

# Advertising Effectively Influences Older Users: How Field Experiments Can Improve Measurement and Targeting

Randall A. Lewis and David H. Reiley\*

First Draft: 16 January 2009

This Version: 5 April 2012

## Abstract

Does advertising affect sales in a measurable way? New technologies for tracking both sales and advertising at the individual level are used to investigate the effectiveness of brand advertising for a nationwide retailer. A controlled experiment on 1,577,256 existing customers measures the *causal* effect of advertising on actual purchases, overcoming the major hurdles regarding attribution typically encountered in advertising effectiveness research by exogenously varying exposure to the ads. Online brand advertising generated a statistically and economically significant effect on in-store sales for this retailer. The design of the experiment permits a demographic breakdown of the advertising's heterogeneous effects. Somewhat surprisingly, the effects are especially large for the elderly. Customers aged 65 and older, comprising only 5% of the experimental subjects, exhibited a 20% average increase in sales due to the advertising campaign, which represents 40% of the total effect among all age groups.

Major Classification: Social Sciences

Minor Classification: Economic Sciences

---

\* [randallalewis@gmail.com](mailto:randallalewis@gmail.com) and [david@davidreiley.com](mailto:david@davidreiley.com), Yahoo! Research, 4401 Great America Parkway, Santa Clara, CA 95054. We thank Meredith Gordon, Sergei Matushevych, and especially Taylor Schreiner for work on the experiment and data. Yahoo! Incorporated provided financial and data assistance, as well as guaranteeing academic independence prior to our analysis, so that the results could be published no matter how they turned out. We acknowledge the helpful comments of Avi Goldfarb, Glenn Ellison, Jerry Hausman, Stephen Ryan, Duncan Watts, and two anonymous referees.

## I. Introduction

Retailing pioneer John Wanamaker (1838-1922) famously remarked, “Half the money I spend on advertising is wasted; the trouble is I don’t know which half.” Measuring the impact of advertising on sales has remained a difficult problem for more than a century (Bagwell 2008). The econometric problems of endogeneity and omitted-variable bias make it very difficult to establish causality, rather than mere correlation, from observational data. In the present paper, we overcome this problem with a large-scale controlled experiment for an online advertising campaign and decompose the estimated effects of the advertising by age and gender. We find the surprising result that this advertising has its largest impact on older people.

As an example of the opportunity to draw mistaken conclusions from non-experimental data, we consider a recent state-of-the-art study by marketing practitioners (Abraham 2008). We quote this study, which estimates large positive effects of advertising on sales, to describe its methodology, “Measuring the online sales impact of an online ad or a paid-search campaign—in which a company pays to have its link appear at the top of a page of search results—is straightforward: We determine who has viewed the ad, then compare online purchases made by those who have and those who have not seen it.”

With observational data, this straightforward technique can give spurious results. The population who sees a particular ad may be very different from the population who does not see it. For example, those people who see an eTrade ad when searching for “online brokerage” are very different from those who do not see that ad, even in the absence of advertising. Almost certainly, those who search for “online brokerage” are much more likely to sign up for an eTrade account than those who do not search for “online brokerage.” So the observed difference in sales may not represent a causal effect of ads at all (Lewis and Reiley 2010).

Many studies of advertising, both in industry and in academia, suffer from similar problems of establishing correlation rather than causality. In general, advertisers do not systematically vary their levels of advertising to measure the effects on sales. Advertisers often change their levels of advertising over time, running discrete “campaigns” during different calendar periods, but this variation does not identify causal effects because other relevant variables also change. For example, if a retailer both advertises more and sells more

during December than in other months, we do not know how much of the increased sales to attribute to the advertising, and how much to increased holiday demand. Similarly, observational academic studies providing regressions of sales data on advertising data often produce positive relationships, but they might be due to reverse causality, when advertisers set advertising budgets as a fraction of sales (Berndt 1991; Dorfman and Steiner 1954; Schmalensee 1972).

A time-honored scientific method for establishing causal effects is to run a controlled experiment, with treatment and control groups. In the eTrade search-advertising example, we would ideally like to experiment on the population who searches for “online brokerage,” showing the eTrade ad only to a treatment group, but not to a control group. The idea of replacing observational studies with field experiments has recently gained importance in economics and the social sciences (Levitt and List 2008).

Measuring the effects of brand advertising presents particular challenges, because the effects are likely to be quite diffuse. A given advertising campaign may be a very small factor in a given customer’s purchase decision, and the results may not be as immediate as in direct-response advertising. A large sample size should enable detection of an advertising signal from the sales noise, but historically it has been nearly impossible to measure both sales and advertising for each individual.

Past attempts to use experiments to measure the effects of brand advertising on sales have found mixed results because of the low signal-to-noise ratio. In the 1960s, Ackoff and Emshoff designed controlled experiments for Anheuser-Busch to evaluate the effects of increased or decreased quantities of advertising on beer purchases, with six geographical areas per treatment. In the 1970s, a series of 19 advertising experiments for Campbell’s Soup measured aggregate sales separately for up to 31 different geographical areas (Eastlack and Rao 1989). In the 1980s and 1990s, IRI’s BehaviorScan technology enabled advertisers to deliver differing quantities of television ads and measure the effect on sales for 3,000 households at a time (Abraham and Lodish 1990; Lodish, et al. 1995a,b; Hu, Lodish, and Krieger 2007). These early experiments generally lacked sufficiently large samples to reach statistically conclusive results.<sup>†</sup>

---

<sup>†</sup> Ackoff and Emshoff (1975) made no attempt to quantify their statistical uncertainty, but Allaire (1975) pointed out that the interesting effects they observed were not statistically significant.

In the present study, we take advantage of recent technological developments in online advertising and customer database management in order to overcome this signal-to-noise problem. Our controlled experiment tracks the purchases of over one million individuals in order to measure the effectiveness of a retailer's nationwide display-advertising campaign on Yahoo! In the companion paper, Lewis and Reiley (2011), we reported results from this experiment. In this paper, we only pay particular attention to looking for heterogeneous treatment effects by age and gender. We find much larger effects for older than for younger customers, demonstrating that in this campaign, the younger customers were the Wanamaker's "wasted half." This case study demonstrates the value of experiments for improving the targeting of online advertising.

## **II. Experimental Design**

We executed a field experiment using a sample of a nationwide retailer's existing customers, identified via matching either postal or email addresses between the retailer's customer database and Yahoo!'s user database. This match yielded a sample of 1,577,256 individuals for whom we could track weekly purchases, both online and in the retailer's physical stores. The retailer believes that they successfully attribute at least 90% of all transactions to the correct customer record. They do this by a variety of techniques, including keeping track of a customer's credit-card number used at the cash register, or asking the customer for their phone number at checkout. To the extent that they do not correctly attribute every single purchase to the correct customer, this will bias our estimates of the effects of advertising towards zero.

Of the matched users, we randomly assigned 81% to a treatment group who were exposed to an advertising campaign for this retailer on Yahoo! The remaining 19% of users were assigned to the control group, who viewed no ads in this advertising campaign, which advertised a clothing line sold by the retailer. The creative content consisted of photographs of models wearing the clothing, with very little text other than the prominently featured name of the store.

The advertising campaign delivered 32 million advertisements to the treatment group during 14 days of October 2007, reaching approximately 800,000 unique users. These rectangular, graphical advertisements appeared on various Yahoo! websites such as mail.yahoo.com, groups.yahoo.com, and maps.yahoo.com. Fig. 1 shows a similar ad on

Yahoo! Shopping. The large, 300-by-250-pixel Capital One graphic ad is an example of a typical “run-of-network” ad of the size and shape used by the retailer in our experiment. These ads are distributed across all pages on the Yahoo network, typically located on the right, near the top of the page. The ads appeared only to treatment-group users when logged in to Yahoo!; none of the retailer’s ads were shown to users who were not logged in. No other users saw this advertising campaign’s ads, whether logged in or not, so we are confident that control-group is valid for identifying customer behavior in the absence of the ads.

**Fig. 1. Example of a typical graphical advertisement on Yahoo!**



These represent the only display ads shown by this retailer on Yahoo! during this time period. However, Yahoo! ads represent a small fraction of the retailer's overall advertising budget, which include other media such as newspaper and direct mail. Yahoo! advertising turns out to explain a very small fraction of the variance in weekly sales, but because of the randomization, the treatment is uncorrelated with any other influences on shopping behavior.

Following the experiment, Yahoo! and the retailer sent data to a third party, who matched the retail sales data to the Yahoo! ad-view data. The third party then anonymized the data to protect the privacy of customers. At the retailer’s request we disguised sales amounts by multiplying by an undisclosed number between 0.1 and 10. All financial numbers

involving sales revenues and cost of advertising will be reported in units of “Retail Dollars” (denoted by R\$) rather than actual US dollars.


### III. Data

Demographic statistics available to us indicate a valid randomization between treatment and control groups. Chi-squared tests generate no significant differences in gender ( $p=0.234$ ), age ( $p=0.790$ ), or U.S. state of residence ( $p=0.232$ ). Nor are there significant differences in Yahoo! usage statistics, including the total number of pages viewed, properties visited, and categories of keywords searched.

Fig. 2 displays the distribution users in the experiment by age and gender, as self-reported by users to Yahoo!‡ The age distribution for this retailer is roughly representative of the general working population, with most customers falling in the range of 25 to 60 years of age. Approximately 93% of customers are at least 25 years of age, 53% of users are at least 40, and 6% of users are at least 65. Both genders are well represented in every age group, with women composing 59.6% of the experimental population.

**Fig. 2. Distribution of customers by age and gender.**

The age population for this retailer is roughly representative of the general working population, with most customers falling in the range of 25 to 60 years of age. Both genders are well represented in every age group, with women composing 59.6% of the experimental population.



‡ We do not know for sure how accurate the self-reported age and gender data are. We know some users lie about their demographics while filling out online forms, but the histogram 2 does not raise any obvious red flags about deceptive reporting. In any case, to the extent that age data are inaccurately reported, that would only make it harder for us to find differences among differently aged consumers.

Treatment-group members saw ads from this campaign on 7.0% of the Yahoo! pages they viewed during the campaign period. The probability of seeing this retailer’s ad on a given page view depended on several factors, including user characteristics as well as the type of page they visited on Yahoo! On the other 93% of pages visited during the same time period, treatment-group members saw a variety of ads from other advertisers.

The number of ads viewed by each Yahoo! user in this campaign is quite skewed. The majority of users in the treatment group viewed fewer than 100 ads in the two-week campaign. By contrast, 1.0% of the treatment group viewed more than 500 ads, with a maximum number of 6050 ad views for one person in two weeks. Such extreme numbers suggest the possibility that a few of our users were automated browsing programs, or “bots,” despite the fact that these user accounts were known to match real people who have previously shopped at the retailer. Indeed, Yahoo! engages in fraud-prevention activities that identified some impressions as not being valid deliveries to humans and did not charge the retailer for those.<sup>§</sup>

The distribution of customer purchases includes a large number of zeroes: only 4.8% of the customers in our data engage in a transaction during the two weeks of the campaign. The standard deviation of sales across individuals, R\$19.49, is much larger than the mean of \$1.88. Though there are some very large outliers, over 90% of positive purchase amounts lie between –R\$100 and +R\$200, where negative purchases represent returns to the retailer. Total purchases occur 14% online versus 86% offline (in-store).

## IV. Results

The treatment group purchased R\$1.89 per person during the campaign versus R\$1.84 per person for the control group, indicating a positive effect of the ads of approximately R\$0.05 per person in the treatment group. Since only 64% of the treatment group actually viewed ads, this represents R\$0.08 per person treated. This treatment effect is only marginally statistically significant ( $p=0.092$ , one-sided)<sup>\*\*</sup>; below we shall see that the

---

<sup>§</sup> We were not able to exclude these potentially invalid observations from our dataset, but they appear to represent a very small part of the sample. To the extent that we include treatment-group members who could not actually have perceived the ads, we will have a slight bias towards measuring zero effect of the advertising.

<sup>\*\*</sup> While a 10% one-sided test is a much lower standard for statistical significance than is usual in the economics literature, we note that this is a more stringent standard than the threshold used for success in the BehaviorScan experiments of Lodish *et al.* (1995a), who used a standard of 20% significance, one-sided. Because of the high variance in sales, we have relatively imprecise estimates of the effects of advertising even with sample sizes in

statistical significance is much greater for older customers. Multiplying by the total number of customers treated, our best estimate of the aggregate, contemporaneous increase in sales due to advertising is R\$83,000, compared with a cost of R\$25,000. Assuming a profit margin of 40% (over the retail price of goods sold) for this class of retailer and ignoring any long-term effects of the ads on shopping behavior, our point estimate of the ad campaign's return on investment (ROI) is well over 30% (though this estimate is quite noisy).

Next, we decompose the treatment effect by age (Fig. 3, top panel). This plots the difference between two locally linear regressions, each using an Epanechnikov kernel with a bandwidth of four years. We find substantial heterogeneity: individuals between ages 20 and 40 experienced little or no effect from the advertising, while individuals aged 50 to 80 experience a sizeable positive effect on sales. Baseline rates of offline and online sales (Fig. 3, bottom panel, Epanechnikov bandwidth of two years) indicate that older customers spend no more than younger customers. The effect of advertising on sales increases with age. During the 14-day advertising campaign, treated individuals over 40 years old increased their purchases by an average of R\$0.15 per person relative to the control group. Customers aged 65+ responded the most, with an average increase in sales of R\$0.37.

Presented for comparison are baseline purchase levels, averaged across treatment and control. Offline purchases are nearly five times as large as online purchases. Teenagers outspent all other age groups, but they are a very small fraction of the customer population. Customers aged 40-65 purchase more than customers 20-40, with a slight decline in purchase amounts after age 65.

A simple model would assume that advertising's effects scale proportionally with sales volume. For example, if middle-aged customers purchase 50% more than the elderly, then the impact of advertising on the middle-aged customers should experience an effect that is 50% larger. However, as a percentage of total purchases, the treatment-effect difference between older and younger customers, shown in Fig. 4, is even more pronounced than in dollar amounts. The treatment effect is 35% of total purchases for customers aged 70 to 75, but only about 5% for customers aged 40 to 65.

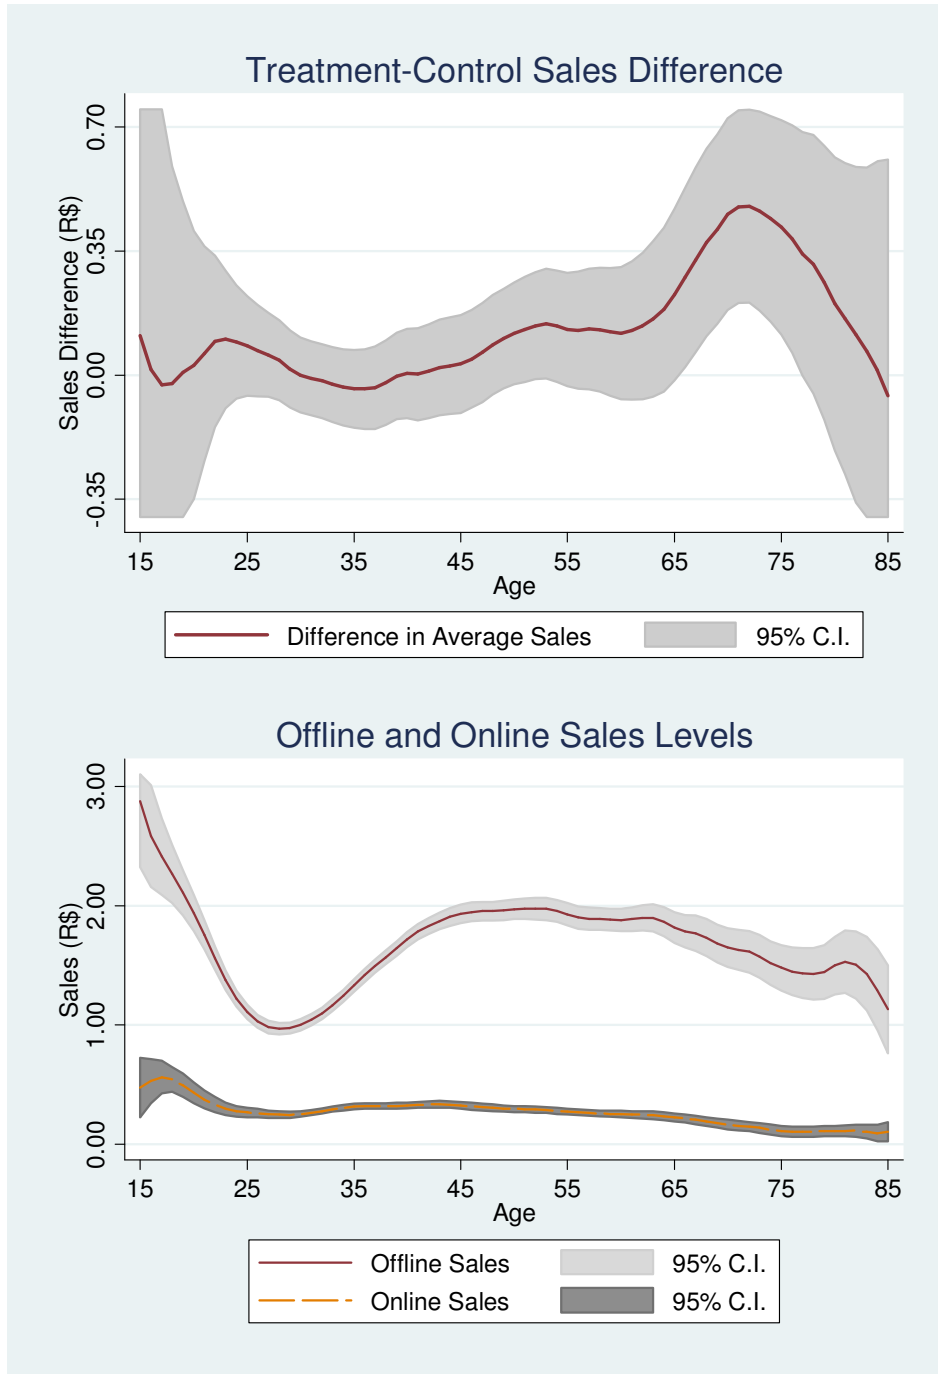
---

the hundreds of thousands. For more detailed power calculations concerning this experiment, see Lewis and Reiley (2011).

This figure uses the two panels of Fig. 3 to compute treatment effects as a percentage of sales, rather than an absolute dollar difference, with pointwise error bands computed by the Delta Method. The 95% confidence interval gives an upper bound on the treatment difference: no more than 10% of average sales for individuals under age 65. By contrast, individuals aged 70 to 75 have an estimated effect of approximately 35%, three times the upper bound for younger individuals.

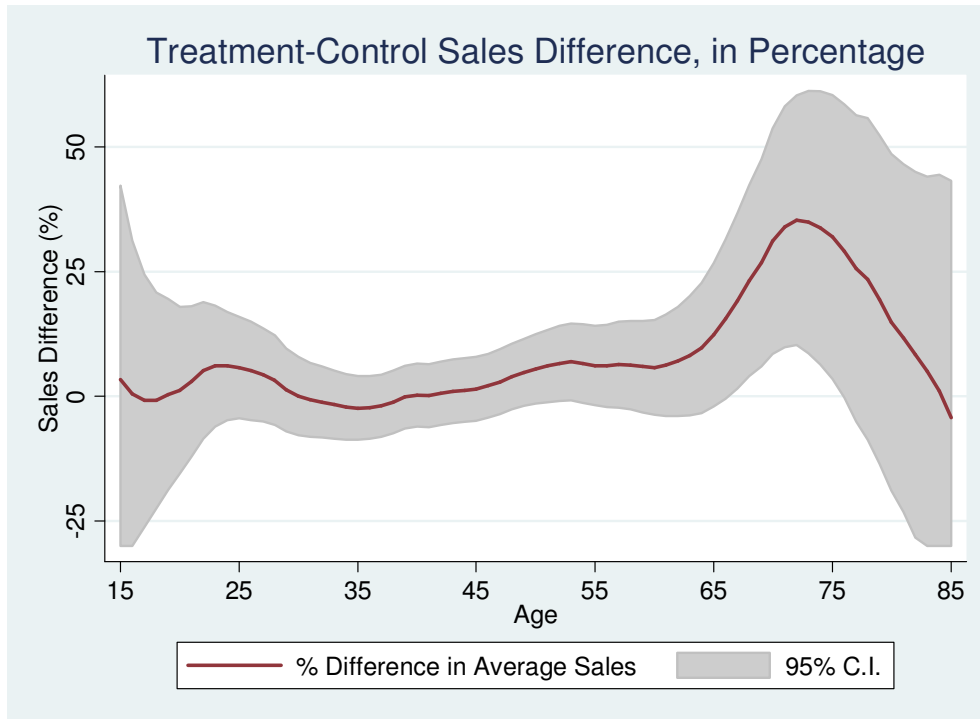
**Fig. 3. Advertising's effects as a function of customer age.**

Estimates are displayed together with 95% confidence-interval bands. *Top panel:* The effect of advertising on sales increases with age. During the 14-day advertising campaign, treated individuals over 40 years old increased their purchases by an average of R\$0.15 per person relative to the control group. Customers aged 65+ responded the most, with an average increase in sales of R\$0.37. *Bottom panel:* Presented for comparison are baseline purchase levels, averaged across treatment and control. Offline purchases are nearly five times as large as online purchases. Teenagers outspent all other age groups, but they are a very small fraction of the customer population. Customers aged 40-65 purchase more than customers 20-40, with a slight decline in purchase amounts after age 65.



**Fig. 4. Advertising's effects as a function of customer age, in percentage terms.**

This figure combines the data in the panels of the previous figure to denominate treatment effects as a percentage of sales, rather than an absolute dollar difference. The 95% confidence interval gives an upper bound on the treatment difference: no more than 10% of average sales for individuals under age 65. By contrast, individuals aged 70-75 have an estimated effect of approximately 35%, three times the upper bound for younger individuals.




We consider two trivial reasons why the advertising might influence the purchases of older customers more than those of younger customers. Do the older customers view more advertisements? The data shows the opposite: conditional on seeing any of the retailer's ads during the 14-day campaign, customers aged 25 to 35 saw nearly 50 ads on average, while customers aged 50+ saw fewer than 35 ads on average (Fig. 5, top panel). Are older customers more likely to see any of the retailer's ads? The answer is again no (Fig. 5, middle panel). The probability of viewing any ads at all also decreases with age: the 25 to 40 age group were about 5% more likely to see ads than those over 40.

In addition to reacting more in terms of purchases, older customers also turn out to be more likely to react to ads by clicking on them (Fig. 5, lower panel). Conditional on seeing at least one ad, the probability of clicking at least one ad increases monotonically until age 45, and is approximately constant for customers over 45 years old. More than 8% of customers

over 45 clicked on ads they saw, compared with only 4% of 20-year-olds. Over the entire sample, 7.2% of ad viewers chose to click.

**Fig. 5. The propensity to view ads declines with age, but the propensity to click increases with age.** Younger users in the dataset browse Yahoo more heavily than older users, and hence see more ads in this 14-day campaign. *Top panel:* Conditional on viewing any ads at all, the number of ads viewed in this campaign decreases from more than 45 to less than 30 ads per person as the age increases from 30 to 75 years old. *Middle panel:* Similarly, the probability of viewing any ads at all also decreases with age: the 25-40 age group were about 5% more likely to see ads than those over 40. *Bottom panel:* However, for those who did see at least one of the retailer’s ads, older users had a much higher probability of clicking on them: over 9% for 70-year-olds, versus 4% for 25-year-olds. All three graphs employ Epanechnikov kernel smoothing with a bandwidth of one year, and include 95% confidence-interval bands.



We have seen evidence that the effects of advertising vary by age. We next quantify these treatment effects by age group, using two simple categorizations of “older” versus “younger” customers, providing standard errors in parentheses. Customers at least 40 years old have an average increase in purchases of R\$ 0.092 (0.055), six times higher than the estimate of R\$ 0.015 (0.051) for those under 40. Using this split, the treatment effect is statistically significant for older customers ( $p=0.047$ , one-sided), but statistically insignificant for younger customers ( $p=0.388$ , one-sided). Customers at least 65 years old exhibit an average increase due to advertising of R\$ 0.373 (0.127), ten times higher than the estimate of R\$ 0.037 (0.039) for those under 65. Senior citizens’ increased purchases are statistically significantly different both from zero ( $p=0.002$ , one-sided) and from the treatment effect of younger customers ( $p=0.012$ , two-sided), while the estimated effect for younger customers is statistically insignificant ( $p=0.175$ , one-sided).

In addition to age, we further decomposed our estimates by gender. While not statistically significantly different ( $p=0.118$ , two-sided), the point estimates are suggestive of the advertising being much more effective for women with an estimated effect of R\$ -0.018 (0.062) for men and R\$ 0.105 (0.047) for women. The treatment difference between women and men is R\$ 0.082 (0.111) for those at least 40 years of age, and R\$ 0.390 (0.252) for those at least 65. For this particular campaign, women account for most of the total effect of the advertising, with the greatest effect on women around 70 years of age.

We also investigated the effects of the advertising separately for online versus offline sales. Online sales increased by only R\$0.007 (0.013) per person in response to the advertising, while offline sales increased by R\$0.046 (0.035) per person. The effect on online sales is quite imprecisely estimated, and when we break it down by age we do not see any obvious difference between older and younger customers. Only the offline sales show the striking trend of more effect of advertising on older customers. (See the Online Supporting Materials for more detailed graphs of the treatment effects separately for each combination of age, gender, and online/offline sales.) This result seems quite surprising when we compare it to the above result on ad-clicking behavior. Older customers who saw ads in this campaign were much more likely to click them than younger customers, but they did not increase their sales much in the online store where the ad click took them. Rather, the ads appear to have

induced them to do more research online and more actual shopping in brick-and-mortar stores.

To check robustness of these results, we subsequently (November 2007 and January 2008) carried out two additional online display ad campaigns with the retailer, using the same treatment-control assignment as in the original campaign. These two campaigns were each much smaller and shorter than the one just analyzed: 10 million and 17 million impressions, respectively, over 10 days each, compared with 32 million impressions over 14 days in the first campaign. The advertiser chose the size and timing of each campaign. Because these campaigns were both smaller and later than the original campaign, we might expect their effects to be smaller. While qualitatively similar to the results for the first campaign—advertising has positive effects on sales, more so for women and older customers—the effects are indeed less statistically significant. More details can be found in the Online Supporting Materials.

## **V. Conclusion**

Our striking result is that this advertising especially influenced older users. We caution that these results are for one particular campaign for one particular retailer, and we do not know to what extent they will generalize. It is entirely possible that other types of advertising campaigns will produce opposite results, with younger customers responding more than older ones. Nevertheless, we find it scientifically meaningful to be able to use our technology to document at least one case where, in contrast to many marketers' belief in the desirability of advertising to younger customers (NAMC 2009), we see clearly that older users respond more to advertising than younger users.

Why might this be the case? We propose four possible explanations, two from economics and two from psychology. From an economic point of view, it might be the case that retired customers have a lower opportunity cost of time than younger customers, so they are more willing to spend time paying attention to display advertisements. Older customers may also have higher wealth available for discretionary spending.

From a psychological point of view, previous research has shown that repetition of a message tends to induce beliefs more strongly in older individuals than in younger ones (Law, Hawkins, and Craik 1998). As this campaign did involve dozens of repetitions of brand-advertising messages to individual customers, this could be a cause of our age-based

differences. Second, previous research has shown older individuals to process information less quickly than younger individuals, and particularly have lower recall of time-compressed television ads (Stephens 1982). To the extent that older customers respond relatively better than younger users when allowed to process information at their own pace, online display ads may produce better effects among older users.

In addition to raising interesting scientific questions, our large-scale, controlled experiment also has practical implications. We close with a back-of-the-envelope calculation of costs and benefits to the retailer of advertising to the different age segments of the population. Each advertising impression cost the retailer approximately R\$ 0.001 on average. In the youngest half of our sample (i.e., under 40 years of age), the average user saw 45 impressions, so the cost of advertising to those customers was about R\$0.045 per person, by comparison with a revenue increase of only R\$0.015. Thus, advertising to the younger customers looks like an unprofitable investment, at least in the short run, although we caution that our lack of statistical precision does not allow us to rule out profitability. The older half of the sample, however, was a much better investment, viewing 35 advertising impressions per person, for a cost to the advertiser of R\$0.035 per customer. This cost compares quite favorably with a revenue increase of R\$0.092 per customer 40 and older, and R\$0.372 per customer 65 and older, even when we consider that the retailer's marginal profit is only about 50% of its marginal revenue. We thus find an extremely large rate of return to advertising to the older customers in the experiment. These numbers suggest a return on investment of over 30% for customers 40 and over, compared with nearly 1,000% for customers 65 and over.

By these calculations, we demonstrate that large-scale, controlled experiments like this one show great promise for improving the targeting of advertising. Improving targeting might improve not only advertiser profits, but also social welfare, as consumers receive ads that are more specifically relevant to each individual (see Bergemann and Bonetti (2011)). We imagine a future in which online field experiments would be implemented at the beginning of an advertising campaign. When heterogeneous treatment effect emerge, such as older customers being more influenced than younger users, or more infrequent past purchasers being more influenced than more frequent purchasers, the results could influence the targeting of the later stages of the campaign. With the development of software for automating online-advertising experiments, we can hope to collect data on a variety of

advertising campaigns. Such data could allow us to advance the science of advertising in various ways, such as better understanding long-run versus short-run effects, and especially in establishing patterns of which types of advertising are most effective for which types of users. We believe that experimental methods will allow advertisers to eliminate the “wasted half” of advertising theorized by John Wanamaker over a hundred years ago.

## VI. References

- Abraham, M. “The Off-Line Impact of Online Ads.” *Harvard Business Review*, April 2008, p. 28.
- Abraham, M. and L. M. Lodish. “Getting the Most out of Advertising and Promotion.” *Harvard Business Review*, May-June 1990, pp. 50-60.
- Ackoff, R.L., and J. R. Emshoff, “Advertising Research at Anheuser-Busch, Inc. (1963-68),” *Sloan Management Review*, vol. 16, no. 2, Winter 1975, pp. 1-15.
- Ackoff, R.L., and J.R. Emshoff, “A Reply to the Comments of Yvan Allaire,” *Sloan Management Review*, vol. 16, no. 3, Spring 1975, pp. 95-98.
- Allaire, Y. “A Comment on ‘Advertising Research at Anheuser-Busch, Inc. (1963-68),’” *Sloan Management Review*, vol. 16, no. 3, Spring 1975, pp. 91-94.
- Bagwell, K. “The Economic Analysis of Advertising.” *Handbook of Industrial Organization*, vol. 3, Mark Armstrong and Robert Porter, eds. Amsterdam: Elsevier B.V., 2008, pp. 1701-1844.
- Bergemann, D., and Alessandro Bonatti, “Targeting in advertising markets: implications for online versus offline media,” *RAND Journal of Economics*, vol. 42, no. 3, Fall 2011, pp. 417-443.
- Berndt, E. R. *The Practice of Econometrics: Classic and Contemporary*. Reading, Massachusetts: Addison-Wesley, 1991, chp. 8.
- Dorfman, R. and P. O. Steiner. “Optimal Advertising and Product Quality.” *American Economic Review*, vol. 44, no. 5, December 1954, 826-836.
- Eastlack, J. O. and Rao, A. G. “Advertising Experiments at the Campbell Soup Company,” *Marketing Science*, 8(1) (Winter, 1989), pp. 57-71.
- Levitt, S. and J. A. List. “Field Experiments in Economics: The Past, the Present, and the Future.” NBER Working Paper 14356, 2008.
- Lewis, R., and D. Reiley, “Does Retail Advertising Work? Measuring the Effects of Advertising on Sales via a Controlled Experiment on Yahoo!” Working paper, Yahoo! Research, 2011.
- Lodish, L. M., Abraham, M., Kalmenson, S., Livelsberger, J., Lubetkin, B., Richardson, B., and M. E. Stevens. “How T.V. Advertising Works: a Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments.” *Journal of Marketing Research*, vol. XXXII, May 1995a, pp. 125-139.
- Lodish, L. M., Abraham, M., Kalmenson, S., Livelsberger, J., Lubetkin, B., Richardson, B., and M. E. Stevens. “A Summary of Fifty-Five In-Market Experiments of the Long-Term Effect of TV Advertising.” *Marketing Science*, vol. 14, no. 3, part 2, 1995b, pp. G133-140.
- Hu, Y., L. M. Lodish, and A. M. Krieger. “An Analysis of Real World TV Advertising Tests: a 15-Year Update.” *Journal of Advertising Research*, vol. 47, no. 3, September 2007, pp. 341-353.

Law, S., S. A. Hawkins, and F. I. M. Craik. "Repetition-Induced Belief in the Elderly: Rehabilitating Age-Related Memory Deficits." *Journal of Consumer Research*, 1998, 25, 91-107.

Schmalensee, Richard. *The Economics of Advertising*. Amsterdam: North-Holland, 1972.

Stephens, N. "The Effectiveness of Time-Compressed Television Advertisements with Older Adults." *Journal of Advertising*, 1982, 11, 48-55.