TV Advertising Effectiveness and Profitability: Generalizable Results from 288 Brands*

Bradley T. Shapiro
University of Chicago – Booth

Günter J. Hitsch
University of Chicago – Booth

Anna E. Tuchman
Northwestern University – Kellogg

January 29, 2021

Abstract

We estimate the distribution of television advertising elasticities and the distribution of the advertising return on investment (ROI) for a large number of products in many categories. Our results reveal substantially smaller advertising elasticities compared to the results documented in the literature, as well as a sizable percentage of statistically insignificant or negative estimates. The results are robust to functional form assumptions and are not driven by insufficient statistical power or measurement error. The ROI analysis shows negative ROIs at the margin for more than 80% of brands, implying over-investment in advertising by most firms. While the overall ROI of the observed advertising schedule is only positive for one third of all brands, statistical uncertainty provides the possibility that advertising may be valuable for a larger number of brands if advertising is reduced at the margin.

*All three authors contributed equally although not listed in alphabetical order. We acknowledge the superb research assistance of Albert Kuo, Jihong Song, Ningyin Xu, and Joe Kook. 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 participants at several seminars and conferences. 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.

1 Introduction

We estimate the effect of television advertising on sales and profitability, using data across 288 consumer packaged goods (CPG) in different categories. Our goal is to provide economists and industry practitioners with a general understanding of the effectiveness and economic value of TV advertising. Knowledge on the effect of advertising is important to the economic analysis of advertising, including work on the impact of advertising on market structure, competition, and concentration. A particularly relevant example is the long-run viability of the current media market model. Advertising is a large industry, with total U.S. spending of $256 billion and TV advertising spending of $66 billion in 2019.[1] In traditional broadcast markets, content and advertising are bundled, and advertising acts as an implicit price that consumers pay to subsidize the cost of producing content.[2] The survival of this business model depends on the effectiveness of advertising and firms' willingness to purchase advertising.

Our first set of results pertains to advertising elasticities. We estimate advertising stock elasticities, which are a form of long-run elasticity that represents the total current and future change in sales volume resulting from a one-percent increase in current advertising. In general, advertising is not randomly assigned, and thus, in the presence of unmeasured confounders, the estimated advertising effects do not have a causal interpretation. In order to obtain causal estimates, we employ an identification strategy that relies on the specific institutions of the ad buying process.

We find that the mean and median of the distribution of estimated long-run own-advertising elasticities are 0.023 and 0.014, respectively, and two thirds of the elasticity estimates are not statistically different from zero. These magnitudes are considerably smaller than the results in the extant literature. The results are robust to controls for own and competitor prices and feature and display advertising, and the advertising effect distributions are similar whether a carryover parameter is assumed or estimated. The estimates are also robust if we allow for a flexible functional form for the advertising effect, and they do not appear to be driven by measurement error. As it would be impractical to include all sensitivity checks in the paper, we created an interactive web application that allows the reader to explore all model specifications. The web application is available at https://advertising-effects.chicagobooth.edu.

Our second set of results relates to the profitability of advertising. Using the elasticity estimates and data on the cost of advertising, we compute the implied return on invest-

[1] Winterberry Group: *The Outlook for Data Driven Advertising & Marketing 2020*

[2] The estimates in Brynjolfsson et al. 2019 imply that consumers would be willing to pay $135 billion per year to avoid losing TV access.

ment of advertising. The results show that the ROI of advertising in a given week, holding advertising in all other weeks constant, is negative for more than 80% of the brands in our sample. The implication is that many firms make systematic mistakes and over-invest in advertising at the margin. Further, we predict that the ROI of the observed advertising schedule, compared to not advertising at all, is positive only for one third of all brands. However, the evidence leaves open the possibility that advertising may be valuable for a larger number of brands, in particular if they reduce advertising on the margin.

Our results imply a misallocation of resources by firms. There are multiple potential explanations for these systematic mistakes. Agency problems may be present, such that managers expect a negative impact on their careers if advertising is shown to be unprofitable, or because optimizing advertising strategies requires costly private efforts from the managers. Alternatively, managers may have incorrect priors on the effectiveness of advertising. Such incorrect priors could originate from analyses that insufficiently adjust for confounding factors, or from published meta-analyses that suffer from selection issues, for example due to publication bias affecting the original estimates. Our discussions with managers suggest that all these explanations may be relevant contributing factors to the documented sub-optimal advertising levels.

Ultimately, we hope that our results will motivate managers to critically assess the status quo and encourage firms and researchers to invest in new data and measurement techniques that can improve the effectiveness of TV advertising. Further, together with research documenting ineffective advertising in digital advertising markets (Blake et al. 2015, Lewis and Rao), our work should motivate economists to further study the managerial and agency issues in advertising markets.

The rest of the paper is organized as follows. In Section 2 we survey the literature on the economics of advertising and the measurement of advertising effectiveness. In Section 3 we describe our research design. We describe the data in Section 4 and show summary statistics and describe the identifying variation in Section 5. In Sections 6 and 7 we provide our elasticity and ROI results, respectively. Finally, in Section 8, we conclude.

2 Literature review

Our work adds to an important literature on the economics of advertising. One strand of that literature has investigated the effect of advertising on market structure, competition, and concentration. Sutton's (1991) endogenous sunk cost theory of market structure and concentration, applied to the case when advertising creates vertical product differentiation, assumes that advertising affects consumer demand. The degree to which advertising is effective and has a long-run impact on demand or "brand equity" determines if entry deterrence is possible (Borkovsky et al. 2017, Ellison and Ellison 2011, Bar and Haviv 2019). Second, a long line of research in economics investigates if advertising is primarily persuasive, informative, or effectively a complement to product consumption (see the survey by Bagwell 2007). This study of the mechanism by which advertising affects demand would be moot if advertising were ineffective.

Within the empirical literature, our work is related to a set of papers that perform meta-analyses of published advertising elasticities with the objective of drawing generalizable conclusions about advertising effectiveness (Assmus et al. (1984), Sethuraman et al. (2011), Henningsen et al. (2011)). These studies report meta-analytic means on long-run elasticities in excess of 0.15. This type of work has two main limitations. It relies on published estimates of advertising effectiveness, and differences in the analytic approach may create spurious differences across studies of ad effectiveness. We overcome these limitations by using a single source of data and the same model for all brands in our sample.

Most closely related to our study is the seminal work by Lodish et al. (1995), which summarizes advertising elasticity estimates for 141 brands. The estimates are based on split-cable experiments conducted between 1982 and 1988 in which advertising treatments were randomized across households. Lodish et al. (1995) documents an average advertising elasticity of 0.05 for established products. These results provide a relevant comparison to our work because (i) the Lodish et al. (1995) results were almost certainly not selected based on size, sign or statistical significance, (ii) robustness is ensured given the split-cable RCT design, and (iii) the population of consumer packaged goods is likely similar to our population.

We aim to build on the work of Lodish et al. (1995) in several ways. First, the IRI BehaviorScan test markets in which their experiments were conducted are no longer in use and cannot be used for advertising measurement today. In contrast, our work evaluates television advertising effects using currently available data and methods that are widely employed in the industry and by researchers. Second, although not reported by Lodish et al. (1995), the power of the tests was likely low. (Abraham and Lodish (1990) reports a total of about 3,000 households in the BehaviorScan markets, and thus 3,000 is the maximum sample size in each test.) Compared to Lodish et al. (1995), our study covers a longer time series and many more markets, through which we obtain better statistical power and greater external validity.

Our work is also complementary to some cross-category studies that relate television advertising to various upper-funnel outcomes (Clark et al. (2009), Du et al. (2018), Deng and Mela (2018)) and to some recent multi-product studies of online advertising (Goldfarb and Tucker (2011), Johnson et al. (2016), Kalyanam et al. (2018)).

3 Research design

3.1 Basic model structure

Our goal is to measure the effect of advertising on sales. For each product or brand, we specify a constant elasticity model with advertising carryover. The basic model structure, not including fixed effects and other covariates that we will introduce below, is:

$$\log(Q_{st}) = \boldsymbol{\beta}^T \log(1 + \boldsymbol{A}_{d(s)t}) + \boldsymbol{\alpha}^T \log(\boldsymbol{p}_{st}) + \epsilon_{st}. \tag{1}$$

Q_{st} is the quantity (measured in equivalent units) of the product sold in store s in week t. $\boldsymbol{A}_{d(s)t}$ is a vector of own and competitor advertising stocks in DMA $d(s)$ in week t. $\boldsymbol{p}_{st}$ is a corresponding vector of own and competitor prices. We specify the *advertising stock* as:

$$\boldsymbol{A}_{d(s)t} = \sum_{\tau=t-L}^{t} \delta^{t-\tau} \boldsymbol{a}_{d(s)\tau}. \tag{2}$$

$\boldsymbol{a}_{d(s)t}$, also a vector, is the flow of own and competitor advertising in DMA $d(s)$ in week t, and δ is the advertising carryover parameter. L indicates the number of lags or past periods in which advertising has an impact on current demand. In our empirical specification we set $L = 52$. This stock formulation is frequently used in the literature as a parsimonious way to capture dynamic advertising effects. We assume that $\boldsymbol{A}_{d(s)t}$ captures all dynamics associated with advertising, including the standard carryover effect (current advertising causes future purchases) and structural state dependence (current purchases caused by current advertising cause future purchases). Variation in current advertising that affects future sales will be captured via the distributed lag structure in equation (2), regardless of the specific mechanism.

We measure own advertising using two *separate* variables. The first own advertising variable captures advertising messages that are specific to the focal product. Such advertising is likely to have a non-negative effect on sales. The second own advertising variable captures advertising messages for affiliated products that, ex ante, could have either a positive effect through brand-spillovers or a negative effect through business stealing. For example, an increase in advertising for Coca-Cola soft drinks could increase demand for regular Coca-Cola, but it could also decrease demand for regular Coca-Cola if sufficiently many consumers substitute to Coke Zero or Diet Coke. We will discuss the corresponding data construction more thoroughly in Section 4.

We include advertising and prices for up to three competitors in the model. The competing brands are selected based on total revenue. Some stores do not carry all brands. If a competing brand that is included in the model is not sold at a store, all observations for that store need to be excluded from the analysis. Therefore, for each brand we determine the number of competitors to be included in the model based on the percentage of observations that would be lost if we added one additional competitor.

As the demand function is specified as a log-log model, $\boldsymbol{\alpha}$ includes the own and cross-price elasticities of demand, and the advertising stock elasticity is given by[3]

$$\frac{\partial Q_t}{\partial A_t}\frac{A_t}{Q_t} = \beta\frac{A_t}{1+A_t} \approx \beta. \tag{3}$$

Thus, β is an approximation of the advertising stock elasticity. As shown in Appendix A, the advertising stock elasticity captures the long-run effect of a change in advertising on demand, and in particular measures the total change in current and future demand resulting from a one-percent increase in current advertising.

As our default specification may work better for some applications than others, we provide robustness to the functional form specification, both allowing δ and the shape of the ad response curve to vary by brand.

Demand model choice and functional form

To obtain generalizable and robust advertising effect estimates we need to scale the computations to obtain the estimates across a large number of brands and a large number of different model specifications. The log-linear demand model makes these computations feasible.

Absent computational constraints, we would ideally estimate the relationship between advertising and sales using a micro-founded, structural demand model, such as Berry et al. (1995). Indeed, for the purpose of specific policy evaluations, such as the effect of a merger on equilibrium prices (e.g., Nevo 2000) and advertising levels, or to assess the welfare effect of new product introductions (e.g., Petrin 2002), a structural demand model would be indispensable. The main goal of this paper, however, is to document the distribution of the overall effectiveness of TV advertising across many brands, and we do not conduct policy evaluations that require a prediction of the change in equilibrium advertising. We consider our demand specification to be a log-linear approximation to

[3]For simplicity, we drop the store and market indices and focus on one component of $\boldsymbol{A}_t$.

a micro-founded, structural demand model. To assess the robustness of our results to the specific functional form, we also present flexible semi-parametric estimates that are regularized using the Lasso. Our main results are unchanged by the additional flexibility.

3.2 Identification strategy

The main challenge when estimating equation (1) is that advertising is not randomly assigned. Firms may target their advertising in DMAs and periods when they believe that advertising will be most effective. As a result, firms may advertise more in markets and periods where consumers are positively disposed towards the product even in the absence of advertising. When conducting research across brands, endogeneity may also be induced by the fact that larger brands tend to advertise more. We solve this problem by separately estimating each model brand-by-brand. There may also be unobserved and hence omitted factors that are correlated with both advertising and sales. In the presence of such confounding factors, the statistical relationship between advertising and sales does not have a causal interpretation.

3.2.1 Institutions of the ad buying process

Our identification strategy is based on the institutions of the ad buying process. Television ads are purchased through negotiations between advertisers (or ad agencies) and television stations. As much as 80% of advertising is purchased well in advance of the ad being aired in an *upfront market*.[4] In addition to being purchased in advance, there is considerable bulk buying. That is, an agency will buy a large quantity of advertising to be divided between many clients in exchange for discounts from the stations. The remaining advertising inventory is sold throughout the year. The *scatter market* allows for last minute purchases of individual ads, typically sold at higher rates than upfronts (Hristakeva and Mortimer 2020). Additionally, local networks sometimes sell unsold remnant advertising space to the national networks or other aggregators, which bundle the ads and sell them to advertisers at a discount.

These institutions of the ad buying process make precise targeting difficult. In the upfront market, advertisers may target demand based on differences across local markets, seasonal factors, and trends in demand that can be predicted in advance. Advertisers may also attempt to target demand based on more concrete information about local demand factors that becomes available over time. However, as the majority of inventory is sold upfront, ad buys in the scatter and bundled remnant markets close to a target date are

[4] https://digiday.com/marketing/upfrontses-wtf-upfronts/

constrained by slot availability. Hence, if ad slots in a given week are unavailable in some local markets, the advertiser may buy air time in a previous or subsequent week or not buy additional ad slots in these local markets at all. Even if ad slots are available, the cost of advertising may differ across local markets. In particular, in some local markets advertising inventory may be available in the relatively cheap bundled remnant market, whereas in other markets ad slots may only be available in the relatively expensive scatter market. Because of these cost differences, ad buys may occur in the cheaper markets but not in the more expensive markets. Further, when purchasing in the bundled remnant market, an advertiser may incidentally purchase an ad slot that was of little interest due to the fact that it was bundled with a more desired ad slot.

These institutions of the ad buying process suggest two plausible sources of quasi-random variation in advertising: (i) variation in the exact timing of advertising in a local market and (ii) variation in the amount of advertising across local markets within a time period.

3.2.2 Baseline specification

Our baseline specification includes various fixed effects and controls to isolate this quasi-random variation:

$$\log(Q_{st}) = \boldsymbol{\beta}^T \log(1 + \boldsymbol{A}_{d(s)t}) + \boldsymbol{\alpha}^T \log(\boldsymbol{p}_{st}) + \gamma_s + \gamma_{\mathcal{S}(t)} + \gamma_{\mathcal{T}(t)} + \boldsymbol{\eta}^T \boldsymbol{x}_{st} + \epsilon_{st}. \tag{4}$$

γ_s is a store fixed effect that subsumes persistent regional differences in demand, $\gamma_{\mathcal{S}(t)}$ is a week-of-year fixed effect that captures seasonal effects, and $\gamma_{\mathcal{T}(t)}$ captures aggregate changes or trends in demand. In our preferred specification, $\gamma_{\mathcal{T}(t)}$ is a month fixed effect corresponding to week t, but we also estimate specifications with quarter or week fixed effects and a specification where $\gamma_{\mathcal{T}(t)}$ represents a linear time trend. $\boldsymbol{x}_{st}$ is a vector of other controls, including feature and display advertising, that is included in some of the model specifications. Also, recall that $\boldsymbol{A}_{d(s)t}$ and $\boldsymbol{p}_{st}$ are vectors that include the advertising stocks and prices of competing brands, which may be correlated with the focal brand's advertising activity.

Advertisers are likely to choose different levels of advertising across markets and seasons based on systematic, predictable differences in demand. The baseline specification incorporates store and season fixed effects to adjust for the corresponding confounds. Advertisers may also predict more idiosyncratic fluctuations or trends in demand at the aggregate, national level. We adjust for the resulting confounds using time fixed effects or a trend. Furthermore, advertising campaigns may be coordinated with national or local feature and display advertising, and thus we include data on displays and features in some of the model specifications.

The baseline specification allows for targeting of national demand fluctuations, but assumes that advertisers do not engage in more sophisticated targeting of transient demand shocks at the local market level. Under these assumptions and given the variability induced by the institutions of the ad buying process (see Section 3.2.1), the residual variation in advertising, conditional on all fixed effects and controls, can plausibly be considered quasi-random.[5][6]

While our primary results are based on the baseline strategy, we employ an alternative strategy for additional robustness based on changes in advertising at TV market borders (Shapiro, 2018, Tuchman, 2019, Spenkuch and Toniatti, 2018, Huber and Arceneaux 2007). The border strategy focuses on counties located on DMA borders. The central assumption is that consumers on different sides of a DMA border exhibit parallel trends in their purchasing behavior. The border strategy utilizes the institutional constraint that DMAs are the smallest geographic unit at which ad buys can be made. Hence, advertising that is targeted to local, DMA-specific demand shocks will be based on the overall, population-weighted *average* demand shocks at the DMA level. Consequently, if the average demand shocks across two bordering DMAs differ, then advertising across the DMAs will differ, too. In particular, the level of advertising at the border will generally be different from the optimal, equilibrium level of advertising that firms would set if they could micro-target more locally within a DMA.[7]

[5]The baseline specification results presented in this paper use monthly time fixed effects. We also estimated specifications using week and quarter fixed effects. Utilizing week fixed effects considerably decreases the statistical power for many brands, due to the large share of national advertising, and it also makes the week-of-year dummies redundant.

[6]Other papers have proposed identification strategies that use instrumental variables to estimate a causal effect of advertising on sales (Gordon and Hartmann, 2013, Sinkinson and Starc, 2019, Thomas 2020, Li et al. 2019). We chose not to implement these strategies because instruments are case-specific and, hence, impractical for a study that estimates advertising effects for 288 brands.

[7]The border strategy also requires the stable unit treatment value assumption (SUTVA) that consumers who are exposed to advertising do not cross the border and purchase in the neighboring DMA. If this assumption is violated, the advertising effect estimate will be biased towards zero. Tuchman (2019) provides empirical evidence consistent with this assumption. She analyzes cigarette and e-cigarette purchases made by Nielsen Homescan panelists that reside in border counties and finds that of the more than 84,000 transactions made by these households, only 3% were made at a store outside the household's DMA of residence. Thus, any bias from violation of SUTVA is likely to be small. This is consistent with our finding that our baseline estimates are statistically indistinguishable from our border strategy estimates.

4 Data

To estimate the effect of advertising on sales we use data on purchase volumes, advertising intensities, and other components of marketing, in particular prices. We construct a data set by merging market (DMA) level TV advertising data with retail sales and price data at the brand level. The data and our matching procedure are described in more detail below. Detailed information on how we construct the data is available in Shapiro et al. (2020).

4.1 RMS retail scanner data

The Nielsen RMS (Retail Measurement Services) data include weekly store-level information on prices and quantities sold at the UPC level. The RMS data include information for about 40,000 stores, including grocery stores, drug stores, mass merchandisers, and convenience stores. Despite covering a large number of stores and retailers, the data available for research from the Kilts Center for Marketing constitute a non-random subset of all retail chains in the U.S. The data cover more than 50% of all market-level spending in grocery and drug stores and one-third of all spending at mass merchandisers.

The sample used in our analysis includes data from 2010 to 2014. We focus our analysis on the top 500 brands in terms of dollar sales. These brands account for 45.3% of the total observed RMS revenue, even though there are more than 300,000 brands in the data. We define a brand as all forms of the same consumable end product, as indicated by the brand code or brand name in the RMS data. That is, Coca-Cola Classic includes any UPC that was composed entirely of Coca-Cola Classic, including twelve ounce cans, two-liter bottles, or otherwise. Because advertising is generally at the brand level, rather than the UPC level, we aggregate across UPCs, calculating total volume sold in equivalent units and average price per equivalent unit. We have 12,671 stores in the final estimation sample.

The price of a UPC is only recorded in weeks when at least one unit of the UPC was sold. To impute these prices that are missing from the data, we follow the approach detailed in Hitsch et al. (2019). This approach uses an algorithm to infer the base price, i.e., regular, non-promoted shelf price of a product, and assumes that weeks with zero sales occur in the absence of a promotion, such that the unobserved price corresponds to the base price.

4.2 Homescan household panel data

The policy experiments and ROI calculations in Section 7 make use of the Nielsen Homescan household panel data as an additional source of purchase information. The Homescan data capture household-level transactions, including purchase quantities and prices paid. Data for more than 60,000 households are available each year. Nielsen provides weights called *projection factors* for each household. Using these weights, transactions can be aggregated across all households to be representative at the national level, i.e., estimate a product's total purchase volume for in-home use. We utilize these estimates of total sales for the policy experiments and ROI calculations because the RMS data do not capture all transactions and would hence underestimate the incremental value of advertising.

4.3 Advertising data

Product-level television advertising data for 2010–2014 come from the Nielsen Ad Intel database. The advertising information is recorded at the occurrence level, where an occurrence is the placement of an ad for a specific brand on a given channel, in a specific market, at a given day and time. Four different TV media types are covered in the data: Cable, Network, Syndicated, and Spot. Occurrences for each of these different media types can be matched with viewership data, which then yields an estimate of the number of impressions, or eyeballs, that viewed each ad. In the top 25 DMAs, impressions are measured by set-top box recording devices. For all other DMAs, impressions are measured using diaries filled out by Nielsen households. These diary data are only recorded in the four "sweeps months," February, May, July, and November. We impute the impressions for all other months using a weighted average of the recorded impressions in the two closest sweeps months.

For Cable ads, which are aired nationally, viewership data are available only at the national level. Spot ads are bought locally, and viewership measures are recorded locally, separately for each DMA. Network and Syndicated ads are recorded in national occurrence files that can be matched with local measures of viewership in each DMA. Thus, in our data, variation in a brand's aggregate ad viewership across markets is due to both variation in occurrences across markets (more Spot ads were aired in market A than in market B) and variation in impressions (eyeballs) across markets (a Network or Syndicated ad aired in both markets A and B, but more people saw the ad in market A than in market B).

Using the occurrence and impressions data, we calculate gross rating points (GRPs), a widely used measure of advertising exposure or intensity in the industry. We first calculate the GRP for a specific ad occurrence, defined as the number of impressions for the ad as a percentage of all TV-viewing households in a DMA (measured on a scale from 0 to 100). To obtain the aggregate, weekly GRPs in a given DMA, we obtain the sum of all occurrence-level GRPs for a brand in a given week in the DMA.

4.4 Matching advertising and retail sales data

We merge the advertising and sales data sets at the store-brand-week level. Our merging procedure warrants discussion because the brand variables in the Ad Intel and RMS data sets are not always specified at the same level. We include three types of advertising variables in our models. First, we include advertising that directly corresponds to the RMS product in question. Second, we create a variable that captures advertising for affiliated brands, including potential substitutes, that may affect the focal RMS product. Third, we include advertising for the top competitor. For example, for the Diet Coke brand, own advertising includes ads for Diet Coke, whereas affiliated advertising includes advertising for Coca-Cola soft drinks, Coke Zero, Coca-Cola Classic, and Cherry Coke. Furthermore, we include advertising for Diet Pepsi, the top competitor of Diet Coke.

We separately estimate the effect of own brand and affiliated brand advertising because own brand advertising is likely to have a positive effect on sales, whereas the sign of the effect of affiliated brands' advertising is ambiguous. Hence, lumping own and affiliated brand advertising together might result in small and uninterpretable elasticity estimates.

Full details of the matching approach are provided in Shapiro et al. (2020), and the estimated affiliated brand and competitor ad effects are reported in Appendix B.

5 Data description

5.1 Brand-level summary statistics

Using the process described in Section 4.4, we match 288 of the top 500 brands in the RMS data to TV advertising records in the Ad Intel database. These products are typically established products, and hence the results from our empirical analysis need not apply to new products.

In Table 1, we provide brand-level summary statistics. Total yearly revenue is larger when based on the spending records in the Homescan data compared to the measured revenue in the RMS retail sales data because the reported RMS revenue is calculated using the subset of stores used in our estimation sample. The Homescan revenue, on the other hand, is predicted using the transaction records and household projection factors

Table 1: Brand Level Summary Statistics

	Median	Mean	Percentiles							
			1%	5%	10%	25%	75%	90%	95%	99%
RMS revenue	113.1	170.8	28.8	51.2	61	75.4	190	332.4	522.4	698.9
Homescan revenue	341	546.3	74.3	119.7	150.8	220	647.7	1046.3	1544.6	3190.4
Advertising spending	10.5	18.6	0.9	2.2	3.6	5.9	22.2	44.7	61.3	106.7
Mean weekly GRPs	35.5	59.4	2.2	4.7	8.4	19	71.8	149.7	184.8	333.8
Adv./sales ratio (%)	2.8	5.3	0.2	0.5	0.8	1.4	5.6	12.4	17.8	33.6
% of Adv. Spending										
Cable	50.9	52.8	10	20.9	29.9	38.4	65.6	81	88.2	99.8
Network	34.5	34.1	0.9	4.2	8.6	19.6	47.5	56.6	66.6	82.4
Spot	3.7	8.7	0.1	0.4	0.7	1.5	8.6	20.9	37.8	82.5
Syndicated	5.3	6.6	0	0	0	1.6	9.8	16.6	19.1	22.9

Note: The sample includes 288 brands. Revenue and advertising spending are expressed in millions of dollars. The advertising/sales ratio is calculated using Homescan revenue.

supplied in the Nielsen data, and is thus designed to be representative of total national spending.

The data reveal a large degree of heterogeneity in both advertising and advertising/sales ratios. In Table 1, we document that total yearly TV advertising spending for the median brand is 10.5 million dollars, with a 90% range of 2.2 to 61.3 million dollars. A similar degree of cross-brand heterogeneity is evident in the advertising/sales ratios, with a median of 2.8 and a 90% range from 0.5 to 17.8. The median of average weekly GRPs across brands is 35.5, with a 90% range from 4.7 to 184.8.

5.2 Temporal and cross-sectional variation at the brand level

The degree of temporal and cross-sectional variation in brand-level advertising is of particular relevance for the goal of estimating advertising effects on demand. We document the extent of this variation in the data. First, separately for each brand, we regress weekly DMA-level advertising, measured in GRPs, on a set of DMA, week-of-year (season), and month fixed effects. Additional covariates included in this regression are own and competitor prices, and competitor advertising. We then calculate the ratio of the residual standard deviation from this regression relative to average DMA/week advertising. This

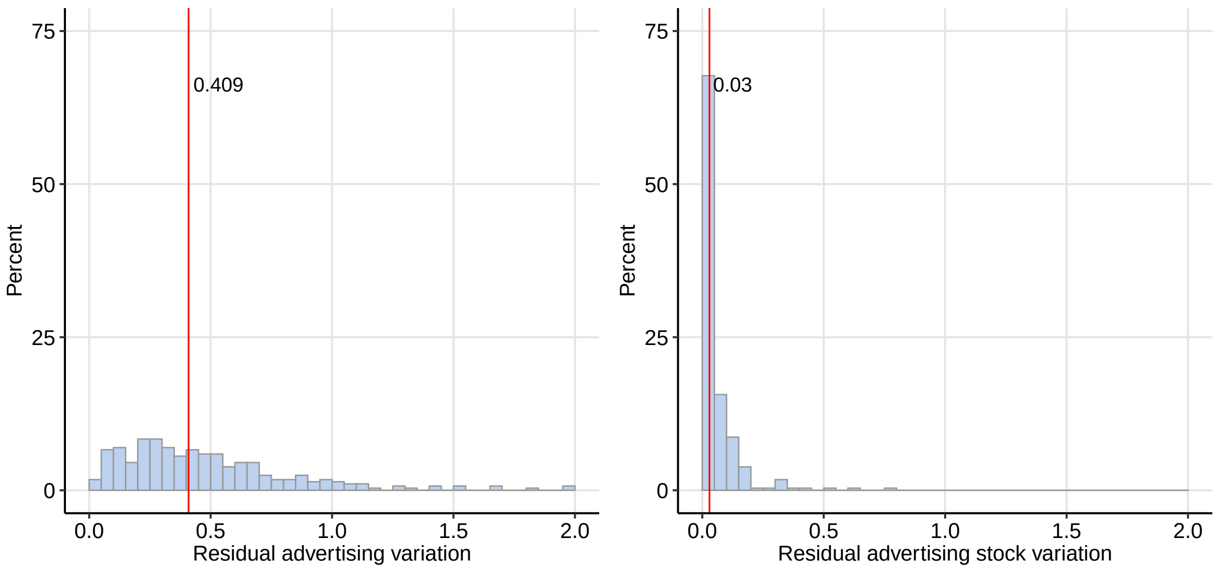

Figure 1: Residual Variation in Advertising

Note: The residual variation measures are based on the residuals from a regression of advertising or advertising stock ($\delta = 0.9$) on DMA, time (month), and seasonal (week-of-year) fixed effects, own and competitor prices, and competitor advertising. The residual variation is the ratio of the standard deviation of these residuals relative to the mean advertising or advertising stock. The measure is calculated separately for each brand, and these graphs show the distribution across brands.

measure is similar to a coefficient of variation and serves as a parsimonious way of quantifying the degree of variation in advertising that is not explained by the fixed effects and the other covariates.

In Figure 1, we present a histogram of the measure across brands and also show a similar measure of the residual variation in *advertising stock* relative to the average DMA/week advertising stock. The advertising stock is calculated assuming a carryover parameter of $\delta = 0.9$. The "coefficient of variation" of advertising flows is 0.41 for the median brand. I.e., for the median brand, the standard deviation of the residuals is substantial, at 41% of average weekly advertising. For advertising stocks the relative residual variation is substantially smaller, 0.03 for the median brand.

6 Results

We first present the results of the baseline specification discussed in Section 3.2 and then analyze the robustness of these results. The baseline specification includes store, week-of-year (season) fixed effects, and common time fixed effects. The estimation results are initially obtained assuming a carryover parameter $\delta = 0.9$, which is similar to other specifications in the literature (e.g., Dubé et al. (2005)).

6.1 Main results

We present the estimation results for the own-advertising stock elasticities, i.e., the coefficients for the focal brand in the vector $\boldsymbol{\beta}$.[8] As discussed in Section 3.1, the advertising stock elasticities can be interpreted as long-run advertising elasticities. For the sake of brevity, from now on we refer to the own-advertising stock elasticities as advertising elasticities or advertising effects. We provide summary statistics for the model estimates in the *Main Results* panel of Table 2, and we display the full distribution of brand-level advertising elasticities, arranged from smallest to largest elasticity together with 95% confidence intervals, in Figure 2.[9]

First, the advertising elasticity estimates in the baseline specification are small. The median elasticity is 0.0140, and the mean is 0.0233. These averages are substantially smaller than the average elasticities reported in extant meta-analyses of published case studies. Second, two thirds of the estimates are not statistically distinguishable from zero.

[8] We discuss the estimated affiliated brand and cross-advertising elasticities in Appendix B.

[9] In Appendix C, we show that product categories and grocery aisles do not predict which brands have relatively large elasticity estimates.

Table 2: Own-Advertising Stock Elasticity Estimates

	Median	Mean	% p ≥ 0.05	% p < 0.05		Percentiles			
				> 0	≤ 0	10%	25%	75%	90%
Main Results									
Naive	0.0299	0.0415	38.89	41.67	19.44	-0.0713	-0.0178	0.0833	0.1827
+ Store FE	0.0218	0.0467	33.68	50.69	15.62	-0.0349	-0.0051	0.0682	0.1494
+ Season FE	0.0152	0.0251	28.82	51.04	20.14	-0.0454	-0.0107	0.0534	0.1051
+ Time trend	0.0110	0.0171	41.67	42.36	15.97	-0.0360	-0.0053	0.0381	0.0772
Baseline specification	0.0140	0.0233	66.32	26.39	7.29	-0.0406	-0.0082	0.0426	0.0919
Robustness									
Border strategy	0.0136	0.0258	68.40	24.31	7.29	-0.0321	-0.0055	0.0472	0.1015
Semi-parametric	0.0140	0.0261	-	-	-	-0.0755	-0.0195	0.0670	0.1409
Estimated δ	0.0090	0.0116	51.04	35.42	13.54	-0.1102	-0.0149	0.0530	0.1733
Prices excluded	0.0118	0.0221	70.83	21.88	7.29	-0.0607	-0.0168	0.0522	0.1359
+ Own price	0.0130	0.0231	66.67	26.39	6.94	-0.0397	-0.0082	0.0458	0.0983
+ Top 1 competitors	0.0132	0.0234	65.97	26.04	7.99	-0.0420	-0.0084	0.0449	0.0984
+ Up to top 2 competitors	0.0142	0.0232	66.32	25.69	7.99	-0.0406	-0.0084	0.0450	0.0931
Feature & display included	0.0086	0.0221	70.14	23.61	6.25	-0.0307	-0.0068	0.0352	0.0809
50% power to detect 0.05	0.0085	0.0098	64.97	26.11	8.92	-0.0237	-0.0067	0.0266	0.0428

Note: Descriptive statistics of estimated advertising elasticities reported for 288 brands. The naive model includes own and competitor advertising stocks and prices but no additional controls. The baseline model includes store, week-of-year (season), and month fixed effects. The border strategy restricts the sample to stores in border counties and includes store, week-of-year (season), and border-month fixed effects. All robustness results are obtained using the baseline strategy. Standard errors are two-way clustered at the DMA and week level in the naive and baseline specifications, and two-way clustered at the border-side and week level in the border strategy.

We show in Figure 2 that the most precise estimates are those closest to the mean and the least precise estimates are in the extremes.

6.2 Robustness

In the *Main Results* panel in Table 2, we show how the advertising elasticity estimates change as we adjust for different sources of confounding. Correspondingly, we provide a visual comparison of the brand-level estimates between different model specifications in Figure 3. We start with a naive specification that includes no fixed effects, and then incrementally add market (store), season (week-of-year), and time fixed effects. The medians of the advertising effects decrease substantially when we adjust for market and season fixed effects, which is consistent with firms targeting predictable demand differences

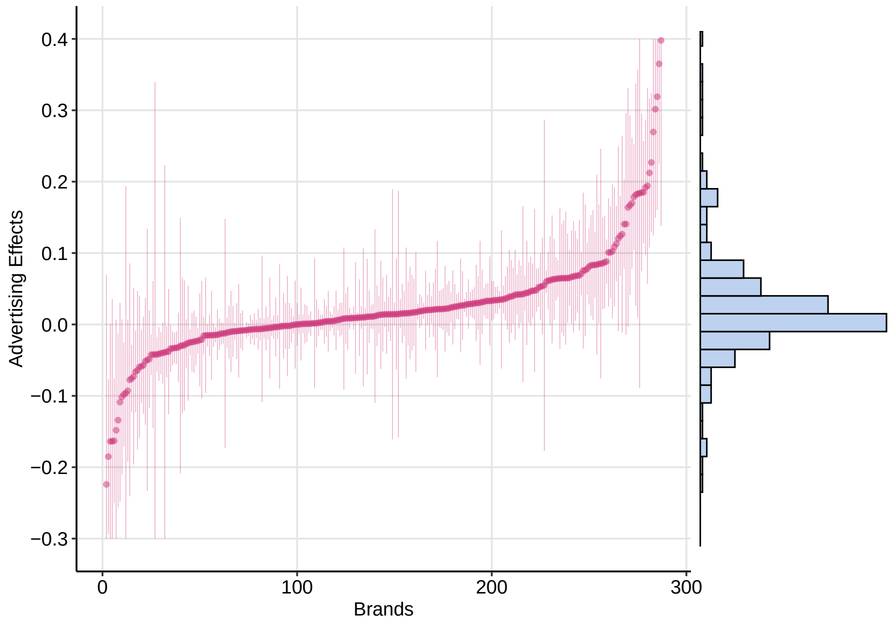

Figure 2: Advertising Effects and Confidence Intervals using Baseline Strategy

Note: Brands are arranged on the horizontal axis in increasing order of their estimated ad effects. For each brand, a dot plots the point estimate of the ad effect and a vertical bar represents the 95% confidence interval. Results are from the baseline strategy model with $\delta = 0.9$ (equation 4).

at the market and season level. Additionally, adjusting for a parametric or flexible time trend has little effect on the estimates, and the estimated advertising elasticities appear to have stabilized once market and season fixed effects are accounted for. Hence, there is little evidence that firms target advertising to more specific temporal demand shocks. These results bolster our confidence in the assumptions underlying the baseline identification strategy.

In addition, we performed an extensive analysis to ensure that the results are robust to different model specifications, with all details available in Appendix D and through the interactive web application, https://advertising-effects.chicagobooth.edu. We provide a concise summary of the main results and report the corresponding summary statistics in the *Robustness* panel of Table 2.

Border strategy Estimating the advertising elasticities using a border strategy (Shapiro 2018), we obtain a distribution of estimates that is nearly identical to the baseline specification. In the bottom left panel of Figure 3, we show that at the brand level, estimates obtained from the border strategy are highly correlated with estimates obtained using the baseline strategy.

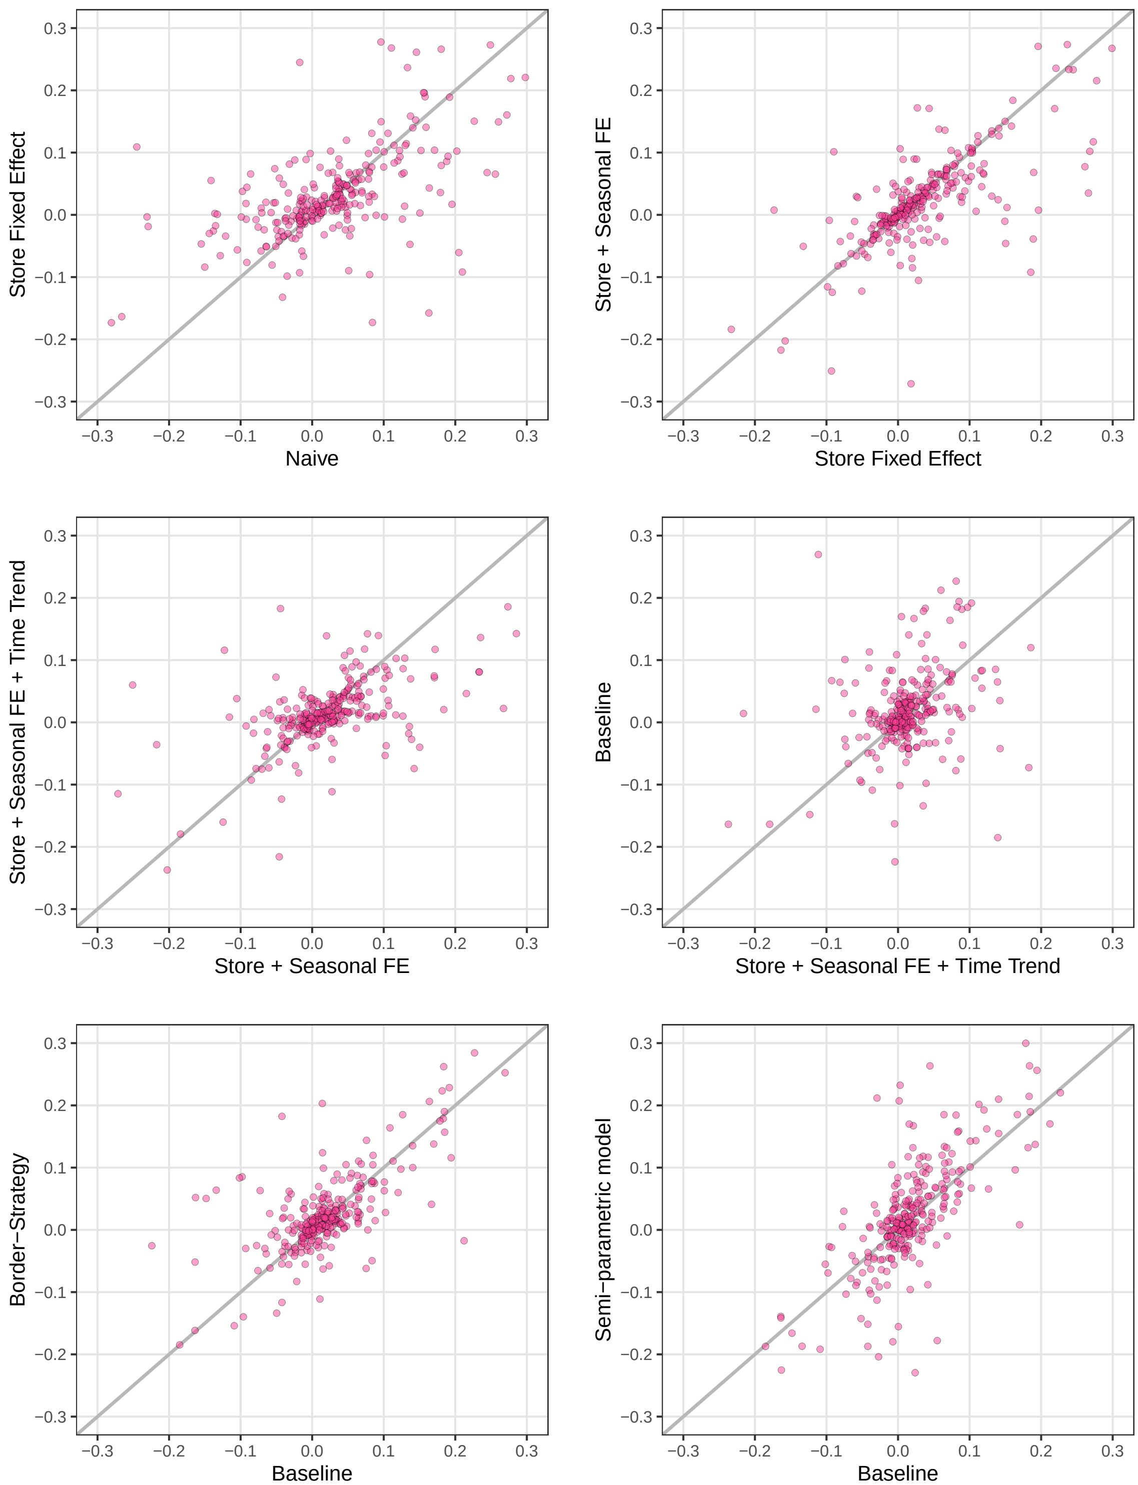

Figure 3: Advertising Stock Elasticities by Specification

Note: Each panel contains the results of two specifications plotted against each other, with the vertical axis having incrementally more adjustment for confounding factors. Each dot represents a brand. The 45 degree line is presented with the plots. Advertising stocks are computed assuming a carryover parameter of $\delta = 0.9$.

Semi-parametric functional form To ensure that our estimates are not biased due to an incorrect choice of the functional form of the model, we allow for a semi-parametric, flexible functional relationship between the advertising stock and sales. We use a linear basis expansion with a basis that includes polynomials of A_j and $\log(1 + A_j)$, and basis B-splines. To prevent over-fitting, we estimate the model using a cross-validated Lasso. For each brand we calculate a summary measure of the advertising elasticity that can be compared to the estimates from the parametric model specification. Full details are provided in Appendix E.

The advertising elasticity estimates from the flexible and parametric model specifications are highly correlated (bottom right panel of Figure 3), and the overall summary statistics of the flexible model estimates are similar to the baseline specification.

Estimation of the carryover parameter The baseline specification results are obtained assuming a carryover parameter, $\delta = 0.9$. When we estimate δ using a grid search we obtain a similar median and mean as in the baseline specification, although the percentage of positive and negative statistically significant estimates is somewhat larger.

Prices and promotions The results remain robust irrespective of the inclusion of own prices or competitor prices. The results also remain unchanged if we include two types of promotions, feature and in-store display advertising, in the model. This rules out confounding if feature and display advertising were coordinated with the TV advertising campaigns. For details on how prices, promotions, feature, and display are correlated with advertising, see Appendix F.

Statistical power and measurement error If we restrict the analysis to brands with 50% power to detect an elasticity of 0.05, the distribution of the advertising elasticity estimates tightens around zero, implying that low statistical power is not the reason for the large share of insignificant estimates. In Appendix D, we detail evidence against the hypothesis that attenuation due to classical measurement error is driving our small estimates.

Further, the data we analyze is among the best available data for advertising measurement, and it is identical to the commercially available data used in the industry. Thus, it is implausible that firms observe different advertising effects because of access to better data sources with less measurement error, a fact that is important for the analysis of the economic value of the observed advertising levels in the next section.

7 Economic implications

We now discuss the implications of the reported advertising elasticities for the economic value of advertising. For each brand we conduct two policy experiments to evaluate the change in profits that results from a change in advertising. We report the impact on profitability as the return on investment (ROI) that results from a modification of the brand manufacturer's advertising policy.[10]

The results reported below are based on the estimated elasticities from the baseline specification with carryover parameter $\delta = 0.9$. To predict total national sales volumes, we scale the RMS sales quantities to the total national level using the Nielsen Homescan data.[11] Because we do not have wholesale price and production cost data, we report the results for manufacturer margins between 20% and 40% (the margins are defined as the difference between the wholesale price and the marginal production cost expressed as a percentage of the retail price). This range of margins is consistent with industry reports. In all ROI calculations we hold constant observed prices, as well as advertising for affiliated and competitor brands.

Standard errors are computed using the delta method. A full description of the data and the approach used to compute the ROIs is presented in Appendix G.

7.1 Average ROI of advertising in a given week

In the first policy experiment we measure the ROI of the observed advertising levels (in all DMAs) in a given week t relative to not advertising in week t. For each brand, we compute the corresponding ROI for all weeks with positive advertising, and then average the ROIs across all weeks to compute the average ROI of weekly advertising. This metric reveals if, *on the margin*, firms choose the (approximately) correct advertising level or could increase profits by either increasing or decreasing advertising.

We show the distribution of the predicted ROIs in Figure 4, and we provide key summary statistics in Table 3. The average ROI of weekly advertising is negative for most brands over the whole range of assumed manufacturer margins. At a 30% margin, the median ROI is -88.15%, and only 12% of brands have positive ROI. Further, for only

[10] We do not attempt to address by how much advertising should be reduced or how the overall advertising schedule should change. Answering these questions requires solving for the dynamically optimal advertising schedule, such as in Dubé et al. (2005), which is beyond the scope of this paper.

[11] 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 quantities. Beer and soft drinks are particularly likely to be affected by this issue. Separating out the 24 beer and soft drink brands does not significantly alter the distribution of ROIs. The results are available by request.

Table 3: Advertising ROI

	Median	**Mean**	**% ROI** > 0	% **p** $\geq$ **0.05**	% **p** $<$ **0.05**	
					ROI > 0	**ROI** < 0
Average ROI of Weekly Advertising						
20% Margin	-92.10	-77.15	7.72	19.65	2.11	78.25
30% Margin	-88.15	-65.72	11.93	29.12	2.81	68.07
40% Margin	-84.20	-54.30	17.19	35.09	3.86	61.05
ROI of All Observed Advertising						
20% Margin	-71.56	-81.24	24.21	48.07	8.07	43.86
30% Margin	-57.34	-71.85	33.68	57.89	11.93	30.18
40% Margin	-43.13	-62.47	40.00	60.00	15.44	24.56

Note: The estimates are obtained using the baseline strategy and assuming a carryover parameter $\delta = 0.9$.

3% of brands the ROI is positive and statistically different from zero, whereas for 68% of brands the ROI is negative and statistically different from zero.

These results provide strong evidence for over-investment in advertising at the margin.[12]

7.2 Overall ROI of the observed advertising schedule

In the second policy experiment we investigate if firms are better off when advertising at the observed levels versus not advertising at all. Hence, we calculate the ROI of the observed advertising schedule relative to a counterfactual baseline with zero advertising in *all* periods.

We present the results in Figure 4 and Table 3. At a 30% margin, the median ROI is -57.34%, and 34% of brands have a positive return from the observed advertising schedule versus not advertising a all. Whereas 12% of brands only have positive and 30% of brands only negative values in their confidence intervals, there is more uncertainty about the sign of the ROI for the remaining 58% of brands. This evidence leaves open the possibility that advertising may be valuable for a substantial number of brands, especially if they reduce advertising on the margin.

[12] In Appendix G.3, we assess how much larger the TV advertising effects would need to be for the observed level of weekly advertising to be profitable. For the median brand with a positive estimated ad elasticity, the advertising effect would have to be 5.37 times larger for the observed level of weekly advertising to yield a positive ROI (assuming a 30% margin).

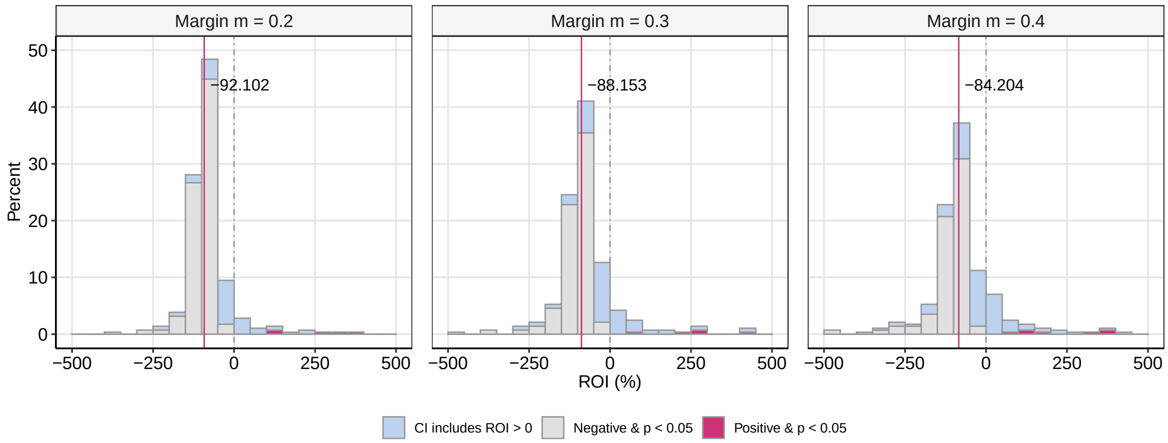

(a) Average ROI of Weekly Advertising

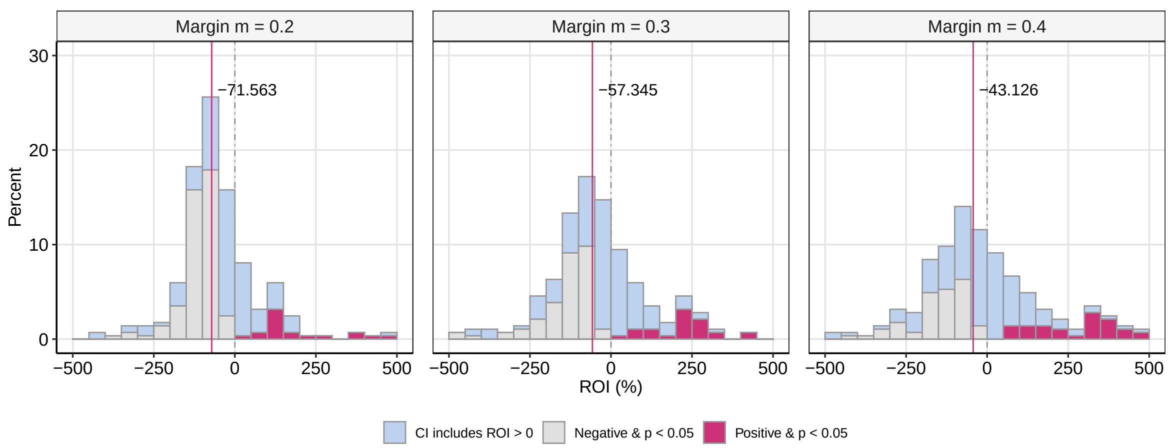

(b) ROI of all Observed Advertising

Figure 4: Predicted ROIs

Note: The top panel provides the distribution of the estimated ROI of weekly advertising and the bottom panel provides the distribution of the overall ROI of the observed advertising schedule. Each is provided for three margin factors, $m = 0.2$, $m = 0.3$, and $m = 0.4$. The median is denoted by a solid vertical line and zero is denoted with a vertical dashed line. Gray indicates brands with negative ROI that is statistically different from zero. Red indicates brands with positive ROI that is statistically different from zero. Blue indicates brands with ROI not statistically different from zero.

8 Conclusions

In this paper, we provide a generalizable distribution of television advertising elasticities for established products that can serve as a prior distribution for firms and researchers. Providing generalizable estimates of TV advertising effects necessitates transparent and replicable estimation methods and an a priori relevant population of products. Our analysis is based on a sample of 288 large, national CPG brands that are selected using a clear research protocol, and our data sources (Nielsen Ad Intel and RMS scanner data) are widely used by marketing managers and academic researchers. We find that the median of the distribution of estimated long-run advertising elasticities is between 0.0085 and 0.0142, and the corresponding mean is between 0.0098 and 0.0261.

We draw two main lessons from these results. First, the estimated advertising elasticities are small, and two thirds of the estimates are not statistically distinguishable from zero. The estimates are also economically small, in the sense that more than 80% of all brands have a negative ROI of advertising at the margin. The estimates are roughly half the size of the most comparable prior study, Lodish et al. (1995), which used data from the 1980s. This difference is consistent with an overall decline in TV advertising effectiveness over the last three decades.

Second, our results are robust. In particular, across a wide range of specifications, the overall distribution of advertising elasticities is stable. Due to this fact, together with the institutional details of the ad buying process that underlie our identification strategy, it appears implausible that our results are affected by any remaining confounds.

Our results have important positive and normative implications. A central finding is the over-investment in advertising for more than 80% of brands, a significant misallocation of resources by firms. Our data are identical to the commercially available data used by firms, and hence it is unlikely that firms observe larger advertising effects because of access to alternative data sources. This raises an economic puzzle. Why do firms spend billions of dollars on TV advertising each year if the return is negative? There are several possible explanations. First, agency issues, in particular career concerns, may lead managers (or consultants) to overstate the effectiveness of advertising if they expect to lose their jobs if their advertising campaigns are revealed to be unprofitable. Second, an incorrect prior (i.e. conventional wisdom that advertising is typically effective) may lead a decision maker to rationally shrink the estimated advertising effect from their data to an incorrect, inflated prior mean. Third, the estimated advertising effects may be inflated if confounding factors are not adequately adjusted for. The last two explanations do not assume irrational behavior, but may simply represent a cost of conducting causal inference to acquire accurate information on the effect of advertising. We view this explanation as plausible given that unified, formal approaches to causal inference have only recently been widely adopted. These proposed explanations are not mutually exclusive. In particular, agency issues may be exacerbated if the general effectiveness of advertising or a specific advertising effect estimate is overstated.[13] While we cannot conclusively point to these explanations as the source of the documented over-investment in advertising, our discussions with managers and industry insiders suggest that these may be contributing factors.

This brings us back to a key motivating question for this research, the long-run viability of traditional media markets. The documented over-investment in advertising suggests a threat to the survival of media markets in their current form, once knowledge about the small degree of TV advertising effectiveness becomes common knowledge. But our results also indicate that for a substantial number of brands (34% based on the point estimates), the observed advertising schedules are valuable compared to the counterfactual of no advertising. There is a large degree of statistical uncertainty about the exact ROIs, and only for 12% of brands the predicted ROIs from the observed advertising schedules are positive and statistically different from zero. This suggests a large option value from adopting improved methods or research designs, such as A/B tests, to estimate the causal effect and ROI of advertising. Our results also do not foreclose the possibility that advertising can be profitable with alternative scheduling, targeting, or advertising copy strategies. The rise of addressable television, in particular, should allow advertisers and researchers to experiment with individual level targeting in the future. These approaches for improving advertising measurement, scheduling, and targeting may well ensure the long-run viability of media markets.

While improvements in targeting technology may theoretically increase the potential for higher advertising returns, they do not solve the underlying agency problems that allow sub-optimal advertising decisions to persist in the traditional TV advertising model we evaluate in this paper. Together with past research documenting similar results in digital advertising markets (Blake et al. 2015; Lewis and Rao 2015), our work should motivate economists to further study the managerial and agency issues in advertising markets.

[13] Another explanation is that many brands have objectives for advertising other than stimulating sales. This is a nonstandard objective in economic analysis, but nonetheless, we cannot rule it out.

References

Abraham, M. M. and L. M. Lodish (1990): "Getting the Most Out of Advertising and Promotion," *Harvard Business Review*.

Assmus, G., J. U. Farley, and D. R. Lehmann (1984): "How Advertising Affects Sales: Meta-Analysis of Econometric Results," *Journal of Marketing Research*, 21, 65–74.

Bagwell, K. (2007): "The Economic Analysis of Advertising," in *Handbook of Industrial Organization*, ed. by M. Armstrong and R. Porter, Elsevier B.V., vol. 3, chap. 28, 1701–1844.

Bar, R. N. and A. M. Haviv (2019): "Brand Building to Deter Entry and its Impact on Brand Value," *manuscript*.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

Blake, T., C. Nosko, and S. Tadelis (2015): "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment," *Econometrica*, 83, 155–174.

Borkovsky, R. N., A. Goldfarb, A. M. Haviv, and S. Moorthy (2017): "Measuring and Understanding Brand Value in a Dynamic Model of Brand Management," *Marketing Science*, 36, 471–499.

Clark, C. R., U. Doraszelski, and M. Draganska (2009): "The effect of advertising on brand awareness and perceived quality: An empirical investigation using panel data," *Quantitative Marketing and Economics*, 7, 207–236.

Deng, Y. and C. F. Mela (2018): "TV viewing and advertising targeting," *Journal of Marketing Research*, 55, 99–118.

Du, R. Y., M. Joo, and K. C. Wilbur (2018): "Advertising and Brand Attitudes: Evidence from 575 Brands Over Five Years," *Working Paper*.

Dubé, J.-P., G. J. Hitsch, and P. Manchanda (2005): "An Empirical Model of Advertising Dynamics," *Quantitative Marketing and Economics*, 3, 107–144.

Ellison, G. and S. F. Ellison (2011): "Strategic Entry Deterrence and the Behavior of Pharmaceutical Incumbents Prior to Patent Expiration," *American Economic Journal: Microeconomics*, 3, 1–36.

GELMAN, A. AND J. HILL (2007): *Data Analysis Using Regression and Multilevel/Hierarchical Models*, Cambridge University Press.

GOLDFARB, A. AND C. TUCKER (2011): "Online Display Advertising: Targeting and Obrusiveness," *Marketing Science*, 30, 389–404.

GORDON, B. AND W. HARTMANN (2013): "Advertising Effects in Presidential Elections," *Marketing Science*, 32, 19–35.

GRILICHES, Z. AND J. A. HAUSMAN (1986): "Errors in variables in panel data," *Journal of econometrics*, 31, 93–118.

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," .

HENNINGSEN, S., R. HEUKE, AND M. CLEMENT (2011): "Determinants of Advertising Effectiveness: The Development of an International Advertising Elasticity Database and a Meta-Analysis," *Business Research*, 4, 193–239.

HITSCH, G. J., A. HORTAÇSU, AND X. LIN (2019): "Prices and Promotions in U.S. Retail Markets: Evidence from Big Data," *manuscript*.

HRISTAKEVA, S. AND J. H. MORTIMER (2020): "Impacts of Legacy Discounts in the Market for National Television Advertising," *Working Paper*.

HUBER, G. AND K. ARCENEAUX (2007): "Identifying the Persuasive Effects of Political Advertising," *American Journal of Political Science*, 51, 957–977.

JOHNSON, G., R. LEWIS, AND E. NUBBEMEYER (2016): "The Online Display Ad Effectiveness Funnel & Carry-Over: A Meta-Study of Ghost Ad Experiments," *Working Paper*.

KALYANAM, K., J. MCATEER, J. MAREK, J. HODGES, AND L. LIN (2018): "Cross Channel Effects of Search Engine Advertising on Brick & Mortar Retail Sales: Meta Analysis of Large Scale Field Experiments on Google.com," *Quantitative Marketing and Economics*, 16, 1–42.

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

Lewis, R. A. and J. M. Rao (2015): "The Unfavorable Economics of Measuring the Returns to Advertising," *Quarterly Journal of Economics*, 130, 1941–1973.

Li, X., W. R. Hartmann, and T. Amano (2019): "Identification Using Border Approaches and IVs," *Available at SSRN 3402187*.

Lodish, L. M., M. Abraham, S. Kalmenson, J. Livelsberger, B. Lubetkin, B. Richardson, and M. E. Stevens (1995): "How T.V. Advertising Works: A Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments," *Journal of Marketing Research*, 32, 125–139.

Nevo, A. (2000): "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, 31, 395–421.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110, 705–729.

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

Sethuraman, R., G. Tellis, and R. Briesch (2011): "How Well Does Advertising Work? Generalizations from Meta-Analysis of Brand Advertising Elasticities," *Journal of Marketing Research*, 48, 457–471.

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.

Shapiro, B. T., G. J. Hitsch, and A. E. Tuchman (2020): "Research Note: Procedure for Building Nielsen Ad Intel Data and Merging with Nielsen RMS Scanner Data," Tech. rep.

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

Spenkuch, J. L. and D. Toniatti (2018): "Political advertising and election results," *Quarterly Journal of Economics*, 133, 1981–2036.

Sutton, J. (1991): *Sunk Costs and Market Structure*, MIT Press, Cambridge, MA.

Thomas, M. (2020): "Spillovers from Mass Advertising: An Identification Strategy," *Marketing Science (forthcoming)*.

Tuchman, A. E. (2019): "Advertising and Demand for Addictive Goods: The Effects of E-Cigarette Advertising," *Marketing Science*, 38, 994–1022.

Appendix (Supplementary Online Material)

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 $\tau \in \{t-L, \dots, t\}$ 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 advertising flow a is

$$\frac{dQ_t}{da}\frac{a}{Q_t} = \beta\frac{\rho a}{1+\rho a}. \tag{5}$$

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 $t, \dots, t+L$:

$$\left(\frac{\partial}{\partial a_t}\sum_{\tau=t}^{t+L} Q_\tau\right)\frac{a_t}{Q_t} = \beta\frac{\rho a}{1+\rho a}. \tag{6}$$

The effect of permanent percentage increase in advertising (5) is equivalent to the cumulative, long-run increase in demand (6). 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 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.[14] 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 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.

We show summary statistics for the estimated affiliated brand and top competitor

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

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%
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 using the baseline strategy and assuming a carryover parameter $\delta = 0.9$. Standard errors are two-way clustered at the DMA level and the week level.

advertising elasticities in Table 4 and histograms of the corresponding distributions of advertising effects in Figure 5. We also report own price elasticities and top competitor price elasticities in Table 4.

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.

Own price elasticities are centered around -1.6, 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 Hitsch et al. (2019).

C Systematic patterns in advertising elasticities

While a distribution of ad effects serves as a useful prior, many managers and researchers may want to know if advertising elasticities are systematically predictable given data on brand or industry characteristics. As noted in the previous section, the estimates with the best statistical power to detect an advertising effect tend to be those closest to the mean of the distribution. We illustrate this point in Figure 2. We arrange the brands on the x-axis in order of the elasticity estimates. On the y-axis we display the elasticities and the corresponding 95% confidence intervals. The estimates near the mean have the

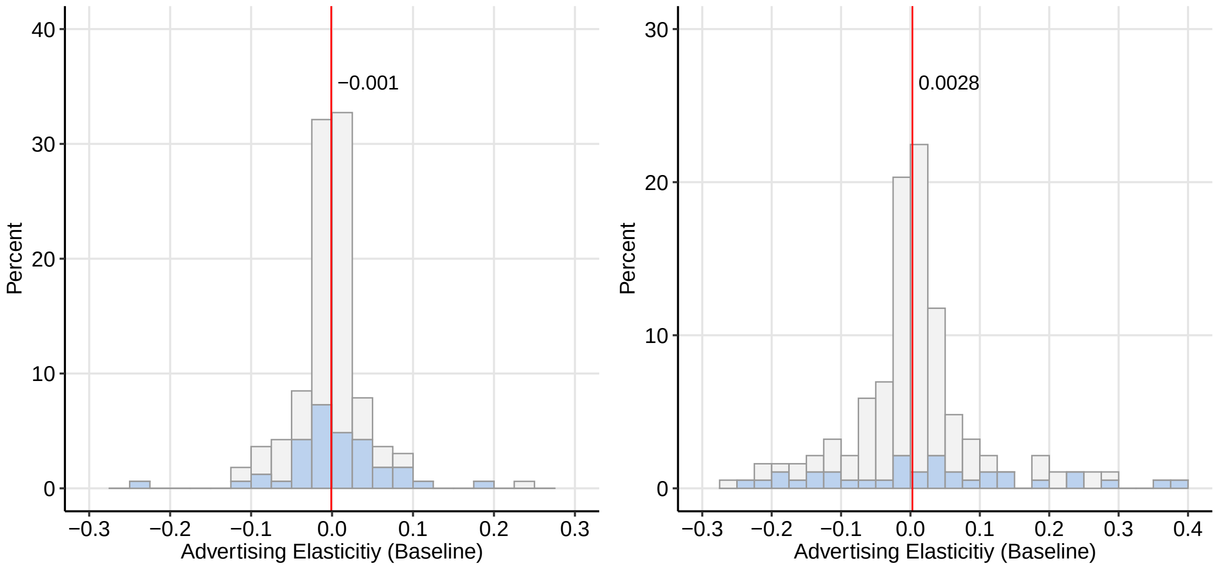

Figure 5: Affiliated Brand and Competitor Advertising Stock Elasticities

Note: The estimates are obtained assuming a carryover parameter $\delta = 0.9$. The left panel shows the distribution of affiliated brand advertising stock elasticities. The right panel shows the distribution of top competitor advertising stock elasticities. Bars highlighted in blue indicate statistically significant estimates. The vertical red line denotes the median of the distribution.

smallest confidence intervals, on average, whereas the large estimates in the right and left tails tend to be imprecisely estimated.

With this in mind, we study if there are systematic differences in the advertising elasticities across the 65 product categories or 8 grocery store departments in our data.[15] The average elasticity is significantly different from the median at the 5% level in two out of 65 product modules, and in one module the average elasticity is significantly different from the mean. Furthermore, the average elasticity is significantly different from the median (but not mean) in 1 out of 8 departments. Hence, as the incidence of statistically significant differences is roughly consistent with results that are obtained by chance, we do not find evidence for a systematic association between product categories or departments and advertising elasticities. This is not to say that other data, such as detailed brand characteristics, would not be able to predict systematic differences in the advertising elasticities. The ability to do so, however, is limited in our data.

[15] We use the Nielsen product module code to define categories.

D Robustness of Elasticity Estimates

D.1 Robustness to other modeling choices

In Table 2, we provide the medians, means and quantiles of the distributions corresponding to different robustness analyses. In all variations, the estimated distributions are very similar.[16] Details of each analysis are discussed below.

Border Strategy To provide a set of estimates using an alternative identification strategy, we estimate a specification that takes advantage of DMA boundaries. The specification is similar to the baseline specification, but focuses on counties at the borders of DMAs and includes border-time specific fixed effects (Shapiro (2018), Tuchman (2019), Spenkuch and Toniatti (2018)).

It is important to note that the border strategy presents a trade-off with the baseline strategy. On the negative side, the border strategy drops nearly 80% of observations, which will reduce statistical power and potentially provide results that are only relevant to the border counties. Additionally, while the baseline and border strategies both require the stable unit treatment value assumption (SUTVA) that consumers who are exposed to advertising do not cross the border and purchase in the neighboring DMA, the border strategy may more plausibly violate that assumption. On the positive side, the border strategy theoretically allows for more sophisticated targeting of advertising by firms without confounding estimates. The border strategy will also sweep out noise in the left hand side variable caused by geographic differences in preferences, due to the geographic-time specific fixed effects

If the baseline and border strategies provide different results, there are two ways to interpret the differences. First, we could assume that both the baseline and border strategies are unbiased, which would lead us to interpret the difference as evidence of treatment effect heterogeneity where the border sample has different ad effectiveness. Alternatively, we could assume that there is no treatment effect difference between the samples, and the difference in the estimates would be interpreted as evidence of confounding in the baseline strategy.

The results of the border strategy specification are not distinguishable from the results in the baseline specification. Hence, the interpretation of the difference is moot in our case.

[16] More results on sensitivity and robustness are available here: https://advertising-effects.chicagobooth.edu/.

Calibration of carryover parameter In Table 5, we present the estimation results for various values of the carryover parameter, $\delta = 0, 0.25, 0.5, 0.75, 0.9, 0.95, 1$. The mean and median of the estimated coefficients change when we change the assumed δ. However, the percentage of statistically insignificant coefficients, the percentage of positive and statistically significant coefficients, and the percentage of negative coefficients is robust to any of the assumed δs.

Estimation of carryover parameter In our main specification we calibrate the carryover parameter to $\delta = 0.9$. Here we estimate δ using a grid search from 0 to 1 in increments of 0.05. For each point in the grid, we calculate the advertising stock using equation (2) and then estimate the remaining model parameters via OLS. For each brand, the estimated δ is the carryover parameter that minimizes the predicted mean squared error.

Estimating δ will yield more accurate advertising effects if the assumption that $\delta = 0.9$ is false or if there is heterogeneity across brands in the degree of advertising carryover. A downside is that if the advertising elasticity is zero ($\beta = 0$), then δ is not identified. In this case, if δ is not restricted, the estimates will be uniformly distributed on $(-\infty, \infty)$. However, since we impose the constraint that $0 \leq \delta \leq 1$, the estimated carryover parameter will likely be at the bounds of the grid, $\delta = 0, 1$. Similarly, in cases where the advertising elasticity β is not precisely estimated, it is likely that δ is also hard to pin down and takes values on the bounds of the grid.

We report the results in Table 5. The distribution has a similar mean and median but exhibits a larger spread compared to the case when we set $\delta = 0.9$.

Own price The basic model includes own prices. However, it is possible that changes in prices are an "outcome" of advertising, making them bad controls. It is also possible that net of the fixed effects in the model, residual prices are correlated both with the error term and the ad stock. We find that the distribution of ad effects is similar if we do not control for prices. We show in Appendix F that this should be expected, as net of the fixed effects in the model, price and advertising are uncorrelated.[17]

Included competitor prices In our main specifications, we include the prices of up to three competing products that have the largest total sales revenue in the category.[18] The results are unchanged whether we include one, two or three of the largest competitors.

[17]We show this is true for temporary promotions as well as general price changes.

[18]Categories are defined based on the Nielsen product module code.

Table 5: Own-Advertising Stock Elasticities by Carryover, δ

		Median	Mean	% p $\geq$ 0.05	% p < 0.05		Percentiles			
					> 0	≤ 0	10%	25%	75%	90%
Baseline Specification										
Specified δ	0.00	0.0023	0.0030	69.10	23.96	6.94	-0.0079	-0.0031	0.0070	0.0135
	0.25	0.0042	0.0040	69.10	24.31	6.60	-0.0087	-0.0024	0.0092	0.0171
	0.50	0.0050	0.0062	66.32	26.39	7.29	-0.0136	-0.0036	0.0123	0.0254
	0.75	0.0064	0.0115	63.54	29.17	7.29	-0.0173	-0.0052	0.0208	0.0439
	0.90	0.0140	0.0233	66.32	26.39	7.29	-0.0406	-0.0082	0.0426	0.0919
	0.95	0.0145	0.0283	67.36	22.92	9.72	-0.0667	-0.0181	0.0661	0.1483
	1.00	0.0067	0.0137	69.79	18.40	11.81	-0.1498	-0.0292	0.0797	0.2000
Estimated δ		0.0090	0.0116	51.04	35.42	13.54	-0.1102	-0.0149	0.0530	0.1733

Note: Descriptive statistics of estimated advertising elasticities reported for two model specifications and 288 brands. Elasticities derived from regressions of log quantity on log advertising GRP stock (own and competitor) and log prices (own and competitor). The baseline model includes store, month, and week-of-year fixed effects. Regressions are estimated separately for each brand. The unit of observation in each regression model is a store-brand-week. Standard errors are two-way clustered at the DMA level and the week level.

Hence, we conclude that the omission of additional competitor prices is unlikely to bias the results.

Feature and display advertising Feature and in-store display advertising by retail chains is typically funded by the brand manufacturers and may hence be coordinated with the advertising campaigns. The results are unchanged when we include feature and display advertising. We show in Appendix F that this should be expected, as net of the fixed effects in the model, feature, display and advertising are uncorrelated.

Statistical power 157 of the 288 brands have at least 50% ex ante power to detect an advertising elasticity of 0.05 at the 5% level.[19] Within this set, the median advertising elasticity is 0.0085, and the mean is 0.0098. 65% of the elasticities are not statistically significant. Hence, the large incidence of estimates that are not statistically significant in the full sample of brands is not simply due to noise. The distribution of advertising effects in the smaller set of 157 brands is compressed—all estimates are less than 0.1 in absolute value. The 90$^{\text{th}}$ percentile of the distribution is 0.0428, compared to 0.0919 in the full

[19]Specifically, we identify the set of brands for which the standard error of the brand's estimated ad effect is less than or equal to $^{0.05}/_{1.96}$ (Gelman and Hill, 2007).

sample. Hence, the large estimates in the full sample seem to indicate a significant degree of noise rather than a truly large advertising effect.

D.2 Measurement error

If there is classical measurement error in advertising net of fixed effects, the advertising elasticity estimates will be biased towards zero. We measure advertising using GRPs, which are constructed from advertising occurrences and the corresponding television viewership, i.e., number of households who watched the program where an ad was aired. Nielsen uses an automated process involving pattern recognition technology to measure occurrences. Correspondingly, advertising occurrences and durations are likely to be measured accurately. The advertising viewership is predicted based on the viewing behavior of a sample of households, the "Nielsen Families." In the top 25 markets, viewership is measured using electronic devices called People Meters. In all other markets, viewership information is collected in the form of self-reported diaries. Hence, the viewership data are the most likely source of measurement error in advertising.

The timing of an ad exposure relative to a shopping trip may induce measurement error. In any given week, some households will be exposed to ads after they have completed the relevant shopping trips. As a result, the measured amount of advertising in the concurrent week may overstate the amount of advertising that could have affected purchasing behavior.

Below we summarize the four approaches we use to assess if the relatively small advertising elasticities that we documented are due to attenuation bias resulting from measurement error.[20] We report the results in Table 6.

Advertising stocks based on durations We use only the occurrence data, which are unlikely to be measured with error, to construct the advertising stocks. The estimated advertising elasticities are similar to the original estimates that are based on GRPs.

LPM markets only Advertising impressions measured using People Meters may be more accurate than the self-reported diary entries. We hence re-estimate the advertising elasticities using only the 25 LPM (Local People Meter) markets. The estimates are similar compared to the estimates using all markets.

[20] Note, however, that previous studies of TV advertising effectiveness utilize similar data and are hence subject to similar concerns regarding measurement error in advertising.

Table 6: Measurement Error Analysis: Own-Advertising Stock Elasticities

	Median	**Mean**	**% p $\geq$ 0.05**	**% p $<$ 0.05**		**Percentiles**			
				> 0	≤ 0	**10%**	**25%**	**75%**	**90%**
Baseline Specification									
GRP, All markets, Time FE	0.0140	0.0233	66.32	26.39	7.29	-0.0406	-0.0082	0.0426	0.0919
Occurrence	0.0170	0.0315	63.89	29.51	6.60	-0.0390	-0.0079	0.0482	0.1289
1-week-lagged GRP	0.0118	0.0207	69.44	23.96	6.60	-0.0458	-0.0091	0.0411	0.1050
LPM markets	0.0118	0.0224	74.65	19.79	5.56	-0.0388	-0.0090	0.0428	0.1030
Time Trend	0.0110	0.0171	41.67	42.36	15.97	-0.0360	-0.0053	0.0381	0.0772

Note: Elasticities derived from regressions of log quantity on log advertising stock (own and competitor) and log prices (own and competitor). All use modifications to the baseline strategy as described.

Fixed effects and measurement error Fixed effects may exacerbate attenuation due to classical measurement error (Griliches and Hausman 1986). In Table 2 we show how the estimates change when we incrementally add fixed effects to the model. The most granular fixed effects are the time fixed effects. The inclusion of these fixed effects does not yield smaller estimates than those in a specification using a time trend in lieu of time fixed effects.

One-week-lagged advertising stock We estimate the model using advertising stocks that are lagged by one week. This ensures that all advertising exposures occur before a shopping trip. To obtain elasticities that are comparable to the original estimates, we divide the coefficient on the lagged ad stock by 0.9 to account for the one-week decay. The results are similar to the original estimates.

Based on these analyses, we conclude that the robust small magnitudes of estimated ad elasticities are unlikely due to measurement error.

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 $10, 20, \dots, 90$ of the advertising stock.

This basis includes the main parametric model specification and a cubic B-spline as special cases. 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 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.[21]

We provide the advertising elasticity estimates from the flexible model specification in Figure 6, and in the bottom right panel of Figure 3, we show that these results are highly correlated with the parametric model results. We provide estimates of the flexible functional form for all brands in the interactive online appendix. As a specific example, in Figure 7 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

[21] Alternatively, we could calculate the advertising elasticity at the mean or median of the observed advertising stock. However, this approach yields noisy results due to wiggles in the estimated advertising response function, i.e. deviations from the overall slope that are local around the mean and median of the advertising stock that do not reflect the slope of the overall advertising response curve.

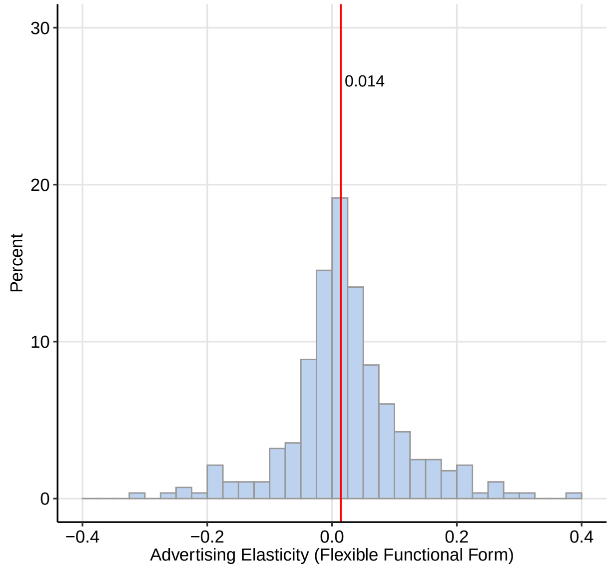

Figure 6: Advertising Elasticities using Semi-Parametric Estimation

Note: The graph shows the distribution of the estimated brand-level summary advertising elasticities obtained using the flexible functional form for the advertising response. The vertical line denotes the median of the distribution.

get relatively precise estimates of the flexible advertising response curve, too. For both brands 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 models indicates that the main results are not driven by the specific functional form assumptions. 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.

F 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.

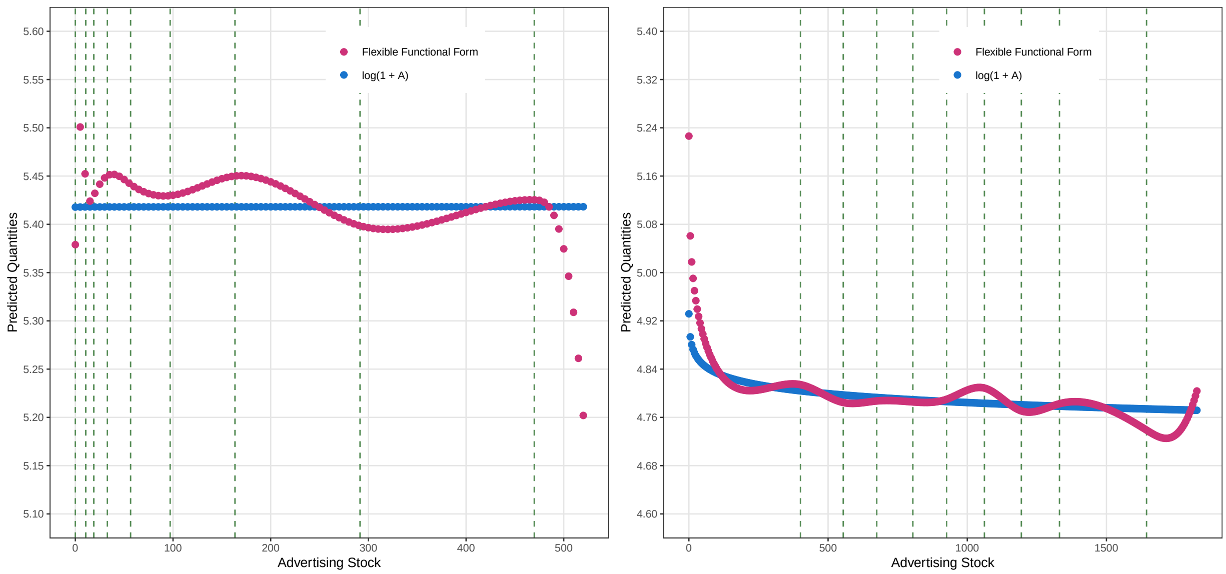

Figure 7: Predicted Quantity using Baseline and Semi-Parametric Estimation

Note: The left panel is for Chobani, while the right panel is for Gatorade. In both panels, we use the baseline strategy model (red dots) with month, seasonal, and store fixed effects, and $\delta = 0.9$. Predicted quantity is plotted against advertising stock, with the semi-parametric advertising response curves in red and the parametric advertising response curves in blue.

In Table 7, we show 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. Without the fixed effects in the model, the median regression coefficient is –0.083, 36% of estimates are negative and significant, and 18% are 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 –0.007, with less than 10% of estimates negative and significant and less than 5% positive and 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*, feature advertising, and display advertising. The results are similar to the price results. Full results are available in Table 7.

We also conduct the same analysis, but using advertising stocks instead of advertising flows. 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

Table 7: 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 Specification									
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

lower magnitude of correlation to these variables than does advertising flow, net of the fixed effects.

Overall, we conclude that net of the fixed effects in our main specifications, price, feature and display are uncorrelated with advertising. Hence, the choice to omit or include each in the model should not substantively alter the results.

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:

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

Table 8: 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 Specification									
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

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:

$$\lambda_{s\tau} \equiv \frac{Q'_{s\tau}}{Q_{s\tau}} = \frac{(1+A'_{d\tau})^{\beta}}{(1+A_{d\tau})^{\beta}} = \left(\frac{1+A_{d\tau}+\delta^{\tau-t}\Delta a_d}{1+A_{d\tau}}\right)^{\beta}. \tag{7}$$

Notably, all store, season and time-specific components cancel out, and thus equation (7) 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:

$$\Delta\pi_{d\tau} = \sum_{s\in\mathcal{S}_d}(\lambda_{d\tau}-1)Q_{s\tau}\cdot m\cdot p_{s\tau}, \tag{8}$$

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.[22] Summing across all DMAs and all periods $\tau \in \{t, ..., t+L\}$

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

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:

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

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 baseline strategy with the carryover parameter $\delta = 0.9$.[23] In order to calculate incremental profits using equation (8), 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 4.2).[24] 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, $\bar{Q}_d^R$.[25] 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:

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

[23] 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.

[24] 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. Separating out the 24 beer and soft drink brands does not significantly alter the distribution of ROIs. Additionally, assuming that all beer and soft drink brands have sales volumes that are twice the volumes that we predict does not significantly alter the distribution of ROIs.

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

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%.[26] In Table 3, we show 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.[27] 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.

In Figure 8, we summarize 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.

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

[27] Shapiro et al. (2020) provides more detail about the advertising expenditure data.

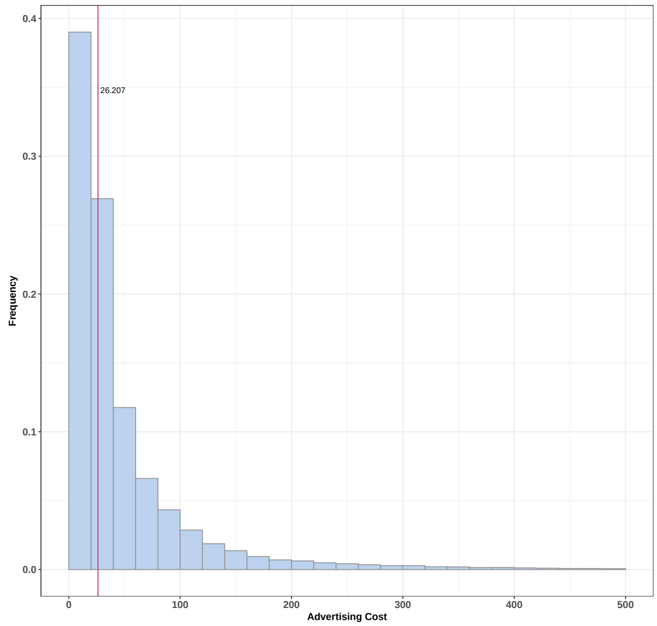

Figure 8: Distribution of Cost per GRP across DMA-Brand-Years

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 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, we show how the break-even ad effect varies as a function of the assumed margin factor m and the chosen ROI metric in Figure 9.

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.[28] In Figure 10, we show 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 with a positive estimated ad elasticity, the estimated ad effect would need to be 5.368 times larger in order for the observed level of weekly advertising to be profitable (assuming a margin factor of $m = 0.3$). In contrast, in the right panel of Figure 10, we show the results when considering the ROI of all observed advertising.

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

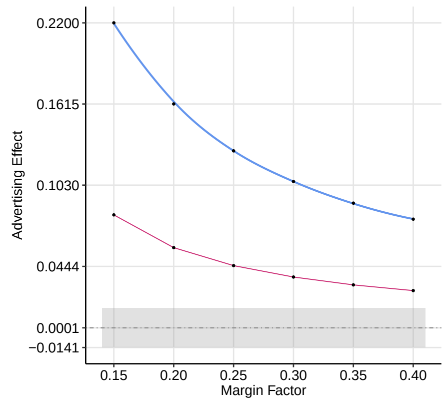

Figure 9: Break-Even Advertising Effect (Chobani)

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 0.0001 (gray dashed line) and the shaded area marks the 95% confidence interval.

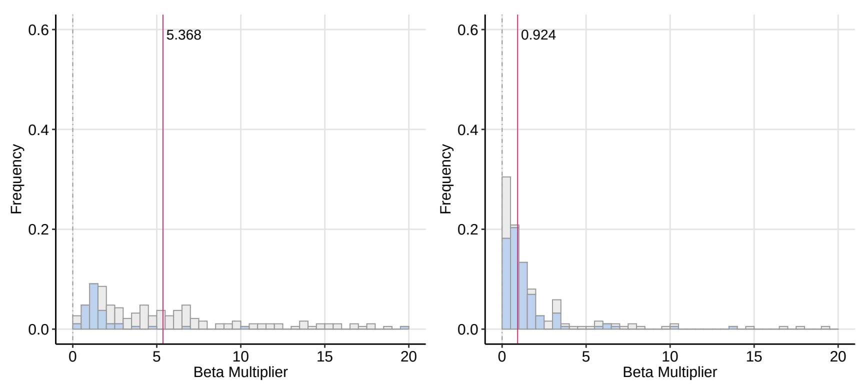

Figure 10: Ratio of Break-Even Ad Effect to the Estimated Ad Effect

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 187 brands with a positive estimated ad effect.