Journal of Marketing Research
Article Postprint

Asymmetric Social Interactions in Physician Prescription Behavior: The Role of Opinion Leaders

Harikesh S. Nair
Associate Professor of Marketing, Graduate School of Business
Stanford University, 518 Memorial Way, Stanford, CA 94305
Phone: 650-723-9675; Fax: 650-725-7979

Puneet Manchanda
Professor of Marketing, Stephen M. Ross School of Business
University of Michigan, 701 Tappan St., Ann Arbor, MI 48109
Phone: (734) 936-2445; Fax: 734-936-8716

Tulikaa Bhatia
Assistant Professor of Marketing, Rutgers Business School
94 Rockafeller Road, Piscataway, NJ 08854, Phone: 732-445-5274

Past versions: June 2006, May 2008, December 2008
This version: August 31, 2009

* The authors are listed in reverse alphabetical order. The authors would like to thank Ron Burt, Pradeep Chintagunta, Tim Conley, Wes Hartmann, Peter Reiss, Christophe Van den Bulte, Raphael Thomadsen and seminar participants at Berkeley (Haas), Chicago GSB (O&M), Christian Albrechts University, Erasmus, Michigan (Information Systems; College of Pharmacy) and participants at Frank M. Bass (Dallas, 2007), INFORMS Practice of Marketing Science (Wharton, 2007), IOFest (Stanford, 2006), Marketing Science (Pittsburgh, 2006) and SICS (Berkeley, 2006) conferences, and at the 7th Choice Symposium (Wharton, 2007) for feedback; and an anonymous pharmaceutical company for providing the data. Manchanda would like to thank the Kilts Center for Marketing and the True North Communications faculty research fund at the University of Chicago for research assistance. The authors can be contacted via e-mail at harikesh.nair@stanford.edu (Nair), pmanchan@bus.umich.edu (Manchanda) and tulikaa@rci.rutgers.edu (Bhatia). The usual disclaimer applies.

Asymmetric Social Interactions in Physician Prescription Behavior: The Role of Opinion Leaders

Abstract

The authors quantify the impact of social interactions and peer effects in the context of prescription choices by physicians. Using detailed individual-level prescription data, along with self-reported social network information, they document that physician prescription behavior is significantly influenced by the behavior of research-active specialists, or "opinion leaders" in the physician's reference group. The authors leverage a natural experiment in the category, whereby new guidelines released about the therapeutic nature of the focal drug generated conditions where physicians were more likely to be influenced by the behavior of specialist physicians in their network. They (1) find important, statistically significant peer effects that are robust across model specifications, (2) document asymmetries in response to marketing activity across nominators and opinion leaders; and (3) measure the incremental value to firms of directing targeted sales-force activity to these opinion leaders, and present estimates of the social multiplier of detailing in this category.

Key-words: *Social Interactions, Peer effects, Social Multiplier, Contagion, Physician Prescription Behavior, Pharmaceutical Industry.*

Marketers, sociologists and economists have traditionally been interested in the role of interpersonal communication (i.e., communication outside the firm's control) on consumer choice and consumption behavior. These interactions have been variously labeled as "peer effects," "contagion" and "word-of-mouth effects." In this paper, we test and provide empirical support for asymmetric peer effects. These effects arise when some consumers exert a stronger influence on the attitudes and behavior of other consumers than vice versa. Such consumers have typically been labeled "opinion leaders" in the literature (Rogers 2003, Chapter 8). There is little research in marketing that has tested for the existence of these asymmetric peer-effects.

The context of our analysis is prescription drug choice by physicians. An asymmetric social interaction or "peer effect" arises in this setting because non-specialist physicians may rely on prominent physicians, the "opinion leaders," to help reduce the uncertainty around their prescription choices. The role of opinion leaders becomes most salient when changes occur in the therapeutic environment, as these typically lead to increased uncertainty about drug efficacy among the non-specialist physicians. The pharmaceutical industry believes in the existence of such opinion leaders, and has invested in targeting marketing activities at opinion leaders (Cutting Edge Information 2004). However, to date, there is little empirical evidence that opinion leaders "matter" i.e., significantly influence the opinions and behavior of other physicians. Coleman et al. (1966) found no asymmetries in peer effects between nominators and their opinion leader's adoption pattern for a new drug. More recent work using the same data as that study found no peer effects at all (Van den Bulte and Lillien 2001). Finally, using computational models of network tipping, Watts and Dodds (2007) also find little or no role for opinion leaders.

Asymmetric social interactions have important implications for the allocation of marketing effort by firms. If present, they increase the return-on-investment to marketing activity targeted at agents having stronger influence. In the pharmaceutical context, if actions of opinion leaders have a true causal effect on the prescription behavior of other physicians, then marketing effort directed at the opinion leader will generate a multiplier effect. The multiplier arises because an incremental sales-call to an opinion leader increases the opinion leader's prescriptions, and induces the physicians s/he influences to prescribe more. The extent to which net prescriptions are higher due to these cross-physician spillovers is the *social multiplier* (c.f. Becker and Murphy 2000). Given that pharmaceutical firms in the US currently tend to set physician-level detailing based on past prescription volume (Manchanda et al. 2004), the presence of significant social multipliers would imply that the return-on-investment of detailing to opinion leaders may be much higher than is suggested by just the opinion leader's prescription volume. We use our estimates to measure the social multiplier of detailing in our data.

The identification and management of opinion leaders is an important component of marketing in the pharmaceutical industry. Marketing to opinion leaders is typically managed by direct contact with physicians through detailing, and in several instances, using special teams consisting of higher caliber detailers (called "Medical Scientific Liaisons"). About half of the large pharmaceutical companies are reported to have MSL-s, with an average team size of about 45 (Cutting Edge Information 2004). Opinion leaders are typically believed to be physicians who have an academic title with the department of a medical school and have contributed peer-reviewed publications (Tan 2003). Industry reports suggest the pharmaceutical industry spends an estimated 24% of their new product commercialization budget on opinion leader activities (Cutting Edge Information 2004). Opinion leader activity is also stepped up when environmental changes occur, for instance, during the launch or withdrawal of drugs or the issue of new guidelines by the Department of Health and Human Service and/or the National Institutes of

Health (NIH). While firms try and manage their relationships with these opinion leaders via marketing, anecdotal evidence suggests that the identification of opinion leaders and the extent to which they impact other physicians are issues the industry grapples, especially as there is little published work quantifying the return on investment from targeting these opinion leaders.

To test for these effects, we leverage a novel dataset that is based on a combination of primary (survey) and secondary (behavioral) data. Broadly speaking, there are five major challenges that arise in measuring these effects. First, opinion leaders that constitute the reference group for a given physician need to be identified. Second, once these opinion leaders have been identified, some change in the environment needs to take place in a manner that it affects the attitudes and/or behavior of the opinion leader. Third, these changes then needs to be transmitted to the agents whose opinions and/or behavior is affected by the opinion leader's behavior. Fourth, there should be a resultant change in the behavior of these consumers. Finally, we need to be able to distinguish between correlation and causation in the observed behavior of physicians and their opinion leaders. As we discuss below, correlation in behavior can arise from three possible sources – endogenous group formation, correlated unobservables and/or simultaneity – and we need to be able to control for these explanations. As the past literature has pointed out (c.f. Manski 1993, Moffitt 2001), solving this latter identification problem is not easy. Our data enable us to formulate empirical strategies that address most, if not all, of these issues. We believe our identification strategy is novel to the literature, and is relevant across a broad range of situations involving the analysis of data arising from social interactions.

Our key contributions are as follows. First, we test for and find the existence of asymmetric peer effects amongst a specific social network (prescribing physicians). We find that opinion leaders' behavior significantly affects physician behavior after an exogenous change in the market that resulted in a change in the therapeutic environment. Our empirical results also find that peer effects in this category are asymmetric in the sense that opinion leaders' prescriptions are not statistically significantly affected by the prescription pattern of the physicians they influence. These effects are robust to functional form and alternative specifications. Further, we document the finding that peer influence can significantly impact behavior even in stable, mature categories. These are novel findings that add to the literature on peer effects in the presence of marketing, especially in the pharmaceutical industry.

Second, we also document asymmetries in response to marketing activity (detailing) across opinion leaders and nominators. We find that nominators tend to be less responsive than opinion leaders. Third, we measure empirically the effect size of these peer effects for both physicians and opinion leaders. We then combine our results on the response to marketing activity and peer effects to derive implications for marketing resource allocation for firms, and present estimates of the social multiplier effects of detailing. We find that for the average opinion leader, who influences 1.56 physicians, social interactions alone provide an additional 5% increase in prescription revenue i.e., a social multiplier of detailing of about 1.05. For the top opinion leader, the social multiplier is 1.35. This large difference underscores the importance of estimating both the average multiplier and the identification of top influencers to make optimal resource allocation decisions.

In addition, we also discuss and clarify how the identification issues that arise in measuring and testing for causal peer effects may be overcome for data-rich settings such as ours. The increasing salience of (online) social networks in the real world makes our methodology and findings particularly relevant for practicing managers and academics who are interested in understanding the return on investment of marketing activity to opinion leaders (e.g. Godes and Mayzlin 2009).

Our paper is related to the sociology, economics and marketing literature on testing for social interactions using micro-level data. A subset of this work has focused on asymmetric peer effects, modeling for example, the role of ties in referrals (Reingen and Kernan 1986) or the characteristics of opinion leaders (e.g., Summers 1970), but without quantifying their effect on outcomes. A few other papers from the medical and public health literature have used surveys and/or field experiments to test for opinion leader effects (e.g., Lomas et al. 1991 and Celentano et al. 2000) or peer effects in general (Christakis and Fowler 2007). In economics, researchers have investigated social interaction effects more (e.g. Conley and Udry 2008; Sorensen 2006; Duflo and Saez 2003 to name a few). A small, but growing number of recent papers in the marketing literature has also investigated the potential role of peer-effects in new product adoption (e.g. Van den Bulte and Lilien 2001, Manchanda et al. 2008 and Iyengar et al. 2008 on prescription drugs; Bell and Song 2007 on Internet grocers; Nam et al. 2006 on video-on-demand technology). We refer the interested reader to Hartmann et al. 2008 for a recent and broad overview of the social interactions literature, which also discusses approaches from several related fields.

Broadly speaking, relative to the previous literature cited above, our approach has several distinguishing characteristics. These include documenting the asymmetric nature of peer interactions, distinguishing causal peer effects as opposed to correlated outcomes that do not rely on peer effects, and the determination of peer effects in mature product categories i.e., using post-adoption behavior. In terms of the causal effect determination, we believe this paper is one of the first to comprehensively outline and address the identification issues related to endogenous network formation, correlated unobservables and simultaneity, and to include specific controls for targeted marketing activity in the analysis of social interactions in the presence of marketing.

The rest of the paper is organized as follows. The next section presents the model and describe the data. We then present results from estimation. The last section concludes.

MODEL

We now discuss our model framework and empirical strategy. Our empirical framework is a descriptive linear model of prescription behavior, which we interpret as the reduced form of the behavioral process generating prescriptions for physicians and their opinion leaders (for structural approaches see Brock and Durlauf 2001; Hartmann 2008). In the "robustness" section later, we discuss some extensions of this linear model that accommodate alternative specifications of the effect of peers as well as relax the linearity assumption (via the use of a count model). We index physicians by *i*, *i*'s opinion leader by *j*(*i*), and time by *t*. Let *D* denote detailing, and *y* and *x* denote continuous variables representing new prescriptions for physicians and opinion leaders respectively. The starting point of our empirical specification for physician prescriptions is a linear regression:

$$y_{it} = \beta D_{it} + \delta x_{j(i),t} + \upsilon_{it} \qquad (1)$$

Here v_{it} denotes unobserved factors that shift prescriptions of physician *i* over time. While ideally we would like to include the actual opinions of the opinion leader as a covariate to capture the social interaction, these are unavailable in our data. Here, we think of the prescriptions of the opinion leader as a proxy for these opinions (later, we present sensitivity checks to different proxies for leader opinions). Formally, our test for peer effects in prescription behavior is whether δ is statistically significantly different from zero. An alternative model that uses the share, rather than levels of prescriptions is equivalent to (1), since the overall volume of prescriptions written for the disease condition remained roughly constant across the months in our data. Identification of peer effects in this model requires us to resolve five issues described earlier. We discuss these in sequence.

Reference group/peer determination: First, we need to identify the proper reference group or reference peer for each agent. Manski (1993) discusses in detail the need for exogenously defined social network information to identify peer effects from behavioral data. Intuitively, one cannot use behavior itself to define reference groups, if the goal is to obtain the effect of a reference group's behavior on an agent's actions. By grouping agents with ex-post similar actions together, a researcher attempting this approach essentially produces an upward bias in any peer effects unearthed through subsequent analysis. Similarly, geographic or location specific proxies for reference groups cannot sort between peer effects and common un-observables that affect the actions of all agents in the location similarly. We overcome these challenges by using a new dataset that contains detailed social network information obtained via a "sociometric" approach (e.g., Coleman et al. 1966; Conley and Udry 2008; Valente and Pumpuang 2007; Christakis and Fowler 2007). Here, individuals units are directly surveyed to obtain information about others who exert a peer effect on their behavior.[1] Each physician in the survey self-reports the doctor whose opinions he incorporates in his prescription decisions, thus identifying his social network. This provides us an exogenous measure of the physician's reference group or peer, circumventing the need to rely on behavior, location or geography-based proxies. Our use of the term "opinion leader" is to be interpreted in this sense as referring to doctors nominated by physicians in this survey (described in detail in a later section). In the

[1] In contrast, some studies follow the "key informant" approach, where a few individuals are polled to determine the identity of individuals with social influence (e.g., Celentano et al. 2000). Interestingly, Iyengar et al. (2008) find that the set of self-reported opinion leaders are different from those identified via a sociometric approach.

absence of such data, other researchers have often defined networks based on geographical location (Bell and Song 2007; Manchanda et al. 2008); dorm/work location (Sacerdote 2001; Duflo & Saez 2003; Sorensen 2006); and ethnic/cultural proximity (Bertrand et al. 2000).

Change in the External Environment: In stable drug categories, general practitioners may have little uncertainty about the usage and efficacy of the drugs they prescribe. Peer effects may be hard to uncover in such settings. Changes in the environment add exogenous variation that assist in unearthing the peer effect. An advantage of our data is that it covers a time-period where there was a significant change in the recommended usage of drugs in the therapeutic category. For the therapeutic category that we study, this environmental change relates to new treatment guidelines issued by the National Institutes of Health (NIH) regarding appropriate treatment for specific disease indications (we describe the new guidelines later in the paper.) This environmental change occurs around the mid-point of the data, and is exogenous to behavior as it arises from the behavior of a third party that is not affected by the actions of physicians and their opinion leaders, which aids identification. Hence, our analysis exploits how changes in prescription behavior of opinion leaders (x in equation 1) before and after the issuance of the guidelines, generates variation in changes in prescription behavior of their nominating physicians (y). In the survey, physicians also report their mode of interactions with their opinion leaders. Hence, our data also allow us to provide some insights into the mechanism through which the opinion leader effect manifests itself.

Distinguishing Causality from Correlation: As mentioned above, peer effects imply that the behavior of agents in the same reference group would tend to be correlated. However, correlation in the behavior of agents *per se* does not imply that any one agent's action has a causal effect on the actions of others in the group. In addition to peer effects, such correlation in behavior could arise due to three other factors, viz. endogenous group formation, correlated unobservables and simultaneity (see Moffitt 2001 for a discussion.) Only a causal peer effect implies a social multiplier; hence it is important to sort out causal effects from each of these sources of correlation. In our application, another factor that could lead to correlation is targeted marketing.

Endogenous group formation: Endogenous group formation arises in our context if physicians choose doctors with similar "tastes" for prescriptions as their opinion leaders (i.e. homophily). For instance, physicians who face patient bases requiring treatments using a specific class of therapeutic drugs may meet experts in that therapeutic category at conferences organized by drug companies. If physicians choose these experts as opinion leaders, it is likely that such physician-opinion leader pairs tend to prescribe more in the therapeutic category than average. In this case, the observed correlation in the behavior of the physician and his opinion leader could arise from omitted individual characteristics that are correlated within the group. In equation (1), such endogenous group formation implies that physician *i*'s unobserved tastes (v_{it}) and opinion leader *j*'s tastes for prescriptions could be correlated – if opinion leader *j*'s tastes also drive his prescriptions, x_{jt}, this generates a correlation between x_{jt} and v_{it} leading to a upward bias in the estimates of δ.

The solution to the group formation problem is facilitated by the availability of panel data (as noted by Manski 1993, the prospect for identification of peer effects in cross-sectional data are poor). Panel data enables us to include physician-specific fixed effects in the regression (1). In terms of our model, we write,

$$\upsilon_{it} = \alpha_i + \eta_{it} \qquad (2)$$

where α_i is a fixed effect specific to physician *i*, which controls for unobserved (to the econometrician) time invariant tastes for prescriptions. By controlling for physician *i*'s tastes, we control for the portion of v_{it} that is correlated with x_{jt} via correlation with opinion leader *j*'s tastes, thus accommodating the endogenous group formation problem. The identifying assumption is that, we assume that group selection is fixed over time, and that physician group formation is not influenced by changes in the external environment.

***Correlated unobservables*:** A second concern is whether there exist correlated unobservables that drive prescriptions of both the physician and the opinion leader similarly. If uncorrected, these manifest themselves as peer effects. An obvious source of correlation is sales-force activity (i.e. detailing) directed at physicians and opinion leaders by drug companies. One can partly control for this source of correlation by including time-period fixed effects that pick up common trends in marketing activity to physicians (e.g. Van den Bulte and Lilien 2001). In our setting, we fully control for such marketing activity by obtaining direct data on detailing to physicians, which we include as explanatory variables in the regression. A potential concern with using this variation arises because detailing may be targeted to physicians. As documented in the recent literature (e.g. Manchanda, Rossi and Chintagunta 2004), many pharmaceutical companies in the US, including our firm, decide detailing allocations based on a volume-based rule. Under this rule, physicians are allocated detailing levels corresponding to their position in deciles of past prescription volume in the focal category (we find evidence for this detailing pattern in our data.) This volume-based detailing rule implies that *D* may be correlated with v_{it}.

Our control for this potential endogeneity derives from the nature of the targeting rule. In effect, on account of stable patient bases, physicians rarely move across deciles (we find this in our data as well). Thus, the inclusion of physician fixed effects pick out the across-physician variation in detailing, and controls for the endogeneity concern. Thus, only within-physician, across-time variation in detailing is used for identification. Fixed effects do not fully absorb all detailing variation however, as in practice, actual detailing levels are centered around, but not exactly equal to top-down allocated levels due to several unanticipated factors that affect visits. These include shocks to physicians' schedules (e.g. the physician is not in his office during a detailing visit), or unanticipated detailer time-constraints (a patient is taking too long, requiring postponement of the visit). This deviation from pre-allocated levels is orthogonal to physician unobservables, and is used for identification. Thus, the underlying identifying assumption is that after controlling for α_i, within-physician detailing is independent of other physician and time-period specific unobservables, η_{it}.

We also consider the possibility that there may be additional correlated unobservables that generate co-movement in prescriptions. Candidates for such unobservables include trends in overall prescriptions across all physicians in the category, as well as any spatially correlated region/location specific shocks to prescription behavior that are captured by η_{jt}. We address these as follows. First, we include a full set of time-period fixed effects. These control for any time-trends common across all physicians and opinion leaders. Second, recall we include physician fixed effects. As none of the surveyed physicians in our data share a zip code, physician fixed effects are equivalent to including a full set of zip code fixed effects. Hence, time-invariant spatially correlated unobservables are also fully controlled for. A final issue is whether there are unobservables that are correlated at the level of the zip code *and* time.

To consider this issue, we discuss a potential difference-in-difference approach.[2] We have access to the prescription behavior of all physicians in the country. We use this data to compute the mean prescription of all other physicians in physician *i*'s zip-code, denoted by $z_{-i,t}$,

[2] Christakis and Fowler (2007) use the ex-ante symmetry in the network to rule out correlated unobservables. In other words, if both parties in a dyad are exposed to a common unobservable and it leads to a peer effect, then the only reason that the effect is asymmetric is because the link between the two parties is asymmetric. While this is a valid identification strategy, the difference-in-differences approach is more general as it does not depend on the ex-ante knowledge of asymmetry in the network.

which we include as a covariate in the regression. Essentially $z_{-i,t}$ proxies for all unobserved time-period and location-specific shocks to prescriptions that are common to all physicians in *i*'s location.[3] By including these in the regression, we essentially use the prescription behavior of other physicians in *i*'s location as a control. Thus, we further decompose (2) as,

$$\eta_{it} = \gamma z_{-it} + \varepsilon_{it} \tag{3}$$

where ε_{it} is a mean zero error term. Note this strategy is subject to the implicit caveat that a given physician's opinion leader does not influence other physicians in his zip code. Unfortunately, data on the social networks of the universe of physicians is not available to test this. Hence, our approach will be to present extensive sensitivity checks in which $z_{-i,t}$ is included or excluded from the regression. Tests for correlated unobservables (presented later) suggest that most of the spurious correlation is along the temporal dimension, which is fully picked up by time-period fixed effects. Hence, the effect of $z_{-i,t}$ on our results is small.

Simultaneity: Finally, we are careful in considering potential simultaneity. Simultaneity implies that physician *i*'s actions and opinion leader *j(i)*'s actions may be contemporaneously interdependent. If peer effects exist, the fact that opinion leaders affect physicians while physicians simultaneously affect them leads to an upward bias in the estimation of the interactions. In the context of our model, if physician *i* and opinion leader *j*'s prescriptions are simultaneously determined, high values of ε_{it} would tend to induce high values of $x_{j(i),t}$, thus leading to an upward bias.

We control for the simultaneity problem via exclusion restrictions. In our context, detailing to the opinion leader, $D_{j(i),t}$, as well as the mean prescriptions of all *other* physicians in the opinion leader's zip-code, $z_{-j(i),t}$, form excluded variables that affect the prescriptions of the

[3] Less than 10% of opinion leaders and nominators are in the same zip code in our data. Our results do not change if we exclude these physicians from our sample.

opinion leader (the endogenous variable), and can be excluded from the prescription equation for physician i. Both $D_{j(i),t}$ and $z_{-j(i),t}$ impact the opinion leader's prescriptions, and are thus correlated with the endogenous variable $x_{j(i),t}$, but uncorrelated with ε_{it}. Thus they serve as instruments for $x_{j(i),t}$ thus addressing the potential simultaneity concern. An alternative approach would be to assume that only past opinion leader prescriptions affect the physician's current prescriptions (i.e. there is no contemporaneous linkage in behavior). We explore model sensitivity to such specifications in the "Robustness" section.

Note that if we assumed that, given their "expert" status (details in the "Data" section below), opinion leaders were not affected by physicians, there would be no simultaneity problem by construction (e.g. Sorensen 2006). Rather than assume away simultaneity concerns a priori in this manner, we use the data to check whether peer effects are truly asymmetric. We run the analogous regressions for opinion leader *j(i)* (i.e. $x_{j(i),t}$ regressed on $\alpha_{j(i)}$, $D_{j(i),t}$, y_{it} and $z_{-j(i),t}$) to check whether physician-prescriptions have a significant effect on the prescription behavior of their opinion leaders. Analogously, D_{it} and z_{-it} are excluded variables for the opinion leader's prescription equation, and serve as instruments for y_{it} in the opinion leader's prescription equation.

Final Specification: Based on the above discussion, our final specification is,

$$y_{it} = \alpha_i + \gamma_t + \beta D_{it} + \delta x_{j(i),t} + \gamma z_{-i,t} + \varepsilon_{it}, \quad i = 1,..,N; t = 1,..,T \qquad (4)$$

The corresponding specification for i's opinion leader is:

$$x_{j(i),t} = \alpha_{j(i)} + \tau_t + \varpi D_{j(i),t} + \varsigma y_{it} + \zeta z_{-j(i),t} + \varepsilon_{j(i),t}, \quad t = 1,..,T; \qquad (5)$$

We estimate both specifications via fixed-effects panel data linear instrumental variables regression. We now describe our data.

DATA DESCRIPTION

Our data pertain to physician prescription behavior in a large therapeutic class (we cannot reveal the name of this class due to confidentiality concerns). The drugs in this class address a serious chronic disease condition that affects about a quarter of all adults in the United States. We consider a combination drug sub-category for the treatment of this chronic disease.

The dataset we use in our analysis is a combination of primary and secondary data. The primary data come from a market survey data carried out in Jan-Feb 2004, of 1500 physicians chosen randomly from a set of 56,000 regularly prescribing physicians across the United States in this therapeutic category. The survey was commissioned by a large pharmaceutical company and carried out by a market research firm with the pharmaceutical company bearing all costs (confidentiality reasons preclude us from naming the companies). The main objective of the survey was to obtain names of those doctors whose actions influence the nominating physician's approach to the treatment of the chronic disease treated by combination drugs. Nominating physicians were encouraged to name doctors who were known to them (by reputation or otherwise) and then queried about the mechanisms by which they were able to obtain information about the opinion leader's beliefs and actions.[4] Note that the answers to the questions detailed in the survey (as described in the footnote), were deemed sufficient by the company to identify opinion leaders. While this direct elicitation method is somewhat different from methods followed in the literature (e.g., self-reports, key informants or sociometric measures such as "indegree"), our results suggest the survey has indeed identified relevant opinion leaders. Under the null of biased elicitation we can only explain correlation between prescriptions of the nominator and the nominee (via correlated unobservables), but not the robust

[4] The specific questions asked during the survey were (a) "Whose opinions do you value most regarding the treatment and/or management of [disease condition] among [disease condition] patients?", and, (b) "How do you obtain information from that Influencer about the treatment/management of [disease condition]?" It is possible that the wording of this question could "encourage" respondents naming peers; however, this would only attenuate the effects we calibrate using (independently) obtained behavioral data, making our results conservative.

asymmetry we uncover. From this survey, we have access to information on 290 physician-opinion leader pairs. Note that our use of the term "opinion leader" is to be interpreted in this sense as referring to physicians nominated in this survey. The opinion leaders so *identified reflect each nominating physician's subjective opinion regarding who in the field he considers an expert*, and whose opinion he incorporates while making prescription decisions in this therapeutic category. We believe this individual-specific measure is the appropriate one for identifying peer effects in such settings.

These data were then supplemented by secondary data – also collected with the help of the company – on the prescription behavior and the marketing activity directed at both the physicians and the opinion leaders. These data span a period of 24 months (from April 2002 to March 2004 inclusive) and contain the prescription counts for the combination-drug category and the count of details received for each month for the universe of physicians for the company's drug. Interestingly, for this drug, the pharmaceutical firm had decided to rely almost entirely on regular detailing to manage relationships with the opinion leaders. Our interviews with the managers of the firm indicated that while the firm would have liked to consider other forms of marketing, such as the setting up of an MSL team, they were not doing that during the period of our data.[5] We also learnt from the firm that only about 50% of the doctors identified as opinion leaders in the survey were on the firm's own list of opinion leaders.

Descriptives: Primary data: We have demographic and location information for 290 opinion leader-physician pairs (including primary affiliation, zip-code and specialty). There are 267 unique nominating physicians. There are 182 unique nominated physicians ("opinion leaders"). The distribution of nominations in the survey is presented in Tables 1 and 2. Interestingly, over

[5] Our data do not contain information on sampling. However, a sample cannot be dropped unless the sales person calls on the physician. This usually results in a very high positive correlation between detailing and sampling and the exclusion of this data is unlikely to change our results, especially those that pertain to the peer effects.

91% of physicians reported being influenced by only one opinion leader. About 38% of the nominated doctors were named as opinion leaders by more than one physician. We do not see any overlap between opinion leaders and nominating physicians.[6]

The typical opinion leader is a research-active specialist physician in the therapeutic category and is associated with a university-based hospital. 97.4% of the opinion leaders are specialists. Over 90% of the opinion leaders in the sample are associated with hospitals, and about 30% are affiliated with University hospitals. The average opinion leader has published about 7.2 refereed papers (st. dev. 9.68, min 0, max 40) in this therapeutic class, confirming his status as an "expert." The survey also queried the nominating physicians about their mode of interaction with the opinion leader. The dominant mechanism of information transfer as reported by the physicians was direct contact, with about 94.5% of nominating physician mentions. This provides some support for our model formulation (equation 1) in which the nominating physician is assumed to respond to the prescriptions of the corresponding opinion leader. Other mechanisms of interaction included symposia/conferences (78%), meeting in clinical and/or hospital settings (67%) and via scientific articles published by the opinion leader (32%).

Prescription/detailing data: The secondary dataset contains information on 24 months of new prescriptions for the combination-drug category for the entire universe of physicians in the therapeutic class of the disease. The data also contains information on monthly physician-level detailing activity by the focal firm in the category. Unfortunately, detailing activity for the other competing drugs in the category is not available at the individual physician-level. We focus on the two largest drugs viz. drug 2, the focal drug, and drug 1, the main competitor (the other two

[6] The company asked the market research firm for a list of only those physicians who nominated an opinion leader. Thus, we can only measure the "treatment on the treated", where the "treatment" is the effect of an opinion leader.

have small market share). We supplemented our data with monthly national aggregate detailing for drug 1 (detailing for drugs 3 and 4 was negligible).

Descriptive statistics for the sample are provided in Table 3. The table shows that nominating physicians typically write a larger number of prescriptions (almost twice that of the opinion leaders) and also receive a higher level of detailing (about 50% higher). This is consistent with anecdotal evidence that opinion leaders tend to be focussed more on medical research and academic publication rather than practice. As can be seen from the table, opinion leaders are detailed less than the nominating physicians. This is likely to be a function of the category volume rule followed in the industry.

In computing z_{-it} in table 3, we had to make a decision whether to average all physicians in a zip-code, or only "active prescribers" in this disease category. We adopt the former approach as (a) it is possible that the guidelines affected all doctors in the zip-code, not just active prescribers; and, (b) since we did not have access to the database of physicians from which the company picked the random sample, we would otherwise have to make arbitrary guesses about who to include when computing z_{-it}. Given this, we present several robustness checks in which we drop z_{-it} from the regressions.

Change in drug-usage guidelines: An important aspect of the data is that it covers a time-period where there was a significant change in the guidelines for the usage of drugs in this therapeutic category. This change in the treatment environment is important for us since it is then that family and general practitioners are most likely to seek and value the opinion of specialists in the category. In the Web Appendix we provide excerpts from published sources and the summary findings from survey of physicians attesting that, in general, a change in guidelines usually increases the uncertainty in terms of physician prescription decisions. In our context, an exogenous change in the market occurred in May 2003, in the form of an announcement by the National Institutes of Health (NIH) affiliate releasing new treatment guidelines for the disease. Thus, we have 13 months of behavioral data before the guidelines were released and 11 months after. The guidelines suggested that, as against the prevailing norm, the initiation of treatment for severe cases of this condition should comprise of at least two agents (or molecules). These guidelines tended to favor the so called combination drugs in this category. A combination drug typically had the two agents in the same pill and results in "polytherapy". This had the obvious advantage of increasing compliance amongst patients and had higher efficacy. Thus, we expect that all combination drugs, including drug 2, should show an increase in prescriptions after these guidelines were issued. Prior to the issuance of new guidelines, combination drugs were generally considered "aggressive" therapy. [7]

We now document the changes in prescription and detailing behavior before and after the release of the guidelines. Figure 1 presents the distribution of mean monthly prescriptions in the combination-drug category for nominating physicians and opinion leaders before and after May 2003 (when the new treatment guidelines were introduced). For each physician, we compute the mean monthly new prescriptions before and after and present them in a box plot. As can be seen from the figure, both sets of physicians prescribe more of the combination drugs category. The mean increase in new prescriptions across both groups is about 10%. We now turn our attention to the distribution of monthly detailing for drug 2 across all physicians and months (Figure 2). Interestingly, the firm seems to deviate from the detailing allocation rule cited above just after May 2003. This can be seen from the figure as the firm details more to opinion leaders after the change even though they write fewer prescriptions in the combination-drug category (see Figure

[7] The firm also showed us survey data collected from 319 physicians after the guidelines were released. These physicians noted that, based on the guidelines, they would expect polytherapy to become more prevalent, leading to an increase in the prescription levels of combination drugs.