Econometrica Supplementary Material

SUPPLEMENT TO "TV ADVERTISING EFFECTIVENESS AND PROFITABILITY: GENERALIZABLE RESULTS FROM 288 BRANDS"

(*Econometrica*, Vol. 89, No. 4, July 2021, 1855–1879)

BRADLEY T. SHAPIRO
Booth, University of Chicago

GÜNTER J. HITSCH
Booth, University of Chicago

ANNA E. TUCHMAN
Kellogg, Northwestern University

APPENDIX SA: ROBUSTNESS OF ELASTICITY ESTIMATES

SA.1. *Robustness to Other Modeling Choices*

IN TABLE II, we provide the medians, means, and quantiles of the distributions corresponding to different robustness analyses. In all variations, the estimated distributions are very similar.[1] Details of each analysis are discussed below.

Calibration of Carryover Parameter. In Table SI, 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 SI. The distribution has a similar mean and median but exhibits a larger spread compared to the case when we set $\delta = 0.9$.

Bradley T. Shapiro: Bradley.Shapiro@chicagobooth.edu
Günter J. Hitsch: guenter.hitsch@chicagobooth.edu
Anna E. Tuchman: anna.tuchman@kellogg.northwestern.edu

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

 https://doi.org/10.3982/ECTA17674

TABLE SI

OWN-ADVERTISING STOCK ELASTICITIES BY CARRYOVER, δ[a]

		Median	Mean	% $p \geq 0.05$	% $p < 0.05$ > 0	% $p < 0.05$ ≤ 0	Percentiles 10%	Percentiles 25%	Percentiles 75%	Percentiles 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

[a]*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.

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 Online Appendix SC that this should be expected, as net of the fixed effects in the model, price and advertising are uncorrelated.[2]

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.[3] The results are unchanged whether we include one, two, or three of the largest competitors. 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 Online Appendix SC that this should be expected, as net of the fixed effects in the model, feature, display and advertising are uncorrelated.

Statistical Power. One hundred fifty-seven of the 288 brands have at least 50% ex ante power to detect an advertising elasticity of 0.05 at the 5% level.[4] Within this set, the median advertising elasticity is 0.0085, and the mean is 0.0098. Sixty-five percent 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

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

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

[4]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 $\frac{0.05}{1.96}$ (Gelman and Hill (2007)).

TABLE SII

MEASUREMENT ERROR ANALYSIS: OWN-ADVERTISING STOCK ELASTICITIES[a]

	Median	Mean	% $p \geq 0.05$	% $p < 0.05$ > 0	% $p < 0.05$ ≤ 0	Percentiles 10%	Percentiles 25%	Percentiles 75%	Percentiles 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
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

[a]*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.

are less than 0.1 in absolute value. The 90th percentile of the distribution is 0.0428, compared to 0.0919 in the full sample. Hence, the large estimates in the full sample seem to indicate a significant degree of noise rather than a truly large advertising effect.

SA.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 of 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. 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.[5]

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

Advertising Stocks Based on Durations. We use only the occurrence data, which are unlikely to be incorrectly recorded, to construct the advertising stocks. The estimated advertising elasticities are similar to the original estimates that are based on GRPs.[7]

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.

[5]It could also be that GRP is an imperfect proxy for a more relevant measure of advertising for some brands. However, advertising contracts with television networks are typically specified in terms of GRPs delivered. As a result, for most brands, GRPs are particularly relevant when evaluating the return on dollars spent.

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

[7]Like with GRPs, occurrences may not be the most relevant measure of advertising to the brand, so this analysis does not rule out all possible measurement error.

Fixed Effects and Measurement Error. Fixed effects may exacerbate attenuation due to classical measurement error (Griliches and Hausman (1986)). The most granular fixed effects in the baseline model are the time fixed effects. We compare the baseline specification that includes time fixed effects (replicated in row 1 of Table SII) with a more parsimonious specification that replaces the time fixed effects with a time trend (row 4 of Table SII). The inclusion of the more granular time fixed effects does not yield smaller estimates than the specification using a time trend in lieu of time fixed effects.

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

APPENDIX SB: 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, $A^2, \dots, A^{10}$,
2. Transformations of $\log(1+A)$: $(\log(1+A))^{\frac{1}{2}}$, $\log(1+A)$, $(\log(1+A))^2, \dots, \log(1+A))^{10}$,
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, $\tilde{\log(Q)}$. 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 a summary statistic for each brand.[8]

In Figure S1, we provide the advertising elasticity estimates from the flexible model specification and 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 S2 we plot the predicted advertising stock response function for two brands, using both the flexible, semi-parametric model and the $\log(1 + A)$ functional form. We chose the two brands, Chobani and Gatorade,

[8]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, that is, 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.

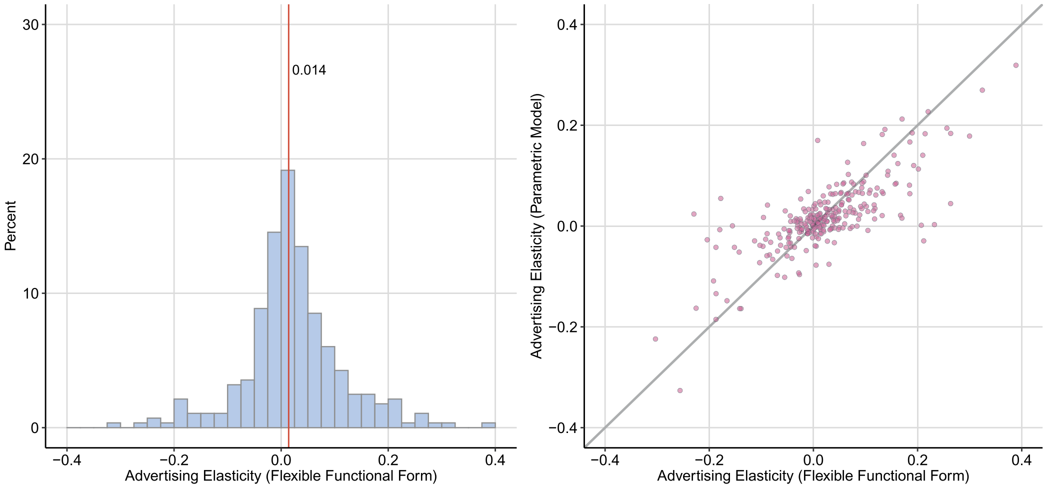

FIGURE S1.—Advertising elasticities using semi-parametric estimation. *Note*: The left graph shows the distribution of the estimated brand-level summary advertising elasticities obtained using the flexible functional form for the advertising response. The vertical line denotes the median of the distribution. The right graph plots the results of the baseline (parametric) and semi-parametric results against each other. Each dot represents a brand. The 45 degree line is presented with the plots. Advertising stocks are computed assuming a carryover parameter of $\delta = 0.9$.

because for both of them we obtain reasonably precise parametric estimates of the advertising effect, making it plausible that we could get relatively precise estimates of the flexible advertising response curve, too. For both brands, 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.

APPENDIX SC: 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.

In Table SIII, 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. Thus, it appears that many firms are coordinating advertising with price reductions and many others with price increases. 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

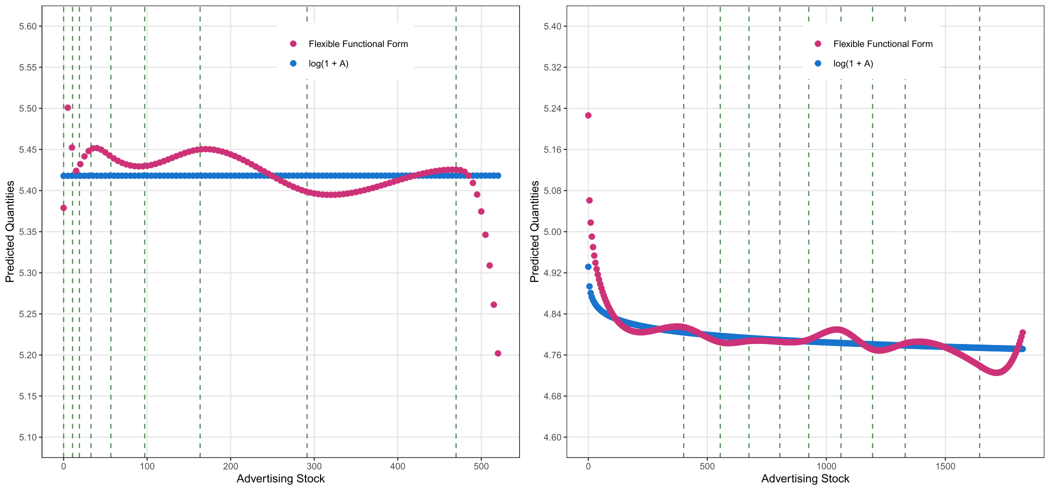

FIGURE S2.—Predicted quantity using baseline and semi-parametric estimation. *Note*: The left panel is for Chobani, while the right panel is for Gatorade. In both panels, we use the baseline strategy model (red dots) with month, seasonal, and store fixed effects, and $\delta = 0.9$. Predicted quantity is plotted against advertising stock, with the semi-parametric advertising response curves in red and the parametric advertising response curves in blue.

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 SIII.

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

TABLE SIII

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

TABLE SIV

CORRELATIONS BETWEEN AD STOCK 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	−0.0222	−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

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.

APPENDIX SD: DATA CONSTRUCTION

This appendix details exactly how we construct our data. We need the following data for each brand:

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

We construct the data using the following steps:

1. Build Ad Intel data
   (a) The ad occurrences and viewerships are separate in the raw Ad Intel data. We need to merge them in order to find the GRPs for each advertisement.
   (b) There are some discrepancies between the national and local records of Network TV ads. We need to resolve those discrepancies.
2. Create brand map between Ad Intel and RMS data sets
   (a) Ad Intel and RMS use different brand definitions, so for each RMS brand, we need to find all the corresponding Ad Intel brands.
3. Aggregate data
   (a) RMS data come in UPC-Store-Week level, so we need to aggregate it to Brand-Store-Week level.
   (b) Ad Intel data come in {Ad Intel Brand}-Market-Channel-Second level, so we need to aggregate it to {RMS Brand}-Market-Week level.
4. Identify RMS stores to be used in estimation
5. Identify products to be used in estimation

Each of these steps is described in more detail below.

SD.1. *Build Ad Intel Data*

SD.1.1. *General Concepts*

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

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

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

- The Network TV and Syndicated TV occurrence files record them at a national level,
  - that is, the date and time each ad is supposed to be broadcast at every local station.
- The Network Clearance Spot TV and Syndicated Clearance Spot TV occurrence files record them at a local channel level,
  - that is, the date and time each ad is actually broadcast at every local station.
- It is important to emphasize that the Network (Syndicated) TV occurrences and the Network (Syndicated) Clearance Spot TV occurrences *are meant to represent the exact same ad occurrences*. The main distinction here is that the "clearance" occurrences are measured at the local DMA level, allowing us to match them with local impressions data. They do not represent two distinct forms of advertising. For our purposes, the Network (Syndicated) TV occurrences (measured nationally) allow us to (i) measure the estimated cost of the ad, which is purchased nationally, and (ii) audit the advertising occurrence schedule in the "clearance" data to detect times when a local station might override a nationally purchased ad and distinguish it from local measurement error.
- To further explain the previous point, local channels have some authority to replace or move nationally scheduled ads, and it is possible that the Nielsen local measurement devices are not perfect. Hence, not every "Network TV" advertising occurrence will match perfectly with a "Network Clearance Spot TV" occurrence, even though they represent the same ad buys. Any incidence of "Network TV" occurrences that do not have associated "Network Clearance Spot TV" occurrences in a local market will be referred to as the "Missing Network Discrepancy."

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

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

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

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

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

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

SD.1.2. *Build the Regular Parts*

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

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

SD.1.3. *Resolve the "Missing Network" Discrepancy*

As mentioned previously, each Network ad that is actually shown on TV in a local market should be associated with a Network Clearance Spot TV ad. These represent the exact same showing of the exact same ads. For the most part, ads purchased as Network TV ads are shown "simultaneously" (in the same slot during the same television program) across the country. However, in rare circumstances, a Network TV ad may not actually run in one or more local markets. Additionally, in even rarer circumstances, a Network TV ad may run in a local market but not get recorded in the Network Clearance occurrences due to an error in the recording devices at the local level. Our goal in this part of the build is to distinguish the explanation when Network TV and Network Clearance Spot TV ads do not match in a given market at a given time.

We do this by reconstructing the TV schedule, down to the second. From the observed schedule, we infer whether or not it was a recording error by the local measurement device or whether the local station likely displaced the Network ad. If, for example, there is a "gap" in the schedule that is 30 seconds long and a 30 second long Network TV ad purportedly ran in that time interval, we infer that the Network TV ad actually ran on the local station in that slot but the recording device failed in that moment. In these cases, we "insert" the Network TV ad occurrence into the advertising schedule in the "gap" where we inferred that it belonged. Alternatively, if there is no "gap" in the schedule (i.e., we observe ads back-to-back every second followed by programming), we then infer that the local station displaced the Network ad either for extended programming (e.g., local news alert, sporting event gone long) or for an additional locally purchased Spot TV ad. Network Clearance TV ads also are always missing between 2AM and 5AM, as local ratings are not recorded during those time intervals.

While in principle this exercise is intuitively simple, in practice, this procedure is complicated to implement. We take the following steps:

1. Find the information for each local station, including:
    (a) The market (MarketCode) and network (Affiliation) for each local station (DistributorCode).
    (b) The DistributorID for each DistributorCode.
        i. This is in fact a one-to-one relationship, but we have to record that because the "Station Affiliation" data only have DistributorCode, while the impressions data only have DistributorID.
2. For each network and each local station, we stack all the monthly data.
    (a) We cannot use the raw monthly data because the national and local files have different dates.
    (b) Stacking also prevents errors at month boundaries. For example, a national ad at 2012/05/31 23:30:00 may be distributed locally at 2012/06/01 00:30:00. This will not be captured if we process the data month-by-month.
3. For each local station, we find the "unexpectedly missing" occurrences. In short, we categorize all the national ads as following:
    (a) A national ad is directly matched to the local data if its closest local occurrence has the same primary brand code.
    (b) A national ad is indirectly matched to the local data if there is a local occurrence that is aired within some time limit before or after the scheduled air-time. This step accounts for the ads that are moved around. The time limit is 3 hours for ETZ/CTZ, 6 hours for MTZ, and 7 hours for PTZ.
    (c) A national ad is replaced by another ad if another spot / network clearance / syndicated clearance ad runs into its scheduled time slot.
    (d) A national ad is not captured locally if its scheduled air-time is between 2AM and 5AM.
    (e) We mark all remaining national ads as unexpectedly missing at this local station. These are the "gaps" described above.
4. We get all the "unexpectedly missing" occurrences at each station, and we reorganize them into monthly files. We then merge those monthly files with the monthly local impressions data.

Note: We must be careful to account for the "broadcast delay" for Mountain and Pacific Time zones.

- A nationally scheduled program or ad can be broadcast with a delay of 0/1/2/3 hours in Pacific Time markets or 0/1 hours in Mountain Time markets. This delay can be seemingly arbitrary.
- In step 3, we say a national ad is "unexpectedly missing" only if it is "unexpectedly missing" under all the possible delays, that is, 0/1 hour in MTZ and 0/1/2/3 hours in PTZ.
- In step 4, for PTZ/MTZ markets, we average the impressions at the airtime and 3/1 hour(s) after the airtime.

SD.2. *Create Brand Map Between RMS and Ad Intel*

We merge the advertising and sales data sets at the store-brand-week level. This merge is non-trivial and non-obvious. In particular, the brand variables in the Ad Intel and RMS data sets are not always specified at the same level. For example, UPCs in the RMS data are sometimes much more specific than the generic brands in the Ad Intel data. Our matching procedure results in three distinct "types" of advertising variables to be used in our models. First, we specify 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.

The procedure for creating this match is as follows. First, we create a map between the brands in the RMS and Ad Intel data sets using an automated string matching procedure. Second, the three authors and two research assistants hand-audited the matches to classify them into tiers that are associated with the above description. Finally, any brands for which the authors could not come to an agreement on classification of matches were thrown out as "failed matches" and not included in the sample. We classify the matches in four "tiers," which are described below. In theory, tier-1 and tier-2 advertising should have a positive effect on sales, while the effect of tier-3 and tier-4 ads can be either positive or negative. Based on this logic, we construct our measure of own advertising by grouping tier-1 and tier-2 advertising together and our measure of "affiliated brands" advertising by grouping tier-3 and tier-4 advertising together.

Own Advertising.

- Tier 1: RMS and Ad Intel brand names are exact matches.
- Tier 2: Ad Intel brand is more specific than the RMS brand.
  - Example: Ad Intel brand LAYS POTATO CHIPS CHICKEN AND WAFFLE is a tier-2 match to RMS brand LAY'S.
- For the median brand in our data, tier-1 matches make up 43% of own advertising GRPs, while tier-2 matches make up the remaining 57% of own advertising GRPs. The distribution of GRPs coming from identical matches is shown in Table SV.

Affiliated Brands Advertising.

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

TABLE SV

FRACTION OF OWN ADVERTISING GRPS FROM DIFFERENT MATCHES

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

  - Example: Ad Intel brand COCA-COLA ZERO DT is a tier-4 match to RMS brand COCA-COLA R.

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

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

SD.3. *Aggregate Data*

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

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

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

RMS. The RMS data build comes at the UPC-Store-Week level, and we want to aggregate it to the Brand-Store-Week level. As mentioned before, UPCs are generally much more specific than brands, as they reflect many different sizes and presentations.

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

SD.4. *Store and Border Selection*

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

For the implementation of the border strategy, we use the Nielsen provided mapping between counties and DMAs. From this, we construct a data set that flags the counties

TABLE SVI

FREQUENCY OF DEPARTMENTS AND REVENUE SHARE[a]

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

[a]*Note*: Three brands in our sample have products in two departments.

that lie on a border between DMAs. However, some counties change DMAs over time, since the borders are redrawn periodically. Therefore, we remove all the counties that did not stay in a single DMA, and we remove the borders that were redrawn.

SD.5. *Product Selection*

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

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

SD.6. *Advertising Cost*

We estimate the cost of buying an ad GRP in DMA d in week t for each manufacturer using data on advertising expenditure, impressions, and audience size contained in the Nielsen Ad Intel data set. We use these cost estimates, together with advertising effect estimates, to compute advertising return on investment (ROI).

Expenditure Data.

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

TABLE SVII

FREQUENCY OF CATEGORIES[a]

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

[a]*Note*: Four brands in our sample have products in two categories.

  - Expenditure data are originally at the {Month}-{Network}-{Daypart} level for national and cable ads. Ad Intel further prorates expenditure and records the data at the {AdTime}-{Network}-{Daypart}-{Program}-{Duration} level.
- For Spot TV, ads are purchased at the DMA level.
  - Nielsen estimates spot TV expenditures by blending cost-per-point data supplied by SQAD with Nielsen's local market ratings data. SQAD's cost-per-point data are based on actual spot television buys reported by contributing ad agencies.

SD.6.1. *Build Advertising Cost*

For each manufacturer, we do the following:

1. Merge expenditure with impressions for each ad occurrence.

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

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

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

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

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

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

REFERENCES

GELMAN, A., AND J. HILL (2007): *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press. [2]

GRILICHES, Z., AND J. A. HAUSMAN (1986): "Errors in Variables in Panel Data," *Journal of Econometrics*, 31, 93–118. [4]

Co-editor Aviv Nevo handled this manuscript.

Manuscript received 17 September, 2019; final version accepted 16 March, 2021; available online 22 March, 2021.