

Positive Spillovers and Free Riding in Advertising of Prescription Pharmaceuticals: The Case of Antidepressants

Bradley T. Shapiro*

This Version October 13, 2014

Abstract

Television advertising of prescription drugs is controversial, and it remains illegal in all but two countries. Much of the opposition stems from concerns that advertising directly to consumers may inefficiently distort prescribing patterns toward the advertised product. Despite the controversy surrounding the practice, its effects are not well understood. Exploiting a discontinuity in advertising along the borders of television markets, I estimate that television advertising of prescription antidepressants exhibits significant positive spillovers on rivals' demand. I then construct and estimate a multi-stage demand model that allows advertising to be pure category expansion, pure business stealing or some of each. Estimated parameters indicate that advertising has strong market level demand effects that tend to dominate business stealing effects. Spillovers are both large and persistent. Using these estimates and a simple supply model, I explore the consequences of the positive spillovers on firm advertising choice. I solve for the Markov Perfect Equilibrium and compare that to counterfactuals whereby firms work together to set advertising. In a full industry cooperative advertising campaign, simulations suggest that the co-operative would produce four times as much advertising as the competitive equilibrium, resulting in a 18 percent increase in category size and a 14 percent increase in category profits.

1 Introduction

How does television advertising affect the consumer choice problem? After a consumer watches a commercial, internalizes its message and decides a product is desirable, she must take further action to obtain the product. With groceries, she must go to the supermarket. With many consumer products, a computer with internet will allow the consumer to make the purchase. With prescription drugs, the consumer must go to the physician to obtain a prescription and then to the pharmacy to purchase the drug. With many steps between the advertising incidence and purchase, at some stage of the process, the consumer might well choose a different product from the one advertised. This may be

*University of Chicago Booth School of Business. I thank Nancy Rose, Ernst Berndt, and Stephen Ryan for all of their advice. I would also like to extend my gratitude to Andrew Ching, Sara Fisher Ellison, Glenn Ellison, Günter Hitsch, Greg Leiserson, Sarah Moshary, Jesse Shapiro, Xiao Yu May Wang, and two anonymous referees for their helpful suggestions. This paper also benefited from conversations with Paul Snyderman and Len Tacconi; as well as participants at the MIT Industrial Organization seminar and field lunch. I thank Cindy Halas at IMS Health, Jackie Allen and Leslie Walker at Kantar Media Intelligence, and LuAnn Patrick and Patrick Angelastro at ImpactRx for their help with data resources. This research was supported by the National Institute on Aging, Grant Number T32-AG000186 and the National Science Foundation Graduate Research Fellowship under Grant Number 1122374. All mistakes are my own.

due to difficulty in remembering advertisements, agency problems in obtaining products or simply because advertising convinces a consumer to go to a retailer, computer or physician. In short, an advertisement could affect the choice process without leading the consumer to buy the advertised product.

In this paper, I identify the existence of positive spillovers of television advertising in the market for antidepressants. Given this, I construct and estimate a demand model which allows such spillovers. To quantify the effect of spillovers on firm behavior, I conduct a supply side analysis supposing that the firms are able to jointly decide advertising, and I compare this outcome to the realized advertising outcome in the antidepressant market.

Television advertising of prescription drugs is contentious and has been condemned by many as inefficiently distorting prescriptions to the advertised products. In fact, it is legal in only two countries: New Zealand and the United States. In light of the controversy, it is important to understand the impact of these advertisements. In particular, understanding spillovers is crucial to regulators, firms and econometricians. From a regulatory perspective, the Food and Drug Administration (FDA) regulates the content of advertisements. To the extent that advertising content is made more informative and less brand specific, content regulation could exacerbate spillovers. Firms may lose individual incentives to advertise as spillovers intensify. This could be either good or bad for social welfare depending on whether or not category expansion is a public good or a public bad. However, it is an important consideration for the regulator in either case. From a firm strategy perspective, understanding possible channels for revenue improvement is vital. While cooperation is often difficult to enforce and non-contractible due to antitrust laws, advertising cooperatives are preceded in other industries such as orange juice, milk and beef. Finally, from a technical perspective, failure to model spillovers in advertising can distort estimated parameters, leading to incorrect inferences about supply and demand.

Previous research incorporating advertising into demand analysis has frequently treated advertising of a product as affecting its probability of being in the choice set (Goeree 2008), or has incorporated advertising into a production of goodwill that enters directly into the utility function (Dube et. al. 2005). However, such specifications also typically exclude the possibility of positive spillovers of advertising onto rivals. While this eliminates the complexity of modeling behavior in the presence of possible free riding, such an exclusion may lead the researcher to miss important strategic considerations. When deciding how much to advertise, firms do not internalize the benefit they provide to other firms and have an incentive to free ride on their rivals' advertising efforts. Understanding these considerations is important for marketing decision makers as well as policy makers potentially seeking to regulate advertising.

Prescription drugs in general, and antidepressants in particular, have many characteristics which facilitate positive spillovers in television advertising. First, the FDA regulates what firms can and cannot say in advertisements. While the name of the product is typically prominently displayed throughout the commercial, most of the time in each commercial is spent explaining the ailment, the mechanism of action of the drug and its side effects. When there are several therapeutic products available, those treating the same ailments tend to share common characteristics. A consumer might remember all of the things being said but forget the name of the product. Agency problems further disrupt this link. A consumer must see a doctor to get a prescription. A physician might have different preferences or opinions about which drugs, if any, work best for a given condition or patient. The advertisement may lead a patient to the physician, but the physician is still the ultimate arbiter of whether and what to prescribe.

My strategy for evaluating the extent of positive spillovers in advertising for antidepressants proceeds in three steps. First, I use discrete television market borders to determine the extent to which advertising does affect rival demand,

positively or negatively. Next, I construct and estimate a model of the antidepressant market, allowing advertising to have positive spillovers on demand of horizontally differentiated products, a feature excluded by typical discrete choice specifications. Positive spillovers are allowed, but not imposed by the model. Given the demand estimates and a model of supply, I back out the marginal costs of advertising from a model of dynamic strategic firm interaction. Finally, given estimates of the demand effects and marginal costs of advertising, I quantify the importance of free riding by re-simulating the supply model, assuming that a co-operative sets advertising for the entire industry.

While most models incorporating advertising into demand have not allowed for positive spillovers, there are studies of Direct-to-Consumer (DTC) advertising in pharmaceuticals with varying credibility of identification strategies that have shown some evidence that cross advertising elasticities could be positive, but results have been mixed. In contrast to the demand analyses mentioned above that typically do not allow for cross advertising elasticities to be positive, these studies tend to find patterns that are consistent with spillovers, rather than allowing for spillovers within a demand model. In particular, [Iizuka & Jin (2005), Berndt et. al. (2004), Wosinska (2002)], find very small estimates of advertising effects on market shares conditional on being in the market and conclude that it might be exhibiting positive spillovers, though spillovers are neither modeled nor tested. Wosinska (2005) and Donohue et. al. (2004) find that advertising has positive spillover effects onto drug compliance and duration of treatment. Other studies find that advertising drives consumers to the doctor (Iizuka & Jin 2004) or has class level effects (Rosenthal et. al. 2003, Avery et. al. 2012), but they do not model any product level own or cross-elasticities of advertising. In work that is more closely related to this study, Berndt et. al. (1995, 1997) estimate the effect of marketing on both the size of the market and on brand shares, focusing mostly on physician detail advertising and academic journal advertising since DTC was extremely limited and unbranded at the time, and found some effects at both category and product levels. In studying detailing effects, Ching et. al. (2012) take advantage of the fact that identical molecules are sometimes marketed by different firms under different names in Canada to separate out brand versus category effects. Narayanan et. al. (2004) estimate a two level model using only time series variation for antihistamines and do not find positive spillovers. In experimental work, Kravitz et. al. (2005) find mixed results for patients going to their physicians asking for products they saw on television. In a structural model, Jayawardhana (2013) imposes that television advertising must only affect class level demand and finds significant effects. Many of these studies either only model a category level response or only model a conditional share level response. Those that model both rely solely on time series data. This paper will model the full decision process and use data with both spatial and time series variation. Stremersch et. al. (2013) and Liu & Gupta (2011) also examine the various effects of DTC on aspects of demand. Stremersch et. al. can explain variation across geography using demographic characteristics and both find heterogeneous effects. This study will differ from both of those in that fixed effects will be used to partial out the reasons for persistent differences in DTC across markets and focus on variation just across the borders.

The supply side of advertising in pharmaceuticals has been much less explored. If advertising helps rivals' demand, there might well be an incentive to invest less in advertising. Iizuka (2004) finds that as the number of competitors increases, firms advertise less, leading him to suggest the existence of a free riding problem. Ellison and Ellison (2011) find evidence that pharmaceutical firms decrease advertising just prior to patent expiration in order to make the market smaller and deter generics from entering. The possibility of such strategic deterrence implies the existence of positive spillovers, at least from brand to generic. However, no research that I am aware of uses a supply model to quantify the magnitude of the potential positive spillover effects on advertising expenditure decisions. Ching (2010), Filson (2012), Liu et. al. (2013) all use a Markov Perfect Equilibrium concept to model the supply side of pharmaceutical markets, but they all focus on different aspects of the pharmaceutical industry rather than television advertising.

Outside of the pharmaceutical literature, Sahni (2013) finds experimental evidence of positive spillovers to rivals

in online restaurant advertising in India. Additionally, Lewis & Nguyen (2012) and Anderson & Simester (2013) find evidence of positive spillovers in a number of categories for online and mail advertising, respectively. Non-experimentally, Ching et. al. (2009) show evidence from scanner data that advertising of an individual brand with a display or feature could have spillover effects for the whole category.

The contributions of this paper are threefold. First, I improve upon the literature that seeks to identify the causal effect of advertising on own and rivals' demand with observational data by using an identification at the border approach. That is, I will identify advertising elasticities by comparing households that are very near to each other geographically but get different advertisements due to the way the television market borders are drawn. I show that advertising has significant positive effects on rivals' sales, though smaller than its effects on own firm sales. Next, I construct and estimate a consumer choice model, which allows advertising to influence the size of the category, the conditional share of each subcategory in the category, and the conditional share of each product in a subcategory. I will consider the category, the subcategory and the product levels as three separate stages of a joint physician-consumer decision making process. At each stage, I will allow for some inter-temporal influence. Results indicate that advertising of antidepressants affects all levels of choice. The category effects are larger and more persistent over time than are business stealing effects, leading to a net positive spillover. Finally, I conduct a supply side analysis to evaluate to what extent positive spillovers suppress the incentive to advertise. Given the demand parameters, I compute the Markov Perfect Equilibrium in television advertising. Next, I find that if firms that advertise work together, removing the need for strategic response, those firms would combine to advertise 50% more than in competitive equilibrium. A co-operative deciding all advertising expenditure levels taking full industry profits into account would advertise four times as much as is observed in competitive equilibrium, increase the category size by 18% and category profits by 14%. No other research that I am aware of conducts such a supply side analysis of the provision of advertising that exhibits positive spillovers. This paper helps move us toward understanding the effects of advertising and the incentives facing the firms who provide it, and understanding both are essential to firm profit maximization and to efficient regulation.

2 Empirical Setting

2.1 Prescription Drugs and Advertising

Television advertising of prescription drugs did not appear in the United States until 1997. While technically not forbidden by law, advertising was required to have much more risk information included on all advertisements than is required today. This required risk information was similar to the package inserts that come with prescriptions. Reading those aloud in the context of a thirty second spot was prohibitively time consuming and costly. In the fall of 1997, the FDA issued a draft memorandum clarifying their stance on advertising risk information, allowing advertisements to air so long as they had a 'fair balance' of risk information, even if abbreviated. Firms had the opportunity to submit their advertisements to the FDA for pre-approval to ensure that the 'fair balance' condition was met. In 1999 the final copy of the FDA memorandum was circulated. The first advertisements on television for antidepressants were seen in 1999 when GlaxoSmithKline's brand, Paxil, began airing its first campaigns.

Figure 1 suggests that the FDA regulation was binding prior to 1999, and advertising did not begin until that point.

2.1.1 Antidepressants

Prescription antidepressants are indicated for treatment of major depressive disorder and dysthymia, which is a more minor version of depression. Traditionally, depression was treated with what are called tricyclic antidepressants (TCAs), which were discovered in the 1950s, but those came with significant side effects and risks. Treatment of depression took a great leap forward in the late 1980s with the innovation of selective serotonin reuptake inhibitors (SSRIs), the first of which was Prozac. Newer generation antidepressants are more tolerable than the older generation TCAs and allow patients to be more safely treated and with fewer side effects (Anderson 2000). This allows easier management of antidepressant treatment by primary care physicians, and makes seeing a specialist less necessary.

Diagnosis and treatment of depression can be rather complicated, as with many mental disorders. As the class of drugs has grown, so have the number of people being treated. In 1996, the industry pulled in around \$5 billion in revenue. By 2004, it was up to \$13 billion. In 2004, an FDA black box warning was instituted suggesting that antidepressants might lead to an increase in suicidality among adolescents (Busch et. al 2012). Around the same time, many widely selling molecules began to go off patent. Figure 2 shows the revenues of the antidepressant industry from 1996 through 2004. Since the discovery of Prozac, ten other brands, some with slightly different mechanisms, have been discovered and have entered the market. Some of those have developed extended release versions which allow patients to have fewer doses per day.

There are six main subcategories of antidepressants: the old style TCAs, Tetracyclic (TeCA), Serotonin Antagonist and Reuptake Inhibitors (SARI), Serotonin-norepinephrine Reuptake Inhibitors (SNRI), Norepinephrine Reuptake Inhibitors (NDRI) and SSRI. While the specific differences between these are not important to this study, it is worth noting that each subcategory has somewhat different mechanisms, interactions and side effect profiles from the others. Deciding which subcategory of antidepressant is appropriate for a given patient is largely up to the physician, and often is related to other medications the patient is taking. The decision between drugs within a subcategory might depend on what is included on the patient's insurance formulary or physician preferences. Antidepressants are characterized by a high degree of experimentation to find a good fit between treatment and patient, as well as a low compliance rate due to the many side effects (Murphy et. al. 2009).

Many physicians see depression as an under treated condition and some research has concluded that restricting access to antidepressants has been associated with negative health outcomes (Busch et. al. 2012). Given this information, it is plausible that market expansive advertising could play a role in this market.

2.1.2 The Market for 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. Advertising purchased in the upfront market cannot be "returned" and typically has minimal flexibility in terms of timing. Next, there is a spot market that is called the 'scatter' market, where firms can purchase advertising closer to the date aired.

Additionally, there are both national and local advertisements. National advertisements are seen by everyone in the country tuned into a particular station, while 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 AC Nielsen, a global marketing research firm. 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 reside¹. 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 DMA². There are 210 DMAs in the United States, the largest 101 of which are included in my data.

From informal conversations with individuals in industry, I learned that pharmaceutical companies participate almost exclusively in the up front market. Like most consumer goods, the majority of antidepressant spending is on national advertising, but there is a significant amount of local advertising as well as significant variation across DMAs in the amount of local advertising.

Prices for advertisements typically are determined by projected volume and type of viewership. A single airing of a national advertisement for antidepressants ranges from \$1,600 to \$23,000 from 1999-2003 and a single airing of a local advertisement ranges from \$0 to \$7,600 for the same time period. Looking at each advertisement in terms of expenditure per capita, I observe that the distribution of local advertising expenditure per capita on a single commercial looks similar to the distribution of national advertising expenditure per capita on a single commercial. National advertisements range from \$0.0002 per 100 to \$0.04 per 100 and 93% of local advertisements fall within that range as well, with a few outliers going down to zero and up to \$0.20 per 100 capita. By scaling expenditures by potential viewing population, local and national advertising expenditures are comparable.

2.2 Data

2.2.1 Prescribing Data

Sales data for this market comes from the Xponent data set of IMS Health, a health care market research company. The prescribing behavior of a 5% random sample of physicians who prescribe antidepressants is followed monthly from 1997 until 2004. The data include a rich set of physician characteristics including address of the primary practice, which is then linked to county. The data used in this study is aggregated to the county level and ends with 2003, thereby avoiding confounding market changes in 2004 including the FDA black box warnings and wave of patent expirations. The sample is partially refreshed annually.

2.2.2 Advertising Data

Product level monthly advertising data at the national and Designated Market Area (DMA) level for the top 101 DMAs comes from Kantar Media. In addition to advertising expenditures, the data includes number of commercials. The

¹<http://www.sbca.com/dish-satellite/dma-tv.htm>

²<http://www.fcc.gov/encyclopedia/evolution-cable-television>

unit of advertising used in this study will be expenditures per 100 capita in the viewing area. Scaling expenditures by population in the viewing area allows me to have a comparable measure of advertising volume between national and local advertising. Total advertising for a county is defined as the national advertising expenditure scaled by the national population plus the local advertising expenditure scaled by the population of the DMA.³ Table 1 provides some descriptive statistics for the DMA level advertising variables at the product, subcategory and category level for the period of the data where advertising is allowed: September 1999 through December 2003. The statistics are also only on the products that ever advertise: Paxil, Paxil CR, Prozac, Prozac Weekly, Wellbutrin SR, Wellbutrin XL and Zoloft.

Figure 3 depicts local advertising expenditures per 100 capita in Boston, New York and Austin as well as national as examples of what local advertising expenditures look like over time. Local advertising for Paxil is higher in New York than it is in Boston, which in turn is higher than it is in Austin, suggesting that there is non-trivial variation across markets in this measure. National advertising makes up the bulk of the advertising that households see, but the local additions to the national advertising vary a great deal. The pattern is very similar in number of commercials, indicating that expenditures per capita are a reasonably comparable object across localities and national. However, as commercials are likely to be priced by reach, using expenditures per 100 capita should be a reasonably comparable way to measure reach of the ads. Figure 4 shows that national commercials and national expenditures per 100 capita are highly correlated.

2.2.3 Detailing Data

In addition to DTC data, I have collected physician level detailing data from ImpactRx, a market research firm. In the data, a panel of general practice physicians are followed monthly through time from 2001 through the end of 2003, and a panel of psychiatrists are followed monthly for 2002 and 2003. This panel is a national and geographically representative sample of physicians, most of whom are in the 40th percentile or greater in terms of total prescriptions written. The data in this sample will be aggregated in two different ways for analysis. First, it is aggregated to the DMA level to see if detailing and DTC decisions are correlated. Second, it is aggregated to the county level and directly included in the analysis using a restricted sample of only when the detailing data is available.

2.2.4 Other Data Sources

I observe prices from Medicaid reimbursement data, collected by the Centers for Medicare & Medicaid Services (CMS). Duggan and Scott-Morton (2006) argue that the average price that Medicaid pays per prescription prior to Medicaid rebates is a good measure of the average price of a drug on the market. As my measure of price, I use the total Medicaid prescriptions dispensed divided by the total Medicaid reimbursements during a quarter for a particular product, deflated to 2010 dollars using the consumer price index.

CMS also collects data on the average pharmacy acquisition cost for all pharmaceutical products (NADAC). As I will not be estimating marginal production costs empirically, these average pharmacy acquisition costs may be used as an

³A possible alternative measure would be to use the number of commercials at the national level plus the number of commercials at the local level. I explored using that measure and the results were not qualitatively different. However, as a commercial during the evening news is likely to capture far more eyeballs than a commercial during a 1:00 AM rerun of MacGyver, using expenditures per 100 capita would seem to do a better job at measuring quality adjusted advertising than number of commercials.

effective upper bound on marginal production costs. While there are markups from branded drugs, pharmacies are typically able to obtain generics at much lower rates, particularly when there are several generic competitors (as is the case in this market), often as low as ten cents per pill. As of 2013, all products in the sample have generic versions available. For an upper bound on the marginal cost of each drug, I use average pharmacy acquisition cost for those generic version of the product, deflated to 2010 dollars using the consumer price index. In Figure 5, the quantity weighted average margin in Boston is plotted versus time for the purposes of illustration. The average margin in the market rises as more popular branded antidepressants replace old generic TCAs and falls as more newer generation generics become widely prescribed.

Yearly county population and income data are drawn from the Current Population Survey (CPS).

Notable in the data is that there is both national and local advertising. While national advertising makes up the majority of advertising, there is significant spatial variation in the local additions to what households see.

Additionally, only four brands from three firms in this market advertise at all. Eli Lilly (Prozac, Prozac Weekly), Pfizer (Zoloft) and Glaxosmithkline (Paxil, Paxil CR, Wellbutrin SR, Wellburin XL) are the only firms advertising in this market. Notably, those firms, along with Merck, are some of the largest advertisers among all of the pharmaceutical industry (Berndt et al. 2003). The lack of advertising from all firms could be indicative of fixed costs of advertising at all or of free riding. Those branded products which do not advertise either have low market share (Effexor XR, Remeron, Serzone) or have a very small parent company which might be less likely to have an advertising division (Celexa, Lexapro).

Finally, own firm and rival firm advertising are negatively correlated, which could be an indicator of the positive spillovers of advertising in this setting.

3 Reduced Form Evidence

In this section, I explore the data to see if spillover effects exist and how they interact with own effects. This exercise has been difficult to implement in previous research, largely due to data limitations. Estimates show that rivals' and own advertising have a positive effect on sales, while rivals' advertising has a smaller effect than own advertising. In addition, the cross partials indicate that rivals' advertising makes own advertising less effective, but own advertising has a larger negative effect on the marginal own advertisement due to decreasing returns to scale.

In particular, I model sales of quantities Q of product j in time t for market m as a function of own advertising, a^{own} , and advertising of rivals, a^{cross} :

$$\log(Q_{jmt}) = \lambda \log(Q_{jmt-1}) + \gamma_1 a_{jmt}^{own} + \gamma_2 a_{jmt}^{cross} + \gamma_3 (a_{jmt}^{own})^2 + \gamma_4 (a_{jmt}^{cross})^2 + \gamma_5 a_{jmt}^{own} a_{jmt}^{cross} + \epsilon_{jmt} \quad (1)$$

This provides insight on whether rivals' advertisements help or hurt own demand, the nature of decreasing returns to scale, and persistence in advertising effects.

3.1 Empirical Identification Strategy - Border Strategy

The endogeneity of advertising and the absence of obvious instruments pose challenges to causal identification of the effect of advertising on demand.

I address the endogeneity concerns associated with advertising decisions by taking advantage of the discrete nature of local advertising markets. That is, two households which are directly across the television market border from one another will see different advertisements despite being otherwise very similar households. I take advantage of this comparison. This approach is similar in spirit to that used by Card and Krueger (1994) and Dube et. al. (2012) to identify the effects of minimum wage increases and that used by Holmes (1998) to identify the effect of right-to-work laws. Similar spatial strategies have been used by Black (1999) and Bayer et. al. (2007). Typically, these types of border analyses rely on state borders, across which any number of laws, market conditions or preferences may vary. A nice feature of television market borders is that they were set with television in mind and have very little correspondence with anything else in the world. As such, we might think the location of DMA borders is far more exogenous to consumer characteristics than are state borders.

Advertising is purchased both nationally and locally. The level of total advertising that a household gets to its television is determined by the Designated Market Area (DMA) that the household's county belongs to, as defined by AC Nielsen. Nielsen places counties into markets by predicting which local stations the households will be most interested in. As such, DMAs tend to be centered at metropolitan areas. A map of all of the DMAs included in the advertising data is presented in Figure 6.

To get an idea of how advertising is distributed across the country, consider the example of the Cleveland and Columbus DMAs. Figure 7 depicts the state of Ohio with each DMA in a different color. Every county in the mustard color Cleveland, Ohio DMA gets the same amount of the same advertising as every other county in the Cleveland DMA. Meanwhile, every county in the green color Columbus, Ohio, DMA gets the same amount of the same advertising as every other county in the Columbus DMA, though this might be different from the advertising in the Cleveland DMA. Meanwhile, these two DMAs border each other. There are five counties in the Cleveland DMA which share a border with at least one county in the Columbus DMA and five counties in the Columbus DMA which share at least one border with a county in the Cleveland DMA. My strategy will be to consider these ten counties as an experiment with two treatment groups (Cleveland and Columbus) in each time period.

The data contain 153 such borders. The map of all of the counties included in this border sample is presented in Figure 8. Each of these borders will be considered a separate experiment, with the magnitude of the treatment determined by the advertising in each DMA at a given time. Only the counties bordering each other will serve as controls for each other to partial out any local effects that may be increasing or decreasing for both sides of the border. The level of an observation is a product-border-DMA-month. This means that a group of counties along a particular border but in the same DMA are aggregated together, as they see they each see the same advertising and they are each being compared with a similar group across the border. In each 'experiment,' one such set of counties will be compared with an adjacent set of counties across the DMA border. For each border experiment in each time period, there will be two observations- one for the group of counties on one side of the border and one for the group of counties on the other side of the border. Each of these observation groups will constitute a market.

To estimate the effects of advertising in this experiment, I will use a modified difference-in-differences estimator. My identification assumption is that along the border of two DMAs, any differential trends in demand between the

two sides of the DMA border stem from differences in advertising. In particular, I use panel data with fixed effects. Border-time fixed effects will ensure that the common trend assumption is only enforced locally at the border between two DMAs, allowing for spatial heterogeneity. Border-DMA fixed effects will allow systematically different demand levels across the border. I will also include a lagged dependent variable to get at the dynamic effects of advertising. Consider the log of quantity $\log(Q_{jbmt})$, at the product-border-DMA-month level. Advertising, a_{jmt} , as mentioned before lives at the product-DMA-month level and affects $\log(Q_{jbmt})$ through some function f :

$$\log(Q_{jbmt}) = f(a_{jmt}) + \varepsilon_{jbmt}$$

Each product-border pair will constitute an experiment with border-markets being treatment groups. My fixed effects specification is:

$$\log(Q_{jbmt}) = \lambda \log(Q_{jbmt,t-1}) + f(a_{jmt}) + \alpha_{jbq} + \alpha_{jbm} + \varepsilon_{jbmt}$$

where the subscripts j and b indicate which experiment is being considered (product and border specific), α_{jbq} is a time effect which is used to control the experiment, which in this case will be a quarter fixed effect, α_{jbm} is a treatment group fixed effect, and $f(a_{jmt})$ is the magnitude of the treatment. The magnitude of the treatment is zero everywhere prior to 1999, as the FDA memo had not yet gone into effect. To investigate persistence in demand, a lagged dependent variable is also included. It should be noted here that the inclusion of α_{jbm} in the specification means that I am focusing on market level deviations from trend. That is, each market has a fixed effect. While Stremersch et. al. (2012) find that the distribution of DTC across markets may be explained by region-specific demographic composition, such cross-market level variation in advertising is accounted for in this specification with the fixed effect. The remaining variation being used is within market, within quarter deviations from the border experiment specific common time effect.

For further intuition, again consider the Cleveland-Columbus example and the case of Zolofit advertisements. In the equation above, $\log(Q_{jbmt})$ is log number of prescriptions of Zolofit in the Cleveland-Columbus border, indexed by month and which side of the border it is on. The magnitude of the treatment, $f(a_{jmt})$ is a function of the Zolofit's advertising in each market. The time effect, α_{jbq} , is a common quarter fixed effect between the Cleveland and Columbus sides of this border and is used to subtract out contemporaneous macro effects. The fixed effect, α_{jbm} , allows the different sides of the border to have systematically different levels in the outcome.

For this strategy to be valid, the Cleveland and Columbus sides of the border may differ by a fixed level, but they must have common trends absent advertising differences. Is this plausible? These counties are bordering, so they are very similar in geography. Both are sufficiently far from their central cities. The counties on the Cleveland side are only slightly closer to Cleveland than they are to Columbus and vice versa.

Also worth noting is that if Columbus always had a high, constant level of advertising and Cleveland always had a low, constant level of advertising, this estimation strategy would have no power to identify the effects of interest, as the border-DMA fixed effect would subtract out this variation, even though that advertising in Columbus might well have had an effect. Since prior to 1997, no DMAs had any advertising, there will be at least some variation in each experiment over time.

While it is clear that advertising is a firm choice rather than completely random, it is instructive to think about the potential sources of endogeneity and how the border strategy addresses those specific sources. Since there are market level fixed effects, endogeneity that comes from, say, the fact that winters in Florida being milder than winters in Wisconsin is not a concern. Those types of concerns are absorbed in the market level fixed effects. Potential bias can only come from within market, time specific demand shocks that affect the firm choice of advertising. Those shocks could come from two main sources: unobserved events (unseasonably bad weather, a large local employer laying off a large number of workers, an important medical seminar that discusses the virtue of these drugs) or rule of thumb based decision making.

First, consider the possibility of unobserved events. Since firm advertising decisions are made at the DMA level, the unobserved shocks of interest are the average across the DMA. Consider an unseasonably cold month that makes people more depressed, boosting both advertising and prescriptions of antidepressants. Weather patterns are continuous phenomena in that there is no reason that the weather should be significantly different on one side of a county border versus another. However, over larger distances, weather tends to be very different. As such, the average temperature over the DMA might be much colder than it is at the border of the DMA, but at the border the temperature will be very similar on both sides. The border strategy takes care of this type of endogeneity. Similarly, consider a large shock to employment in a given month in a DMA. This might simultaneously lead to a large increase in depression as well as an increase in advertisements, potentially biasing any estimated effect of advertising. Employment tends to be more concentrated in cities, which tend to be at the center of DMAs. The further away a person is from a place of employment, the less likely he or she will work there, just due to costs of commuting long distances. The distances just across the DMA borders to a central city are pretty similar and do not discontinuously jump as the border is crossed. As such, at the border, counties bordering but on opposite sides of the border are similar in their potential to be affected by any particular employment shock, but they are much less likely to be affected than those close to the center of the DMA. Again, the border strategy should be able to handle this source of endogeneity. Finally, consider a seminar meant to educate physicians about any particular course of treatment. This might increase the use of some antidepressant while also increasing advertising to the DMA where it occurs. Since these seminars also tend to be in center of DMAs, those at the outskirts are less likely to attend due to transportation costs, but those transportation costs do not discontinuously change at the DMA border. These three unobserved shocks seem to be the most likely to drive advertising decisions at the DMA level, and all are addressed by the border approach. It is important to note that there might be other potential unobserved shocks. However, so long as these shocks are reasonably continuous in distance from the center of the DMA, dissipate fast enough that the border is different from the center and do not discontinuously change at the border, the border approach will be valid. All of these requirements hold for the aforementioned examples.

Next, consider the possibility that firms use rules of thumb to allocate advertising based on the demand in the previous period. If this is the case, advertising in a DMA in the current period is determined by some function of last period's demand across the DMA. If previous period demand is correlated with current period demand, estimates will be biased. This is a classic reverse causality issue in advertising noted by Berndt (1991) and Bagwell (2007), among others. How does the border sample help this problem? The border counties comprise a fraction of the population, sales and counties in a DMA. Even if the trends for demand are exactly identical between the border areas and the DMA as a whole, the covariance between last period's demand in a DMA and current demand in a border area is a small fraction of the covariance between last period's demand in the whole DMA and this period's demand in the whole DMA. Further, the demand trends are likely to be different between the border and the full DMA, further reducing that covariance, and reducing any omitted variables bias. Again, by comparing the counties along the border

to their counterparts on the other side of the border via the common time trend, the omitted variable of last period's demand will largely get soaked into the product-border-time fixed effect, as demand is very similar immediately across the DMA borders.

In principle, the rule of thumb reverse causality problem could be solved using the full DMA but controlling for the previous period's demand. That is a problem in this setting since the previous period demand is already in the model and has an interpretation of its own. If previous period demand is effective as a control for rule of thumb advertising, it will inflate the estimate on the lagged dependent variable and its interpretation as a persistence parameter will be incorrect. The border approach, by making comparisons between similar counties that constitute a small fraction of total DMA demand, alleviates the concern of reverse causality.

3.1.1 Limitation of the Border Strategy

The main limitation of the border strategy is similar to that of a usual regression discontinuity design in that the estimated treatment effects are identified at the border and not elsewhere. It might be the case that the true treatment effect in the interior of the DMA is different from that at the border. If this is the case, the interpretation of the supply analysis will be limited as I assume that the estimated advertising effects hold in both border and non-border counties. This concern is partially addressed in a robustness check in the appendix. In particular, I separate out borders that are closer to urban centers and estimate the model separately for those borders than from those that are further out. The effects are very similar and not statistically distinguishable from the full border sample. While this might not fully establish that the effect in the interior of the DMAs is the same as the effect at the borders, it provides a small piece of evidence that the effects seem to hold up similarly across different types of counties. Further, I compare measurable characteristics for in-sample counties (those at the borders) with those out of the sample (those at the interior). Those comparisons are available in Appendix F.2. A t-test fails to reject that border counties and interior counties are the same in average population, average income, average number of physicians and number of non-federal physicians. To the extent that we continue to worry that the estimated treatment effects are different at the interiors, that will be a limitation of this analysis. Similar limitations apply to other regression discontinuity designs as well as to instrumental variables methods that reveal only local average treatment effects.

3.1.2 Potential Threats to the Border Strategy

One potential worry is that there would be little variation net of the fixed effects. This would be the case if too much of the advertising were national and not enough were local. Figure 9 displays a histogram of advertising net of these fixed effects showing significant variation. Net of fixed effects, the log of advertising expenditures per 100 capita has a mean of zero and a standard deviation of 0.25, so there is substantial variation net of fixed effects.

Also potentially problematic is the lagged dependent variable, which can generate omitted variables bias in the presence of small T , as differencing mechanically induces correlation between the lagged dependent variable and the error term. However, as $T \rightarrow \infty$, the mechanical correlation with the error term diminishes to zero and the fixed effects estimator is consistent. As my data is monthly from 1997 through 2003, $T=84$ should be sufficiently large that any

bias will be minimal⁴.

Additionally, we might be concerned about measurement error. There are a two main possibilities that could lead to measurement error and biased estimates:

1. Consumers watch advertisements in one DMA, but drive across the border to see their physicians.
2. Consumers watch advertisements in one DMA, but drive towards the center of the DMA to a county not included in the border sample to see their physicians.
3. Consumers watch advertisements over the air and sometimes see the advertisements from the DMA that is on the other side of the border.

All of these scenarios would lead me to understate the effect of the advertisements. To the extent that we think that these biases are present, we can look at my estimates as lower bounds on the true parameters.⁵

It should also be noted that very few consumers watch over the air. According to the Consumer Electronics Association, fewer than seven percent of households rely on over the air signals for their television⁶. Further, at the DMA border, TV signals tend to be less reliable over the air, as stations tend not to locate in the outskirts of DMAs. It is likely that consumers would be even less likely to rely on over the air signals at the DMA border.

Omitted variables bias could also be a source of bias. Prices and detailing are omitted from this estimation. As prices tend not to vary geographically due to a very low transport cost and ease of obtaining drugs through the mail, prices are absorbed in the product-border-time fixed effect. Any national average effects of detailing are also controlled for by the product-border-time effect. Where there might be trouble, however, is if firms strategically raise (or lower) detailing at the product-market-time level in exactly the same places where DTC is concentrated to take advantage of any complementarities or substitutabilities. For a panel of physicians and a fraction of the time period, ImpactRx data has physician specific detailing information on the number of minutes sales representatives spend with physicians. In appendix A, I show that for any given time period and market, DMA totals of detailing minutes are uncorrelated with DTC. Even if detailing plays a significant role at the local level, the fact that it is close to orthogonal to DTC will make the omitted variables bias close to zero. I confirm this intuition by including county level detailing minutes into a restricted specification of the main model in the paper and show that the inclusion of detailing does not affect estimates of the effect of DTC. Since the number of time periods are greatly reduced when using detailing, my preferred specifications will omit detailing and use the long time series. These results are perhaps unsurprising given that in the literature Manchanda et. al. (2004) find that detailing is largely determined by practice size, which is effectively controlled for in the market fixed effects.

A further concern is that various policies or cost inputs could discontinuously change right at the DMA border, causing the more impressionable physicians to locate on a particular side of the DMA border. As DMAs are in general only

⁴In Nickell (1981)'s paper that describes the bias induced when doing fixed effects and lagged dependent variables, he analytically solves for the bias as a function of T. The bias is: $\frac{plim}{N \rightarrow \infty} (\hat{\lambda} - \lambda) = \left\{ \frac{2\lambda}{1-\lambda^2} - \left[\frac{1+\lambda}{T-1} \left(1 - \frac{1}{T} \frac{(1-\lambda^T)}{(1-\lambda)} \right) \right]^{-1} \right\}^{-1} \approx \frac{-(1+\lambda)}{T-1}$, where the approximation holds for "reasonably large" T. With T=84, that approximation is bounded above by $-\frac{2}{83} \approx -0.02$. If we do not wish to concede that 84 is reasonably large, plugging in 0.7 as the true λ , the exact bias formula gives the bias at -0.025.

⁵However, the Dartmouth Institute has drawn primary care commuting zones which describe how far Medicare patients travel to see their physicians. It is very rare for a commuting zone to cross DMA lines- only about 1% of primary care commuting zones cross DMA borders at all, and those that do tend to be predominantly in only one DMA. This should minimize the measurement error worry. Further explanation of the Dartmouth Institute commuting zones is provided in the appendix.

⁶<http://www.tvtechnology.com/default.aspx?tabid=204&entryid=9940>

relevant to television markets, it is hard to imagine why any tax laws would systematically vary across DMA borders. Almost all business tax policies are set by state governments or potentially large city governments. Being on the border of the DMA typically leaves those counties out of reach of large city specific taxes. However, as many DMA borders coincide with state borders, state tax policies could be a problem. To address that concern, I have removed the DMA borders that coincide with state borders and re-estimated the whole model. The results of the estimation are available in Appendix C. The estimated parameters are not statistically different from those if all borders are left in the sample.

Data on corporate rental rates by county is not available for this paper, so it is impossible to assess whether or not there is a discontinuity in the rental rates of physician offices across a DMA border. However, it is hard to imagine that within a state for two very similar counties which border each other and are similar distances from major cities that rental rates would differ significantly. Furthermore, if they did, it is hard to imagine that advertising decisions for the full DMA would hinge on rental rates at the border of the DMAs. The only worry is if the physicians that select into cheaper rent are systematically those which have different responsiveness to advertising than those who select into more expensive rent. To further examine the issue of selection across the border, data was collected from the Area Resource File to see if the number of physicians, the average income or the population is significantly different on the higher advertising side of the border from the lower advertising side of the border. T-tests cannot reject that all of these variables are the same across the borders. Results of these tests of balance are available in Appendix F.

Finally, the identifying assumption of difference-in-differences could be violated. It could be that the difference-in-differences model fails the parallel trends assumption, invalidating the difference-in-differences design. To address this concern, I have conducted a placebo test. Using data on DMA level television advertising of over-the-counter sleep aids as a placebo treatment, I find no economically significant effects. Details for this robustness check are in the appendix.

3.1.3 Why the border strategy?

A more conventional identification strategy in the discrete choice literature is to use an instrumental variables approach, as in Berry, Levinsohn and Pakes (1995). The main identifying assumption for the validity of the BLP instruments is that the characteristics of competing products within a market are exogenous, thus the changing competitive structure of the market may be used as a supply side instrument for demand side choice variables. In the market for prescription drugs, entry happens in all markets simultaneously by all products, thus use of the BLP instruments would eliminate any spatial variation, which is a main attribute of the data I am using. Furthermore, it might be unreasonable to think that competitor characteristics are exogenous in this setting. It stands to reason that as consumers demand more antidepressants with fewer of some kind of side effects that firms might well focus research and development on that kind of product. I am more comfortable with assumptions required by the border strategy in addition to the fact that it allows me to use the spatial variation in the data.

3.2 Results

Using the identification strategy at the border outlined, the estimating equation including fixed effects becomes:

$$\log(Q_{jmt}) = \lambda \log(Q_{jmt-1}) + \gamma_1 a_{jmt}^{own} + \gamma_2 a_{jmt}^{cross} + \gamma_3 (a_{jmt}^{own})^2 + \gamma_4 (a_{jmt}^{cross})^2 + \gamma_5 a_{jmt}^{own} a_{jmt}^{cross} + \alpha_{jbq} + \alpha_{jbd} + \varepsilon_{jmt} \quad (2)$$

where α_{jbq} is a product-border-quarter fixed effect and α_{jbd} is a product-border-DMA fixed effect. The α_{jbq} effect will sweep out all variation that is not between two areas that are on opposite sides of a DMA border. The product-border-DMA fixed effects sweep out all variation that is due to persistent differences between different markets (e.g. people are generally more depressed in New York than Wisconsin).

Partialing these fixed effects out makes the identifying variation within product j local advertising that is over and above the average on its side d of the border b and over and above the average local advertising of product j in time period t in all counties on that either side of border b .

Results of the above regression are provided in Table 2. Most notable is that both rivals' and own advertising has a positive and significant effect on demand. Rivals' advertising hits decreasing returns to scale more slowly than does own advertising. Also, the cross partial indicates that rivals' advertising works a firm down its marginal revenue curve with respect to advertising, but not as much as own advertising does. Finally, there is evidence of persistence, though the persistence parameter is not especially large. This is consistent with the idea that there is much experimentation to find the correct fit between patient and treatment in the depression space. If the category expansive effect of advertising is more persistent, it will be another potential avenue for firms to under-invest relative to a co-operative.

4 Model

4.1 Demand

I propose a multi-stage choice model where advertising may affect the consumer's choice at each stage. A consumer arrives at her desired end product through a sequence of choice problems. First, the consumer chooses between entering the category (inside option) and the outside option. If she chooses to enter the category, she chooses which subcategory of product she wants. Finally, given her choice of subcategory, she chooses which product to purchase. This process can be extended, in principle, to have any number of stages.

In the specific case of prescription antidepressants, this is plausible. A consumer first decides whether she has a problem with depression, goes to the physician and together with the physician, determines which class of drugs would be most suitable (perhaps considering interactions with other drugs taken) and which product in particular is the best choice (perhaps having to do with what is on her formulary).

I define 'utility' u of consuming the inside option, as a function of total advertising stock as well as other market level factors:

$$u_{ilmt} = \Gamma_1(A_{lmt}) + \beta_1 X_{lmt} + \alpha_t + \alpha_{lm} + \xi_{lmt} + \varepsilon_{ilmt} = \delta_t + \varepsilon_{ilmt}. \quad (3)$$

In this specification, l denotes the inside option versus outside option, m denotes market and t denotes time period. I define Γ_1 as an increasing function of A_{lmt} , total advertising stock of all inside option products in market m at time t , α_{lt} is a time specific taste for the inside option, α_{lm} is a market specific taste for the inside option, and X_{lmt} are market-time characteristics.

For the next stage, I define the relative utility v^l of subcategory n conditional upon the choice of the inside option as a function of the total advertising stock in subcategory n , A_{nmt} , as well as other subcategory-market-time level factors:

$$v_{nmt}^l = \Gamma_2(A_{nmt}) + \beta_2 X_{nmt} + \alpha_{nt} + \alpha_{nm} + \xi_{nmt} + \varepsilon_{nmt} = \delta_{n|l} + \varepsilon_{nmt} \quad (4)$$

Finally, relative utility w^n of product j conditional upon the choice of subcategory n , is defined as a function of advertising stock of product j , A_{jmt} , and other product-market-time level factors:

$$w_{jmt}^n = \Gamma_3(A_{jmt}) + \beta_3 X_{jmt} + \alpha_{jt} + \alpha_{jm} + \xi_{jmt} + \varepsilon_{jmt} = \delta_{j|n} + \varepsilon_{jmt} \quad (5)$$

Dynamics enter the model through advertising carry-over. That is, a consumer may remember an advertisement from a previous period, and that advertisement may affect current period demand. In general, advertising stock is a function of current period advertising (measured in expenditure per 100 capita) in choice stage s , a_s , where $s \in \{l, n, j\}$, last period's advertising stock, A_{smt-1} and a parameter governing depreciation over time, λ_s .

$$A_{smt} = f(\lambda_s, A_{smt-1}, a_{smt}) \quad (6)$$

I set each disturbance term, ε , to be iid extreme value type I. Given the logit errors, I compute a closed form solution for shares. The unconditional share of product j in subcategory n is a product of conditional shares, where market and time subscripts have been suppressed:

$$s_j = (s_{j|n})(s_{n|l})(s_l) \quad (7)$$

Those conditional shares take logit form:

$$s_{j|n} = \frac{\exp(\delta_{j|n})}{1 + \sum_{j \in n} \exp(\delta_{j|n})} \quad (8)$$

$$s_{n|l} = \frac{\exp(\delta_{n|l})}{1 + \sum_n \exp(\delta_{n|l})} \quad (9)$$

$$s_l = \frac{\exp(\delta_l)}{1 + \exp(\delta_l)} \quad (10)$$

I note here that the equations at each level are independent of each other.⁷I allow each level to have a different persistence λ_s , and different effects of advertising, γ . I also note that while I call the latent variables at each level ‘utilities’, it is not essential to interpret them literally as such. In this paper, I will not be computing consumer welfare, and it is likely that the latent variables contain a combination of patient and physician utility, information and persuasion. The purpose of the choice model is to guide the firm decision problem. While it is possible that these parameters could be related across levels by some kind of summing up identity (as they would if each of the equations were only utility and consumers maximized utility), I do not restrict them to be, as discovering the relative magnitudes of advertising effects at each level is a main question of this paper.

I also note that this model incorporates unobserved heterogeneity through the inclusion of fixed effects, both for the market and for the comparison group time effect. That is, there might be different effects for each market and time period, and the average of those effects will be the reported coefficient on advertising. Including heterogeneity over and above these fixed effects resulted in no significant findings, perhaps because so much variation is explained by these fixed effects.

For general intuition of the model, consider what happens if a single product, Zoloft, raises advertising in a market while everything else remains constant. That advertisement may have three effects. First, it may raise the probability that a consumer purchases any antidepressant. That effect is expressed through the top level equation, increasing a_{lmt} , which increases A_{lmt} which then in turn increases $\Gamma_1(A_{lmt})$. Next, the information in the advertisement may push the consumer towards the subcategory of antidepressants that Zoloft is in over another, as the commercials often contain information about mechanisms and side effects, which are highly correlated within subcategory. The Zoloft advertisement increases a_{nmt} , which increases A_{nmt} , which in turn increases $\Gamma_2(A_{nmt})$. The marginal revenue will depend on the shape of the curve and the amount of advertising done by other products in the same subcategory. Finally, the advertisement may have a pure business stealing effect. By increasing a_{jmt} , A_{jmt} and $\Gamma_3(A_{jmt})$ increase to take share away from other products within the subcategory.

It is also worth noting that the model does not explicitly examine the various possibilities for forward looking consumers as in Dickstein (2014) or Crawford & Shum (2005). However, the difference in persistence parameters from the category level to the product levels allows for consumers to purchase one brand, decide that it does not work satisfactorily and move to another brand. Further, the product-time fixed effects allow differences in market conditions to be taken into account. That is, in the year prior to Prozac going off patent, consumers may wish to be prescribed Prozac rather than Zoloft because they know a cheaper generic will be available in the following period. This effect is absorbed into the product-time effects.

4.1.1 Derivatives and Elasticities

Given product shares in equation (4) and the logit structure, we can get the derivative of s_j which is in subcategory n with respect to new advertising, a_k , of product k which is in subcategory n' by using the chain rule and the typical logit derivatives:

⁷As each equation is a ‘conditional’ statement, the independence of the error terms seems reasonable. That is, the error term at the business stealing level indicates “conditional on already having chosen to get an antidepressant and having decided that an SSRI is appropriate, what is my idiosyncratic taste for Prozac versus Zoloft.” It is hard to imagine why a relative preference between Prozac and Zoloft should affect a consumer’s absolute taste for antidepressants.

$$\frac{\partial s_j}{\partial a_k} = s_{j|n} \left[s_{n|l} \frac{\partial s_l}{\partial a_k} + s_l \frac{\partial s_{n|l}}{\partial a_k} \right] + s_{n|l} s_l \frac{\partial s_{j|n}}{\partial a_k} \quad (11)$$

solving this out using our specification on shares, we get derivatives,

$$\frac{\partial s_j}{\partial a_k} = \begin{cases} s_j \left[\frac{\partial \Gamma_1}{\partial a_k} (1 - s_l) + \frac{\partial \Gamma_2}{\partial a_k} (1 - s_{n|l}) + \frac{\partial \Gamma_3}{\partial a_k} (1 - s_{j|n}) \right] & j = k \\ s_j \left[\frac{\partial \Gamma_1}{\partial a_k} (1 - s_l) + \frac{\partial \Gamma_2}{\partial a_k} (1 - s_{n|l}) - \frac{\partial \Gamma_3}{\partial a_k} s_{k|n} \right] & j \neq k, \& n = n' \\ s_j \left[\frac{\partial \Gamma_1}{\partial a_k} (1 - s_l) - \frac{\partial \Gamma_2}{\partial a_k} s_{n'|l} \right] & j \neq k \& n \neq n' \end{cases} \quad (12)$$

and advertising elasticities equal to,

$$\eta_{jk} = \begin{cases} a_k \left[\frac{\partial \Gamma_1}{\partial a_k} (1 - s_l) + \frac{\partial \Gamma_2}{\partial a_k} (1 - s_{n|l}) + \frac{\partial \Gamma_3}{\partial a_k} (1 - s_{j|n}) \right] & j = k \\ a_k \left[\frac{\partial \Gamma_1}{\partial a_k} (1 - s_l) + \frac{\partial \Gamma_2}{\partial a_k} (1 - s_{n|l}) - \frac{\partial \Gamma_3}{\partial a_k} s_{k|n} \right] & j \neq k \& n = n' \\ a_k \left[\frac{\partial \Gamma_1}{\partial a_k} (1 - s_l) - \frac{\partial \Gamma_2}{\partial a_k} s_{n'|l} \right] & j \neq k \& n \neq n' \end{cases} \quad (13)$$

From these equations, we can see that firm benefits from own advertising may flow through expansion of the category, as is denoted by the term $s_j \frac{\partial \Gamma_1}{\partial a_j} (1 - s_l)$, through expansion of the subcategory in $s_j \frac{\partial \Gamma_2}{\partial a_j} (1 - s_{n|l})$ and through business stealing within the nest in $s_j \frac{\partial \Gamma_3}{\partial a_j} (1 - s_{j|n})$. Firm benefits from rivals' advertising in the same subcategory may flow through expansion of the category in $s_j \frac{\partial \Gamma_1}{\partial a_k} (1 - s_l)$ or through expansion of the subcategory in $s_j \frac{\partial \Gamma_2}{\partial a_k} (1 - s_{n|l})$, while this same advertising may hurt through business stealing within the subcategory in $-s_j \frac{\partial \Gamma_3}{\partial a_k} s_{k|n}$. Advertising from rivals in other nests may benefit the firm only through the expansion of the inside option, but may hurt through expansion of the other subcategory at the expense of the firm's subcategory. It is worth noting that this structure fully allows for advertising that is a pure category expansion (i.e. if $\frac{\partial \Gamma_2}{\partial a_j} = \frac{\partial \Gamma_3}{\partial a_j} = 0 \forall j$), for advertising that is pure business stealing (i.e. if $\frac{\partial \Gamma_2}{\partial a_j} = \frac{\partial \Gamma_1}{\partial a_j} = 0 \forall j$), or anything in between, including cross subcategory substitution. It is also possible that rival advertising outside of the subcategory could help more than inside of the subcategory if $\frac{\partial \Gamma_2}{\partial a_j}$ is sufficiently small and $\frac{\partial \Gamma_3}{\partial a_j}$ is sufficiently large or vice versa. What is restricted is that a firm's own advertisements may not help another firm more than it helps itself in elasticity terms. In the most extreme scenario, it is pure category expansion and helps all firms equally. Whether advertising provides positive or negative spillovers depends on the relative strength of the market expansion and the business stealing channels and is a result of estimation rather than an assumption of the model.

Notable is that through the category expansion channel, rivals' advertising moves a firm's marginal revenue with respect to advertising downward. However own advertising must move a firm's residual marginal revenue curve even further downward, as there are decreasing returns at the conditional share level as well. Assuming that the effect of advertising is positive at all levels, the primary effect of own advertising is stronger than that of rivals' advertising and decreasing returns to own advertising are more severe than decreasing returns to rival advertising. This is a testable implication of the model.

4.2 Supply

Firm free riding may be an optimal strategy in a game with positive spillovers. To investigate the incentives generated by the demand problem above, I assume that firms play a simultaneous game, choosing advertising each period while taking into account expectations of rival behavior and the dynamic effects of advertising. This enables me to analyze the magnitude of potential under-provision levels on average generated by positive spillovers.

4.2.1 The Firm's Problem

A forward looking firm maximizes a discounted stream of future profits with respect to advertising. Suppose advertising has a constant marginal cost k_{jm} , the market size is μ_m , prices are p_j , marginal production costs are mc_j and the discount rate is β . Further, suppose the exogenously given set of products that a firm f has in the market is denoted by Φ_f and the full set of products in the market is denoted by $\cup_f \Phi_f$.

Per period profit for the firm will be a function of advertising stock at the product, subcategory and category level $\mathbf{A}_{mt} = \{A_{jmt}, A_{nmt}, A_{lmt}\}$, which is a function of the vector of current advertising for all products, \mathbf{a}_{mt} , advertising stocks in the previous period $\mathbf{A}_{m,t-1}$, and persistence parameters λ_j, λ_n , and λ_l . Per period profit is also a function of other variables X , a constant marginal cost of advertising k_{jmt} , and product-market-time specific iid demand shocks ξ :

$$\pi_{fjmt} = \sum_{j \in \Phi_f} (p_{jt} - mc_{jt}) \mu_{mt} s_{jmt}(\mathbf{A}_{mt}, X, \xi) - k_{jmt} a_{jmt}.$$

The firm's problem is to maximize the stream of future profits for all products in its portfolio. Advertising is set prior to the realization of any demand shocks, so only expected profits matter for firm choices. Expected per period profits are:

$$\pi_{fjmt} = \sum_{j \in \Phi_f} \int (p_{jt} - mc_{jt}) \mu_{mt} s_{jmt}(\mathbf{A}_{mt}, X, \xi) p(\xi) d\xi - k_{jmt} a_{jmt}. \quad (14)$$

The timing of the game is as follows. At the start of each period, the state of the market \mathbf{g}_{mt} , is revealed to all firms. Based on that state, firms will make advertising decisions $\sigma_{jm}(\mathbf{g}_{mt}) = a_{jmt}$. The state variables could include anything in the past history of the game, but I will restrict attention to games where payoff-relevant state variables are the only ones that matter. Given the demand system provided above, state variables will include only the current advertising stock $\mathbf{A}_{mt} = \{A_{jt}, A_{nmt}, A_{lmt}\}$. After the state variables are observed, the firm makes its advertising decision, the current period demand shocks are realized and the firms collect their profits for the period.

The strategy profile $\sigma = (\sigma_1, \dots, \sigma_f)$ contains the decision rules of the firms. Firm f 's expected present discounted value of all profits given the current state \mathbf{g}_{mt} and the full strategy profile σ is

$$V_{fm}(\mathbf{g}_{mt} | \sigma) = \mathbb{E} \left[\sum_{j \in \Phi_f} \sum_{\tau=t}^{\infty} \beta^{\tau-t} \pi_f(\mathbf{g}_{m\tau}, \sigma_f(\mathbf{g}_{m\tau})) | \mathbf{g}_{mt} \right]. \quad (15)$$

Firm f chooses a sequence of advertising levels σ_f , which are state dependent, to maximize $V_{fm}(\mathbf{g}_{mt}|\sigma)$. I will assume that firms form expectations regarding future states and thus future strategies given current period state variables. As such, firm f will have a *Markov strategy*, $\sigma_{fm} : \mathbf{g} \rightarrow \sigma_f(\mathbf{g}) = (a_{jm})_{\forall j \in \Phi_f}$. To assess the strategies of competitors, firm f makes an assumption about σ_{-f} and chooses σ_f .

The equilibrium concept used will be *Markov perfect equilibrium*. A *Markov perfect equilibrium* is a set of strategies that form a subgame perfect equilibrium where strategies may only be conditioned on payoff relevant state variables and on the current state of the game. In particular,

$$\sigma_f^* \in \operatorname{argmax}\{V_{fm}(\mathbf{g}_{mt}|\sigma)\} \quad (16)$$

for all states \mathbf{g}_{mt} , firms f , and actions a_{jm} . That is, each firm will maximize the discounted sum of expected profits given beliefs about competitors behavior, and each firm's beliefs are mutually consistent in equilibrium. Using MPE makes the allowable strategy space relatively simple. I will further restrict attention to pure strategies in advertising. Firms will use the Bellman equation to solve the dynamic programming problem:

$$V_j(\mathbf{g}|\sigma) = \sup_{a \in \mathbb{R}^+} \{\pi_j(\mathbf{g}, a, \sigma_{-j}(\mathbf{g})) + \beta \mathbb{E}V(\mathbf{g}'|\sigma)\} \quad (17)$$

4.2.2 Other Choice Variables

Other choice variables are excluded from my model for reasons of expositional clarity and data limitation. Prices are observed only at the product-quarter level and detailing only at the product-month (national aggregate) level. While it is indeed conceivable that firms would maximize profit jointly with respect to both price and advertising, the purpose of this study is to isolate the advertising decision. In many settings this may not be possible. However, in the pharmaceutical market that I study, the institutional detail of insurance as an intermediary makes pricing largely decided by a collection of bilateral negotiations between the firm, insurance companies and health care providers. The resulting prices observed in the data are highly persistent over time and show no correlation with observed advertising. Persistence in brand pricing along an increasing trend, even after patent expiration is noted throughout the literature (Aitken et. al. 2013, among others). As such, I will view the pricing decision as orthogonal to the DTC decision, conditional on product-market and product-time fixed effects. Under this assumption, there will be no bias in the estimation of the effects of DTC rising from omitting price from the estimation. A discussion of pricing in this market as well as the analysis showing that it is not correlated with advertising is available in Appendix A.⁸

Detail advertising to the physician is another choice variable and type of marketing pursued in the pharmaceutical world. It is possible that firms could strategically detail in response to DTC in a particular market. In Appendix A, I show that in the antidepressant data for this time period, DTC and detailing decisions are, in fact, orthogonal. This may be due to some organizational constraint or unobserved cost, but it appears as though in this time period, detailing levels do not adjust to increases or decreases in DTC. As such, I will assume that detailing decisions are

⁸Without including price in the analysis, I will not be estimating a price elasticity. As such, the computation of optimal advertising will be a bit different in spirit than the classic Dorfman-Steiner (1954) problem which suggests that the advertising to sales ratio should equal to the advertising elasticity to price elasticity ratio. This theorem was modified to allow for dynamic advertising in Nerlove and Arrow (1962). Neither of these formulations considers positive spillover effects, and free riding incentives may make these policies sub-optimal.

made separately from DTC decisions. This will allow me to focus on the incentive effects of the positive spillover effects of DTC in isolation.

Omission of pricing and detailing may still affect some conclusions in counterfactual simulations and estimations and discussion of those concerns will be explored in those sections.

5 Empirical Specification and Estimation of the Model

5.1 Demand Specification

I define the advertising stock at each level s , where $s \in \{l, n, j\}$ is either the category level, subcategory level or product level, to be a lag of a nonlinear function of current advertising, similar to Dube et. al. (2005).

$$A_{smt} = \sum_{\tau=0}^t \lambda_s^{t-\tau} \log(1 + a_{sm\tau}) \quad (18)$$

Specifying advertising stock as a concave function of each period's advertising allows the firm's problem to have a well behaved optimum. Other functional forms were explored and none changed the results in any significant way.

The advertising stock enters into the utility specification linearly at each level.

$$\Gamma_s(A_{smt}) = \gamma_s A_{smt} \quad (19)$$

I account for all product characteristics other than advertising with a rich set of fixed effects, as the only pieces of data that vary at the choice level, DMA and time levels are shares and advertising.

Substituting equations (18) and (19) into equations (1)-(3), for the market level l obtain:

$$u_{ilmt} = \gamma_l [\sum_{\tau=0}^t \lambda_l^{t-\tau} \log(1 + a_{lm\tau})] + \alpha_{lt} + \alpha_{lm} + \xi_{lmt} + \varepsilon_{ilmt} \quad (20)$$

The conditional utilities for the subcategory and product levels are defined analogously. From here, it is notable that current period advertising enters the utility function in a concave manner, so the firm maximization problem is well behaved.

5.2 Transforming to a Linear Problem

Following Berry (1994), at each level of the problem, I specify an 'outside good', take the log of the market share and subtract from it the log of the outside option share. This results in a linear form.

At the category level the outside good is naturally defined as the population not filling a prescription for an antidepressant in month t in market m :

$$\log(s_{lmt}) - \log(s_{omt}) = \gamma_1 [\sum_{\tau=0}^t \lambda_l^{t-\tau} \log(1 + a_{lm\tau})] + \alpha_{lt} + \alpha_{lm} + \xi_{lmt} \quad (21)$$

At the subcategory level, the outside good will be defined as the subcategory of older style TCA antidepressants. The share of a subcategory conditional on being in the inside option follows:

$$\log(s_{nmt|l}) - \log(s_{omt|l}) = \gamma_2 [\sum_{\tau=0}^t \lambda_n^{t-\tau} \log(1 + a_{nm\tau})] + \alpha_{nt} + \alpha_{nm} + \xi_{nmt} \quad (22)$$

At the product level, the outside option in each nest will be the set of all products that never advertise on television. The product share equation conditional on already having chosen subcategory n is:

$$\log(s_{jmt|n}) - \log(s_{omt|n}) = \gamma_3 [\sum_{\tau=0}^t \lambda_p^{t-\tau} \log(1 + a_{jm\tau})] + \alpha_{jt} + \alpha_{jm} + \xi_{jmt} \quad (23)$$

Now, using these to solve for inside option shares in time $t - 1$, and substituting that back into the expression for time t shares yields,

$$\log(s_{lmt}) - \log(s_{omt}) = \lambda_l (\log(s_{lmt-1}) - \log(s_{omt-1})) + \gamma_1 \log(1 + a_{lmt}) + \theta_{lt} + \theta_{lm} + v_{lmt} \quad (24)$$

where

$$\theta_{lt} = \alpha_{lt} - \lambda_l \alpha_{lt-1} \quad (25)$$

is a inside option-time specific taste or quality parameter.

and

$$\theta_{lm} = \alpha_{lm} - \lambda_l \alpha_{lm} \quad (26)$$

is the category-market specific taste parameter. Finally,

$$v_{lmt} = \xi_{lmt} - \lambda_l \xi_{lmt} \quad (27)$$

is a market-time specific demand shock. Equation (24) is precisely a lagged dependent variable with fixed effects specification as described above, making possible the use of the border identification strategy.

Similarly, subcategory and product level share equations may be specified as:

$$\log(s_{nmt|t}) - \log(s_{0mt|t}) = \lambda_n(\log(s_{nm,t-1|t}) - \log(s_{0m,t-1|t})) + \gamma_2 \log(1 + a_{nmt}) + \theta_{nt} + \theta_{nm} + v_{nmt} \quad (28)$$

$$\log(s_{jmt|n}) - \log(s_{0mt|n}) = \lambda_p(\log(s_{jm,t-1|n}) - \log(s_{jm,t-1|n})) + \gamma_3 \log(1 + a_{jmt}) + \theta_{jt} + \theta_{jm} + v_{jmt}. \quad (29)$$

5.3 Identification and Estimation Strategy

Since the share equations have been transformed to a linear form, estimation may be done by OLS. A notable problem in estimating this equation is that advertising is a firm choice variable determined in equilibrium and is thus endogenous. As such, I will take advantage of the discrete nature of DMAs to make use of spatial variation as described in section (3).

In particular, I specify the estimation equation as:

$$\log(s_{lmt}) - \log(s_{0mt}) = \lambda_l(\log(s_{lm,t-1}) - \log(s_{0m,t-1})) + \gamma_1 \log(1 + a_{lmt}) + \theta_{lbq} + \theta_{lbn} + v_{lbn} \quad (30)$$

where θ_{lbq} is a border-time fixed effect and θ_{lbn} is a border, DMA fixed effect. Partialing these fixed effects out makes the identifying variation at the market level the total advertising in market m that is over and above the average on its side d of the border b and over and above the average local total advertising in quarter q in all counties on that either side of border b . The fixed effects will also control for the product quality terms θ_t and θ_m .

I identify the effects at the other two levels similarly. In the subcategory level, I include fixed effects α_{nbq} and α_{nbn} and at the product level, I include fixed effects α_{jbq} and α_{jbn} . Identifying variation will come at the subcategory level from total subcategory advertising that is above and beyond the historical advertising in its market and above the border average in the current time period. At the product level, identifying variation will be advertising for product j that is above and beyond advertising for product j on the border in quarter q and above and beyond the average over all time in market m . No between product variation in advertising will be used to identify the advertising parameter.

Table 3 has variable definitions and summary statistics for those variables that will enter the estimation.

5.4 Demand Results

5.4.1 Effects at Each Level

Results are presented in Table 4. The effect of advertising stock on demand at each stage of the decision is positive. The strongest effects are at the category level, deciding between inside and outside option and at the product business stealing level. Effects at the subcategory level are not significant, but it is notable that there is only advertising in two subcategories, with most of the advertising happening in the SSRI subcategory. The small and insignificant effect

at the subcategory level is not surprising, as it seems unlikely that patients would have good information about what separates the subcategories. Table 5 presents short run demand elasticities of current advertising showing that the category expansive properties of advertising dominate the business stealing effects and all cross advertising elasticities are positive. This finding is consistent with the identified positive spillovers in the reduced form.

5.4.2 Persistence

Persistence is highest at the category level- that is getting someone into consuming antidepressants at all. The persistence parameter of 0.68 implies that 90% of the effect dissipates within six months. Meanwhile, the 0.33 persistence parameter on the bottom level implies that 90% of the business stealing effect of an advertisement dissipates within only two months. This makes advertising in the long run more of a category expansion than a business stealing tool. This is consistent with the common wisdom that antidepressants are subject to a high degree of experimentation. If a patient tries one and finds the side effects unbearable, she might well switch to another one rather than quitting altogether. It is also consistent with a limited memory view of advertising. Since advertising for pharmaceuticals on television usually contain a lot of information about the condition, the mechanisms of action and the side effects and these characteristics are highly correlated within category, a consumer might well remember seeing an advertisement about depression without remembering which brand was advertised. This high persistence at the category level relative to the product level is another source for potential underinvestment in advertising relative to a co-operative.

6 Supply and Counterfactual

6.1 Supply Implications of Positive Spillovers

The demand results above imply that the incentive to invest in advertising is dampened by positive spillovers for two reasons. First, advertising provides benefits to rival firms which are not internalized by the advertising firm. Second, rival advertising lessens the incentive to advertise through the incentivize to free ride on the efforts of rivals.

6.1.1 Internalization

To further illustrate the effects of advertising over time on rivals, consider an impulse response graph in Figure 10. The purpose of this graph is to follow the effect of a marginal dollar per 100 capita spent by Zoloft in January of 2002 on both Zoloft and total market prescriptions for the subsequent year. The top downward sloping curve is the marginal effect of Zoloft advertising on total market prescriptions, while the bottom downward sloping curve is the marginal effect on Zoloft prescriptions. The upward sloping dashed line is the ratio of the total market effect to the Zoloft effect. A marginal dollar per 100 capita of Zoloft advertising would lead to a contemporaneous increase of about 70,000 antidepressant prescriptions, only about 20,000 of which would be for Zoloft. Further, as we follow that effect through time, the effect on the total market is more persistent, and the marginal effect of Zoloft advertising goes more and more to other products. There is a large contemporaneous positive spillover that intensifies through time. Zoloft has no incentive to internalize the benefits it bestows upon other firms, and thus will under invest in advertising relative to a co-operative controlling advertising in the whole market.

6.1.2 Free Riding

To illustrate the free riding incentive, I consider three marginal revenue curves and a horizontal marginal cost curve in Figure 11 given the demand parameters estimated in the previous section, but for a single point in time and for a single product. In the figure, I consider the perspective of Zolofit in January 2002 in the Boston DMA.

The top curve is the marginal revenue for Zolofit if all competitors set advertising equal to zero. Notably far below that curve, the middle curve is the marginal revenue curve of Zolofit if competitors combine to advertise \$3 per 100 capita, which is about the average competitor advertising Zolofit sees in the Boston DMA during the time that it advertises. Finally, the lowest curve depicts the marginal revenue with respect to advertising of Zolofit when its competitors advertise \$10 per 100 capita, about the maximum it ever faces from competitors in the Boston market. Notable from the curves is that the marginal revenue curve of Zolofit takes a significant hit as its competitors advertise more. In fact, when competitors advertise up to \$10 per 100 capita, it is almost not worthwhile for Zolofit to advertise at all. There is a clear incentive for Zolofit to free ride as competitors advertise more and more.

In the next subsection, I will more systematically explore these incentives for our realized antidepressant market.

6.2 Supply Simulation

Given the demand estimation above, it is clear that there is an incentive both for firms to free ride off of their competitors and for firms not to internalize the positive spillovers that their advertising generates for rivals. To explore this systematically, my strategy will be first to solve the Markov Perfect Equilibrium (MPE) given the demand parameters estimated and to assume that the marginal cost of advertising one dollar is one dollar plus a customary 15% service fee that has historically been charged by advertising agencies. That is, I do not assume that firms are already behaving optimally. Rather, I solve for their optimal decisions given an assumption about the marginal costs of advertising and the estimated demand parameters. Having solved this equilibrium, I then compute a counterfactual whereby the advertising firms work together in an advertising co-operative to internalize the positive spillovers and eliminate free riding between them.

One complication of solving a MPE in this setting is that the market conditions change frequently. That is, products enter throughout the sample, both from new innovation and from patent expiration which leads to generic entry. Similarly, with patent expiration, products effectively leave the market (that is, their market shares get very low and they cease to participate in the advertising game). In the demand estimation, these issues are dealt with using product-time fixed effects. In principle, we might view these effects as additional state variables. Since entry and exit are not the main focus of this paper, my strategy will be to focus on a period within my sample where the market is stable. From December 1999 until December 2000, there are no product entries or exits. The next major product entry is July of 2001 when Prozac goes off patent and generic Prozac enters. Indeed, over the year from December 1999 until December 2000, the estimated product-time fixed effects for the advertised products remain relatively constant. Further, since the effects of advertising are almost entirely dissipated over the course of seven months, the infinite horizon dynamic programming problem assuming stable market conditions from December 1999 through December 2000 is a reasonable approximation to one that includes the important market changes in July of 2001.

For the period described, only two products advertise on television: Paxil and Prozac. Paxil is produced by Glaxo-SmithKline, which also produces Wellbutrin and Wellbutrin SR during this period. Prozac is produced by Eli Lilly,

which produces only Prozac in this window. As such, Paxil will advertise in order to maximize discounted future profits over Paxil, Wellbutrin and Wellbutrin SR while Prozac will advertise to maximize discounted future profits for Prozac. During this time period, Prozac has a larger market share than Paxil and, in fact, has the largest market share of any antidepressant. As such, it will have the largest incentive to advertise.

Each firm will solve the Bellman equation (17) in order to best respond to its rival in each period in terms of television advertising. Prices and marginal costs of production are given in the data. For a reference point, the average price of a prescription of Paxil over the course of the period is \$65 while the average production cost is \$3.68. The average price for a prescription of Prozac is \$72 and the production cost is \$3.53. Details of the computation of the equilibrium are in Appendix E. The profit maximization problem is concave, so an equilibrium is found relatively easily. Given the lack of S-shaped demand as in Dube et. al. 2005, pulsing strategies do not arise from the demand parameters.

For the sake of exposition, results are shared for the Atlanta, Georgia DMA. In particular, prior to the period in question, the only advertising done was by Paxil. That is included as an initial condition. Figure 12 shows the evolution of MPE advertising for Paxil and Prozac over the course of the period. Paxil begins the period advertising about \$2.00 per 100 capita, and increases its advertising to about \$3.00 per 100 capita by the third month, holding constant there for the remainder of the period. Prozac advertises just under \$8.00 per 100 capita in the first period and increases advertising up to nearly \$9.00 per 100 capita by the seventh period and holds constant from there. Given these levels of advertising, the category share rises from an initial condition of 4.9% of the population to about 6.7% of the population over the course of the year. Profit evolution for each firm is provided in Figure 14. Eli Lilly's profits increase from \$3.2 million per month in Atlanta in the pre period to about \$4.1 million per month by the end of the period, while GlaxoSmithKline's profits (the combined profits of Paxil, Wellbutrin and Wellbutrin SR) go from about \$2.7 million per month to about \$3.6 million per month. Combined, by the end of the sample, GlaxoSmithKline and Eli Lilly are advertising about \$11 per 100 capita and earning \$7.7 million per month by the end of the period.

Comparing this with what is observed in the data (shown in Figure 13), Prozac advertises far less than would be predicted from the demand model. This could be a result of organizational costs or budget constraints, or simply a matter of learning the game, as DTC was relatively new to the pharmaceutical industry during this period. Paxil, on the other hand, advertises a little more closely to what might be expected, though it is unclear if Paxil would respond differently to a different Prozac strategy. It is notable, however, that in the times when Prozac has a large advertising push, Paxil advertises less, as would be predicted as a response to positive spillovers.

6.2.1 Pulsing Strategies

As is noted by Dube et. al. (2005) and is evident from the figures of observed advertising, firm advertising is often highly variable and unpredictable over time. The model of Dube et. al. rationalizes the spikes and zeros in advertising with an advertising stock function that is s-shaped in current advertising. From other specifications, I found no evidence of S-shaped demand as in Dube et. al. . As such, the supply side equilibrium simulated unsurprisingly does not show any pulses.

6.3 Counterfactual Assumptions

Given the positive spillovers of advertising, we should expect that the incentive to invest in advertising is lessened by both a failure to internalize the benefits of advertising on rivals and by an incentive to free ride. As such, I will

consider two counterfactual scenarios. First, I assume that Eli Lilly and GlaxoSmithKline co-operate on advertising that maximizes the combined profits of their products, which takes away the need for strategic response. Next, I'll consider a scenario whereby the entire market is allowed to set advertising in a single optimization problem to think about the potential of a category-wide advertising co-operative. For such a scenario to work, cooperation would need to be not only allowed, but enforced in some way. Antidepressants (and other pharmaceutical products) might be able to cooperate similarly through non-profit organizations called patient advocacy groups. Such groups are focused on educating patients on specific diseases and treatments. The specific mission of these types of organizations is to educate patients on how, when and why to seek treatment for various health care needs. Some examples of these types of organizations are the American Cancer Society, the Alzheimer's Association and the Depression and Bipolar Support Alliance. While they do not tend to advertise on television currently, they might be an ideal facilitator for category level advertising of antidepressants. As many physicians see depression as an under-treated condition, commercials that encouraged patients to seek their physicians' advice if they encounter the symptoms of depression might be allowed by regulators. The ability to contract on coordination would allow firms to overcome the free riding problem and provide advertising, even without a brand level component.

For the purposes of the counterfactuals, I assume that the advertising firms in the antidepressant market cooperate to make a common non-branded category advertisement for antidepressants, facilitated by a patient advocacy group. The effect of those advertisements is equal to the category level effect of the branded advertisements estimated above.

The co-operative solves the firm's problem in each month and in each market. In the Eli Lilly/GSK counterfactual, the co-operative includes Paxil, Prozac, Wellbutrin and Wellbutrin SR as part of the portfolio. In the full co-operative counterfactual, all products in the category are included as part of the portfolio in the firm's problem. The marginal co-operative advertisement dollar has cost equal to \$1.15, as before. The co-operative solves the Bellman equation, but since all products in the market are included in Φ_f , strategic response is not necessary.

While a monopoly may wish to raise prices with new advertising, I assume that the co-operative does not. I do this for two reasons. First, the thought experiment in this scenario is a cooperation only on advertising, not a full monopoly situation. Two firms cooperating on market expansive advertising but competing on other dimensions are unlikely to want to raise their prices- indeed they might want to lower their prices. Given that, we might be worried that I am not allowing firms to lower their prices in the counterfactual. However, an important industry detail is discovered in the data. In 1999, we observe the antidepressant market going from zero television advertising to a significant amount of television advertising. This switch had no effect on product prices. Indeed, in Appendix A I show that prices are unresponsive to DTC and that brand prices are very tightly predicted by a product fixed effect and a time trend. This is likely because most price flexibility happens in negotiations with managed care organizations which were less prevalent in the time of this sample and those negotiations happen infrequently and in a disjointed way from the advertising decision. So while DTC may respond to markups in the model, prices do not respond to DTC conditional on product-market and market-time effects. This is consistent with research (Aitken et. al. 2013) that show very stable brand pricing along a trend line.

Detailing, on the other hand, could be more troublesome. If detailing were mainly a business stealing device, as is suggested in Narayanan et. al. (2004), it would be complementary to co-operative DTC. A firm would be happy for the co-operative to make the market large and it would then try to compete over brand share in detailing space. Anticipating this, the co-operative would shade down its advertising, anticipating that profits would be competed away with detailing. If detailing were, alternatively, also mainly category expansive, it would be a substitute for DTC. In response to the DTC, firms would be glad to reduce detailing and let the co-operative foot the bill for market expansion.

The co-operative would anticipate this and shade up its advertising. To address this question, I go to the physician specific detailing data from 2001-2003. Particulars of the detailing analysis may be found in Appendix A. Detailing is not statistically or economically significantly correlated with DTC at the DMA level. In the competitive world, even those firms which do not advertise on television are not ramping up detailing in response to large television advertising campaigns by rivals. While it is difficult to predict exactly what firms would do under the co-operative advertising scenario, it appears from the data that the detailing strategy part of the firm operates reasonably independently of the DTC part. As the main purpose of the counterfactuals is to illustrate the incentive effects of positive spillovers on DTC decisions, I will assume that firms will continue not to adjust detailing to category DTC increases or decreases, just as is observed in the realized world.

Given the fact that detailing and prices are observed not to be affected by DTC in the data, the counterfactual world described has not been observed, and it is still possible that under this scenario, firms would be more likely to re-visit how they decide on detailing and pricing in response to DTC. To the extent that I cannot provide any insight on these dimensions, that is a limitation of these counterfactual experiments. However, the counterfactuals continue provide a useful illustration of the direction of the incentive effects of positive spillovers.

Also worth noting, since in the observed world firms under invest in advertising, the counterfactual results in some values of advertising that are not observed in the sample. Since we only estimate the demand curve for values of advertising within the sample, the out-of-sample specific numbers are driven by the assumed functional form of the advertising effect in the model.⁹ While differences in the level of the effect at these points will affect the magnitude of the conclusion (exactly how much more the co-operative will advertise), it will not affect the directionality, as long as advertising still has a positive category effect.

6.4 Advertising

As the business stealing incentive grows, observed total advertising is expected to increase relative to the counterfactual advertising. As the business stealing incentive dwindles, the free riding incentive associated with the positive spillovers should lead to lower observed advertising relative to the co-operative's ideal. As in the demand estimation, the business stealing effects of advertising are swamped by the category expansive effects, cooperation should lead to an increase in total category advertising.

For illustrative purposes Figure 15 shows the computed Markov Perfect Equilibrium versus co-operative advertising choice in the Atlanta DMA in the co-operative scenarios described above, with the left panel being the Eli Lilly/GSK co-operative and the right panel being the full category advertising co-operative. The Eli Lilly/GSK co-operative advertises about 50% more in total than the sum of Prozac and Paxil in the competitive equilibrium. The full category co-operative advertises almost four times as much as the competitive equilibrium.

6.5 Quantities and Profits

Figure 16 illustrates the difference in profits between the MPE and the co-operative situations. In the left panel, we see that in the Eli Lilly/GSK co-operative, the included products see roughly 10% higher profits than they do combined

⁹Alternative specifications with quadratic and square root functions of advertising were tried and produced qualitatively and quantitatively very similar results.

in the MPE. In the right panel, we see that the full category profits increase by about 14% by the end of the period. Figure 17 compares the category share of the population in the Atlanta DMA between the co-operatives and the MPE. In the Eli Lilly/GSK co-operative the category is about 6% larger than it is under MPE. In the full co-operative, the category is almost 18% larger by the end of the period than it is under MPE.

6.6 Discussion

While market expansion and increasing profits in the counterfactual would be viewed as welfare increasing in many consumer goods markets, there are a few reasons for us to take caution in drawing conclusions about social welfare in the context of antidepressants. First, many prescriptions are covered by insurance. While many people are getting prescribed and incurring minimal if any cost, the insurance system pays out a significant price. It is possible that the new prescriptions are not justified by the societal cost. However, as many physicians see depression as an under treated condition, the total welfare benefits could also be very high. Second, if I have missed important price or detailing complementarities, it might be the case that all increased profits are competed away after the co-operative sets higher advertising. If this is the case, the welfare effect is also ambiguous. However, as neither prices nor detailing are not correlated with advertising expenditures conditional on product effects, this concern might not be very large.

The balance of societal costs and benefits, while very interesting and important, is not identified in this study and is certainly worthy of further research. Category expansion is typically thought of as consumer welfare improving, as consumers would not purchase something for which the costs exceeded the benefit. However, in health care, many levels of agency mask the costs and benefits of treatments. As such, this study remains agnostic on consumer welfare.

7 Conclusions

Using data from the antidepressant market and an identification strategy taking advantage of both policy and spatial discontinuities, I find that television advertising has significant positive spillovers. I construct and estimate a model to systematically explore this fact and its implications on the supply decisions of firms. In particular, I find that the spillovers induce a free-riding and internalization problem whereby observed advertising is significantly lower than the optimal strategy that a co-operative would set if it controlled the entire market. If the advertising firms worked together, they would advertise significantly more, increase the size of the category by 6% and their own profits by 10%. Meanwhile, a full industry co-operative would set advertising four times as high as is observed in equilibrium and would increase industry shares by 18% and profits by 14%.

These findings also speak to some of the controversy surrounding the practice of advertising pharmaceuticals on television. Contrary to some of the criticism, this type of advertising drives consumers into the market and helps all products in the category, including the low cost generics. It is the proverbial rising tide that lifts all ships. While there is a brand effect, it is short lived while the category expansion effect is persistent. Especially for conditions that are seen as under treated, this type of advertising could be beneficial.

These findings are potentially relevant to firms, regulators, econometricians and marketers. Firms might be able to realize gains from cooperation that might be allowed by regulators. In the absence cooperation, it is important for

firms to properly take account of spillovers when deciding advertising policy. Regulators should take into account that content regulation might reduce or eliminate the firms' incentives to advertise. Finally, it is important for marketers and econometricians to consider the possibility of positive spillovers when building models of advertising impacts on supply and demand.

References

- [1] Aitken, M.L., Berndt, E.R., Bosworth, B., Cockburn, I.M., Frank, R., Kleinrock, M., & Shapiro, B.T. (2013). The regulation of prescription drug competition and market responses: patterns in prices and sales following loss of exclusivity. NBER Working Paper #19487.
- [2] Anderson, E. T., & Simester, D. (2013). Advertising in a competitive market: the role of product standards, customer learning, and switching costs. *Journal of Marketing Research*, 50(4), 489-504.
- [3] Anderson, I., (2000). Selective serotonin reuptake inhibitors versus tricyclic antidepressants: a meta-analysis of efficacy and tolerability. *Journal of Affective Disorders*, 58, 19-36.(1)
- [4] Avery, R. J., Eisenberg, M., & Simon, K. I. (2012). The impact of direct-to-consumer television and magazine advertising on antidepressant use. *Journal of Health Economics*, 31(5), 705-718.
- [5] Bagwell, K. (2007). *The economic analysis of advertising*. Handbook of industrial organization, 3, 1701-1844.
- [6] Bayer, P., Ferreira, F., & McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4), 588-638.
- [7] Berndt, E. R. (1991). *The practice of econometrics: classic and contemporary*. Reading, MA: Addison-Wesley.
- [8] Berndt, E. R., Bui, L., Reiley, D. R., & Urban, G. L. (1995). Information, marketing, and pricing in the US antiulcer drug market. *The American Economic Review*, 85(2), 100-105.
- [9] Berndt, E. R., Bui, L., Reiley, D. R., & Urban, G. L. (1997). The roles of price, quality, and marketing in the growth and composition of the antiulcer drug market. *The economics of new goods* (NBER Studies in Income and Wealth, vol. 58), Bresnahan, T. F., & Gordon, R. J. (Eds.). (1997), (Vol. 58). University of Chicago Press., pp. 277-322.
- [10] Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2), 242-262.
- [11] Black, S. (1999). Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics*, 114(2), 577-599.
- [12] Busch, S., Golberstein, E., & Meara, E. (2011). The FDA and ABCs: The unintended consequences of antidepressant warnings on human capital. NBER Working Paper # 17426.
- [13] Card, D., & Krueger, A. (1994). Minimum wages and employment: a case study of the New Jersey and Pennsylvania fast food industries. *American Economic Review*, 84(4), 772-793.
- [14] Ching, A. (2010). A dynamic oligopoly structural model of the prescription drug market after patent expiration. *International Economic Review*, 51(4), 1175-1207.

- [15] Ching, A., Erdem, T., & Keane, M.P. (2009). The price consideration model of brand choice. *Journal of Applied Econometrics*, 24(3), 393-420.
- [16] Ching, A., & Ishihara, M. (2012). Measuring the informative and persuasive roles of detailing on prescribing decisions. *Management Science* 58(7), 1374-1387.
- [17] Dickstein, M. (2014). Efficient provision of experience goods: evidence from antidepressant choice. Stanford University working paper.
- [18] Donohue, J. M., & Berndt, E. R. (2004). Effects of direct-to-consumer advertising on medication choice: The case of antidepressants. *Journal of Public Policy & Marketing*, 23(2), 115-127.
- [19] Donohue, J. M., Berndt, E. R., Rosenthal, M., Epstein, A. M., & Frank, R. G. (2004). Effects of pharmaceutical promotion on adherence to the treatment guidelines for depression. *Medical Care*, 42(12), 1176-1185.
- [20] Dorfman, R., & Steiner, P. O. (1954). Optimal advertising and optimal quality. *The American Economic Review*, 44(5), 826-836.
- [21] Dube, Arindrajit, T. William Lester, and Michael Reich. "Minimum wage effects across state borders: Estimates using contiguous counties." *The Review of Economics and Statistics* 92.4 (2010): 945-964.
- [22] Dubé, J. P., & Hitsch, G. J., & Manchanda, P. (2005). An empirical model of advertising dynamics. *Quantitative Marketing and Economics*, 3(2), 107-144.
- [23] Duggan, M., & Morton, F. M. S. (2006). The distortionary effects of government procurement: evidence from Medicaid prescription drug purchasing. *The Quarterly Journal of Economics*, 121(1), 1-30.
- [24] Ellison, G., & Ellison, S. F. (2011). Strategic entry deterrence and the behavior of pharmaceutical incumbents prior to patent expiration. *American Economic Journal: Microeconomics*, 3(1), 1-36.
- [25] Ellison, S. F., & Wolfram, C. (2006). Coordinating on lower prices: pharmaceutical pricing under political pressure. *The RAND Journal of Economics*, 37(2), 324-340.
- [26] Filson, D. (2012). A markov-perfect equilibrium model of the impacts of price controls on the performance of the pharmaceutical industry. *The RAND Journal of Economics*, 43(1), 110-138.
- [27] Goeree, M. S. (2008). Limited information and advertising in the US personal computer industry. *Econometrica*, 76(5), 1017-1074.
- [28] Holmes, Thomas J. "The effect of state policies on the location of manufacturing: evidence from state borders." *Journal of Political Economy* 106.4 (1998): 667-705.
- [29] Iizuka, T. (2004). What explains the use of direct-to-consumer advertising of prescription drugs?. *The Journal of Industrial Economics*, 52(3), 349-379.
- [30] Iizuka, T., & Jin, G. Z. (2005). The effect of prescription drug advertising on doctor visits. *Journal of Economics & Management Strategy*, 14(3), 701-727.
- [31] Iizuka, T., Jin, G. Z. (2005). Direct to consumer advertising and prescription choice. Working Paper Available at SSRN 700921.

- [32] Jayawardhana, J. (2013). Direct-to-consumer advertising and consumer welfare. *International Journal of Industrial Organization*, 31(2), 164-180.
- [33] Kravitz, R. L., Epstein, R. M., Feldman, M. D., Franz, C. E., Azari, R., Wilkes, M. S., Hinton, L., & Franks, P. (2005). Influence of patients' requests for direct-to-consumer advertised antidepressants. *JAMA: the journal of the American Medical Association*, 293(16), 1995-2002.
- [34] Lewis, R., Nguyen, D., (2012). Wasn't that ad for an iPad? display advertising's impact on advertiser- and competitor-branded search. Yahoo Research Working Paper.
- [35] Lewis, R. A., & Rao, J. M. (2012). On the near impossibility of measuring advertising effectiveness. Yahoo Research Working paper.
- [36] Liu, Q. & Gupta, S. (2011). The impact of direct-to-consumer advertising of prescription drugs on physician visits and drug requests: empirical findings and public policy implications. *International Journal of research in Marketing*, 28, 205-217.
- [37] Liu, Q., Gupta, S., Venkataraman, S., & Liu, H. (2013). An empirical model of drug detailing: dynamic competition and policy implications. Purdue University working paper.
- [38] Manchanda, P., Rossi, P., & Chintagunta, P. (2004). Response modeling with non-random marketing mix variables. *Journal of Marketing Research*, 41(4), 467-478.
- [39] Murphy, M. J., R. L. Cowan, and L. I. Sederer (2009): *Blueprints: Psychiatry*, Philadelphia, PA: Lippincott Williams Wilkins, 5 ed.
- [40] Narayanan, S., Desiraju, R., & Chintagunta, P. K. (2004). Return on investment implications for pharmaceutical promotional expenditures: The role of marketing-mix interactions. *Journal of Marketing*, 68(4), 90-105.
- [41] Nerlove, M., & Arrow, K. J. (1962). Optimal advertising policy under dynamic conditions. *Economica*, 29(114), 129-142.
- [42] Rosenthal, M. B., Berndt, E. R., Donohue, J. M., Epstein, A. M., & Frank, R. G. (2003). Demand effects of recent changes in prescription drug promotion. Henry J. Kaiser Family Foundation.
- [43] Sahni, N. (2013). Advertising spillovers: field experimental evidence and implications for the advertising sales-response curve. Stanford University working paper.
- [44] Stremersch, S., Landsman, V., & Venkataraman, S. (2012). The relationship between DTCA, drug requests, and prescriptions: uncovering variation in specialty and space. *Marketing Science*, 32(1), 89-110.
- [45] Wosinska, M. (2002). Just what the patient ordered? Direct-to-consumer advertising and the demand for pharmaceutical products. Direct-to-Consumer advertising and the demand for pharmaceutical products (October 2002). HBS Marketing Research Paper, (02-04).
- [46] Wosinska, M. (2005). Direct-to-consumer advertising and drug therapy compliance. *Journal of Marketing Research*, 42(3), 323-332.

Tables and Figures

Figure 1: Antidepressant Commercials Relative to FDA Memo



Figure 2: Antidepressant Revenues 1996-2008




Table 1: Descriptive Statistics: Advertising

	Mean	Q25	Median	Q75
DTC per 100 capita	0.782	0	0	1.358
Subcategory DTC per 100 Capita	2.012	0	1.496	3.505
Category DTC per 100 Capita	4.035	2.284	3.515	5.534
	Mean	Q25	Median	Q75
DMAs	101			
DMA Population	2340774	903090	1469823	2622567

Figure 3: Variation Across Three Markets in Advertising

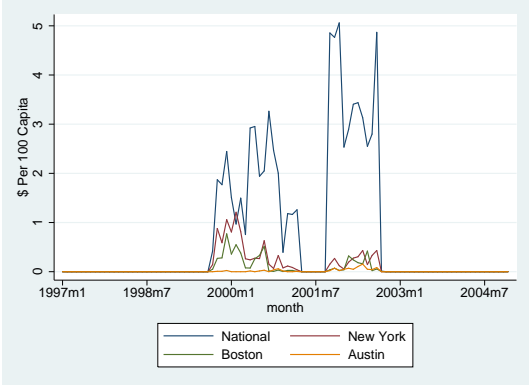


Figure 4: Two Different Advertising Measures: National




Figure 5: Quantity Weighted Average Margins in Boston




Figure 6: Full Sample: Top 101 DMAs

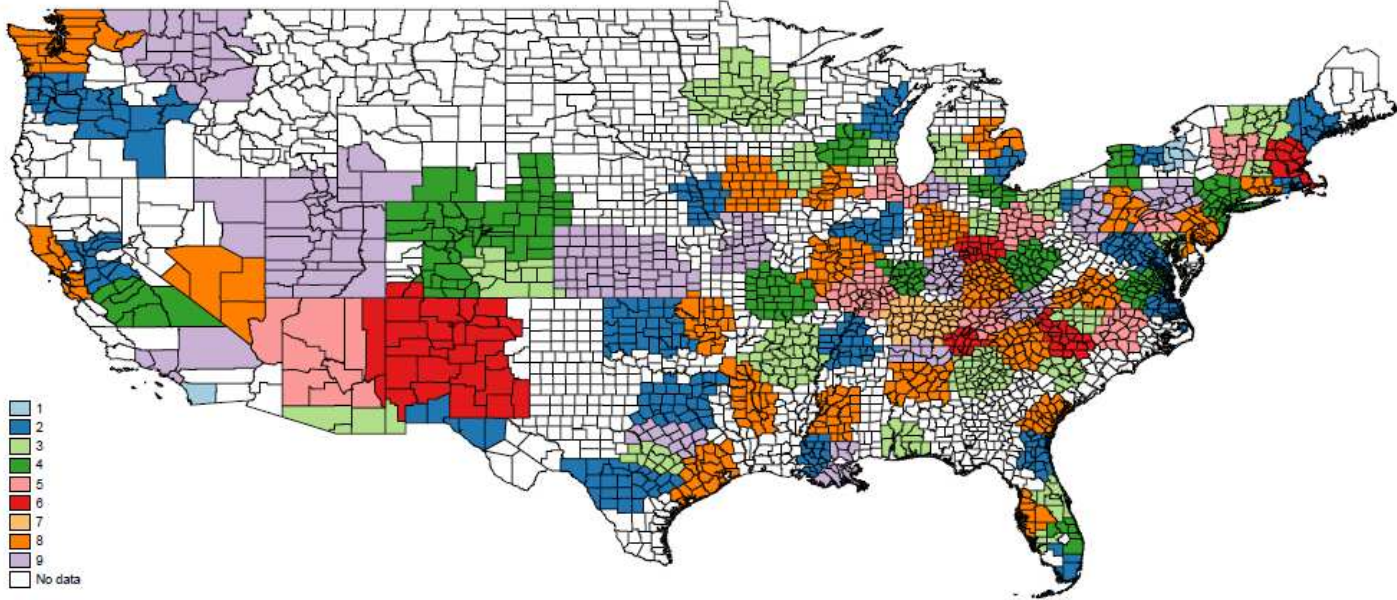


Figure 7: Ohio and Its DMAs




Figure 8: Border Sample: Counties on the Borders of the Top 101 DMAs




Figure 9: Variation in Log DTC Net of Fixed Effects, 14% Zeros




Table 2: The Effect of Own and Rival Advertisements on Sales

VARIABLES	(1) $\log(Q)$
lagged $\log(Q)$	0.334*** (0.00746)
DTC	0.0240*** (0.00621)
DTC^2	-0.00216* (0.00113)
DTC_{rival}	0.0164*** (0.00266)
DTC_{rival}^2	-0.000938*** (0.000252)
$DTC \times DTC_{rival}$	-0.00134** (0.000631)
Product-Border-Time	yes
Product-Border-DMA	yes
Observations	316,428
R-squared	0.955

DMA clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3: Descriptive Statistics: Border Sample, 1997-2003

	Mean	Q10	Median	Q90
Number of Border Experiments	153			
Number of DMAs	97			
$LOGDTC_{product}$	0.190	0	0	1.033
$LOGDTC_{nest}$	0.447	0	0	1.682
$LOGDTC_{market}$	0.817	0	0.921	1.987

$LOGDTC$: log of one plus dtc expenditures per 100 capita

All are defined at the experiment-DMA-month level

Table 4: Results of Base Model

VARIABLES	Category Level	Subcategory Level	Product Level
$adstock$	0.0472* (0.00665)	0.0093 (0.00719)	0.0223* (0.00756)
$persistence, \lambda$	0.684* (0.0299)	0.282* (0.0116)	0.330* (0.0140)
Observations	23,091	93,284	140,687
R-squared	0.948	0.935	0.960

DMA clustered standard errors in parentheses

* p<0.01

Figure 10: Impulse Response Effect of Zoloft Advertisement on Own and Total Prescriptions




Table 5: Short Run Advertising Elasticities

Products	Paxil	Paxil CR	Prozac	Prozac Weekly	Wellbutrin SR	Wellbutrin XL	Zoloft	Outside Option
Paxil	0.037	0.019	0.021	0.019	0.020	-	0.021	-0.023
Paxil CR	0.016	0.029	0.016	0.016	0.012	0.010	0.016	-0.015
Prozac	0.0092	-	0.020	0.0080	0.0097	-	0.0092	-0.011
Prozac Weekly	0.0088	-	0.0088	0.018	0.0068	-	0.0088	-0.0080
Wellbutrin SR	0.014	-	0.014	0.012	0.021	-	0.014	-0.014
Wellbutrin XL	0.017	0.017	0.017	0.017	0.019	0.035	0.017	-0.018
Zoloft	0.013	0.013	0.013	0.013	0.013	0.010	0.027	-0.015

Figure 11: Marginal Revenue Curves Under Various Scenarios

Marginal Revenue of Advertising for Zoloft




Figure 12: MPE Simulation Advertising - Atlanta DMA




Figure 13: Realized Advertising - Atlanta DMA

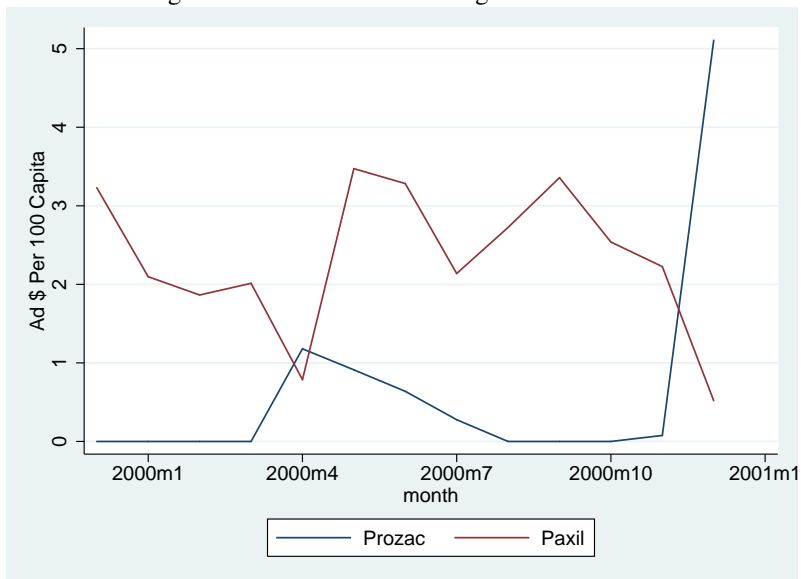


Figure 14: MPE Simulation Profits - Atlanta DMA




Figure 15: Paxil/Prozac and Full Co-operative versus Competitive Total Advertising - Atlanta DMA




Figure 16: Paxil/Prozac Co-operative versus Competitive Total Profits - Atlanta DMA





Figure 17: Paxil/Prozac Co-operative versus Competitive Category Share - Atlanta DMA



Appendix A - Orthogonality of Pricing and Detailing Decisions

A.1 Pricing

The regression of television commercials on prices is shown in Table 6. As can be seen the point estimate is very small and insignificant with respect to both own and cross advertising. Just for perspective, the average unit price of a branded drug is about \$3.60 over the course of the sample and the average DTC per capita of those drugs that advertise over the course of the sample is \$0.0055. Raising DTC per capita by \$0.01 is associated with a price decrease of

Table 6: Predicting Prices with Advertising

VARIABLES	(1) realunitprice
<i>DTC</i>	-0.0518 (0.0437)
<i>DTC_other</i>	-0.0332 (0.0248)
time	0.00963*** (0.00313)
expired	-0.210 (0.149)
Product FEs	yes
Observations	1188
R-squared	0.994

Product Clustered Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

\$0.03. This is very small economically. Also interesting is that just a time trend, a product fixed effect and a dummy for patent expiration can explain prices with R squared bigger than 0.99. Prices seem quite sticky, especially relative to advertising in this market.

A.2 Detailing

To address the potential omitted variables bias problem associated with physician detailing, I acquired data from ImpactRx that follows a panel of physicians through time, measuring the amount of detailing time was spent on them by drug. For general practitioners, the panel begins in 2001 and for psychiatrists, the panel begins in 2002. As such, I am able to match this data with a portion of the data used in the main analysis to test the robustness of the identification to the inclusion of detailing. Detailing is measured in minutes per product per physician per month.

The first strategy employed to assess the potential for bias is to aggregate detailing to the DMA level. Since DTC advertising is decided at the DTC level and no more local, I want to see if detailing decisions are sensitive to DTC decisions at the DMA level. I regress product detailing minutes on own local DTC and rival local DTC, measured in dollars per hundred capita. Similar to the main specification, I include product-quarter and product-DMA fixed effects. I want to know within time period and within DMA, is higher DTC for one product associated with higher detailing? The results of the regression are presented in Table 9. As it turns out, there is no statistically significant relationship between either own or rival DTC and detailing decisions. This is consistent with conversations with managers who said that detailing and DTC typically have different managers making decisions and coordination very difficult. While the estimates are not statistically significant, it is worth noting that even if they were, the magnitudes of the point estimates are very small. The standard deviation of local DTC campaigns is 0.17. Taken at face value, an increase in own local DTC of one standard deviation is associated with an increase in detailing in the entire DMA by twenty seconds. The point estimate on rival advertising is even smaller and negative. This suggests that widespread strategies to undercut rival DTC with detailing campaigns do not seem to be happening.

Table 7: Are Detailing Campaigns Coordinated with DTC?

VARIABLES	(1) Detailing Minutes
Local DTC	2.081 (1.605)
Rival Local DTC	-0.284 (0.568)
Product-DMA FEs	yes
Product-Quarter FEs	yes
Observations	22,589
R-squared	0.673

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 8: Results of Limited Model Including Geographic Detailing from 2001 On

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0370*** (0.00619)	0.0194** (0.00922)	0.0310*** (0.0108)
<i>logdetail</i>	0.00227 (0.00464)	-0.00179 (0.00388)	0.00358* (0.00216)
<i>persistence</i>	0.481*** (0.0723)	0.148*** (0.0247)	0.280*** (0.0183)
Observations	4,903	21,975	41,539
R-squared	0.967	0.951	0.972

DMA clustered standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

The second strategy used to think about this problem is to directly incorporate the detailing data into the main analysis. Detailing data was aggregated to the product-border-DMA-time level as with the main specification. Since the physician panel in the detailing data is smaller in both number of physicians and in length of time, the matched sample will be substantially smaller and over a shorter time frame than the main specification. As such, there is no longer sufficiently large T to ensure that the lagged dependent variable Nickel bias is small. The persistence parameters will all be biased downward in these specifications. The main goal of this exercise is to see if the coefficients on DTC are robust to the inclusion of detailing.

The effect of including detailing in this model is demonstrated in the next four regressions. The first two include all matched data from 2001 onwards. The regression in Table 8 includes the log of detailing minutes while the regression in Table 9 does not. As we can see, the effects of DTC are not significantly affected by the inclusion of detailing. As expected, the persistence parameters are all much smaller than in the main model due to the smaller number of time periods.

The second two regressions include all matched data just from 2002 onwards. This is to address the potential concern that in 2001 the detailing data is coming only from general practitioners while from 2002 onwards, psychiatrists are also included. Table 10 shows the results from the regression including log detailing from the years in which there is

Table 9: Results of Limited Model Without Geographic Detailing from 2001 On

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0375*** (0.00612)	0.0190** (0.00911)	0.0312*** (0.00108)
<i>persistence</i>	0.481*** (0.0723)	0.148*** (0.0247)	0.280*** (0.0183)
Observations	4,903	21,975	41,539
R-squared	0.967	0.951	0.972

DMA clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 10: Results of Limited Model Including Geographic Detailing from 2002 On

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0351*** (0.00806)	0.00474 (0.0180)	0.00958 (0.0103)
<i>logdetail</i>	0.00535 (0.00355)	-0.00399 (0.00449)	0.00362 (0.00248)
<i>persistence</i>	0.257*** (0.0967)	0.207*** (0.0269)	0.124*** (0.0205)
Observations	2,742	12,394	24,464
R-squared	0.972	0.954	0.983

DMA clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

data for both general practitioners and psychiatrists. Table 11 omits detailing. Once again, the persistence parameters shrink, as the length of the panel has again shortened. Also, at all three levels, the DTC parameters are not significantly affected by the inclusion of detailing.

Table 11: Results of Limited Model Without Geographic Detailing from 2002 On

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0359*** (0.00819)	0.00388 (0.0178)	0.00973 (0.0103)
<i>persistence</i>	0.256*** (0.0967)	0.209*** (0.0269)	0.123*** (0.0205)
Observations	2,742	12,394	24,464
R-squared	0.972	0.954	0.983

DMA clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

These two exercises show that over the course of the time period where detailing data is available in this class, detailing is not significantly correlated with DTC at the DMA level. Further, including detailing in the main model does not change the main parameters of interest. While this paper makes no claims of identifying the causal effects of detailing, it seems as though it poses a minimal threat to the identification of the causal effect of DTC using the border strategy.

Table 12: Results of Base Model with Placebo

VARIABLES	Category Level	Subcategory Level	Product Level
<i>OTCSLEEPAD</i>	0.0000565 (0.000397)	0.000471 (0.000482)	0.00218* (0.00113)
<i>persistence</i>	0.684*** (0.0304)	0.280*** (0.0116)	0.323*** (0.0139)
Observations	22,826	93,280	147,443
R-squared	0.948	0.935	0.960
DMA clustered standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1			

Appendix B - Placebo Test

With a difference-in-differences model, the assumption of parallel trends in the outcome variable absent the treatment is required for a valid estimation. One way to assess the validity of this assumption is through the use of a placebo test. In this case, I will use advertising for over-the-counter (OTC) sleep aids as a placebo treatment. This is an ideal placebo for two reasons: first, it varies at the same level as antidepressant advertising- at the DMA month. Next, OTC sleep aids need not be prescribed by a physician, so we should not expect a “going to the doctor” effect of advertising to be present in OTC advertising. I will use the same identification strategy, but I will also include OTC sleep aid advertising as a treatment. The results are below. None of the coefficients on OTC sleep aid advertising is statistically significant at the 5% level or economically important at any level.

Appendix C - Alternative Sample Selection

One might worry that the effect of advertising in the border sample counties differs systematically from the non-border counties or that the effect of advertising in rural areas would be much different than the effect of advertising in urban areas. To think about these concerns, I have repeated the analysis with several alternative sample selections.

C.1 The Urban Rural Divide

C.1.1 Without the Northeast Corridor and Other Urban Areas

It seems the less urban areas show very similar results to the full border sample. The category and product level effects are larger, but not significantly different from the full sample of borders.

C.1.2 Only the Urban Border Counties

It seems the effects of advertising are a bit smaller in the more urban areas, except at the nest level. However, it seems more likely that the identifying assumption might fail in the more urban areas, as the borders are much closer to the central cities than in the more rural borders.

Table 13: Results of Base Model without Northeast Corridor and Other Urban Areas

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0508*** (0.00825)	0.00874 (0.00904)	0.0279*** (0.00939)
<i>persistence</i>	0.689*** (0.0322)	0.255*** (0.0129)	0.313*** (0.0159)
Observations	17,400	67,976	99,101
R-squared	0.946	0.927	0.955

DMA clustered standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 14: Results of Base Model with only Northeast Corridor and Other Urban Areas

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0348*** (0.00951)	0.0191** (0.00940)	0.00707*** (0.0111)
<i>persistence</i>	0.648*** (0.0678)	0.387*** (0.0215)	0.381*** (0.0284)
Observations	5,775	25,550	41,586
R-squared	0.963	0.956	0.972

DMA clustered standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

C.2 Are the Borders Special?

C.2.1 Full Sample Estimation

Next, I assume that the border sample approach is not necessary and that advertising may be viewed as predetermined with respect to demand shocks. This might be plausible since a large fraction of television advertising in the pharmaceutical market is purchased in the upfront market. Using a difference-in-differences approach with a lagged dependent variable and a common trend for all DMAs, I estimate the model. The point estimates for the effects of advertising are not statistically different from those in the border approach. However, the persistence parameters are larger. If we believe these persistence parameters over the ones estimated in the border approach, the spillover problem will be even larger than was estimated. However, there is reason to believe that these persistence parameters are overstated if firms are using rule of thumb strategies, as will be explained in the next subsection. I continue to prefer the border approach, as the assumptions of common trends among similar geographies is more plausible than a nationally common trend. In addition, it gives more modest estimates of the size of the spillover, working against, my main finding of spillover effects.

C.2.2 “Anti”-Border Sample

Here, I use all of the counties that are not in the border sample. For this sample, endogeneity should be the most severe. The results of this estimation should reveal the main drivers of the endogeneity problem that necessitates use of the

Table 15: Results of Base Model Without Using the Border Approach

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0436*** (0.00140)	0.0030 (0.00248)	0.0187*** (0.00438)
<i>persistence</i>	0.747*** (0.0209)	0.721*** (0.0124)	0.595*** (0.0116)
Observations	8,466	42,181	73,628
R-squared	0.976	0.984	0.984

DMA clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 16: Results of Base Model Using the “Anti”-Border Approach

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0508*** (0.00164)	0.0049 (0.00251)	0.0182*** (0.00439)
<i>persistence</i>	0.751*** (0.0259)	0.721*** (0.0128)	0.594*** (0.0115)
Observations	8,466	42,181	73,628
R-squared	0.981	0.984	0.984

DMA clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

border sample. If the main effects of advertising are much larger in this estimation, it would lead us to believe that advertising was strongly correlated with unobserved local demand shocks such as weather shocks, employment shocks or special seminars that are unobserved. These shocks, if positively correlated with advertising decisions would lead us to falsely conclude that the effect of advertising is high. If, on the other hand, the reverse causality concern- where firms target an advertising to sales ratio based on sales in the previous period, we would expect that in the anti-border sample the persistence parameter would be much higher. That is, lagged shares would not only be measuring the persistence of prescribing associated with advertising carry-over, they would also be controlling for the rule of thumb decision of firms to target DMAs that had large shares in the previous period. Results of the estimation are in Table 14. As can be seen, while the point estimates on advertising do not significantly change from the border sample, they are larger, and the persistence parameters are much (and significantly) larger. This suggests that the main source of endogeneity is reverse causality- if firms are following a rule of thumb advertising to sales ratio based on previous period sales, the lagged share variable will be biased upward as it controls for the firm’s rule of thumb rather than estimating persistence.

C.3 Removing the State Borders

As described in the text, if tax rates cause a selection problem of physicians into DMAs, then results could be biased. To assess this possibility, DMA borders which coincide with state borders are removed from the sample. That removes roughly one third of the borders in the sample. The analysis is re-run using the set of DMA borders that do not coincide

Table 17: Results of Base Model Without State Border Overlap

VARIABLES	Category Level	Subcategory Level	Product Level
<i>adstock</i>	0.0376*** (0.00764)	0.0131* (0.00783)	0.0256*** (0.00881)
<i>persistence</i>	0.683*** (0.0368)	0.271*** (0.0133)	0.327*** (0.0166)
Observations	15,975	65,097	97,720
R-squared	0.981	0.984	0.984

DMA clustered standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

with state borders. Results are available in Table 17. The results are very similar and not statistically distinguishable from those in the full sample.

Appendix D - Primary Care Service Areas

As mentioned in section (3.1.1), measurement error could bias my estimates towards zero if patients are going to the doctor in different counties than where they watch television advertisements. The Dartmouth Center for Health Policy Research has developed a Primary Care Service Area (PCSA) project which is the first national database of primary care resources for small areas. These areas were defined using Medicare claims data from 1999 and Census data from 2000. The service areas include a ZIP area with one or more primary care providers and any bordering ZIPs where the population largely gets their primary care from those physicians.

This database allows me to ask how many patients travel across DMA borders to seek their primary care. In particular, I can match this data to my prescribing data at the ZIP level. I can then see what percentage of each PCSA falls into a single DMA. Doing this I find that only about 1% of PCSAs cross DMA borders at all. Of those that do cross DMA borders, they do so only minimally. That is, the DMA holding the majority of a PCSA which crosses a border on average contains 97% of that PCSA. As such, measurement error bias should be minimal.

Appendix E - Computational Details

E.1 MPE Simulation

I solve the Markov Perfect Equilibrium (MPE) of the advertising game using a policy iteration algorithm, similar to Dube et. al. 2005. I take an initial guess of the strategy profile $\sigma^0 = (\sigma_1^0, \dots, \sigma_J^0)$, and then take the following steps:

1. For the strategy profile σ^n , calculate the value functions for each of the J firms, V_j^n . The value functions are defined by the Bellman equation (17). In this step, the maximization on the right hand side is not carried out, but the current guess of σ^n is plugged in.

2. For all $n > 0$, check whether the value functions and policy functions satisfy pre-determined convergence criteria, $\|V_j^n - V_j^{n-1}\| < \varepsilon_V$ and $\|\sigma_j^n - \sigma_j^{n-1}\| < \varepsilon_\sigma$. If so, stop the process.
3. Update each firm's strategy using the Bellman equation (17). Do this by carrying out the optimization on the right hand side. Denote the resulting policies and value functions by σ^{n+1} and V^{n+1} . Return to step 1.

The value functions are represented on a grid. As in the supply simulation, there are two firms, there are two firm advertising stocks and a category advertising stock. That makes three state variables. The grid is 5x5x5. Outside of the grid, the value function is determined using bilinear interpolation.

For the two-firm co-operative counterfactual, the MPE simulation turned into basically a monopoly game simulation where both firms' profits are considered and no other firms advertise. For the full co-operative counterfactual, computation is the same, only the advertising firm considers the profits of all products on the market.

Using the derived policies in each of these computations, the equilibria are simulated forward given the market initial conditions to generate the figures.

Appendix F - Selection of Households and Physicians to Counties

F.1

It might present a problem if households and physicians pre-disposed to be more affected by advertisements systematically select into border counties on the side of the border that receive more advertisements. This is a difficult level of selection to address. However, I do check for balance in characteristics across the borders. First, I check to see if physicians are more likely to bunch on the "high DTC" side of the border. To do this, I collected data from the Area Resource File from 2001 that contained information on the number of physicians in a particular county- I collected this data for both total number of physicians and non-federal physicians. To determine the "high DTC" side of the border, I totaled the local ads from 2001 for each side of a particular border. If one side saw more advertising in total in 2001, it is called the "high DTC" side. I then ran a t-test to see if the means of the numbers of physicians were the same on either side of the border. The t-test fails to reject that the number of physicians is the same on either side of the border, with $p=0.21$ for total doctors and $p=0.25$ for non-federal doctors.

Next, we might be concerned that the counties on the "high DTC" side of the border are systematically higher income or higher population than those on the "low DTC" side of the border. A t-test fails to reject that average income is the same across the border, with $p=0.47$. Similarly, a t-test for population fails to reject that the populations are the same across the border with $p=0.17$.

Results for all of these tests are presented in Table 18. It should be noted that the market fixed effects in the model allow for different levels in all of these variables so long as the trends are parallel and that the effectiveness of advertising is the same across the border. That there does not seem to be systematic differences in characteristics across the border further solidifies that this level of comparison is a reasonable one.

Table 18: Tests for Selection Across Borders

variable	high DTC	low DTC	t-stat	p-value
Physicians	855 (136)	715 (124)	1.25	0.215
Non-Fed Physicians	98.8 (12.9)	88.0 (11.3)	1.15	0.252
Income	24,162 (460)	23,881 (399)	0.723	0.471
Population	1,150,642 (168,114)	979,142 (180,925)	1.37	0.173
N = 152 border clusters				
Means Above Standard Errors in Parentheses				

F.2

We might also worry that the border counties are systematically different from non-border counties in observable ways. This would hurt the generalizability of the results. As such, I compare border counties to interior counties in terms of the average number of doctors, the average number of non-federal doctors, the average population and the average income in Table 19. These comparisons are at the county level rather than at the border cluster as in the previous comparison, so there are many more observations.

Table 19: Tests for Observable Selection into Border Sample

variable	Border Sample	Non Border Sample	t-stat	p-value
Physicians	297.404 (864.491)	330.548 (1175.56)	0.792	0.428
Non-Fed Physicians	34.827 (81.697)	36.932 (99.333)	0.553	0.580
Income	24,203 (6249)	24,413 (6201)	0.763	0.446
Population	127,310 (287,262)	123,482 (346,498)	0.286	0.774
N=664 border, N=2221 non-border				
Means Above Standard Errors in Parentheses				