

Perception Spillovers Between Bundled Goods: A Model of Physician Learning from Chemotherapy Choice

Joseph LeCates
Cornell University

7 November 2011

JOB MARKET PAPER

Abstract

When consumers evaluate a product that is part of a bundle of goods, they can learn about the value of other, similar bundles. Such perception spillovers affect the adoption rate of new products, particularly for entrants that are strong complements to available goods. I extend current models of consumer behavior by developing a Bayesian model of learning which allows across-alternative perception spillovers in the context of chemotherapy for colorectal cancer. Using cancer surveillance data with patient and physician characteristics, I estimate a model of treatment choice to test whether there are informational spillovers across chemotherapy regimens. Despite sharing their major component drug, regimens entering the market in 1991 were not initially considered by physicians to be significantly correlated ($\rho = 0.02$). In contrast, physicians' initial beliefs for those regimens with irinotecan, introduced in 1996, exhibited considerable correlation ($\rho = 0.71$); physicians updated their beliefs for both regimens from experiences with either. Simulated counterfactuals demonstrate that learning and spillovers mitigate the costs of uncertainty by more than 60 percent. These results provide a richer perspective for consumer learning; in particular, physicians use patient feedback to update their beliefs about regimen quality, even across non-prescribed regimens.

Keywords: consumer learning, perception spillovers, new products, pharmaceutical demand

I owe special thanks to Sean Nicholson, Claudio Lucarelli, and Don Kenkel for their guidance and suggestions. I also thank participants at the Department of Policy Analysis and Management's Research In Progress seminar for helpful comments.

E-mail: JRL73@cornell.edu; Web: www.JosephLeCates.com; Google Voice: (607) 699-1006; 404 Uris Hall, Cornell University, Ithaca, NY 14853.

1 Introduction

When new products are launched, or when consumption does not fully reveal the quality of an experience good, consumers face uncertainty about the expected utility value of available products. The natural formation of quality beliefs aids their decision making and provides a framework for individuals to revise their choices in response to new information. Empirical models of learning have shown that information gleaned from experience with a particular good strongly impacts subsequent consumption decisions. A growing literature demonstrates that physicians in particular, acting as agents for their patients, prescribe a specific treatment using previous patients' feedback as a source of information about the value of that therapy.¹ When treatments are combinations of two or more drugs, as is increasingly common in pharmaceutical markets, a physician's experience with any component drugs can also influence her beliefs about the quality of a combination regimen, so that information spills over across treatment options. Such perception spillovers have implications for new treatment adoption, optimal patient care, and product market share whenever physicians evaluate products through joint consumption with other goods and use the quality signals to update their beliefs across all potential bundles.

This research extends empirical models of consumer behavior to allow individuals to learn about quality across all consumption experiences. Specifically, I develop a dynamic model of physician beliefs wherein patient feedback from a particular treatment regimen is used by the physician to update her beliefs about the quality of all potential treatment regimens, not just the prescribed regimen. Starting with Erdem and Keane (1996) and continuing with recent research, learning models have examined consumers' changing beliefs as they incorporate information from noisy quality signals arising from consuming a particular good.² That is, consumers update their beliefs about the

¹The influence of physician experience has been demonstrated by Coscelli and Shum (2004) and Crawford and Shum (2005) with anti-ulcer drugs, Narayanan et al. (2005) for prescription antihistamines, Chintagunta et al. (2009) with Cox-2 inhibitors, Janakiraman et al. (2009) and Dickstein (2011) for antidepressants, Narayanan and Manchanda (2009) in the treatment of erectile dysfunction, Ching and Ishihara (2010) with ACE-inhibitors, and Camacho et al. (2011) for asthma and COPD medications.

²For an example of recent research, see Osborne (2011).

true quality of a good from repeated consumption of the good. Existing models, however, allow beliefs to change only for the particular product being consumed, regardless of whether the revealed information might alter beliefs about other products as well. I generalize a model originating with Coscelli and Shum (2004) to allow feedback from one regimen to inform the perceived value of all regimens, and thereby capture learning across choice options.

Numerous studies in the empirical literature of economics and marketing have demonstrated that consumer perceptions respond to market conditions and experiential information across a variety of settings. Erdem and Keane (1996) originally estimated structural models of brand choice to explain endogenous purchase behavior by allowing both consumption experience and advertising to act as sources of information for uncertain brand attributes. Their model of changing consumer beliefs as Bayesian learning has since been used in diffusion models of new product adoption (Ackerberg, 2003; Coscelli and Shum, 2004; Chintagunta et al., 2009), to study the influence of the extent and timing of advertising (Narayanan et al., 2005; Ching and Ishihara, 2010; Narayanan and Manchanda, 2009), empirical evaluations of branding (Erdem, 1998; Erdem and Sun, 2002; Janakiraman et al., 2008), and models of consumers' behavioral inconsistencies (Lovett, 2008; Camacho et al., 2011).

Absent from this literature is an empirical model allowing for broad perception spillovers across horizontally differentiated products. Information transfers across goods has been investigated as it relates to branding. Erdem and Sun (2002) expanded Erdem (1998) to show that both experience and advertising for a particular brand influences consumer beliefs across products (toothbrushes and toothpaste) of the same brand. The common component, the brand name, is the channel through which information flows. Janakiraman et al. (2009) look for evidence of perception spillovers across competing brands by comparing models of learning in which the quality beliefs of similar and dissimilar products (SSRI and TCA antidepressants, respectively) are used separately as initial beliefs for new entrants (novel SSRI drugs).³ They conclude that perception spillovers occur only when products are sufficiently similar without including a channel for spillovers between similar and dissimilar products simultaneously. In the context of patient-drug matches, patient responsiveness

³Although Janakiraman et al. (2009) focus on the branding implications of perception spillovers, the market for antidepressants may instead be suited to an analysis of horizontally differentiated goods.

to one treatment is informative about the effectiveness of other treatments via correlated quality beliefs, as shown specifically within the learning literature by Crawford and Shum (2005) and Dickstein (2011). A correctly specified choice model must therefore allow for informational spillovers between all products concurrently if consumers indeed update beliefs across multiple alternatives from experience with one.

When goods are consumed jointly, the potential to learn about the value of alternative choices is significant because of the shared components. A variety of consumer experiences are characterized by bundled consumption, such as computer hardware and software. This type of learning is increasingly important to physicians because many new treatments, especially for cancer, HIV/AIDS, and cholesterol, are combinations of drugs. Since 1991, several new molecules have been approved by the Food and Drug Administration (FDA) for the chemotherapeutic treatment of colorectal cancer, most as part of combination therapy. Today nearly all new cases are treated with a multi-drug regimen. The intriguing adoption patterns of early entrants to the colorectal chemotherapy market provide evidence that physicians have particular beliefs about how information spills over to alternatives. The first new drug was adopted slowly despite being a strong complement in combination with the incumbent drug. In contrast, the second entrant, although approved as a stand alone treatment, was adopted primarily as part of a combination regimen with existing chemotherapy drugs. Broadening the scope of information spillovers can explain some of this behavior.

A physician's first choice of chemotherapy regimen for the treatment of colorectal cancer is an ideal setting for a test of perception spillovers. An oncologist's experience is the primary source of information about regimen quality because only a fraction of chemotherapy patients meet the specific conditions for which each of the available drugs was approved, and because pharmaceutical companies can only market a drug for approved uses. When the FDA approved new colorectal chemotherapy drugs in 1990 and 1991, the first time since 1962, physicians faced uncertainty regarding which regimen was best suited for a new patient. As they learned about the quality distributions of the regimens through use, the market leading regimen was supplanted by the entrant, a regimen that would eventually hold more than 80 percent of the market share. Another chemotherapy drug introduction in 1996 repeated the need for learning regarding regimen qualities,

particularly since it was primarily adopted in combination with the other available drugs. The later entrant also faced a greater tradeoff between efficacy and severity of side effects, which additionally required physicians to learn about balancing these drug characteristics.

Finding empirical evidence of perception spillovers requires observing repeated choices by consumers as they encounter new products and learn about their qualities. I hypothesize that consumers have a joint distribution of quality beliefs for the available products and update those beliefs in a Bayesian fashion after experience with a particular bundle. If consumers hold nonzero beliefs about the covariances of quality, feedback from one product experience will influence the expected quality of every alternative choice option and thereby exhibit perception spillovers. By allowing those initial covariances to be determined within an econometric model, the existence of perception spillovers, even across differentiated products, can be tested empirically.

I estimate the model of physician learning using linked cancer surveillance and Medicare claims data on new colorectal cancer diagnoses. Covering an area representing 26 percent of the U.S. population, the data include detailed medical information, patient demographics, and a consistent physician identifier. I construct the complete Medicare patient load of physicians using that identifier and estimate a learning model with perception spillovers for their chemotherapy choices between 1991 and 2001. Identification of the model comes from the changing regimen shares within physician as they treat a sequence of patients over the decade.

The results confirm that physicians hold joint distributional beliefs about regimen qualities. Physicians update their beliefs for each regimen from all chemotherapy experience, a conclusion eluding the previous literature. In accordance with intuition from the adoption rates, regimens entering the market in 1991 were not initially considered by physicians to be significantly correlated ($\rho = 0.02$). Although the regimens shared their major component drug, the new molecule entered after 28 years without innovation and was not proven superior in a large study until years after introduction. In this case, perception spillovers were not a significant source of information for physicians. In contrast, the initial beliefs for those two regimens with irinotecan, introduced in 1996, exhibited considerable correlation ($\rho = 0.71$); physicians updated their beliefs for both regimens from experiences with either. However, neither of the new regimen qualities was strongly correlated

with the incumbent regimens' qualities. These results together provide evidence that physicians believed that any new drug was sufficiently different from the existing drugs during the period that information spilled over only between regimens sharing the new drug. This type of learning behavior slows the adoption rate of complementary drugs such as those entering in the years after the sample.

To explore the magnitude of these results, I simulate three counterfactual scenarios. First, to evaluate the importance of spillovers to the dynamics of learning, I simulate regimen shares when the covariances between the beliefs is set to zero. The counterfactual regimen shares indicate that learning from experience is primarily informative about the prescribed drug, but that spillovers to other alternatives do accelerate adoption mildly. The second and third counterfactuals demonstrate the influence of learning on regimen use. In the second simulation I force physicians to maintain their initial quality beliefs regarding the regimens and I find that the adoption of the dominant regimen is delayed by roughly one and a half years. In contrast, when physician beliefs are set at their ex post levels, estimated as the final posterior in the model, the third counterfactual shows that learning only partially mitigates the delayed adoption resulting from quality uncertainty. Together, these results provide a richer perspective for consumer behavior; namely, physicians form beliefs regarding how information spills over between bundled goods and use patient feedback to update their beliefs across treatment options.

This paper proceeds by first presenting an overview of the chemotherapy market for colorectal cancer. I then introduce the model in two parts: prescription choice and physician learning. A description of the data then allows for a discussion of the identification and estimation of the model. Finally I present the results and discuss simulated counterfactuals before concluding.

2 Chemotherapy for Colorectal Cancer

In 2001, colorectal cancer was third most diagnosed and fatal cancer with an estimated 135,400 new cases and 56,700 deaths in the United States (Greenlee et al., 2001). It has since become the second most diagnosed and fatal cancer among cancers affecting both men and women, with diagnosis

expected in one of twenty people born today (Centers for Disease Control and Prevention, 2011).⁴ The extent of disease is determined by the tumors' increasing invasion of local tissue, involvement of lymph nodes, and metastasis, the spread of cancer cells to other parts of the body. Between 1995 and 2000, the median age at diagnosis with colorectal cancer was 70 years old and patients had a 63.4 percent chance of five year survival overall (Ries et al., 2004). The probability a patient survived for five years ranged from 89.9 percent among those with localized cancers, to 9.6 percent for patients with metastatic cancer.

The three general methods of cancer treatment are surgical removal, radiation therapy, and chemotherapy, the total cost for which has more than doubled in the U.S. since 1987, and exceeded \$20 billion in 2009 (Caplan, 2011). The primary treatment for any colorectal cancer is resection, removing the tumor or sections of the colon with cancer cells, if possible. Chemotherapy is generally considered adjuvant, or supplementary, to resection before metastasis, but is the standard of care after spread has occurred. The National Comprehensive Cancer Network (NCCN) maintains treatment guidelines which update the standard of care with results from new medical research. Most oncology drugs are infused intravenously into patients within their doctors' offices. Hospitals and physicians directly purchase the individual drugs, storing and administering the drugs to their patients as chosen, creating any cocktail treatment regimens themselves. For patients 65 and older, Medicare reimburses the physician for the purchase and administration of intravenous drugs and any oral equivalents if they exist. Particularly as new drugs later entered the market, physician profits varied by regimen, altering treatment incentives and leading to Medicare reimbursement reform in 2003.

Chemotherapy for colorectal cancer is primarily dichotomized into the treatment of metastatic disease and adjuvant therapy for localized cancer. FDA approvals for new chemotherapy drugs, whether the drug is used alone or in combination, indicate for which type of therapy the approval is issued: metastatic or adjuvant. To the extent a drug's use represents a legitimate medical practice, physicians legally may use any FDA-approved (ethical) drug for therapy; however, pharmaceutical firms are only allowed legally to market the drug for its indicated, "on-label" use. Those firms may

⁴The change in rank is due to increased screening for colorectal cancer and improved screening and therapies for prostate and breast cancers.

later apply for other indications to be added, such as the other of metastatic or adjuvant treatment, or in a different combination regimen, by providing supporting clinical evidence of effectiveness for the new indication. All currently available colorectal chemotherapy regimens were initially approved for metastatic therapy, despite only 19 percent of new diagnoses being for metastatic cancer and only 25 percent of chemotherapy patients having such extensive disease. Not surprisingly, the National Comprehensive Cancer Network (NCCN) estimates that 50 to 75 percent of drugs and biologics used to treat all cancers are used “off-label,” for conditions not approved by the Food and Drug Administration (Soares, 2005). These circumstances leave physician experience as the primary source of information for 75 percent of all colorectal chemotherapy patients.

The FDA’s first approved molecule for the treatment of metastatic colorectal cancer was fluorouracil (5FU) in 1962, predating the current era requiring controlled clinical trials in oncology (Ibrahim, 2003). Before 1962 and well into the 1980s, the nonspecific anti-cancer drug methotrexate, approved in 1953, was prescribed in some circumstances, but chemotherapy remained an adolescent practice.⁵ The use of 5FU was still sporadic through the 1980s, when oncologists tried using 5FU for adjuvant therapy. Since the drug had passed patent protection, no clinical trials were performed so that no pharmaceutical firm applied to the FDA to include adjuvant use as an approved indication.

Levamisole, a veterinary deworming drug, was approved in 1990 for the adjuvant treatment of colorectal cancer after two successful studies. On the theory it produced an immune response to the tumor, the National Institutes of Health (NIH) published a consensus statement recommending it in 1990 as the first new therapy in 28 years (NIH Consensus Conference, 1990). Scientific evidence for levamisole was weak, however, and continued to be debated in the medical literature. Levamisole never became an orthodox treatment option and by 1998 it had definitively been shown to produce no benefit while causing unwanted side effects, prompting its withdrawal from the market.

Concurrent with research on levamisole, leucovorin (LV) was tested with 5FU and shown to be effective against metastatic cancer, receiving FDA approval in 1991. LV alone was definitively not chemotherapeutic, but like levamisole, 5FU combined with leucovorin (5FU/LV) failed to consistently produce benefits in small samples in the early 1990s. However, as larger clinical trials

⁵As one medical oncologist described early chemotherapy, “When I graduated from residency in 1972 there were five chemotherapy drugs, and you wouldn’t wish any of them on your worst enemy.”

developed, 5FU/LV demonstrated superior benefits and would comprise nearly 90 percent of colorectal chemotherapy use over the middle of the decade. As an adjuvant treatment, the regimen quickly became the standard of care and continued to be the NCCN's preferred adjuvant treatment into the next decade (National Comprehensive Cancer Network, 2004).

In 1996 Pfizer won accelerated approval for its molecule irinotecan (IRI) under the brand name Camptosar, with full approval coming in 1998. The FDA indicated approval for the drug by itself as second line therapy for metastatic cancer; physicians, however, adopted the drug more quickly in combination with 5FU/LV for patients with advanced disease. The first NCCN practice guidelines for colorectal cancer in 1996 did not endorse IRI (National Comprehensive Cancer Network, 1996), although the combination IRI+5FU/LV was recommended in 2000. The bundled therapy was added as an approved indication in 2000 because it was demonstrated to extend patients' lives by an average 3 months despite having more significant side effects. In addition to effectiveness, the introduction of IRI required physicians learn about the trade-offs between greater toxicity, which promotes efficacy in cancer treatment, and more severe side effects.

By the end of the decade, there were four main regimens for use in the treatment of colorectal cancer: 5FU, 5FU/LV, IRI, and IRI+5FU/LV. A very small fraction of patients received other regimens, typically chemotherapies approved for other types of cancer. In the analysis, this non-standard treatment regimen forms the outside option available to physicians. Table 1 provides an approval time line, approved indications, and some facts for each regimen, and Figure 1 illustrates the share of prescriptions for each regimen within the data.

As noted earlier, a significant portion of learning research has focused on physician choice and pharmaceutical markets because of their impact on individuals and the broader national accounts. Considering generic prescription drugs cost an average 40 percent less than brand name products, Ching (2010) simultaneously models aggregate consumer learning and firm pricing policies to show that slow diffusion of generics can at least be partially explained by learning. Coscelli and Shum (2004) and Crawford and Shum (2005) use panel data for the Italian anti-ulcer market to study within-patient and across-diagnosis learning, respectively, showing that each can substantially reduce consumer uncertainty and influence choice probabilities. Other, ongoing research investigates

the impact of supply-side behavior stemming from product bundling in the colorectal chemotherapy market with results showing a less competitive equilibrium resulting from bundling (Lucarelli et al., 2010). Although firm profits are subject to competitors' behaviors, their effect on information production and dissemination to physicians is uncertain.

3 Model

My model of physician treatment behavior consists of two parts: chemotherapy choice, and physician learning. The choice problem follows the standard conditional logit framework with assumptions reflecting the context of chemotherapy. Physician learning occurs between choices, so that treatment decisions are made using feedback from all previous patients.

In the model, physicians act as consumer agents for their patients and maximize their own utilities by optimally choosing among the available regimens to treat their patients. I assume physicians are myopic, risk-neutral decision-makers, so that in any period, a doctor chooses the regimen with the highest expected utility given her current information and without regard to the dispersion of her beliefs. In contrast to a myopic physician, a forward-looking physician would consider the potential gain to future patients by experimentation with regimens on the current period's patients. The risk of malpractice litigation and high cost of liability insurance are significant deterrents to such behavior. A form of single-patient experimentation, "N-of-1 trials," has been established since 1981 but remains extremely rare (Vohra et al., 2011). Interviews with practicing oncologists confirm that physicians prefer to prescribe regimens well known through practice, research, or communication with pharmaceutical companies, and to practice "defensive medicine" to decrease the probability of a lawsuit (Janakiraman et al., 2008; Kessler and McClellan, 1996). Oncologists, in particular, face greater than average risk of malpractice suits, though payments are sufficiently rare to obscure statistics about the amount (Jena et al., 2011).

Risk-neutrality is an assumption I must make to allow initial quality perceptions to differ from their final values. In a learning model similar to that described below, Coscelli and Shum (2004) demonstrate the inability to separately identify a risk aversion parameter from the initial quality perceptions without additional assumptions. Namely, physicians' initial prior means of the

quality perceptions must coincide with the “true” qualities across all patients. In the context of an expanding choice set, however, quality perceptions should change to reflect the changing relative values of regimens. As a result, no constant “true” quality ranking exists across patients of different times. Instead I assume risk neutral physicians, allowing initial and final quality perceptions to differ and quality rankings to change after new regimen introductions.

Finally, as described above, the set of possible regimens to treat colon cancer has experienced significant growth. Because the set of available regimens grows over time, define an “era” as the time span over which the regimen choice set is constant. The introduction of a new drug will increase the choice set of regimens and thus start a new era. Notationally, time and eras will pass separately.

3.1 Chemotherapy Choice

Let the utility of doctor i treating patient j with drug regimen $g \in [1, \dots, G_p]$ at time t (within era p) be given by the following:

$$\begin{aligned}
 U_{ijg}^* &= \beta_1 \text{Profit}_{gt} + \beta_2 \text{OnLabel}_{jgt} + \beta_3 \text{Recommended}_{jgt} + \beta_4 \text{NumVisits}_g \\
 &\quad + \gamma_{1,g} \text{Metastatic}_j + \gamma_{2,g} \text{Comorbidity}_j + \gamma_{3,g} \text{Age}_j + \delta_g^* + \xi_g^*(t) + \varepsilon_{ijg}^* \\
 &\equiv V_{ijg}^* + \varepsilon_{ijg}^*
 \end{aligned} \tag{1}$$

where

- Profit_{gt} is the average reimbursement for regimen g at time t less the physician’s acquisition costs of its components,⁶ OnLabel_{jgt} indicates if the use of regimen g on patient j falls within the official FDA indication for the regimen at the time of treatment, Recommended_{jgt} indicates whether the use is recommended by the NIH or NCCN, and NumVisits_g is the minimum number of office visits a patient would need to make to receive the regimen.
- Metastatic_j indicates if the patient has metastatic cancer, Comorbidity_j is the Charlson comorbidity weight of patient j derived from his Medicare claims as developed by Klabunde

⁶Chemotherapy is traditionally dosed according to the patient’s body surface area since it is believed to reduce variability in both drug exposure and side effects between patients (Gurney, 1996; Baker et al., 2002). Since patient height and weight is unknown, this variable has been constructed using a representative patient based on the mean body surface area and kilograms from patient-level IntrinsicQ data.

et al. (2000), and Age_j is the patient’s age at treatment.

- δ_g^* parameterizes the core quality perception of regimen g . Although completely unobserved by the econometrician, physicians learn imperfect information about the values both within and across prescriptions as described below.⁷
- $\xi_g^*(t)$ is a flexible, drug-specific time trend summarizing those aspects of period t which affect the perception of regimen g but are outside the available data. Namely, $\xi_g^*(t)$ captures learning from medical journals, participation in clinical trials, medical congresses, and detailing by pharmaceutical companies. The drug-specific trend is constant across patients and doctors.
- ε_{ijg}^* are independent and identically distributed (across patients, doctors, regimens, eras, and time) shocks to patient-drug match observed (perfectly) by the doctor but not the analyst.

If the outside option, the “other” regimen is chosen, let the physician’s utility be

$$U_{ij0}^* = \delta_0^* + \xi_0^*(t) + \varepsilon_{ij0}^* \quad (2)$$

where δ_0^* is known to the physician. $\xi_0^*(t)$ and ε_{ij0}^* are consistent with the above.

Because the utility in Equation 1 has a stochastic component, a physician chooses that regimen g which has the greatest expected utility: $E_t(U_{ijg}^*) > E_t(U_{ijh}^*)$ for all $h \in [0, \dots, G_p]$. Since only differences in utility matter, I normalize the value of the “other” regimen to be zero without loss of generality. The expected utility for any regimen $g \in [1, \dots, G_p]$ may therefore be given as

$$\begin{aligned} E_t(U_{ijg}) &= \beta_1 \text{Profit}_{gt} + \beta_2 \text{OnLabel}_{jgt} + \beta_3 \text{NCCN}_{jgt} + \beta_4 \text{NumVisits}_g \\ &\quad + \gamma_{1,g} \text{Metastatic}_j + \gamma_{2,g} \text{Comorbidity}_j + \gamma_{3,g} \text{Age}_j + E_t(\delta_g) + \xi_g(t) + \varepsilon_{ijg} \quad (3) \\ &\equiv V_{jgt} + \varepsilon_{ijg} \end{aligned}$$

where $\delta_g = \delta_g^* - \delta_0^*$, $\xi_g(t) = \xi_g^*(t) - \xi_0^*(t)$, and $\varepsilon_{ijg} = \varepsilon_{ijg}^* - \varepsilon_{ij0}^*$. Letting g_{ij} denote the chosen regimen and assuming ε_{ijg} is distributed i.i.d. with the type 1 extreme value distribution, the probability

⁷Within the utility model, the core quality perception is constant across patients. While patient response to particular regimens is indeed variable, the included covariates, taken together, form a particular patient-drug match value.

of choosing a specific regimen g is

$$\begin{aligned} \text{Prob}(g_{ij} = g | X, \Theta) &= \int \frac{\exp(V_{ijg})}{1 + \sum_{h=1}^{G_p} \exp(V_{ijh})} dF(\Theta) \\ &\equiv \int P_{ijg}(X, \Theta) dF(\Theta) \end{aligned} \quad (4)$$

for any $g \in [1, \dots, G_p]$. In the standard conditional logit decision model, the choice probability is simply the integrand P_{ijg} above. Since physician perceptions δ_g vary with experience, however, integration recovers the unconditional probability being modeled.

3.2 Physician Learning

The limited information sources available to physicians make experience and learning the primary driver of the treatment and health of future patients. As discussed earlier, no approved FDA therapy existed for nearly 75 percent of colorectal chemotherapy patients over the period. Although the FDA had approved three regimens before 2000 for metastatic cancer, clinical trials provide a limited amount of information to physicians for at least the reason that an individual's patient load will differ from the trial group. Nevertheless, a physician begins the era with initial perceptions about each regimen's quality. In the second and all subsequent eras, these initial beliefs will be the posterior beliefs from the previous era. For clarity, let t denote the beginning of the time period.

I assume that at the beginning of time $t = 1$, doctor i has the following initial beliefs about $\vec{\delta}$, the vector of regimen quality perceptions:⁸

$$\vec{\delta}^i \sim N \left(\vec{\delta}_1^i \equiv \begin{bmatrix} E_1 \delta_1^i \\ \vdots \\ E_1 \delta_{G_p}^i \end{bmatrix}, \Sigma_{\delta,1}^i \equiv \begin{bmatrix} \sigma_{1,1} & \sigma_{1,2} & \dots \\ \sigma_{1,2} & \sigma_{2,2} & \dots \\ \vdots & \vdots & \ddots \end{bmatrix} \right). \quad (5)$$

Her beliefs consist of average regimen qualities and their variances, as well as covariances between the regimens. Those covariances capture the extent to which regimens may be related because of

⁸I present the general case for era $p = 1$; the same structure will hold true for subsequent eras by taking the posteriors of the previous era, $\vec{\delta}_{t-1}^i$ and $\Sigma_{\delta,t-1}^i$, and extending them by the appropriate number of rows (and columns) for the new regimens available.

shared component drugs and, like the mean quality beliefs, will be updated based on the physician's experiences. The Bayesian learning framework formalizes this.

Each time a physician prescribes a regimen, she receives feedback from the patient's experience. That feedback comes in the form of a quality signal. That is, for each patient receiving regimen g , she observes a signal which she characterizes by

$$\mu_{gt} = E_t \delta_g^i + \nu_{gt}, \quad (6)$$

with $E_t \delta_g^i$ coming from her prior and $\nu_{gt} \sim \text{i.i.d. } N(0, \omega^2)$ over j , g and t , resulting from patient idiosyncrasies.

Supposing doctor i begins period t with beliefs on $\vec{\delta}^i$ as her priors and receives signal vector $\vec{\mu}_t$, the joint distribution of beliefs and signals in t is given by

$$\begin{pmatrix} \vec{\delta}^i \\ \vec{\mu}_t \end{pmatrix} \sim N \left(\begin{bmatrix} \vec{\delta}_t^i \\ \vec{\delta}_{\mu t}^i \end{bmatrix}, \begin{bmatrix} \Sigma_{\delta t} & \Sigma_{\delta \mu t} \\ \Sigma'_{\delta \mu t} & \Sigma_{\mu t} \end{bmatrix} \right), \quad (7)$$

where (i) $\vec{\delta}_t^i$ and $\Sigma_{\delta t}$ are the expected mean and variance-covariance matrix of $\vec{\delta}^i$, conditional on the previous signals received before t ; (ii) $\vec{\delta}_{\mu t}^i$ and $\Sigma_{\mu t}$ are the mean and variance-covariance matrix of the physician's signals in period t ; and (iii) $\Sigma_{\delta \mu t}$ is the covariance matrix between $\vec{\delta}^i$ and $\vec{\mu}_t$ which allows for across-regimen learning. Given the form of a signal μ_{gt} and the beliefs $\vec{\delta}_t^i$, the elements of the arrays above derive from the following facts:

- $E(\mu_{gt}) = E_t \delta_g^i$,
- $\text{Var}(\mu_{gt}) = \Sigma_{\delta t(g,g)} + \omega^2$,
- $\text{Cov}(\mu_{gt}, \mu'_{gt}) = \Sigma_{\delta t(g,g)}$, that is, for two signals on the same regimen,
- $\text{Cov}(\mu_{gt}, \mu_{g't}) = \Sigma_{\delta t(g,g')}$, for two signals of different regimens, and
- $\text{Cov}(\delta_g^i, \mu_{g't}) = \text{Cov}(E_t \delta_g^i, E_t \delta_{g'}^i + \nu_{g't}) = \Sigma_{\delta t(g,g')}$

with $\Sigma_{\delta t(x,y)}$ giving the (x, y) coordinate of $\Sigma_{\delta t}$. Thus $\Sigma_{\delta \mu t}$ consists of the elements of $\Sigma_{\delta t}$ ordered according to the number of received signals for each regimen in each period.

I model physician learning from experience as Bayesian belief updating by use of the best linear predictor of $\vec{\delta}_t^i$ given $\vec{\mu}_t$. The distribution of quality beliefs conditioned on the signals received in

period t is the posterior distribution for period t . This posterior then becomes the prior belief distribution for the next period. Following Amemiya (1985), the conditional distributions are calculated as

$$\begin{aligned}\vec{\delta}_{t+1} &\equiv \text{E}(\vec{\delta}|\vec{\mu}_t) = \text{E}_t\vec{\delta} - \Sigma_{\delta\mu t}\Sigma_{\mu t}^{-1}\vec{\delta}_{\mu t} + \Sigma_{\delta\mu t}\Sigma_{\mu t}^{-1}\vec{\mu}_t \\ &= \text{E}_t\vec{\delta} + \Sigma_{\delta\mu t}\Sigma_{\mu t}^{-1}(\vec{\mu}_t - \vec{\delta}_{\mu t}) \\ \Sigma_{\delta,t+1} &\equiv \text{Var}(\vec{\delta}|\vec{\mu}_t) = \Sigma_{\delta t} - \Sigma_{\delta\mu t}\Sigma_{\mu t}^{-1}\Sigma'_{\delta\mu t}.\end{aligned}\tag{8}$$

so that the updating behavior weights the new information by the inverse of the signal's variance, just as in ordinary least squares. Physicians use these updated beliefs in their therapy choice decisions in the following period.

A standard concern with logit choice models is the restrictive substitution patterns resulting from functional form. The independence from irrelevant alternatives (IIA) exhibited by the logit model implies proportional substitution across alternatives; the model is misspecified if the choices exhibit varying degrees of substitutability because of unobserved correlation between them. Because the learning model here explicitly accounts for correlations between the choices, proportional substitution must only hold within physician at any time's information set. When the physician updates her beliefs and their covariances from a new feedback signal, the same proportionality need not hold, reducing the restrictiveness of IIA. The incorporation of learning into the choice model thus additionally provides for more flexible substitution patterns in applied choice analysis.

In summary, the sequence of events in the models is as follows. The physician begins by holding initial beliefs about the value of all regimens available. When the first patient needs treatment, she chooses the regimen with the greatest expected utility, which is a function of her beliefs, the patient's characteristics, and national recommendations. After treatment, she receives a signal of the value of the chosen regimen. Because the quality values across regimens is correlated, she uses the information from the signal to update her beliefs for all regimens. When the next patient arrives, she uses these updated beliefs in her decision. Therefore, from the model, the parameter vector Θ consists of β , γ , and $\xi(t)$ from the choice problem, the initial mean and variance-covariance

matrix of the regimens' quality perceptions ($E_1\delta_1, \dots, E_1\delta_G; \sigma_{1,1}, \dots, \sigma_{G,G}$), and the signal variance ω^2 .

4 Data

The Surveillance Epidemiology and End Results (SEER) Program of the National Cancer Institute (NCI) coordinates the collection of the universe of cancers within several cancer registries covering 26 percent of the US population. This analysis uses the cases reported in the SEER-13 registries from 1991 to 2001: San Francisco, CA; Connecticut; Detroit, MI; Hawaii; Iowa; New Mexico; Utah; Atlanta, GA; San Jose, CA; the Arizona Indian System; Los Angeles, CA; and Rural Georgia.⁹ In addition to information about the tumor, SEER collects demographic and socioeconomic information about the patient and, through a unique patient identifier, follows patients after their diagnosis and surgery. For those in the Medicare population, the SEER-Medicare linked database provides all Medicare claims linked with the patient before and after diagnosis. These claims also contain a consistent physician identifier. Because SEER collects the universe of cancer patients, the physician identifier enables me to construct each physician's complete Medicare patient load.

The analysis sample consists of 411 doctors treating 10,283 cases of colorectal cancer between 1991 and 2001. The sample of physicians was limited to those who treat at least 10 Medicare patients, aged 65 or older, over the period to select a sample for which learning could significantly enhance a physician's treatment patterns. Using drug-specific codes on physician claims, I deduce the regimen first chosen after diagnosis following Warren et al. (2002). The regimen specific characteristics come from Thomson Reuters/Medical Economics Red Book, the FDA's *Drugs@FDA* web site,¹⁰ the NIH and NCCN practice guidelines, and individual drug package inserts.

Physicians in the sample had a median Medicare patient load of 19 colorectal cancer cases, with a mean of 25 cases.¹¹ They used an average of 2.38 regimens in chemotherapy, with the number of

⁹Data requests including ZIP code information must be approved by both the NCI and each individual registry. All registries except Seattle, WA, approved this request.

¹⁰Available: <http://www.accessdata.fda.gov/scripts/cder/drugsatfda/index.cfm>. Approval information for levamisole and leucovorin were not available on the web site but were obtained through a Freedom of Information Act request to the FDA.

¹¹Industry publications from the American Society of Clinical Oncology suggest these numbers are representative of a typical oncologist's Medicare patient load for colorectal cancer.

those choosing the “other” regimen roughly equal between physicians who used 2 and those who used 3 regimens. Table 4 shows how quickly each regimen was adopted by physicians. Without controlling for patient arrival, 75 percent of physicians had prescribed 5FU within the first three years, as compared to the nearly 5 years until the same percentage had used 5FU/LV. Even after 6 years, however, only 35 percent of physicians treated with IRI+5FU/LV, and less than two percent ever prescribed IRI alone. Given the FDA approval patterns, such adoption rates imply calculated behavior by individual physicians.

The median age of the patients in the sample is 73 years old, slightly older than the median of the entire population, and 25 percent have metastatic cancer. Figure 2 shows how patients are distributed over time and by metastatic versus adjuvant therapy. Since the number of patients is evenly distributed, and the proportion with metastatic disease is constant, patient composition alone cannot explain the changes in regimen shares over time. In estimation, age is scaled down so that the Medicare entry age of 65 corresponds to 0, easing later calculations. Using each patient’s claims for the year preceding the month of diagnosis, the Charleson comorbidity weight, a measure of potentially confounding factors to the treatment of cancer, shows no other major comorbidities for 71 percent of the patients in the sample. Controlling separately for age and comorbidity is necessary since the elderly are absolutely less likely to receive chemotherapy than younger patients (Schrag et al., 2001).

Despite the broad coverage of the SEER-Medicare data, the specifics of each program prevents some information from being available. The SEER data identify cases of colorectal cancer based on a physician’s location. If a patient is diagnosed or receives therapy outside of the registry, his data may be incomplete. Additionally, because levamisole was an oral medication, physicians could not submit reimbursement claims for its use in chemotherapy.¹² Thus, the patients receiving the regimen labeled “5FU” may be receiving 5FU alone or in combination with levamisole. The bias from this measurement error should be small: leucovorin was an unproven therapy and used infrequently over the period, declining even more rapidly than 5FU.

¹²The newest data available from the SEER-Medicare linkage does include Medicare Part D claims, made available after the program’s 2006 expansion to prescription drug coverage.

5 Identification

The estimated parameters from the model fall into two groups: the learning parameters, consisting of the initial quality perceptions, their variance-covariance matrix, and the signal variance; and the choice parameters, which apart from the doctor’s expectation of the regimen quality at time t , affects the choice probabilities. The choice parameters are identified from variation in the choice characteristics across patients, regimens, and time as in a standard conditional logit model. The identification of the learning parameters comes from the physicians’ choice patterns and signaling mechanism.

Over the sample period, each of the regimens experiences variation across patients and time. The FDA approvals and NIH and NCCN recommendations are matched by a patient’s stage and quarter of treatment, so that each doctor faces variation across her patient load as to the approval and recommendation status of the regimens.¹³ Further, the NIH recommendations are updated by the NCCN in 1996 and 2000. The number of visits for a regimen, the metastatic indicator, comorbidity weight, and age variables are all regimen specific to satisfy the conditional logit requirements of variation across alternatives.

The initial quality values are identified by both the choice model for the initial period of prescription, as well as by the subsequent pattern of a physician’s prescriptions. In the first instance of treatment, all physicians share the same initial beliefs on 5FU and 5FU/LV, so that a conditional logit estimated on that subset of the data would identify those parameters as regimen-specific constants. In 1996, at the introduction of IRI, the beliefs on 5FU and 5FU/LV would have updated, but again all physicians then share the same beliefs for IRI and IRI+5FU/LV so that their values act as regimen-specific constants. Beyond these initial conditions, however, the learning process itself provides identification. Given a physician’s prescription choices and any covariance matrix for the beliefs, every value of $E_t \vec{\delta}_g$ within the choice model is a function of the initial conditions because of the Bayesian updating process. Every choice within the data, therefore, provides information on the initial beliefs, albeit with different, decreasing signal to noise ratios as time progresses.

The variance-covariance matrix of the initial beliefs and the variance of the feedback signals are

¹³Of the 411 physicians, 5 do not treat any metastatic cancer patients.

determined across physician prescription patterns. Since only the expected value of the regimen quality enters the choice problem, the variances affect only the degree to which those means change given the quality signals from a prescription. It is thus necessary to observe a sufficient number of prescriptions both within and across regimens by the same doctor to learn about how the updating process occurs. Limiting the sample to only those physicians who treat at least ten people ensures sufficient information to capture that process. Figure 3 shows the distribution of the number of regimens used by physicians in the sample. The signal variances are a measure of how responsive physicians are to patient feedback. A small signal variance will discourage physicians from updating their beliefs since signals are centered at their perceived mean. As that variance grows larger, physicians will update their beliefs toward the value of the signal as well as update the regimen’s perceived variance. Repeatedly noisy signals result in persistently high belief variances over the period.

6 Estimation

The data set contains observations on patients, their physicians, and the regimen prescription for treating colorectal cancer. If I observed the feedback signal from each treatment episode I could calculate the revised belief distribution for each physician and substitute their expectations into each of her treatment choice problems. That signal, however, is unobserved.¹⁴ Nevertheless, since the quality perception of a regimen, conditional on the covariates, is homogeneous across patients and normally distributed for a given physician, I can simulate a sequence of physician signals before substituting the quality perception values into the treatment choice model. Utilizing the Bayesian assumption on updating along with the physician’s prescription pattern, I use simulated maximum likelihood estimation to estimate the model parameters as suggested by Train (2009).

Given some initial beliefs on $\vec{\delta} \sim N(\vec{\delta}_1, \Sigma_{\delta_1})$, I can forward construct the signals and updated beliefs as follows. For a given period’s treatment choices, signal vector \vec{m}_t has variance-covariance matrix M made up of parameters from $\Sigma_{\delta t}$ and ω^2 according to the equations on page 14. Just

¹⁴The SEER-Medicare data are rich enough to measure overall survival, progression-free survival, and potentially even severe adverse events. Incorporating these measures into the value of the signal stands at the front of future work.

as random variables are transformed to match known distributions, I use $\vec{\delta}_t$, M , and a vector \vec{z}_m , with the appropriate dimension and randomly drawn from $N(0, I)$, according to

$$\vec{m}_t = \vec{\delta}_t + \check{M}\vec{z}_m \quad (9)$$

where \check{M} is the lower triangular matrix of strictly positive values which results from Cholesky decomposition of M . With $\vec{\delta}_t$ and \vec{m}_t for a period, I use the conditional distribution of the learning model in Equation 8 to create $\vec{\delta}_{t+1}$, the period's posterior and following period's prior, and its variance-covariance matrix $\Sigma_{\delta, t+1}$. Starting with the first period and running through each of the eras, I simulate the beliefs by doctor for all periods. The beliefs are then matched by period to the treatment choice decision for each patient and a conditional logit analysis is performed, resulting in an estimated $\hat{\Theta}$ conditional on the randomly drawn \vec{Z} .

Any one draw of \vec{Z} and simulation of signals is unlikely to approximate the actual signal well. To prevent bias from entering through the signal construction, simulating multiple signal sequences and averaging their outcomes is necessary. Simulating signals R times results in the overall simulated log-likelihood function to be estimated:

$$\mathcal{S}\mathcal{L}(\Theta|X) = \sum_{i=1}^N \sum_{j=1}^{J_i} \log \left(\frac{1}{R} \sum_{r=1}^R [1[g_{ij} = g]P_{ijg}(X, \Theta^r)|\vec{m}^r] \right). \quad (10)$$

That is, for any randomly drawn \vec{Z}^r , I calculate the sequence of signals and update beliefs for all physicians before substituting into the logit choice probabilities for each choice. Finally I average the probabilities across the R draws and then calculate the likelihood as in a standard conditional logit using the averaged probabilities; they are consistent estimates of the true probabilities (Train, 2009). For the results in this paper, $R = 10$ simulations, as in Coscelli and Shum (2004).

7 Results

Tables 5 and 6 contain the learning and choice parameter estimates, respectively, from maximum simulated likelihood estimation. The standard errors are derived from finite differences estimates

of the Hessian at the convergent value of the parameter vector. As a measure of model fit, Figure 4 plots the actual and predicted market shares from the data and model.

Since the initial perceptions estimates are relative to the outside “other” regimen’s value of zero, the initial beliefs are estimated on the order we would expect: 5FU was initially the most popular regimen used and hence has the highest estimate, followed by 5FU/LV. When IRI is introduced, physicians hesitate to adopt it so that its initial quality value is large and negative. IRI+5FU/LV offers more promise, however, and thus has a small but positive value. Each of these parameters is statistically significantly different from zero so that physicians were clearly differentiating these regimens from the “other” option. Although the large negative value on IRI holds that the outside option is conditionally preferred to an FDA approved therapy, this result is only indicative that IRI was not considered a legitimate treatment option for most patients despite its approval. In particular, estimates from the choice model show that IRI was considered only a treatment for metastatic cancer on patients with no comorbidity. Because these characteristics match patients likely to enter a clinical trial, physicians may be attempting to match clinical trial populations for this regimen whereas the absence of clinical trial data for IRI+5FU/LV made it more appealing for more general use.

Turning to the estimates of the initial variance-covariance matrix of the physicians’ beliefs, the variance for the 5FU regimen is the smallest; that physicians were relatively certain of the quality value of the oldest drug on the market is reassuring. As entrants approved on small samples and short-run endpoints, the beliefs around 5FU/LV and IRI are quite uncertain, and only the latter are precisely estimated. Leucovorin’s (LV) initial approval was based on its marginal superiority in a number of small studies when researchers were desperate to find more effective therapies. The benefits of LV continued to be debated in the literature until large, widespread clinical trials finished years after its introduction. Similarly, the large clinical trials for IRI alone concluded after its accelerated approval. In comparison to the two other new regimens, the initial variance on IRI+5FU/LV is much smaller but also statistically significant. Why physicians appear to have been so much more confident about the quality of IRI+5FU/LV at its launch is unclear. Finally, the variance of the signal is less than that of 5FU and IRI, and its standard deviation in comparison

to the size of the initial quality perceptions is not so large as to prevent inference by physicians on the value of regimens.

The estimates of the covariances between the initial beliefs of the regimens show to what extent physicians initially believed the regimens to be correlated, and thus to what extent information spilled over across alternative regimens. Results indicate that physicians do hold joint distributional beliefs across bundles in the form of statistically significant covariances. Although most covariances are small relative the large variances of 5FU/LV and IRI, they are all statistically significantly different from zero. Between 5FU and 5FU/LV, the correlation is 0.0233, indicating an initial perception of only weak similarity. If 5FU was the standard regimen in 1991 when LV was added under weak medical evidence, then physicians could reasonably have considered their effectiveness to be related only mildly, with the new drug providing an uncertain degree of complementarity. 5FU/LV remains only weakly correlated to IRI and IRI+5FU/LV. The exception to the small magnitudes of the correlations, however, is that between the regimens containing IRI. The initial covariance estimates from the model imply an initial belief correlation between IRI and IRI+5FU/LV of 0.726. The strong correlation between IRI and IRI+5FU/LV implies that information from the use of either regimen would considerably inform the beliefs of the other simultaneously. Given the pattern of sizable correlations, physicians seem to believe that any new drug's addition to a regimen considerably alters the quality of the regimen, but regimens that share a new drug are initially perceived to be very similar.

Like the learning parameters, the choice parameter estimates in Table 6 reinforce the anecdotal evidence. Being FDA approved is positively related to the probability a regimen is chosen, although the model indicates that being recommended by the NIH or NCCN is negatively related to a regimen being chosen. This latter result is not surprising given the current knowledge of efficacy for these drugs; that is, the early recommendations were incorrect. The NIH recommended 5FU as adjuvant therapy in its 1990 publication, only mentioning ongoing research into 5FU/LV. When the NCCN updated standard of care recommendations in 1996, it gave the two regimens equal recommendation, not settling on 5FU/LV until 2000. In several large clinical trials publishing later in the decade, 5FU/LV was been demonstrated to be superior in every dimension, so that the

recommendations stood wrong for over most of the sample period. IRI+5FU/LV, although later shown to be superior to IRI alone, was not recommended either. In sum, the “recommendation” variable is signed appropriately. The remaining covariates suitably show that IRI regimens were favored for metastatic patients over those with localized diseases, but that IRI alone was intended for otherwise healthy individuals. Finally, although age has been shown to affect whether or not chemotherapy is prescribed for any colorectal cancer patient, these results show that age has a small impact on which regimen is ultimately chosen.

To assess the importance of learning and perception spillovers to physician treatment practice, I simulate regimen shares over the sample period under three counterfactual scenarios. First, to test the empirical importance of perception spillovers, the initial covariances between regimen beliefs are set equal to zero to prevent information from experience to spill over between regimens in the updating process. Given the small correlations from the estimates, and as Figure 5 shows, the difference between strict product learning and across-alternative learning is small with respect to how regimen shares shift. Although the market for colorectal chemotherapy used a variety of single- and multiple-drug regimens for therapy between 1991 and 2001, physicians’ beliefs about how those drugs were related reduces the magnitude of the effect of spillovers. Nevertheless, since physicians do hold statistically significant beliefs about the initial covariances of the quality of regimens, models of treatment choice, and consumer choice in general, should capture these covariances to preclude estimation bias from misspecification.

Although the market for colorectal chemotherapy displayed limited use of perception spillovers between 1991 and 2001, the second and third counterfactual simulations confirm that experience was a significant source of information over the period. As presented in Figure 6, if physicians do not learn from experience, but maintain their initial beliefs over the entire period, a substantial fraction of patients continue to be prescribed 5FU despite its inferiority. Because physicians are not learning about 5FU/LV at the same time, IRI is briefly adopted more quickly than if learning occurs, but later falls below the actual rate because physicians are so hesitant to move away from 5FU. Although learning from experience greatly speeds up the adoption rate, it is not a substitute for greater knowledge. In Figure 7, regimen shares are simulated under the assumption

that physicians start the treatment period with the ex post beliefs estimated in the learning model. 5FU and 5FU/LV reach their actual market shares roughly a year before they are predicted in the learning model. By the end of the data, IRI is also predicted to have a 16 percent greater market share. These counterfactuals require the assumption that the extensive margin of chemotherapy would have remained the same under the alternative belief regimes, though the data suggest the extensive trend does not change over the period.

Using overall survival data from clinical trials and the FDA's original drug approvals, I can estimate the survival value of learning within each of these counterfactuals based on market share. Given the prescribed regimens and their average survival duration in clinical trials, the sample patients survived an estimated 6925.951 years, or 8.082 months per patient. Disallowing perception spillovers as in the first counterfactual shows that patients would live 0.046 months less if spillovers were not a source of information to physicians. The bigger survival difference is a 237.004 year increase learning gives over the second counterfactual of maintained priors, amounting to a 0.277 month, or 3.43 percent, increase in survival per patient from learning. If instead, physicians had their ex post beliefs at the beginning, the sample patients would have survived an additional 145.082 years, or 0.169 months, under the third counterfactual. In sum, changes to a physician's practice pattern from learning reduces the cost of uncertainty by 62.1 percent. Learning from experience therefore enables physicians to overcome their initial uncertainty about a regimen due to imperfect information, a relatively small portion of which comes from perception spillovers. Nevertheless, shares would still be greater for the superior regimens if physicians had even more information available.

8 Conclusions

In this paper, I identify a market in which knowledge gleaned from experience with a particular bundle of goods is informative about the value of other alternative bundles. I estimate a model of learning allowing for these perception spillovers using physicians and their chemotherapy choices during a period of expanding treatment alternatives. As a particular example of consumer learning, this work contributes to the understanding of physician decision-making and consumer information

processing. If information flow across alternatives indeed influences subsequent choices, empirical models which ignore this behavior will be misspecified and suffer from unobserved correlation between choices. Since a restricted spillover model can be nested within my proposed framework, however, this model may serve as a starting point for future research.

Employing a data set of the complete Medicare patient load for a set of physicians treating colorectal cancer between 1991 and 2001, I estimate the model of Bayesian learning and find physicians do learn across consumption alternatives, evidenced by statistically significant non-zero covariances within their belief distributions. The results show that physicians are initially skeptical about the spillovers between incumbent goods and market entrants even when those goods are used in combination, as seen with the low correlation between 5FU and 5FU/LV, and 5FU/LV and IRI+5FU/LV. Regimens which shared the new drug IRI were nevertheless considered highly correlated by physicians, indicating that quality beliefs for both regimens were updated when new information for either regimen was received.

Together these results present a picture of how physicians reacted to pharmaceutical innovation: they considered the addition of a new drug to an established regimen as a completely distinct therapy, although regimens sharing the a new drug were informative for each other. After the sample period, the market for colorectal chemotherapy continued to grow with the approval of two additional drugs over the next three years and two more two years later, bringing the total number of treatment regimens to 14 by the end of 2006. In contrast to the regimens introduced in the current sample, one of subsequent entrants, oxaliplatin, is used exclusively in combination with previously available regimens. Estimation is currently in progress to include this other type of entrant in the model to further understand the adoption patterns of physicians.

The importance of learning and perception spillovers are explored in several counterfactual simulations after estimation of the model. From these, results indicate that learning about the individual product prescribed accounts for most of the changes to regimen shares over time. Perception spillovers do accelerate adoption slightly in this market, but the difference between a model which account for learning and one holding prior beliefs constant is far greater. In fact, comparing this difference to the difference between the learning model I employ and simulated market shares

using physician's ex post beliefs over the entirety of the treatment period reveals that learning accelerates regimen share use by more than 60 percent.

Despite the rich dataset used in this analysis, however, not all potential sources of variation are captured within the model. The sample is limited to chemotherapy after a patient's first cancer diagnosis to reduce informational heterogeneity between cases. Physicians may also learn from recurrent cancer in patients, but how the information available to physicians in such cases would influence beliefs for initial patient treatment is unclear. An additional limitation is the absence of information on younger patients with private insurance. Previous research has shown that older patients, that is, those on Medicare, are treated differently on the extensive margin of chemotherapy. If they are treated differently on the intensive margin as well, the estimates presented here for physician learning on older Americans are not biased by the exclusion of younger patients as sources of information for physicians. Finally, a potential source of unobserved information is likely to come from detailing. If physicians experience differential detailing so that heterogeneity is missed in the time trend then pharmaceutical advertising remains a concern. Future research could address these issues and further explain how consumers learn and process information.

References

- Ackerberg, Daniel (2003). “Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination.” *International Economic Review*, 44(3), 1007–1040.
- Amemiya, Takeshi (1985). *Advanced Econometrics*. Harvard University Press, Cambridge, MA.
- Baker, Sharyn, Jaap Verweij, Eric Rowinsky, Ross Donehower, Jan Schellens, Louise Grochow, and Alex Sparreboom (2002). “Role of Body Surface Area in Dosing of Investigational Anticancer Agents in Adults, 1991–2001.” *Journal of the National Cancer Institute*, 94(24), 1883–1888.
- Camacho, Nuno, Bas Donkers, and Stefan Stremersch (2011). “Predictably non-Bayesian: Quantifying Salience Effects in Physician Learning About Drug Quality.” *Marketing Science*, 30(2), 305–320.
- Caplan, Arthur (2011). “Will Evidence Ever Be Sufficient to Resolve the Challenge of Cost Containment?” *Journal of Clinical Oncology*, 29(15), 1946–1948.
- Centers for Disease Control and Prevention (2011). “Vital Signs: Colorectal Cancer Screening, Incidence, and Mortality – United States, 2002–2010.” *Morbidity and Mortality Weekly Report*, 60, 884–889.
- Ching, Andrew (2010). “Consumer Learning and Heterogeneity: Dynamics of Demand for Prescription Drugs After Patent Expiration.” *International Journal of Industrial Organization*, 28(6), 619–638.
- Ching, Andrew and Masakazu Ishihara (2010). “The Effects of Detailing on Prescribing Decisions Under Quality Uncertainty.” *Quantitative Marketing and Economics*, 8(2), 123–165.
- Chintagunta, Pradeep, Renna Jiang, and Ginger Jin (2009). “Information, Learning, and Drug Diffusion: The Case of Cox-2 Inhibitors.” *Quantitative Marketing and Economics*, 7(4), 399–443.
- Coscelli, Andrea and Matthew Shum (2004). “An Empirical Model of Learning and Patient Spillovers in New Drug Entry.” *Journal of Econometrics*, 122(2), 213 – 246.
- Crawford, Gregory and Matthew Shum (2005). “Uncertainty and Learning in Pharmaceutical Demand.” *Econometrica*, 73(4), 1137–1173.
- Dickstein, Michael (2011). “Efficient Provision of Experience Goods: Evidence from Antidepressant Choice.” Job Market Paper, Harvard Business School, Boston, MA 02163.
- Erdem, Tülin (1998). “An Empirical Analysis of Umbrella Branding.” *Journal of Marketing Research*, 35(3), 339–351.

- Erdem, Tülin and Michael Keane (1996). “Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets.” *Marketing Science*, 15(1), 1–20.
- Erdem, Tülin and Baohong Sun (2002). “An Empirical Investigation of the Spillover Effects of Advertising and Sales Promotions in Umbrella Branding.” *Journal of Marketing Research*, 39(4), 408–420.
- Greenlee, Robert, Mary Beth Hill-Harmon, Taylor Murray, and Michael Thun (2001). “Cancer Statistics, 2001.” *CA: A Cancer Journal for Clinicians*, 51(1), 15–36.
- Gurney, Howard (1996). “Dose Calculation of Anticancer Drugs: A Review of the Current Practice and Introduction of an Alternative.” *Journal of Clinical Oncology*, 14(9), 2590–2611.
- Ibrahim, Amna (2003). “Past FDA Approvals in Colorectal Cancer.” In “Clinical Trial Endpoints in Colorectal Cancer,” Center for Drug Evaluation and Research, Food and Drug Administration, Washington, DC.
- Janakiraman, Ramkumar, Shantanu Dutta, Catarina Sismeiro, and Philip Stern (2008). “Physicians’ Persistence and Its Implications for Their Response to Promotion of Prescription Drugs.” *Management Science*, 54(6), 1080–1093.
- Janakiraman, Ramkumar, Catarina Sismeiro, and Shantanu Dutta (2009). “Perception Spillovers Across Competing Brands: A Disaggregate Model of How and When.” *Journal of Marketing Research*, 46(4), 467–481.
- Jena, Anupam, Seth Seabury, Darius Lakdawalla, and Amitabh Chandra (2011). “Malpractice Risk According to Physician Specialty.” *New England Journal of Medicine*, 365(7), 629–636.
- Kessler, Daniel and Mark McClellan (1996). “Do Doctors Practice Defensive Medicine?” *Quarterly Journal of Economics*, 111(2), 353–390.
- Klabunde, Carrie, Arnold Potosky, Julie Legler, and Joan Warren (2000). “Development of a Comorbidity Index Using Physician Claims Data.” *Journal of Clinical Epidemiology*, 53(12), 1258–1267.
- Lovett, Mitchell (2008). *Unstable Consumer Learning Models: Structural Models and Experimental Investigation*. Ph.D. thesis, Duke University.
- Lucarelli, Claudio, Sean Nicholson, and Minjae Song (2010). “Bundling Among Rivals: A Case of Pharmaceutical Cocktails.” Working Paper 16321, National Bureau of Economic Research.
- Narayanan, Sridhar and Puneet Manchanda (2009). “Heterogeneous Learning and the Targeting of Marketing Communication for New Products.” *Marketing Science*, 28(3), 424–441.

- Narayanan, Sridhar, Puneet Manchanda, and Pradeep Chintagunta (2005). “Temporal Differences in the Role of Marketing Communication in New Product Categories.” *Journal of Marketing Research*, 42(3), 278–290.
- National Comprehensive Cancer Network (1996). “Colorectal Cancer Practice Guidelines.” *Oncology*, 10, 140–175.
- National Comprehensive Cancer Network (2004). “Colon Cancer.” In “Clinical Practice Guidelines in Oncology,” NCCN, Fort Washington, PA.
- NIH Consensus Conference (1990). “Adjuvant Therapy for Patients with Colon and Rectal Cancer.” *Journal of the American Medical Association*, 264(11), 1444–1450.
- Osborne, Matthew (2011). “Consumer Learning, Switching Costs, and Heterogeneity: A Structural Examination.” *Quantitative Marketing and Economics*, 9(1), 25–70.
- Ries, L.A.G., M.P. Eisner, C.L. Kosary, B.F. Hankey, B.A. Miller, L. Clegg, A. Mariotto, E.J. Feuer, and B.K. Edwards (Eds.) (2004). *SEER Cancer Statistics Review, 1975–2001*. National Cancer Institute, Bethesda, MD.
- Schrag, Deborah, Laura Cramer, Peter Bach, and Colin Begg (2001). “Age and Adjuvant Chemotherapy Use After Surgery for Stage III Colon Cancer.” *Journal of the National Cancer Institute*, 93(11), 850–857.
- Soares, Michael (2005). “Off-Label Indications for Oncology Drug Use and Drug Compendia: History and Current Status.” *Journal of Oncology Practice*, 1(3), 102–105.
- Train, Kenneth (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, New York, NY, 2nd edn.
- Vohra, Sunita, Larissa Shamseer, and Margaret Sampson (2011). “Efficacy Research and Unanswered Clinical Questions.” *Journal of the American Medical Association*, 306(7), 709.
- Warren, Joan, Linda Harlan, Angela Fahey, Beth Virnig, Jean Freeman, Carrie Klabunde, Gregory Cooper, and Kevin Knopf (2002). “Utility of the SEER-Medicare Data to Identify Chemotherapy Use.” *Medical Care*, 40(8), IV55–IV61.

Figure 1: Regimen Shares for Colorectal Chemotherapy Prescriptions




Figure 2: Patients by Year and Extent of Disease




Figure 3: Number of Regimens Used by Sample Physicians




Figure 4: Actual and Predicted Regimen Shares




Figure 5: Counterfactual Regimen Shares Without Spillovers




Figure 6: Counterfactual Regimen Shares With Maintained Prior Beliefs




Figure 7: Counterfactual Regimen Shares With Ex Post Knowledge




Table 1: FDA Drug Approvals for Colorectal Chemotherapy

Approval Year	Drug, In Combination	Indication	Type of Approval	Infusions
1962	Fluorouracil, 5FU	Metastatic, First Line	Regular	30
1990	Levamisole	Adjuvant	Regular	30
1991	Leucovorin, 5FU/LV	Metastatic, First Line	Regular	30
1996	Irinotecan, IRI	Metastatic, Recurrent	Accelerated	8
2000	IRI+5FU/LV	Metastatic, First Line	Regular	24

‘Infusions’ are the expected number of scheduled visits at the outset of treatment to receive therapy.
Sources: Ibrahim (2003); National Comprehensive Cancer Network (2004).

Table 2: Colorectal Chemotherapy Regimen Shares

	Patients	5FU	5FU/LV	IRI	IRI+5FU/LV	Other
1991	710	69.44	30.00			0.56
1992	929	57.48	41.98			0.54
1993	1033	55.66	43.56			0.77
1994	988	55.87	43.93			0.20
1995	1052	53.23	46.67			0.10
1996	1030	45.63	52.91	0.00	1.17	0.29
1997	1000	34.70	63.30	0.00	2.00	0.00
1998	983	24.52	74.16	0.10	0.61	0.61
1999	904	15.60	81.31	0.22	2.43	0.44
2000	859	9.43	81.84	0.12	7.33	1.28
2001	795	2.39	85.53	0.50	10.57	1.01

Regimen shares are calculated from estimation sample.

Table 3: Sample Statistics

		Mean	Std. Dev.	Median
Physicians				
	Patient Load	25.017	17.453	19
	Number of Regimens Used	2.380	0.693	2
Patients				
	Metastatic Cancer	0.251		
	Comorbidity Weight	0.422	0.791	0
	Age at Therapy	73.404	5.419	73
Regimens				
	Prescribed Regimen			
	FDA Approved	0.633		
	Recommended	0.276		
	5FU	0.390		
	5FU/LV	0.584		
	IRI	0.001		
	IRI+5FU/LV	0.020		
	Other	0.005		

$N = 411$ Physicians, 10,283 Patients, 11 Years

Table 4: Percentage of Physicians Who Prescribe A Regimen for the First Time in a Particular Year

All Patients					
	5FU	5FU/LV	IRI	IRI+5FU/LV	Other
1991	44.28	26.03			0.97
1992	21.90	27.01			0.73
1993	9.00	12.17			0.97
1994	5.35	6.08			0.49
1995	4.14	6.33			0.24
1996	4.14	3.89	0.00	2.19	0.49
1997	0.97	4.38	0.00	4.14	0.00
1998	0.73	4.62	0.24	1.22	1.22
1999	0.49	3.41	0.49	4.62	0.97
2000	0.00	1.22	0.24	10.95	2.43
2001	0.24	2.68	0.97	11.44	1.70
Never Rx	8.76	2.19	98.05	65.45	89.78
Metastatic					
	5FU	5FU/LV	IRI	IRI+5FU/LV	Other
1991	14.29	14.29			0.00
1992	11.58	23.40			0.00
1993	7.14	14.53			0.00
1994	4.68	8.87			0.00
1995	3.45	8.37			0.00
1996	4.19	4.68	0.00	0.74	0.00
1997	4.19	3.94	0.00	2.22	0.00
1998	1.23	5.42	0.25	1.23	0.00
1999	0.74	3.69	0.49	3.94	0.00
2000	0.74	2.71	0.25	7.88	0.00
2001	0.00	1.23	0.99	10.34	0.00
Never Rx	47.78	8.87	98.03	73.65	100.00
Adjuvant					
	5FU	5FU/LV	IRI	IRI+5FU/LV	Other
1991	43.07	17.76			0.97
1992	21.65	17.03			0.73
1993	9.25	12.41			0.97
1994	5.84	9.00			0.49
1995	3.89	6.08			0.24
1996	3.89	8.52	0.00	2.19	0.49
1997	1.22	9.00	0.00	2.19	0.00
1998	1.22	5.84	0.00	0.00	1.22
1999	0.24	4.14	0.00	0.97	0.97
2000	0.00	3.16	0.00	4.62	2.43
2001	0.24	3.65	0.00	3.65	1.70
Never Rx	9.49	3.41	100.00	86.37	89.78

'Never Rx' indicates the group of physicians within the sample who never prescribe the regimen over the sample period.

Table 5: Learning Parameter Estimates from SMLE

	Estimate	Std. Errors
Initial Perceptions		
5FU	7.2332 *	0.2128
5FU/LV	5.8231 *	0.2127
IRI	-54.3400 *	1.6539
IRI+5FU/LV	3.0278 *	0.1706
Initial Covariance Matrix		
Var(5FU)	1.4017 *	0.0068
Cov(—, 5FU/LV)	0.2401 *	0.0023
Cov(—, IRI)	0.5140 *	0.0031
Cov(—, IRI+5FU/LV)	0.3092 *	0.0057
Var(5FU/LV)	75.4800	86.5414
Cov(—, IRI)	-0.1837 *	0.0097
Cov(—, IRI+5FU/LV)	0.5287 *	0.0044
Var(IRI)	76.9792 *	2.5571
Cov(—, IRI+5FU/LV)	9.7357 *	0.0337
Var(IRI+5FU/LV)	2.3374 *	0.0228
Signal Variance	48.7669 *	0.3422
Simulated Log Likelihood	-6484.387	

* $t > 1.96$. $R = 10$ Simulations. $N = 411$ Physicians, 10,283 Patients

Table 6: Choice Parameter Estimates from SMLE

	Estimate	Std. Errors
FDA Approved	0.7652 *	0.0012
Recommended	-0.4724 *	0.0008
Number of Visits	-0.0764 *	0.0071
Metastatic		
5FU	-3.1380 *	0.0032
5FU/LV	-1.3330 *	0.0030
IRI	52.7961 *	1.7007
IRI+5FU/LV	0.3996 *	0.0034
Comorbidity Weight		
5FU	0.2016 *	0.0022
5FU/LV	0.2470 *	0.0022
IRI	-24.6524	159.4879
IRI+5FU/LV	0.2383 *	0.0023
Age		
5FU	0.0516 *	0.0003
5FU/LV	0.0557 *	0.0003
IRI	0.0048 *	0.0007
IRI+5FU/LV	-0.0189 *	0.0003
Time		
5FU	0.8047 *	0.0015
5FU/LV	0.5861 *	0.0015
IRI	0.3670 *	0.0031
IRI+5FUL/LV	0.1196 *	0.0010
Time ²		
5FU	-1.2174 *	0.0014
5FUL/LV	-0.5648 *	0.0014
Simulated Log Likelihood	-6484.387	

* $t > 1.96$. $R = 10$ Simulations. $N = 411$ Physicians, 10,283 Patients