

NBER WORKING PAPER SERIES

AN EMPIRICAL MODEL OF HETEROGENEOUS CONSUMER
SEARCH FOR RETAIL PRESCRIPTION DRUGS

Alan T. Sorensen

Working Paper 8548
<http://www.nber.org/papers/w8548>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2001

I am indebted to Ernst Berndt, Peter Davis, Glenn Ellison, Whitney Newey, and Nancy Rose for numerous discussions and insightful comments during the early stages of this research. Tom Hubbard, Florian Zettelmeyer and participants at workshops at MIT, UCLA, the FTC, and the NBER Winter IO Program also offered useful suggestions. Financial support from the National Science Foundation (Grant SES-0079201) is gratefully acknowledged. Special thanks are due to Joanne Nodecker and Sue Capps at IMS, Inc. for their help in assembling the data. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

© 2001 by Alan T. Sorensen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

An Empirical Model of Heterogeneous Consumer Search for Retail Prescription Drugs
Alan T. Sorensen
NBER Working Paper No. 8548
October 2001
JEL No. D83, L65

ABSTRACT

This paper uses detailed data on retail pharmacy transactions to make inferences about the nature and intensity of consumer search for prescription drugs. Prescription prices exhibit patterns that should, in principle, induce search: in particular, prices vary widely across stores, and stores' price rankings are inconsistent across drugs (so the low-price pharmacy is different for one prescription vs. another). Estimates from a model of pharmacy choice suggest that search intensities are generally low: I estimate that for a typical prescription, the fraction of consumers that price-shops is approximately 5-10 percent. However, variation in this estimated search intensity across drugs is substantial and appears to be consistent with explanations based on rational search; for instance, price-shopping is more prevalent for maintenance medications than for one-time purchases, presumably because the benefits of finding a low price are magnified for prescriptions that are purchased repeatedly. Under some relatively strong assumptions imposed by the empirical model, the data also identify parameters of a search cost distribution, suggesting that the cost of conducting an exhaustive price search is approximately \$15 for the average consumer.

Alan T. Sorensen
Department of Economics
University of California, San Diego
La Jolla, CA 92093-0508
and NBER
asorensen@ucsd.edu

1 Introduction

Consumers commonly purchase goods without first surveying firms' prices. Even when competing firms' prices differ widely, such behavior may be optimal if the costs of conducting a price search are high relative to the potential savings. Since search costs presumably are heterogeneous in any population of consumers, search decisions will also be heterogeneous: for a given product, some consumers will shop around, and others will not. The equilibrium "intensity" of search that results from the aggregation of individual search decisions has important implications for pricing and profitability.

The objective of this paper is to address questions about the nature of consumer search in a specific, important market: the retail market for prescription drugs. The analysis builds on previous work showing that price distributions for retail prescription drugs are consistent with search-based models (Sorensen 2000). Using detailed transactions data from competing retail pharmacies in local markets, this paper addresses three principal questions: (1) How intense is consumer search for prescription drugs (i.e., what fraction of consumers appears to price-shop before purchasing)? (2) Is variation in search intensity across drugs consistent with theories of consumer search? (3) Assuming search decisions are made optimally in some sense, what levels of underlying search costs are consistent with the estimated search intensities?

The transactions data are used to estimate a model of pharmacy choice that explicitly incorporates consumers' search decisions. The results suggest very low levels of consumer search: estimated search probabilities for cash-paying consumers generally fall between 0 and 0.37, with a mean of approximately 0.10. Propensities to search do appear to vary across drugs; in particular, search appears to be most intense for maintenance medications and for prescriptions purchased pre-

dominantly by females. These findings are consistent with principles of consumer search (e.g., that search incentives are magnified for repeatedly purchased goods) and with reports in the pharmacy trade literature (e.g., that oral contraceptives are among the most heavily price-shopped items). When the transactions of insured customers are included in the estimation, the estimated search intensities for such customers is essentially zero, which is precisely what we would expect since insurance coverage eliminates consumers' incentives to price-shop.

The estimated empirical model posits that consumers make search decisions by weighing the potential savings from a price-search against an idiosyncratic search cost, which is heterogeneous in the population of consumers. The estimated search intensities therefore imply an underlying distribution of search costs. In particular, the estimates suggest the cost of conducting an exhaustive price search is roughly \$15 for the average consumer. However, it should be noted that these estimates are based on strong assumptions about the functional form of the search cost distribution..

Understanding the fundamentals of consumer search decisions is clearly important if we ultimately want to evaluate policies aimed at affecting search costs (e.g., price-posting legislation, information campaigns, etc.). Section 5 offers a preliminary discussion of the potential impact of such policies by estimating the partial equilibrium effects of exogenous reductions in average search costs.

2 Background and Data

The widely observed phenomenon of price dispersion seems to disaffirm the proverbial law of one price. However, equilibrium price differences can be explained by a variety of economic models. Consumer search models are particularly attractive, since they typically generate equilibrium price

distributions even if products are homogeneous and firms' costs are identical. Moreover, search models fit well with some aspects of reality: for instance, they typically predict that in equilibrium some consumers will search for low prices, while others will not—a common feature of actual markets.¹

An important element of consumer search models is that search decisions are endogenous; that is, whether a consumer chooses to price-shop or not depends on her expectations about the distribution of prices, which in turn depend on the intensity of search. Equilibrium requires not only that the price distribution be optimal from the firm's perspective, but that consumers' search decisions be based on correct expectations about that price distribution.² An intuitive result that arises from such models is that exogenous increases in consumers' propensities to search (for instance, due to a decrease in search costs) will constrain prices to be lower and less dispersed. In that sense, consumers can be thought of as "policing" the market through search.

Consider a market with perfectly homogeneous firms charging unequal prices. Intuitively, if search is very intense in equilibrium (i.e., a large proportion of consumers shops around before purchasing), firms charging relatively low prices will have disproportionately high sales volumes. Firms charging high prices earn a higher profit-per-sale, but make correspondingly fewer sales. On the other hand, if only a very small fraction of consumers searches, sales volumes will be roughly equal across stores in spite of price differences. In principle, therefore, we can learn about the intensity of search by looking at data on prices and quantities sold at competing firms within a market: for a given product, search intensity will be reflected in the observed sensitivity of market shares to price differences.

¹For studies that discuss this feature of search models, see for instance Stahl (1996) or Rob (1985).

²See Burdett and Judd (1983) for an early exposition of this rational expectations approach to modeling equilibrium price dispersion.

In retail prescription drug markets, search intensities vary across drugs due to differences in therapeutic characteristics of the drugs themselves. The heterogeneity of search intensities should lead to variation in observed market share patterns. For prescriptions that are heavily price-shopped, we expect to see the pharmacy charging the lowest price in the market to have a relatively high market share. Conversely, market shares for prescriptions that are not heavily price-shopped will appear to be independent of price differences. Figure 1 shows two examples from the data. The first panel illustrates the sensitivity of market shares to differences in pharmacies' prices for the anti-ulcer drug Cimetidine in zip code A. Most of the pharmacies' prices were clustered around \$125 for this prescription; one pharmacy sold the prescription at a discount (\$108) and captured 23 percent of the total sales to cash-payers for the six-month period; another pharmacy charged a high price of \$139, and made only one sale to a cash-payer during the six-month period. In contrast, sales of the antibiotic Cephalexin in zip code B (plotted in the second panel of figure 1) were much less responsive to price. Although prices differed by as much as \$15, sales appear to be roughly independent of the price differences.

The contrasting graphs of figure 1 show search intensities (and their determinants) are empirically identified. The sales patterns for Cimetidine suggest that some consumers price-shopped, with the prevalence of price-shopping attributable in part to the magnitude of the price differences and in part to the fact that Cimetidine is often purchased repeatedly as a maintenance medication. Sales patterns for Cephalexin (an antibiotic that is almost always purchased as a one-time prescription) suggest very little price-shopping, since if anything the low-price pharmacies make *fewer* sales than their high-price competitors. The estimated model of section 4 predicts search probabilities of 0.19 and 0.01 for Cimetidine and Cephalexin, respectively.

Interpreting the graphs of figure 1 is not entirely straightforward, since in reality not all pharmacies are homogeneous. In prescription drug markets, pharmacies with superior services or locations will have high market shares even if their prices are relatively high. In the presence of differentiation (vertical or spatial), separately identifying the intensity of search and the degree of differentiation will generally be more difficult.³ For a particular drug product, a pattern of prices and market shares may as easily result from a relatively undifferentiated market with relatively little search as from a highly differentiated market with relatively intense search.

In the analysis conducted here, the availability of data across many drug products helps overcome this identification problem. A critical assumption implicit in the empirical approach is that pharmacies are equally differentiated for all prescriptions. This essentially implies that firms' average market shares (over all prescriptions) provide information about the degree of differentiation in the market, while deviations from those market shares for a particular prescription (resulting from across-firm price differences) provide information about search intensity. I also make use of data on customers with insurance coverage: since there are fundamental *a priori* reasons to expect such customers *not* to price-shop, market share differences for insured transactions should reflect only the effects of pharmacy differentiation. In other words, market shares for purchases reimbursed by third parties reflect the market shares that would prevail in the absence of search, so that for a given prescription, intensity of search can be revealed by the degree of deviation from these shares.

³For the drugs represented in figure 1, one might consider plotting *deviations* from average market shares against price differences to crudely account for pharmacy effects; for these two drugs, doing so produces graphs that are very similar to the ones shown.

2.1 Data

As indicated in the foregoing discussion, the principal data components needed to address questions about equilibrium search intensities are prices and quantities. The data used in this study come from transaction records at pharmacies in each of two zip codes for the first six months of 1995. Prices are recorded for each transaction, and quantities can be obtained by summing transactions for any given prescription. For each transaction, the following information is available: date of transaction, drug name, drug form (e.g., tablets or caplets), prescription strength, prescription quantity, payment type (cash, third-party, or Medicaid), price, and an indicator for whether the prescription was new or a refill.

The data were provided by IMS, Inc., a pharmaceutical market research firm that collects electronic data from individual retail pharmacies to construct its annual National Prescription Audit. For confidentiality reasons, the actual zip codes and the pharmacies' identities were withheld.⁴ However, the markets came from a set of markets identified by the author as being good candidates for the study.⁵

The data cover transactions for 111 drug products.⁶ What I will refer to as a “prescription” will be a narrower definition than “drug”: for instance, $\langle \textit{Prozac}, 10\textit{mg}, 60 \textit{tablets} \rangle$ and $\langle \textit{Prozac}, 10\textit{mg}, 90 \textit{tablets} \rangle$ and $\langle \textit{Prozac}, 5\textit{mg}, 60 \textit{tablets} \rangle$ are all separate prescriptions for the drug Prozac. Since there are many combinations of quantity and strength, the data contain thousands of separate

⁴I will refer to the zip codes simply as “A” and “B” throughout the paper.

⁵I offered a list of roughly fifty cities that were geographically isolated (so that the set of pharmacies in the zip code is likely to comprise all of the available options to local consumers) but still large enough to have several pharmacies. (Some examples of markets in the set of candidates are Grand Junction, CO, Sioux City, IA, and Galveston, TX.) From the list I provided, IMS selected the markets with the best data coverage—i.e., those for which nearly all of the pharmacies provide transactions data to IMS.

⁶The drugs were selected by the author primarily based on their (high) sales volumes, and the set was assembled to cover a wide variety of drug types (e.g., antihypertensives, antidepressants, pain relievers, contraceptives, etc.).

prescriptions for which to consider price and quantity patterns. However, after eliminating low-volume prescriptions and outliers (see the data appendix for details), the dataset was reduced to 75,974 transactions for 341 prescriptions.

Although the markets were selected such that the data would comprise a near census of pharmacy transactions, in each market there were some pharmacies for which data were not available. Table 1 lists the coverage rates in each market, along with additional summary information about the number of prescriptions covered in each market and the relative shares of cash, third-party, and Medicaid transactions.

In addition to price and quantity data from pharmacy transactions, data on prescription characteristics were obtained from IMS' National Disease and Therapeutic Index (NDTI). These data come from monthly telephone surveys of over 2,000 physicians, and describe the prescribing patterns associated with each drug. In particular, the survey identifies the demographic characteristics of the patients for whom the drugs are being prescribed. In the analysis, variables on prescription characteristics related to search (for instance, measures of whether the drug is typically used for new or continued treatment) are taken from this source.

3 Price patterns and search incentives

Why should we expect consumer search to be relevant in retail markets for prescription drugs?

Retail pharmacy markets are typically local markets consisting of a relatively small number of pharmacies, since consumers tend not to travel long distances to purchase prescriptions.⁷ Since

⁷The recent rise of mail-order pharmacies has diminished the “localness” of pharmacy markets; however, at the time the data for this study were collected, mail-order purchases were a very small fraction of total prescription purchases. (Even in more recent years, mail-order purchases account for a very small proportion of total prescription sales: in 1997, 4.2 percent of prescription sales were through mail order, and in 1998, 4.7 percent.) Another exception (presumably somewhat rare) is the practice among the elderly of taking “field trips” to Canada or Mexico to purchase large amounts of chronic medications at bargain prices.

prices are not posted centrally, price comparisons are not costless: consumers interested in finding a low price must call each individual pharmacist for a price quote. Discussions in the pharmacy trade literature and the author's communications with pharmacists indicate that such phone calls occur regularly. Consumer advocacy groups actively encourage consumers to shop around for their prescriptions.⁸

In principle, at least two price patterns must be present for search to be a rational behavior. First, and most obviously, there must exist some dispersion in prices, or else consumers couldn't expect to save anything by shopping around. Retail price dispersion for prescription drugs is a well-documented phenomenon, and the pricing patterns in the data used here are generally consistent with previous surveys. Table 2 shows summary statistics for three measures of price dispersion in zip codes A and B. Numbers in the first panel reflect dispersion for all prescriptions covered in the data (a different set of prescriptions for each market). In zip code A, dispersion appears to be rather substantial: for instance, the average price range is \$10, and some prescriptions have price ranges over \$18. For the set of drugs covered in the data for zip B, prices are not as dispersed; however, the average price range is still nearly \$7. The potential savings from price-shopping are nontrivial, suggesting that some consumers may indeed choose to search before purchasing.

The second panel of Table 2 lists the same dispersion measures for the subset of 65 prescriptions that are covered in the data for both markets, so that the numbers are directly comparable across zip codes. The measures are strikingly similar in the two markets, suggesting that the dispersion phenomenon is not just a byproduct of some anomalous pricing pattern unique to one zip code or set of prescriptions. In fact, measures of dispersion are highly correlated across the two zip

⁸See, for example, the studies by the New York City Department of Consumer Affairs (1997) and the Massachusetts Public Interest Research Group (1998).

codes: the simple correlation of the standard deviations is .89, and the simple correlation of the coefficients of variation is .66, suggesting that dispersion depends on features of the prescriptions that are consistent across markets.

Even if price dispersion is present, consumer search will only be important if pharmacies' locations in the price distributions are not perfectly predictable. The presence of dispersion means that consumers expect some pharmacies to have high prices and others to have low prices; however, if consumers can accurately predict *which* stores will have high prices and which will have low prices, then search is unnecessary. The simple evidence shown in Tables 3 and 4 indicates that pharmacies' price ranks are not entirely consistent across drugs, suggesting that consumers would in fact have to search to determine which pharmacy has the low price for any given prescription. The tables group the prescriptions at each pharmacy into one of three price groups: low, middle, or high. For instance, the price at pharmacy 1 in zip A is in the lowest third of the price distribution for 31 prescriptions, in the middle third for 42 prescriptions, and in the highest third for 40 prescriptions. Especially in zip A, pharmacies' prices are not consistently low or high; instead, prices are low for some drugs and high for others. The implication is that for any given prescription, a consumer will be generally unable to predict whether a given pharmacy's price is high or low relative to the rest of the market. Of course, exceptions are present in both markets: pharmacy 7 in zip A and pharmacy 2 in zip B both appear to have consistently low prices. In general, however, pharmacies do not appear to be strictly sorted into high-price and low-price pharmacies.

Stability of prices over time also influences search decisions, since price stability affects the returns to price-shopping for repeat-purchase drugs. For prescriptions that are purchased repeatedly for chronic treatment, the benefits of finding a low-price pharmacy can be realized multiple

times if prices remain stable over time. On the other hand, if prices (and price rankings) change more frequently than prescriptions are purchased, information gained from a price search is only useful once. Table 5 presents a basic summary of price changes during the six-month period covered by the data. For every prescription-pharmacy pair, we can ask how often the price changed, and by how much.⁹ The table indicates that prices are rather stable: in both markets, prices didn't change at all over the six-month period for over half of the prescription-pharmacy pairs. Of the prices that did change, most changed only once—usually a small (less than 5%) increase.¹⁰ The fact that prices are somewhat stable over time suggests an increased role for search—especially for maintenance medications—for the reasons alluded to previously.¹¹

4 Empirical Model and Results

Based on the arguments outlined above, this study takes as given that consumer search is an important feature of the cash-paying segment of retail prescription drug markets. The aim of the empirical analysis is to learn about the nature and intensity of consumer search in these markets. The approach will be to use transactions data to estimate a discrete-choice demand model that embeds a simple search decision. That is, before purchasing a prescription, consumers will be modeled as first choosing whether or not to conduct a price search. The pharmacy at which to purchase the prescription is then selected, conditional on the information gained (or not gained) as

⁹The table presents changes in cash prices only.

¹⁰Among the prices that changed more than once, a small subset represent instances in which pharmacies were simultaneously charging two separate prices. This is presumably due to some form of discount (e.g., to senior citizens), and the differences in price are generally 10% or less.

¹¹It should be noted that if prices almost *never* change, the role of search may in fact be reduced. If prices never change, consumers need only search once, and information sharing may become more prevalent (since the information gained by one consumer from a search 8 months ago would still be useful to another consumer today). The steady-state intensity of search would therefore be lower. The data used here don't cover a long enough time frame to make any precise statements about the volatility of price rankings. Since pharmacies do occasionally change prices, and evidently do so at different times, it seems reasonable to expect price rankings to also change occasionally, and presumably often enough to induce search.

a result of the search decision.

4.1 Basic Framework

In principle, an individual consumer’s decision of whether or not to price-shop involves a complex comparison of the expected benefits and costs, which are at least partly idiosyncratic to the consumer. The model here will be somewhat simplified. The indirect utility for consumer i from purchasing prescription k at pharmacy j is given by

$$u_{ijk} = \theta_j - \alpha p_{jk} + \epsilon_{ijk}, \quad (1)$$

where θ_j is a pharmacy fixed effect intended to capture the mean (across prescriptions) of pharmacy “quality”,¹² p_{jk} is pharmacy j ’s price for prescription k , and ϵ_{ijk} is a random disturbance representing preference shocks idiosyncratic to the individual.

Notice that this specification does not include unobservable prescription-specific pharmacy effects. The utility specified in equation 1 could in principle include an extra stochastic term (say, η_{jk}) to allow for the possibility that some pharmacies are preferred for some prescriptions but not for others (for reasons unobservable to the econometrician). This may be important if pharmacies “specialize” in certain prescriptions.¹³ The implicit assumption is that non-price pharmacy characteristics are common across prescriptions for any given pharmacy, and can therefore be captured by the θ_j terms. Since most pharmacy characteristics that affect demand are not prescription-specific (e.g. location, service, selection of over-the-counter products, etc.), this restriction is probably reasonable. Also, previous research on price distributions for retail prescription drugs has suggested

¹²This term can also be regarded as capturing the effects of spatial differentiation. For instance, a pharmacy located in a densely populated area will have a relatively high θ .

¹³For example, a pharmacy located in a retirement community may have a high η for antihypertensives or anti-arthritis.

that prescription-specific effects appear to play a very limited role: for instance, Sorensen (2000) shows (in a different dataset) that interacting pharmacy effects with drugs' therapeutic categories adds little explanatory power in price regressions.

To analyze individuals' search decisions, we must clarify our assumptions about the information available to them. I will assume consumers are aware that prices are dispersed. For any given prescription, consumers assume pharmacies' prices are drawn from a common, nondegenerate distribution. Before purchasing, a consumer knows the actual minimum price in the market, p_{min} , and she assumes a pharmacy's deviation from that minimum, $p - p_{min}$, is distributed as an exponential random variable with parameter λ . This distributional assumption is made primarily for computational reasons—it obviates the computation of numerical integrals in the empirical simulations. Importantly, however, this assumption can also be regarded as data-driven: as Figure 2 illustrates, the exponential distribution can be fitted to the empirical distribution of $p - p_{min}$ remarkably well.

Rational consumers will expect that different prescriptions have different price distributions—that is, the parameter of the price distribution (λ) will differ across drugs. Rather than assume consumers are fully rational and know the *actual* λ for each prescription, in the estimation I assume that consumers form an estimate of λ based on the price level of the prescription, since consumers may expect dispersion to depend on the expensiveness of the drug. In particular, a consumer searching for prescription k takes $\hat{\lambda}_k$ to be the predicted value coming from a regression of $(p_{jk} - p_{min,k})$ on a constant and $\ln(p_{min,k})$. The implied means for $p - p_{min}$ lie between \$0.87 and \$7.83.

I will assume that search is “all or nothing”—that is, consumers either search exhaustively, learning all pharmacies' prices, or not at all. In the former case, the consumer will simply choose the pharmacy that yields the highest utility, given the *known* prices. In the latter case, the consumer

elects not to price-shop, and therefore will choose the pharmacy that yields the highest utility based on non-price considerations (i.e., the pharmacy that would be chosen if all prices were equal). Note that the search process might be more realistically described as sequential: consumers obtain price quotes, and after each successive quote assess the net benefit of obtaining another one. In that case, search “intensity” would be reflected not only in the fraction of consumers who choose to search, but also in the distribution of the number of price quotes obtained by those consumers. For purposes of estimation, however, simulation of choices based on sequential search is impractical. The simplifying assumption employed here reduces the computational burden by at least an order of magnitude. Furthermore, while the estimated search intensities may be understated due to the imposition of an all-or-nothing search technology, the basic pattern of the results is probably not sensitive to this assumption (as will be discussed in more detail below).

If a consumer chooses not to price-search, she selects the pharmacy that best matches her non-price preferences, and expects to pay the market average price. Letting u_{ik}^* denote the utility achieved without a price-search by consumer i for prescription k , we can write the following:

$$E[u_{ik}^*] = \text{Max}_j \{ \theta_j + \epsilon_{ij} \} - \alpha (p_{\min,k} + \lambda_k) . \quad (2)$$

If u_{ik}^{**} denotes the utility achieved after conducting an exhaustive price search, then the expected utility from price-shopping can be written as

$$E[u_{ik}^{**}] = E [\text{Max}_j \{ \theta_j + \epsilon_{ij} - \alpha p_{\min,k} - \alpha (p_{jk} - p_{\min,k}) \}] , \quad (3)$$

where the expectation is over prices. In Appendix B I show how the above expectation can be computed analytically, thus avoiding costly numerical integrations in the estimation routine.

Note that the dependence of the search decision on the idiosyncratic preference terms (ϵ) significantly complicates estimation, since the distribution of the ϵ terms conditional on the search decision will not generally be tractable. If instead search decisions were independent of these preference terms, then the conditional choice probabilities (given search or given no search) could take familiar multinomial logit or multinomial probit forms, depending on the assumed distribution of the ϵ 's. However, incorporating the dependence allows for more realistic (and potentially quite important) effects: namely, that individuals with relatively intense preferences for a particular pharmacy (i.e., with a very high ϵ for that pharmacy) will be unlikely to search, since they'd be unlikely to find prices low enough at other pharmacies to compensate them for foregone idiosyncratic utility that could have been obtained from the preferred pharmacy.

A consumer will choose to search if the expected gain from doing so exceeds the costs of conducting the search. The expected benefit of a price search is simply the difference in expected utilities under search vs. no search, potentially inflated to reflect the fact that information may be used repeatedly. Formally, consumer i will choose to conduct a price search if

$$\psi_{ik} (E[u_{ik}^{**}] - E[u_{ik}^*]) > \tau_{ik} \quad (4)$$

where ψ_{ik} is an inflation factor and τ_{ik} is the idiosyncratic search cost.

In estimating the model, I specify the following functional forms for the inflation factor and the distribution of search costs:

$$\psi_{ik} = \exp \{ \beta_0 \text{CHRONIC}_k + \beta_1 \text{INSURED}_{ik} \} \quad (5)$$

$$\tau_{ik} \sim \chi^2, \text{ with degrees of freedom} = \delta_0 + \delta_1 \text{FEMALE}_k \quad (6)$$

where the the prescription characteristics (from the NDTI database) are as follows:

CHRONIC: The national proportion of prescriptions for the drug that were issued as prescriptions for continuing treatment (as opposed to new prescriptions), between 0 and 1. This serves as a measure of whether the drug is used for acute or chronic conditions (a larger value implies the prescription is more often purchased on a maintenance basis). This is intended to capture the effect of repeat purchasing on search incentives: the savings from shopping for drugs that tend to be purchased repeatedly (i.e., with a high value of *CHRONIC*) can be realized multiple times, which should magnify the perceived benefit of conducting a price search.

INSURED: A binary variable equal to one if the purchase was paid or reimbursed by a third-party. For reasons explained above, insured consumers have no incentive to price-shop. In the context of the model, we expect the benefit of conducting a price search to be “deflated” to zero for insured consumers.

FEMALE: The proportion (between 0 and 1) of prescriptions that are issued to females, included to capture the possibility that mean search costs differ across men and women.

Summary statistics for these variables are shown in Table 6.

The search cost (τ) is parameterized as a chi-square random variable for two reasons. First, the chi-square distribution’s nonnegative support is consistent with the notion that search costs must (in principle) be nonnegative.¹⁴ Secondly, the fact that the chi-squared distribution has only one parameter makes it a parsimonious choice for purposes of estimation.¹⁵

Letting S_{ik} denote the event $\{\tau_{ik} < \psi_k (E[u_{ik}^{**} - u_{ik}^*])\}$, the probability that individual i will choose to purchase prescription k at pharmacy j is given by

¹⁴One can argue that some consumers may derive utility from the process of shopping (see Stahl’s (1989) discussion of bargain-hunters, for example), in which case we could think of search costs being negative for these consumers. By assuming search costs are nonnegative, we are simply asserting that any utility derived from shopping is due to price savings, not due to any enjoyment of calling pharmacists.

¹⁵Ideally, one would like to separately identify both the location and shape of the search cost distribution. The data used in this study are not rich enough to enable identification at this level, however. For example, the estimation algorithm is generally unable to converge to an optimum when the search cost distribution is parameterized as lognormal. Heuristically, the intuition is this: the important variation in the data is the observed sensitivity of market shares to price differences. As explained previously, what this variation really identifies is the intensity of search in equilibrium. Roughly speaking, for a distribution like the lognormal with separate parameters governing the mean and variance, the distribution can generate a particular search intensity either by holding the mean fixed and changing the variance, or by holding the variance fixed and changing the mean. So separately identifying the two parameters is very difficult (if not impossible).

$$\begin{aligned}
P_{ik}(j) &= P(S_{ik}) \cdot P(\{\theta_j - \alpha p_{jk} + \epsilon_{ijk} > \theta_h - \alpha p_{hk} + \epsilon_{ihk} \forall h \neq j\} | S_{ik}) + \\
&[1 - P(S_{ik})] \cdot P(\{\theta_j + \epsilon_{ijk} > \theta_h + \epsilon_{ihk} \forall h \neq j\} | \neg S_{ik})
\end{aligned} \tag{7}$$

The log-likelihood for the sample can then be written as

$$\ell(\theta, \alpha, \beta, \delta) = \sum_k \sum_j n_{jk} \ln P_{ik}(j) \tag{8}$$

where n_{jk} is the number of times prescription k was purchased at pharmacy j .

Estimation by maximum likelihood is complicated by the fact that $P_{ik}(j)$ cannot be computed in closed form.¹⁶ The parameters are therefore estimated by maximizing a simulated likelihood function, $\hat{\ell}$.¹⁷ The simulation proceeds as follows: (1) for each prescription, a sample of r consumers is drawn, with each consumer characterized by a vector of idiosyncratic preferences and a search cost: (ϵ_i, τ_i) . The preference term ϵ is drawn from the Type I Extreme Value distribution, and τ is drawn from the chi-square distribution as described in equation 6; (2) each simulated consumer's search decision is calculated according to equation 4; (3) conditional choice probabilities (conditional on the search decision) are calculated using a kernel-smoothed frequency simulator.¹⁸ The conditional choice probabilities and simulated search decisions are used to construct the total choice probabilities for each prescription, and these simulated probabilities are then plugged into the log-likelihood function to construct $\hat{\ell}$. The simulations are performed at each iteration of the

¹⁶Note that a closed form for the choice probabilities does not exist even if we make the otherwise convenient assumption of a type-I extreme value distribution for the idiosyncratic preference terms, ϵ . Since these terms affect the search decision, the *conditional* distribution of ϵ 's given search (or no search) doesn't have the same easily integrable form as the unconditional distribution.

¹⁷See Lerman and Manski (1981) for another example and discussion of this approach.

¹⁸I use the logit kernel suggested in McFadden (1989) and Geweke, Keane, and Runkle (1994) with smoothing parameter $\rho = 0.15$.

search algorithm; separate simulation samples are drawn for each prescription, but the samples stay fixed across iterations for obvious computational reasons.

For the results to be reported in the next section, the model was estimated without an outside option. The model therefore represents choice probabilities *conditional* on having chosen to purchase from one of the pharmacies in the dataset. The outside option would consist of either purchasing from one of the few pharmacies not covered in the data, or not purchasing at all. Existing evidence in the medical literature on prescription compliance suggests the fraction of consumers who elect not to purchase a prescription even after its authorization by a physician may be as large as 40 percent. Importantly, compliance is likely to vary considerably across prescriptions, depending for instance on the cost of the prescription and the health status of the patient. This is the primary motivation for omitting the outside option from the estimated model: the “market size” should be expected to vary dramatically across prescriptions, and estimates of market size based on currently available data would tend to be very unreliable. The outside option is omitted so that results regarding search behavior would not be driven by ad hoc estimates of compliance rates. Note that the omission should not pose a problem for the analysis of search; the model effectively examines intensity of search among only those pharmacies covered by the data. The primary cost of not including an outside option is that price elasticities become difficult to interpret: the estimated price elasticities reflect only the price-induced change in demand from consumers who have chosen to purchase, not the change in total demand.

4.2 Results

Table 7 presents model results using cash transactions only. For the estimates reported, the simulation size (r) is 500. Because simulation noise can bias estimates of the standard errors, the reported

standard errors were computed using the robust covariance matrix estimator recommended by McFadden and Train (2000) for Maximum Simulated Likelihood estimation (MSLE).

The parameter estimates are precisely estimated and broadly consistent with predictions of consumer search theory. The positive coefficient on the *CHRONIC* variable indicates that the expected benefits from price-searching are greatest for prescriptions issued predominantly for continuing treatment. For example, the estimates imply that if two prescriptions (say X and Y) are prescribed for continued treatment 90 percent of the time and 80 percent of the time (respectively), then the expected benefits of search will be “inflated” 15.7 percent more for X than Y in zip A (10.0 percent more in zip B). The bottom panel of the table shows the minimum, median, and maximum multipliers (ψ) implied by the estimates for each zip code. For prescriptions that tend to be purchased strictly on a one-time basis (low values of *CHRONIC*), the inflation factor is close to one, whereas for repeatedly purchased medications the inflation factor can be as high as 3.94 or 2.45 in zips A and B, respectively. These numbers are not implausible: assuming no discounting, they suggest consumers expect to capture the benefits of a price search 2 or 3 times before the information “expires” (e.g., due to changes in the distribution of prices).

The means of the search cost distributions depend on the parameters δ_0 and δ_1 , the estimates of which are remarkably similar across the two markets. Dollar equivalents of the implied search costs are shown in the bottom panel. Given the estimated parameters, I simulated the model and computed the average search costs separately among consumers who searched and those who did not. (Presumably, the former number is more cleanly identified by the data, since it is driven by the implied location of the search vs. no-search margin. The latter number may be driven largely by the chi-square functional form assumption for the search cost distribution.) Among consumers

who price-shopped, the estimated average search cost was \$12.48 in zip code A and \$4.96 in zip code B. Recall that this represents the implied cost of an exhaustive price search—that is, the cost of obtaining price information from all of the pharmacies in the market. Given that price searches in these markets are conducted by phone-calling for price quotes, one can imagine that the costs consist of opportunity costs of time and perhaps psychic costs incurred from having to harass busy pharmacists for price quotes.

The estimated coefficient on the *FEMALE* variable suggests that average search costs are lower for females than males. Much of this effect is apparently attributable to contraceptives: when the model is estimated with contraceptives dropped from the sample, the coefficient on *FEMALE* is reduced in magnitude by about half (but is still statistically significant). This result is consistent with pharmacists' reports that oral contraceptives tend to be among the most heavily price-shopped prescriptions.¹⁹ The result could also be a reflection of differences across men and women in the average opportunity cost of time.

The search probabilities implied by the model (shown in the bottom panel of the table) fall within a plausible range. The estimates suggest that the vast majority of prescriptions are sold to consumers who do not shop around for price. In zip A, the prescription least likely to be price-shopped has an estimated search probability of 0.000008, while the most heavily price-shopped has a search probability of 0.224. In zip B, search probabilities fall between 0.000002 and 0.267.

These estimates fit well with reports from a few pharmacists from whom I collected anecdotal

¹⁹For example, one pharmacy newsletter reports that contraceptives “are not advertised in in-store circulars or newspaper inserts, but women frequently call pharmacies in their area to do some comparative shopping for the best price on a particular oral contraceptive.” (*Supermarket News*, 1995).

data.²⁰

Importantly, the model embedding the parameterized search decision appears to fit the data substantially better than a simple logit model without search. Likelihood-ratio tests strongly reject the hypothesis that the search probability can simply be constrained to equal one (that is, that the model can simply be written as a multinomial logit model with pharmacy and price effects only).²¹

The results shown in table 7 represent the transactions of cash-paying customers only. There is good reason to expect that such customers are the only ones for whom price-shopping is relevant: consumers with insurance coverage for prescriptions typically pay a flat copayment regardless of which pharmacy they purchase from, as long as the pharmacy is within the network of coverage. Moreover, even if the insurance requires a percentage coinsurance payment, third-party payers typically negotiate equal reimbursement rates across pharmacies within a local market, so that prices facing an insured customer don't vary across pharmacies.²²

Even if insured consumers don't have any incentive to price-shop, the transactions of these consumers contain useful information for estimating the parameters of the model. In particular, these transactions aid in estimation of the pharmacy effects (θ), since the purchases of insured consumers are presumably based on differences in pharmacy quality or location, not on price.²³

Table 8 shows estimates analogous to the ones described previously, using data on both cash

²⁰I queried three pharmacists about the rate at which they typically receive phone calls for price quotes and the rate at which they process prescriptions. Their answers implied that they receive between 2-8 phone calls for every 100 cash prescriptions they process.

²¹The χ^2_3 test statistics and p -values are 74.6 and 0.000, and 139.7 and 0.000, for zips A and B respectively.

²²This information was confirmed in a conversation with Kristine Zawicki, coordinator of the clinical pharmacy program at Blue Cross of Massachusetts.

²³One caveat is that insured customers as a group may have different preferences than cash-paying customers, in which case the pharmacies should be thought of as having separate fixed effects for the two groups. Also, if some pharmacies are not included in some insurance plans' networks, inclusion of the third-party transactions may distort the estimates of the pharmacy effects. The contention here is that neither of these effects is likely to be important, so that transactions of insured customers can be appropriately used to help identify a set of common pharmacy fixed effects.

and third-party transactions. Since the arguments outlined above suggest the search decisions of insured and uninsured consumers are intrinsically different, the model is parameterized to allow insured customers to perceive the potential benefits of search differently than cash-paying customers. Specifically, the multiplicative term (ψ) on the expected benefit contains an indicator variable for whether or not the transaction was covered by a third-party (equation 5). As the table shows, the estimated coefficient on this *INSURED* dummy is large and negative for both zip codes, indicating that for insured customers the perceived benefits of price-shopping are essentially zero. Note that this feature of the estimates is not imposed by the model; rather, it is a direct consequence of insured customers' purchase patterns being roughly independent of price in the data.

Estimates of the coefficients on other search-related prescription characteristics are consistent with those from the cash-only specification: the expected benefits from search are magnified for maintenance medications, leading to higher search probabilities for such medications; and females appear to have lower average search costs than males (as revealed by the negative coefficient on the *FEMALE* variable.) Taking the estimates at face value, we would infer that the average female in zip code A (B) has a search cost \$11.12 (\$11.76) lower than the average male. The implied search costs (in dollars), shown in the bottom panel of the table, are lower than the search costs implied by the cash-only estimates. For consumers who chose (in a simulation) to shop for prices, the average search cost was \$5.97 in zip A and \$3.95 in zip B. The lower estimated search costs among price-shoppers in zip B most likely reflect the lower levels of price dispersion in zip B (see table 2): less dispersion implies lower expected savings from search, so (other things equal) the search cost threshold must also be lower in zip B. The unconditional average search cost (i.e., the dollarized τ averaged across prescriptions for both shoppers and non-shoppers) is \$15.47 in zip

code A and \$14.07 in zip code B. Given that these search costs should primarily reflect the time costs of calling several pharmacists for price quotes (a process that would probably take 15-30 minutes in these markets), their magnitudes seem plausible. Also, the apparent similarity of search costs across markets is rather remarkable given that the model is estimated separately for each market.

The implied search probabilities, also shown in the bottom panel of Table 8, are slightly higher than those implied by the cash-only estimates. However, they still suggest relatively low levels of search: the median search probability in zip code A is 10.1 percent, and in zip code B this probability is 5.0 percent. Even the prescriptions for which search is apparently most intense have implied search probabilities around 35 percent.

As mentioned above, one of the principal advantages of including insured transactions in the estimation is to aid in identifying pharmacy-specific “quality” effects (the θ 's in equation 1) separately from the effects of search. It is therefore worth noting that differences between the estimated pharmacy effects (not shown in the tables) across the two specifications reported in Tables 7 and 8 are very small. The implied ordinal rankings of the pharmacies change only slightly, and (as would be expected) the standard errors shrink substantially when third-party transactions are included in the estimation.

As an additional illustration of the conceptual basis for the estimation, Table 9 lists a sample of prescriptions and their estimated search probabilities. The last two columns of the table show information about the sales volumes at the lowest-price stores. For low-search prescriptions like Cephalexin (a common antibiotic), the market shares of the low-price stores deviate little from their average market shares across all drugs. Allocation of sales across pharmacies for low-search

drugs appears to be approximately random, with probability weights equal to the pharmacies' average overall market shares. On the other hand, for high-search prescriptions like Triphasil (oral contraceptive) or Mevacor (cholesterol reducer), low-price stores capture a disproportionate share of the market. For Mevacor in zip code B, for example, the two stores with the lowest prices ordinarily combine for a 21 percent share of the market, but their low prices for this particular prescription attract price-shoppers, and the combined market share is 42 percent.

4.3 Robustness

The empirical model described and estimated above is obviously stylized, so in this section I briefly discuss the robustness of the results to changes in various assumptions. One strong assumption of the model is that search is all or nothing—that is, consumers who elect to search obtain price quotes from all pharmacies. A more realistic model would describe search as a sequential process: after each price quote is obtained, the consumer reevaluates the expected net gains from obtaining another one. Changing the empirical model to incorporate sequential search is in principle straightforward; however, the computational burden of doing so is prohibitive.²⁴ Though I cannot say for certain, it seems likely that the results would be more or less unaffected if the model were one of sequential search. Ultimately, the “best” two or three pharmacies drive the consumers' search and purchase decisions even when search is all-or-nothing, and these are exactly the firms that would be contacted first in a sequential search. A simpler (and more easily tested) modification of the search technology is to limit the number of price quotes obtained, since in reality consumers are unlikely

²⁴ Among other issues, modeling search as sequential requires assumptions about the order in which price quotes are obtained. The most obvious approach would be to have consumers sample pharmacies in the order of preference for non-price characteristics. Since this preference ordering is different across consumers, the search process would have to be simulated separately for each consumer, which for reasonably large simulation sizes increases computation time dramatically. (For the model described above, consumers' search decisions are computed simultaneously as matrix calculations.)

to sample all pharmacies. When the model is modified so that price-shoppers obtain price quotes from only a subset of pharmacies, the results are very similar to those reported above.²⁵ Again, this is not surprising: since the model allows for heterogeneous pharmacy quality and idiosyncratic consumer preferences, only the prices of the “best” pharmacies will ultimately be relevant for any given consumer.

The results are also relatively robust to alternative functional form assumptions about the distribution of consumers’ idiosyncratic preferences (ϵ). Columns one and two of table 8 show the estimated parameters for zip B when ϵ is assumed to be uniformly or normally distributed, respectively. (The parameters of the distributions are chosen so that the variances match the variance of the Type I Extreme Value distribution.) The most notable change from the results of table 8 is the diminution of the price coefficient (α) and the constant (δ_0), indicating that these parameters are somewhat sensitive to functional form assumptions. However, the estimated search intensities and implied search costs are very similar to those reported in table 8.

5 Discussion

If the fraction of consumers who price-shop is generally low even in the face of substantial price dispersion (as the estimates described above seem to suggest), an implication is that search costs prevent most consumers from fully optimizing their purchase decisions. That is, many consumers will fail to find “bargains” because the cost of price-shopping is too high—they rationally resign themselves to making decisions based on incomplete information. An obvious question that arises, therefore, is how much utility is lost due to this inhibiting effect of search costs? If we could reduce or eliminate search costs, what would be the expected increase in consumer welfare?

²⁵Parameter estimates assuming 3,5, and 7 sampled pharmacies are available from the author upon request.

In most markets where consumer search is important, one can imagine public policies or market institutions that would reduce or eliminate the costs of search. For instance, the rise of internet price-comparison sites for computer hardware and software, books, home electronics, and other consumer goods illustrates the role of easily accessible, centralized price information in facilitating search and (presumably) intensifying competition. Recent work by Goolsbee and Brown (2000) suggests that internet price-shopping for life insurance has intensified competition and led to substantial gains in consumer welfare. In retail markets for prescription drugs, some public policies may inhibit search. For instance, most states prohibit price advertising by pharmacies. While valid motivations for such a proscription presumably exist, one undesirable consequence may be a dearth of price information and an increased cost of search.

Using the estimated model described in the previous section, we can easily simulate one counterfactual: What would be the partial equilibrium consequences of an exogenous decrease in the costs of search? Table 11 summarizes the results from simulating the purchase decisions of consumers assuming varying degrees of search cost reductions. Importantly, note that the numbers reflect only partial equilibrium effects; price distributions are assumed fixed. The results are shown separately for each zip code. The first row indicates how the fraction of consumers who price-shop increases as search costs decrease. Even after an exogenous 50 percent reduction in search costs, fewer than 42 percent of prescription transactions in zip code A would be price-shopped (31 percent in zip code B). The second row lists the average equivalent variation, or the price reduction that would yield an equivalent change in the average utility obtained per transaction. In zip code A, for instance, a 50 percent reduction in search costs would yield the same increase in average utility as would a rebate of \$1.49 on every transaction (\$0.63 in zip code B). Note that the numbers

reflect changes in gross utility—i.e., not including the costs of search themselves.

In general, the numbers shown in Table 11 suggest relatively modest consumer welfare gains, especially for partial reductions of search costs. To understand this, consider the composition of partial equilibrium welfare changes caused by a small search cost reduction. The gross utility of consumers who would have searched even prior to the reduction does not change—they price-shop both before and after the reduction in search costs, so their purchase decisions are unaffected.²⁶ Similarly, the change will not affect non-shoppers for whom the search cost reduction is too small to induce them to shop. The only affected group is the “new” price-shoppers—those who choose to search only because of the reduced search costs. These consumers will on average achieve higher utility than before (as their searches lead to the discovery of better “matches” or “bargains”). However, the gains experienced by this group will necessarily be small: since they were the consumers at the margin of the search/no-search decision, their expected benefits from search are by definition very nearly equal to their search costs, so that small reductions in search costs can lead to only small expected net benefits for these consumers.

More importantly, the numbers shown in the table may understate the potential impact of search cost reductions because they only reflect partial equilibrium effects. In a general equilibrium framework, decreases in search costs will not only change the search decisions of consumers, but will also lead to changes in the price distributions themselves. In particular, a typical search model would predict that search cost reductions result in lower average prices and less disperse price distributions. The empirical model employed here does not explicitly incorporate the supply-side price-setting equilibrium, and therefore cannot address these general equilibrium considerations directly. However, one can speculate that the welfare changes resulting from the

²⁶A policy leading to search cost reductions can be thought of as a pure windfall for these consumers.

intensification-of-competition effect of lowering search costs would be substantially larger than the partial equilibrium effects described above. Accommodating general equilibrium effects in the empirical model is a high priority for future extensions of this research.

6 Conclusion

What does this study tell us about the nature of consumer search for retail prescription drugs? First of all, the estimates suggest that price-shopping is relatively uncommon, even though prices are sometimes very disperse. Secondly, consumers' propensities to price-shop vary substantially across prescriptions, and patterns in this variation are consistent with what we would expect from search theory. For instance, maintenance medications (e.g., antihypertensives) are searched more heavily than one-time prescriptions (e.g., antibiotics), since the benefits of finding a low price are magnified for prescriptions that will be purchased repeatedly. Also, patterns in the data indicate that consumers with insurance coverage do not price-shop at all, which is exactly what we would expect. Finally, under some relatively strong assumptions about the search technology and consumers' beliefs about price distributions, we can draw inferences about the level of search costs that would generate the estimated price-shopping patterns. The estimates imply average costs of approximately \$15 for conducting an exhaustive price search, with these costs being somewhat lower on average for females than for males.

Simulations of the partial equilibrium effects of exogenous search cost reductions suggest that such reductions would lead to modest gains in consumer welfare. Ideally, the model here could be extended to incorporate pharmacies' equilibrium price-setting decisions in the estimation, in which case estimates of the model could be used to predict the welfare impact of search cost reductions

in a general equilibrium framework. This is a potentially interesting path for future research; the difficult trick will be how to model the supply side of equilibrium price dispersion in an empirically tractable way.

More generally, this work points to the usefulness and relevance of consumer search models. Amid claims of an “information revolution,” in which information is disseminated with increasing ease, speed, and breadth (via electronic or other means), important economic questions arise about the effects of such changes on market outcomes. Applications of consumer search models will likely play a fundamental role in answering such questions.

Appendices

A Data

This section describes the process by which the data used in the analysis were selected. The original raw data contain 259,328 transactions over the six-month period, covering 111 drug products and 10,397 possible prescriptions for those products. The transactions from one pharmacy were dropped, because IMS discovered they were reporting prices incorrectly after I pointed out some anomalies in their price patterns. Prescriptions with very low sales volumes were also dropped; to be included in the final data, the number of prescription sales had to be more than three times the number of pharmacies in the zip code. Prescriptions with very low prices were also dropped: the cutoff was a mean price (across stores) of \$2.00.

A nontrivial number of anomalous transactions existed in the raw data. After fixing obvious key-punch errors (e.g., decimal place errors), “outliers” were dropped if they met either or both of the following criteria:

- $\text{price} < (.5\text{AWP})$
- $\left\{ \frac{\text{price}}{\text{AWP}} \notin [.6, 2] \right\} \text{ and } \left\{ |\text{price} - \text{AWP}| > 20 \right\}$

Roughly 5 percent of the total number of transactions were eliminated on this basis.

A requirement for estimating the model of pharmacy choice is that price data be available at all relevant pharmacies. For many prescriptions, although many sales were made during the quarter, there were one or two pharmacies that did not make sales. The absence of any transaction for a prescription at a pharmacy meant that there was no price to associate with the zero quantity. Rather than lose such prescriptions, I filled in the “missing” prices with estimates based on the pharmacies’ prices for prescriptions sold to insured customers. That is, I estimated an equation of cash prices as a linear function of third-party prices, with each pharmacy having its own separate slope coefficient. (The R^2 from this regression is 0.996.) Therefore, if a pharmacy had zero cash sales but one or more insured sales for a prescription, I assign the *predicted* cash price to the zero cash quantity. I do this only for prescriptions where less than a quarter of the pharmacies in a zip code had zero cash sales. This method of “filling in” missing price information increases the

number of prescriptions available (over the alternative of using only prescriptions that were sold at least once at *every* pharmacy) by roughly 25 percent.

B Computing the expected utility under search

This section describes how to compute the expected utility given search, assuming the vector of idiosyncratic preferences (ϵ) is known to the consumer (and therefore incorporated into the expectation). The model estimated in section 4 does not include an outside good; for the sake of generality, the analysis here includes an outside good (but of course is easily translated to the no-outside-good case).

Recall the assumption that the minimum price in the market is known to consumers, and that the deviation from the minimum is distributed exponentially with parameter λ_k . Let Y denote the random variable representing $p - p_{min}$, and let Z denote the random variable representing the maximum of the utilities associated with each option (drop the i and k subscripts for notational convenience):

$$Z = \text{Max}_j \{[\theta_j + \eta_j + \epsilon_j - \alpha p_{min} - \alpha Y], [\theta_0 + \eta_0 + \epsilon_0]\} \quad (9)$$

Note that $E[Y] = \frac{1}{\lambda}$ and

$$1 - F_Y(y) = \begin{cases} e^{-\lambda y} & \text{if } y > 0 \\ 1 & \text{otherwise} \end{cases}$$

Letting $\tilde{u}_j = \theta_j + \eta_j + \epsilon_j - \alpha p_{min}$ and writing $G(\cdot)$ for the CDF of Z , we have that

$$\begin{aligned} E[Z] &= \bar{z} - \int_{\underline{z}}^{\bar{z}} G(z) dz \\ &= \bar{z} - \int_{\underline{z}}^{\bar{z}} \mathbf{I}(u_0 \leq z) \prod_{j=1}^J \left(\mathbf{I}(z < \tilde{u}_j) \cdot e^{-\frac{\lambda}{\alpha}(\tilde{u}_j - z)} + [1 - \mathbf{I}(z < \tilde{u}_j)] \right) dz \end{aligned} \quad (10)$$

where $\mathbf{I}(\cdot)$ denotes the indicator function and the extremal values of Z are given by

$$\begin{aligned} \underline{z} &= u_0 \\ \bar{z} &= \text{Max}_j \{u_0, \tilde{u}_j\} . \end{aligned} \quad (11)$$

Since $\underline{z} = u_0$, the first indicator function in equation 10 is always equal to one over the range of integration. However, the integrand changes as we move over the range of integration, so the integral must be computed piece-wise. Noting this, and incorporating all the information above, equation 10 can be written more directly:

$$E[Z] = \text{Max}_j \{u_0, \tilde{u}_j\} - \sum_{\ell=0}^{J-1} \left[\mathbf{I}(\tilde{u}_j > u_0) \cdot \int_{\tilde{u}_\ell}^{\tilde{u}_{\ell+1}} \exp \left\{ -\frac{\lambda}{\alpha} \left[\left(\sum_{j=\ell+1}^J \tilde{u}_j \right) - (J-\ell)z \right] \right\} dz \right], \quad (12)$$

where the embedded integral can be computed analytically as

$$\begin{aligned} \int_{\tilde{u}_\ell}^{\tilde{u}_{\ell+1}} \exp \left\{ -\frac{\lambda}{\alpha} \left[\left(\sum_{j=\ell+1}^J \tilde{u}_j \right) - (J-\ell)z \right] \right\} dz = \\ \frac{\alpha}{\lambda(J-\ell)} \exp \left\{ -\frac{\lambda}{\alpha} \left(\sum_{j=\ell+1}^J \tilde{u}_j \right) \right\} \left[\exp \left\{ \frac{\lambda}{\alpha} (J-\ell)\tilde{u}_{\ell+1} \right\} - \exp \left\{ \frac{\lambda}{\alpha} (J-\ell)\tilde{u}_\ell \right\} \right] \end{aligned} \quad (13)$$

The expected utility conditional on conducting a price-search is just $E[Z] - \tau$, where τ is the (individual-specific) search cost. In the estimation algorithm, this expectation is compared with the expected utility conditional on no search to determine an individual's search decision. The distributional assumption on $Y = p - p_{min}$ enables analytical computation of the relevant integrals (equation 13), dramatically reducing computation time.

References

- (1995, October 9). "The Pull of the Pill". *Supermarket News*.
- (1997). 1997 Prescription Drug Survey. New York City Department of Consumer Affairs.
- (1998). "Prescription for Savings". Massachusetts Public Interest Research Group Education Fund, September 1998.
- Burdett, K. and K. Judd (1983, July). Equilibrium Price Dispersion. *Econometrica* 51(4), 955–969.
- Geweke, J., M. Keane, and D. Runkle (1994, November). Alternative Computational Approaches to Inference in the Multinomial Probit Model. *Review of Economics and Statistics* 76(4), 609–632.
- Goolsbee, A. and J. Brown (2000). Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry. Mimeo, October 2000.
- Lerman, S. and C. Manski (1981). *On the Use of Simulated Frequencies to Approximate Choice Probabilities*, pp. 305–319. MIT Press.
- McFadden, D. (1989, September). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica* 57(5), 995–1026.
- McFadden, D. and K. Train (2000, September-October). Mixed MNL Models for Discrete Response. *Journal of Applied Econometrics* 15(5), 447–470.
- Rice, J. (1984). Boundary Modification for Kernel Regression. *Communications in Statistics: Theory and Methods* 13, 893–900.
- Rob, R. (1985, July). Equilibrium Price Distributions. *Review of Economic Studies* 52(3), 487–504.
- Sorensen, A. (2000, August). Equilibrium Price Dispersion in Retail Markets for Prescription Drugs. *Journal of Political Economy* 108(4), 833–850.
- Stahl, D. O. (1989, September). Oligopolistic Pricing with Sequential Consumer Search. *American Economic Review* 79(4), 700–712.
- Stahl, D. O. (1996). Oligopolistic Pricing with Heterogeneous Consumer Search. *International Journal of Industrial Organization* 14(2), 243–268.

Table 1: Summary of transactions data

| Zip | Coverage | Transactions | #Prescriptions | % Cash | % Third | % Med |
|-------|----------|--------------|----------------|--------|---------|-------|
| A | 14/20 | 41,969 | 160 | 37.4 | 54.9 | 7.7 |
| B | 10/14 | 34,005 | 181 | 48.9 | 42.4 | 8.7 |
| Total | 24/34 | 75,974 | 341 | 42.6 | 49.3 | 8.1 |

Numbers of transactions and prescriptions are from the reduced dataset (after eliminating outliers); frequencies of payment types (cash/third/Medicaid) reflect percentages from the original raw data.

Table 2: Summary statistics for measures of price dispersion

| All prescriptions in each zip code | | | | | | |
|---|---------------|-------------|-----------------|------------------|------------|------------|
| | | Mean | Std.Dev. | Quantiles | | |
| | | | | .10 | .50 | .90 |
| Zip A | Range | 10.02 | 8.92 | 2.85 | 7.50 | 18.74 |
| | Std. Dev. | 3.39 | 3.10 | .96 | 2.47 | 7.05 |
| | Coef. of Var. | .08 | .04 | .04 | .07 | .13 |
| Zip B | Range | 6.81 | 4.92 | 1.87 | 5.49 | 13.63 |
| | Std. Dev. | 2.53 | 1.68 | .83 | 2.28 | 4.82 |
| | Coef. of Var. | .09 | .05 | .05 | .07 | .14 |

| Prescriptions covered in both zips A and B (65 total) | | | | | | |
|--|---------------|-------------|-----------------|------------------|------------|------------|
| | | Mean | Std.Dev. | Quantiles | | |
| | | | | .10 | .50 | .90 |
| Zip A | Range | 7.57 | 5.21 | 2.80 | 5.65 | 14.78 |
| | Std. Dev. | 2.51 | 1.71 | .81 | 2.12 | 4.90 |
| | Coef. of Var. | .08 | .03 | .04 | .07 | .13 |
| Zip B | Range | 6.98 | 4.48 | 2.01 | 5.65 | 13.02 |
| | Std. Dev. | 2.61 | 1.68 | .79 | 2.32 | 4.82 |
| | Coef. of Var. | .08 | .03 | .05 | .07 | .11 |

Table 3: Price rankings in zip code A

| Pharmacy | Price Group | | |
|----------|-------------|--------|---------|
| | Lowest | Middle | Highest |
| A.1 | 31 | 42 | 40 |
| A.2 | 18 | 30 | 65 |
| A.3 | 18 | 51 | 44 |
| A.4 | 31 | 37 | 45 |
| A.5 | 42 | 38 | 33 |
| A.6 | 43 | 49 | 21 |
| A.7 | 111 | 1 | 1 |
| A.8 | 40 | 38 | 35 |
| A.9 | 30 | 28 | 55 |
| A.10 | 36 | 33 | 44 |
| A.11 | 14 | 31 | 68 |
| A.12 | 82 | 13 | 18 |
| A.13 | 40 | 29 | 44 |
| A.14 | 29 | 32 | 52 |

Groupings based on price orderings across stores.

Table 4: Price rankings in zip code B

| Pharmacy | Price Group | | |
|----------|-------------|--------|---------|
| | Lowest | Middle | Highest |
| B.1 | 15 | 54 | 59 |
| B.2 | 126 | 2 | 0 |
| B.3 | 7 | 70 | 51 |
| B.4 | 10 | 54 | 64 |
| B.5 | 93 | 23 | 12 |
| B.6 | 5 | 83 | 40 |
| B.7 | 16 | 72 | 40 |
| B.8 | 15 | 62 | 51 |
| B.9 | 88 | 23 | 17 |
| B.10 | 9 | 69 | 50 |

Groupings based on price orderings across stores.

Table 5: Summary of price changes over time

| | Zip A | Zip B |
|------------------------|--------------|--------------|
| No price change (%) | 56.54 | 58.83 |
| One price change (%) | 30.70 | 30.07 |
| Decrease by >5% | 1.87 | 2.19 |
| Decrease by <5% | 7.69 | 9.84 |
| Increase by <5% | 76.30 | 81.97 |
| Increase by >5% | 14.14 | 6.01 |
| >one price changes (%) | 12.76 | 11.09 |

Table 6: Summary statistics for search-related prescription characteristics

For unique prescriptions ($N=176$) in zips A & B:

| | Mean | Std. Dev. | Quantiles | | |
|----------------|-------------|------------------|------------------|------------|------------|
| | | | .10 | .50 | .90 |
| <i>CHRONIC</i> | .60 | .29 | .14 | .71 | .90 |
| <i>FEMALE</i> | .62 | .15 | .51 | .56 | .84 |

Weighted by transaction volume:

| | Mean | Std. Dev. | Quantiles | | |
|----------------|-------------|------------------|------------------|------------|------------|
| | | | .10 | .50 | .90 |
| <i>CHRONIC</i> | .66 | .26 | .16 | .77 | .90 |
| <i>FEMALE</i> | .65 | .18 | .51 | .56 | 1.00 |

Table 7: Model Estimates: Cash Transactions Only

| | Zip A | Zip B |
|----------------------------|-------------------|-------------------|
| Price (α) | 0.261 (0.065) | 0.525 (0.026) |
| Pct. Chronic (β_0) | 1.458 (0.254) | 0.955 (0.119) |
| Constant (δ_0) | 13.482 (1.809) | 13.001 (0.851) |
| Pct. Female (δ_1) | -6.892 (1.174) | -6.741 (0.729) |
| Multiplier (ϕ): | | |
| Min | 1.124 | 1.059 |
| Median | 3.165 | 1.723 |
| Max | 3.939 | 2.453 |
| Prob[search]: | | |
| Min | 0.000008 | 0.000002 |
| Median | 0.031 | 0.026 |
| Max | 0.224 | 0.267 |
| Search costs: | | |
| Average search | 12.48 | 4.96 |
| Average no search | 36.16 | 16.45 |
| Observations | 21,655 | 21,388 |
| Avg. Log-likelihood | -2.422 | -2.074 |

Estimates obtained by maximizing the simulated log-likelihood function described in the text. Standard errors (in parentheses) were computed using the robust covariance matrix estimator recommended by McFadden and Train (2000) to account for simulation noise. The minima (medians/maxima) of the second panel are the minima (medians/maxima) across prescriptions in the sample.

Table 8: Model Estimates: Cash and Insured Transactions

| | Zip A | Zip B |
|------------------------------------|-------------------|-------------------|
| Price (α) | 0.524 (0.024) | 0.484 (0.020) |
| Pct. Chronic (β_0) | 0.482 (0.068) | 0.826 (0.096) |
| Insured (β_1) | -3.199 (0.301) | -4.760 (0.379) |
| Constant (δ_0) | 11.726 (0.590) | 10.416 (0.532) |
| Pct. Female (δ_1) | -5.826 (0.498) | -5.693 (0.467) |
| Multiplier (cash prescriptions): | | |
| Min | 1.039 | 1.051 |
| Median | 1.463 | 1.601 |
| Max | 1.573 | 2.174 |
| Prob[search] (cash prescriptions): | | |
| Min | 0.00009 | 0.00004 |
| Median | 0.101 | 0.050 |
| Max | 0.366 | 0.348 |
| Search costs: | | |
| Average search | 5.97 | 3.95 |
| Average no search | 16.57 | 14.19 |
| Observations | 41,969 | 34,005 |
| Avg. Log-likelihood | -2.467 | -2.082 |

Estimates obtained by maximizing the simulated log-likelihood function described in the text, using data on transactions paid in cash and transactions reimbursed by third parties. Standard errors (in parentheses) were computed using the robust covariance matrix estimator recommended by McFadden and Train (2000) to account for simulation noise. The minima (medians/maxima) of the second panel are the minima (medians/maxima) across prescriptions in the sample, and reflect only prescriptions paid for in cash. For insured transactions, estimated search probabilities are uniformly less than 0.0001 in both zip codes.

Table 9: Sample of prescriptions and estimated search probabilities

| Prescription | Primary Use | Average Price | Estimated Search Prob. | Mkt. Share of two lowest-price stores | Avg. Share for those two stores |
|--------------------------------|--------------------|---------------|------------------------|---------------------------------------|---------------------------------|
| Lanoxin 0.125mg (30 tabs) | cardiac glycoside | \$2.53 | 0.0002 | 0.25 | 0.24 |
| Tylenol+Codeine 30mg (30 tabs) | pain reliever | \$7.22 | 0.0025 | 0.25 | 0.28 |
| Cephalexin 500mg (28 caps) | antibiotic | \$25.36 | 0.0073 | 0.07 | 0.12 |
| Lotrisone 15g (cream) | antifungal | \$16.35 | 0.0077 | 0.22 | 0.22 |
| Tussionex 120ml (oral susp.) | antitussive | \$18.73 | 0.028 | 0.26 | 0.28 |
| Biaxin 500mg (14 tabs) | antibiotic | \$41.26 | 0.046 | 0.21 | 0.24 |
| Premarin 0.9mg (30 tabs) | conj. estrogens | \$13.76 | 0.137 | 0.32 | 0.26 |
| Capoten 25mg (100 tabs) | antihypertensive | \$65.71 | 0.166 | 0.21 | 0.12 |
| Triphasil 28 (28 tabs) | oral contraceptive | \$23.87 | 0.187 | 0.35 | 0.24 |
| Prozac 20mg (30 caps) | antidepressant | \$64.37 | 0.209 | 0.35 | 0.28 |
| Mevacor 20mg (100 tabs) | chol. reducer | \$202.85 | 0.327 | 0.42 | 0.21 |
| Procardia XL 60mg (100 tabs) | C-C blocker | \$207.42 | 0.357 | 0.38 | 0.21 |

Table 10: Estimates using alternative error distributions for Zip B

| | $\epsilon \sim U[-k, k]$ | $\epsilon \sim N(0, \sigma^2)$ |
|------------------------------------|--------------------------|--------------------------------|
| Price (α) | 0.287 (0.029) | 0.216 (0.039) |
| Pct. Chronic (β_0) | 0.953 (0.154) | 1.309 (0.132) |
| Constant (δ_0) | 7.916 (0.626) | 6.944 (0.498) |
| Pct. Female (δ_1) | -3.960 (0.495) | -4.491 (0.477) |
| Prob[search] (cash prescriptions): | | |
| Min | 0.00009 | 0.0003 |
| Median | 0.048 | 0.078 |
| Max | 0.269 | 0.449 |
| Search costs: | | |
| Average search | 3.91 | 3.68 |
| Average no search | 19.11 | 19.32 |
| Observations | 34,005 | 34,005 |
| Avg. Log-likelihood | -2.095 | -2.087 |

Maximum Simulated Likelihood estimates using data for both cash and insured transactions. Standard errors (in parentheses) were computed using the robust covariance matrix estimator recommended by McFadden and Train (2000) to account for simulation noise. Parameters of the preference term distributions were chosen to match the variance of the Type I Extreme Value distribution ($k = 2.222$ and $\sigma = 1.283$).

Table 11: Partial Equilibrium Effects of Search Cost Reductions

| Zip A | | | | |
|---------------------------|-------|--------|--------|--------|
| Search cost reduction: | 0% | 20% | 50% | 100% |
| Avg. search probability | 0.115 | 0.185 | 0.419 | 1.000 |
| Avg. equivalent variation | — | \$0.38 | \$1.49 | \$2.69 |

| Zip B | | | | |
|---------------------------|-------|--------|--------|--------|
| Search cost reduction: | 0% | 20% | 50% | 100% |
| Avg. search probability | 0.093 | 0.143 | 0.308 | 1.000 |
| Avg. equivalent variation | — | \$0.15 | \$0.63 | \$1.65 |

Results from model simulations using the parameters from Table 8 and simulated samples of 500 consumers for each prescription. Averages are weighted by transaction volumes, and utility measures are gross utility (i.e., not including the costs of search). The average equivalent variation is the average reduction in transaction price that would achieve the same gross utility as the corresponding search cost reduction.

Figure 1: Market shares and price differences: two examples

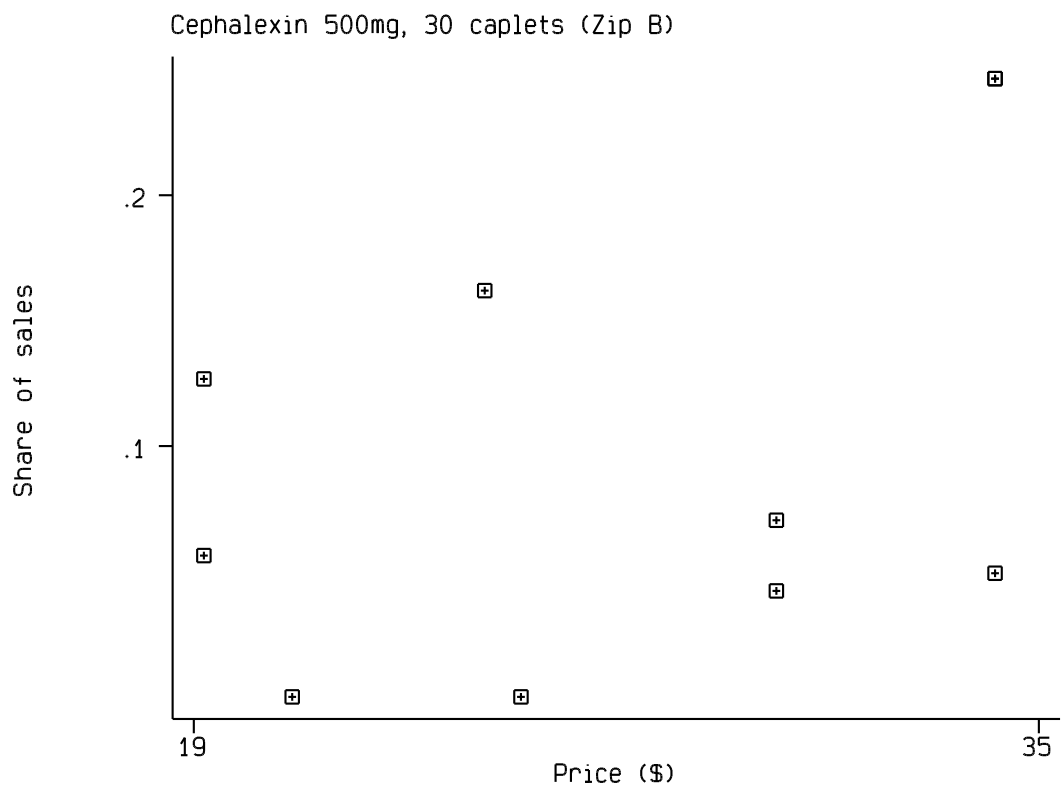




Figure 2: Exponential distribution and $(p - p_{min})$



The kernel density estimate is of the distribution of $p - p_{min}$, where these deviations are normalized by the corresponding prescription mean ($\hat{\lambda}$) to make observations comparable across prescriptions. A quartic kernel was used, and the boundary correction suggested by Rice (1984) was employed to account for the nonnegative support.