

Patient Learning and Advertising in the Diffusion of Cox-2 Inhibitors *

Pradeep K. Chintagunta [†]

Renna Jiang [†]

Ginger Z. Jin[‡]

October 11, 2006

Preliminary and Incomplete
Comments Welcome

^{*†} Graduate School of Business, University of Chicago. Email: pradeep.chintagunta@ChicagoGSB.edu, rjiang1@ChicagoGSB.edu. [‡]Department of Economics, University of Maryland & NBER. Email: jin@econ.umd.edu. We wish to thank Greg Crawford, Sean Nicholson, Marta Wosinska, Puneet Manchanda, and the participants at the 2006 American Society of Health Economists Conference for constructive comments and suggestions. All errors remain ours.

Abstract

The recent withdrawal of Cox-2 Inhibitors has generated debates on the role of information in drug diffusion: can the market learn the efficacy of new drugs, or does it solely depend on manufacturer advertising and FDA updates? In this study, we use a novel data set to quantify the diffusion of three Cox-2 Inhibitors – Celebrex, Vioxx and Bextra. From 1999 to 2003, IPSOS, a marketing research company, tracked a representative sample of patients that recorded every Cox-2 prescription received. Starting January 2001, the company also collected information on how satisfied a patient is after she takes a specific drug. This data set, together with direct-to-doctor and direct-to-consumer advertising intensity from drug manufacturers, allows us to model patient learning and drug advertising separately. In the model, we also distinguish the learning of a drug's general efficacy from the learning of the specific match between a drug and a patient. The former entails learning across patients, while the latter is primarily within patient.

Preliminary results suggest that patient learning plays a much more important role in explaining drug diffusion than does advertising. At the beginning of 2001 and upon the Bextra entry in January 2002, doctors held a strong prior belief about the relative efficacy of Celebrex, Vioxx and Bextra. We find that patient satisfaction signal is much noisier than the prior. Hence, doctors learn from patient satisfaction information but the learning is gradual. In comparison, none of the advertising variables have significant and positive impact on prescription choice in the 2001 to 2003 time period. We also find that learning across patients and learning within patients are both important: a model with one of them generates much worse fit with the real data.

1 Introduction

Information plays a critical role in evaluating a prescription drug: before approval, the Food and Drug Administration (FDA) must gather data from clinical trials to determine the drug's overall benefits and risks. These clinical trials are often conducted on a limited number of human volunteers, in comparison with a placebo, and during a period that is appropriate to determine the short-run efficacy and safety. After the FDA approval, more information is generated from long-term clinical trials and actual everyday usage. Such information, especially that from clinical trials, often leads the FDA to reevaluate the drug.

While it is important to count on the FDA to evaluate whether a drug should enter or exit the market, it is equally important to understand how doctors prescribe the drug while it is available on the market. At one extreme, if every doctor strictly follows the FDA guidelines and ignores the information generated from daily practice, society must ensure that the FDA gathers all the relevant information and processes it in a timely manner. At the other extreme, doctors may have learned from the experience of their own patients, made inferences from drug advertising, talked to other doctors, and followed all the updates from the medical literature. If so, the FDA regulation on drug withdrawal is more of a follow-up from the market rather than a new guidance for the market. In that case, it would be very important to ensure that the information reaching doctors is unbiased and the information flow is frictionless throughout the market.

The purpose of this paper is to provide a detailed empirical account of how doctors have learned about Cox-2 Inhibitors from 2001 to 2003. The information issues arising from Cox-2 Inhibitors are dramatic and unusual, but the lessons learned from this class are potentially relevant for every drug.

Between 1998 and 2001, the FDA approved three Cyclooxygenase-2 (Cox-2) Inhibitors: Celebrex (Dec. 1998), Vioxx (May. 1999), and Bextra (Nov. 2001). All of them were heavily advertised as safer alternatives to the existing pain killers. By September 2004, the class had more than 10 million patients, the annual sales reached \$6 billion in 2003, and the total advertising dollars spent in 2003 were as high as \$400 million. After a clinical trial associated Vioxx with severe cardiovascular (CV) risks, Merck withdrew the blockbuster drug in September 2004. CV risks and enhanced concerns on skin irritation led to the withdrawal of Bextra in April 2005. As of today, Celebrex is the only Cox-2 Inhibitor remaining on the market, with warnings added in April 2005.

The ups and downs of Cox-2 Inhibitors raise serious questions regarding information: how does the information from FDA, drug advertising, and patient experience affect prescription behavior? Who is responsible for providing and updating the information?

This paper uses a unique data set to partially answer the first question. Specifically, a marketing research company, IPSOS, tracked a national representative sample of Cox-2 patients from 1999 to 2003. Not only did IPSOS report every Cox-2 prescription received by the sampled patients, it started to keep a longitudinal record of patient satisfaction since January 2001. These satisfaction measures, together with the advertising intensity from drug manufacturers, allow us to associate Cox-2 prescriptions with various sources of information. More precisely, we assume doctors held a prior about Cox-2 at the end of 2000, which summarizes the information from FDA approved labels, advertising, and patient experience up to 2000. Starting Jan. 2001, doctors received new advertising and new reports of patient satisfaction, revised their prior in a Bayesian fashion, and made new choices of Cox-2 based on the updated belief.

Our model is different from the existing literature in two ways: first, to our best knowledge, all the existing studies on drug learning have no direct data on patient satisfaction. Instead, authors assume that the unobserved patient satisfaction data conforms to an assumed statistical distribution. They then model prescription choice as a result of random draws from that distribution. Equipped with detailed satisfaction data, we are able to minimize and validate these structural assumptions.

Second, probably due to the lack of satisfaction data, existing studies focus on either learning across patients (Ching 2005, Coselli and Shum 2003, Narayanan et al. 2005) or learning within patients (Crawford and Shum 2005). We believe the two types of learning are linked: doctors are not only uncertain about the average quality of a drug, they also have imperfect information on the specific match between a drug and a patient. As a result, they try to learn both. The average drug quality can be learned from the experiences of all patients while the patient-drug match must be learned from a specific patient's own experience. Our data are better-suited but not ideal for integrating the two types of learning. Because we have no data on physician identities, we assume doctors in the same geographic area exchange opinions extensively and that they learn from each others' patients experiences. To mitigate the effect of arbitrary assumptions regarding the geographic area of information exchange, we test the scope of information pooling by changing the definition of geographic area and assessing the sensitivity of our empirical results. To separate the role of patient learning from the role of advertising, we obtain data on four types of advertising: direct-to-consumer advertising (DTCA), detailing,

professional journal advertising, and free samples (the latter three target doctors). We control for the effects of these factors on patients' prescription decision.

Results suggest that patient learning plays a much more important role in drug diffusion than does advertising. At the beginning of 2001 and upon the Bextra entry in January 2002, doctors held a strong prior belief about the relative efficacy of Celebrex, Vioxx and Bextra. We find that patient satisfaction signal is much noisier than the prior. Hence, doctors learn from patient satisfaction information but the learning is gradual. In comparison, none of the advertising variables have significant and positive impact on prescription choice in the 2001 to 2003 time period. We also find that learning across patients and learning within patients are both important: a model with one of them generates much worse fit with the real data.

The rest of the paper is organized as follows. Section 2 provides detailed information on the background of Cox-2 Inhibitors. Section 3 describes and summarizes the data. Section 4 presents the econometric model, while section 5 reports results and robustness checks. Conclusion is offered in Section 6.

2 Background

Cox-2 Inhibitors were initially introduced to reduce the gastrointestinal (GI) risks of conventional non-steroidal anti-inflammatory drugs (NSAIDS) while maintaining the same efficacy in pain relief. NSAIDS, such as Aspirin, ibuprofen (Motrin) and naproxen (Naprosyn), block Cox-1 and Cox-2 enzymes and therefore impede the production of the chemical messengers (prostaglandins) that cause inflammation. However, since some Cox-1 enzyme exists in stomach and its production of chemical messengers protects the inner stomach, blocking Cox-1 enzymes tends to reduce the mucus lining of the stomach, causing GI problems such as stomach upset, ulceration, and bleeding. In comparison, Cox-2 enzyme is located specifically in the areas that cause inflammation and not in the stomach. By selectively blocking Cox-2 enzyme, Cox-2 inhibitors reduce the GI risks.¹

Before FDA approval, clinical trials presented evidence that all three Cox-2s reduce the incidence of GI ulcers visualized at endoscopy compared to certain non-selective NSAIDS. But up to April 2005, only Vioxx demonstrated a reduced risk for serious GI bleeding in comparison with naproxen (FDA 2005). After FDA approval, all three Cox-2s were heavily marketed as

¹For a complete layman-description of Cox-2 inhibitors, readers can refer to www.medicinenet.com.

being equally effective as traditional NSAIDS but with less adverse effects on the GI system.

The diffusion of Cox-2 inhibitors was very fast: according to the National Ambulatory Medical Care Survey (NAMCS) and the National Hospital Ambulatory Medicare Care Survey (NHAMCS), in 1999 (the first year of Cox-2 introduction), the number of ambulatory visits resulting in Cox-2 prescriptions was totaled at 15 millions, slightly more than half of the visits that resulted in traditional NSAIDS. By the end of 2000, the number of Cox-2 visits has exceeded that of traditional NSAIDS, reaching an estimate of 31.5 million. This growth continued in 2001, but with a much slower speed (Dai et al. 2005, Table 2).

In terms of prescriptions, according to the New Product Spectra², the total number of new Cox2 prescriptions grew sharply from 61,066 in January 1999 to 2 millions in December 2000, but leveled off since January 2001. The number of all Cox2 prescriptions (including new and old) demonstrated a similar pattern. Since Bextra was not approved until November 2001, its introduction was mainly market stealing (from Celebrex and Vioxx) rather than market expanding. Figure 1 plots the number of new Cox2 prescriptions by drug and month.

All three drugs were heavily marketed. As shown in Figure 2, the amount of detailing expenditure was comparable across drugs and even across time. In comparison, direct-to-consumer advertising (DTCA) was less intense, more fluctuating, and concentrated in Celebrex and Vioxx (Figure 3). This is probably because Pfizer produced both Celebrex and Bextra and decided to only promote Celebrex in front of consumers.

At the first glance, the trends of advertising are inconsistent with the prescription pattern: although the total amount of detailing has increased after 2002 (largely because of the entry of Bextra), the number of (new and total) prescriptions remains stable since the beginning of 2001. Moreover, the ups and downs of DTCA do not seem to create any significant fluctuation in the number of prescriptions.

After a three-year placebo-controlled clinical trial³ showed that taking Vioxx 25 mg once daily doubles the risk of serious adverse cardiovascular (CV) events, Merck withdrew Vioxx on September 30, 2004. In April 2005, FDA's Arthritis and Drug Safety and Risk Management Advisory Committees reviewed the available data and concluded that (1) the increased CV risk

²A database provided by IMS Health that tracks monthly number of prescriptions (new and refill) dispensed by pharmacists and monthly advertising activities of pharmaceutical manufacturers upto 60 months after initial launch.

³The trial is named the Adenomatous Polyp Prevention on Vioxx (APPROVe).

is a class effect applying to all the Cox-2s and traditional NSAIDs; (2) Aside from the CV risk, Bextra is associated with an increased rate of serious and potentially life-threatening skin reactions and should be withdrawn from the market; (3) the overall benefits of Celebrex exceeds its potential risks, which allows Celebrex to remain on the market but the label must be revised to carry explicit warning on potential CV and GI risks (FDA 2005). The FDA does not rank the three Cox-2s by their CV risks, but the evidence underlying the withdrawal requests suggests that the overall quality of Celebrex is better than the other two, while Vioxx might be better than Bextra given the fact that skin irritations only apply to Bextra.

3 Data Summary

We combine two data sources: the patient-level prescription and satisfaction data come from the IPSOS patient diary database (IPSOS-PD), and the monthly advertising expenditures are derived from the New Product Spectra (NPS) database.

In 1997, the IPSOS created a national representative sample of 16,000 households and tracked their drug purchasing month by month.⁴ The patient-level data used in this paper include all the individual records that the IPSOS have collected on Cox-2 Inhibitors from January 1999 to August 2003. Since we do not know the patients' other drug purchases, the data do not allow us to examine the choice between Cox-2s and traditional NSAIDs. For this reason, we focus on prescription choices within Cox-2s.

Each record provides information on the patient's prescription date, age, sex, race, copay, insurance status, and residential location defined by region and DMAs (Designated Market Areas). Since over 90% of patients have health insurance and the self-reported copays are noisy and sometimes inconsistent with the reported drug insurance, we do not use insurance and copay information in the current analysis. There is a possibility that insurers use formulary to indicate certain preferences across the three Cox-2s. Unfortunately, we do not have this information, nor do we have information on the identity of insurers. While we may infer certain formulary status from the reported copays, we cannot infer the formulary status for the non-prescribed brands.

Starting from January 2001, the data also provide five satisfaction measures, reflecting patients' self report on the effectiveness of the prescribed drug, its side effects, whether the drug

⁴Detailed description is available at <http://www.ipsos.ca/product.cfm?id=66&name=Healthcare&fn=health&fl=reid> and Bowman et al. (2004).

works quickly, how long it lasts, and whether it is easy to take. Each satisfaction measure is scaled from 1 to 5 in integer, with 1 for extremely satisfied and 5 for extremely dissatisfied.

The full sample from 1999 to 2003 in the IPSOS data involves 3,781 patients and 14,532 prescriptions of Cox-2s. To ensure that this sample is indeed national representative, we calculate the number of prescriptions and drug-specific market shares from the sample and compare their trends with those reported in the NPS. They are similar. We also regress the number of new patients in our sample and the number of new prescriptions in the NPS on various advertising variables, the regression coefficients and significance are comparable. These results motivate us to focus on the IPSOS data.

In the IPSOS data, 47.7% of patients received only one Cox-2 prescription, and the vast majority (91.5%) received no more than 10 prescriptions. Figure 4 provides more details on prescription frequency. Although the focus on Cox-2s does not allow us to say anything about the learning within these single-prescription patients, we do allow their satisfaction information to be used for other patients.

Conditional on the records with non-missing values in all five satisfaction questions, the sample is reduced to 2,512 patients and 9,067 prescriptions. This reduction is largely due to the fact that IPSOS did not collect satisfaction data until 2001. Between 2001 and 2003, the reporting rate for satisfaction measures is 87.8%.⁵ We define a “run” as a sequence of one or more prescriptions to a single drug. For example, if a patient receives a prescription sequence A,A,A,B,C, we say that he has three runs, the length of each being 3, 1, 1. By this definition, we classify the 9,067 prescriptions into 3,091 runs. An average run consists of 2.93 prescriptions, and an average patient has 1.23 runs in our data. The corresponding numbers are 2.37 and 1.23 in Crawford and Shum (2005).

Conditional on the 9,067 prescriptions that have non-missing values in patient satisfaction, Table 1 presents the number of prescription switches between the three drugs. By definition, switch does not occur unless a patient has at least two Cox-2 prescriptions. On average, the switching rate is 8.83%, slightly lower if a patient started with Celebrex (7.92%).

Table 2 summarizes satisfaction scores by drug and survey questions. Celebrex and Vioxx are comparable in effectiveness and side effects, with Vioxx slightly better in terms of working quickly, long-lasting and easy to take. As a result, the average across all five questions

⁵From 2001 to 2003, there are 10,327 prescriptions, of which 9,493 report some satisfaction scores and 9,067 report all five.

($satisf_{12345}$) is slightly better for Vioxx than for Celebrex. Among the three, Bextra is the worst in every measure except side effects. This is inconsistent with the fact that only Bextra is associated with skin irritations. However, it is inconclusive as to whether Bextra is better or worse in CV and GI risks. Figure 5 shows the distribution of $satisf_{12345}$ by drug. As expected, the satisfaction score is more dispersed for Bextra, because it entered the market later than the other two.⁶

If doctors learn anything from patient satisfaction, satisfaction scores should correlate with drug market shares and drug switches within patient. On drug switches, we run a logit regression. The dependent variable is whether the drug prescribed to patient p in the current period is different from that in the last period. The independent variables are the satisfaction scores patient p reports up to the last period. We use 7,953 observations on 1,333 patients who still have at least two prescriptions after deleting missing values on relevant satisfaction scores. As shown in Table 3, results suggest that side effects and easiness to take are not important but the more satisfied patients are with drug efficacy (i.e. the lower score on $satisf_{134}$), the less likely they switch.

The cleanest way to associate satisfaction with market shares is to check whether the number of new patients correlates with the satisfaction scores available up to the study period. We conduct such a check in two ways. First, with the patient-diary data, we count the number of new patients that first appear in month t and start with drug j . We regress this number on the cumulative average of $satisf_{12345}$ and all four types of advertising expenditures up to month $t - 1$. As expected, $satisf_{12345}$ is significantly correlated with the number of new patients (Table 4), implying that the more satisfied the market is with drug j , the more new patients start with drug j . Among advertising variables, only DTCA has a significant and positive correlation with the number of new patients. All the direct-to-doctor advertising are non-distinguishable from zero, some even show a negative sign. This is somewhat surprising, given the fact that drug companies spend more advertising dollars on doctors than on consumers. In the second check, we use the above number of new patients in an aggregate logit model, following the methodology outlined in Berry (1994). Results are similar to Table 4: higher historical satisfaction attracts more new patients⁷, and none of the advertising coefficients are significant.

Our data are truncated in the sense that many patients “exit” the data set before the end of the sample period. There are two possibilities: these patients may be cured, or they

⁶These satisfaction differences are statistically insignificant. The correlation across the five satisfaction measures ranges from 0.5 to 0.8.

⁷The coefficient for $satisf_{12345}$ is -1.9431 with t-statistics -3.8844 .

may switch to other anti-inflammatory drugs such as traditional NSAIDs. We have no direct information to distinguish the two explanations. But one may argue that, if patients are cured, they should be more satisfied with the drug than the continuing patients. On the other hand, if they switch away from Cox-2s, they may be disproportionately dissatisfied. The data seem to support the second scenario: patients with single prescriptions during the data period are less satisfied in every dimension.

Overall, these summary statistics suggest that patient satisfaction have a significant correlation with prescription choice within Cox-2s. The main task of our econometric model is to interpret such correlation as a process of how doctors learn from patient satisfaction and update their prescription choices accordingly. To best fit a learning model, our empirical analysis focuses on the new patients that first appear in the data set on or after January 1, 2001. The main reason for discarding old patients is because doctors may have formed patient-specific priors based on their satisfaction before 2001, on which we have no information. Fortunately, there are not too many of them: 2,062 out of the 2,512 patients (with non-missing satisfaction scores) are new since 2001, and these new patients account for 5,688 prescriptions. By definition, new patients are likely to have fewer runs and fewer prescriptions per run, which explains why the number of prescriptions has declined 37% in this step of data cleaning while the number of patients only goes down by 18%.

The experience of the old patients may have contributed to doctor beliefs about the average drug quality as of January 1, 2001. Such contribution will be captured in the model since we estimate the prior as of January 1, 2001. What is ignored is how the old patients' satisfaction reported after 2001 contributes to the across-patient learning after 2001. We will address this part of learning in the future.

4 Econometric Model and Identification

4.1 Model

Consider a situation in which doctor d has concluded that patient p needs a Cox-2 prescription of a fixed length starting from time t , but has not determined which drug of Cox-2 is the best choice. In making such choice, the doctor intends to maximize the patient's expected utility for this *single* prescription.

Here we make three assumptions: in reality the doctor-patient relationship involves a number of information and incentive issues, and the doctor may not act as a perfect agent for the patient. We ignore such imperfection because we have no data on individual doctors. This is the first simplification. Second, we focus on prescription choice within Cox-2s, largely because our data do not allow us to consider the potential tradeoff between Cox-2s and traditional NSAIDs. Third, we assume each doctor is myopic enough to focus on the current prescription. As detailed below, we assume a doctor considers all the drug information available to her up to t , but she does not consider how experience learned from the current prescription would affect her future prescription choice on the same or other patients. We choose to ignore forward-looking, not only because it simplifies the econometric model, but also because a large proportion of patients exit the data set with one prescription and the potential risk of mal-practice is likely to prevent doctors from experimenting.

We assume that patient p 's CARA utility from a prescription of drug j can be written as:

$$\tilde{V}_{pjt} = -e^{-\gamma(\tilde{Q}_{pjt} + \beta_{xj}X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt})}$$

where

- \tilde{Q}_{pjt} = doctor's belief about drug j 's quality for patient p at time t ;
- γ = risk aversion parameter which is greater than 0;
- X_{pt} = patient p 's characteristics at time t , including age and gender;
- Z_{jt} = drug j 's characteristics at time t , including various advertising expenditures up to t ;
- ϵ_{pjt} = logit errors.

In theory, advertising could have both informative and persuasive effects on doctor prescription. Under the assumption that only high quality firms can recoup heavy advertising through repeated purchases (Nelson 1974), doctors may infer quality from advertising expenditure. In this case, it is tempting to model advertising as a factor that contributes to the learning of \tilde{Q}_{pjt} . Such model has been explored in Narayanan et al. (2005). On the other hand, one may argue that doctors derive utility from advertising and such utility persuades them to prescribe the advertised drug. By this argument, advertising intensity enters the utility function directly.

Although Akerberg (2001, 2003) has provided a method to distinguish informative and persuasive advertising, it relies on the assumption that consumers can learn about product quality in one purchase and therefore any effect of advertising on experienced consumers is solely persuasive. The learning of drug quality is often noisy and time-consuming, which hardly

satisfies the assumption of experience good. Mathematically, we could still allow advertising to be both informative and persuasive at the wake of the literature, but the two are only identified from functional forms. To avoid arbitrariness, we try each separately. Since results are similar, we only present the specification that includes advertising directly in the utility function. However, this specification does not imply that we interpret advertising as persuasive. Rather, we view advertising variables as necessary controls so that we can separate advertising from patient learning.

The process of patient learning is modeled as follows. Doctors are uncertain about \tilde{Q}_{pjt} . Specifically, doctors are uncertain about two things: the general quality of drug j that applies to every patient, which we refer to as Q_j ; and the specific match value between drug j and patient p , which we refer to as q_{pj} . The true effect of drug j on patient p is therefore

$$Q_{pj} = Q_j + q_{pj}.$$

This term is fixed but unknown to any one. Over the whole population, q_{pj} is independent and identically distributed by a normal density $N(0, \sigma_{q_0}^2)$.

When drug j was first introduced to the market (or at the beginning of our data set), all doctors shared two priors: for the general quality of drug j , the prior is

$$\tilde{Q}_{j0} \sim N(\bar{Q}_{j0}, \sigma_{Q_{j0}}^2).$$

The prior for the patient-drug match (q_{pj}) is mean independent of Q_{j0} and can be written as:

$$\tilde{q}_{pj0} \sim N(0, \sigma_{q_0}^2).$$

Together, the prior for the specific quality of drug j on patient p is

$$\tilde{Q}_{pj0} = \tilde{Q}_{j0} + \tilde{q}_{pj0} \sim N(\bar{Q}_{j0}, \sigma_{Q_{j0}}^2 + \sigma_{q_0}^2).$$

We allow both \bar{Q}_{j0} and $\sigma_{Q_{j0}}$ to be drug-specific. This reflects the fact that the initial information about the average drug quality, whether it is from FDA guidelines, medical research, or patient experience, may differ across drugs. For example, the prior on Celebrex and Vioxx is defined as of January 1, 2001 and the prior on Bextra is defined as of March 1, 2002 (the first date that Bextra appears in our data set). Since doctors may have learned about Celebrex and Vioxx before 2001, the prior should be less dispersed for them than for Bextra. Since we put no restrictions on $\sigma_{Q_{j0}}$, we can test this conjecture in the data. For simplicity, we assume the amount of patient heterogeneity (captured by σ_{q_0}) is the same across all three drugs. Note we

assume that doctors prior belief on the distribution of patient heterogeneity coincides with the actual distribution.

Over time, doctors receive signals from patient experience. We assume doctors located in the same geographic area (say a census region or a DMA) share information immediately and extensively. Assuming each prescription generates one signal, patient p 's satisfaction with drug j at time t , denoted as R_{pjt} , is a noisy but unbiased indicator of the true quality:

$$R_{pjt} = \alpha_0 + \alpha_R \cdot (Q_j + q_{pj}) + v_{pjt}$$

where α_0 and α_R equalizes the scales of R and Q , and the signal noise v conforms to $N(0, \sigma_v^2)$.

Let n_{pjt}^R denote the number of satisfaction reports from patient p on drug j up to time t in a certain area, and \bar{R}_{pjt} denote the average satisfaction across these n_{pjt}^R reports. At time t , doctors in that area will use all the n_{pjt}^R signals across all local patients to update their beliefs on the average drug quality Q_j . However, because patients are independent from each other, the other patients (patients other than p)' experiences do not contain any information about q_{pj} .

With all the patient satisfaction information up to t , doctors posterior on the effect of drug j on patient p can be decomposed into two parts: doctor's posterior about the general quality of drug j , and doctor's posterior about the specific match between drug j and patient p . That is:

$$\tilde{Q}_{pjt} = \tilde{Q}_{jt} + \tilde{q}_{pjt}$$

According to the Bayes rule (DeGroot 1970):

$$(1) \quad \begin{pmatrix} \tilde{Q}_{jt} \\ \tilde{q}_{pjt} \end{pmatrix} \sim N\left(\begin{pmatrix} \bar{Q}_{jt} \\ \bar{q}_{pjt} \end{pmatrix}, \Sigma\right)$$

where

$$\begin{aligned} \bar{Q}_{jt} &= \frac{\sum_p \frac{n_{pjt}^R \cdot \alpha_R \cdot (\bar{R}_{pjt} - \alpha_0)}{\sigma_v^2 + n_{pjt}^R \cdot \alpha_R^2 \cdot \sigma_{q_0}^2} + \frac{\bar{Q}_{j0}}{\sigma_{Q_{j0}}^2}}{\sum_p \frac{n_{pjt}^R \cdot \alpha_R^2}{\sigma_v^2 + n_{pjt}^R \cdot \alpha_R^2 \cdot \sigma_{q_0}^2} + \frac{1}{\sigma_{Q_{j0}}^2}} \\ \bar{q}_{pjt} &= \frac{\sigma_{q_0}^2 \cdot n_{pjt}^R \cdot \alpha_R \cdot (\bar{R}_{pjt} - \alpha_0 - \alpha_R \cdot \bar{Q}_{jt})}{\sigma_v^2 + n_{pjt}^R \cdot \alpha_R^2 \cdot \sigma_{q_0}^2} \\ \Sigma^{-1} &= \begin{pmatrix} s_{11} & s_{12} \\ s_{21} & s_{22} \end{pmatrix} \end{aligned}$$

$$\begin{aligned}
s_{11} &= \sum_p \frac{n_{pjt}^R \cdot \alpha_R^2}{\sigma_v^2} + \frac{1}{\sigma_{Q_0}^2} \\
s_{12} = s_{21} &= \frac{n_{pjt}^R \cdot \alpha_R^2}{\sigma_v^2} \\
s_{22} &= \frac{n_{pjt}^R \cdot \alpha_R^2}{\sigma_v^2} + \frac{1}{\sigma_{q_0}^2}
\end{aligned}$$

Note that the two posterior beliefs, \tilde{Q}_{jt} and \tilde{q}_{pjt} , are correlated because both make use of the satisfaction information from patient p . While the above formula focuses on one patient, we can show that the across-patient terms in Σ^{-1} are zero. Inverting Σ^{-1} implies that the posterior \tilde{q}_{pjt} is no longer independent across patients. This is because all the updates of q_{pj} rely on the update of Q_j , which in turn relies on satisfaction reports from all patients.

Equipped with the posterior updates, the expected utility is given by:

$$\begin{aligned}
E[\tilde{V}_{pjt}] &= -e^{-\gamma(\beta_{xj}X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt})} E\left[e^{-\gamma\tilde{Q}_{pjt}}\right] \\
&= -e^{-\gamma(\beta_{xj}X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt})} e^{-\gamma\bar{Q}_{pjt} + \frac{1}{2}\gamma^2\sigma_{\tilde{Q}_{pjt}}^2} \\
&= -e^{-\gamma\left(\bar{Q}_{pjt} - \frac{1}{2}\gamma\sigma_{\tilde{Q}_{pjt}}^2 + \beta_{xj}X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt}\right)} \\
&= -e^{-\gamma\left(\bar{Q}_{pjt} - \frac{1}{2}\gamma\sigma_{\tilde{Q}_{pjt}}^2 + \beta_{xj}X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt}\right)}
\end{aligned}$$

Thus, to maximize the expected utility is equivalent to maximize the following,

$$U_{pjt} = \bar{Q}_{pjt} - \frac{1}{2}\gamma\sigma_{\tilde{Q}_{pjt}}^2 + \beta_{xj}X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt}$$

The standard logit probability (McFadden 1973) for patient p getting drug j at time t is:

$$PR_{pjt} = \frac{\exp(U_{pjt})}{\sum_{j=1}^J \exp(U_{pjt})}$$

From the prescribing probabilities, we can estimate parameters by maximizing the log likelihood function:

$$\ln(L) = \sum_{p,j,t} 1_{data=pjt} \cdot \ln(PR_{pjt}).$$

4.2 Identification

The model employs four sets of parameters: $[\beta_{xj}, \beta_z]$ capture the effects of individual demographics and drug advertising, $[\bar{Q}_{j0}, \sigma_{Q_{j0}}, \sigma_{q_0}]$ capture doctors' prior, $[\alpha_0, \alpha_R, \sigma_v]$ capture the

importance of patient satisfaction, and γ captures doctors' risk preference.

The identification of β_{xj} comes from the time-invariant prescription heterogeneity across patients. For example, if Celebrex prescriptions are always concentrated in the elderly, it translates into a significant and positive coefficient corresponding to the interaction of Celebrex and age. Similarly, β_z is identified from the co-movements of drug market shares and drug advertising. Causality could go either way: on the one hand, advertising may trigger sales; on the other hand, historical or predicted sales patterns may motivate changes in advertising intensity. Since the main purpose of the paper is to detect patient learning, we use the advertising variables as controls, with β_z reflecting the correlation between advertising and sales.

The prior means of drug quality, \bar{Q}_{j0} , are identified from initial market shares. Because we focus on choices within Cox-2, the initial market shares only identify the relative magnitudes, leaving the default drug's \bar{Q}_{j0} under-identified. In the analysis, we define $\bar{Q}_{Vioxx,0} = 0$. However, patient satisfaction R is reported in absolute terms. Apparently, the noise in R , denoted by σ_v , is determined by the heterogeneity in R . Since we assume R equals to a linear function of true quality Q_{pj} plus noise, we can derive σ_v by regressing R_{pjt} on a full set of patient-drug dummies and calculating the standard deviation of the residuals. This procedure does not require any prescription data, so we complete it before estimating the full model.

In theory, the other two parameters, α_0 and α_R , can be identified by the scale difference between satisfaction R and true quality Q_{pj} . However, since we do not know Q_{pj} , they must be proxied by the posteriors. Note that evolving market shares define evolving posteriors about the general efficacy of the three drugs. Plotting these posteriors against the priors, we get the diffusion path of each drug. Such diffusion path helps us identify a number of parameters. If the path is flat for each drug, the lack of updating implies that patient satisfaction has little impact, which amounts to $\alpha_R = 0$. If drug j 's diffusion path is positively related to drug j 's average satisfaction over time, it implies a significant, positive α_R , and α_0 is simply an intercept shift that is derived from the relative scale of R and Q .

The dispersion on the prior of the average quality of drug j , namely $\sigma_{Q_{j0}}$, is identified by the speed of diffusion. According to the Bayes formula, the mean of the posterior, \bar{Q}_{jt} , is essentially a weighted average between R and the prior mean \bar{Q}_{j0} , while the weights are inversely related to the amount of noise in the two terms. Since we already identify the noise of R , a relatively small (large) $\sigma_{Q_{j0}}$ implies that doctors believe the prior is relatively precise (noisy) and therefore put less (more) weight on patient satisfaction, which results in slow (fast) learning.

Similarly, the dispersion on the prior of patient-drug match, namely σ_{q_0} , is identified by how fast doctors update their patient-specific beliefs. Small (large) σ_{q_0} implies that patient p 's doctor is reluctant (eager) to revise her prior after she receives p 's satisfaction report, because she thinks the report is relatively noisy (precise).

The risk aversion parameter, γ , is identified by a functional form restriction. As noted in Crawford and Shum (2005), the data only identify the term $\bar{Q}_{pjt} - \frac{1}{2}\gamma\sigma_{Q_{pjt}}^2$. To identify γ , we must decompose this term, which entails a functional form restriction on $\sigma_{Q_{pjt}}^2$. This restriction comes from the Bayes formula for the variance of the posterior.

5 Results

As described in Section 3, we focus on the patients that first appear in the data on or after January 1, 2001. The analysis sample ends at August 31, 2003 and is conditional on the prescriptions that come with valid answers for all five satisfaction questions. In total, the sample involves 2,062 patients and 5,688 prescriptions.

Before heading into the structural model, we check two benchmark models. These benchmarks utilize a discrete choice framework but do not incorporate a learning structure. Comparing them with our structural model will help us understand the importance of the learning structure. Specifically, Benchmark I estimates the prescription choice within Cox-2s, assuming that the utility of patient p using drug j is:

$$U_{pjt} = \beta_{j0} + \beta_s \text{satisf}_{12345} + \beta_{xj} X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt}.$$

To capture the fundamental difference across drugs, we also include a set of drug dummies, whose impacts on utility are captured by coefficients β_{j0} .

Benchmark II omits patient satisfaction in the utility function so that a comparison of the two benchmark models would highlight the role of patient satisfaction. Specifically, the utility function for Benchmark II is:

$$U_{pjt} = \beta_{j0} + \beta_{xj} X_{pt} + \beta_z Z_{jt} + \epsilon_{pjt}.$$

Assuming logit errors, we can write out the probability of patient p choosing drug j and maximize the overall likelihood. In both models, we normalize the satisfaction measure

as $6 - \text{satisf}_{12345}$ so that a positive coefficient on patient satisfaction implies that the more satisfied the better the drug choice. Since the benchmark models do not incorporate learning structure, in order to capture all the information available up to the study period, we compute the satisfaction variable as the average of all satisfaction reports up to one month before the prescription month.

For advertising variables, we report DTCA and detailing. Adding the other two advertisements (professional journal and free samples) do not change the results. To account for the fact that advertising may be long-lasting and the effect of advertising may diminish over time, we define DTCA and detailing as the log of their cumulative sum from the day of drug entry up to one month before the prescription month. We have tried other definitions, including the cumulative sum itself (without log), the advertising flow (instead of cumulative sum), and the monthly average of the cumulative sum. Results are similar.

As shown in Table 5, when we include both patient satisfaction and advertising in Benchmark I, patient satisfaction has an important impact on prescription choice, but all the advertising variables are indistinguishable from zero. The impact of satisfaction has a greater magnitude for Bextra, probably because Bextra is newer than the other two drugs. Furthermore, Benchmark I shows that, on average, Celebrex is comparable to Vioxx but Bextra is significantly worse than both. This reflects the fact that Bextra has the smallest market share among the three. In terms of demographics, female patients are more likely to get Celebrex and less likely to get Bextra, as compared to Vioxx.

In comparison, omitting patient satisfaction leads to a worse fit in Benchmark II. A likelihood test between the two benchmark models clearly favors Benchmark I. Moreover, in Benchmark II, DTCA has a positive and weakly significant impact on Bextra, which is counterintuitive because the intensity of DTCA is much less (in fact close to zero) for Bextra than for the other two drugs. In the mean time, the coefficient of the Bextra dummy is no longer worse than Celebrex and Vioxx. This contradicts the fact that Bextra has much smaller market shares than the other two. These findings, as well as the comparison with Benchmark I, suggest that a discrete choice model without patient satisfaction is probably subject to omitted variable bias.

The results on the two benchmark models encourage us to think more systematically about patient satisfaction. In accordance, the structural model adds a Bayesian learning structure on top of the classical discrete choices. As discussed in Section 4.2, we estimate the structural model in two steps: first, we regress R_{pjt} on a full set of patient-drug (pj) dummies, and compute the residuals' standard deviation. According to our model, this standard deviation gives us an

unbiased estimate of σ_v . With R-square 0.697, the regression produces $\sigma_v = 0.496$. Figure 6 plots the histogram of these residuals. Although a normality test rejects the hypothesis that the residual conforms to a normal distribution, the shape of the distribution is not too far away from normal.

In the second step, we set σ_v at 0.496 and search for the best parameters that maximize the overall log likelihood.⁸ To ensure a valid comparison with the benchmark models, we adopt the same variable definitions for patient satisfaction and advertising, except that we now give a separate consideration to each satisfaction report rather than pooling them into a simple average.

Results reported below assume that doctors talk to each other within a census region. The model assuming nationwide information pooling generates a significantly worse fit with the data. In the other direction, one may narrow the market definition from region to DMA. We have not implemented such a model because some small DMAs have very few observations. One possibility is to apply DMA market definition to large DMAs while keep the regional definition for small DMAs, but this involves arbitrary classification of DMAs. We will explore this more in the future.

Table 6 presents three sets of structural results: Column 1 focuses on the structural model but assumes that doctors are risk neutral (i.e. setting $\gamma = 0$). This is the simplest version that allows structure on across-patient and within-patient learning. The full model with risk aversion requires much longer time to estimate. We are working on it now. To gauge the relative importance of the two types of learning, Column 2 mimics Column 1 but ignores within-patient learning (by setting $\sigma_{q_0} = 0$), Column 3 mimics Column 1 but ignores across-patient learning (by setting $\sigma_{Q_{j_0}} = 0$).

Three findings stand out of Table 6. First, there is significant learning from patient satisfaction. On the one hand, the positive, significant estimate of α_R suggests that doctors believe the satisfaction reports from patients are correlated with drug efficacy and therefore use them to update the prior. On the other hand, the magnitudes of $\sigma_{Q_{j_0}}$ are much smaller than both the noise in satisfaction report (i.e. σ_v) and the dispersion of patient-drug match (i.e. σ_{q_0}). This suggests that doctors hold strong priors about the average efficacy of the three drugs. As a result, although they value satisfaction report, the updating on the general drug quality is slow. In comparison, the learning on the specific match between a drug and a patient is faster, because the magnitude of σ_{q_0} is much closer to that of σ_v .

⁸We use both gradient and simplex search, and the two always report the same results.

This interpretation is consistent with the comparison across Columns 1, 2 and 3. The overall likelihood of Column 1 is significantly better than that of Columns 2 and 3, suggesting that both across- and within-patient learnings are important in our data. However, the likelihood (and point estimates) of Column 3 is much closer to Column 1. This implies that a majority proportion of the data variations are driven by within-patient learning, the same conclusion as we have inferred from the relative magnitudes of σ_{q_0} , $\sigma_{Q_{j0}}$, and σ_v . Note that Column 1 also fits the data much better than the two benchmark models in Table 5, suggesting that the Bayesian learning structure captures a more significant amount of data variations.

The second finding is that no advertising variable has a significant, positive coefficient in the model that incorporates both types of learnings (Column 1). In fact, the coefficient for DTCA is even negative and significant. We suspect this unexpected sign captures something else that correlates with advertising but we do not observe. For this reason, we do not interpret these advertising coefficients as the causal impact of advertising on prescription choice. Rather, they are passive controls that allow us to single out the role of patient satisfaction.

Lastly, the prior estimates are largely as expected: the prior mean of Bextra is smaller than that of Vioxx and Celebrex, which is consistent with the relative market shares of the three drugs; the dispersion in the prior of Bextra is greater than that of the other two, which is consistent with the late entry of Bextra.

Overall, the structural model suggests that patient learning plays a much more important role in drug diffusion than does advertising. Doctors learn from patient satisfaction information but the learning, especially the learning on the general drug quality, is gradual. In comparison, none of the advertising variables have significant and positive impact on prescription choice in the 2001 to 2003 time period. We also find that learning across patients and learning within patients are both important, but more data variations are concentrated in within-patient learning.

6 Conclusion

Acquiring information about drug efficacy is not only the center of FDA regulations, but also the key element driving each prescription decision in doctor's office. Using a unique data set of patient diary, this paper describes how patient satisfaction and drug advertising affect the diffusion of Cox-2 inhibitors from 2001 to 2003. Preliminary results suggest that learning across patients and learning within patients are both important but manufacturer advertising has little

power explaining prescription patterns.

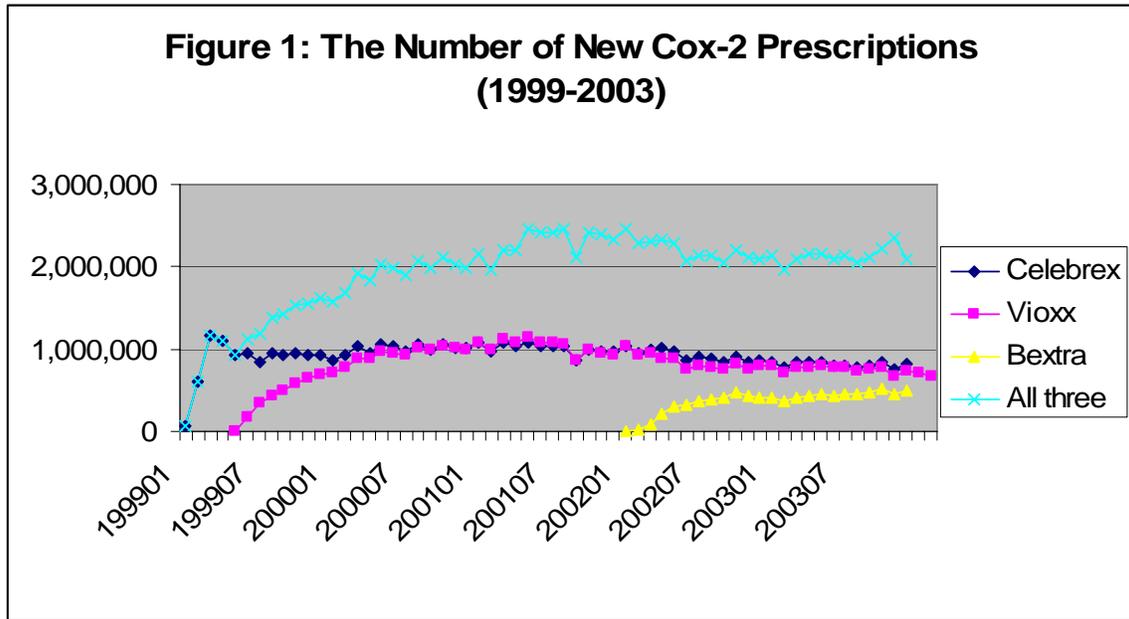
Before jumping into any conclusion, we would like to test whether the same results hold for several extensions: first, doctors may obtain information via sources other than patient report and drug advertising. These sources include news coverage, professional articles and FDA warnings. We are collecting data in these avenues. Second, Cox-2 inhibitors are (imperfect) substitutes to traditional NSAIDS and therefore a study on Cox-2 alone does not speak to the role of information in the choice between Cox-2 and traditional NSAIDS. We have expanded the patient diary data to include every patient that has ever consumed any NSAIDS since 1999. In the ongoing work, we modify the model to include traditional NSAIDS as the outside good. The third and more difficult extension is to make a better distinction between learning and the patient heterogeneities that are observable to doctors but unobservable to researchers. To carry out this task, we add patient random effects to the current model. Patient random effects may also help us identify patients into (doctor-observable) groups thus narrowing down the extent of across-patient learning. Lastly, it is possible that manufacturers advertise a specific drug according to some patient characteristics that they observe but we don't. This may hamper our interpretation of the advertising effects. We are searching for methodologies to address such endogeneity.

7 References

- Akerberg, Daniel (2001) "Empirically Distinguishing Informative and Prestige Effects of Advertising" *RAND Journal of Economics* 32, pp 316-333.
- Akerberg, Daniel (2003) "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination" *International Economic Review*, 44, pp. 1007-1040.
- Berry, Steve (1994) "Estimating Discrete Choice Models of Product Differentiation" *RAND Journal of Economics* 25, pp. 242-262.
- Ching, Andrew (2005?) "Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs After Patent Expiration" *University of Toronto Working Paper*.
- Coselli, Andrea and Matthew Shum (2003) "An Empirical Model of Learning and Patient Spillovers in New Drug Entry" *Journal of Econometrics*.
- Crawford, Greg and Matthew Shum (2005) "Uncertainty and Learning in Pharmaceutical Demand" *Econometrica*.

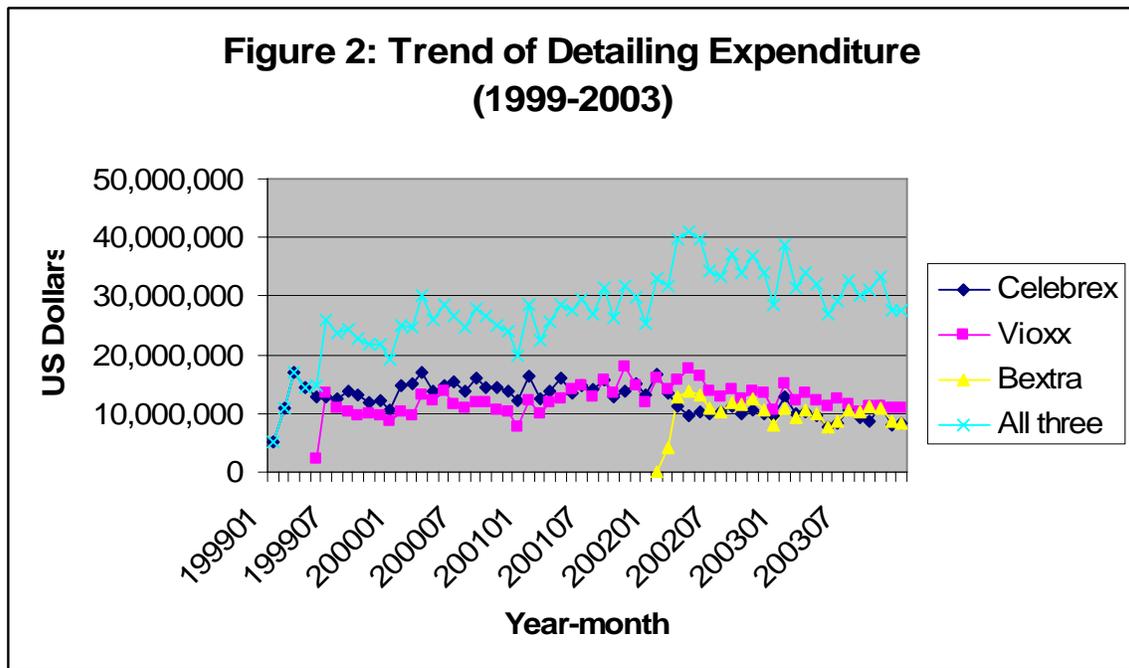
- Dai, Carolanne; Randall S. Stafford and Caleb G. Alexander (2005) National Trends in Cyclooxygenase-2 Inhibitor Use Since Market Release: Nonselective Diffusion of a Selectively Cost-Effective Innovation, *Achieves of Internal Medicine* Jan. 24, 2005.
- DeGroot, M. (1970) *Optimal Statistical Decisions* New York: mcGraw-Hill.
- Food and Drug Administration (2005) "Analysis and recommendations for Agency action regarding non-steroidal anti-inflammatory drugs and cardiovascular risk" FDA Memorandum as of April 6, 2005.
- McFadden, Daniel (1973) "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka (ed.) *Frontiers in Econometrics*, Academic Press, New York.
- Narayanan, Sridhar; Puneet Manchanda; and Pradeep K. Chintagunta (2005) "Temporal Differences in the Role of Marketing Communication in New Product Categories" *Journal of Marketing Research*.
- Nelson, P. (1974) "Advertising as Information" *Journal of Political Economy* 82, 729-753.

Figure 1: The Number of New Cox-2 Prescriptions (1999-2003)

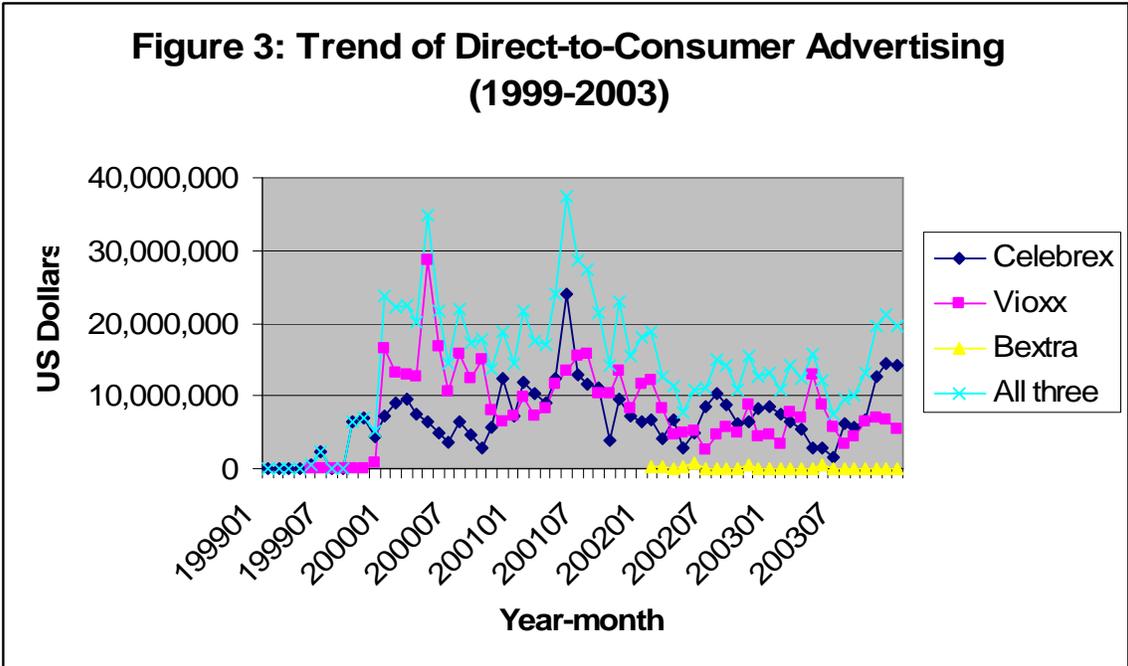


Source: New Product Spectra, 1999-2003.

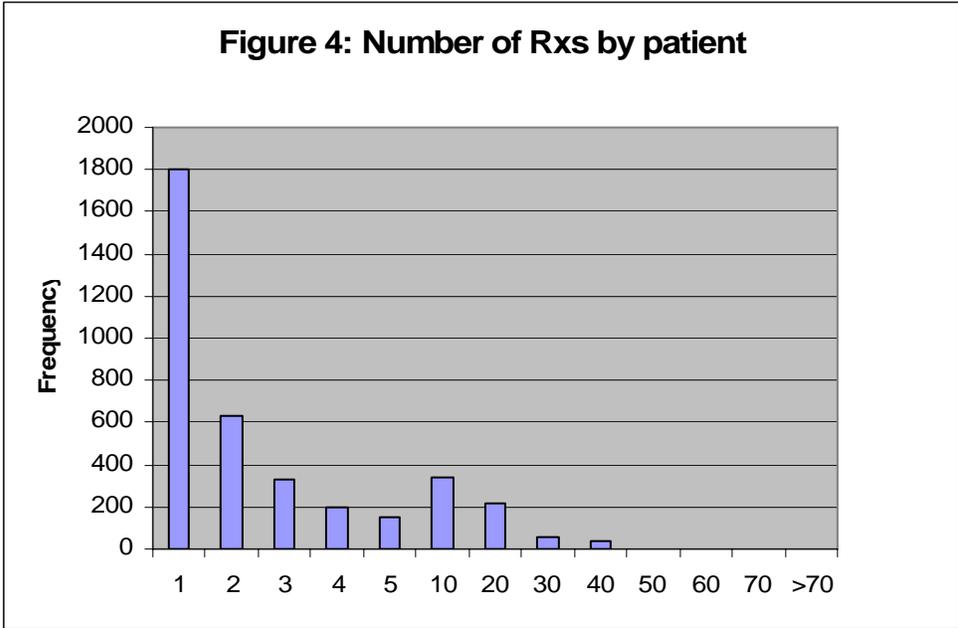
Figure 2: Trend of Detailing Expenditure (1999-2003)



Source: New Product Spectra, 1999-2003.



Source: New Product Spectra, 1999-2003.



Source: IPSOS patient diary data on Cox-2 Inhibitors.

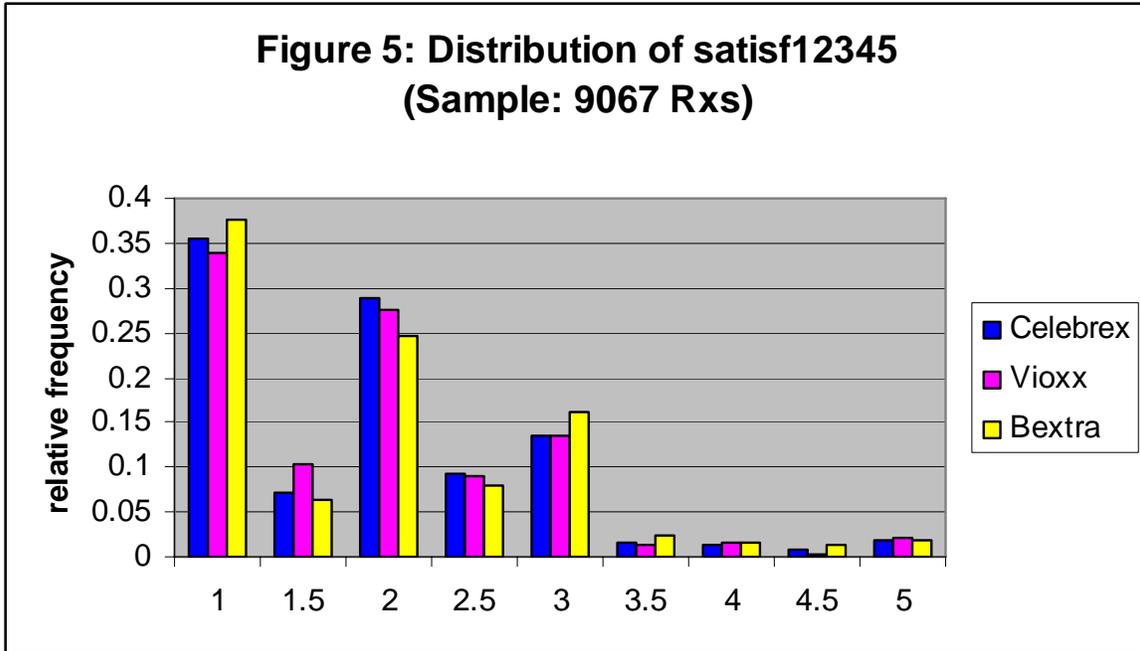


Figure 6: Histogram of residuals from the first step estimation (i.e. regress satisf12345 on a full set of drug-patient dummies)

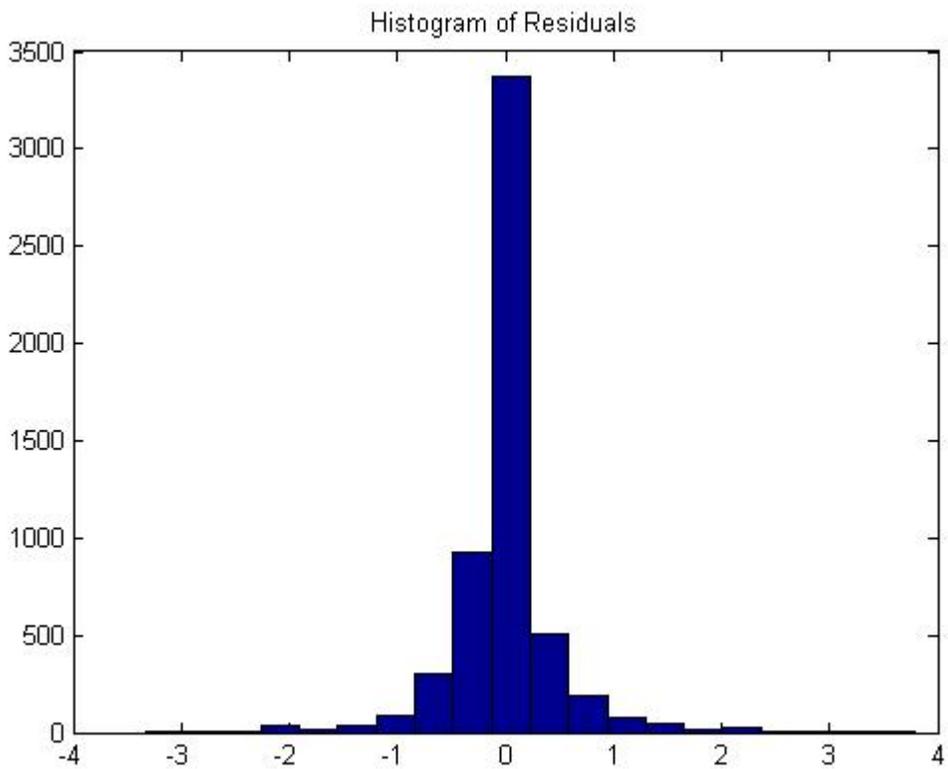


Table 1: Switching Matrix

conditional on non-missing satisfaction report, total 9,067 prescriptions

	Celebrex	Vioxx	Bextra	Switching rate
Celebrex	3036	197	64	7.92%
Vioxx	197	2506	69	9.60%
Bextra	34	18	434	10.70%
All three				8.83%

Note: Switches are from row to column.

Table 2: Summary of Satisfaction Scores conditional on non-missing
(total 9,067 prescriptions)

(1=extremely satisfied, 5=extremely dissatisfied)

	Celebrex	Vioxx	Bextra
Effectiveness	1.840	1.841	1.949
Side effects	1.839	1.845	1.835
Works quickly	1.975	1.961	2.012
How long does it last?	1.973	1.969	2.003
Easy to take	1.397	1.353	1.414
Average effectiveness (satisf134)	1.929	1.924	1.988
Average across five (satisf12345)	1.805	1.794	1.843

Table 3: Logit model on brand switching

Dependent Variable: switch=1 if switch brand from the last period

Independent Variable: satisfaction scores (1=extremely satisfied, 5=extremely dissatisfied)

	Coefficients	t Stat	
Drug effectiveness (satisf134)	0.2545	4.0345	***
Side Effects	0.0548	0.9332	
Ease to Take	-0.0341	-0.5909	
Intercept	-2.8992	-29.0476	****
OBS: 7,953 prescriptions on 1,333 patients			

Source: IPSOS patient diary data, Jan. 2001 – Aug. 2003. *** p<0.01.

Table 4: Relating satisfaction and advertising data with the number of new patients

Dependent Variable = the number of new patients that first appear in month t and start with drug j

	Coefficients	t Stat	
Intercept	46.5466	2.2143	**
Satisf12345_lag	-19.3156	-1.7053	*
DTCA_lag	9.4222	3.1997	***
Detailing_lag	-4.8884	-0.7777	
Journal Advertising_lag	-25.9789	-0.5328	
Free Samples_lag	38.4098	1.5816	
OBS (month-drug)	79		
R2	0.3674		

Source: IPSOS patient diary data, Jan. 2001 to Aug. 2003.

Table 5: Benchmark models -- Discrete Choice Model without Learning Structure

	(1)	(2)
Dummy of Celebrex	-1.2584 (3.2144)	-2.1166 (3.2165)
Dummy of Bextra	-10.7258 ** (4.1557)	-1.0962 (4.0269)
(6-Satisf12345) for Celebrex	0.2933 *** (0.0636)	
(6-Satisf12345) for Vioxx	0.2134 *** (0.0579)	
(6-Satisf12345) for Bextra	1.7873 *** (0.1821)	
Log Cum DTCA for Celebrex	-0.5164 (0.8296)	-0.4868 (0.8250)
Log Cum DTCA for Vioxx	-0.7381 (1.0956)	-0.5631 (1.0885)
Log Cum DTCA for Bextra	0.1949 (0.3509)	0.5931 * (0.3337)
Log Cum Detailing for Celebrex	-0.1874 (1.8953)	0.7747 (1.9183)
Log Cum Detailing for Vioxx	-0.2122 (1.1399)	0.3506 (1.1502)
Log Cum Detailing for Bextra	0.085 (0.2787)	0.2767 (0.2889)
Patient age * Celebrex	0.001 (0.0020)	0.0014 (0.0020)
Patient age * Bextra	-0.0063 (0.0037)	-0.0049 (0.0036)
Patient female * Celebrex	0.2235 *** (0.0603)	0.2126 *** (0.0601)
Patient female * Bextra	-0.2242 (0.1136)	-0.2729 ** (0.1121)
Log L	-5008.7	-5071.9
# of patients	2,062	2,062
# of Rxs	5,688	5,688

Notes: The default drug is Vioxx. Standard errors in parentheses. Satisfaction is measured by 6-satisf12345, computed as the average of all patient satisfaction up to the month before prescription. Advertising variables are measured as the log of cumulative sums of advertising expenditures up to the previous month. ***p<0.01, **p<0.05, * p<0.1.

Table 6: Results with Learning Structure

	(1) Risk Neutral		(2) Risk Neutral Across- patient learning only		(3) Risk Neutral within- patient learning only	
α_0	-8.1931 *** (3.0990)		-471.0103 *** (99.5136)		-4.4973 ** (2.0580)	
α_R	2.0675 *** (0.4682)		112.2335 *** (22.3689)		2.1473 *** (0.4957)	
σ_v	0.4960		0.4960		0.4960	
Q0_celebrex	-0.1974 0.7425		0.3003 (1.0350)		-0.2760 (0.6342)	
Q0_bextra	-1.3771 * 0.7934		1.2422 (0.8542)		-2.1873 *** (0.6965)	
σ_{Q0} celebrex	0.0270 *** (0.0064)		0.0002 *** (0.0000)			
σ_{Q0} vioxx	0.0269 *** (0.0065)		0.0002 *** (0.0000)			
σ_{Q0} bextra	0.0398 *** (0.0097)		0.0010 *** (0.0002)			
σ_{q0}	0.3068 *** (0.0703)				0.2682 *** (0.0626)	
Log cum DTCA	-0.3246 *** (0.0603)		0.5632 *** (0.1310)		-0.4522 *** (0.0596)	
Log cum Detailing	0.1340 (0.1371)		-0.2806 * (0.1542)		0.5680 *** (0.1066)	
Patient Age * Celebrex	0.0079 *** (0.0028)		0.0013 (0.0019)		0.0076 *** (0.0028)	
Patient Age * Bextra	0.0000 (0.0049)		-0.0049 (0.0036)		0.0007 (0.0049)	
Patient Female * Celebrex	0.1391 * (0.0825)		0.2253 *** (0.0592)		0.1390 * (0.0837)	
Patient Female * Bextra	-0.2714 * (0.1475)		-0.2678 ** (0.1127)		-0.2804 * (0.1470)	
Log Likelihood	-2738.1		-5036.5		-2816.7	
# of patients	2062		2062		2062	
# of Rxs	5688		5688		5688	

Notes: The default drug is Vioxx. Standard errors in parentheses. Satisfaction is measured by 6-satisf12345. Advertising variables are measured as the log of cumulative sums of advertising expenditures up to the previous month. Across-patient learning is assumed to be within each of the 9 census regions. ***p<0.01, **p<0.05, * p<0.1.