

Advertising in Health Insurance Markets*

Bradley T. Shapiro[†]

March 5, 2017

Abstract

The effects of television advertising in the market for health insurance are of distinct interest to both firms and regulators. Regulators are concerned about firms potentially using ads to “cream skim,” or attract an advantageous risk pool, as well as the potential for firms to use misinformation to take advantage of the elderly. Firms are interested in using advertising to acquire potentially highly profitable seniors. Meanwhile, health insurance is a useful setting to study the mechanisms through which advertising could work. Using the discontinuity in advertising exposure created by the borders of television markets, this study estimates the effects of advertising on consumer choice in health insurance. Television advertising has a small effect on brand enrollments, making advertising a relatively expensive means of acquiring customers. Heterogeneous effects point to advertising being more effective in less healthy counties, which runs opposite to the concern of cream skinning. Leveraging the unilateral cessation of advertising by United Healthcare, evidence is provided that the small advertising effect is not explained by a prisoner’s dilemma equilibrium. An analysis of longer-run effects of advertising shows that advertising effects are short-lived, further decreasing the potential of advertising to create long run value to the firm. KEYWORDS: Advertising, Insurance, Health, Medicare Advantage, Adverse Selection

*I thank Neale Mahoney, Sanjog Misra, Brett Hollenbeck, Ken Wilbur, Dan Bartels, Garrett Johnson, Carl Mela, Elisabeth Honka and seminar participants at the 2015 Junior Health Economics Conference, the 2016 Choice Symposium, Duke University, UCSD Rady, the University of Naples, and the Harris School of Public Policy for helpful comments. I am grateful to Mariel Schwartz for excellent research assistance.

[†]Chicago Booth. Email: bradley.shapiro@chicagobooth.edu

1 Introduction

Television advertising by health insurance plans is large and growing, rising from about \$250 million in 2004 to \$500 million in 2012. With the implementation of the Affordable Care Act (ACA) Marketplaces and a more broad-based shift towards health plan choice, television advertising by health insurance plans is expected to continue to grow.

Relative to other markets, health insurance advertising faces increased regulatory scrutiny. Historically, the majority of advertising has been for Medicare Advantage (MA) plans that provide seniors with a private, government-subsidized alternative to Traditional Medicare (TM). Regulators in this market have expressed concern that advertising might be used to “cream skim” lower-cost enrollees, inflating government costs, and they have expressed concern about the potential for misleading advertisements, which might induce seniors into purchasing plans they did not need or were poorly suited to their preferences.¹

There also exist reasons to think that MA advertisements—and health insurance advertising more generally—might have important benefits. Health insurance advertising could, of course, serve the standard informative function, helping seniors choose plans that better reflect their own preferences. In addition, because MA, and many other health plans, are purchased during an open enrollment periods, advertising might have the added benefit of alerting consumers about enrollment deadlines. And because of the well-documented inertia in health plan choice (e.g., [Handel, 2013](#)), advertising might be particularly useful in making consumers aware of other plans, thereby intensifying competition in the market.

This paper takes a first step towards understanding the role of advertising in health insurance markets by estimating the impact of television advertising on MA enrollments, characterizing heterogeneity in these effects and testing mechanisms through which advertising might work. In particular, this study tests whether advertising is especially useful in expanding the category by moving seniors out of TM and into MA and characterizes whether such expansion happens disproportionately by county characteristics, including health status. Additionally, this paper tests whether competitive advertising cancels out in a kind of prisoner’s dilemma equilibrium and whether advertising has long-lived effects.

Television advertising is measured using the AC Nielsen Media Database, which provides spot-level information on television advertising from 2004 to 2012. MA enrollment is measured at the

¹<https://www.cms.gov/Medicare/Health-Plans/ManagedCareMarketing/FinalPartCMarketingGuidelines.html>

contract-year level using administrative data from the Center for Medicare & Medicaid Services (CMS) over 2007 to 2012. These data provide the number of potential MA enrollees for each county, enrollment totals for each plan in each county, and information on premiums, co-insurance, and other characteristics for each plan and county.

Sharp discontinuities in the level of advertising at the borders of geographically based television markets provide exogenous variation in advertising, as in [Shapiro \(2016\)](#), to assess the extent to which advertising increases the demand for an MA plan over either TM or another MA plan as well as observable heterogeneity in advertising effects by health status as well as other factors that are informative to both regulators and managers.

Using this variation, I find a small average lift of brand MA advertising on demand. The point estimate from the preferred specification is statistically significant but small, with an implied cost per conversion (CPC) of about \$1300. The 95% CPC confidence interval ranges from \$664 to \$25,100, which rules out very profitable effects of advertising. Meanwhile, a more naive approach would suggest advertising is a very inexpensive means of converting enrollees and rival advertising leads to a material reduction in demand.

Additionally, I find that the effect of advertising on moving seniors from TM into MA (category expansion) is small in magnitude. The preferred specification suggests that removing all MA advertising would only decrease MA demand by 0.23 percentage points, from 11.95% to 11.72% of eligible seniors. At the right edge of the 95% confidence interval, eliminating MA advertising results in a drop in MA demand of only 0.53 percentage points. A more naive approach would produce an effect more than an order of magnitude larger. Given the small main effects, any heterogeneous effects that alter the risk pool in MA must be small.

Indeed, when examining how ad effects vary with observable characteristics, no statistically meaningful relationship is detected between advertising effectiveness and the average health risk of a county. The point estimates suggest that advertising works slightly better on less healthy counties, which works in the opposite direction of cream skimming. Advertising effectiveness is also correlated with lower income, a higher share of elderly population and a higher share of Asian population.

If advertising works primarily to steal business, firms must keep advertising to maintain the status quo share and avoid competitors stealing their enrollees in a kind of prisoner's dilemma. I address this question by leveraging the unilateral cessation of advertising from 2008-2010 by United Healthcare, one of the largest players in this market, providing a direct test of the consequences of removing

advertising on brand share. The estimates suggest that no material loss of brand share occurred, providing evidence against the prisoner's dilemma hypothesis.

Finally, given the sticky nature of health insurance purchase, advertising might be expected to have longer run effects, as seniors who sign up in one year due to advertising are likely to stick around for subsequent years, making advertising more valuable than it would appear in a static sense. Using a stock conception of advertising and for all assumed rates of ad stock persistence, the advertising effect is small, statistically insignificant and precisely estimated, providing evidence against large long run effects of advertising on demand.

Combined, these results suggest that while firms might be trying to cream skim using advertising, they are not particularly successful at drawing a large number of healthy (or any, for that matter) consumers. As such, concerns over the cream skimming and misinformation may be overblown. Of course, these small effects might be due to current regulatory attention. Advertising might affect consumer choices in potentially undesirable ways if regulatory attention were to be reduced. However, given advertising has a very small effect, additional regulatory scrutiny seems unwarranted. On the firm side, these estimates show that advertising is an expensive means of customer acquisition and provide some guidance on targeting.

While the regulatory implications might be clear, a puzzle remains. If television advertising is ineffective relative to other means of customer acquisition in the short- and the long-run, why are firms spending hundreds of millions of dollars per year on television advertising? That advertising increases from 2004-2012 suggests that firms are not learning over time that advertising is ineffective.² It could be that firms have a high cost measuring their own advertising effectiveness or that there is a difficult-to-overcome principal-agent problem with advertising agencies,³ though it is difficult to say definitively using these data. Both the difficulty in measuring advertising effects for firms (e.g., [Lewis and Rao \(2015\)](#)) and that firms as a consequence could make systematic mistakes in advertising strategy (e.g., [Blake, Nosko and Tadelis \(2015\)](#)) have been documented in the advertising effectiveness literature.

The contribution of this paper is an empirical one, finding a small effect of TV advertising in the market for health insurance for the elderly. Health insurance is an important market in particular, but the results of this study should also inform our priors on the usefulness of advertising for

²In fact, after three years of zero advertising, United Health care re-enters the advertising market in a significant way in 2011.

³<http://www.wsj.com/articles/ad-agencies-earn-rebates-from-media-companies-for-ad-spending-probe-finds-1464888210?mod=djemCMOToday>

selective targeting more broadly as well as the usefulness of advertising in markets for goods with complicated, infrequent choices. While there is some recent research studying advertising targeting in MA markets, this is the first paper that estimates the causal effect of advertising using a natural experiment in the market for health insurance. [Aizawa and Kim \(2015\)](#) shows that firms tend to target advertisements towards healthier consumers while [Mehrotra, Grier and Dudley \(2006\)](#) find that ad content is targeted towards healthy patients, giving some credence to the regulatory concern. [Duggan, Starc and Vabson \(2016\)](#) find that firms advertise more in markets where the government pays them more per enrollee. Through a structural model, [Aizawa and Kim \(2015\)](#) explore how market equilibria are affected if firms can use marketing levers to risk select.⁴

This paper also contributes to literatures on competition in health insurance ([Dafny, 2010](#)) and in the Medicare Advantage market (e.g., [Curto et al., 2015](#); [Cabral, Mahoney and Geruso, 2014](#); [Town and Liu, 2003](#)). In particular, [Ericson \(2014\)](#) finds that default plans are persistent, so firms are more likely to offer new plans rather than lower prices. [Cooper and Trivedi \(2012\)](#) find that firms try to gain advantageous selection by offering plan characteristics such as gym memberships. Given these results, there might be concern that equilibrium outcomes are not allocatively efficient, and in fact, [Abaluck and Gruber \(2011\)](#) find that consumers could save a significant amount of money by switching to the lowest-cost prescription drug plan for them. Despite evidence of cream skimming and social mis-allocation, this paper shows evidence that further regulation of television advertising is unlikely to solve any of these problems.

This paper also adds to a growing literature on advertising effectiveness, some of which documents similarly small advertising effects (e.g., [Blake, Nosko and Tadelis, 2015](#)), while others document significant effects, even in the presence of the aforementioned principle-agent issues in the industry (e.g., [Johnson, Lewis and Reiley, 2015](#); [Shapiro, 2016](#)). On the measurement side, all of these studies have shown that failure to consider the endogeneity of advertising can lead to large biases in estimated advertising effects, usually in the upward direction. In particular, as firms often target areas of historic and recent strength, reverse causality is a large concern in the estimation of advertising effects. Documentation of this problem along with attempts at a solution date back at least to [Lodish et al. \(1995\)](#), which uses split cable experiments to think about advertising effectiveness. Additionally, a recent stream of literature has used either field experiments online (e.g., [Blake, Nosko and Tadelis, 2015](#);

⁴In its empirical application, it finds advertising to be more than twice as effective overall and more effective on the healthy from 2001-2005, which is at odds with this study. This difference could be attributed to differences to research design or differences in sample period, as [Aizawa and Kim \(2015\)](#) use the period from 2001-2005, before sophisticated risk adjustment policies were put into place to try and remove the incentive to cream skim.

Johnson, Lewis and Reiley, 2015; Sahni, 2015a,b; Lewis and Nguyen, 2015), instrumental variables (e.g., Sinkinson and Starc, 2015; Hartmann and Klapper, 2014) or natural experiments (e.g., Shapiro, 2016; Tuchman, 2015; Spenkuch and Toniatti, 2015). This paper will follow the latter strategy and exploit the random nature of TV market borders.

The use of borders as a source of variation is related to a small but growing literature. Spatial strategies have been used to identify the effects of minimum wages (Dube, Lester and Reich, 2010), the effects of right-to-work laws (Holmes, 1998), the effects of schools on home values (e.g., Black, 1999; Bayer, Ferreira and McMillan, 2007) and the response of households to changes in electricity prices (Ito, 2014). While many of these studies exploit state borders, that is unattractive in this setting, as many health-policy related factors vary across states. As such, this study will focus on within-state comparisons across the borders of television markets. Additionally, since some DMAs have few counties, we will only use those border areas that make up less than 35% of the counties in the DMA.

Since the main result of this paper is a very small effect of advertising, this paper might be seen as presenting a null effect. However, it is an important and informative result with regulatory and managerial implications. In particular, the documented advertising effect is precisely estimated. Indeed, the confidence interval shows that even viewed optimistically, advertising is a relatively expensive means to acquire business. As regulators work with limited resources to find interventions that work, knowing which ones will not work or that will only work in a very limited fashion is necessary. Additionally, advertising spending is in the hundreds of millions of dollars per year in health insurance, making the documentation of a negative result important for firms. Finally, failure to report zero results on important questions when the research design is sound could potentially bias our understanding of the world, with false positives that occur by chance being published and true negatives being left out, making any meta-analysis biased in favor of the false positives.

The rest of the paper proceeds as follows. Section 2 describes the markets for advertising and health insurance for the elderly. Section 3 describes the data. Section 4 explains the research design in detail. Section 5 documents the results, and Section 6 provides general discussion and concludes.

2 Background

2.1 Health Insurance for Seniors

Nearly all seniors at or above sixty-five years old in the U.S. receive health insurance coverage under the Medicare program. Historically, most seniors have enrolled in what is now called Traditional

Medicare (TM), which is a public insurance program administered by the Center for Medicare and Medicare Services (CMS). Beneficiaries can go to any provider who is willing to see them. The program is fee-for-services, meaning the providers are paid according to the medical services they provide. In addition to premiums, beneficiaries have to pay deductibles and coinsurance, or purchase supplemental Medigap coverage to cover this cost-sharing.

Medicare Advantage (MA) was established in the early 1980s to provide a private alternative to TM coverage.⁵ MA plans are differentiated from TM in having restricted provider networks, alternative cost-sharing arrangements, and additional benefits, such as vision and dental coverage. MA plans have historically been offered by health maintenance organizations (HMOs). Plans receive a capitation payment from Medicare for each enrolled beneficiary and often charge beneficiaries a supplemental premium. Premiums are determined by the relationship between plan bids to the government and statutory benchmark rates. MA enrollment has skyrocketed in the last 15 years, fueled in part by legislation that has that increased payment to plans and lifted restrictions on entry by non-HMO plans. Since 2000, MA enrollment has risen considerably from nearly 0% to over 30% of Medicare Beneficiaries. See [McGuire, Newhouse and Sinaiko \(2011\)](#) for an in-depth history of the MA program.

There are six large national firms that make up around 65% of the total MA share: United Healthcare, Aetna, Humana, Cigna, Kaiser Permanente, and the Blue Cross & Blue Shield (BCBS) plans. While United and BCBS have strength in many markets across the country, other plans have more geographically concentrated historical strength. In addition to these large national firms, there exist more geographically concentrated local plans in many markets.

For most of Medicare's history, very few enrollees had prescription drug coverage. Some had coverage through their Medicare Advantage plan and some purchased supplemental insurance with drug coverage, but the majority of seniors paid for most prescription drugs out of pocket. The 2004 Medicare Modernization Act changed this with the creation of Medicare Part D. Starting in 2006, seniors with TM could enroll in subsidized, private Part D plans and seniors who enrolled in MA could use their subsidies for MA plans that provided drug coverage.

2.2 Television Advertising

Firms can purchase advertising space on television in two ways. First, there is an upfront market each summer where advertising agencies and firms make deals for the upcoming year of television. Ad-

⁵MA was previously known as Medicare Part C or Medicare+Choice. I use the current naming convention throughout the paper.

vertising purchased in the upfront market cannot be “returned” and typically has minimal flexibility in terms of timing, though there is a secondary market that firms sometimes use to offload unneeded advertising space. There is also a spot market, where firms can purchase advertising closer to the date aired.

Ads may be purchased for local or national television. Local advertisements are only seen by households within a particular designated market area (DMA). A DMA is a collection of counties, typically centered around a major city, and it is defined by the global market research firm, AC Nielsen. The DMAs were first defined to allow for the sale of advertising in a way that was straightforward to the advertisers. The DMA location of a county determines which local television stations that a consumer of cable or satellite dish gets with his or her subscription. The original idea was to place counties into the same DMA with the local television station that most people wanted to watch, which often times was just the station that was easiest to pick up over the air. That is, if a county picks up the Cleveland stations over the air more easily than the Columbus stations, it would be placed in the Cleveland DMA. Existing laws and regulations in most circumstances do not allow satellite or cable operators to provide broadcast signals from outside of the DMA in which they reside.⁶ Even for over the air signals, the FCC moderates the signals to try to keep the signal from each station localized only in its own DMA.⁷ There are 210 DMAs in the United States. National advertisements are, in principle, meant to be seen by everyone in the country tuned into a particular network station. However, local affiliate stations do have leeway to bump national ads in favor of additional programming or local ads, generating some local variation in national ads.

3 Data and Summary Statistics

3.1 Advertising

Advertising data from AC Nielsen’s Media database from 2004-2012 is used in this study. The database tracks television advertising at the spot-time-DMA level for every product which advertises on television. The top 130 out of 210 DMAs are indicated as "full discovery markets" by AC Nielsen, meaning all television advertising occurrences are measured using monitoring devices. In many of the smaller DMAs, only advertising occurrences that match ads in the larger markets are included. This study uses each of these full discovery markets which has a monitoring device on every major network affiliate (ABC, NBC, CBS and FOX), which is 120 DMAs.

⁶<http://www.sbca.com/dish-satellite/dma-tv.htm>

⁷<http://www.fcc.gov/encyclopedia/evolution-cable-television>

In the top 25 DMAs, household impressions are measured from set top viewing information that is recorded in a random subset of households. In DMAs ranked 26-210, advertising impressions are estimated from quarterly diaries filled out by a random subset of households. While impressions are the main advertising measure of interest, there is some concern that the infrequent and self-reported viewing data may be measured with error. In the appendix, all analysis will be repeated using ad occurrences as an alternative measure to see if the results are consistent.⁸ The data also include the total estimated expenditure of the firm on the advertisement, the duration of the advertisement and very coarse age, race and gender demographic breakdowns of the impressions data. The data include the parent company of the product advertised, a description of the product being advertised and a very brief description of the content of the advertising copy.

The largest six firms account for 62% of the total advertising impressions. While United Healthcare and Humana, have some national ads, overall 90% of advertising for MA is local. Descriptive statistics about firm-level advertising and shares are presented in Table 1. There is also considerable variation in ad spending within a year, with ads heavily focused during the open enrollment period, which runs from October 15 through December 7 each year. Figure 2 shows this dynamic.

Pairing these data with population data from the U.S. Census, the total number of Gross Rating Points (GRPs) that each advertisement constituted is computed. A GRP is the typical unit of sale between a firm and a television network for advertising space: the total number of households who watched an ad divided by the population in the DMA. As such, a yearly increase of one GRP can be interpreted as the average person viewing the ad one additional time over the course of that year. Figure 1 shows the evolution of health insurance advertising over the course of our sample as a fraction of total television advertising. In 2004, health insurance advertising makes up about 0.25 percent of a \$100 billion total of television advertising. By 2012, that number has roughly doubled.

3.2 Enrollment

MA enrollments and plan characteristics are measured using data from the Centers for Medicare and Medicaid Services (CMS). Enrollments at the plan-county-month level from 2007-2012 are observed. However, with few exceptions, seniors choose their MA plans once per year during open enrollment. The enrollments decided upon in open enrollment translate into enrollments effective January of the

⁸Additionally, conditional on the fixed effects in the model, ad occurrences predict the measure of impressions precisely and nearly identically in the top 25 DMAs and the DMAs ranked 26-130, suggesting that measurement error is not random, but systematic by DMA. Conditional on the DMA fixed effect, changes in impressions over time appear to be reasonably well measured. Please contact the author if you would like further details.

following year. As such, enrollments are measured at the yearly level in February of each year, and those enrollments will be paired with advertising from the prior year when seniors make their choices about the upcoming year's insurance. Plan characteristics such as premiums at the plan-county-year level are also observed.

Finally, information from the Census on demographics such as population, Medicare eligible population and race are merged at the county level. Data on the Medicare risk scores of each county and statutory benchmark rates, which help to determine the capitation payment rates to MA plans, and plan-level characteristics from CMS are also included. Risk scores are centered at 1 with a standard deviation of 0.1, with higher scores indicating worse health. Combining all of the data, this study uses enrollment and advertising at the county-plan-year-level from 2007-2012, using only data from after the introduction of Medicare's Part D prescription drug benefit.

4 Research Design

4.1 Endogeneity of Advertising

Identifying the effects of advertising can be difficult, both in terms of statistical power and bias induced by various forms of endogeneity. In terms of power, [Lewis and Rao \(2015\)](#) shows that due to small true advertising effects and often large amounts of noise in purchases, it can be very difficult to pin down advertising effects with precision. In terms of bias, advertising is a firm choice, subject to equilibrium forces and firm maximization. These forces may cause advertising to be correlated with sales for reasons other than a treatment effect of advertising. Firms might also use rules of thumb based on targeting past or expected sales rather than perceived treatment effects, leading to potential concerns about reverse causality. Indeed, most plausible confounds would bias the researcher in favor of finding a larger advertising effect where none (or a smaller one) exists. In the case of correlated firm behaviors, this is nicely illustrated in [Lewis, Rao and Reiley \(2011\)](#).

4.2 Identification Strategy

In this study, sharp discontinuities in the level of advertising at the borders of geographically-based television markets provide exogenous variation. This design was first used in [Shapiro \(2016\)](#) to study the effects of television advertising on antidepressant demand, but is also used in [Tuchman \(2015\)](#) to study e-cigarette advertising, as well as in [Spenkuch and Toniatti \(2015\)](#) to study political advertising. Consumers who live on different sides of DMA borders face different levels of advertising, due to

market factors elsewhere in their DMA, but they have similar observable characteristics and choice sets of products. In this way, at the borders, observed advertising is ‘out of equilibrium’ and simulates an experiment.

Capturing this intuition, I estimate the casual effect of advertising on MA enrollment controlling for unobservable geographic characteristics with border-specific brand-time fixed effects. This allows unobservables to be spatially correlated in ways that are consistent with the take-up of MA across the country. To both improve precision and to control for any unobservables that are persistent within counties over time, the panel nature of the data is leveraged using brand-county fixed effects. As regulatory regimes may differ across state lines, I focus on DMA borders that are within a state. The identifying assumption is that there are no unobserved differences in trends across these borders which are simultaneously correlated with changes in advertising and the MA share. At the brand level, I only conduct these comparisons where there exists a matched pair across the border. That is, if United is present in a county on one side of a DMA border, it is only included in the analysis if it is present in at least one county on the opposite side of the DMA border.

The top 120 DMAs contain 210 such within-state borders, 164 of which where the border areas make up no more than 35% of the total DMA population. At the brand level, this creates 772 brand level border experiments where cross border matches can be made, 573 of which are in border areas that make up no more than 35% of the DMA population. In the main analysis, attention will be restricted to these borders, but sensitivity analysis around the 35% cutoff will be conducted in the appendix. Each of these brand-border pairs will be considered a separate experiment, with the magnitude of the treatment determined by the advertising in each DMA at a given time, measured in GRPs. Only the counties bordering each other while being in the same state will serve as controls for each other to partial out any local effects that may be increasing or decreasing MA enrollments for both sides of the border, including any national advertising. The level of an observation is a county-year in the category level analysis and a brand-county-year in the brand level analysis. In each ‘experiment,’ one such set of (brand-)counties will be compared with an adjacent set of (brand-)counties directly across the DMA border.

For a leading example, figure 3 shows the Cleveland and Columbus DMAs in the state of Ohio. The border experiment considered is outlined in bold. I will be comparing how outcomes on one the Cleveland side of the border change when when the Cleveland DMA receives a change in advertising GRPs relative to the Columbus DMA. Figure 4 shows this dynamic graphically. In the first panel,

the time series of advertising for Humana for both the Cleveland and Columbus DMAs are pictured and in the second panel, the time series of Humana demand is shown for a county on either side of the border. As the Columbus side gets a larger bump in GRPs than the Cleveland side of the border, the county on the Columbus side of the DMA border sees an increase in demand. This is the type of variation that the border approach exploits.

4.3 Econometric Model

To model the main effects of advertising on brand demand, let i index counties, b index borders, j index brand and t index time. Let s_{bjt} indicate the percentage of Medicare beneficiaries with MA coverage through brand j and let GRP indicate level of advertising. The effect of an increase in advertising GRPs on MA brand j demand is estimated with regressions of the form

$$s_{bjt} = \gamma_1 GRP_{jit}^{own} + \gamma_2 GRP_{jit}^{rival} + \alpha_{bjt} + \alpha_{ij} + X_{bjt} \alpha_X + \epsilon_{bjt}, \quad (1)$$

where α_{bjt} are border-brand-time fixed effects, α_{ij} are brand-county fixed effects, and X_{bjt} is a vector of brand and county control variables, including demographic, competitive environment and plan characteristics. In this case, the coefficients of interest, γ_1 and γ_2 , capture the casual effects of an increase in own and rival advertisements on brand share, respectively. Since variation is at the brand-DMA level and includes repeated measurements over time for each brand-county, standard errors are clustered by brand-DMA in all brand-level analysis. This base model is easily augmented to study category-level effects, to include interactions with advertising effects or to use a different measure of either advertising exposure or demand.

For this approach to be useful in identifying advertising effects, two conditions must hold. First, there must be sufficient variation in advertising across the borders in the data. If all advertising variation were at the national level over time and local stations rarely used their discretion to displace national ads, the border-specific time fixed effects would sweep away all variation in advertising. Figure 5 shows that there is significant advertising across the borders in both total-market and brand-level GRPs, as well as rival GRPs by plotting histograms of total, brand and rival GRPs net of the fixed effects included in the border approach. In this way, the variation in the histogram reflects the identifying variation in the border approach.

Second, the placement of the borders must be quasi-random with respect to preferences for health

insurance. As policies related to health insurance and health care often vary at the state level, DMA borders that coincide with state borders are excluded, as many policies that may affect preferences change at state borders. The location of DMA borders were determined historically by AC Nielsen and have been changed rarely over time. To provide evidence that observables are not predicted by advertising across the borders, table 2 shows the estimated coefficient from a regression with observable county and plan characteristics as the dependent variable and DMA-level MA advertising as the independent variable, controlling for brand-border-year fixed effects. As such, this regression tests whether there are discontinuous changes in observables as the DMA border is crossed. None of the county demographics or plan level variables, including population, share of population over 65, average income, race or average premium are predicted by advertising at the border. The only observable that is predicted by advertising differences across the border is in the number of brands competing in the market. An additional GRP is associated with 0.0023 more competitors across the border. As the average number of competitors is 2.52, this is a 0.09% difference. In terms of unobservables, the DMA boundaries were set by AC Nielsen long ago based on which local news station a family was more likely to want to watch. This was largely determined by which television stations were reachable over the air by households at the time. As very few households watch TV over the air in the present, these boundaries are as good as random. The maintained identifying assumption throughout is that trends in any unobservables that correlate with demand must be parallel across the DMA borders.

While great care is taken to control for as many potential confounds as possible, it is possible that some contamination could still occur if firm strategies are mechanically correlated with one another. In particular, direct mail advertising or online advertising are unobserved. The maintained assumption is that firms do not alter their direct mail or online ad strategies discontinuously at DMA borders as television advertising is altered. Doing so would likely be prohibitively costly technically given the potential return. Additionally, using a different identification strategy, [Aizawa and Kim \(2015\)](#) show that direct mail advertising has no significant effect on MA shares and any omitted variables bias requires that the omitted variable be driving the outcome as well as being correlated with the endogenous variable. Finally, as pointed out by [Lewis, Rao and Reiley \(2011\)](#), firms tend to correlate their various advertising strategies positively, so if firms were engaged in such highly detailed targeting of other strategies such that they changed discontinuously at borders, these would tend to bias advertising effects upward.

4.4 Features and Limitations

Perhaps the largest feature of this approach is that the observed advertising levels at the border are out of equilibrium. That is, variation is driven by the equilibrium in other markets. At the border of the Cleveland, OH DMA, viewers see Humana ads that were driven by metro Cleveland, despite the fact that at the border these residents can be quite different. If ads were micro-targeted to the county level, these consumers would likely see different ads. Similarly, on the Columbus, OH side of the DMA border, the advertising is largely driven by metro Columbus, which is away from the border, again giving rise to rather different advertising at the Columbus border than if ads could be micro-targeted. If metro Columbus and metro Cleveland are sufficiently different from each other, these very similar consumers right on the border will get very different ads, even though their equilibrium micro-targeted ads would have been very similar. This gives a reasonable amount of variation away from what would be the equilibrium in the micro-targeted world while using the fact that these consumers across the border from one another are very similar to control for unobservable factors driving demand.

By conducting the analysis in this way, it is possible to see advertising levels that are likely to be both well above and well below what would be optimal if firms micro-targeted each county individually, which makes the estimated treatment effect approximate an average treatment effect across the advertising response curve for this population. In an experimental approach where the researcher injects some noise into a pre-existing equilibrium or targeting rule, the estimated effect will only be local to levels of advertising that are near that equilibrium or targeting rule. If the firm is already optimally allocating advertising spending, the incremental effect from a small amount of noise being added to the equilibrium might be hard to pin down and smaller than the average effect. In an instrumental variables (IV) approach, the estimated effect will only be local to those people who are affected by the instrument (i.e. the 'compliers')- a group which is not always straightforward to characterize for policy or managerial purposes. Additionally, because the border approach does not require the use of instruments, it is not subject to potential weak instrument bias as well as some of the less desirable finite sample properties of IV estimators.

The border approach falls victim to the familiar local average treatment effect issues that are also common to experiments and IVs. In this case, the estimated effect will be local to those consumers who live in border areas. That is the 'compliers' will be the set of people that live within the border sample, which is a group that can be characterized and compared with the population at large in a

straightforward way to assess whether sample selection is a problem. Table 3 shows how consumers in the border sample are systematically different from consumers outside of the border sample. Most notably, the average population in a border county is considerably smaller than in a county outside of the border sample. The border sample also has a lower percentage Asian and Hispanic population, slightly less competition and a higher percentage Medicare eligible population. If anything, intuition suggests the larger percentage of Medicare eligible in the population (though a small difference) would lead to a higher estimated advertising effectiveness for a product that only Medicare eligible consumers purchase. While there are these systematic differences, there is considerable overlap in the support of the distributions of these characteristics between the border sample counties and other counties. As such, the extent to which these characteristics are important to advertising effects may be estimated directly by interacting them with advertising. Additionally, specifications will be run that are identical using the full sample and the border sample to see the likely effect of sample selection in this context.

4.5 Computing Cost of Conversion

Using the above approach produces the average treatment effect of advertising GRPs on demand. A more managerially relevant metric that also helps to frame the economic significance of the estimates is the cost of a conversion (CPC), which is the acquisition cost of a customer using advertising. To convert the estimates from average lift to CPC, I use the cost of observed advertising from the data, computed at the average cost of a health insurance GRP in the sample. This is inflated to account for the fact that most brands do not operate in every county in a DMA and as such, 'waste' some percentage of eyeballs.

I note here that the computed CPC from the estimated advertising effects represents a lower bound on the true average CPC for at least two reasons. First, the CPC is a convex function of both the effect size and cost of advertising. Since CPC is computed at the average value of cost per GRP and the ad effect represents an average, Jensen's inequality says that the average CPC is greater than the CPC computed at the average values of effect size and cost per GRP. Next, from talking with managers in the industry, nearly all sign-ups for MA plans are consummated by brokers, who extract fees as high \$450 per enrollment with continuing fees if the enrollee renews in subsequent years. Even if the advertising caused the conversion, it's likely that some of the converted enrollees used brokers that were paid commissions, increasing the cost of the average conversion.

CPC estimates are provided for both ends of the 95% confidence interval of average lift. Where

the confidence interval in average lift is less than zero, the upper end of the CPC confidence interval is infinity. External information on the cost of acquiring enrollees through price reductions as well as information on average static profitability provide context on whether or not a particular CPC is relatively high or low.⁹

4.6 Mechanisms

Four hypotheses about the mechanisms of advertising effectiveness are tested. First, advertising could primarily drive category enrollments. Next, advertising could work differentially on different populations. Advertising might also be a prisoner's dilemma whereby firms only advertise to 'cancel out' the advertising of their competitors. Finally, advertising might have longer-run effects due to the sticky nature of insurance enrollments as well as potential brand building effects.

Category demand and heterogeneity are straightforward extensions of the original model. To study category level demand, the same model is applied, slightly augmented to include only category demand and only total category GRPs. To study heterogeneity, GRPs are interacted with observable characteristics of interest. All of these variables are normalized to have a mean of zero and standard deviation of one for ease of interpretation on the main effect of advertising and are assessed at both the category- and brand-levels. In particular, regulators are interested in the category-level interaction between GRPs and health status. If advertising disproportionately works on healthy consumers, it might induce adverse selection for TM plans, increasing the costs to the government. For the purposes of targeting and potentially increasing profit, firms should be interested in any heterogeneity in the treatment effect that they may legally be able to target. As scientists, heterogeneous treatment effects may help us to understand something about how exactly advertising works.

While the effects of rival advertising on brand demand are shown in the main effects, a more direct test of the prisoner's dilemma hypothesis is presented, leveraging the fact that United Healthcare unilaterally stopped advertising during 2008 and 2009. In 2006, a *Wall Street Journal* article detailed a scandal involving backdated stock options.¹⁰ This led to an SEC investigation over the course of many years and the eventual resignation of CEO William W. McGuire. In the midst of the commotion and regulatory attention, United spent almost nothing on health insurance advertising between 2008 and

⁹The Government Accountability Office (GAO) computed that average revenues per enrollee were \$9,893 with an average profit margin of 4.5% in 2011. Of note is that advertising cost is included in non-medical expenses that have already been subtracted out of revenues to obtain the 4.5% number. While advertising costs are small relative to other non-medical expenses such as administrative expenses at about \$59 per enrollee, they are added back in to the 4.5% to compute incremental static profit. This amounts to an average one-year profit of \$504 per incremental enrollee. More details are available at <http://www.gao.gov/products/GAO-14-148>

¹⁰<http://www.wsj.com/articles/SB114265075068802118>

2009, as can be seen in Figure 6. If the prisoner's dilemma hypothesis holds, then rival advertising should deteriorate United's brand share over the years in which it does not advertise. To control for direct effects of the scandal common to similar counties, the cessation is interacted with the border strategy to provide the causal effect of rival advertising on United brand share.

To address the long-run effects of advertising, a goodwill stock conception of advertising is used instead of yearly GRP levels. Given the sticky nature of health insurance purchase, advertising might be expected to have longer run effects, as enrollees who sign up in one year due to advertising are likely to stick around for subsequent years, making advertising more valuable than it would appear in a static sense. In fact, [Curto et al. \(2015\)](#) find that roughly 77% of MA enrollees stick with the same plan each year. Advertising could also have effects on brand equity that would only manifest themselves over a number of years. These two facts combined suggest that the effects of past advertising stock is important to assessing profitability. I measure the long-run effect of advertising using a stock conception assuming different rates of ad stock persistence. I use data from the earlier years of the data (2006-20010) to initialize the ad stock and estimate the model on the last two years of data (2011-2012) to maximize the number of periods I can use to build the ad stock while still allowing for the use of fixed effects. If the advertising marginal patient were equally as sticky as the average patient and there are zero brand equity effects, assuming an ad stock persistence parameter of 0.77 should produce roughly the same ad effect as in the static model. Brand equity effects should further increase the ad stock effect. Advertising could fail to show long-run effects if either advertising marginal enrollees were less sticky than average enrollees or if advertising were an essential part of utility for the advertising marginal and needs to be maintained each year for the enrollee to stick.

5 Results

5.1 Main Effects

Table 4 presents the estimation of the effect of brand and rival GRPs on the brand demand using equation 1. To provide economic intuition for the size of the advertising effects, cost per conversion (CPC) estimates and confidence intervals are presented below the advertising effects.

Column (1) presents the naive regression, where all counties are included, and no fixed effects or controls are used. It suggests a positive and significant effect of advertising on brand demand, with an increase of one GRP associated with an increase in brand demand 0.0971 percentage points, implying an average CPC of \$123.72. Effects of this size would make advertising a relatively inexpensive way

to acquire customers, if true. Additionally, column (1) shows a small but significant negative impact of rival advertising on demand. Column (2) adds observable control variables for demographics, premiums, and year fixed effects as well as the number of the competitors in the market, average risk scores of the counties and statutory county benchmarks rates. The effect size is almost unchanged and implies a CPC of \$121.72 while the magnitude of the rival effect increases to -0.027, suggesting significant damage to the firm from rival advertising. If firms systematically target advertising to markets that are strong for reasons other than advertising, the first two columns would provide spurious positive estimates of the own ad effect as well as spuriously negative effects from rival ads. In column (3), county fixed effects are added to control for time-invariant factors affecting MA strength in a particular county. In this case, identifying variation comes from deviations from the average brand advertising in a market. Controlling for these factors makes the ad effect shrink considerably. The average lift from one GRP is about 0.01 percentage points and implies a CPC of \$1133.44 and the effect of rival ads goes to zero. Meanwhile, the estimated effect is not especially precise, as the 95% confidence interval of the CPC in this specification is [$484.19, \infty$]. While \$484.19 might be worth spending to acquire a customer, zero average effect (and infinite CPC) of advertising cannot be ruled out. Column (4) runs exactly the same specification as column (3) but limits the sample to that which will be included in the border approach, to highlight the role of sample selection in this context. In this specification, counties in the border sample are included, but a general brand-year fixed effect is included as in column (3) rather than controlling for local demand shocks using the brand-border-year fixed effects. The estimated effect of advertising goes up, suggesting that the sample selection effect to the border would tend to over-state the true average effect in the population. However, if firms target markets of not only historic strength but also of recent strength or target unobserved concurrent local demand shocks, columns (3) and (4) will also over-estimate the advertising effect. In column (5), the border approach is employed to control for all remaining endogeneity issues. The effect of advertising on brand demand is now more precisely estimated and is statistically significant, though small in magnitude. The confidence interval around the ad effect is roughly two thirds as wide as the fixed effects estimation in column (3). However, despite finding a statistically significant ad effect, the implied CPC is difficult to bound above at a reasonable rate, as the CPC approaches infinity as the ad effect approaches zero. The 95% CPC confidence interval ranges from \$664.15 to \$25,100 with an average CPC of \$1294. The estimated rival effect is positive, suggesting a positive spillover, but is neither statistically nor economically significant. While a simple correlation would lead the re-

searcher to conclude that advertising is a highly profitable way to shift the elderly into a branded MA plan, a more careful analysis using plausibly exogenous variation shows much more limited advertising effectiveness. Additionally, while the simple correlation would lead the researcher to conclude significant harm from rival advertising, the more careful approach shows no material effect.

To provide a graphical illustration of these results, a bin scatter, using 100 bins, of brand MA percent and brand GRP is presented in figure 7 as well as one of brand MA percent and rival GRP in figure 8. These amounts are residualized by a brand-county fixed effect and a brand-border-year fixed effect, so they reflect variation across TV market borders, as in column (5).

5.1.1 Discussion

The estimated effect of advertising, while detectable in this setting, is small and, as discussed, corresponds with high CPC. It implies an advertising elasticity of about 0.027. The estimates in [Aizawa and Kim \(2015\)](#) suggest that a 1% increase in advertising leads to a 0.066% increase in enrollments, more than twice as large. It is also smaller than estimates of advertising effectiveness in other context. For example [Shapiro \(2016\)](#) finds advertising elasticities of about 0.04 in the context of antidepressant demand and [Sethuraman, Tellis and Briesch \(2011\)](#) find an average ad elasticity of about 0.12 in a meta-analysis.

To provide further context on the magnitude CPC, an estimate of how much an incremental customer would be worth to the firm is useful. An incremental customer provides static profits and potentially continuing lifetime value as customers often stick to plans for many years after initial sign-up. From the GAO report cited previously, the average static profit for a MA enrollee in 2011 was about \$504. The cost of conversion is bounded below by \$664.15 in the preferred specification, which rules out short run return on investment using advertising. However, firms might value customers above \$504 if they also provide future profits.

Firms may also reveal how much they are willing to pay for customers through other means of conversion. One way to convert additional customers would be to lower premiums. The cost of gaining a marginal customer is that all inframarginal customers will also receive the premium reduction. The best evidence on price elasticity in MA plans comes from [Curto et al. \(2015\)](#), which documents MA customers have an average premium semi-elasticity of 0.012. This implies that a reduction in monthly premium of \$1 increases brand enrollments by 1.2%. This implies an average acquisition cost using price reductions of about \$1000, which is less expensive than advertising, on

average.

Overall, the effect of advertising on demand is small relative to previous literature on advertising effects, and even though it is statistically significant, prohibitively large CPC cannot be ruled out. Even at the lower end of the CPC confidence interval, advertising is not profitable, at least in the short run. As explained in the previous section, the average CPC computed here is a lower bound. Given that and the small estimated static effects, firms need advertising marginal customers to stay for multiple years, at minimum for advertising to be profitable.

5.2 Mechanisms

5.2.1 Category Expansion

From a regulatory perspective, the concern about cream skimming comes at the category level. If MA plans can draw the healthy out of TM, it leaves the government with unhealthy consumers. Since the capitation rates received by MA plans depend on the average cost of providing health care to a TM consumer, such cream skimming would increase the per customer medical cost for TM and increase the amount the government pays to MA plans, even holding total expenditures on health care provision in the county fixed. Such an effect can only exist insofar as advertising moves seniors out of TM and into MA and the magnitude of the selection effect is limited to the size of the category expansion effect.

Table 5 shows the effects of category total advertising GRPs on the percent of Medicare eligible seniors who choose any MA plan over TM. The columns correspond with the specifications in the previous section and the CPC estimates are interpreted as the cost of moving a senior out of TM and into any MA plan. This is lower than the firm relevant CPC given that a senior converted from TM to MA might choose a rival plan. In the column (1), neither fixed effects nor controls are used in a completely naive approach. This specification suggests that an increase of one GRP corresponds with an increase in category demand of 0.075 percentage points and is statistically significant. The estimate corresponds with a CPC of \$160.20. Column (2) adds in controls for average county-level premiums, demographics, competition, risk scores and benchmark rates. The estimate shrinks by a factor of three, to 0.0257 and is no longer statistically significant. It corresponds with a CPC of \$467.51 and infinite CPC is now part of the confidence interval. Column (3) adds in county fixed effects and the effect size decreases by an order of magnitude to 0.0022, implying a CPC of \$5399.25. Column (4) uses the same specification as column (3), but only on the sample of counties at the borders of DMAs. The estimate

increases slightly to 0.0073, as in the brand-level analysis, but remains statistically insignificant and corresponds with an average CPC of \$1635.11. Column (5) is the preferred specification and employs the border approach with county fixed-effects and border-year fixed effects. The estimate shrinks to 0.0042 with a corresponding average CPC of \$2874.09.

Using advertising to convert TM customers into MA plans has limited potential to create regulatory harm. The estimate of in column (5) implies that the elimination of all MA advertising would move the average percentage of seniors who choose MA over TM from 11.95% to 11.72%. At the right edge of the confidence interval, the elimination of all advertising would move the average MA share from 11.95% to 11.42%. If regulatory bodies relied on the naive approach in column (1), effects would appear more than an order of magnitude larger and indeed be cause for concern. However, the more careful approach suggests limited potential for category expansion to drive significant shifts in the distribution of seniors picking MA over TM. In addition, with an average CPC of \$2874.09 and no guarantee that the customer will pick the advertised brand, using advertising to bring customers into the category is a managerially costly strategy.

5.2.2 Selection Through Ad Effect Heterogeneity

While the overall effect of advertising on category expansion is limited, it could be that the regulatory concern over using advertising to cream skim is still well founded directionally. Such cream skimming using advertising might be possible if less healthy consumers were less able to focus on advertising or if advertising copy were particularly attractive to healthy rather than unhealthy people. Meanwhile, heterogeneity in the treatment effect is very important to firms in addition to regulators. If advertising is more useful on more profitable patients than average, the estimated CPC could be more easily justifiable. Additionally, the heterogeneity should inform targeting decisions about which markets should receive more advertising attention from firms.

To assess heterogeneity in the treatment effect, advertising GRPs are interacted with variables about the average health risk in the county, the number of brands competing in a county, average premiums and demographic variables. To assess ad timing, advertising GRPs are also interacted with whether or not the ad took place during the open enrollment period. To assess whether observed advertising is potentially in the flat part of the advertising response curve, a quadratic term of GRPs is also included. These interactions are all included using the border approach with the border sample, (brand-)county fixed effects, and (brand-)border-year fixed effects.

The results are presented in table 6. Column (1) shows heterogeneity in the treatment effect of total GRP on category demand. First, advertising is not disproportionately moving lower-health-risk patients into MA, which was the main basis for regulatory concern. In fact, the point estimate suggests that advertising works better on less healthy counties, though the effect is small and not significant. Advertising in the open enrollment period appears to expand the category less than advertising during the rest of the year in moving customers from TM to MA. It could be that direct brand comparisons rather than general feelings about health insurance brands are more intense during open enrollment, but given that there are many interactions tested, it may also be a false positive.

Column (2) shows heterogeneity in brand demand response to advertising. Again, the interaction with average risk status in a county is positive, suggesting that advertising works better in less healthy counties, which works against the cream skimming story, though it is not statistically significant. In terms of managerially actionable targeting, advertising works better in counties with a higher share of elderly population, lower average income and a higher percentage Asian population. Given that poor health tends to be correlated with low income, the income effect may reflect further health effects that are not captured in the risk scores. It is also worth noting that this analysis is at the county level, so I cannot definitively rule out that advertising works only on the unusually healthy individuals within the unhealthy counties. However, if advertising works better on the healthy either due to the advertising copy or the ability of the healthy to engage with the ads, that should be reflected in advertising working better in healthier counties. It is difficult to tell a story that advertising on television can be specifically targeted to the healthy individuals in unhealthy counties while not affecting healthy individuals in healthy counties.

The interaction effects also inform the degree to which sample selection to the borders might affect the estimated average treatment effect of advertising. Recall that the border counties had a lower percentage Asian population but also lower income and higher elderly population. On net, these sample selection effects combined with the interactions roughly offset, suggesting that the average overall effect of advertising should be roughly equal in the border sample and the interior of the DMA.

5.2.3 Prisoner's Dilemma

The main results suggest that rival advertising has a small positive effect on brand demand, but that could possibly be true only if the own brand also advertises in equilibrium. To provide a direct test

of the prisoner's dilemma theory of advertising, United Healthcare's two-year cessation of advertising is leveraged. It is notable that United Healthcare does reasonably well in terms of brand share over the time in which it is not advertising, as can be seen in figure 9. Table 7 presents the results of the regression analysis using United brand share as the dependent variable. Column (1) presents the correlational results when not using the border strategy, controls or fixed effects. While the point estimate implies a negative effect of a rival advertising GRP on United's brand share of -0.0778, it is small and implies a CPC, or a cost of stealing an enrollee from United, of \$1465.43, which is rather high. However, the bad press could have directly lowered shares. United also could have lowered premiums to compensate for its lack of advertising. To address these concerns, column (2) includes control variables, including United's premium, the average premiums of other brands, demographics, and year fixed effects. The effect size shrinks and becomes insignificant, with average effect of -0.024 and corresponding CPC of \$4750.38. The sign of the estimate flips once county fixed effects are included in column (3), making the average CPC infinite. Column (4), which moves to the border sample but maintains the specification, leaves the effect unchanged. Moving to using the border approach in column (5), the point estimate is still positive, but now closer to zero at 0.0123. This leaves the average CPC at infinity and the left edge of the CPC confidence interval at \$1895.73 to steal an enrollee from United while it is not advertising. These results provide direct evidence against the prisoner's dilemma explanation theory of advertising and show that other firms in the market did not gain brand share from United's cessation of advertising.

5.2.4 Long Run Effects

While all of the above analysis provides evidence about the usefulness of advertising in the year it airs, it could still be the case that advertising works over a longer-term horizon. The effects of advertising might be expected to be long-lasting, as [Curto et al. \(2015\)](#) find that roughly 77% of MA enrollees stick with the same plan each year. If advertising were responsible for converting an enrollee, it is possible that the enrollee would stick longer than one year, even if advertising were to cease. To systematically think about the question of long-run effects of advertising, the border approach is interacted with a stock conception of advertising. In this case, advertising is assumed to persist at a geometric rate, δ , per year for four candidate values of δ : 0.5, 0.6, 0.7 and 0.8. If enrollees acquired by advertising stuck to plans as frequently as the average enrollee, using a $\delta=0.77$ should leave the short run advertising effect previously estimated unchanged. If, alternatively, advertising did not have a long-lived effect,

making the independent variable a function of past advertising would be similar to introducing left hand side measurement error and would attenuate the estimate towards zero. Advertising effects could fail to be long-lived if either advertising marginal enrollee were less sticky than average or if the advertising had to be maintained in order to keep the enrollee, as would be the case if advertising were a direct utility shifter.

Table 8 provides the results of the effects of advertising stock on brand demand. Columns (1), (2), (3) and (4) show the analysis with persistence parameters of 0.5, 0.6, 0.7 and 0.8, respectively. In all four columns, the estimated advertising effect is smaller than the estimated short-run effect from before. Additionally, the four columns cannot be statistically distinguished from one another, nor from zero. These estimates provide evidence against long-lived effects of advertising, either from particularly sticky advertising marginal enrollees or from some other long-run effect, such as the building of brand equity. It could be that advertising marginal enrollees are less sticky than the average customers, or it could be that for the advertising marginal enrollees, advertising goes into the yearly utility function, so the firm must keep advertising up in order to keep them.

5.2.5 Discussion

The exploration of mechanisms provides evidence against both the long-run effectiveness of advertising in this setting and the prisoner's dilemma hypothesis of advertising effectiveness. It also suggests that the effects on the category overall are small and directionally go against the regulator concerns over cream skimming. From a targeting perspective, there is some heterogeneity in the advertising effect on brand demand that might be actionable to managers. Advertising works better in counties that have a larger share of elderly population as well as counties that have a larger percent Asian population. However, some of the heterogeneous effects provide murkier implications. For example, advertising works better in lower income counties, which tend to be less healthy and might be less profitable.

5.3 Overall Results

All of the results together potentially raise a puzzle. Why exactly does broadly targeted advertising persist in equilibrium in this market? One potential explanation is that even when using numerous control variables, the advertising effect appears to be economically quite large. It is not until carefully employing quasi-exogenous variation that the ad effect shrinks considerably. Managers could have difficulty in finding good variation to measure the causal effects and mistakenly use spurious corre-

lations as their indicator of how well advertising works. Previous research, such as [Blake, Nosko and Tadelis \(2015\)](#), has found instances where managers struggled to make correct decisions with regards to advertising due to poor measurement of causal effects.

Additionally, managers in health insurance could give little attention to television advertising by outsourcing decision making to ad agencies. The market for health care is a three trillion dollar market, and advertising spend is only around five hundred million. This implies a relatively small advertising to sales ratio compared with other industries. Figure 10 shows that relative to the size of the industry, health advertising is reasonably small. Advertising agencies might well have very different incentives from the firms contracting their services, which could lead to poorly executed advertising strategy.

Even if all incentives were perfectly aligned and firms used quasi-exogenous variation that improves both accuracy and power, as it does here, it is possible that for some values in the confidence interval, advertising could provide positive long run ROI. Since I cannot see incremental firm profits, it is impossible to say for sure using these data. However it is clear that advertising is a relatively expensive way to acquire customers.

6 Conclusion

In this paper, the effect and mechanisms of television advertising for health insurance are explored. Policy makers are very concerned with advertising in the MA market, as evidenced by their numerous publications governing the proper conduct of advertising by firms. They are mostly concerned with firms trying to cream skim a favorable risk pool through advertising, as well as attempting to mislead seniors. Conversely, there is a concern that regulation may cause inefficiency in the marketplace, as consumers may be exposed to a lower-than-optimal level of information about each plan.

Leveraging the discrete borders of television markets, the effect of advertising is estimated to be small. The 95% confidence intervals suggest that even at the most optimistic, advertising effects imply large customer acquisition costs. Meanwhile, a more naive approach, even including a large number of control variables would lead the researcher to conclude that advertising is more than ten times as effective as the more careful approach. The naive correlation suggests that advertising is a very inexpensive way to acquire customers relative to static profits and relative to acquiring customers through price reductions.

Further, this study provides evidence on mechanisms. Perhaps easing regulatory concerns, advertising has a small overall effect on the share of seniors who choose MA over TM and a careful look

at how that effect varies by health status reveals no systematic relationship between ad effect and the average health risk of a county. Advertising is found to work better on some types of counties than others, particularly those that have lower income as well as those with high elderly and Asian shares of the population. United Healthcare's cessation of advertising reveals no significant negative effect from rival advertising when own advertising is stopped for an extended period of time, providing evidence against a prisoner's dilemma theory of advertising. Finally, evidence is provided that the even though this is a particularly sticky market, advertising effects are not particularly persistent. Combined these results suggest that advertising is a relatively costly way to acquire additional seniors, and regulating advertising would provide limited change in market outcomes.

Puzzles remain for future research. First, the reason for the small advertising effect is unclear. Whether the lack of advertising efficacy is due to the current set of regulations or poor advertising copy is an interesting question with important firm implications. If the regulations make it unlikely for advertising to work better, then managers should scale back advertising on television considerably and focus on other areas of marketing strategy. If poor ad copy can explain the small effects, then managers might be able to improve the effectiveness of television ads. Second, given that advertising is estimated to have limited effectiveness in both the short and long run, it remains a puzzle that firms are spending hundreds of millions of dollars on this form of promotion. Since estimating advertising effects could be costly, it might simply be worth taking the gamble that advertising might work given the industry's large revenues. These results further highlight the results of [Lewis and Rao \(2015\)](#) describing the difficulty in accurately measuring ad effects given that the CPC confidence intervals can be quite large and advertising might be worthwhile at the lower end of these confidence intervals. Conversely, agency issues could cause firms extra difficulty in setting optimal advertising budgets.

The estimates from this study imply that concerns about cream skimming and deception due to advertising may be overblown. However, with advertising having little effect on enrollments, the concerns about the deadweight loss due to regulating advertising are also mitigated. Finally, estimates in this study suggest that firms are potentially making systematic mistakes in their advertising strategy and highlight the difficulty in making accurate assessments.

References

- Abaluck, Jason, and Jonathan Gruber.** 2011. "Heterogeneity in choice inconsistencies among the elderly: evidence from prescription drug plan choice." *The American Economic Review*, 101(3): 377.
- Aizawa, Naoki, and You Suk Kim.** 2015. "Advertising competition and risk selection in health insurance markets: Evidence from Medicare Advantage." *University of Minnesota working paper*.

- Bayer, Patrick, Fernando Ferreira, and Robert McMillan.** 2007. "A unified framework for measuring preferences for schools and neighborhoods." *Journal of Political Economy*, 115(4): 588–638.
- Black, Sandra E.** 1999. "Do better schools matter? Parental valuation of elementary education." *Quarterly Journal of Economics*, 114(2): 577–599.
- Blake, Thomas, Chris Nosko, and Steven Tadelis.** 2015. "Consumer heterogeneity and paid search effectiveness: a large-scale field experiment." *Econometrica*, 83(1): 155–174.
- Cabral, Marika, Neale Mahoney, and Mike Geruso.** 2014. "Does privatized health insurance benefit patients or producers? Evidence from Medicare Advantage." *NBER Working Paper*, w20470.
- Cooper, Alicia L, and Amal N Trivedi.** 2012. "Fitness memberships and favorable selection in Medicare Advantage plans." *New England Journal of Medicine*, 366(2): 150–157.
- Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya.** 2015. "Can managed competition work? The US Medicare Advantage program." *Mimeo*.
- Dafny, Leemore S.** 2010. "Are health insurance markets competitive?" *The American Economic Review*, 100(4): 1399–1431.
- Dube, Arindrajit, T William Lester, and Michael Reich.** 2010. "Minimum wage effects across state borders: estimates using contiguous counties." *The Review of Economics and Statistics*, 92(4): 945–964.
- Duggan, Mark, Amanda Starc, and Boris Vabson.** 2016. "Who benefits when the government pays more? Pass-through in the Medicare Advantage program." *Journal of Public Economics*, 141: 50–67.
- Ericson, Keith M Marzilli.** 2014. "When Consumers Do Not Make an Active Decision: Dynamic Default Rules and Their Equilibrium Effects." *NBER Working Paper*, w20127.
- Handel, Benjamin R.** 2013. "Adverse selection and inertia in health insurance markets: when nudging hurts." *The American Economic Review*, 103(7): 2643–2682.
- Hartmann, Wesley R, and Daniel Klapper.** 2014. "Super bowl ads." *Available at SSRN: <http://ssrn.com/abstract=2385058>*.
- Holmes, Thomas J.** 1998. "The effect of state policies on the location of manufacturing: evidence from state borders." *Journal of Political Economy*, 106(4): 667–705.
- Ito, Koichiro.** 2014. "Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing." *The American Economic Review*, 104(2): 537–563.
- Johnson, Garrett A, Randall A Lewis, and David Reiley.** 2015. "When Less is More: Data and Power in Advertising Experiments." *Available at SSRN 2683621*.
- Lewis, Randall A, and Justin M Rao.** 2015. "The unfavorable economics of measuring the returns to advertising." *The Quarterly Journal of Economics*, 130(4): 1941–1973.
- Lewis, Randall A, Justin M Rao, and David H Reiley.** 2011. "Here, there, and everywhere: correlated online behaviors can lead to overestimates of the effects of advertising." 157–166, ACM.
- Lewis, Randall, and Dan Nguyen.** 2015. "Display advertising competitive spillovers to consumer search." *Quantitative Marketing and Economics*, 13(2): 93–115.
- Lodish, Leonard M, Magid Abraham, Stuart Kalmenson, Jeanne Livelsberger, Beth Lubetkin, Bruce Richardson, and Mary Ellen Stevens.** 1995. "How TV advertising works: a meta-analysis of 389 real world split cable TV advertising experiments." *Journal of Marketing Research*, 32(2): 125–139.

- McGuire, Thomas G., Joseph P. Newhouse, and Anna D. Sinaiko.** 2011. "An economic history of Medicare Part C." *Milbank Quarterly*, 89(2): 289–332.
- Mehrotra, Ateev, Sonya Grier, and R Adams Dudley.** 2006. "The relationship between health plan advertising and market incentives: evidence of risk-selective behavior." *Health Affairs*, 25(3): 759–765.
- Sahni, Navdeep S.** 2015a. "Advertising spillovers: field-experiment evidence and implications for returns from advertising." *Journal of Marketing Research*, 53(4): 459–478.
- Sahni, Navdeep S.** 2015b. "Effect of temporal spacing between advertising exposures: evidence from online field experiments." *Quantitative Marketing and Economics*, 13(3): 203–247.
- Sethuraman, Raj, Gerard J Tellis, and Richard A Briesch.** 2011. "How well does advertising work? Generalizations from meta-analysis of brand advertising elasticities." *Journal of Marketing Research*, 48(3): 457–471.
- Shapiro, Bradley.** 2016. "Positive spillovers and free riding in advertising of prescription pharmaceuticals: the case of antidepressants." *Journal of Political Economy*, forthcoming.
- Sinkinson, Michael, and Amanda Starc.** 2015. "Ask your doctor? Direct-to-consumer advertising of pharmaceuticals." *NBER Working Paper*, w21045.
- Spenkuch, Jörg L, and David Toniatti.** 2015. "Political advertising and election outcomes." *Available at SSRN 2613987*.
- Town, Robert, and Su Liu.** 2003. "The Welfare Impact of Medicare HMOs." *RAND Journal of Economics*, 34(4): 719–736.
- Tuchman, Anna E.** 2015. "Advertising and Demand for Addictive Goods: The Effects of E-Cigarette Advertising." *Mimeo*.

Table 1: Advertising and Shares by Brand per Year

Brand	Average GRP	% Local	Brand Share (%)	Number of Counties
Aetna	4.14	99.82	11.54	171.7
BCBS	23.23	99.99	30.06	962.9
Cigna	0.041	99.26	8.777	296.8
Humana	10.78	81.39	45.11	1989
Kaiser	20.24	99.98	30.69	97.61
United	5.00	59.47	28.13	1024
Other Insurers	21.22	100	44.08	1770
Total	53.67	90.20	100	2412

Table 2: "Balance" Test

	Est.	Standard Error	P-Value	Mean
Avg. Premium	-0.0134	0.0218	0.5404	49.25
County Benchmark	0.0587	0.0771	0.44779	757
HCC Score	-0.0026	0.0056	0.6404	94.81
% White	-0.0076	0.0088	0.3868	86.03
% Black	0.0076	0.0088	0.3680	9.738
% Hispanic	-0.0116	0.0071	0.1067	4.914
% Asian	-0.0007	0.0007	0.3584	0.912
Avg Income	-1.5021	4.3982	0.7334	31,800
% Medicare Eligible	0.0026	0.0037	0.4827	16.75
Population	7.5015	56.3943	0.8944	67,200
Brands Present	0.0023	0.0007	0.0024	2.52

Regressions reflect a regression with the listed variable as the dependent variable and total GRP as the independent variable, controlling for the border-year fixed effects from the border strategy. This tests whether there is a discontinuous change in the observable across the border that is predicted by advertising.

Table 3: Selection into the Border Sample

	Difference	Mean	%Difference	P-Value
Avg. Premium	0.9007	48.61	1.853%	0.4497
County Benchmark	-21.43	772	-2.776%	<0.01
HCC Score	0.1127	95.04	0.119%	0.7357
% White	0.4374	85.45	0.512%	0.5228
% Black	0.382	9.6947	3.940%	0.5616
% Hispanic	-3.1858	7.1323	-44.667%	<0.01
% Asian	-0.5278	1.3032	-40.500%	<0.01
Avg Income	-2,760	33500	-8.239%	<0.01
Medicare Eligible	0.7774	16.0474	4.844%	<0.01
Population	-75,200	122,000	-61.639%	<0.01
Brands Present	-0.0857	2.6909	-3.185%	0.016

"Difference" reflects the estimate from a regression with the listed variable as the dependent variable and an indicator for "in the border sample" as the independent variable. This shows how the observable demographics are systematically different for the estimation sample.

Table 4: Brand Level Demand(MA%)

	(1)	(2)	(3)	(4)	(5)
<i>GRP</i>	0.0971*** (0.0228)	0.0987*** (0.0210)	0.0106 (0.0073)	0.0207* (0.0098)	0.0093* (0.0045)
<i>RivalGRP</i>	-0.0163* (0.0075)	-0.0267*** (0.0074)	-0.0012 (0.0030)	-0.0013 (0.0037)	0.0025 (0.0027)
Controls		x	x	x	x
County-rand FEs			x	x	x
Border Sample				x	x
Border Approach					x
CPC	\$123.72	\$121.72	\$1133.44	\$579.95	\$1294
CPC CI	[84.74, 229.10]	[85.90, 208.82]	[484.19, ∞]	[300.98, 7930.12]	[664.16, 25100]
Mean MA %	5.4456	5.4365	5.4959	5.0191	5.0191
R-squared	0.070	0.155	0.895	0.853	0.918
Observations	36764	36489	35919	10651	10651

*** p<0.001, ** p<0.01, * p<0.05

Brand-DMA clustered standard errors in parentheses. Dependent variable is the percentage of Medicare eligible seniors who choose brand a brand j MA plan. Included in controls: average brand premium, indicator for zero average brand premium, number of competitors, share white, share black, share hispanic, share elderly, MA county benchmark rates, and average county Medicare risk scores.

Table 5: Category Level Demand (MA%)

	(1)	(2)	(3)	(4)	(5)
<i>TotalGRP</i>	0.0750*** (0.0195)	0.0257 (0.0145)	0.0022 (0.0063)	0.0073 (0.0074)	0.0042 (0.0027)
Controls		x	x	x	x
County FEs			x	x	x
Border Sample				x	x
Border Approach					x
CPC	\$160.20	\$467.51	\$5399.25	\$1635.11	\$2874.09
CPC CI	[106.10, 326.85]	[222.18, ∞]	[826.43, ∞]	[550.97, ∞]	[1255.86, ∞]
Mean MA %	12.5388	13.5378	13.5535	11.9524	11.9524
R-squared	0.061	0.302	0.963	0.944	0.978
Observations	17136	15535	15511	4951	4951

*** p<0.001, ** p<0.01, * p<0.05

DMA clustered standard errors in parentheses. Dependent variable is the percentage of Medicare eligible seniors who choose any MA plan instead of a TM plan. Included in controls: average category premium, indicator for zero average category premium, number of competitors, share white, share black, share hispanic, share elderly, MA county benchmark rates, and average county Medicare risk scores.

Table 6: Heterogeneity and Targeting

	(1) MA %	(2) Brand MA %
<i>GRP</i>	0.0069 (0.0068)	0.0126 (0.0110)
<i>xGRP</i>	0.00004 (0.00003)	0.00002 (0.00009)
<i>xRisk</i>	0.0009 (0.0035)	0.0058 (0.0039)
<i>xPremium</i>	-0.0002 (0.0026)	-0.0021 (0.0032)
<i>xCompetition</i>	-0.0015 (0.0020)	0.0017 (0.0036)
<i>xOpenEnrollment</i>	-0.0276*** (0.0080)	-0.0145 (0.0112)
<i>xElderly</i>	0.0052 (0.0030)	0.0107* (0.0042)
<i>xIncome</i>	-0.0056 (0.0029)	-0.0115** (0.0044)
<i>x%Hispanic</i>	-0.0044 (0.0076)	-0.0013 (0.0120)
<i>x%Asian</i>	0.0073* (0.0032)	0.0212** (0.0068)
<i>x%Black</i>	0.0021 (0.0041)	0.0048 (0.0038)
R-squared	0.967	0.918
Observations	4591	10651

*** p<0.001, ** p<0.01, * p<0.05

Brand-DMA clustered standard errors in parentheses. All specifications use the border approach with a full set of control variables and county fixed-effects.

Table 7: United Share (%)

	(1)	(2)	(3)	(4)	(5)
<i>RivalGRP</i>	-0.0778* (0.0371)	-0.0240 (0.0270)	0.0548 (0.0378)	0.0555 (0.0606)	0.0123 (0.0369)
Controls		x	x	x	x
County-Brand FEs			x	x	x
Border Sample				x	x
Border Approach					x
CPC	\$1465.43	\$4750.38	\$∞	\$∞	\$∞
CPC CI	[757.74, 22200.00]	[1484.04, ∞]	[5946.12, ∞]	[1803.70, ∞]	[1895.73, ∞]
Mean United %	25.37	25.30	25.60	22.079	22.079
R-squared	0.015	0.215	0.947	0.900	0.907
Observations	1867	1850	1405	288	288

*** p<0.001, ** p<0.01, * p<0.05

Brand-DMA clustered standard errors in parentheses. Only counties with at least two competitors present included (a firm cannot steal or lose brand share when it is fixed at 100%). Included in controls: average brand premium, indicator for zero average brand premium, average category premium, indicator for zero average category premium, number of competitors, share white, share black, share hispanic, share elderly, MA county benchmark rates, and average county Medicare risk scores.


Table 8: Brand Level Demand - Stock (%)

	(1) $\delta=0.5$	(2) $\delta=0.6$	(3) $\delta=0.7$	(4) $\delta=0.8$
<i>GRPStock</i>	-0.0012 (0.0043)	0.00026 (0.0040)	0.0020 (0.0037)	0.0040 (0.0034)
<i>RivalGRPStock</i>	0.0019 (0.0029)	0.0013 (0.0028)	0.0008 (0.0026)	0.0003 (0.0023)
R-squared	0.972	0.972	0.972	0.972
Observations	3717	3717	3717	3717

*** p<0.001, ** p<0.01, * p<0.05


Brand-DMA clustered standard errors in parentheses. All specifications use the border approach with a full set of control variables and county fixed-effects. Included in controls: average brand premium, indicator for zero average brand premium, number of competitors, share white, share black, share hispanic, share elderly, MA county benchmark rates, and average county Medicare risk scores.

Figure 1: Health Insurance Advertising by Year




Note: Figure shows spending on health insurance television advertising as a percentage of total television advertising by year. Total television advertising is approximately \$100 billion. The source is the AC Nielsen's Media Database.

Figure 2: Health Insurance by Calendar Month



Note: Figure shows that significant amounts of advertising is concentrated in the open enrollment period from October 15-December 7. The source is the AC Nielsen's Media Database. Dashed lines mark the month of November each year.

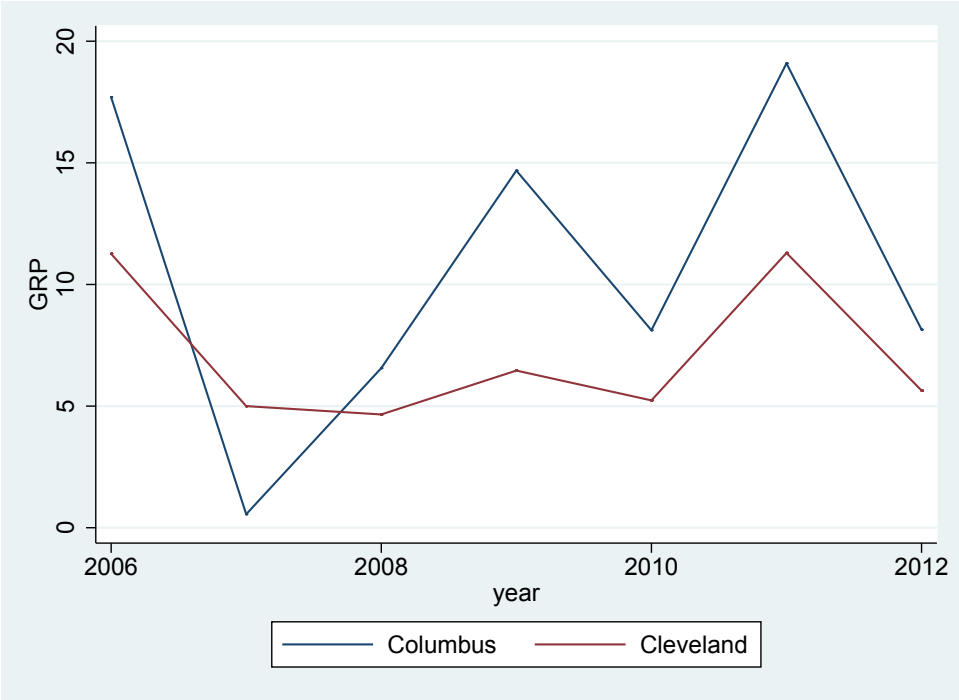
Figure 3: Ohio and its DMAs



Note: Figure shows the Cleveland and Columbus DMAs, highlighting the border region used in identification.

Figure 4: Border Experiment Illustration

((A)) Humana Advertising Acorss the Border



((B)) Humana Demand Across the Border

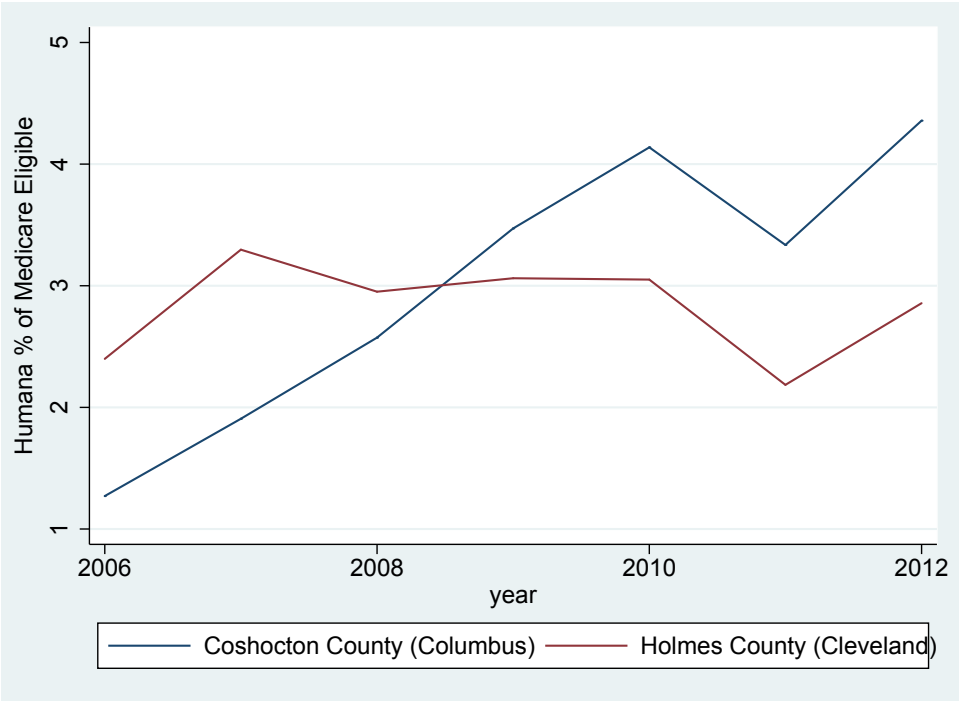
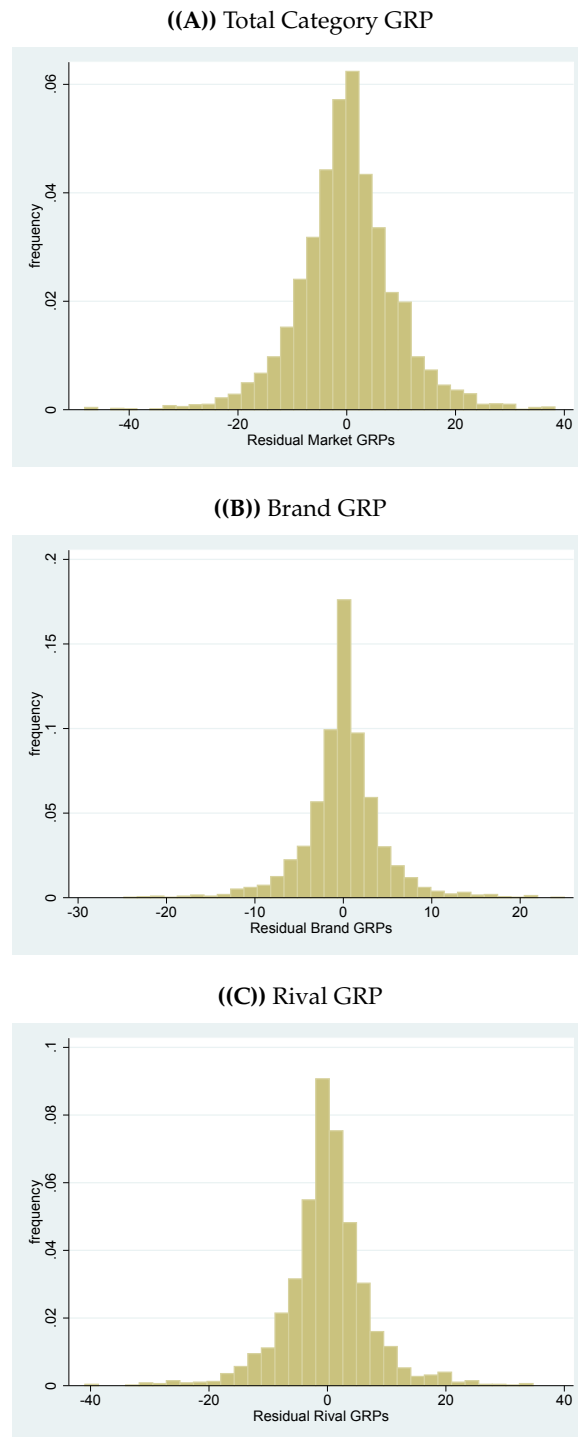


Figure 5: Variation in GRP Changes across DMA Borders



Note: Figure shows that there is significant variation in total, own and rival advertising across DMA borders, conditional on (brand-)county and (brand-)border-year fixed effects.

Figure 6: United Advertising Expenditure Over Time



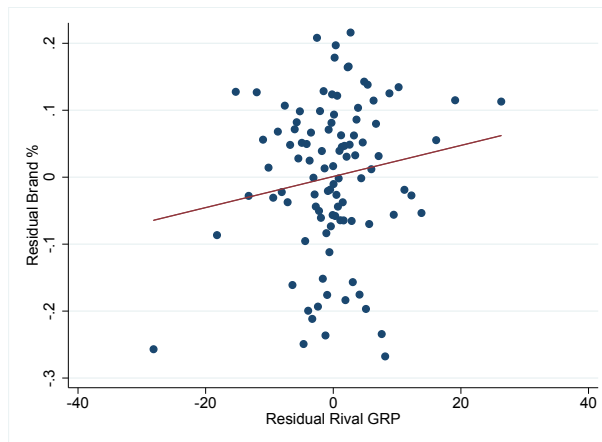


Figure 7: Brand Bin Scatter - Own GRP




Note: This bin scatter has 100 bins and reflects brand % of the Medicare eligible population on the vertical axis and brand GRP on the horizontal axis, net of brand-county and brand-border-year fixed effects.

Figure 8: Brand Bin Scatter - Rival GRP




Note: This bin scatter has 100 bins and reflects brand % of the Medicare eligible population on the vertical axis and rival GRP on the horizontal axis, net of brand-county and brand-border-year fixed effects.

Figure 9: United Shares Around Cessation



Note: This shows the average United brand share over time. Vertical dashed lines represent the years with advertising turned off.

Figure 10: Health Insurance Advertising in Context



Appendix A - Sensitivity to Border Size

In the main analysis, the sample is restricted to those borders which make up no more than 35% of the DMA as a whole. In this section, I test sensitivity to this cutoff in the analysis of the main effects. Extending this sensitivity check to every regression in the paper is straightforward, but is not provided here for the sake of brevity, with the intention of eventually including that in an online appendix. Results are in table A.1. While the point estimates get noisier as the threshold moves down to 10% (as fewer observations are used), both the direction and economic significance of the results remain unchanged. The average CPC remains similar and large across all specifications. While the point estimates are not all statistically significant, the confidence intervals around those point estimates are all very near to those in the main analysis.

Appendix B - Using Log Enrollments Rather than MA% and Brand %

In the main analysis, the dependent variable in the brand demand analysis is the percentage of eligible seniors who choose the advertised brand. In this section, the dependent variable is changed to the natural log of brand enrollments, with the CPC computation adjusted to reflect the change in dependent variable. The analysis of main effects is provided in table B.1. All results are consistent with the main analysis and produce nearly identical CPC confidence intervals. Again, extending this to all of the analysis in the paper is straightforward, but for the sake of brevity is not provided here, with the intention of eventually including that in an online appendix.

Appendix C - Using Occurrences Instead of GRP

As noted in the data section, the main analysis focuses on gross rating points (GRPs) as the relevant measure of advertising. This generates some concern about measurement error. Outside of the top 25 DMAs, impressions are measured using self-reported diaries, which may be measured with considerable error. Meanwhile, advertising occurrences (the number of instances of an ad) are mechanically measured in the top 130 DMAs. In this section, occurrences are used as the relevant measure of advertising, and the analysis of main effects is repeated. For the sake of brevity, the exploration of mechanisms using occurrences as the relevant measure of advertising is not in this paper but available from the author upon request, with the intention of eventually including that in an online appendix.

More noise might be expected when using occurrences as the relevant measure of advertising. That is, each thirty seconds of ad air time is coded as one occurrence, regardless of how many people see it. As such, an ad on a midnight re-run of I Love Lucy is coded exactly the same as an ad during the nightly news. Response to these two ads would be expected to be very different, potentially generating extra noise in the estimation. This is the reason why the main part of this study uses GRPs: it provides a theoretically reasonable way to weight each ad by how many people actually saw it.

Table C.1 present the results. ROI measures are adjusted to reflect the average cost of an occurrence instead of the average cost of a GRP, as was the case in the main analysis. Results are very similar to the main analysis, though with more noise. Without using the border strategy, advertising appears useful in lifting brand demand. When the border strategy is used, the average lift approaches zero. In using this alternative measure of advertising, moving from the simple fixed effects to the border approach provides a more dramatic contrast, highlighting the importance of a careful research design. The preferred specification using the border approach in column (5) shows a very similar CPC and CPC confidence interval as when GRP is used in the main analysis.

Table A.1: Brand Share (%) - Sensitivity to Border Size

	(1) 10%	(2) 20%	(3) 30%	(4) 40%	(5) 50%
<i>GRP</i>	0.0025 (0.0063)	0.0050 (0.0058)	0.0082 (0.0044)	0.0090* (0.0045)	0.0073 (0.0046)
<i>RivalGRP</i>	0.0052 (0.0051)	0.0018 (0.0029)	0.0026 (0.0027)	0.0026 (0.0026)	0.0025 (0.0025)
CPC	\$4760	\$2420	\$1460	\$1340	\$1640
CPC CI	[812, ∞]	[732, ∞]	[714, ∞]	[675, 99800]	[734, ∞]
Mean Brand %	4.21	4.50	4.73	4.98	5.05
R-squared	0.871	0.895	0.900	0.920	0.925
Observations	1194	4544	9374	11452	13353
*** p<0.001, ** p<0.01, * p<0.05					

DMA clustered standard errors in parentheses. All specifications use the border approach with all controls.

Table B.1: Brand Share (Log Enrollments)

	(1)	(2)	(3)	(4)	(5)
<i>GRP</i>	0.0119** (0.0041)	0.0147*** (0.0033)	0.0031** (0.0011)	0.0044** (0.0015)	0.0017* (0.0008)
<i>RivalGRP</i>	0.0007 (0.0018)	-0.0053*** (0.0012)	0.0001 (0.0008)	0.00004 (0.0010)	0.0011 (0.0006)
Controls		x	x	x	x
County-Brand FEs			x	x	x
Border Sample				x	x
Border Approach					x
CPC	\$194.00	\$156.28	\$737.81	\$527.51	\$1380.39
CPC CI	[115.73, 599.06]	[109.00, 276.00]	[430.99, 2561.15]	[315.31, 1613.05]	[699.03, 54600]
Mean Brand LogQ	5.6342	5.6401	5.6652	5.3289	5.3289
R-squared	0.014	0.391	0.932	0.898	0.951
Observations	36004	35741	35154	10406	10406
*** p<0.001, ** p<0.01, * p<0.05					

Brand-DMA clustered standard errors in parentheses. Included in controls: average brand premium, indicator for zero average brand premium, number of competitors, share white, share black, share hispanic, share elderly, MA county benchmark rates, and average county Medicare risk scores.

Table C.1: Brand Share (%) - Using Occurrences

	(1)	(2)	(3)	(4)	(5)
<i>Ads</i>	0.4013*** (0.0905)	0.3866*** (0.0989)	0.1023** (0.0330)	0.1487*** (0.0420)	0.0335 (0.0215)
<i>Rival Ads</i>	-0.0384 (0.0439)	-0.1279** (0.0461)	-0.0151 (0.0174)	-0.0306 (0.0250)	0.0038 (0.0157)
Controls		x	x	x	x
County-Brand FEs			x	x	x
Border Sample				x	x
Border Approach					x
CPC	\$108.30	\$112.42	\$424.66	\$292.25	\$1297.03
CPC CI	[75.10, 194.14]	[74.87, 225.52]	[260.08, 1156.42]	[188.18, 653.81]	[574.15, ∞]
Mean Brand %	5.4456	5.4365	5.4959	5.0191	5.0191
R-squared	0.031	0.124	0.896	0.854	0.918
Observations	36764	36489	35919	10651	10651
*** p<0.001, ** p<0.01, * p<0.05					

Brand-DMA clustered standard errors in parentheses. Included in controls: average brand premium, indicator for zero average brand premium, number of competitors, share white, share black, share hispanic, share elderly, MA county benchmark rates, and average county Medicare risk scores.