGS, ACCESSORIES	1
HAIR CARE	4	STATIONERY, SCHOOL SUPPLIES	1
FRESHENERS AND DEODORIZERS	4	TOBACCO & ACCESSORIES	1
BREAD AND BAKED GOODS	4	FIRST AID	1
SOUP	3	PASTA	1
GUM	3	VEGETABLES - CANNED	1
BREAKFAST FOODS-FROZEN	3	DOUGH PRODUCTS	1

Note: Four brands in our sample have products in two categories.

Table 4: Affiliated Brand, Top Competitor Advertising Stock Elasticities and Other Controls

	Median	**Mean**	**% Brands**	**% p $\geq$ 0.05**	**% p $<$ 0.05**		**Percentiles**			
					$>$ **0**	$\leq$ **0**	**10%**	**25%**	**75%**	**90%**
Border Strategy										
Affiliated Brand Advertising	-0.0012	0.0167	58.68	73.37	11.24	15.38	-0.044	-0.017	0.016	0.049
Top Competitor Advertising	0.0082	-0.0087	66.32	78.01	10.47	11.52	-0.106	-0.021	0.044	0.094
Own Price Elasticity	-1.6104	-1.6558	100.00	4.17	3.82	92.01	-3.159	-2.359	-1.016	-0.284
Top Competitor Price Elasticity	0.0938	0.1256	87.85	39.92	44.66	15.42	-0.216	-0.052	0.321	0.618
Baseline										
Affiliated Brand Advertising	-0.0010	0.0050	58.68	70.41	14.79	14.79	-0.051	-0.019	0.016	0.050
Top Competitor Advertising	0.0028	-0.0025	66.67	79.17	10.42	10.42	-0.125	-0.036	0.035	0.109
Own Price Elasticity	-1.5760	-1.6447	100.00	2.08	3.82	94.10	-3.040	-2.327	-1.061	-0.422
Top Competitor Price Elasticity	0.1025	0.1372	87.85	37.15	45.45	17.39	-0.254	-0.052	0.335	0.623

Note: The estimates are obtained assuming a carryover parameter $\delta = 0.9$. Standard errors are two-way clustered at the border-side level and the week level in the border strategy specification.

competitor brand's ads may also bring new customers into the category and could therefore lead to an increase in sales for the focal brand. The net effect of these different forces depends on the relative strength of these two advertising effects.

Table 4 shows summary statistics for the estimated affiliated brand and top competitor advertising elasticities corresponding to the baseline and borders model specifications in equations (**??**) and (**??**), and Figures 1 and 2 show histograms of the corresponding distributions of advertising effects. Table 4 also shows own price elasticities and top competitor price elasticities.

The distributions of both the affiliated brand and competitor advertising elasticities are centered at zero and the competitor advertising elasticity distribution is relatively disperse. That is, the particulars of what causes affiliated and competitor advertising to help or hurt own demand is likely case dependent. Results from past case studies are unlikely to be a good guide for predicting whether any particular affiliated brand or competitor advertising elasticity will be positive or negative. The location and shape of the distributions is similar between the baseline and the border strategy approaches.

Own price elasticities are centered around -1.6 for both the border and baseline strategies, with almost all of the mass less than zero, as expected. Top competitor price elasticities are centered around 0.1 in each strategy. These results largely replicate those in .

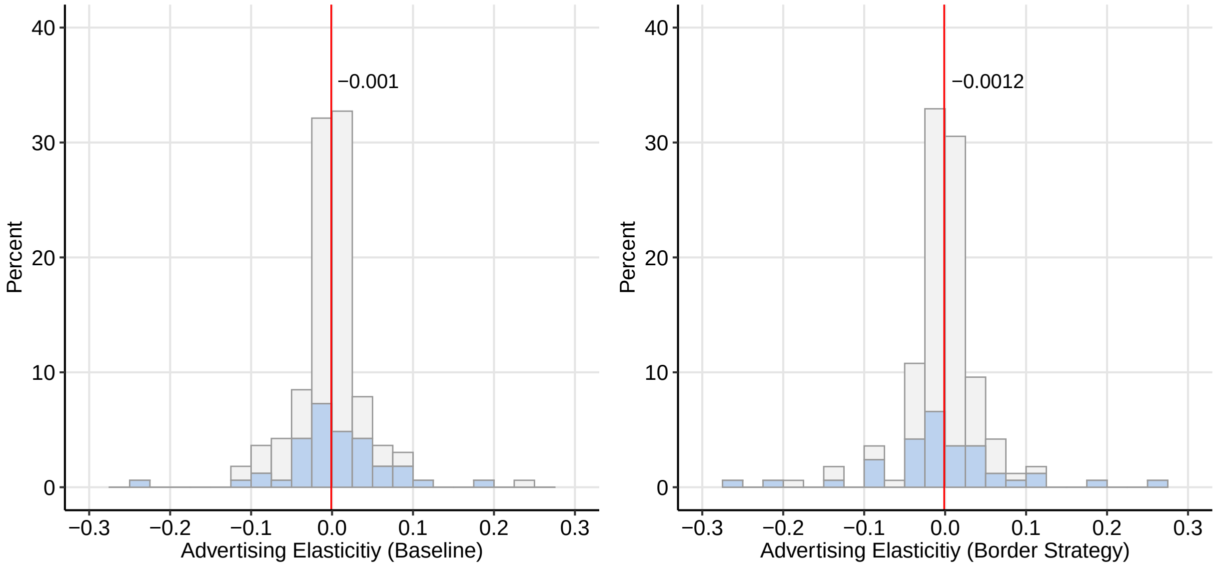

Figure 1: Affiliated Brand Advertising Stock Elasticities

Note: The estimates are obtained assuming a carryover parameter $\delta = 0.9$. Bars highlighted in blue indicate statistically significant estimates. The vertical red line denotes the median of the distribution.

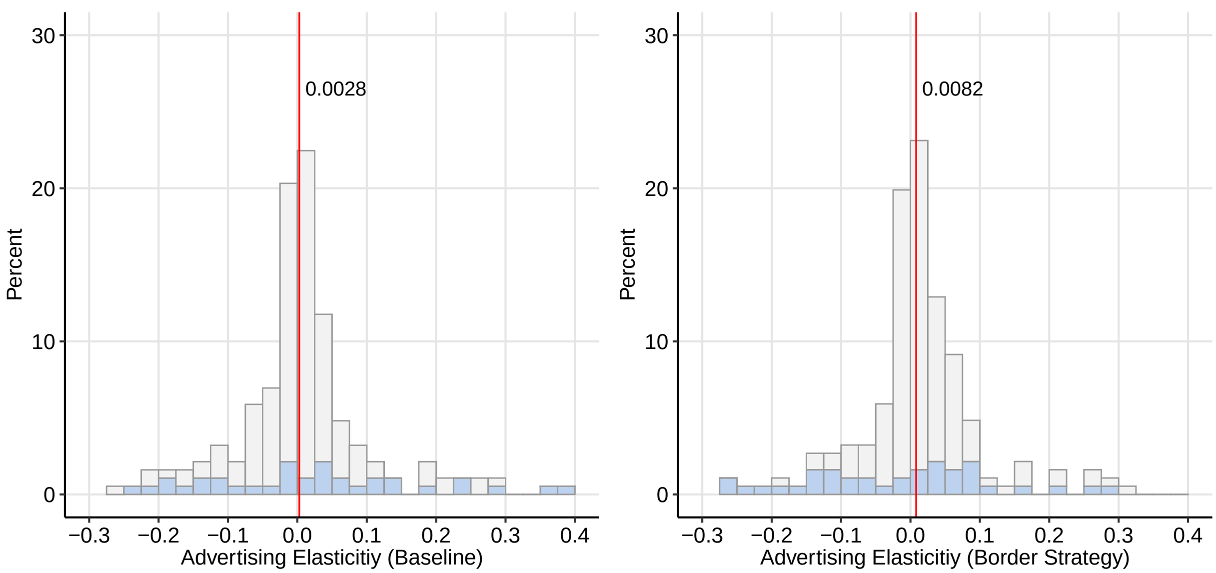

Figure 2: Top Competitor Advertising Stock Elasticities

Note: The estimates are obtained assuming a carryover parameter $\delta = 0.9$. Bars highlighted in blue indicate statistically significant estimates. The vertical red line denotes the median of the distribution.

D Correlation between advertising and other variables

In our main specification, we include own price but exclude feature and display advertising. We effectively treat price, feature and display as exogenous conditional on the fixed effects and other covariates in the model. These choices could be problematic if net of fixed effects, these variables are correlated with both advertising and the error term. In this appendix, we show the degree to which advertising is correlated with prices (both in general and temporary price reductions in particular) and feature and display advertising. We show these correlations both unconditionally and conditional on the fixed effects and covariates in our models.

Figure 3 shows that while the unconditional correlation between price and advertising is non-zero, the correlation conditional on the fixed effects in the model is approximately zero. The right panel shows the regression coefficients for a regression of advertising flow on price for the baseline and border samples, with no fixed effects. The left panel shows the regression coefficients for a regression of advertising flow on price plus the fixed effects included in the baseline and border strategies, respectively. Without the fixed effects in the model, the median regression coefficients are –0.09885 and –0.08295 in the baseline and border samples, respectively, and each has about 35% of estimates negative and significant and about 17% positive and significant. Thus, it appears that many firms are coordinating advertising with price reductions and many others with price increases. However, when we add the fixed effects, the median regression coefficient is about –0.007 on each, with less than 10% negative significant and less than 5% positive significant. Conditional on the fixed effects, advertising and price are approximately uncorrelated. This is consistent with the fact that our ad elasticity distribution does not change significantly depending on whether or not price is included.

We conduct similar analysis for temporary price reductions, which we call *promotions* (Figure 4), feature advertising (Figure 5), and display advertising (Figure 6). The results are similar to the price results. Summary statistics describing all of the figures are available in Table 5.

We also conduct the same analysis, but using advertising stock instead of advertising flow. While the concern about coordination relates to flows, concerns of bias relate to the correlation of the potentially troublesome variables to the error term and the treatment variable of interest, which is advertising stock. We find that advertising stock has an even lower magnitude of correlation to these variables than does advertising flow, net of the fixed effects in the two main strategies.

Finally, we estimate a regression with advertising on the left-hand side, all of the fixed effects and covariates of the border strategy on the right-hand side, as well as price, promotion, feature and display. From that regression, we compute the partial R^2 of price, promo, feature and display combined to determine how much of the residual variation in advertising is explained by those four variables. The result is reported in Figure 11. The partial R^2 for almost all brands is less than 0.005, with most being much less than that. Conducting the same exercise with advertising stock in Figure 12 shows similar results.

Overall, we conclude that net of the fixed effects in our main specifications, price, feature

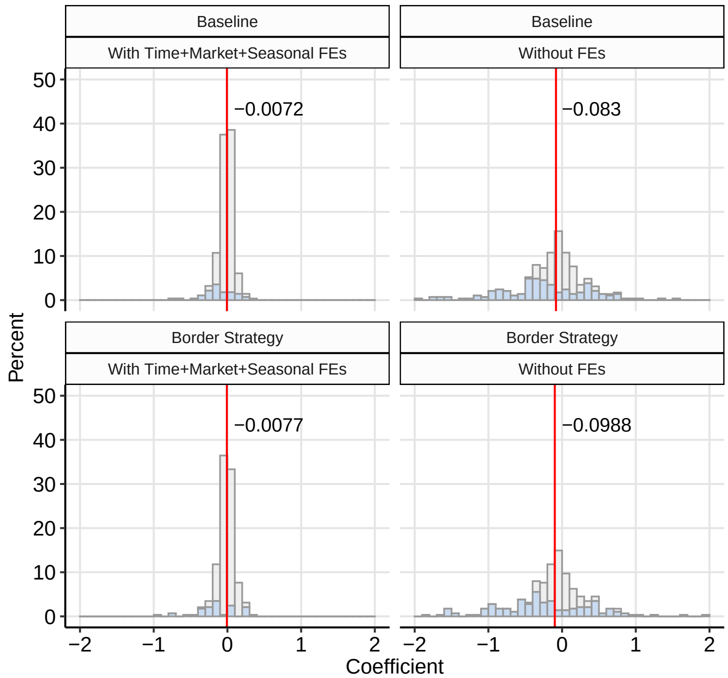

Figure 3: Advertising and Price Correlation

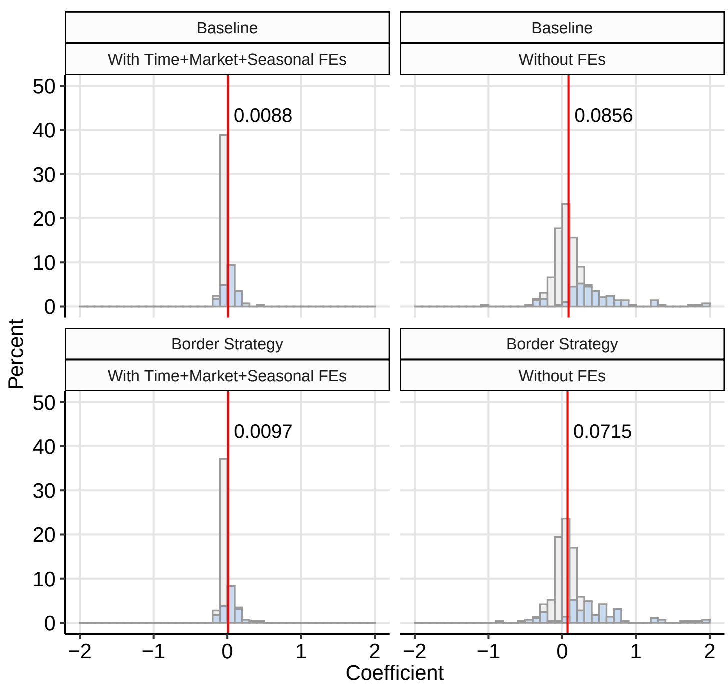

Figure 4: Advertising and Promotions Correlation

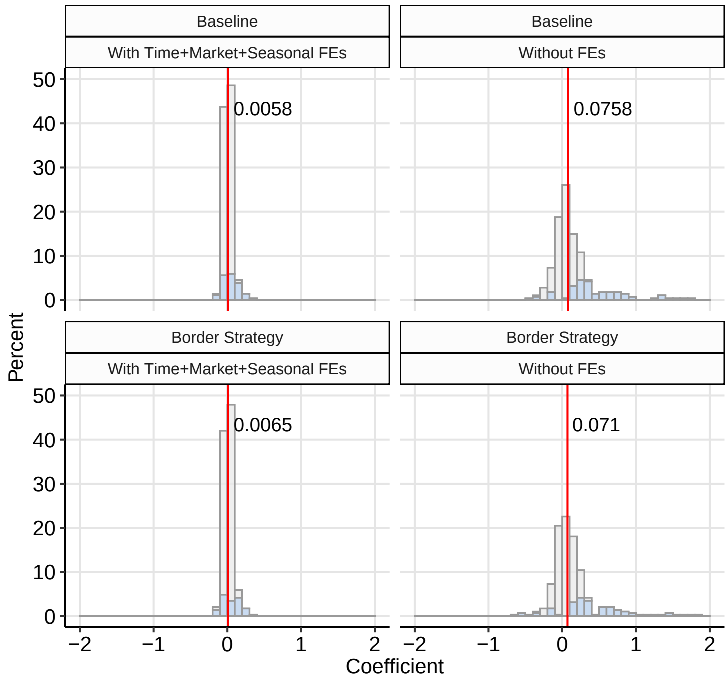

Figure 5: Advertising and Feature Correlation

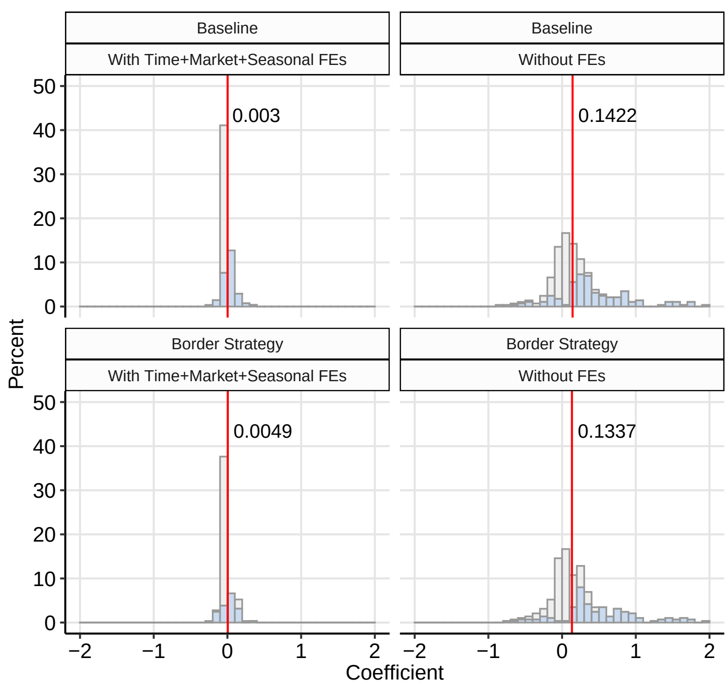

Figure 6: Advertising and Display Correlation

Table 5: Correlations Between Advertising and Price, Promotions Feature and Display

	Median	**Mean**	**% p $\geq$ 0.05**	**% p $<$ 0.05**		**Percentiles**			
				> 0	≤ 0	**10%**	**25%**	**75%**	**90%**
Baseline									
price									
without FEs	-0.0830	-0.1698	46.53	17.71	35.76	-0.8586	-0.3504	0.1172	0.4003
with FEs	-0.0072	-0.0190	86.43	3.93	9.64	-0.1387	-0.0630	0.0420	0.0864
promo									
without FEs	0.0856	0.2431	63.19	32.64	4.17	-0.1092	-0.0280	0.2768	0.6468
with FEs	0.0088	0.0112	79.51	13.89	6.60	-0.0441	-0.0150	0.0335	0.0645
feature									
without FEs	0.0758	0.2061	71.53	25.69	2.78	-0.1187	-0.0143	0.2226	0.6315
with FEs	0.0058	0.0131	81.94	11.46	6.60	-0.0528	-0.0212	0.0373	0.0788
display									
without FEs	0.1422	0.3658	49.31	42.71	7.99	-0.1491	-0.0082	0.3738	0.8642
with FEs	0.0030	0.0088	74.18	16.73	9.09	-0.0385	-0.0156	0.0280	0.0638
Border Strategy									
price									
without FEs	-0.0988	-0.1807	47.57	17.01	35.42	-0.9192	-0.3596	0.1024	0.4255
with FEs	-0.0077	-0.0250	86.11	4.51	9.38	-0.1645	-0.0777	0.0459	0.1074
promo									
without FEs	0.0715	0.2312	63.19	31.25	5.56	-0.1312	-0.0294	0.2580	0.6637
with FEs	0.0097	0.0128	81.60	12.85	5.56	-0.0514	-0.0150	0.0396	0.0687
feature									
without FEs	0.0710	0.1991	72.22	23.61	4.17	-0.1077	-0.0187	0.2153	0.6501
with FEs	0.0065	0.0124	84.03	9.72	6.25	-0.0639	-0.0230	0.0385	0.0828
display									
without FEs	0.1337	0.3634	55.21	39.24	5.56	-0.1741	-0.0162	0.3778	0.9083
with FEs	0.0049	0.0072	83.62	10.10	6.27	-0.0520	-0.0149	0.0258	0.0745

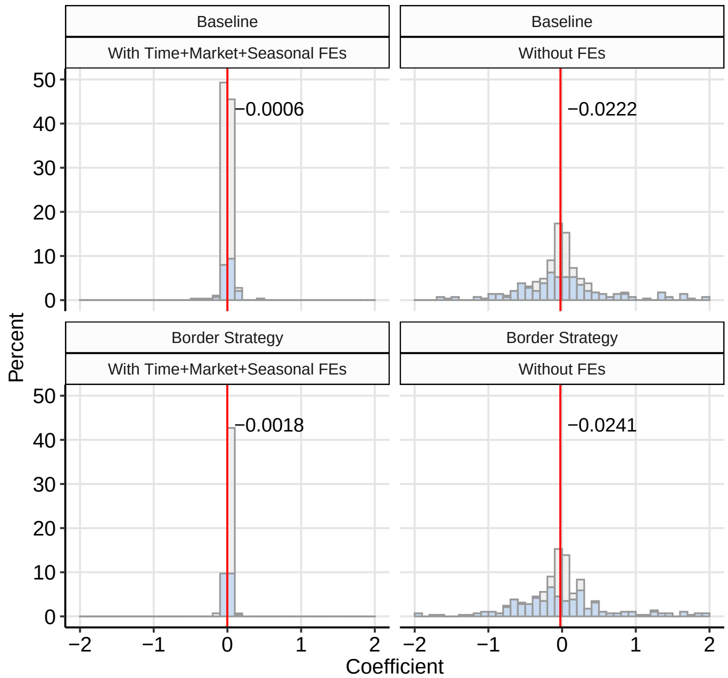

Figure 7: Ad Stock and Price Correlation

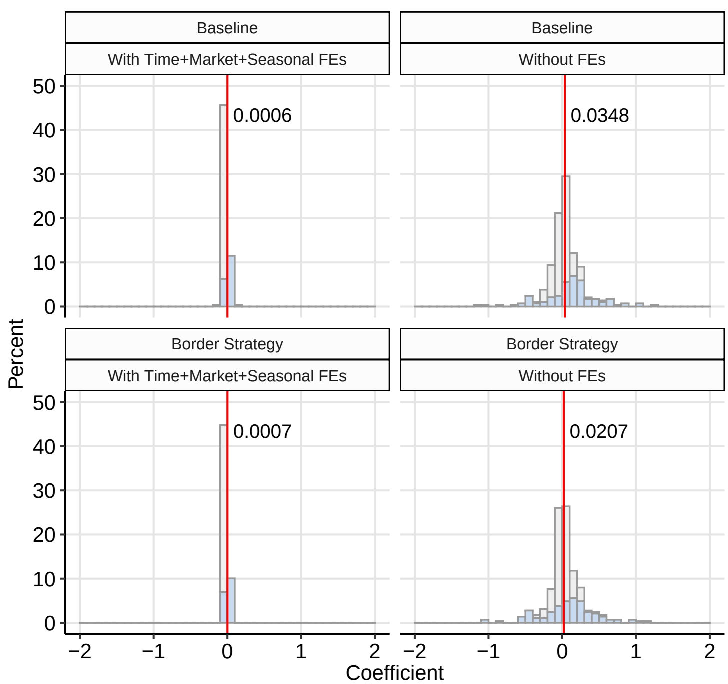

Figure 8: Ad Stock and Promotions Correlation

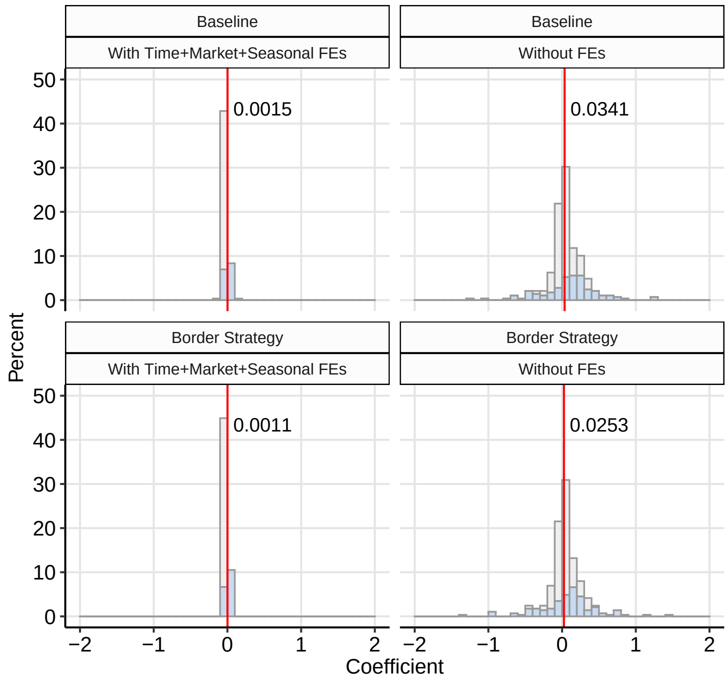

Figure 9: Ad Stock and Feature Correlation

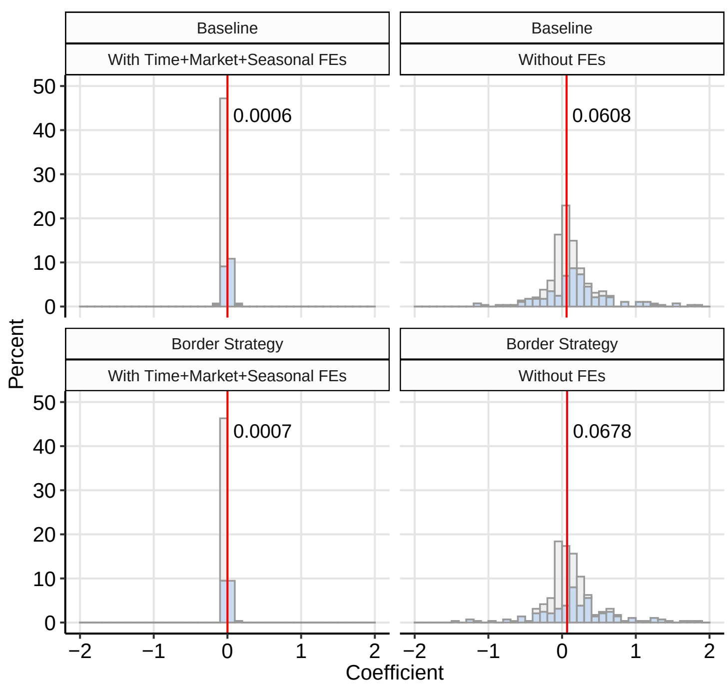

Figure 10: Ad Stock and Display Correlation

Table 6: Correlations Between Ad Stock and Price, Promotions Feature and Display

	Median	Mean	% p ≥ 0.05	% p < 0.05		Percentiles			
				> 0	≤ 0	10%	25%	75%	90%
Baseline									
price									
without FEs	-0.0222	-0.0065	34.38	30.56	35.07	-0.6533	-0.2311	0.1841	0.7899
with FEs	-0.0006	-0.0015	79.51	11.81	8.68	-0.0449	-0.0179	0.0154	0.0513
promo									
without FEs	0.0348	0.0663	62.15	27.08	10.76	-0.1876	-0.0557	0.1399	0.2904
with FEs	0.0006	0.0006	81.88	11.50	6.62	-0.0122	-0.0040	0.0074	0.0150
feature									
without FEs	0.0341	0.0711	63.54	25.00	11.46	-0.1603	-0.0308	0.1509	0.3078
with FEs	0.0015	0.0029	84.67	8.36	6.97	-0.0128	-0.0043	0.0079	0.0217
display									
without FEs	0.0608	0.1430	46.18	39.58	14.24	-0.2206	-0.0281	0.2377	0.5344
with FEs	0.0006	0.0007	79.37	11.19	9.44	-0.0179	-0.0039	0.0084	0.0183
Border Strategy									
price									
without FEs	-0.0241	-0.0169	31.25	30.56	38.19	-0.7137	-0.2881	0.2049	0.8220
with FEs	-0.0018	-0.0029	80.21	10.07	9.72	-0.0342	-0.0152	0.0081	0.0260
promo									
without FEs	0.0207	0.0546	62.15	24.31	13.54	-0.1977	-0.0619	0.1332	0.3007
with FEs	0.0007	0.0017	82.99	10.07	6.94	-0.0114	-0.0032	0.0079	0.0151
feature									
without FEs	0.0253	0.0632	64.24	23.26	12.50	-0.1646	-0.0426	0.1419	0.3046
with FEs	0.0011	0.0036	82.81	10.53	6.67	-0.0109	-0.0037	0.0084	0.0218
display									
without FEs	0.0678	0.1416	52.08	34.03	13.89	-0.2280	-0.0457	0.2418	0.6109
with FEs	0.0007	0.0009	80.70	9.82	9.47	-0.0117	-0.0042	0.0054	0.0151

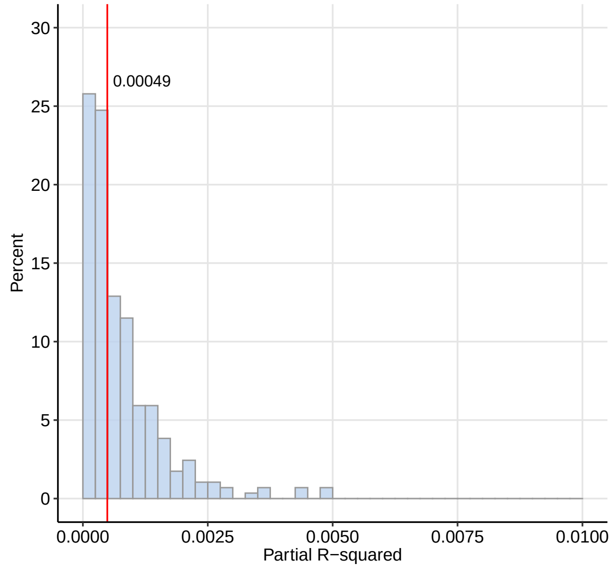

Figure 11: Partial R Squared of Price Variables on Residual Advertising Flow

and display are uncorrelated with advertising. Hence, the choice to omit or include each in the model should not substantively alter the results.

E Flexible functional form for advertising response

In this section, we explore how sensitive our results are to the chosen functional form. In particular, the ROI estimates for small levels of the advertising stock, A, are reliant on the steep slope of the $\log(1+A)$ functional form.

To allow for a flexible functional relationship between each component A_j of the advertising stock vector $\boldsymbol{A}$ and sales we use a linear basis expansion. The basis, $\mathcal{B}$, includes

1. Polynomial (and a square root) transformations of the advertising stock, $A^{\frac{1}{2}}$, A, A^2, ..., A^{10},
2. Transformations of $\log(1+A)$: $(\log(1+A))^{\frac{1}{2}}$, $\log(1+A)$, $(\log(1+A))^2$, ..., $\log(1+A))^{10}$,
3. A cubic B-spline with 9 interior knots, placed at the percentiles $10, 20, \ldots, 90$ of the advertising stock.

We regularize the estimates to prevent over-fitting using a cross-validated Lasso. The Lasso is trained using *residualized* elements of the basis $\mathcal{B}$. In particular, we regress each column $X_k \in \mathcal{B}$ on all fixed effects and covariates (including the competitor advertising stocks) that are included in the original model and then compute the residual, $\tilde{X}_k$, from this regression. Similarly, we obtain a residualized dependent variable, $\widetilde{\log(Q)}$. Using the residualized dependent variable and residualized terms in the linear basis we estimate the cross-validated Lasso. We use this approach

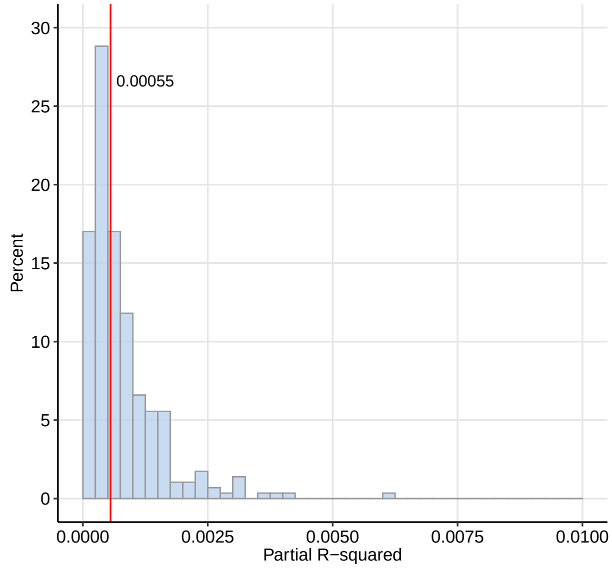

Figure 12: Partial R Squared of Price Variables on Residual Advertising Stock

because we do not want to shrink or eliminate any of the fixed effects or other covariates using the Lasso, because these variables are essential controls to adjust for confounding.

The estimated model allows us to predict $\log(Q)$ and the advertising stock elasticity, $\frac{\partial Q}{\partial A}\frac{A}{Q}$, for any value of the advertising stock. We compute the elasticity at each percentile $25, 26, \dots, 75$ of the observed advertising stock values and use the median of these predicted elasticities as summary statistic for each brand.

As we discussed in Section **??** (Figure **??**), the advertising elasticity estimates from the flexible and parametric model specifications are highly correlated. We provide estimates of the flexible functional form for all brands in the interactive online appendix. As a specific example, in Figure 13 we plot the predicted advertising stock response function for two brands, both using the flexible, semi-parametric model and the $\log(1 + A)$ functional form. We chose the two brands, Chobani and Gatorade, because for both of them we obtain reasonably precise parametric estimates of the advertising effect, making it plausible that we could get relatively precise estimates of the flexible advertising response curve, too. For both brands there are differences between the two response curves in the tails, but overall the semi-parametric and the $\log(1 + A)$ models correspond fairly well.

The overall similarity between the advertising elasticity estimates from the flexible and parametric model does not indicate that the main results are driven by the specific functional form assumption. That being said, the brand-by-brand flexible functional form estimates are available in the online web-application, and further study of why some brands differ significantly from others in the shape of the response curve may be of interest for future research.

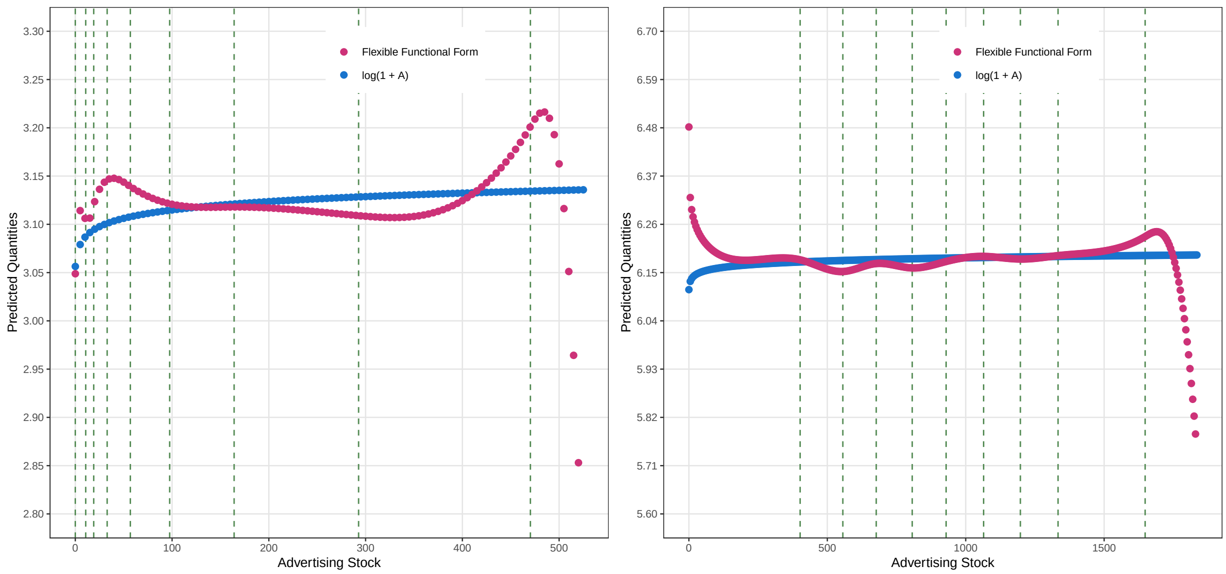

Figure 13: Predicted Quantity using Border-Strategy and Splines Estimation

Note: The left panel is for Chobani, while the right panel is for Gatorade. In both panels, we use the border-strategy model (red dots) with border-month, seasonal, and store fixed effects, and $\delta = 0.9$. For splines estimation (blue dots), we use cubic B-spline with knots placed at the 10th, 25th, 50th, 75th and 90th percentile of observed advertising stock (green dashed lines).

F Advertising cost

In order to calculate a manufacturer's return-on-investment (ROI) from advertising, we need to estimate the cost of buying an ad GRP in DMA d in week t for each manufacturer. We use data on advertising expenditure, impressions, and audience size contained in the Nielsen Ad Intel data set for this purpose.

Expenditure Data

- For Cable, Network, and Syndicated TV, ads are purchased at the national level.
  - For network ads, Nielsen obtains expenditure data from the networks. If expenditure data are unavailable, Nielsen derives estimates of expenditures using supplementary industry data and proprietary models.
  - For cable ads, Nielsen's source for expenditure data is SQAD's NetCosts database. SQAD compiles occurrence-level data on actual purchases reported by contributing ad agencies. The measures SQAD shares with Nielsen are averages at the monthly-network-daypart level. The reported figures are believed to reflect the true weighting of upfront and scatter buys.
  - Expenditure data are originally at the {Month}-{Network}-{Daypart} level for national and cable ads. Ad Intel further prorates expenditure and records the data at the {AdTime}-{Network}-{Daypart}-{Program}-{Duration} level.

- For Spot TV, ads are purchased at the DMA level.
  - Nielsen estimates spot TV expenditures by blending cost-per-point data supplied by SQAD with Nielsen's local market ratings data. SQAD's cost-per-point data are based on actual spot television buys reported by contributing ad agencies.

Impression (Viewership) Data

- For Cable TV, impressions are recorded at the national level.
- For Network and Syndicated TV, while ads are purchased at the national level, the programs are broadcast at local TV stations and viewership is recorded at the DMA level.
- For Spot TV, ads are broadcast at local TV stations and viewership is recorded at the DMA level.

Universe Estimates

- Ad Intel provides the estimated total number of TV households at the national and market levels. These universe estimates are updated yearly.

Build Advertising Cost

For each manufacturer, we do the following:

1. Merge expenditure with impressions for each ad occurrence;
2. Aggregate expenditure and impressions to the {National}-{Year} level. This involves adding up expenditure and impressions across media type, date and markets;
   - We calculate advertising cost at the annual level since expenditure fluctuates across weeks. Hence, advertising cost for all weeks in the same year y remains the same.
3. Calculate national advertising cost per GRP in year y as:

$$\text{adcost per GRP}_{\text{national},y} = \frac{\sum_d \sum_{t\in y} \text{Expenditure}_{dt}}{100 \times \sum_d \sum_{t\in y} \text{Impression}_{dt}/\text{Universe}_{\text{national},y}}$$

4. Calculate DMA-level factor for national advertising cost using:

$$\text{Factor}_{dt} = \frac{\text{Universe}_{dt}}{\text{Universe}_{\text{national},t}}$$

5. Estimate advertising cost per GRP in DMA d in week t:

$$\text{adcost per GRP}_{dt} = \text{adcost per GRP}_{\text{national},y} \times \text{Factor}_{dt}$$

G ROI calculation details and break-even ad effects

G.1 ROI derivation

Consider the impact of changing brand j's advertising by the amount Δa_d in period t. The baseline advertising stock in DMA d in period t is A_{dt}, and the advertising stock resulting from the change in advertising is $A'_{dt} = A_{dt} + \Delta a_d$. Q_{st} denotes the quantity of brand j sold at store s under the baseline advertising stock, A_{dt}. Consistent with our demand specification, Q_{st} can be written as:

$$\log(Q_{st}) = z_{st} + \beta \log(1 + A_{dt}),$$
$$Q_{st} = e^{z_{st}}(1 + A_{dt})^{\beta}.$$

Here, z_{st} contains all other factors besides advertising that affect quantity sales, including prices, competitor advertising, store, season and time intercepts, etc. For any period $\tau \in \{t, ..., t+L\}$, the relative change in sales or sales *lift* that results from the change in advertising in period t is:

$$\lambda_{s\tau} \equiv \frac{Q'_{s\tau}}{Q_{s\tau}} = \frac{(1 + A'_{d\tau})^{\beta}}{(1 + A_{d\tau})^{\beta}} = \left(\frac{1 + A_{d\tau} + \delta^{\tau - t}\Delta a_d}{1 + A_{d\tau}} \right)^{\beta}. \tag{3}$$

Notably, all store, season and time-specific components cancel out, and thus (3) provides the relative increase in overall sales in DMA d that results from the change in advertising. That is, $\lambda_{s\tau} = \lambda_{d\tau}$ for all stores s in DMA d. Hence, the DMA-level change in profits in period τ that results from the increase in advertising in period t is:

$$\Delta\pi_{d\tau} = \sum_{s \in \mathcal{S}_d} (\lambda_{d\tau} - 1) Q_{s\tau} \cdot m \cdot p_{s\tau}, \tag{4}$$

where $\mathcal{S}_d$ includes all stores in DMA d, $Q_{s\tau}$ is the baseline sales quantity in store s, $p_{s\tau}$ is the retail price in the store, and m represents the manufacturer's dollar margin as a percentage of the retail price.[2] Summing across all DMAs and all periods $\tau \in \{t, ..., t+L\}$ yields the total increase in profits that results from the advertising increase Δa_d in period t:

$$\Delta\pi = \sum_{\tau=t}^{t+L} \sum_{d=1}^{D} \Delta\pi_{d\tau}.$$

We denote the cost of buying Δa_d GRPs in DMA d by c_{dt}, such that the total cost of the additional advertising is:

$$C = \sum_{d=1}^{D} c_{dt} \Delta a_d.$$

[2] $m = p^{-1}(w - mc)$, where w is the wholesale price and mc is the marginal cost of production.

Finally, the ROI resulting from the change in advertising is:

$$ROI = \frac{\Delta\pi - C}{C}.$$

G.2 Data sources for ROI calculations

We calculate $\lambda_{d\tau}$, the sales lift that results from changing advertising by Δa_d, using the estimated advertising elasticities from the border strategy with the carryover parameter $\delta = 0.9$.[3] In order to calculate incremental profits using equation (4), we need an estimate of the sales quantities in DMA d in week t (at the observed advertising level, A_{dt}).

The total sales volume from the RMS data under-estimates total market-level sales, because the data available to us do not contain information on all retailers in the market. We correct for this problem as follows. Using the Homescan household panel data and the projection factors provided by Nielsen, we predict market-level quantities, Q_{dt}^H (see Section **??**).[4] We then calculate the weekly average of the Homescan quantities in market d, $\bar{Q}_d^H$. Similarly, we calculate the weekly average of the market-level sales quantities observed in the RMS data, $\bar{Q}_d^R$.[5] We use the ratio $\bar{Q}_d^H/\bar{Q}_d^R$ to scale the weekly store-level RMS sales quantities such that the aggregate quantity across stores predicts the total sales volume at the market level:

$$Q_{st} = \frac{\bar{Q}_d^H}{\bar{Q}_d^R} Q_{st}^R.$$

We use this hybrid of the RMS and Homescan data because the RMS data are likely to provide more accurate information on sales quantity differences across weeks than the Homescan data, whereas the average Homescan volume provides more accurate information on total market-level sales quantities.

To estimate the dollar margin that a manufacturer earns from an incremental sales unit, we use the observed retail prices in the RMS data and multiply by a margin-factor m that represents the manufacturer's dollar margin as a percentage of the retail price. Because we do not observe wholesale prices and manufacturing costs, we need to make assumptions on what margins the manufacturers earn. We consider a range of likely values for the manufacturer margin, $m = 0.2, 0.3, 0.4$. This range corresponds to a range of manufacturer gross margins between 25% and 55% and retail gross margins between 20% and 30%.[6] In the results section

[3]We also calculated the ROIs using different model specifications and carryover parameters. As the estimates of the advertising elasticities are quite robust to the different assumptions, we choose to focus on a single specification here.

[4]For products where on-site purchase and consumption are commonplace (for example at a fast food restaurant or at a sporting event), the Homescan data will understate total quantity. Beer and soft drinks are particularly likely to be affected by this, while other products are much less so. Separating out the 24 beer and soft drink brands does not significantly alter the distribution of ROI that is ultimately estimated. Additionally, assuming that all beer and soft drink brands have twice as many sales as we predict also does not significantly alter the distribution of ROI that is ultimately estimated.

[5]The weekly averages are re-calculated for each year in the data.

[6]To see this, note that m can be expressed as the product of the manufacturer margin and 1 minus the retail

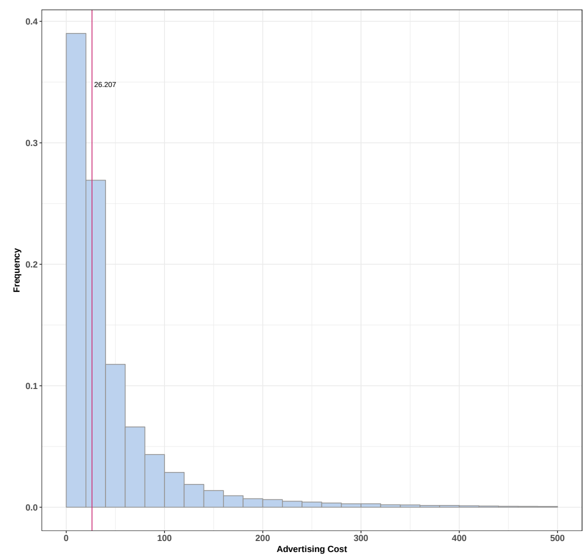

Figure 14: Distribution of Cost per GRP across DMA-Brand-Years

below, we consider how the distribution of estimated ROIs changes under different assumptions about margins.

Finally, we need data on c_{dt}, the cost of buying an incremental advertising GRP in DMA d in week t. The exact marginal advertising cost is not observed by us. Hence, we use data on advertising expenditures in the Nielsen Ad Intel data set and proxy for c_{dt} using the average cost of a GRP in each DMA-year. We calculate the advertising cost separately for each brand and thus capture differences in the campaign costs across brands.[7] We assess the sensitivity of the ROI predictions to this specific advertising cost calculation to ensure that measurement error in the advertising costs does not substantially change the conclusions.

Figure 14 summarizes the distribution of advertising costs. Each observation in the histogram is the average cost of a GRP calculated for a brand, DMA- and year combination. The median cost of buying one additional GRP in a DMA is \$26.21, although there is significant variation in the cost of advertising across brands, media markets, and years.

G.3 Break-even ad effects

In this section, we analyze how much larger TV ad effects would need to be in order for the observed level of advertising to be profitable. To this end, for different assumed values of margin factors and advertising costs, we compute the "break-even" ad elasticity for each brand. That

margin

$$m = \left(\frac{w - mc}{w}\right)\left(1 - \frac{p - w}{p}\right) = \frac{w - mc}{p}.$$

The range of manufacturer gross margins that we consider aligns with industry reports of median manufacturer gross margins of 34% for food companies, 44% for beverage companies, and 50% for companies selling household goods and personal care products (Grocery Manufacturers Association and PricewaterhouseCoopers (2006)).

[7] Appendix F provides more detail about the advertising expenditure data.

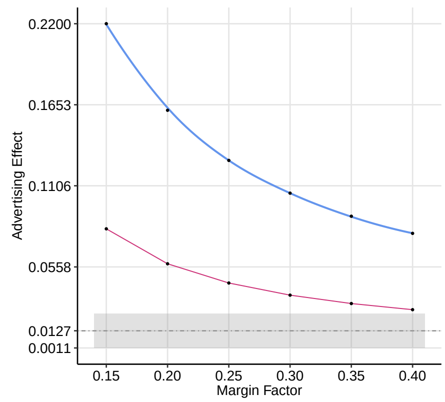

Figure 15: Break-even Advertising Effect (Chobani)

Note: The blue line is the break-even ad effect for the average weekly ROI, while the red line is for the overall ROI. For Chobani, our estimated advertising effect is about 0.0127 (gray dashed line) and the shaded area marks the 95% confidence interval.

is, we solve for the elasticity at which the observed level of weekly advertising would yield an ROI of 0. We calculate the break-even ad effect separately for the average weekly ROI and the overall ROI. Using Chobani as an example, Figure 15 shows how the break-even ad effect varies as a function of the assumed margin factor m and the chosen ROI metric.

For each brand, we compare the break-even ad elasticity to the estimated ad elasticity. To summarize the results across brands, we calculate the ratio of the break-even ad effect to the estimated ad effect.[8] Figure 16 shows the distribution of this multiplier across brands for both the weekly break-even ROI and the overall break-even ROI. The left panel shows that for the median brand in our data, the estimated ad effect would need to be 6.07 times larger in order for the observed level of weekly advertising to be profitable (assuming a margin factor of $m = 0.3$). In contrast, the right panel of Figure 16 shows the results when considering the ROI of all observed advertising.

References

Grocery Manufacturers Association and PricewaterhouseCoopers (2006): "Insights into the Food, Beverage, and Consumer Products Industry: GMA Overview of Industry Economic Impact, Financial Performance, and Trends," .

Hitsch, G. J., A. Hortaçsu, and X. Lin (2019): "Prices and Promotions in U.S. Retail Markets: Evidence from Big Data," *manuscript*.

Lewis, R. and D. Nguyen (2015): "Display advertising's competitive spillovers to consumer search," *Quantitative Marketing and Economics*, 13, 93–115.

[8]Note that we compute this multiplier for the subset of brands with a positive ad elasticity estimate.

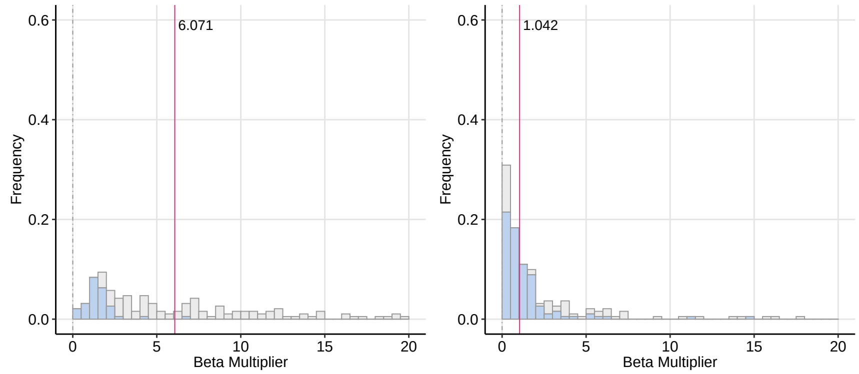

Figure 16: Ratio of Break-even Ad Effect to the Estimated Ad Effect

Note: The left panel shows the distribution of the ratio of the break-even ad effect to the estimated ad effect (multiplier) for weekly break-even ROIs. The right panel shows the multiplier for overall break-even ROIs. The histograms only include the 191 brands with a positive estimated ad effect.

SAHNI, N. S. (2016): "Advertising Spillovers: Evidence from Online Field Experiments and Implications for Returns on Advertising," *Journal of Marketing Research*, 53, 459–478.

SHAPIRO, B. T. (2018): "Positive Spillovers and Free Riding in Advertising of Prescription Pharmaceuticals: The Case of Antidepressants," *Journal of Political Economy*, 126, 381–437.

SINKINSON, M. AND A. STARC (2019): "Ask Your Doctor? Direct-to-Consumer Advertising of Pharmaceuticals," *Review of Economic Studies*, 86, 836–881.