

Learning About New Products: An Empirical Study of Physicians' Behavior

Maria Marta Ferreyra* and Grigory Kosenok†

September 24, 2009

Abstract

We develop and estimate a model of market demand for a new pharmaceutical, whose quality is learned through prescriptions by forward-looking physicians. We use a panel of anti-ulcer prescriptions from Italian physicians between 1990 and 1992 and focus on a new molecule available since 1990. We solve the model by calculating physicians' optimal decision rules as functions of their beliefs about the new pharmaceutical. According to our counterfactuals, physicians' initial pessimism and uncertainty can have large, negative effects on their propensity to prescribe the new drug and on expected health outcomes. In contrast, subsidizing the new good can mitigate informational losses.

Keywords: learning, pharmaceutical demand, dynamic discrete choice (JEL I10, L10)

*Tepper School of Business, Carnegie Mellon University, e-mail: mferrey@andrew.cmu.edu

†New Economic School, Moscow, e-mail: gkosenok@nes.ru

We thank Peter Arcidiacono, Daniele Coen-Pirani, Eric French, George-Levy Gayle, Marty Gaynor, Ron Goettler, Susumu Imai, Ken Judd, Yuichi Kitamura, Steve Lehrer, Robert Miller, Alvin Murphy, Sridhar Narayanan, Mark Roberts, Marc Rysman, and especially Phil Haile for helpful comments on various versions of this paper. Seminar participants at Carnegie Mellon, the New Economic School, Queen's, Stanford, Washington University and Wisconsin, and participants at the 2009 Summer Meetings of the Econometric Society also made valuable comments. We gratefully acknowledge Jeff Reminga's assistance with the computational aspects of this project, and Alexander Migita's research assistance. We especially thank Matthew Shum and Andrea Coscelli for giving us access to the data. All errors are ours.

1 Introduction

An important aspect of economic innovation is the development of new products. Thus, firms often devote large amounts of resources to research and development. The pharmaceutical industry, in particular, is characterized by a high degree of product innovation and significant investments in research and development.¹ Critical to the adoption of a new pharmaceutical, however, is the willingness of physicians to prescribe it. Although clinical trials and advertising convey information about the quality of a new pharmaceutical, considerable uncertainty remains not only about its general quality, but also about its effect on individual patients.² Since physicians ultimately resolve this uncertainty by prescribing the new drug³ to their patients -namely, by experimentation- the adoption of a new pharmaceutical is closely related to the experimentation process by which physicians gain information.

In this paper we focus on the entry of a new anti-ulcer drug, omeprazole, in the Italian pharmaceutical market in June of 1990.⁴ Our data consist of monthly prescriptions of anti-ulcer medications by physicians in the Rome metropolitan area between June 1990 and December 1992. We develop and estimate a dynamic discrete choice model of physician Bayesian learning about the quality of a new drug. In our model, physicians have initial beliefs about the new drug's efficacy, and update them every time that they prescribe the new drug and observe the treatment outcome. When choosing which drug to prescribe to a particular patient, a forward-looking physician considers not only the expected outcome for that patient, but also the opportunity to learn about the new drug's quality and use the information on future patients. Our parameter estimates allow us to disentangle how uncertainty about the new drug's quality, intertemporal preferences, and prices affect the

¹For instance, the US Congressional Budget Office (2006) reports that in 2003, pharmaceutical companies devoted 37.6 billion dollars to research and development of new products.

²As the US Food and Drug Administration has noted (see <http://www.fda.gov/cder/reports/rtn/2005/rtn2005-3.htm>), the limited size of pre-marketing clinical trials means that not everything can be learned about a medicine's safety before the medicine is approved. As a result, uncertainty always exists, at least to a certain extent, both before and after the drug's approval.

³In this paper we use "drug" to refer to molecule or active principle. Thus, a new drug represents a therapeutically new treatment.

⁴In the United States, this drug was first marketed by AstraZeneca in 1989 under the brand names of Losec and Prilosec.

adoption of the new drug and the corresponding health outcomes.

A recent literature has explored the role of learning about pharmaceuticals.⁵ As noted by Manchanda et al (2005) in their review of this literature, some important aspects of the learning mechanism are the channels through which agents learn (prescription, detailing, marketing meetings, and events), their preferences toward risk (risk-averse or risk-neutral), and whether they are short-sighted (myopic) or forward-looking. Myopic agents maximize their current expected utility, while forward-looking agents evaluate their current utility as well as the expected future value of the information they might gain through the current prescription.⁶ The studies also differ in whether they use physician-level or aggregate data.

A central contribution of our paper is the prediction of market demand for the new drug through a parsimonious yet behaviorally rich model. Using the same physician-level data used here, Coscelli and Shum (2004) model the behavior of doctors who learn in a Bayesian fashion yet are myopic. We, in contrast, model the forward-looking behavior of physicians, and reject the null hypothesis that physicians are myopic. According to our estimates, myopic behavior significantly delays the adoption of the new drug, and generates large losses to patients' health. Crawford and Shum (2005) use related patient-level data, and model how forward-looking patients learn about their curative and symptomatic matches with respect to alternative medications. In their model, a patient's learning about a medication is not driven by the fact that the medication is new; rather, it is driven by her need to learn which medication suits her best. Hence, she learns about her match to each medication, not just the new one. In our model, however, learning is driven by the market-

⁵See, for instance, Ching (2008), Chintagunta et al (2009), Coscelli and Shum (2004), Crawford and Shum (2005), Currie and Park (2002), Le Cates et al (2008), Mukherjee (2002), Narayanan et al (2005), and Narayanan and Manchanda (2007). Berndt et al (2003) studied the diffusion of anti-ulcer drugs based on consumption externalities, although they did not specify a behavioral learning model.

⁶In industrial organization and marketing, empirical models of learning in dynamic frameworks have also been applied to other products, such as yogurt (Ackerberg 2003), breakfast cereal (Eckstein et al 1998), laundry detergents (Erdem and Keane 1996), online grocer (Goettler and Clay 2009), and telephone consumption (Narayanan et al 2007). Previously, Miller (1984) had estimated a dynamic Bayesian learning model of job-worker matching. Miller modeled a multi-armed bandit problem in which the return to every option was uncertain, whereas in our case only the return to the new drug is uncertain.

level uncertainty associated with omeprazole as a new product.⁷ While it would be of interest to study both doctors' and patients' learning, our doctor-level data motivate our focus on physician learning. We believe that focusing on the physician is useful, since in practice he chooses which drug to prescribe.⁸ In addition, as he gains knowledge through his experience with one patient, he can apply it to other patients. Finally, Narayanan and Manchanda (2007) have studied the role of prescription experience, detailing, meetings, and events using physician-level data. In their model physicians learn from two sources (detailing and experimentation), yet are myopic.⁹ While the frameworks provided by Coscelli and Shum (2004) and Narayanan and Manchanda (2007) allow for the prediction of market demand, ours is more parsimonious. In addition, physicians in our model are forward-looking.

Whereas the combination of physician-level data and a forward-looking setting lends great power to our framework, it also poses considerable challenges. In the model, we assume that when a physician sees a patient, he observes some aspects of her condition. This condition is one of the state variables of the problem, as are the mean and variance of the physician's beliefs about the true quality of new drug, and the price difference between the new and the incumbent drug. The physician's optimal decision rule is to choose the drug with the largest expected discounted value given the state variables.

Researchers typically calculate value functions using iterative algorithms. Continuous state variables are discretized by using a grid with a finite number of points, with the number of grid points increasing with the number of variables and with the desired level of accuracy. Iterative algorithms working over grids are computationally demanding even for a small number of state variables. In our case, to estimate the model we must calculate the

⁷Another difference between Crawford and Shum (2005) and our work is the modeling of beliefs. Crawford and Shum (2005) assume rational expectations, which means that while a doctor does not know a patient's match with a given medication, his beliefs about her match correspond to the actual distribution of matches in the population. We, in contrast, allow prior beliefs to differ from true quality. See footnote 21 for a related discussion. To accommodate for the fact that omeprazole is a new drug, Crawford and Shum allow the population distribution of matches to omeprazole to change several times over the sample period. They do not model these changes as the outcome of a learning process, whereas we model all changes in beliefs as the outcome of learning.

⁸Direct-to-consumer advertising of pharmaceuticals is only allowed in the United States and New Zealand.

⁹We do not examine detailing in this paper because we have no data on it.

probability of the observed prescription choices (approximately half a million in our dataset) for each parameter point. This poses severe computational challenges.

Faced with these challenges, we exploit theoretical properties of our model and some features of our data to reduce the physician’s dynamic discrete choice problem to the computation of a threshold value for the patient’s condition, net of price difference disutility. The threshold is a function of only two state variables - the mean and variance of physician’s beliefs about the new drug’s true quality. In particular, the threshold is not a function of continuation payoffs. A patient whose net condition is above the threshold is prescribed the new medication, and a patient whose net condition is below the threshold is prescribed the old medication. Exploiting properties of physicians’ Bayesian learning, we have developed a recursive algorithm that finds threshold values for one grid point at a time. Calculating thresholds is thus straightforward and very fast, which leads to dramatic savings in computational time.

Our approach is similar to Rust (1987) because his assumptions of Additive Separability and Conditional Independence hold in our model. If we were to apply Rust’s approach to our problem, these assumptions would enable us to eliminate one dimension of integration when calculating continuation payoffs by assuming a type I extreme value distribution for the patient’s condition. Given our assumption that the physician faces only two choices, we are able to eliminate one dimension of integration for *any* distribution of the patient’s condition, by exploiting the *shape* of the value function with respect to the patient’s condition. Since the physician does not update his beliefs when he prescribes the old medication, we simplify further by avoiding the calculation of continuation payoffs.

Having a parsimonious, rich and computationally tractable model facilitates the prediction of market demand for the new medication. Furthermore, a similar model may be applied to study demand for other experience goods. Whereas the institutional setting in the Italian pharmaceutical market (see next section) has allowed us to abstract away from the supply side without biasing our estimates, in other settings this may not be possible, although computational considerations often preclude the joint study of dynamic demand and supply.¹⁰ By building tractability on the demand side, our approach could facilitate

¹⁰For instance, Ching (2008b) studies the dynamic of the demand for prescription drugs after patent expi-

more comprehensive studies of dynamic demand and supply, and may be applicable to other dynamic discrete choice problems.

Using our parameter estimates, we conduct several counterfactuals. These indicate that uncertainty about the new drug’s quality affects physician prescription behavior, because a physician who does not know the drug’s quality does not prescribe it to all the patients who need it. This leads to inferior expected health outcomes for patients. More damaging than uncertainty, however, is physicians’ myopic behavior. Since a myopic physician places no value on the learning opportunity from prescribing the new drug, he prescribes it much less, which in turn delays his learning. Of the total health-related losses inflicted by a myopic physician, his failure to experiment accounts for a larger share than his uncertainty about the new drug’s quality. The expected discounted payoff over 20 years for a representative, forward-looking physician is 9 percent lower than for a fully-informed physician, and the payoff for a myopic physician is 41 percent lower than for a fully-informed physician. Our simulations show, however, that informational losses can be mitigated by pricing policies that encourage the adoption of the new drug.

The remainder of this paper is organized as follows. Section 2 presents some descriptive statistics and institutional aspects of our data. Section 3 presents the model, and section 4 analyzes it. Section 5 outlines the estimation procedure, and section 6 describes the estimation results. Section 7 presents the counterfactuals, and section 8 concludes.

2 Data and descriptive statistics

The barrier of mucus in the stomach and duodenum has the function of protecting these organs from the powerful acids and enzymes produced by the body to digest food. A number of factors, such as infection with *Helicobacter pylori* bacterium, and long-term use of aspirins and other medicines, can damage the barrier of mucus. Damage to this barrier leads to gastric ulcers, which anti-ulcer drugs purport to cure. More generally, anti-ulcer drugs are usually prescribed for four different purposes: as treatment to pathological hypersecretory conditions, as attack therapy for gastroesophageal reflux disease (GERD) or peptic ulcer, as

ration, and Ching (2008a) studies the dynamics of the supply of prescription drugs after patent expiration.

maintenance therapy for GERD or peptic ulcer, and as treatment for minor heartburn.

The anti-ulcer market is the largest therapeutic drug market worldwide (Coscelli and Shum 2004). Most of the anti-ulcer drugs used around the world before 1990 were based on H₂-receptor antagonists. These block the production of histamine, a substance that stimulates acid secretion.¹¹ Yet the most powerful inhibitors of acid secretion currently available are proton-pump inhibitors, which completely block the production of stomach acid by stopping the final step in acid secretion (Freston 1994, De Giorgi et al 2006). The first proton-pump inhibitor was omeprazole. During its first years in the market, omeprazole was the preferred choice for pathological hypersecretory conditions and as attack therapy for GERD or peptic ulcer (Buhl and Clearfield 1990, Freston 1994, Meijer et al 1989). Omeprazole entered the Italian market in June 1990 under the brand names of Losec and Omeprazen. Coscelli and Shum (2004) document that in their sample, about 25% of the prescriptions were written out to patients who suffered from the conditions for which omeprazole is the preferred choice.

The data used in this study were collected by the Italian National Institute of Health, which recorded all prescriptions for anti-ulcer medications written by randomly sampled doctors in the Rome metropolitan area between June 1990 and December 1992. For each physician we observe the monthly number of anti-ulcer prescriptions, and the monthly number of prescriptions written out for omeprazole. As for the prescriptions written out for other medications, we do not know which particular drugs were prescribed. Since most of them would have belonged to the same pharmaceutical class (H₂-antagonists), we pool them into a single non-omeprazole alternative. We include physicians with at least ten patients in each month, and we exclude physicians who in August of 1991 or 1992 saw 80% fewer patients than their average over July and September of the corresponding year.¹² The resulting dataset contains 256 physicians and 31 months, for a total of 463,199 prescription

¹¹Brands of H₂-receptor antagonists in the US include Tagamet, Zantac, Pepcid, and Axid, among others. See Berndt et al (2003) for a study of the diffusion of H₂-receptor antagonists and inter-brand competition.

¹²Physicians with very few patients may be less interested in learning about the new drug simply because they have fewer opportunities to write prescriptions. In addition, physicians with large seasonal fluctuations may not fit the assumption of a constant patient arrival process, which we invoke in the model (see next Section).

episodes organized into 7,936 doctor-month observations.

Moreover, in Italy, all patients have the same insurance status because their costs are covered by the National Health System. The Drug Commission sets drug prices and copayments. Pharmaceuticals based on the same molecule are usually assigned the same price, and the copayment for anti-ulcer drugs is 50 percent of the price (Crawford and Shum 2005). All patients face the same prices and copays. Over the sample period, the average real price difference between omeprazole and the incumbent for a day's dose was 1,106 liras, which was equivalent to slightly less than one June 1990 dollar. The price difference experienced little variation over the sample period.

TABLE 1
Descriptive Statistics

	No. of Observations	Mean	St.Dev.	Min.	Max.
Monthly number of patients	7,936 doctor-months	58.4	27.2	10	252
Monthly number of omeprazole prescriptions	7,936 doctor-months	5.79	6.08	0	78
Monthly price difference (in Italian liras)*	31 months	1,106.0	71.8	1,001.1	1,237.2
Total number of patients	256 doctors	1,809.4	466.6	937	3,676
Total number of omeprazole prescriptions	256 doctors	179.4	111.7	24	773
Omeprazole share over the whole period	256 doctors	0.098	0.051	0.016	0.290
Month of first omeprazole prescription	256 doctors	2.75	2.03	1	9


* Average exchange rate in January 1990: 1,236 liras = 1 dollar. Price difference for one day of treatment.

Table 1 presents some summary statistics for our data. As the table shows, there is wide variation in the number of office visits received by physicians in the sample, as well as in the number and proportion of omeprazole prescriptions written over the period. Figure 1 depicts the evolution of the in-sample market share of omeprazole over the period. As the figure shows, omeprazole's share rose from 1.6 to 14 percent over the sample period. Furthermore, Coscelli and Shum (2004) document that by mid-1995, omeprazole's market share in Italy had risen to 25 percent.¹³ As Figure 1 shows, over the sample period physicians

¹³Prilosec, based on omeprazole, was the leading drug in terms of retail expenditure in 2000-2001 in the US.

gradually converged in their propensity to prescribe omeprazole. Moreover, by the third month about 75 percent of the doctors had prescribed the new drug at least once, and all doctors had prescribed it at least once by the ninth month.

FIGURE 1
Omeprazole Market Share



Omeprazole market share is the in-sample market share per month. The coefficient of variation is calculated for the monthly fraction of omeprazole prescriptions across doctors.

Unfortunately, our data are limited in several ways, all of which constrain our model and estimation. The main limitations are the following. First, we do not have patient-level information. This means that we do not know whether any given patient receives one or

Protonic-pump inhibitors accounted for more than 75 percent of all anti-ulcer retail expenditure (NIHCM, 2002). From conversations with physicians we learned that it is quite common for a new pharmaceutical to follow this type of trajectory - that is, to enter the market focused on specific uses, and then broaden the set of uses with a concomitant increase in market share. In the case of omeprazole, as time went by new studies documented its safety for longer-term treatments and a broader patient base than initially considered.

multiple prescriptions over the sample period. In the sample used by Crawford and Shum (2005), slightly more than half of the patients receives more than one prescription, which suggests that a substantial number of prescriptions in our sample might have been written out to the same patients. In the absence of patient-level information, however, we assume that each prescription is written out to a different patient.¹⁴ While this assumption may not be suited to all the prescriptions in the data, it is necessary to explain at least some of the evolution of the market share. This is because in the data, the total number of prescriptions remains roughly constant over time. Hence, a rising market share cannot be generated simply from multiple prescriptions of the new drug to the patients for which the drug has been effective; rather, it requires that additional patients be prescribed omeprazole over time. The second important limitation in our data is the absence of doctor-level covariates such as age, training or specialty to help us characterize physician prescription behavior. We do exploit, however, the substantial heterogeneity in the number of patients seen by doctor, probably attributable to differences in doctors' patient bases. Third, we do not have data on other channels of learning such as advertising, detailing, professional interactions, etc. Hence, we model a doctor's learning exclusively as a function of his experimentation with patients.

3 Model

Consider a physician who treats patients afflicted by ulcers. When he examines a patient he observes aspects of her condition and determines whether to prescribe the new or the old medication. He knows the quality of the old medication, but is uncertain about the quality of the new one. The patient returns at the end of her treatment (we assume that she complies with her treatment), and the doctor observes the treatment outcome. If she consumed the new medication, the doctor obtains new information on the new drug's quality, and uses it

¹⁴To our knowledge, no researcher has used both patient- and doctor-level information. Researchers studying doctor-level learning handle prescription episodes as independent from one another. Narayanan and Manchanda (2007) use the first prescription for each patient. Coscelli and Shum (2005) observe the diagnosis treated with each prescription and allow for within-month correlation among the prescription signals for each diagnosis to capture spillovers across patients in different diagnosis groups.

to update his beliefs about the new drug. Thus, when he sees the next patient, his beliefs on the new drug will differ from those he had when he saw the current patient. If, on the other hand, the patient consumed the old medication, his beliefs about the new drug remain unchanged. The doctor is a Bayesian learner: when he sees the outcome of the new drug on a particular patient, he views it as a signal of its true quality. His updated beliefs are a weighted average of his prior beliefs and the signal; the greater weight is given to the more accurate of these two elements.

More formally, we model two drugs in the anti-ulcer market: drug 0 - the “old” or “incumbent” drug (which, in our empirical application encompasses all pre-existing drugs), and drug 1 - the “new” drug, which enters the market at time $t = 0$. The quality of the old drug is known, but the quality of the new drug is unknown.

We assume that physicians only differ in their patients’ arrival processes.¹⁵ Consider doctor i , and let t_k be the time of arrival of his k th patient. Based on the patient’s observable condition, the doctor makes the prescription $d_{t_k}^i$, which is equal to 1 when the new drug is prescribed, and is equal to 0 otherwise. We assume that the doctor always writes a prescription because our data are based on prescription counts.

As is common in this literature, we assume that the doctor is benevolent, in the sense that there is no agency problem between him and the patient.¹⁶ Thus, doctor i ’s instantaneous utility is the same as patient k ’s. We lack the data to identify the utility from the old drug; instead, we can only identify the difference in utilities between the new and old drug. Hence, we normalize the instantaneous utility from the old drug to zero, which means that the utility from prescribing the new drug stands for the new drug’s net utility, or its

¹⁵While we do not model other sources of doctor heterogeneity, Coscelli and Shum (2005) include a doctor-specific shifter. They conclude that neither doctor heterogeneity nor time effects explain much of the observed market share. Furthermore, the variance of the shifter is not significantly different from zero in specifications that include time-specific fixed effects.

¹⁶The institutional framework of health care in Italy mitigates agency problems between doctors and patients, as each enrollee of the National Health Service must list a general practitioner (Coscelli and Shum (2005)), presumably developing a long-term relationship with her doctor. According to our estimates (see Section 6), a higher price for the new medication relative to the old one lowers the probability of prescribing the new medication. This would not necessarily be the case if physicians did not care at all about their patients.

premium over the old drug (see Appendix A1 for further details). We model instantaneous utility as follows:

$$U(d_{t_k}^i) = \begin{cases} \eta_{t_k}^i + \delta + \xi_{t_k}^i - \phi \Delta p_{t_k} & \text{if new drug } (d_{t_k}^i = 1) \\ 0 & \text{if old drug } (d_{t_k}^i = 0) \end{cases}. \quad (1)$$

In this expression, the outcome of the new drug equals $\eta_{t_k}^i + \delta + \xi_{t_k}^i$. The term $\eta_{t_k}^i$ corresponds to the patient's condition observed by the physician during the office visit. We often refer to this term as the patient's match parameter, because it indicates to the physician whether the new drug would be a good choice for the patient given her condition. For instance, consider a patient who comes to her doctor's office with a severe hypersecretory condition. Since omeprazole is the preferred treatment for pathological hypersecretory conditions, this means that she has a high η . Parameter δ is the quality of the new medication. We can also think of it as the true average effect of the new drug in the population of patients - equivalently, the unconditional average outcome in the population of patients. The term $\xi_{t_k}^i$ reflects how patient k 's outcome deviates from the new drug's mean outcome, δ . This term is realized *after* patient k takes the new medication but *before* the arrival of patient $k + 1$.¹⁷ Finally, the term $\phi \Delta p_{t_k}$ captures the disutility due to the price difference between the new and old drugs, Δp_{t_k} , with ϕ being a parameter that measures patients' price sensitivity. The variables $\eta_{t_k}^i$ and $\xi_{t_k}^i$ are random and i.i.d. across patients, and their realizations are independent. Furthermore, $\eta_{t_k}^i$ and $\xi_{t_k}^i$ are normally distributed as follows:

$$\eta_{t_k}^i \sim N(0, \sigma_\eta^2) \text{ and } \xi_{t_k}^i \sim N(0, \sigma_\xi^2) \quad (2)$$

where σ_η^2 and σ_ξ^2 are positive variances.

The price difference Δp_{t_k} is also i.i.d.,¹⁸ and distributed according to the c.d.f. function $W(\cdot)$:

$$\Pr(\Delta p_{t_k} \leq x) = W(x). \quad (3)$$

If the new drug is prescribed, then when patient k returns after her treatment the

¹⁷We assume this timing for belief updating for simplicity. Moreover, omeprazole produces results as quickly as in one day (see, for instance, <http://www.fda.gov/bbs/topics/news/2003/new00916.html>).

¹⁸Since an administrative entity sets prices in Italy, it is reasonable to assume that prices are i.i.d. from the point of view of physicians.

doctor observes the full outcome of the new drug on this patient.¹⁹ The doctor had already observed part of this outcome, $\eta_{t_k}^i$, and now observes the rest. However, the doctor cannot distinguish between δ and $\xi_{t_k}^i$ because he does not know the true quality δ . Instead, he observes a prescription signal for the new drug’s quality, equal to $\varepsilon_{t_k}^i \equiv \delta + \xi_{t_k}^i$. Thus, the doctor performs a signal-extraction exercise, as he extracts information about the new drug’s quality from the noisy signal he observes.

More specifically, we assume that at time 0 doctor i does not know the new drug’s quality. His prior beliefs about it can be described as follows:

$$\delta \sim N(\hat{\delta}_0, \hat{\sigma}_0^2) \tag{4}$$

where $\hat{\delta}_0$ and $\hat{\sigma}_0^2$ are the mean and variance of his initial beliefs. The lower the value of $\hat{\delta}_0$, the more pessimistic he initially is; the greater the value of $\hat{\sigma}_0^2$, the less accurate his initial beliefs are. Prescribing the new drug allows the doctor to update his beliefs. Moreover, we assume that experimentation (prescription) is the only source of physician learning.²⁰ Thus, when he sees patient k , doctor’s beliefs can be described as follows:

$$\delta \sim N(\hat{\delta}_{t_k}^i, [\hat{\sigma}_{t_k}^i]^2). \tag{5}$$

If he prescribes the old drug to patient k , the physician is certain that he will observe an outcome equal to zero. If he prescribes the new drug, he is uncertain about the prescription signal he will observe. However, his beliefs about the new drug lead him to perceive this

¹⁹In our model, σ_ξ^2 can be viewed as the true variance of the random portion of the outcome adjusted by the probability that the doctor sees the outcome. A high σ_ξ^2 could indicate either high volatility in the outcome or low probability that the doctor observes the outcome. These possibilities are empirically not distinguishable, though their implication is the same – other things equal, the doctor will update his beliefs more slowly than he would with a lower σ_ξ^2 .

²⁰Modeling physicians who learn from each other would prove challenging, because each physician would then have to predict the evolution of other physicians’ beliefs and, in equilibrium, these predictions should be consistent with the actual beliefs. The introduction of a simple function of time might be able to proxy for the learning based on word of mouth, professional conferences, journals and news, changes in labels and package inserts, etc. Coscelli and Shum (2004) introduce a time trend that captures learning through these sources. Their counterfactuals suggest, however, that the trend explains little of the market share’s evolution.

signal as having the following distribution:

$$\varepsilon_{t_k}^i \sim N(\hat{\delta}_{t_k}^i, [\hat{\sigma}_{t_k}^i]^2 + \sigma_\xi^2) \quad (6)$$

which follows from (2) and (5). Hence, his updating is based on the following perceived joint distribution:

$$\begin{pmatrix} \delta \\ \varepsilon_{t_k}^i \end{pmatrix} \sim N \left(\begin{pmatrix} \hat{\delta}_{t_k}^i \\ \hat{\delta}_{t_k}^i \end{pmatrix}, \begin{bmatrix} [\hat{\sigma}_{t_k}^i]^2 & [\hat{\sigma}_{t_k}^i]^2 \\ [\hat{\sigma}_{t_k}^i]^2 & [\hat{\sigma}_{t_k}^i]^2 + \sigma_\xi^2 \end{bmatrix} \right).$$

This means that the mean and variance of the doctor's posterior, updated beliefs are as follows:

$$\hat{\delta}_{t_{k+1}}^i = \frac{\sigma_\xi^2}{[\hat{\sigma}_{t_k}^i]^2 + \sigma_\xi^2} \hat{\delta}_{t_k}^i + \frac{[\hat{\sigma}_{t_k}^i]^2}{[\hat{\sigma}_{t_k}^i]^2 + \sigma_\xi^2} \varepsilon_{t_k}^i \quad \text{and} \quad [\hat{\sigma}_{t_{k+1}}^i]^2 = \frac{\sigma_\xi^2}{[\hat{\sigma}_{t_k}^i]^2 + \sigma_\xi^2} [\hat{\sigma}_{t_k}^i]^2. \quad (7)$$

These will be the doctor's prior beliefs the next time he sees a patient. No updating takes place if the old drug is prescribed, or

$$\hat{\delta}_{t_{k+1}}^i = \hat{\delta}_{t_k}^i \quad \text{and} \quad [\hat{\sigma}_{t_{k+1}}^i]^2 = [\hat{\sigma}_{t_k}^i]^2. \quad (8)$$

We assume that doctors are forward-looking and risk-neutral.²¹ Next, let \varkappa be the rate of time discounting, common across all physicians. Then, if patients arrive at times $t_1, t_2, \dots, t_k, \dots$, where the subindex stands for the arrival order of a patient, doctor i 's objective is to maximize the expectation of his discounted utility or

$$W^i = E_{t=0} \left\{ \sum_{k=1}^{+\infty} e^{-\varkappa t_k} U(d_{t_k}^i) \right\} \rightarrow \max_{\{d_{t_k}^i\}} \quad (9)$$

where $E_{t=0} \{.\}$ stands for the doctor's expectation at time zero, given his initial beliefs (4) and future optimal choices $d_{t_k}^i$, based on the information available at time t_k ($\hat{\delta}_{t_k}^i, [\hat{\sigma}_{t_k}^i]^2$,

²¹It is not possible to identify risk aversion in this setting. The logic is quite similar to that of Coscelli and Shum (2004)'s discussion of lack of identification of risk aversion in a myopic setting. Intuitively, a low initial market share could be rationalized either by risk-neutral physicians who are pessimistic about the new drug, or by risk-averse physicians who know the new drug's quality but are uncertain about it (i.e., physicians who have $\hat{\delta}_0 = \delta$ but $\hat{\sigma}_0^2 > 0$). We chose the first alternative because it seems reasonable to assume that physicians would not know the true quality of the new drug on the full population of patients just from clinical trials conducted before the drug's entry.

Δp_{t_k} and $\eta_{t_k}^i$). At times we will refer to the doctor’s (expected) discounted utility as the doctor’s payoff or value.

Since the instantaneous utility in (1) may be linearly transformed without affecting optimal choices, we adopt a normalization for the sake of identification. In particular, we impose the following restriction:

$$\sigma_\eta^2 \equiv 1$$

or

$$\eta_{t_k}^i \sim N(0, 1). \tag{10}$$

Thus, we measure δ , ξ and $\phi\Delta p$ in units of standard deviation of the match parameter η .

Finally, we assume that the time elapsed between two consecutive patients’ arrivals, $\Delta t_k = t_k - t_{k-1}$ follows a Gamma distribution,²² and is independent of the time elapsed between any other consecutive patients and all other random variables:²³

$$\Pr\{\Delta t_k < x\} = T(x; a_i, b_i) = \int_0^x \frac{1}{b_i \Gamma(a_i)} \left(\frac{y}{b_i}\right)^{a_i-1} e^{-\frac{y}{b_i}} dy, \tag{11}$$

In this doctor-specific distribution, a_i is the “shape” parameter of the distribution, and parameter b_i is associated with the intensity of the arrival process - the lower the b_i , the more intensive the arrival rate.

Having stated the primitives of the model, in the next section we characterize the optimal behavior of a doctor.²⁴

²²A Poisson arrival process generates an exponential distribution for the time between consecutive arrivals. An exponential distribution is a special case of a Gamma distribution.

²³We assume that the arrival process is exogenous from the point of view of the physician. It is not affected by treatment outcomes, or by the entry of the new drug (the number of patients seen by each doctor does not change much over the sample period). For the arrival process we have assumed, the mean arrival time is equal to $a_i b_i$.

²⁴One could think of the following alternative model: at the beginning of the month, the doctor chooses the share of patients who receive the new drug; at the end of the month, he updates his beliefs based on all the signals received throughout the month. We do not think this model is plausible, as the doctor would use information from the signals only at the end of the month, instead of using it as soon as it becomes available. In addition, choosing the share of patients who receive the new drug at the beginning of the month would introduce new computational challenges. For instance, if the physician were considering the option of prescribing the new drug to 5 out of 10 patients in the next month, in order to calculate the continuation value from this choice he would need to integrate over 5 unobserved signals.

4 Analysis of the model

To simplify the exposition, in what follows we suppress the doctor index, i . We use t , and $t + \Delta t$, to denote the time of the current and next patient's arrival, respectively. Consider the problem of a physician who sees a patient at time t . He must make a decision based on four state variables: $\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t$ and η_t . Let $V(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t, \eta_t)$ be the doctor's expected continuation payoff at time t when he behaves optimally. Given (9), (7) and (8) the following *Bellman Equation* holds:

$$\begin{aligned}
 V(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t, \eta_t) = & \tag{12} \\
 \max \left\{ \int \int \int e^{-\kappa \Delta t} V(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_{t+\Delta t}, \eta_{t+\Delta t}) dF(\eta_{t+\Delta t}) dW(\Delta p_{t+\Delta t}) dT(\Delta t; a, b), \right. \\
 & \eta_t + \hat{\delta}_t - \phi \Delta p_t + \\
 & \left. \int \int \int \int e^{-\kappa \Delta t} V(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_{t+\Delta t}^2, \Delta p_{t+\Delta t}, \eta_{t+\Delta t}) \right. \\
 & \left. \times dF(\eta_{t+\Delta t}) dW(\Delta p_{t+\Delta t}) dG(\varepsilon_t | \hat{\delta}_t, \hat{\sigma}_t^2) dT(\Delta t; a, b) \right\}
 \end{aligned}$$

$$\text{s.t. } \hat{\delta}_{t+\Delta t} = \frac{\sigma_\xi^2}{\hat{\sigma}_t^2 + \sigma_\xi^2} \hat{\delta}_t + \frac{\hat{\sigma}_t^2}{\hat{\sigma}_t^2 + \sigma_\xi^2} \varepsilon_t \text{ and } \hat{\sigma}_{t+\Delta t}^2 = \frac{\sigma_\xi^2}{\hat{\sigma}_t^2 + \sigma_\xi^2} \hat{\sigma}_t^2, \tag{13}$$

where c.d.f.'s $F(\cdot)$, $W(\cdot)$, $G(\cdot | \cdot, \cdot)$ and $T(\cdot; \cdot, \cdot)$ are given by (10), (3), (6) and (11) respectively. In this equation, the max function has two arguments. The first is the expected payoff to prescribing the old drug. When he prescribes the old drug, the doctor gets zero instantaneous utility and does not update his beliefs; hence, his expected utility for the next patients depends on the same beliefs he has for the current patient. The second argument is the expected payoff to prescribing the new drug. This is equal to the subjective expectation of the instantaneous utility given by (1), plus the expected discounted utility for the next patients. The latter depends on beliefs which will be updated when the doctor sees the outcome of the new drug on the current patient. At the time he writes the prescription, the doctor does not know what his posterior beliefs will be, though he knows their distribution.

The Bellman equation can be solved numerically by solving iteratively for the value function. Since the state space is four-dimensional, the integrations required are computationally very costly for a reasonable degree of accuracy. Nonetheless, our problem has

certain properties that enable us to solve for the doctor's optimal behavior in a computationally simpler way without sacrificing accuracy. We provide the details below.

Since Δt is independent of η , Δp and ε , in (12) we can factor out the integration over Δt and rewrite the Bellman equation in the following, more familiar form:

$$V\left(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t, \eta_t\right) = \max \left\{ \beta \int \int V\left(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_{t+\Delta t}^2, \Delta p_{t+\Delta t}, \eta_{t+\Delta t}\right) dF(\eta_{t+\Delta t}) dW(\Delta p_{t+\Delta t}), \right. \\ \left. \eta_t + \hat{\delta}_t - \phi \Delta p_t + \beta \int \int \int V\left(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_{t+\Delta t}^2, \Delta p_{t+\Delta t}, \eta_{t+\Delta t}\right) dF(\eta_{t+\Delta t}) dW(\Delta p_{t+\Delta t}) dG(\varepsilon_t | \hat{\delta}_t, \hat{\sigma}_t^2) \right\} \quad (14)$$

subject to (13), where

$$\beta = \int e^{-\varkappa \Delta t} dT(\Delta t; a, b) = (1 + \varkappa b)^{-a}$$

is the doctor-specific effective time discount factor between adjacent patient's arrivals. The fact that the effective discount factor differs across doctors implies that the value function differs across doctors even for the same $(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t, \eta_t)$ combination.

It is now convenient to introduce the following notations for the integral terms in (14):

$$C_o\left(\hat{\delta}_t, \hat{\sigma}_t^2\right) = \beta \int \int V\left(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_{t+\Delta t}^2, \Delta p_{t+\Delta t}, \eta_{t+\Delta t}\right) dF(\eta_{t+\Delta t}) dW(\Delta p_{t+\Delta t}) \quad (15)$$

and

$$C_n\left(\hat{\delta}_t, \hat{\sigma}_t^2\right) = \beta \int \int \int V\left(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_{t+\Delta t}^2, \Delta p_{t+\Delta t}, \eta_{t+\Delta t}\right) dF(\eta_{t+\Delta t}) dW(\Delta p_{t+\Delta t}) dG(\varepsilon_t | \hat{\delta}_t, \hat{\sigma}_t^2) \\ \text{s.t. (13),} \quad (16)$$

where $C_o\left(\hat{\delta}_t, \hat{\sigma}_t^2\right)$ and $C_n\left(\hat{\delta}_t, \hat{\sigma}_t^2\right)$ are the doctor's continuation payoffs to prescribing the old and new drug, respectively, given his system of beliefs $(\hat{\delta}_t, \hat{\sigma}_t^2)$.

Now we can write the doctor's value function in the following parsimonious form:

$$V\left(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t, \eta_t\right) = \max \left\{ C_o\left(\hat{\delta}_t, \hat{\sigma}_t^2\right), \eta_t + \hat{\delta}_t - \phi \Delta p_t + C_n\left(\hat{\delta}_t, \hat{\sigma}_t^2\right) \right\}. \quad (17)$$

As (17) suggests, there is a value of the patient's match parameter η_t that makes the doctor indifferent between the old and the new drug given the observed price difference Δp_t . Thus, the doctor's optimal behavior consists of the following threshold rule: prescribe the new drug when the patient's match parameter adjusted for price, $\eta_t - \phi \Delta p_t$, is higher than

the threshold value $\omega(\hat{\delta}_t, \hat{\sigma}_t^2)$, and prescribe the old drug otherwise. According to (17), the threshold value is given by

$$\omega(\hat{\delta}_t, \hat{\sigma}_t^2) = -\hat{\delta} + C_o(\hat{\delta}_t, \hat{\sigma}_t^2) - C_n(\hat{\delta}_t, \hat{\sigma}_t^2). \quad (18)$$

The threshold value allows us to express the value function as follows:²⁵

$$V(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t, \eta_t) = \begin{cases} C_o(\hat{\delta}_t, \hat{\sigma}_t^2) & \text{if } \eta_t - \phi \Delta p_t < \omega(\hat{\delta}_t, \hat{\sigma}_t^2) \\ \eta_t + \hat{\delta}_t - \phi \Delta p_t + C_n(\hat{\delta}_t, \hat{\sigma}_t^2) & \text{if } \eta_t - \phi \Delta p_t \geq \omega(\hat{\delta}_t, \hat{\sigma}_t^2) \end{cases}. \quad (19)$$

FIGURE 2

Value Function and Threshold Rule




Figure 2 illustrates the value function with respect to the patient's match parameter η_t . The function is flat for all values of η_t below the threshold, and linearly increasing (with unit slope) otherwise. In other words, the threshold-rule property lends a very convenient shape to our value function. We exploit this shape when we substitute (19) into (15) and integrate over $\eta_{t+\Delta t}$ to get

$$C_o(\hat{\delta}_t, \hat{\sigma}_t^2) = \beta \int \left\{ \Phi \left[\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right] C_o(\hat{\delta}_t, \hat{\sigma}_t^2) + f \left[\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right] + \left(\hat{\delta}_t - \phi \Delta p_{t+\Delta t} + C_n(\hat{\delta}_t, \hat{\sigma}_t^2) \right) \left(1 - \Phi \left[\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right] \right) \right\} dW(\Delta p_{t+\Delta t})$$

²⁵The event $\eta_t - \phi \Delta p_t = \omega(\hat{\delta}_t, \hat{\sigma}_t^2)$ takes place with zero probability, so the doctor's choice in this case has no effect on the value function. We assume that he prescribes the new drug in this case.

where $\Phi[\cdot]$ and $f[\cdot]$ are the c.d.f. and p.d.f. of the standard normal distribution, respectively.

Next, according to (18) we can exclude $C_n(\hat{\delta}, \hat{\sigma}^2)$ in the above equation as follows:

$$\begin{aligned} C_o(\hat{\delta}_t, \hat{\sigma}_t^2) &= \beta \int \left\{ \Phi \left[\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right] C_o(\hat{\delta}_t, \hat{\sigma}_t^2) + f \left[\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right] \right. \\ &\quad \left. + \left(C_o(\hat{\delta}_t, \hat{\sigma}_t^2) - \left[\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right] \right) \left(1 - \Phi \left[\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right] \right) \right\} \\ &\quad dW(\Delta p_{t+\Delta t}) \end{aligned}$$

This gives us the following one-to-one relationship between $C_o(\hat{\delta}_t, \hat{\sigma}_t^2)$ and $\omega(\hat{\delta}_t, \hat{\sigma}_t^2)$:

$$C_o(\hat{\delta}_t, \hat{\sigma}_t^2) = \frac{\beta}{1-\beta} \int \mathcal{F} \left(\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t} \right) dW(\Delta p_{t+\Delta t}) \quad (20)$$

where the function $\mathcal{F}(\cdot)$ is defined as follows:²⁶

$$\mathcal{F}(x) \equiv f(x) - \Phi(-x)x.$$

From (16), (15) and (20) it follows that

$$\begin{aligned} C_n(\hat{\delta}_t, \hat{\sigma}_t^2) &= \int C_o(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_{t+\Delta t}^2 \mathcal{F}) dG(\varepsilon_t | \hat{\delta}_t, \hat{\sigma}_t^2) \quad (21) \\ &= \frac{\beta}{1-\beta} \int \int \mathcal{F} \left(\omega(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_{t+\Delta t}^2) + \phi \Delta p_{t+\Delta t} \right) dW(\Delta p_{t+\Delta t}) dG(\varepsilon_t | \hat{\delta}_t, \hat{\sigma}_t^2) \\ &\quad \text{s.t. (13)}. \end{aligned}$$

Equations (13) and (6) imply that at time t the physician perceives the distribution of $\hat{\delta}_{t+\Delta t}$ as follows:

$$\hat{\delta}_{t+\Delta t} \sim N \left(\hat{\delta}_t, \frac{\hat{\sigma}_t^4}{\hat{\sigma}_t^2 + \sigma_\xi^2} \right). \quad (22)$$

Hence we can rewrite (21) as

$$C_n(\hat{\delta}_t, \hat{\sigma}_t^2) = \frac{\beta}{1-\beta} \int \int \mathcal{F} \left(\omega \left(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_t^2 \sigma_\xi^2 / (\hat{\sigma}_t^2 + \sigma_\xi^2) \right) + \phi \Delta p_{t+\Delta t} \right) dW(\Delta p_{t+\Delta t}) dH(\hat{\delta}_{t+\Delta t} | \hat{\delta}_t, \hat{\sigma}_t^2) \quad (23)$$

where $dH(\hat{\delta}_{t+\Delta t} | \hat{\delta}_t, \hat{\sigma}_t^2)$ represents the distribution given by (22).

²⁶The function $\mathcal{F}(x)$ has the following properties: $\mathcal{F}(x) > 0$, $\mathcal{F}'(x) < 0$, $\mathcal{F}(x) \approx -x$ when x converges to $-\infty$ and $\lim_{x \rightarrow +\infty} \mathcal{F}(x) = 0$.

By substituting (20) and (23) into (18) we get the following expression for the threshold function $\omega(\cdot, \cdot)$:

$$\begin{aligned} \omega(\hat{\delta}_t, \hat{\sigma}_t^2) &= -\hat{\delta}_t + \frac{\beta}{1-\beta} \int \mathcal{F}\left(\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t}\right) dW(\Delta p_{t+\Delta t}) \\ &\quad - \frac{\beta}{1-\beta} \int \int \mathcal{F}\left(\omega(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_t^2 \sigma_\xi^2 / (\hat{\sigma}_t^2 + \sigma_\xi^2)) + \phi \Delta p_{t+\Delta t}\right) \\ &\quad dW(\Delta p_{t+\Delta t}) dH(\hat{\delta}_{t+\Delta t} | \hat{\delta}_t, \hat{\sigma}_t^2). \end{aligned} \quad (24)$$

As pointed out before, in our data the price difference exhibits very little variation over time. This allows for further simplification of (24) since the following approximation is accurate:

$$\int \mathcal{F}\left(\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \Delta p_{t+\Delta t}\right) dW(\Delta p_{t+\Delta t}) \cong \mathcal{F}\left(\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \overline{\Delta p}\right)$$

where $\overline{\Delta p}$ is the mean of price difference. Hence, we use the following equation for the threshold function:²⁷

$$\begin{aligned} \omega(\hat{\delta}_t, \hat{\sigma}_t^2) &= -\hat{\delta}_t + \frac{\beta}{1-\beta} \mathcal{F}\left(\omega(\hat{\delta}_t, \hat{\sigma}_t^2) + \phi \overline{\Delta p}\right) \\ &\quad - \frac{\beta}{1-\beta} \int \mathcal{F}\left(\omega(\hat{\delta}_{t+\Delta t}, \hat{\sigma}_t^2 \sigma_\xi^2 / (\hat{\sigma}_t^2 + \sigma_\xi^2)) + \phi \overline{\Delta p}\right) dH(\hat{\delta}_{t+\Delta t} | \hat{\delta}_t, \hat{\sigma}_t^2). \end{aligned} \quad (25)$$

Thus, we have reduced the original Bellman equation (12), with four state variables $(\hat{\delta}_t, \hat{\sigma}_t^2, \Delta p_t, \text{and } \eta_t)$ and four-dimensional integration (with respect to $\Delta t, \eta_{t+\Delta t}, \Delta p_{t+\Delta t}$ and $\hat{\delta}_{t+\Delta t}$), to the threshold equation (25), with two state variables $(\hat{\delta}_t \text{ and } \hat{\sigma}_t^2)$ and one-dimensional integration (with respect to $\hat{\delta}_{t+\Delta t}$), by exploiting certain features of our model and data. The Additive Separability (AS) and Conditional Independence (CI) assumptions invoked by Rust (1987) hold in our model. AS holds because η_t and Δp_t enter in an additively linear fashion in the instant utility function. CI holds because the distribution of doctor's beliefs in the next period can be completely characterized by the doctor's current beliefs and his current decision, and because η_t and Δp_t are i.i.d. (in the absence of CI, the distribution of doctor's beliefs in the next period would generally be a function of all the

²⁷We exploit the fact that price difference exhibits little variation to avoid one dimension of integration for the threshold calculations. In the estimation, however, we use the actual price difference, which varies over time, in order to identify ϕ .

current period's state variables, and the current period's decision). CI allows us to write the continuation payoffs C_o and C_n as functions of only two state variables ($\hat{\delta}_t$ and $\hat{\sigma}_t^2$) instead of the original four. AS, in conjunction with CI and the assumption that the doctor only faces two choices (old or new drug), gives us a convenient shape of the value function with respect to η_t (see (19) and Figure 2), a shape whose usefulness will be explained below.

To reduce the dimension of integration from four to one, we rely on: a) the assumption that the patient arrival process is invariant to other variables, which allows us to integrate Δt out; b) the very low variation of Δp over time in our data, so that we can evaluate all functions of Δp at $\overline{\Delta p}$ and avoid integration with respect to Δp ; c) the shape of the value function with respect to η_t . Rust (1987) assumes that i.i.d. shocks follow a Type I extreme value distribution. If we assumed this distribution for η , we would obtain closed-form solutions for the integrals with respect to η in the continuation payoffs (see (15) and (16)). However, we are able to avoid numerical integration with respect to η without assuming a Type I distribution for η . This is because the convenient shape of the value function with respect to η yields analytical solutions for the integrals with respect to η in the continuation payoffs. Although we have assumed that η has a standard normal distribution, it would be possible to find an analytical solution for these integrals for *any* distribution of η , because our ability to find this solution depends entirely on the *shape* of the value function with respect to η .

Since the doctor's beliefs do not change when he prescribes the old drug, we can collapse the equations for C_o and C_n to one equation in one unknown -the threshold- involving just a one-dimensional integral over $\hat{\delta}_{t+\Delta t}$ (see (25)). In particular, the threshold is not a function of the continuation payoffs. In order to solve this (nonlinear) equation for a given $(\hat{\delta}_t, \hat{\sigma}_t^2)$ combination, we have developed a fast algorithm that operates in a recursive fashion. The algorithm rests on the Bayesian updating feature by which the variance of a physician's beliefs falls each time he prescribes the new drug (see (13)). This means that after a sufficiently large number of prescriptions of the new drug the variance of beliefs is almost equal to zero, and the mean of beliefs no longer changes. For a sufficiently small $\hat{\sigma}_t^2$, $C_o(\hat{\delta}_t, \hat{\sigma}_t^2) = C_n(\hat{\delta}_t, \hat{\sigma}_t^2)$, which means that $\omega(\hat{\delta}_t, \hat{\sigma}_t^2) = -\hat{\delta}_t$. We can then apply backward induction to find the threshold for increasingly higher values of $\hat{\sigma}_t^2$. Appendix A2 contains

further details on the algorithm.

To summarize, we have reduced the physician’s dynamic discrete problem to the calculation of threshold values that depend only on two state variables and are not functions of continuation payoffs. The computational burden of our approach is the same as that which would result from assuming a Type I extreme value distribution for η (as in Rust 1987). However, the computational burden of our approach would be larger if the physician faced more than two choices, because in this case it would be harder to characterize the doctor’s decision rule through the use of threshold functions.

5 Estimation

The model includes the following parameter vector:

$$\theta = \left\{ \{a_i, b_i\}_{i=1}^I, \varkappa, \delta, \hat{\delta}_0, \hat{\sigma}_0^2, \sigma_\xi^2, \phi, \overline{\Delta p} \right\} \quad (26)$$

where a_i and b_i are the parameters of doctor i ’s patients’ arrival process, I is the total number of doctors in the sample, \varkappa is the discount factor, δ is the true quality of the new drug, $\hat{\delta}_0$ and $\hat{\sigma}_0^2$ characterize the doctors’ system of initial beliefs (prior beliefs), σ_ξ^2 is the variance of the prescription signal, ϕ measures price sensitivity and $\overline{\Delta p}$ is the expected value of price difference. Below we describe our estimation procedure and the features of the data that identify the parameters of the model.

To describe the estimation, some additional notation is in order. Let M denote the total number of months in the sample. Let n_m^i be the number of patients who visit doctor i in month m , and r_m^i the total number of new drug prescriptions by doctor i in month m . The estimation procedure involves the following steps:

Step 1 (Calibration of \varkappa): We set the value of \varkappa equal to 0.00025, which corresponds to an annual discount factor of 0.997.²⁸

²⁸When we attempted to estimate \varkappa , our point estimate was on the boundary at zero, which corresponds to an annual discount factor of one. Problems with the estimation of the discount factor are pervasive in the literature, as indicated by Erdem and Keane (1996), whose estimate of the annual discount factor is above one. Thus, we calibrate \varkappa to 0.00025, which is very small yet still numerically manageable. Erdem and Keane (1996) and Crawford and Shum (2005) also calibrate the discount factor.

Step 2 (Estimation of $\overline{\Delta p}$): Our estimate of the expected value of price difference is the sample mean of price difference, equal to 1,106 liras per day.

Step 3 (Estimation of the arrival process parameters for each doctor): The patient arrival process for doctor i is characterized by two parameters, a_i and b_i . We estimate these parameters using Maximum Likelihood. For each doctor, we match the empirical distribution of the number of patients per month n_m^i to the distribution implied by (11). The variation in the number of patients across months for a given doctor identifies the arrival process parameters.

Step 4 (Estimation of the remaining parameters): In principle, to estimate the remaining parameters $\tilde{\theta} = \{\delta, \hat{\delta}_0, \hat{\sigma}_0^2, \sigma_\xi^2, \phi\}$, one would apply Maximum Likelihood. However, the lack of some critical data leads us to apply Simulated Maximum Likelihood instead, as we explain in the paragraphs that follow.

Consider the ideal situation, in which we observe for each doctor the sequence of his prescription decisions $\{d_{jm}^i\}$, $m = 1, \dots, M$, $j = 1, \dots, n_m^i$, and his beliefs at the time of writing each prescription $\{\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2\}$, $m = 1, \dots, M$, $j = 1, \dots, n_m^i$. The likelihood function for the whole sample is then equal to

$$L(\tilde{\theta}|\{d_{jm}^i\}, \{\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2\}, \{\Delta p_{jm}\}) = \prod_{i=1}^I \prod_{m=1}^M \prod_{j=1}^{n_m^i} \left\{ \left[\Pr(\eta_{jm}^i - \phi \Delta p_{jm} > \omega(\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2)) \right]^{d_{jm}^i} \right. \\ \left. \times \left[\Pr(\eta_{jm}^i - \phi \Delta p_{jm} < \omega(\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2)) \right]^{1-d_{jm}^i} \right\} \quad (27)$$

where Δp_{jm} is the price difference for patient j at month m . Doctor i 's contribution to the likelihood function (27) is

$$L^i(\tilde{\theta}|\{d_{jm}^i\}, \{\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2\}, \{\Delta p_{jm}\}) = \prod_{m=1}^M \prod_{j=1}^{n_m^i} \left\{ \left[\Pr(\eta_{jm}^i - \phi \Delta p_{jm} > \omega(\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2)) \right]^{d_{jm}^i} \right. \\ \left. \times \left[\Pr(\eta_{jm}^i - \phi \Delta p_{jm} < \omega(\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2)) \right]^{1-d_{jm}^i} \right\}. \quad (28)$$

In this ideal situation, to estimate the model one finds the value of $\tilde{\theta}$ that maximizes (27) given the values of the parameters estimated in the previous steps.

However, our data are not ideal for two reasons. First, we do not observe the sequence of prescription decisions for each doctor; we only observe how many patients he receives per

month, n_m^i , and to how many he prescribes the new drug, r_m^i . For instance, if the doctor sees two patients in a given month and prescribes the new drug to only one of them, the patient getting the new drug could be either the first or the second. In other words, for doctor i and month m , there are $\binom{n_m^i}{r_m^i}$ possible prescription sequences that could have generated the observed r_m^i new drug's prescriptions out of a total n_m^i prescriptions. Second, we do not observe the doctor's beliefs at the time he writes the prescriptions, or the prescription signals that he receives when prescribing the new drug.

Denote as \mathbb{C}_r^n the set of all sequences of prescription decisions where out of n prescriptions the new drug was prescribed r times. There are $\binom{n}{r}$ elements in this set. Then, conditional on $\{r_m^i, n_m^i\}$ and $\{\Delta p_{jm}\}$ the doctor's likelihood is

$$\begin{aligned} & L^i(\tilde{\theta}|\{r_m^i, n_m^i\}, \{\Delta p_{jm}\}) \\ = & \sum_{\{d_{jm}^i\} \in \times_{m=1}^M \mathbb{C}_{r_m^i}^{n_m^i}} \int_{\{\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2\}} L^i(\tilde{\theta}|\{d_{jm}^i\}, \{\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2\}, \{\Delta p_{jm}\}) dZ(\{\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2\}|\{d_{jm}^i\}) \end{aligned} \quad (29)$$

where $\{d_{jm}^i\}$ is a sequence of prescriptions (over the whole period) consistent with $\{r_m^i, n_m^i\}$, $\times_{m=1}^M \mathbb{C}_{r_m^i}^{n_m^i}$ is the set of all such sequences and $Z(\{\hat{\delta}_{jm}^i, [\hat{\sigma}_{jm}^i]^2\}|\{d_{jm}^i\})$ is the distribution of doctor i 's beliefs conditional on the sequence of prescriptions $\{d_{jm}^i\}$.

To calculate (29) directly, one needs to integrate over doctor i 's beliefs. The dimensionality of this integration equals the total number of prescriptions of the new drug, which is approximately equal to 180 for the average doctor. In addition, one needs to sum over all the sequences of prescriptions which are consistent with $\{r_m^i, n_m^i\}$. For a doctor who sees 50 patients each month and prescribes the new drug to 5 of them, the number of these sequences is of order 10^{196} . In other words, direct calculation of (29) is computationally infeasible.

Nonetheless, we can approximate (29) through Monte-Carlo integration. We do this by randomly picking $S = 1,000$ prescription sequences for each doctor that are consistent with $\{r_m^i, n_m^i\}$.²⁹ For each sequence of prescription choices $\{d_{jm}^{is}\}$ and a given value of $\tilde{\theta} = \{\delta, \hat{\delta}_0, \hat{\sigma}_0^2, \sigma_\xi^2, \phi\}$, we calculate the implied beliefs $\{\hat{\delta}_{jm}^{is}, [\hat{\sigma}_{jm}^{is}]^2\}$ according to the distribution $Z(\{\hat{\delta}_{jm}^{is}, [\hat{\sigma}_{jm}^{is}]^2\}|\{d_{jm}^{is}\})$. To do so, we generate a sequence of prescription signals $\{\varepsilon_{jm}^{is}\}$ that

²⁹We examined the sensitivity of the likelihood function to changes in the number of simulation sequences. We found that 1,000 simulation sequences was enough to achieve the desired accuracy in our parameter estimates, which are accurate up to the third decimal place.

is consistent with $\{d_{jm}^{is}\}$. Hence, we only draw signals for prescriptions of the new drug (recall that the physician does not observe a signal if he prescribes the old drug). We draw the signals from their true distribution $N(\delta, \sigma_\xi^2)$. Given this sequence of signals we calculate the implied beliefs $\{\hat{\delta}_{jm}^{is}, [\hat{\sigma}_{jm}^{is}]^2\}$. Then the likelihood of interest (29) is approximately equal to the following average of simulated likelihood values:

$$\begin{aligned} & \tilde{L}^i(\tilde{\theta}|\{r_m^i, n_m^i\}, \{\Delta p_{jm}\}) \\ &= \frac{1}{S} \sum_{s=1}^S \left\{ \Pi_{m=1}^M \binom{n_m^i}{r_m^i} \right\} L^i(\tilde{\theta}|\{d_{jm}^{is}\}, \{\hat{\delta}_{jm}^{is}, [\hat{\sigma}_{jm}^{is}]^2\}, \{\Delta p_{jm}\}). \end{aligned} \quad (30)$$

Note that this formulation takes into account the fact that there are $\Pi_{m=1}^M \binom{n_m^i}{r_m^i}$ possible sequences that are consistent with $\{r_m^i, n_m^i\}$ (i.e., the cardinality of the set $\times_{m=1}^M \mathbb{C}_{r_m^i}^{n_m^i}$ is equal to $\Pi_{m=1}^M \binom{n_m^i}{r_m^i}$).

Finally, the joint likelihood function for all doctors (27) can be approximated by the product of the doctor's individual simulated likelihood functions given by (30). Substituting (28) into (30) and rearranging terms we obtain the following *simulated likelihood function*.³⁰

$$\begin{aligned} & \tilde{L}(\tilde{\theta}|\{r_m^i, n_m^i\}, \{\Delta p_{jm}\}) \\ &= \prod_{i=1}^I \left\{ \frac{1}{S} \sum_{s=1}^S \Pi_{m=1}^M \left(\binom{n_m^i}{r_m^i} \prod_{j=1}^{n_m^i} \left\{ \left[\Pr(\eta_{jm}^i - \phi \Delta p_m > \omega(\hat{\delta}_{jm}^{is}, [\hat{\sigma}_{jm}^{is}]^2)) \right]^{d_{jm}^{is}} \right. \right. \right. \\ & \quad \left. \left. \left. \times \left[\Pr(\eta_{jm}^i - \phi \Delta p_m < \omega(\hat{\delta}_{jm}^{is}, [\hat{\sigma}_{jm}^{is}]^2)) \right]^{1-d_{jm}^{is}} \right\} \right) \right\}. \end{aligned} \quad (31)$$

Thus, in step 4 of the estimation we maximize $\tilde{L}(\tilde{\theta})$ with respect to five parameters (δ , $\hat{\delta}_0$, $\hat{\sigma}_0^2$, σ_ξ^2 , and ϕ) in order to get an estimate of $\tilde{\theta}$. To evaluate $\tilde{L}(\tilde{\theta})$ we first solve for $\omega(\cdot, \cdot)$, and find threshold values for about 40,000,000 combinations of $\hat{\delta}$, $\hat{\sigma}^2$ and β in approximately 15 seconds in a 2.3 Ghz Intel Centrino processor. Then we run through each of the S simulated sequences calculating the probability of prescribing the new drug to each patient, and updating the corresponding beliefs every time that new drug is prescribed. This takes about 15 seconds. Thus, evaluating $\tilde{L}(\tilde{\theta})$ takes only half a minute, which is very little given the complexity of the problem and the total number of prescription episodes in our

³⁰Since we observe price variation across rather than within months, we assume a constant price difference within each month.

data.³¹

We finish this section with an intuitive discussion of the identification of the parameters involved in step 4. For each doctor i and month m , we observe his monthly share of new drug prescriptions, $\rho_m^i = r_m^i/n_m^i$. The sensitivity of ρ_m^i to fluctuations in the price difference identifies the price coefficient ϕ . The average share of the new drug across doctors for the first months identifies the mean of the prior beliefs, $\hat{\delta}_0$. The ratio $\hat{\sigma}_0^2/\sigma_\xi^2$ is directly related to the speed of learning (see (7)): the lower the ratio, the flatter the market share trajectory and the lower the speed of learning. The trajectory, however, is not strictly monotonic (see Figure 1), neither for the market as a whole nor for individual physicians. The variability of the new drug’s share around the expected path identifies the variance of the prescription signal, σ_ξ^2 . This, in turn, allows us to separately identify $\hat{\sigma}_0^2$ and σ_ξ^2 . Finally, given $\phi\overline{\Delta p}$, market share in the last months of the sample identifies the true quality of the new drug, δ .

6 Estimation results

In this section we describe our estimation results. Column 1 of Table 2 shows the parameter estimates for our model (recall that our unit of measure is the standard deviation of the patient’s match parameter, η). All the parameters are accurately estimated. According to our estimates, the new medication has lower quality than the old one, as the estimate for δ is negative. This is probably due to the fact that, as explained in section 2, during our sample period omeprazole is the preferred treatment only for a subset of ulcer pathologies. The fact that the estimate of the mean of initial beliefs, $\hat{\delta}_0$, is lower than the estimate of the true quality δ indicates that physicians are initially pessimistic about the new drug. However, these beliefs are not completely accurate, as the estimated variance of initial beliefs, $\hat{\sigma}_0^2$, is positive. If it were zero, physicians would not experiment or learn at all. The fact that it is positive means that physicians have incentives to learn by experimentation. The estimated variance of the new drug’s signal, σ_ξ^2 , is larger than the estimated variance of initial beliefs.

³¹We use the outer product of the gradient approximation to obtain the variance-covariance matrix for the estimates of the parameters pertaining to step 4. Since estimating the whole model involves sequential steps, we correct this matrix following Newey (1984) to obtain correct asymptotic standard errors. The resulting matrix is the asymptotic approximation of the variance-covariance matrix (Gourieroux and Monford 1996).

Finally, the estimated price coefficient is positive.


TABLE 2
Parameter Estimates

Parameters	Forward-Looking Model	Myopic Model
	(1)	(2)
δ (true quality)	-0.690 (0.0585)	-0.751 (0.0575)
$\hat{\delta}_0$ (mean of initial beliefs)	-2.826 (0.0710)	-1.966 (0.0708)
$\hat{\sigma}_0^2$ (variance of initial beliefs)	0.971 (0.0355)	0.721 (0.0223)
σ_ξ^2 (variance of signal)	12.108 (0.268)	11.678 (0.251)
ϕ (coefficient on price)	0.195 (0.0579)	0.126 (0.0568)
Log-Likelihood	-18,511.86	-18,520.96

Standard errors are in parentheses. Price is in thousands of liras per day.

To examine the fit of our model, we simulate 1,000 patient sequences for each physician in our sample; for each patient we draw a match parameter η and a signal ε . Figure 3 displays predicted and observed market share, and 95 percent confidence bounds for predicted share.³² As the figure shows, the model fits the data reasonably well, and the observed market share falls within the 95 percent confidence interval for each month. Furthermore, as physicians learn and converge in their beliefs, their prescription behavior converges as well. In other words, learning by experimentation provides an empirically relevant explanation for omeprazole’s usage pattern depicted in Figure 1.

³²When drawing the confidence bounds, we take into account the fact that the parameter estimates are also random variables. The number of patients in the simulated sequences for a given doctor and month is taken from the data for that doctor and month.

FIGURE 3*Goodness of Fit: Observed and Predicted Market Shares for Omeprazole*

Predicted market share and confidence bounds are based on 1,000 simulations for each physician.

The model predicts that in the long run, physicians learn the true quality of the new drug and their beliefs are accurate. According to (18), the threshold value for the patient's match parameter adjusted for price, $\eta - 0.195\Delta p$, is equal to $\omega(-0.69, 0) = 0.69$. If the price difference were the same as the average price difference during our sample period, this would predict a long-run market share of 18% for omeprazole. If the price difference were zero, the predicted market share would be equal to 25%. These numbers square quite well with the fact (see section 2) that omeprazole is the preferred choice for patients whose condition accounts for about 25% of all prescriptions.

Column 2 of Table 2 displays the estimates from the myopic model, in which physicians make prescription choices only to maximize instantaneous utility (i.e., $\kappa \rightarrow \infty$, and the

annual discount factor equals zero for each physician). The fit of the data, as measured by the log-likelihood value, is better for the forward-looking than the myopic model. Using these log-likelihood values we test the null hypothesis that physicians are myopic, and the data reject it.³³ Hence, in the counterfactual simulations that follow we use estimates from the forward-looking model. Overall, we are encouraged by the fit of our parsimonious model.

7 Counterfactuals

We can now use our structural estimates to gauge the cost of uncertainty and myopia, and the value of learning. To do so, we consider a representative physician who sees 50 patients per month for 10 years, and whose patients arrive at a uniform rate within each month. We assume that the treatment lasts 14 days, and that the patient pays a 50 percent copay. We simulate 100,000 prescription (patient) sequences. The results we present are averages over the simulated sequences. We assume that the price difference is constant and equal to the observed sample mean.³⁴

We compare three cases. In the first (“forward-looking”), the physician behaves as in our model – that is, he is uncertain about the new drug’s quality, and is forward-looking. In the second (“myopic”), the physician is also uncertain about the new drug’s quality, but only seeks to maximize current utility. In the third (“full information”), the physician knows the true quality of the new drug, i.e., $\hat{\delta}_0 = \delta$, and $\hat{\sigma}_0^2 = 0$; however, the outcome of the new drug continues to be random.³⁵

We begin by comparing prescription behavior in the three cases. Figure 4 displays the probability that a patient arriving at a particular time during the first 10 years following the

³³The p-value of the test is 0.00002.

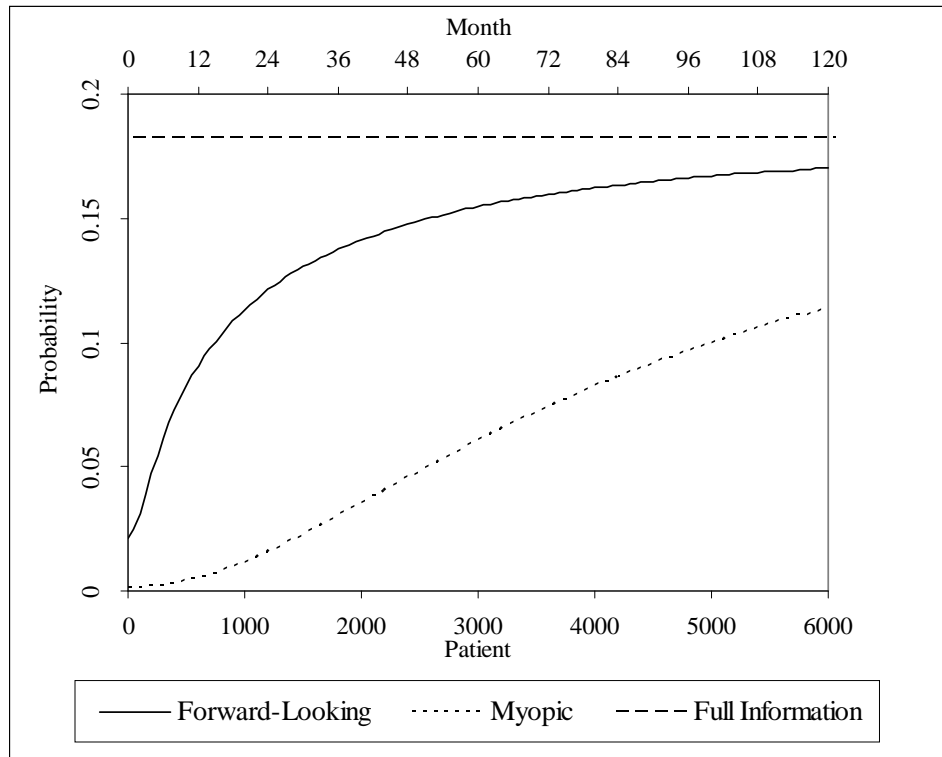
³⁴On treatment length, see <http://www.cdc.gov/ulcer/files/hpfacts.pdf>. The averages depicted in Figures 4, 5 and 6 are kernel-smoothed to reduce simulation noise.

³⁵For the myopic case we use the parameter estimates from the forward looking case, but set the effective discount factor β equal to zero (equivalently, $\kappa \rightarrow +\infty$). Even though we have rejected the null hypothesis that doctors are myopic, it is still of interest to gauge the gains associated with the forward-looking rather than myopic behavior of doctors. The full-information counterfactual is of interest because the Italian policy-maker, who controls prices, might want to choose a price path that induces physicians to make similar choices to those under full information.

new drug's entry receives a prescription for the new drug.³⁶ For a fully-informed physician, this probability is constant at about 0.18, as we saw in the previous section.

FIGURE 4

Probability of Prescribing New Drug



50 patients per month; patients equally spaced within each month. Treatment lasts 14 days; patient's copay equal to 50 percent of drug's price. Constant price difference, equal to the mean in the sample. Average over 100,000 possible sequences of patient and signal realizations.

In contrast, the forward-looking physician is less likely to prescribe omeprazole at any time during the period considered, although he becomes more likely to do so as he learns about it. Learning is faster at the beginning, when he is more uncertain about the true quality. The myopic physician also learns over time, but at a much lower rate. While the forward-looking physician weighs the interest of the current patient with the continuation

³⁶For instance, for patient number 201, who comes at the beginning of the fifth month, this probability is the fraction of simulation sequences in which the patient obtains the new drug, relative to the 100,000 sequences.


value of prescribing the new drug, the myopic physician only considers the first. Thus, he prescribes the new drug less, which means that he has fewer opportunities to learn about it. The difference in behavior between forward-looking and myopic physicians is substantial – forward-looking doctors are about 20 times more likely to prescribe the new drug at the beginning of the period, and 1.5 times more likely at the end. Such a difference provides evidence that prescribing the new drug has indeed a high learning value for forward-looking doctors.

The question, however, is whether these differences in prescription behavior translate into differences in health outcomes. For instance, a physician who believes that omeprazole’s quality is lower than it actually is will not prescribe it to patients who should receive it. Thus, in Figure 5 we seek to answer the following question: for a patient who comes at a particular time, what is her expected treatment outcome? Expected outcomes differ over time because beliefs, which condition prescription choices, differ over time.³⁷

³⁷We express the expected health outcome in thousands of liras. For a patient with match parameter equal to η , we calculate the expected health outcome as $(14/2)(\delta + \eta)/\phi$ when the new drug is prescribed, and zero otherwise. We multiply by 14 (days) to adjust for treatment length, and divide by 2 because patients only pay half of a drug’s actual price.

FIGURE 5

Monetary Value of Expected Health Outcomes




50 patients per month; patients equally spaced within each month. Health outcomes expressed in thousands of liras of January 1990 (1,236 liras = 1 dollar). Treatment lasts 14 days; patient's copay equal to 50% of drug's price. Constant price difference, equal to the mean in the sample. Average over 100,000 possible sequences of patient and signal realizations.

The figure shows that if a patient visits a fully-informed physician, she can expect her health outcome to be more than 3.5 times better at the beginning of the period than if visiting a forward-looking physician, and 0.11 times better at the end of the tenth year. However, it is worth noting that if the new drug were not available at all, then health outcomes in Figure 5 would coincide with the horizontal axis (recall that instant utility is normalized to zero for the old drug). In other words, although uncertainty is costly, it is even more costly not to have the new drug at all. Figure 6 depicts similar patterns for expected instantaneous utility, which adjusts health outcomes for the disutility of paying for the new

drug.³⁸

FIGURE 6

Monetary Value of Expected Utility per Patient



50 patients per month; patients equally spaced within each month. Health outcomes expressed in thousands of liras of January 1990 (1,236 liras = 1 dollar). Treatment lasts 14 days; patient’s copay equal to 50 percent of drug’s price. Constant price difference, equal to the mean in the sample. Average over 100,000 possible sequences of patient and signal realizations.

Expected health outcomes are worse for myopic than forward-looking doctors. Myopic doctors, who do not learn fast enough due to the lack of experimentation, write the “wrong” prescription at a higher rate - namely, they prescribe the old drug to some of the patients who would receive the new drug from a fully-informed physician. Although expected health

³⁸We express the expected instantaneous utility in thousands of liras. For a patient with match parameter equal to η , we calculate the expected instantaneous utility as $(14/2) (\delta + \eta - 1.106\phi)$ when the new drug is prescribed, and zero otherwise.

outcomes for myopic and forward-looking doctors would converge in the long run, when all doctors learn the new drug's true quality, in the short- and medium-run myopic physicians contribute to relatively poor health outcomes. Even after 10 years, the expected outcomes for myopic doctors are 40 percent lower than for forward-looking doctors.

When writing a prescription, a physician could make two types of errors. First, he could prescribe the new drug to a patient who does not need it (in the sense that her match parameter η is below the full-information threshold). Second, he could fail to prescribe the new drug to a patient who does need it (in the sense that her match parameter is above the full-information threshold value). In our model, as long as doctors are initially pessimistic about the new drug's quality, the first type of error will not happen because the full-information threshold is lower than the forward-looking threshold; only the second type of error may happen, affecting the patients whose match parameter is between the full-information and the forward-looking thresholds. In contrast, only the first type of error would happen if doctors were initially optimistic about the new drug's quality.

Since the forward-looking doctor maximizes expected discounted utility, his prescription for a particular patient may not maximize her expected *instantaneous* utility, conditional on the doctor's beliefs. In contrast, the myopic doctor does maximize the patient's expected instantaneous utility, conditional on his beliefs. Thus, it might seem that the forward-looking doctor exploits the current patient on behalf of future patients whereas the myopic doctor does not. However, this is not the case, because the myopic doctor updates his beliefs more slowly. In the case of initial pessimism his threshold is higher than the forward-looking doctor's and hence he prescribes the new drug less. This means that patients that need the new drug (in the sense explained above) are less likely to get it from him than from the forward-looking doctor. This, in turn, entails lower expected health outcomes and utility for any given patient, as illustrated in Figures 5 and 6.

We have also evaluated a scenario in which the physician never learns about the new medication – namely, a situation in which the physician never updates his beliefs. In this scenario, patients who would benefit from the knowledge gained by the physician through previous patients do not reap such benefits (we have not depicted the corresponding results because they are visually indistinguishable from the horizontal axis, although they

are positive). Results from this scenario tell us that the gains from the existence of the new drug are rendered almost mute when pessimistic doctors do not learn.

To provide a summary measure of the distortions induced by myopia and incomplete information, Table 3 shows the present value of the expected discounted utility from prescription choices over 20 years following the introduction of the new drug.³⁹ The table compares payoffs to physicians who are fully informed, forward-looking, myopic, or who do not learn. These payoffs differ substantially: almost 34,000 dollars for the fully-informed physician; 31,000 for the forward-looking physician; 20,000 for the myopic physician, and 950 for the physician who does not learn (recall that the payoff equals zero when the new drug does not exist). That is, uncertainty causes 9 percent losses for the forward-looking relative to the fully-informed physician. The myopic physician, in turn, suffers 41 percent losses relative to the fully-informed. The myopic physician's losses encompass both the effect of uncertainty and short-sightedness, and most of his losses are accounted for by his short-sightedness. In other words, his failure to experiment in order to reduce the uncertainty has more severe consequences than the uncertainty itself. The worst scenario, however, corresponds to the physician who never learns, whose payoff is many times lower than the myopic physician's.

TABLE 3

Expected Discounted Utility

	Full Information	Forward-Looking	Myopic	No Learning
In thousand of liras	41,576	37,911	24,570	1,178
In US dollars	33,637	30,673	19,879	953

Present value of expected discounted utility calculated for a physician who sees 50 patients per month over twenty years following the introduction of the new drug. Patients equally spaced within each month. Annual discount rate = 0.997. Treatment lasts 14 days; patient copay equal to 50 percent of the drug's price. Constant price difference, equal to the mean in the sample. Average over 100,000 possible sequences of patients and signal realizations. Monetary values expressed in real liras of January 1990 and converted to dollars using the January 1990 exchange rate (1,236 liras = 1 dollar). All values rounded to

³⁹To calculate the presented discounted value we use a discount rate of 0.997, for consistency with our estimation.

the nearest unit.

Although we do not examine the production and supply of the new drug in this paper, it is clear that a delayed adoption of the new drug causes revenue losses for the new drug's manufacturer. For instance, some simple calculations show that the present value of the manufacturer's revenues over the first ten years following the new drug's entry is about 23 percent lower when physicians are forward-looking than when they are fully-informed. The manufacturer, then, may want to reduce initial uncertainty, perhaps through advertising and detailing.

Since uncertainty causes inferior health outcomes, it is plausible that the Italian policy-maker, who controls prices, might want to choose the price path that maximizes social welfare by inducing the optimal amount of experimentation on the part of physicians. In other words, a sufficiently low price may stimulate the use of the new drug, otherwise limited by uncertainty.⁴⁰

Our model allows us to find the price that maximizes the expected social welfare. We consider, again, a representative physician who sees 50 patients per month during his infinite lifetime. Importantly, we assume that the new and old drug have the same, constant marginal cost of production. Under this assumption, maximizing social welfare is equivalent to maximizing net expected social surplus, defined as the difference between the new and

⁴⁰Even if the policy-maker cannot implement these policies at $t = 0$, he can still accomplish a large welfare improvement by implementing them early on. For instance, following the new drug's entry he could have waited for 31 months (the length of our sample period) to gather data and estimate the new drug's true quality. Armed with this estimate, he could have designed his pricing policy. Considering that this policy would have raised welfare from that moment onwards, the informational losses incurred during the first 31 months would likely have been relatively small, particularly given the large value of the discount factor.

old drug's expected welfare:⁴¹

$$\Delta S = \sum_{k=1}^{+\infty} \beta^k E_{t=0} \{[\delta + \eta_k + \xi_k]d_k/\phi\}. \quad (32)$$

The price affects the decision d_k and the threshold value for the new drug. Hence, in (32) price also affects the set of values of η for which the new drug is prescribed. We calculate (32) by simulating 100,000 sequences of patients, and averaging over them. We restrict our search to price paths in which the price is constant and deterministic.⁴² According to our calculations, the socially optimal price difference per day of treatment is -15 real liras, which means that the new drug's price should be 15 liras below the old drug's. In contrast, in the sample the new drug costs, on average, 1,106 liras more than the old drug. For the representative doctor, the social surplus associated with the optimal price is 1,043,953 thousands of real liras (\$844,622), which is about 5 percent higher than the net social surplus from the observed mean price.⁴³ In other words, a slight subsidy for the new drug induces the experimentation rate that is optimal from a social perspective. Our calculations also

⁴¹Social welfare for the representative doctor is the discounted sum of social surpluses from all patients. For a specific patient k (k is the patient's order of arrival) the instantaneous social surplus is the difference between the monetary value of the treatment outcome, and the production cost:

$$S_k = [\delta + \eta_k + \xi_k]d_k/\phi + \hat{u} - c$$

where \hat{u} is the monetary value of patient's utility from the old drug treatment and the other terms are familiar (\hat{u} corresponds to the second term in the last equality in Appendix A1), and c is production cost. Thus, the expected discounted social surplus S is

$$S = \sum_{k=1}^{+\infty} \beta^k E_{t=0} S_k$$

where $E_{t=0}$ stands for expectation taken at time zero. Since a price policy does not affect c or \hat{u} , maximization of S is equivalent to maximization of the net discounted social surplus given by (32).

⁴²Our model describes doctors' behavior under i.i.d. price difference paths. It can be shown that no i.i.d. path yields a higher expected social welfare than the maximal social welfare under constant price difference paths. Thus, searching over constant paths is less restrictive than it might seem.

⁴³The percent difference between the optimal social welfare and the social welfare under the observed price may seem small. This is because the high discount factor gives great weight to the distant future, when the truth is (almost) known. The percent difference would be larger if the time horizon were shorter (i.e., 10 or 20 years).

show that even if it were not feasible to implement the socially optimal price, it would still be possible to raise social welfare up to the level reached by fully-informed physicians under the observed price difference. In particular, lowering the price difference from 1,106 to 1,031 liras per day would accomplish this goal.

8 Conclusions

In this paper we have developed and estimated a model to predict the demand for experience goods. In particular, we have investigated how physicians learn about the quality of a new pharmaceutical. We have studied the adoption of omeprazole, an anti-ulcer molecule that entered the Italian market in June of 1990. Using a panel dataset of prescriptions written by physicians in Rome for almost three years following the new drug’s entry, we estimated a dynamic discrete choice model of prescription choice and learning.

Exploiting theoretical properties of our model, we avoided value function and continuation payoff calculations and reduced the dynamic discrete choice problem to the straightforward calculation of threshold functions. Our parsimonious model fits the data well, and provides evidence that physicians indeed learn by experimentation. In particular, prescribing the new drug has a high learning value for forward-looking physicians.

Our counterfactuals reveal that uncertainty has large effects on the propensity to prescribe the new drug, which in turn leads to large negative effects in health outcomes. While uncertainty has negative effects, myopia’s are even worse, as it substantially reduces a doctor’s propensity to prescribe the new drug and thus limits opportunities for learning. The negative effects of uncertainty, however, might be mitigated by the positive effect of a price discount on the new medication.

We believe that our modeling and computational approach are helpful for predicting demand in markets for experience goods. More broadly, the approach may be applicable to other dynamic discrete choice problems with two choices (for instance, whether to search for a new job or not).

9 Appendix

9.1 A1: Instantaneous utility and normalization

Here we show that the optimal behavior of the doctor depends only on the difference between the instantaneous utility from the new and old drug. We begin by assuming a general instantaneous utility specification:

$$U(d_{t_k}^i) = \begin{cases} \eta_{1t_k}^i + \delta_1 + \xi_{1t_k}^i - \phi p_{1t_k} & \text{if new drug } (d_{t_k}^i = 1) \\ \eta_{0t_k}^i + \delta_0 + \xi_{0t_k}^i - \phi p_{0t_k} & \text{if old drug } (d_{t_k}^i = 0) \end{cases}, \quad (33)$$

where the variables have the same meaning as those in the model. Doctor i maximizes his expected discounted utility

$$\begin{aligned} W^i &= E_{t=0} \left\{ \sum_{k=1}^{+\infty} e^{-\rho t_k} U(d_{t_k}^i) \right\} \\ &= E_{t=0} \left\{ \sum_{k=1}^{+\infty} e^{-\rho t_k} \left\{ [\eta_{1t_k}^i + \delta_1 + \xi_{1t_k}^i - \phi p_{1t_k}] d_{t_k}^i + [\eta_{0t_k}^i + \delta_0 + \xi_{0t_k}^i - \phi p_{0t_k}] (1 - d_{t_k}^i) \right\} \right\} \\ &= E_{t=0} \left\{ \sum_{k=1}^{+\infty} e^{-\rho t_k} [(\eta_{1t_k}^i - \eta_{0t_k}^i) + (\delta_1 - \delta_0) + (\xi_{1t_k}^i - \xi_{0t_k}^i) - \phi(p_{1t_k} - p_{0t_k})] d_{t_k}^i \right\} \\ &\quad + E_{t=0} \left\{ \sum_{k=1}^{+\infty} e^{-\rho t_k} [\eta_{0t_k}^i + \delta_0 + \xi_{0t_k}^i - \phi p_{0t_k}] \right\} \\ &= E_{t=0} \left\{ \sum_{k=1}^{+\infty} e^{-\rho t_k} [\eta_{t_k}^i + \delta + \xi_{t_k}^i - \phi \Delta p_{t_k}] d_{t_k}^i \right\} + E_{t=0} \left\{ \sum_{k=1}^{+\infty} e^{-\rho t_k} [\eta_{0t_k}^i + \delta_0 + \xi_{0t_k}^i - \phi p_{0t_k}] \right\} \end{aligned}$$

where in the last step we rearranged terms and replaced the differences with the following notations:

$$\eta_{t_k}^i \equiv \eta_{1t_k}^i - \eta_{0t_k}^i, \delta \equiv \delta_1 - \delta_0, \xi_{t_k}^i \equiv \xi_{1t_k}^i - \xi_{0t_k}^i \text{ and } \Delta p_{t_k} \equiv p_{1t_k} - p_{0t_k}.$$

The expectation in the second term of the last equality does not depend on the physician's choices. Hence, the physician's optimal behavior is the same as the behavior of a physician for whom the utility of the old drug is normalized to zero.

9.2 A2: Numerical solution for the threshold function

We begin by noting that the variance of the doctor's belief⁴⁴ at time t , $\hat{\sigma}_t^2$, is uniquely determined by the number of the new drug prescriptions n that he has written up to that point. Hence, for simplicity we refer to $\hat{\sigma}_n^2$, given by the following expression:

$$\hat{\sigma}_n^2 = \frac{\sigma_\xi^2}{\hat{\sigma}_0^2 n + \sigma_\xi^2} \hat{\sigma}_0^2. \quad (34)$$

We can then rewrite equation (25) as

$$\omega(\hat{\delta}, \hat{\sigma}_n^2) = -\hat{\delta} + \frac{\beta}{1-\beta} \left\{ \mathcal{F}\left(\omega(\hat{\delta}, \hat{\sigma}_n^2) + \phi \overline{\Delta p}\right) - \int \mathcal{F}\left(\omega(\hat{\delta}', \hat{\sigma}_{n+1}^2) + \phi \overline{\Delta p}\right) dH(\hat{\delta}' | \hat{\delta}, \hat{\sigma}_n^2) \right\} \quad (35)$$

where $\hat{\delta}$ is the mean of doctor's belief for the current patient, $\hat{\sigma}_{n+1}^2$ is given by (34) and $H(\hat{\delta}' | \hat{\delta}, \hat{\sigma}_n^2)$ describes the distribution given by (22) or

$$\hat{\delta}' \sim N\left(\hat{\delta}, \frac{\hat{\sigma}_n^4}{\hat{\sigma}_n^2 + \sigma_\xi^2}\right).$$

Consider an accuracy level, Δ , for the calculation of ω .⁴⁵ From (34) it follows that $\hat{\sigma}_n^2 \rightarrow 0$ as $n \rightarrow +\infty$. For any desired accuracy Δ , there is a corresponding $N(\Delta)$ such that⁴⁶

$$|\omega(\hat{\delta}, \hat{\sigma}_N^2) - \{-\hat{\delta}_N\}| < \Delta \text{ for all } \hat{\delta} \text{ of interest,}$$

where we have written $N(\Delta)$ as N for expositional simplicity. N has the following meaning: after N prescriptions the forward-looking doctor behaves as a myopic doctor because there is nothing to learn when $\hat{\sigma}_N^2 \cong 0$.

Next, we specify a uniform grid for $\hat{\delta}$ with values $\hat{\delta}_j$, $j = 1, \dots, J$ ⁴⁷ and calculate all values $\omega(\hat{\delta}_j, \hat{\sigma}_n^2)$ at each grid point $(\hat{\delta}_j, \hat{\sigma}_n^2)$, $j = 1, \dots, J$, $n = 0, \dots, N$ in the following recursive way: for $n = N$ we set $\omega(\hat{\delta}_j, \hat{\sigma}_N^2) = -\hat{\delta}_j$ for all j . We use these values to calculate $\omega(\hat{\delta}_j, \hat{\sigma}_{N-1}^2)$ for all j from (35). Then, based on the values $\omega(\hat{\delta}_j, \hat{\sigma}_{N-1}^2)$ we calculate $\omega(\hat{\delta}_j, \hat{\sigma}_{N-2}^2)$ for all j and so on. The process finishes when we find all the values $\omega(\hat{\delta}_j, \hat{\sigma}_0^2)$.

⁴⁴We omit the doctor's index, i , for notational simplicity.

⁴⁵In our estimation, we chose $\Delta = 10^{-6}$.

⁴⁶For our chosen accuracy $\Delta = 10^{-6}$ and typical values of model parameters, we have $N(\Delta) \cong 12000$.

⁴⁷By experimenting with the grid size we found that 400 points evenly spaced on the interval $[-5, 5]$ are enough to reach the required accuracy of ω when we use quadratic interpolation between points.

In this recursive procedure, we find each value $\omega\left(\hat{\delta}_j, \hat{\sigma}_n^2\right)$ by solving at a time one equation in one unknown, $\omega\left(\hat{\delta}_j, \hat{\sigma}_n^2\right)$, as is clear from (35). We solve each equation using a Newton method for which the starting point is $\omega\left(\hat{\delta}_j, \hat{\sigma}_{n+1}^2\right)$. To solve each equation, a crucial aspect is evaluating the integral in (35). We must numerically approximate this integral, since we only know the "grid values" of $\mathcal{F}(\cdot)$, namely the values of $\mathcal{F}(\cdot)$ at the grid points $\hat{\delta}_j, j = 1, \dots, J$, equal to $\mathcal{F}\left(\omega\left(\hat{\delta}_j, \hat{\sigma}_n^2\right) + \phi\overline{\Delta p}\right), j = 1, \dots, J$. We construct a piecewise quadratic approximation to $\mathcal{F}(\cdot)$ using the grid values. We favor a quadratic over a linear approximation because this allows us to use fewer grid points. From the theory of approximation by piecewise quadratic functions we know that for any degree of accuracy $\bar{\varepsilon}$ there is a sufficiently large number of grid points such that the approximating function differs from $\mathcal{F}(\cdot)$ at most by $\bar{\varepsilon}$ over the whole range of $\hat{\delta}'$ values. Hence, the integral in (35) is approximated with an accuracy of $\bar{\varepsilon}$. Moreover, the integral of the product of the approximation to $\mathcal{F}(\cdot)$ and the pdf of the known normal distribution $dH(\cdot)$ can be calculated analytically, which allows for a fast evaluation of the integral. The integral does not need to be recalculated during the Newton iterations, which further enhances numerical performance.

References

Ackerberg, Daniel (2003). "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination." *International Economic Review*, 44(3): 1070-1040.

Berndt, Ernst, Robert Pindyck and Pierre Azoulay (2003). "Consumption Externalities and Diffusion in Pharmaceutical Markets: Antiulcer Drugs." *Journal of Industrial Economics*, 51(2): 243-270.

Buhl, K. and Clearfield, H.R. (1990). "Omeprazole: A New Approach to Gastric Acid Suppression." *American Family Physician*, 41(4): 1225-1227.

Ching, Andrew (2008a). "A Dynamic Oligopoly Structural Model for the Prescription Drug Market after Patent Expiration." Working Paper, University of Toronto.

Ching, Andrew (2008b). "Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs after Patent Expiration." Working Paper, University of Toronto.

Chintagunta, Pradeep, Renna Jiang and Ginger Jin (2009). "Information, Learning and Drug Diffusion: the Case of Cox-2 Inhibitors." Working Paper, University of Chicago.

Coscelli, Andrea and Matthew Shum (2004). "An Empirical Model of Learning and Patient Spillovers in New Drug Entry." *Journal of Econometrics*, 122: 213-246.

Crawford, Gregory and Matthew Shum (2005). "Uncertainty and Learning in Pharmaceutical Demand." *Econometrica*, 73(4): 1137-1173.

Currie, Gillian and Sangin Park (2002). "The Effects of Advertising and Consumption Experience on the Demand for Antidepressant Drugs." Manuscript, University of Calgary.

De Giorgi, F., M.F. Savarese, E. Atteo, C.A. Leone and R. Cuomo (2006). “Medical Treatment of Gastro-oesophageal Reflux Disease.” *Acta Otorhngolaryngologica Italica*, 26(5): 276-280.

Eckstein, Zvi, Dan Horsky and Yoel Raban (1988). “An Empirical Model of Optimal Brand Choice.” Working Paper, Tel-Aviv University.

Erdem, Tulin and Michael Keane (1996). “Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets.” *Marketing Science*, 15(1): 1-20.

Freston, J.W. (1994). “Emerging Strategies for Managing Peptic Ulcer Disease.” *Scandinavian Journal of Gastroenterology*, Supplement, 21: 49-54.

Goettler, Ron and Karen Clay (2009). “Tariff Choice with Consumer Learning and Switching Costs: Sorting-Induced Biases and Illusive Surplus.” Manuscript, Chicago Booth School of Business.

Gourieroux, Christian and Alain Monford (1996). “Simulation-Based Econometric Methods.” London: Oxford University Press.

LeCates, Joseph, Claudio Lucarelli, and Sean Nicholson (2008). “How Oncologists Learn About New Chemotherapy Drugs.”. Mimeo, Cornell University.

Manchanda, Puneet, Dick Wittink, Andrew Ching, Paris Cleanthous, Min Ding, Xiaojing Dong, Peter Lefflang, Sanjog Misra, Natalie Mizik, Sridhar Narayanan, Thomas Steenburgh, Jaap Wieringa, Marta Wosinska, and Ying Xie (2005). “Understanding Firm, Physician and Consumer Choice Behavior in the Pharmaceutical Industry.” *Marketing Letters*, 16(3): 293-308.

Meijer, J.L., J.B. Jansen, and C.B. Lamers (1989). “Omeprazole in the Treatment of Zollinger-Ellison Syndrome and Histamine H₂-antagonist Refractory Ulcers.” *Digestion*, 44(1): 31-39.

Miller, Robert (1984). "Job Matching and Occupational Choice." *Journal of Political Economy*, 92(6): 1086-1120.

Mukherjee, Prokriti (2002). "A Structural Model of Learning for New Products." Working Paper, University of Minnesota.

Narayanan, Sridhar, Puneet Manchanda, and Pradeep Chintagunta (2005). "Temporal Differences in the Role of Marketing Communication in New Products Categories." *Journal of Marketing Research*, 42 (August): 278-291.

Narayana, Sridhar, and Puneet Manchanda (2007). "Heterogeneous Learning and the Targeting of Marketing Communication for New Products." *Marketing Science*, forthcoming.

Narayanan, Sridhar, Pradeep Chintagunta, and Eugenio Miravete (2007). "The Role of Self-Selection, Usage Uncertainty and Learning in the Demand for Local Telephone Service." *Quantitative Marketing and Economics*, 5: 1-34.

National Institute for Health Care Management (2002) "Prescription Drug Expenditures in 2001: Another Year of Escalating Costs." NICHM Foundation. Washington, DC.

Newey, Whitney (1984). "A Method of Moments Interpretation of Sequential Estimators." *Economic Letters*, 14: 201-206.

Rust, John (1987). "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher". *Econometrica*, 55(5): 999-1033.

United States Congressional Budget Office (2006). "Research and Development in the Pharmaceutical Industry." Washington, DC.