AN EMPIRICAL INVESTIGATION OF ONLINE COMPETITIVE PRICE PROMOTIONS*

Xiaoxun (Cathy) Gao+

JOB MARKET PAPER

First Draft: August, 2008
Current Version: November, 2008

Abstract

Using daily price data collected from a leading price comparison website, this paper empirically examines the role of vertical store differentiation on firms' equilibrium pricing strategies in a clearinghouse model. By allowing for different quality levels across firms the theoretical model used for estimation generates asymmetric pricing strategies. Our estimates show that observable store effects explain about 30 percent of the variation in prices for four digital cameras, a finding that is supported by hedonic regression results. Furthermore, we find that prices adjusted for quality differences between stores have greater intra-distribution mobility than unadjusted prices: stores move up and down the adjusted price distribution more frequently so that consumers cannot tell the identity of the lowest-priced firm over time. Using duration and Markov transition analysis we find that firms keep their prices constant for on average three weeks. We structurally estimate the model using nonlinear least squares and show nonlinear least squares gives consistent parameter estimates. Our general model explains the real world price data reasonably well. Factors that generate price dispersion are similar across the four digital cameras: the maximal profit margin ranges from 16 to 22 percent and the estimated proportion of consumers who visit the price comparison site is 16 to 26 percent. The cost of updating the prices is about two percent of the maximal profit margin. Finally, our estimates suggest that store managers believe to be competing with four to five stores.

Keywords: clearinghouse models, product differentiation, price dispersion, structural estimation

JEL Classification: C14, D43, D83, L13

---

* I am very grateful to Chenguang Li for sharing his data with me. In addition I would like to thank my dissertation committee members: Michael Baye, Michael Rauh, Roy Gardner, Matthijs Wildenbeest and David Jacho-Chávez for their guidance, encouragement, helpful comments and suggestions. I also benefited from discussions with Lan Zhang and Chenguang Li, the industrial organization class on structural estimation taught by Matthijs Wildenbeest and the computational econometrics class that introduces me to nonparametric estimation taught by David Jacho-Chávez. Any errors that remain are my responsibility alone.

+ Indiana University, Department of Economics, E-mail: gaox@indiana.edu.

1. Introduction

With the rising popularity of online shopping, price comparison sites have become more popular as well. Websites like PriceGrabber.com, BizRate.com, and Shopping.com post prices of millions of products across different online retailers, and make it relatively easy for consumers to find the lowest price for the product of interest. A quick look at the prices posted at these sites shows us that even if we are looking at a homogenous good, there is a lot of variation in prices across stores. Moreover, for a given store and product prices seem to fluctuate over time.

In the theoretical literature several explanations for these observed price differences have emerged. In the consumer search literature (e.g., Stahl, 1989; Burdett and Judd, 1983; Reinganum, 1979) it is typically assumed that each single price observations is costly to obtain. Heterogeneity in search costs will result in some consumers searching once, while others will search more often. In the literature on clearinghouse models (e.g., Varian, 1980; Baye and Morgan, 2001) the distinction between consumers is more stark: some consumers have access to a clearinghouse and observe all prices, while others do not have access and observe the price of just a single firm. Both consumer search and clearinghouse models have in common that firms typically try to balance the tradeoff between catering to consumers that do compare prices and consumers that do not. In order to maximize surplus from the first group a low price should be set, while maximizing surplus from the consumers that do not compare prices requires a relatively high price. As a result prices will be dispersed.

Since price quotations from price comparison sites are relatively easy to obtain they have become a major data source for empirical studies on price dispersion. Using price data from price comparison sites, empirical researchers have documented persistent price dispersion in a variety of industries such as electronics, books, and CDs (see e.g. Brynjolfsson and Smith, 2000; Pan, Ratchford and Shankar, 2002; and Baye, Morgan, and Scholten, 2004a, 2004b). Most of these papers look at whether observed pricing patterns are in line with implications derived from theoretical models of price dispersion.

So far very few studies have used the structure theoretical model on price dispersion directly to explain price dispersion. An exception is Hong and Shum (2006), which uses the equilibrium conditions of a consumer search model similar to Burdett and Judd (1983) to structurally estimate search frictions using price data alone. Moraga-González and Wildenbeest (2008) extend their method to an oligopoly market and show how to estimate the model using maximum likelihood. Both papers focus on consumer search models; clearinghouse models have received far less attention in the structural literature on price dispersion. In this paper we fill this gap by proposing a structural estimation method for a clearinghouse model that encompasses the clearinghouse models of both Baye and Morgan (2001) and Varian (1980).

Baye and Morgan (2001) is the first paper to explicitly model online price comparison sites. It extends Varian (1980) by endogenizing the optimal fee that is charged at the information clearinghouse. In a companion paper (Gao, 2008) we propose a more general model that also allows for store differentiation. Moreover, the model in that paper allows for a click-trough fee instead of a fixed fee and is therefore more suitable for empirical application. The goal of the current paper is to structurally estimate the theoretical generalized clearinghouse model presented in Gao (2008) using prices obtained from a price comparison site. Besides investigating how well the model fits real world data, we explicitly investigate the impact of structural parameters like the number of shoppers, unit cost, and willingness to pay on price dispersion in online markets.

Villas-Boas (1995) is the first paper to directly estimate the model of Varian (1980) using weekly prices of coffee and saltine cracker in grocery stores. Villas-Boas finds that Varian's model cannot be rejected for only 34 percent of the coffee brands and 50 percent of the saltine crackers brands. He argues that the strong symmetric conditions the theoretical model imposes on the data together with uncertainty about how a period should be defined are responsible for the bad fit. In addition, the assumption that it is costless for the firms to change prices is not very realistic. Our theoretical model deals with this by allowing for both zero and positive updating costs for firms.[1] Furthermore, by studying the implications of the model using duration analysis and Markov transition

[1] Varian (1980) is a special case of our theoretical model since there it is assumed the updating cost is zero.

matrices, we find that the interval in which firms keep their prices constant is roughly three weeks.[2] Based on these results we keep a three-week period between subsequent store prices selected for estimation. For the prices selected using this rule the hypothesis of no serial correlation cannot be rejected.

Our theoretical model shows that pricing strategies of firms in a seemingly homogenous product market can be asymmetric because of differences in the quality of service provided. The result suggests that simple aggregation of prices across stores is not justified by the model. Wildenbeest (2008) models vertical store differentiation in a consumer search model and proposes a method to distinguish between price dispersion due to store differences and price dispersion due to search frictions. Using his method to correct prices for store differentiation we find that intra-distributional movement is more frequent in adjusted prices than in posted prices. Only around 40 percent of adjusted prices for a given store are still in the same quartile after a three-week period compared to around 60 percent of the advertised prices. Over 20 percent of the adjusted prices move up to two quartiles away while less than 10 percent of the posted prices switch their relative positions within a three-week period. This means it is unlikely that consumers who compare adjusted prices find a pattern of firms persistently setting high or low adjusted prices, which is consistent with our theoretical model.

A paper that also uses price data from price comparison sites to structurally estimate a clearinghouse model is Yang (2008). The key insight of his paper is that not all firms that advertise on price comparison sites are actively competing. He proposes a two-step efficient GMM method to estimate the market parameters that give rise to price dispersion among competing firms. A key difference between that paper and ours is that while we correct for vertical product differentiation, Yang (2008) makes no distinction between competing firms and aggregates all prices into a single observed price distribution. Similar to Yang (2008), we minimize the squared distances between the empirical distribution and the theoretical distribution. However, we make two

[2] Prices are collected over some period of market interaction while the aforementioned theoretical models are one-shot games. It is common practice in the empirical literature to assume firms play a stationary repeated game of finite horizon in order to increase the sample size.

improvements upon his method. First of all, we use a continuous nonparametric data distribution instead of the empirical data distribution. We show using Monte Carlo simulations that a continuous nonparametric data distribution achieves a lower mean squared error value. Secondly, we use bootstrap methods to obtain standard errors of the coefficient estimates. Since our method is a two-step procedure, standard errors from the second step regression are well known to be smaller than the true ones. The bootstrap method takes care of the variability of the estimated parameters in the second step estimation.

The structure of the paper is as follows. In the next section, we review the theoretical model of Gao (2008) assuming both zero and positive updating costs. Next, we examine the implications of this for empirical estimation. In Section 3, after describing the data set, we discuss a couple of assumptions implicitly made throughout our analysis and summarize features of our data that are consistent with previous empirical studies. Section 4 presents evidence of vertical store differentiation. We conclude that the interval in which firms compete is about three weeks for the products of interest. In Section 5 we first present results for a Monte Carlo study using simulated data. Next, we estimate the parameters that explain price dispersion. Finally, Section 6 concludes.

2. The Model

In this section, we summarize the results of our generalized theoretical model that is pertinent to estimation. The focus is on the strategic decisions of firms in the three-stage game of consumers, firms and the gatekeeper. Based on the derived equilibrium strategies, we identify the factors that induce price dispersion using the price data from a leading price comparison site.

Let us explain the setup of our model. Suppose there are N geographically separated towns in a homogenous product market and each town is served by a local store. We normalized a continuum of consumers to be one, each of whom demands a unit of the good. Although the product is homogenous, the service provided by each firm is different. Hence, the reservation price of consumers differs depending on the service quality of each firm.

At the same time, we allow for different marginal cost of providing the service and the product for each firm. However, we impose the restriction that the difference between the reservation price and the marginal cost is the same for every firm, which is called the maximal profit margin (*X*).[3]

There are two types of consumers: shoppers (*S*) and loyal consumers (*L*). Shoppers compare not only prices but also store service qualities at the price comparison site. Nowadays, seller ratings are displayed along with product prices at most of the price comparison sites, making it easy to evaluate both characteristics simultaneously as a bundle. It is different from Baye and Morgan (2001) or Varian (1980) where shoppers compare prices alone. By writing their model in terms of utilities, we obtain asymmetric pricing strategies of firms because of differences in service qualities. Loyal consumers buy from their local firms provided prices are below their reservation price. They are evenly distributed across N firms and hence each firm has $(1-S)/N$ of them.

Firms make decisions on whether or not to advertise via the price comparison site and on prices they advertise. To maximize profit, firms charge their reservation prices (adjusted for service quality differences) to their local loyal consumers when they do not advertise. To win over shoppers, firms compete to offer a lowest adjusted price or a highest utility at the price comparison site.

The gatekeeper has a monopoly control over the information platform. Consumers have free access to the posted prices. Firms pay a click-through fee for each click from the price comparison site that leads a consumer to the website of the firm.[4]

In addition to the advertising fee, we assume a non-negative updating cost k for firms. It can be understood as the opportunity cost of a firm in preparing data feed files conformable to the price comparison site. Alternatively, it can be the equivalent monetary payment of firms to some third party business (e.g. LinkShare and Commission Junction)

[3] This restriction is most applicable to the short run and firms are endowed with their service quality levels.
[4] Click-through fee is shown to generate greater profit in Gao (2008) and is consistent with fee practices of comparison sites.

in managing their sales directed from the price comparison site and in ensuring the accuracy of the click volume the firm pays to the site. By investigating both zero and positive updating costs, we obtain different equilibrium strategies for the firms that guide our empirical analysis.

Before delving into the discussion of the firms' strategies, we explain the optimal behaviors of consumers under our setting. Intuitively, shoppers first visit the comparison site and buy from the firm that offers the best deal (either a lowest adjusted price or a highest utility). If no firm advertises, shoppers buy from their local firms and so do the loyal consumers provided prices are below their reservation prices.

We present the summaries of firms' optimal responses of our generalized theoretical model given the optimal click-through fee charged by the gatekeeper.

2.1 When Updating Cost (*k*) Is Zero

All of the *N* firms advertise on the comparison site when updating cost is zero.[5] They offer utility (equivalent to the negative of adjusted price) at the comparison site according to the following *H(m)* distribution:

$$H(m) = \left[1-\left(\frac{Lm+\frac{S}{N}X}{S(\frac{X}{N}-m)}\right)^{\frac{1}{N-1}}\right] \text{ where } m \in \left[0, \frac{S}{S+L}\frac{X}{N}\right]. \qquad (1)$$

Varian (1980) is shown to be a special case when there are no service quality differences.

2.2 When Update Cost (*k*) Is Positive

Not all of the *N* firms advertise at the price comparison site. Firms advertise with probability:

[5] The optimal click-through fee is *(1-1/N)X* when *k=0*.

$$\alpha = 1 - \left[\frac{k}{S[(1-\frac{1}{N})X - t]} \right]^{\frac{1}{N-1}} . \qquad (2)$$

Firms advertises utility m according to the following *G(m)* distribution:

$$G(m) = 1 - \frac{1}{\alpha}\left(1 - \left[\frac{k \dfrac{X-t}{(1-\frac{1}{N})X - t} + Lm}{S(X - m - t)} \right]^{\frac{1}{N-1}} \right)$$

$$where\ m \in \left[0, \frac{X-t}{S+L}\left(S - \frac{k}{(1-1/N)X - t} \right) \right]. \qquad (3)$$

Baye and Morgan (2001) is shown to a special case when there are no service quality differences in addition to zero click-through fee.

2.3 Implications for Empirical Estimation

First of all, the equilibrium strategies in advertising provide us a way to tell apart whether firms update prices at zero or positive cost. Under zero updating cost, all firms advertise while under positive updating cost, firms advertise with some probability α which is not one. Hence, by examining the summary descriptive statistics, we know which equation to use in the aggregated utility distribution.

Second, our theoretical model shows that consumers maximize utilities instead of minimizing prices. Firms price asymmetrically because of the quality differences but they draw utilities from a common distribution. As a result, simple aggregation of posted prices across stores is not justified by the model.

We can recover the utilities from the observed prices using a store's average price as a proxy for its service quality (See Wildenbeest, 2008). Utilities a consumer receives from purchasing a product from a particular firm *j* is: $U_j = V_j - p_j$ where V_j is the reservation value in dollar terms consumers placed on firm *j*.[6] Rearranging the equation, we have $p_j = V_j - U_j$. Next, the reservation value can be parameterized in terms of the firms' characteristics. If what are observed at the comparison site about the firms covers the characteristics consumers are willing to pay, we can run an OLS regression of a firm's posted prices on the observable characteristics and use the negative residuals as utilities. The coefficient estimates give us the dollar value consumers place in each of the different characteristics. However, if the observable characteristics are only a subset of those features consumers value, the residuals do not reflect the strategic interaction of firms alone and hence are not suitable for estimation.

$$p_j = X_j \beta - U_j . \tag{4}$$

When there is potential interaction between some unobservable characteristics and the observable characteristics, we can use the average price of a firm's price series to proxy its service quality level, which is essentially the cross-sectional fixed effects regression. Although the recovered utilities are smaller than the true ones, they are ordinal in the consumers' evaluation of product and store bundles and hence do not affect consumers' purchasing decision. We rescale the lowest recovered utility to be zero to be consistent with the theoretical distribution. Furthermore, we can adjust the time effect that impacts all N firms in the price data by the two-way fixed effects regression. The recovered residuals correct the time trend of and the store effects on observed prices.

$p_{jt} = V_j + V_t - U_{jt}$ where p_j proxies V_j, p_t proxies V_t and p_j is calculated by taking the average of $p_t - p_{jt}$ for every firm *j*. *(5)*

Last but not least, we deal with the definition of a period from the implications derived from the theoretical model. The common utility distribution we use in estimation is valid for aggregating adjusted prices across firms at a time period. To take advantage of the collected panel data, we assume firms play a stationary repeated game of finite horizon. The implication of the assumption is that adjusted prices of a firm over every period and adjusted prices across firms at a time period are independent and identically distributed from the theoretical distribution. As a result, we expect adjusted prices of a

[6] Recall that reservation price differs for each firm because of their service quality differences in the model.

firm across different periods move up and down the distribution at random. If we sort the adjusted prices of a representative firm into four quartiles, they have each 25 percent probability (or in terms of relative frequency) to the four quartiles the next period regardless of their current quartiles.

3. Data and Assumptions

We describe the data collection and summarize the key features of our dataset that are consistent with the previous empirical studies on the implications derived in line with our theoretical model. In particular, we discuss the short-run timeframe assumption and present some evidence in support.

3.1 Data

The data is collected for digital cameras from a leading price comparison site: PriceGrabber.com. During the data collection period, there are fifty different digital cameras on sale from 132 retailers. Data is collected daily from October 2007 to May 2008 for a total of 27 weeks. On a typical screen display of the search request of a particular model, we have for each retailer its web page link, its advertised price, the seller's service rating, whether the retailer is featured, tax and shipping fee if any, whether the product is in stock, an icon indicating whether the online retailer is hacker safe and if the product is on promotion.[7,8] Seller ratings are supplied by consumers who have bought products from the retailers. The scale is from 1 star (the worst) to 5 stars (the best) with the increment of ½ star. Retailers that advertise under the digital camera category pay $0.75 for every click that leads a consumer to the store website.[9] They prepare updating data files to be uploaded to the price comparison site.[10]

[7] Retailers make additional payments to be featured but the amount is not specified on the comparison website. There are at most four featured retailers which have the first four positions of the initial screen display at the search request of a product. With a mouse click, the order can be easily sorted from the lowest price to the highest.
[8] One example of promotion is to hand out free memory card with the purchase of the camera.
[9] There are two types of retailers on PriceGrabber.com: online merchants which have their own websites and price comparison site storefronts which operate only under the comparison site. Storefronts are typically small-volume individual sellers. They pay commission which is a percentage of their sales value to the comparison site. See the FAQ page of PriceGrabber.com. We delete those storefronts from our analysis because they are sell to far fewer consumers than the online retailers.
[10] A detailed five-step instruction webpage is in the FAQ of PriceGrabber.com.

3.2 The Short-run Timeframe and Assumptions

We assume our price data is obtained from a relative stable market environment. As mentioned above in the model section, the constraint that all firms have the same maximal profit margin (the difference between consumer reservation and the cost of the product and service bundle) is most applicable in the short-run timeframe. In addition, the model imposes the equal share of loyal consumers for each firm. They are important in identifying the factors of the theoretical distribution. We observe that price dispersion is significant and stable over the data collection period for each of the five measures.[11] It is in line with our assumption of a relative stable market environment.

In terms of the varying number of retailers at the price comparison site of our data collection, we assume the primary reason is that firms advertise with positive probability as predicated by the theoretical model (see Table 1 below). On a typical day, we observe an average of 17 firms and the number varies from only two firms to a maximal thirty-six firms. For a typical firm, we observe it advertises 65 out of 189 days.[12] Based on the result, we can rule out the all-firms-advertise equilibrium as discussed under the zero updating cost case. We restrict our attention to the positive updating cost case. Employing positive (but not equal to one) probability in advertising reduces the price competition at the price comparison site and firms earn more profit than the all-firms-advertise equilibrium.

Table1: Price Change Frequency of a Typical Day and of a Typical Firm[13]

	Mean (Std. Dev.)	Min, Max
Number of Firms that Change Price in a Day	1.27 (0.75)	0.07, 3.01
Number of Firms in a Day	17.20 (8.61)	1.97, 35.76
Percent of Firms that Change Price in a Day	0.07 (0.0188)	0.02, 0.13
Number of Days a Firm Changes Price	4.77 (1.87)	0.62, 8.19
Number of Days a Firm Advertises Price	65.26 (20.97)	12.65, 97
Percent of Days a Firm Changes Price	0.07 (0.0146)	0.03, 0.09

[11] See Table I and Figure II of the appendix.

[12] We notice that firms keep their prices constant for some time. When they update prices, the change does not occur synchronously. About seven percent of the firms change prices in one day and a firm changes price five times on average when they advertise (See Table 1).

[13] We have data collected over 189 days for fifty digital cameras. We calculate the six measures of interest for each product and summarize the average, standard deviation, minimum and maximum of fifty markets.

Let us check the competitive entry theory and the product life cycle theory that may explain the variation of the number of retailers. According to the competitive entry theory, as more and more firms enter the market, price competition will drive down the price level. We expect an inverse relationship between the number of firms and product prices. On the other hand, product life cycle theory states that later in the product life cycle, consumers' reservation price will fall. Together with free entry condition and convex average cost curve, fewer firms stay in the market. In addition, the lowest adjusted price a firm may use to sell to shoppers also falls.[14] So, we expect the number of firms is positively related to product prices. The correlation of the number of firms and the average price[15] and that of the number of firms and the minimal price in all products are calculated for the fifty markets. We have mixed results: about half of the markets reject the null that positive relationship exists and about half of the markets reject the negative relationship at five percent. Neither theory explains our data overwhelmingly well. A main reason may be that data is collected for less than 7 months, much less relative to the 18 months in Baye and Morgan (2004b). They use the first 8 months as the base period in an OLS regression and test if the later periods dummies are positively significant for the product life cycle theory.

The second assumption is the symmetric condition of updating cost (k), share of loyal consumers, maximal profit margin (or the difference between reservation price and cost no matter how different the reservation price and cost differ) for all of the N firms. Furthermore, a demand or supply shock at times if any affects all stores equally. The assumption is necessary because we do not observe sales volume of individual firms and hence cannot distinguish how some exogenous event may impact each firm differently. Under this assumption the residuals we obtain from fixed effects are to compare and derive properties in line with the theoretical model.

3.3 Evidence of Mixed Pricing Strategy

[14] The decline in the lowest adjusted price is less than the fall of the reservation price. The main reason of the fall is that the equilibrium number of firms is lower in the market and the pressure of price competition is relaxed. The reduction is smaller because the opportunity cost of giving up the loyal consumers is lower.
[15] We use both daily average price and daily minimum price in measuring the correlation with the number of firms observed daily.

We present evidence that firms play mixed pricing strategies in order to balance the tradeoff between catering to shoppers and loyal consumers. Many empirical papers have documented that prices or adjusted prices are not constant at a high level or a low level (see Lach, 2002; Baye and Morgan, 2004; Lewis, 2008; Wildenbeest, 2008). The pricing patterns are commonly predicted in models of consumers with asymmetric information of product prices which includes both clearing house model and consumer search model.

In order to gauge the relative position of the advertised firms on the observed price distribution, we divide the daily prices of each market into four quartiles (i.e. 25%, 50% and 75%) and count the frequency of a firm's prices on the different quartiles. As noted before, firms usually keep their prices constant for some days and the price adjustment does not occur at the same time, we present two approaches: the first method counts every price quotes of a firm while the second method counts price quotes of a firm only when the price is changed. For the first method, we assume firms compete for shoppers during all of their advertised days. The second method is more restrictive in that a firm considers seriously the competition between firms when it adjusts price.

Results are presented in Figure 2. Both approaches show most of the 132 retailers set prices on three to four price quartiles, meaning they do not persistently set a high or low price. The length of time a firm spends in each price quartile is not equal. Some firms noticeably have low prices more often than high prices. According to the theoretical model, vertical store differentiation leads to asymmetric prices. However, we can correct for the service quality differences and use adjusted prices in our estimation.[16] In addition, we observe a significant difference in the total frequency counts of retailers. It is mainly due to the number of digital cameras a firm advertise at the price comparison site and the number of days they advertise at each market.[17] However, the difference is less stark for individual markets.

[16] We compare the percentage of low-priced, medium-priced and high-priced firms (defined as the adjacent price quartiles accounts for over 80 percent of duration). The result shows that using adjusted prices we find less than 10 percent of stores can be identified as low, medium or high-priced firms as compared to over 50 percent of stores using posted prices. Summary tables are available upon request.

[17] We also observe about 20 firms stay in the same quartile and those are firms that advertise for an average of 13 days in one market and do not change prices as compared to a representative firm in our data that advertises 122 days and changes 7 times in each market.

Figure 2: Evidence of Mixed Pricing of 132 Firms

Part I: Frequency Count of Firms on Different Price Quartiles When Prices Are Advertised

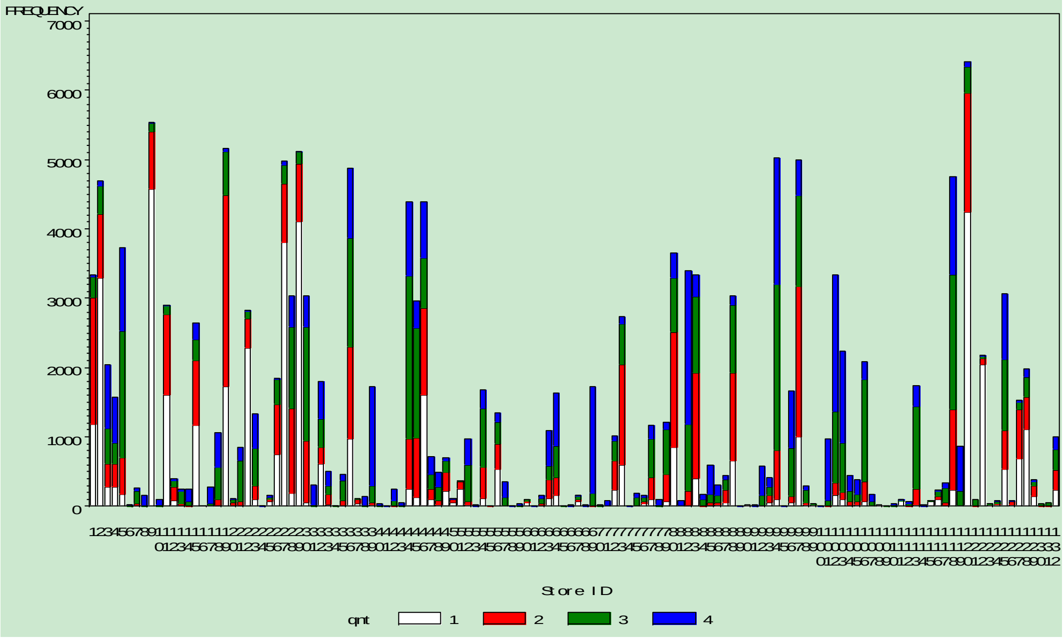

Part II: Frequency Count of Firms on Different Price Quartiles When Prices Are Adjusted

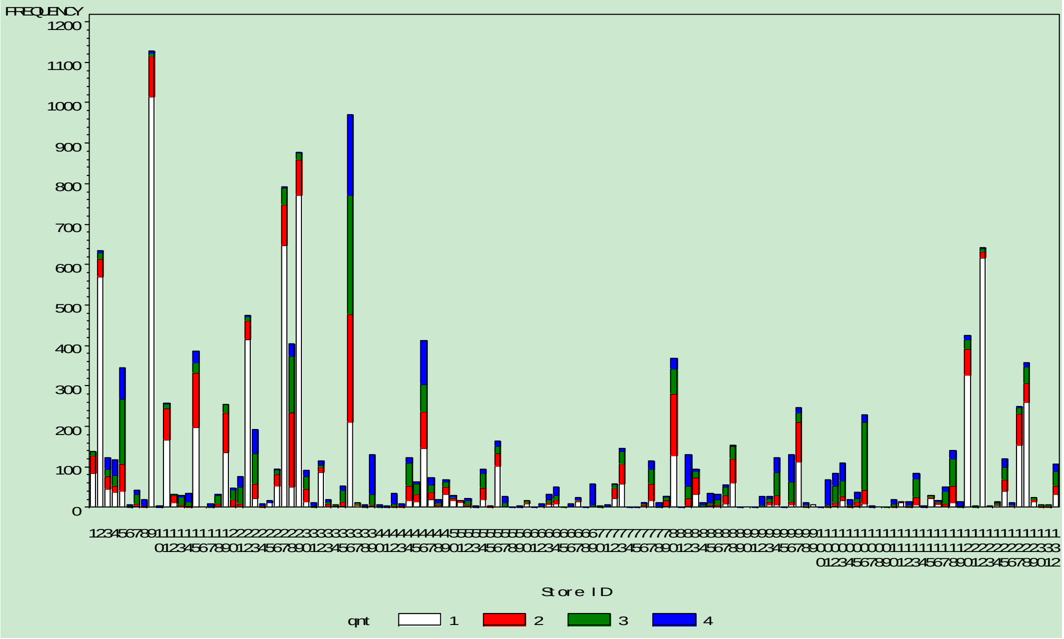

4. Vertical Store Differentiation

In this section, we examine the role of vertical store differentiation on firms' equilibrium pricing strategies. For presentation purpose, we pick four popular brands of digital cameras from different price ranges.[18] Descriptive statistics are presented in Table III of the appendix. We first look at the observable store effects as the sole characteristics that are responsible for differences of consumers' reservation value. Then, we compare the seller ratings of two price comparison sites and study if the observable store effects remain unchanged over the data collection period. After pointing out the limitation of using observed store features, we apply fixed effects regression to take away the unobservable store effects and common time trend in advertised prices. Using adjusted prices, we analyze the length of time a representative firm stays in each price quartile and the transition pattern of a firm in different quartiles. Based on the implication of the theoretical model, we derive the period that firms keep constant of their prices.

4.1 Observable Store Effects

Assuming consumers who compare prices at the price comparison site possess no additional information of the advertised retailers, the observable features are the only information consumers have to differentiate the retailers. We run a hedonic regression of daily prices on the observable features. It is an exercise to see how much variation of the product prices is explained by those observed store characteristics and if service quality is positively correlated as we have built in our theoretical model.

We notice in the data that featured stores with high service ratings price much lower and non-featured stores with high service ratings price much higher than the average store. However the effect of being featured or not is not that strong in stores with average or low service ratings. An interaction term of service rating and the dummy whether a retailer is featured or not is added in to observed features of the explanatory variables. Results are summarized in the Table 3 below.

The observable features have fairly good explanatory power on posted prices. The adjusted R-square ranges from the 9 percent to 46 percent. Prices of the product with the lowest average price (Canon PowerShot A560) seem to have a much greater variation

[18] Factors that induce price dispersion may have different effects depending on the product price level. Venkatesan, Mehta and Bapna (2007) find that market characteristics interact with product characteristics using price data from a leading comparison site: BizRate.com. Retailers with higher service ratings charge higher prices and the service premium increases as the product price level rises.

than what can be explained by the observed store effects. In addition, the lack of variability of the posted prices for days and the constant store characteristics cast some doubt on the variation explained that averages about 30 percent.

Nevertheless, seller ratings of the four digital cameras are positive and significant, suggesting the existence of store service premium. Consistent with the result of Venkatesan, Mehta and Bapna (2007), service premium is positively related to the average price of the product. If a retailer improves his rating by one star, he can raise price on Canon PowerShot A560 by about $4 or 3 percent of the base price which assumes all the control observable features are zeros. The service premiums of a one-star improvement for the other three digital cameras are 11 percent, 3 percent and 7 percent respectively. As we have suspected before the regression analysis, the interaction term of seller rating and being featured is negative and significant across four products. It implies that a high-rating seller offers a deeper discount to price-comparing consumers than a low-rating seller. In terms of seller ratings, a hypothetical seller with zero star charges roughly the same price whether or not he is being featured. However, if the same hypothetical seller is five-star, he charges $19 more when he is not featured than his zero-star counterpart with exactly the same of the other observable features. When being featured, the five-star seller now charges $13 dollars more. In other words, the five-star seller charges $6 less for being featured while the price difference is essentially zero for the zero-star seller. When we look at the coefficient of being featured or not, we can see that if a featured seller increases its rating by one star, holding all other factors constant, he charges about $3 more for Canon PowerShot A560, about $7 more for Canon PowerShot SD800, about $1 more for Sony CyberShot DSC-T200 and about $12 more for Pentax K10D SLR digital cameras. Variable promotion does not have consistent signs, which may reflect the different nature of promotion either being a price reduction or a value addition (e.g. free memory cards or camera bags). Hacker-safe is negatively significant which raises concerns of the validity of the regression because we expect this potentially beneficial feature shall raise consumers’ willingness to pay for the product knowing their online payment is secure. However, the negative sign may reflect omitted variable bias or different expectations of consumers of the retailers’ security features (or simply retailers are not able to make money from this feature). Similarly, shipping fee is not all negative but three out of four products are. Finally, tax is positive for all four products, which may reflect consumer trust on the sellers who charge taxes. A check of the data reveals that those firms are well-known companies like Dell, Apple and Best Buy.

Table 3: OLS of Price on Firm Observable Features with Robust Standard Errors

	Canon PowerShot A560	Canon PowerShot SD800	Sony CyberShot DSC-T200	Pentax K10D SLR
Intercept	123.78*** (1.68)	206.56*** (8.39)	331.73*** (3.32)	664.15*** (16.76)
Seller Rating	3.78*** (0.37)	22.52*** (1.84)	10.50*** (0.73)	46.32*** (4.05)
Seller Rating * Featured	-1.13* (0.65)	-15.41*** (1.76)	-9.79*** (1.48)	-34.23*** (3.45)
Featured	-0.03 (2.85)	28.65*** (7.85)	17.61*** (6.41)	98.74*** (15.08)
Promotion	-4.13*** (1.38)	13.36*** (2.78)	10.07*** (0.70)	7.39 (4.85)
Hacker-safe	-2.12*** (0.35)	-15.15*** (1.43)	-4.78*** (0.88)	-51.95*** (3.51)
Shipping Fee	-0.08* (0.0431)	1.11*** (0.13)	-0.26*** (0.0987)	-4.91*** (0.34)
Tax	1.11*** (0.19)	1.63*** (0.27)	0.71*** (0.0532)	0.52** (0.2177)
N	3321	1804	2056	2120
Adjusted R^2	0.09	0.46	0.34	0.29

Note: Standard errors are in parenthesis. * is 10%, ** 5% and *** 1% significance level.

Although seller ratings are positive and significant across the four products, the reliability may be in question due to concerns that firms hire people to write positive reviews. Furthermore, that some of beneficial features do not have the expected signs may be an indication that the observable features may not be the sole characteristics that consumers value. In order to check the validity of the store service ratings, we collect seller ratings from another major comparison site: BizRate.com which claims to place emphasis on store quality ratings in June 2008.[19,20] We find service ratings of 61 sellers

[19] Our price data from PriceGrabber.com are collected from October 2007 to May 2008.
[20] Empirical papers on online store service ratings such as Pan, Ratchford and Shankar (2002), Clay et al. (2002) and Venkatesan, Mehta and Bapna (2007) collects data from BizRate.com

that match those in our price dataset. The rating system is in terms of the 15 criteria and each item has 0 to 10 points. Consumers are invited to the survey after filling out an order at the store web site. Sellers do not necessarily advertise at BizRate.com but they must agree to participate the service rating program. Table 4 Part I summarizes the 15 criteria and gives a brief explanation of their meanings. Since those 15 items are highly correlated, we apply principal component analysis to extract the unique factors. Results are in Table 4 Part II.

Applying the commonly-accepted rule of keeping the factors with eigenvalue greater than one, we retain one factor which explains 86 percent of the total variation of those 15 items. Then we use correlation matrix to see how close the service rating from BizRate.com and from PriceGrabber.com. We find the correlation is about 51 percent and significant at 1 percent level. The result implies that seller ratings of PriceGrabber.com (which we use in the hedonic regression) contain useful information to the consumers. However, they seem not cover all aspects of consumers' valuation of the sellers.

Table 4: Service Ratings from BizRate.com

Part I: Brief Explanation of the 15-criteria Seller Ratings of BizRate.com

	Explanation		Explanation		Explanation
r1	would shop again	r6	prices relative to other online merchants	r11	availability of product you wanted
r2	overall rating	r7	overall look and design of the site	r12	order tracking
r3	ease of finding what you are looking for	r8	shipping charges	r13	on-time delivery
r4	selection of products	r9	variety of shipping options	r14	product met expectation
r5	clarity of product information	r10	charges stated clearly before submission	r15	customer support

Part II: Principal Component Analysis of 15-point Service Quality (61 sellers):

Fact-or	Eigen-value	Proportion Explained	Fact-or	Eigen-value	Proportion Explained	Fact-or	Eigen-value	Proportion Explained
1	12.92	0.86	6	0.14	0.01	11	0.02	0.00
2	0.77	0.05	7	0.10	0.01	12	0.01	0.00
3	0.38	0.03	8	0.07	0.01	13	0.01	0.00
4	0.33	0.02	9	0.04	0.00	14	0.01	0.00
5	0.16	0.01	10	0.03	0.00	15	0.01	0.00

We also notice that most of the observable store features remain the same throughout the data collection period. See Table 5 below for the percentage of sellers that do not change each of the observed characteristics in the four digital cameras. At least 60 percent of sellers do not change one of the observed store features. The percentage of sellers that have the same service ratings and promotion policy is as high as at least 75 percent. We take a look at the sellers that do change their service ratings: the magnitude is small. The standard deviation of stores with different service ratings averages at 0.21, 0.21, 0.26 and 0.25 respectively in the four products.

Table 5: Percentage of Firms That Do Not Change Their Observable Features[21]

	Canon PowerShot A560	Canon PowerShot SD800	Sony CyberShot DSC-T200	Pentax K10D SLR
Seller Rating	75%	83%	80%	76%
Featured	75%	75%	60%	71%
Promotion	97%	100%	89%	87%
Hacker-safe	81%	69%	77%	73%

In summary, we find some support of store service premium and about 30 percent of the price variation is explained by the hedonic regression. However, by comparing store service ratings from another price comparison site, we discover that the two measures tap into different aspects of store service characteristics that consumers are willing to pay for. Potentially we may have omitted store features that are not captured by the observed features which are not accounted for in our theoretical model. In addition, we show most of the firms keep the same observable features in our data, in support of our short-run assumption. Furthermore, the result gives us some validity to use fixed effects regression to recover the adjusted prices which a firm use in strategic interaction with other firms.

4.2 Unobservable Store Effects

We run fixed-effects regressions and summarize the F-test results in the Table 6 below. Store effects are highly significant and explain 69 to 88 percent of price variation. Adding day dummies in the regressions explains 79 to 93 percent of the total price

[21] Shipping fee depends on the delivery option specified by the retailers. It may be calculated as a percentage of the product price, which is similar as tax. Those two features are not considered in the above table because strategic price setting behavior may change shipping fee and tax paid.

variation. Again, the *p*-values of both time effects and store effects are extremely small. However, there is still unexplained variation in prices that cannot be accounted for.

Table 6: Store Fixed Effects and Store & Time Fixed Effects Regression

	Canon PowerShot A560	Canon PowerShot SD800	Sony CyberShot DSC-T200	Pentax K10D SLR
Store Fixed Effects R^2	0.87	0.88	0.69	0.85
F-test of Store Fixed Effects	0.00	0.00	0.00	0.00
Store and Time Fixed Effects R^2	0.88	0.93	0.79	0.88
F-test of Time Fixed Effects	0.00	0.00	0.00	0.00
F-test of Store and Time Fixed Effects	0.00	0.00	0.00	0.00

We use the residuals from the two-way fixed effects regression to study the strategic interactions of the adjusted prices that are consistent with the theoretical model implication. Following Lach (2002), we divide the daily adjusted price distribution into four quartiles (i.e. 25%, 50% and 75%) in order to investigate how those adjusted prices move inside the distribution. According to our theoretical model, adjusted prices move up and down the distribution at random so that consumers are unlikely to recognize the pricing pattern as opposed to Baye and Morgan (2001) who predict the posted prices are at random. Hence, we use both adjusted prices and advertised prices in the following analysis to compare the relative empirical performance of the two models.

The first question we are interested in is the duration of the adjusted price in one price quartile. The duration spell is counted as the number of days a seller charges the advertised price staying in one price quartile. The spell can be ended by the seller switches from one price quartile to the other price quartile or the seller takes off his advertisement at the price comparison site. As a result, a firm may have several spells. We aggregate the spells of all advertised firms and summarize the result in Table 7. It is clear that firms do not compete on a daily basis: no more than 10 percent of the spells last one day. A general impression we have from comparing the duration table of adjusted prices and advertised prices is that adjusted prices have a shorter duration. For Canon PowerShot A560, about 24 percent of posted price spells stay in the same quartile over one month (greater than 4 weeks) while only about 17 percent of adjusted price spells do so. For the other three products, the percentage of posted prices over one month is 33, 22 and 28 percent respectively. The reduction of duration in adjusted prices over the same period is greater at 13, 5 and 10 percent respectively. Taking the average of the duration days, 80 to 95 percent of adjusted price spells switch their quartiles in three weeks while 65 to 80 percent of raw price spells do so.

Table 7: Number of Days in Each Price Quartiles and Adjusted Price Quartiles

Note: the first four are of raw price quartiles and the last four are adjusted price quartiles

Part I. Duration Spell of Canon PowerShot A560 in Percentage

	q1	q2	q3	q4		q1	q2	q3	q4
1 day	2.51	5.20	4.60	5.55		4.29	7.46	9.51	4.38
< 1 week	25.30	42.47	34.84	26.94		23.25	43.11	40.14	28.49
1 to 2 weeks	17.06	22.76	21.35	20.58		20.24	10.29	30.13	26.26
2 to 3 weeks	12.05	6.36	9.41	9.36		12.62	8.91	5.21	16.48
3 to 4 weeks	8.41	8.51	7.49	14.80		11.83	7.94	8.81	11.92
> 4 weeks	34.68	14.70	22.31	22.77		27.78	22.29	6.21	12.48
Mean (days)	21	14	21	21		18	18	10	14
Median (days)	18	8	10	14		15	7	8	12

Part II. Duration Spell of Canon PowerShot SD800 in Percentage

	q1	q2	q3	q4		q1	q2	q3	q4
1 day	2.48	5.23	4.63	2.97		2.18	9.48	6.13	5.00
< 1 week	9.77	38.57	38.15	22.25		39.56	48.98	40.52	22.29
1 to 2 weeks	13.95	25.39	20.86	13.98		20.60	32.05	33.64	15.63
2 to 3 weeks	6.05	23.26	11.94	8.26		14.05	3.61	13.94	14.38
3 to 4 weeks	0.00	0.00	8.38	10.59		3.68	5.87	0.00	15.42
> 4 weeks	67.75	7.56	16.04	41.95		19.92	0.00	5.76	27.29
Mean (days)	56	11	13	24		15	8	10	21
Median (days)	44	8	10	23		8	6	8	18

Part III. Duration Spell of Sony CyberShot DSC-T200 in Percentage

	q1	q2	q3	q4		q1	q2	q3	q4
1 day	3.13	7.38	6.87	3.97		2.55	6.15	6.05	3.62
< 1 week	18.32	42.30	30.43	30.06		19.94	37.07	40.20	35.99
1 to 2 weeks	18.32	24.10	14.17	13.36		24.05	38.96	27.29	35.44
2 to 3 weeks	16.62	21.80	15.29	11.48		24.05	9.31	15.52	15.73
3 to 4 weeks	6.39	4.43	19.35	4.80		19.94	3.47	10.95	3.98
> 4 weeks	37.22	0.00	13.88	36.33		9.48	5.05	0.00	5.24
Mean (days)	22	9	17	24		16	10	9	11
Median (days)	17	8	13	15		15	8	8	12

Part IV. Duration Spell of Pentax K10D SLR in Percentage

	q1	q2	q3	q4		q1	q2	q3	q4
1 day	1.43	3.77	4.10	2.75		2.66	4.60	4.57	3.09
< 1 week	27.86	33.28	36.90	20.71		45.72	37.16	42.35	39.28
1 to 2 weeks	8.98	21.08	31.33	6.47		26.85	24.39	18.49	24.19
2 to 3 weeks	16.41	12.65	9.81	5.99		12.38	18.26	24.06	2.74
3 to 4 weeks	2.86	3.46	3.66	31.88		3.01	2.94	4.37	18.01
> 4 weeks	42.45	25.75	14.20	31.20		9.38	12.64	6.16	12.69
Mean (days)	27	19	12	24		12	12	11	15
Median (days)	19	14	10	25		8	9	8	11

After learning firms move away from a price quartile after staying there for some days, we ask the question how firms switch to different quartiles: do they move to the adjacent price quartile or do they move to all quartiles with equal probability of 25 percent? Considering the period we use to answer the question may affect the transition probability, we use either one week or three weeks. The Markov transition matrix tracks down the relative price positions of a seller for an equally spaced pre-specified interval (i.e. one week or three weeks). We use the duration spells during which a seller continuously post prices at the price comparison site. Results are summarized in Table 8 and Table 9. Quartiles of the columns are of time *t* (that is, one week or three weeks) and quartiles of the rows are of time *t+1* (that is, the next week or the next three weeks away). Clearly the sum of the percentage on each row is 100.

We see that adjusted prices are more consistently with a symmetric mixed strategy in both the one-week or three-week transition matrices: they move more at random so that consumers are unlikely to guess the pricing pattern of a particular firm. Advertised prices exhibit a strong tendency to stay in the same or the adjacent price quartiles. About 70 to 75 percent of posted prices remain in the same price quartiles after one week. Less than 5 percent moves two quartiles away (i.e. from 1st to 3rd or 4th quartile and from 4th to 1st and 2nd quartile). If we look three weeks ahead, posted prices move more frequently and the change in price quartile is bigger. Price comparing consumers have a 55 to 65 percent to find the seller remains in the same price quartile. However, the switch of price quartiles is still small: less than 10 percent of the prices deviates two quartiles away. On the other hand, adjusted prices move more freely in the distribution. For one-week period, 60 to 65 percent of adjusted prices are in the same quartile and at least 10 percent are away two quartiles. In three weeks, only 35 to 45 percent of adjusted prices keep their price quartiles. Furthermore, we see greater movement inside the price distribution: over 20 percent of the adjusted prices move from top to bottom or from bottom to top. Sellers are hard to be identified as setting consistently high or low prices. Price promotion with unpredictable price cuts at times keeps price comparing consumers in the dark of firms' pricing strategies. As for posted prices, our explanation is that firms play a mixed strategy but they draw adjusted prices instead of posted prices from a common distribution. The posted prices are stagnant in relative price distribution because of the store service quality difference. From the duration and price position transition analysis, we find variation in the price quartiles of the advertised prices but the variation is too small that consumers are likely to tell a seller's prices next period.

Table 8: One-Week Transition Matrix

Note: the first four quartiles are prices and the last four quartiles are adjusted prices; # is the number of price or adjusted price quotes in the initial quartiles.

Part I. Canon PowerShot A560

#		q1	q2	q3	q4		#		q1	q2	q3	q4
1034	q1	0.79	0.18	0.02	0.01		1007	q1	0.70	0.14	0.05	0.11
908	q2	0.24	0.61	0.13	0.02		1031	q2	0.13	0.65	0.18	0.04
1112	q3	0.02	0.13	0.72	0.13		823	q3	0.08	0.21	0.56	0.15
695	q4	0.01	0.02	0.25	0.72		888	q4	0.14	0.09	0.11	0.66

Part II. Canon PowerShot SD800

#		q1	q2	q3	q4		#		q1	q2	q3	q4
583	q1	0.84	0.15	0.01	0.00		585	q1	0.70	0.20	0.06	0.04
402	q2	0.17	0.65	0.18	0.00		339	q2	0.26	0.50	0.21	0.03
422	q3	0.01	0.10	0.72	0.17		423	q3	0.09	0.16	0.57	0.18
356	q4	0.01	0.00	0.18	0.81		416	q4	0.12	0.05	0.15	0.69

Part III. Sony CyberShot DSC-T200

#		q1	q2	q3	q4		#		q1	q2	q3	q4
604	q1	0.77	0.18	0.03	0.02		613	q1	0.69	0.18	0.03	0.10
522	q2	0.13	0.61	0.22	0.04		528	q2	0.15	0.59	0.23	0.03
607	q3	0.04	0.12	0.67	0.17		523	q3	0.08	0.16	0.54	0.22
393	q4	0.02	0.04	0.17	0.77		462	q4	0.16	0.06	0.17	0.61

Part IV. Pentax K10D SLR

#		q1	q2	q3	q4		#		q1	q2	q3	q4
653	q1	0.82	0.15	0.02	0.01		662	q1	0.66	0.19	0.05	0.10
542	q2	0.13	0.72	0.14	0.00		636	q2	0.17	0.65	0.14	0.04
489	q3	0.04	0.14	0.71	0.11		416	q3	0.07	0.22	0.55	0.16
498	q4	0.02	0.00	0.14	0.84		468	q4	0.16	0.06	0.13	0.65

Table 9: Three-Week Transition Matrix

Note: the first four quartiles are prices and the last four quartiles are adjusted prices; # is the number of price or adjusted price quotes in the initial quartiles.

Part I. Canon PowerShot A560

#		q1	q2	q3	q4		#		q1	q2	q3	q4
837	q1	0.72	0.24	0.02	0.02		735	q1	0.49	0.19	0.09	0.23
716	q2	0.31	0.48	0.16	0.05		843	q2	0.15	0.53	0.24	0.08
923	q3	0.06	0.14	0.65	0.15		654	q3	0.14	0.26	0.38	0.22
499	q4	0.02	0.04	0.32	0.62		743	q4	0.29	0.19	0.15	0.37

Part II. Canon PowerShot SD800

#		q1	q2	q3	q4		#		q1	q2	q3	q4
517	q1	0.73	0.23	0.04	0.00		456	q1	0.46	0.30	0.12	0.12
334	q2	0.21	0.43	0.33	0.03		212	q2	0.23	0.34	0.28	0.15
258	q3	0.02	0.13	0.61	0.24		346	q3	0.24	0.19	0.33	0.24
233	q4	0.00	0.01	0.21	0.78		328	q4	0.28	0.11	0.25	0.36

Part III. Sony CyberShot DSC-T200

#		q1	q2	q3	q4		#		q1	q2	q3	q4
532	q1	0.67	0.25	0.05	0.03		504	q1	0.50	0.32	0.05	0.13
428	q2	0.17	0.41	0.35	0.07		417	q2	0.17	0.47	0.31	0.05
482	q3	0.08	0.18	0.51	0.23		413	q3	0.14	0.15	0.37	0.34
322	q4	0.01	0.10	0.20	0.70		430	q4	0.35	0.11	0.16	0.38

Part IV. Pentax K10D SLR

#		q1	q2	q3	q4		#		q1	q2	q3	q4
562	q1	0.72	0.20	0.06	0.02		540	q1	0.50	0.22	0.13	0.15
446	q2	0.22	0.61	0.16	0.01		491	q2	0.17	0.53	0.21	0.09
343	q3	0.06	0.19	0.60	0.15		317	q3	0.21	0.27	0.35	0.17
376	q4	0.02	0.01	0.20	0.77		379	q4	0.31	0.07	0.15	0.47

In sum, we have shown that most of the adjusted prices (which are essentially advertised prices adjusted for store effects and day effects) stay in one quartile for roughly three weeks. The transition pattern of one quartile to the other for adjusted prices is more in line with the theoretical model: three-week Markov transition matrices reveal that only 40 percent of adjusted prices remain in the same quartile and over 20 percent move two quartiles away. Hence we conclude that firms compete in an interval of three weeks when they keep their prices relatively constant (in the same price quartile). We recognize the fact that some firms may revise their prices in the period while others do not. A way to understand the difference is that firms which change prices more often pay more of the updating cost k while those which do not change prices incur less of the updating cost. The three-week interval may be a result that firms balance the updating cost and the benefit of reacting to the rival firms' prices. In addition, we understand the tendency to stay in the same price quartile is still much larger than the model prediction. Our theoretical model only accounts for factors that vertical differentiate stores and time shocks that influence equally all of the sellers due to the data limitation. However, there are many other possibilities in the real world data that cause heterogeneity that cannot be accounted for by the model. Nevertheless, we show that the empirical data is more consistent with our model in terms utility than Baye and Morgan (2001) in terms of posted prices.

5. Structural Estimation

In this section, we impose the theoretical distribution on the price data and estimate the factors that induce price dispersion. We observe in our data that not all firms advertise and a firm does not advertise all the time at the price comparison site, which implies updating cost is not zero according to our theoretical model. Hence, *equation (3)* (refer to section 2) is used as the theoretical distribution. Based on the result we derive that firms keep prices relative constant for three weeks, we select prices every three weeks in the estimation. As a result, the maximal possible observation for one firm is nine meaning it advertises every period. We aggregate those adjusted prices into a single distribution.

5.1 Model Implication Checks

The first implication of a positive updating cost is that firms have a positive probability to advertise, meaning firms appear on and off the price comparison site. We choose several firms at random and present their advertised prices every three weeks in Figure 10. Only Canon PowerShot A560 is shown below because it is representative of the other three products.[22] For the pricing history of each firm, a missing dot means the firm does not advertise at the comparison site. Overall, there seem not be a pattern of missing prices on the graph. In terms of posted prices, when some firms raise prices, the other firms lower prices or keep prices constant. Price changes appear to be at random. We also notice posted prices seem to be moving within some bounds, which may be an indication of service premium differences.

[22] Graphs of the other three products are available upon request.

Figure 10: Mixed Advertising and Pricing of Canon PowerShot A560

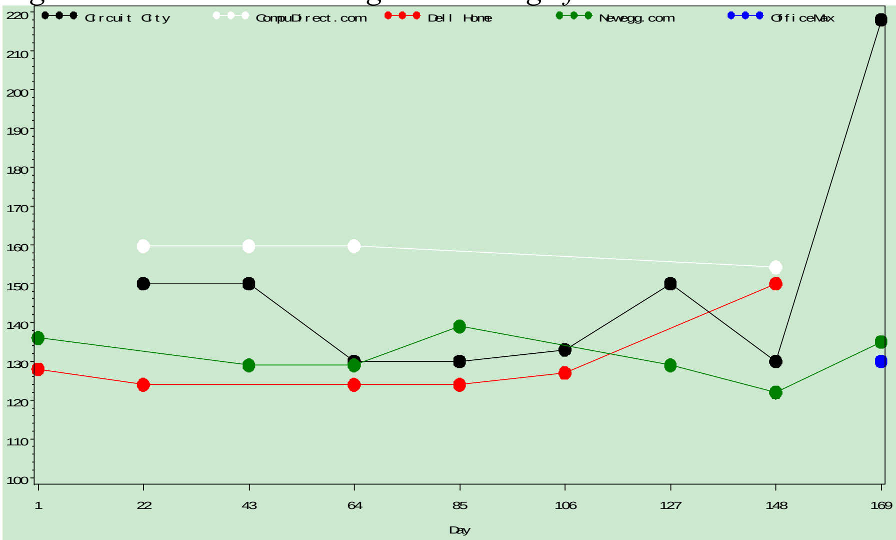

The second implication we check is the pricing history of a firm based on our selection every three weeks. Following Moraga-González and Wildenbeest (2008), we calculate the autocorrelation function of firms that advertise at least seven out of nine periods and use no less than two different prices. Results are in Table 11. We cannot reject the null that prices of the firm are not serial correlated at any significance level.

Table 11: Autocorrelation Function of Firms in Four Product Markets

ACF	Canon PowerShot A560	Canon PowerShot SD800	Sony CyberShot DSC-T200	Pentax K10D SLR
Mean (Std. Dev.)	0.51(0.15)	0.05(0.32)	0.35(0.21)	0.10(0.17)
Min.	0.31	-0.28	-0.05	-0.14
Max.	0.66	0.58	0.67	0.27
# of Firms	4	7	7	5

5.2 Monte Carlo Simulations

Our estimation strategy is to minimize $\sum_{i=1}^{N}\sum_{t=1}^{T}(G_{analytical}(m;\theta)-G_{data}(m))^2$ where the theoretical distribution is from *equation (3)*. The parameters to be estimated (θ) are updating cost (k), proportion of shoppers (S), maximal profit margin (X) and the number of competing firms (N). The theoretical distribution is of complicated functional form. We adopt a couple of strategies to simply and identify the parameters. First, we estimate the probability of advertising (α) and substitute in the distribution. Equivalent to the symmetric condition of the parameters we impose on the firms, we assume that all of the firms have equal probability of advertising which is the predicted α. Second, although the difference (X-t) can be estimated, the theoretical distribution does not allow us to separately identify the maximal profit margin (X) and the click-through fee (t). In our theoretical model, consumers who click through will buy from the retailer. In reality, it may not be the case. One way to get around the problem is to estimate calibrate the click-through fee and plug it in as a known value. Let us use the click-through fee at the time of data collection which is \$0.75 per click and a conversion rate of 30 percent (meaning 30 percent of the interested consumers actually buy from the retailer). Then, the calibrated click-through fee consistent with the model which assumes 100 percent conversion rate is \$0.75/0.30 or \$2.5. Alternatively, we may assume a zero click-through fee and the estimated X is in fact the maximal profit margin net of the click-through fee. The advantage is that it does not require the assumption of a conversion rate.

Through a Monte Carlo study, we show that the other estimated parameters seem to be affected little and the effect of click-through fee is absorbed into the redefined net maximal profit margin. We also compare the performance of using the non-continuous empirical distribution and using the continuous estimated nonparametric distribution as $G_{data}(m)$ in the simulations. Both distributions converge to the true distribution at a rate of root n, the number of observations. Simulation result shows that the continuous nonparametric CDF outperforms the non-continuous empirical CDF in obtaining a lowered value of the mean square errors. In addition, we check the validity of bootstrapped standard errors. Since we estimate the advertising probability and plug it in as the distribution as if it is observed, the standard errors from the minimization problem are not valid. In order to obtain the valid standard errors, we also need to take into account the variability of the estimated advertising probability. Bootstrap provides us a conceptually and computationally easy fix.

Simulation results are summarized in Table 12 below. Part I is to assume that the true click-through fee \$2 is known while in Part II we estimate the net maximal profit margin assuming a zero click-through fee. We can see that the effect of imposing a zero click-through fee on the estimation is small. The impact is mainly reflected in the maximal profit margin X, which seems to absorb the true click-through fee. Nevertheless, estimation in both Part I and Part II are consistent: the value of mean square error falls as the number of observations increase. As expected, the values are higher in Part II than its counterparts in Part I which is mainly due to the underestimated true maximal profit margin. In both parts, smooth nonparametric distribution as the theoretical distribution improves on the mean square value as compared with the empirical distribution with kinks. Furthermore, the standard deviations of the 1000 simulation estimates are very close to the average of the 1000 bootstrap standard errors. We obtain each of the 1000 bootstrap standard error from using 500 bootstrap samples in each set of simulated data. We can conclude that bootstrap standard errors are valid.

Table 12: Simulation Results of Empirical CDF and Nonparametric CDF Methods

Part I. When click-through fee is assumed to be known (t=2)

True Value	Empirical CDF			Smooth Nonparametric CDF		
	60 observations	120 observations	250 observations	60 observations	120 observations	250 observations
k=1	1.23 (1.36) (1.27)	1.12 (0.70) (0.83)	0.96 (0.40) (0.36)	1.43 (1.75) (1.52)	1.16 (0.67) (0.74)	0.99 (0.40) (0.32)
S=0.2	0.23 (0.0783) (0.0810)	0.23 (0.0410) (0.0458)	0.22 (0.0307) (0.0260)	0.24 (0.0739) (0.0614)	0.23 (0.0377) (0.0397)	0.22 (0.0297) (0.0227)
X=100	95.07 (9.65) (9.45)	95.61 (7.93) (8.59)	96.01 (7.81) (8.13)	97.13 (9.48) (9.86)	96.77 (7.88) (8.18)	98.37 (7.64) (7.28)
N=5	5.35 (1.17) (1.43)	5.22 (0.74) (0.95)	5.18 (0.51) (0.62)	5.20 (1.10) (0.96)	5.17 (0.68) (0.79)	5.15 (0.50) (0.45)
MSE	118.81	82.81	79.65	96.46	76.89	57.28

Part II. When click-through fee is not know and t=0 is imposed

	Empirical CDF			Smooth Nonparametric CDF		
	60 observations	120 observations	250 observations	60 observations	120 observations	250 observations
k=1	1.32 (1.36) (1.65)	1.16 (0.77) (0.62)	1.00 (0.41) (0.51)	1.45 (1.75) (1.62)	1.20 (0.71) (0.83)	1.03 (0.42) (0.39)
S=0.2	0.23 (0.0793) (0.0764)	0.23 (0.0415) (0.0348)	0.22 (0.0304) (0.0263)	0.23 (0.0719) (0.0627)	0.23 (0.0374) (0.0399)	0.22 (0.0275) (0.0271)

X=100	94.15 (9.96) (10.05)	93.57 (7.39) (8.27)	95.06 (8.27) (8.16)	95.08 (9.58) (9.67)	94.88 (8.39) (8.91)	95.75 (7.82) (8.18)
N=5	5.30 (1.17) (1.32)	5.20 (0.73) (0.76)	5.13 (0.50) (0.48)	5.18 (1.11) (0.96)	5.14 (0.69) (0.74)	5.09 (0.49) (0.51)
MSE	137.06	101.54	95.75	114.06	96.35	85.72

Note 1: standard deviations of 1000 estimates from stimulations are in the first parentheses and the mean of bootstrap standard deviations of estimates are in the second parentheses.
Note 2: in the estimation of smooth nonparametric CDF, the bandwidth is fixed at the optimal value. 1000 bandwidth values are calculated using least square cross validation method and the optimal value is the average of the 1000 bandwidths. The optimal bandwidths are nearly identical in Part I and Part II. We use 2.89, 1.68 and 1.17 for the three sample sizes respectively.

5.3 Estimation Results

We apply our nonlinear least square method to the adjusted price data to estimate the parameters that reflect the firms' strategic interactions in the theoretical model. Based on our result that adjusted prices move across different quartiles of the distribution every three weeks, we select adjusted prices at a three-week interval. Then we non-parametrically estimate the smooth distribution function using least squares cross validation to choose the appropriate bandwidth. By finding the coefficient estimates that minimize the squared distances between the estimated nonparametric distribution of the observed data and the theoretical distribution, we present our results in Table 13 below. The units of updating cost and maximal profit margin are in terms of dollars. To facilitate comparison across different product prices, we also calculate the updating cost as a percentage of the maximal profit margin and the maximal profit margin as a percentage of the daily average maximal price.

We find that factors that generate price dispersion are similar across the four digital cameras. During a three-week interval, firms incur roughly $2 in adjusting prices, which is about 2 percent of their net maximal profit margin. The estimated updating cost has a positive relationship with the average product prices. It may reflect the fact that firms are more careful hence paying greater attention in posting higher product prices. However, given the relatively big standard deviations, updating costs may not be very different between products. The measure of special interest to store managers is the proportion of shoppers which ranges from 16 to 26 percent in the four digital cameras. To understand the predicted proportion of loyal consumers, let us discuss the number of competing firms now. The estimates of the number of firms are in fact very small as compared to what we observed in a typical day at the price comparison site. One reason for the discrepancy may be that store managers actually believe that they are competing with just a few other online stores when they set product prices. If that is the case, the share of loyal consumers in the market is estimated at about 18 percent.[23] Price dispersion at the comparison exists according to the theoretical model is that a representative firm balances the tradeoff between selling to the price comparing consumers (about 21 percent) and the loyal consumers (about 18 percent). To those 21 percent of shoppers, a firm tries to set a low price to attract. To those 18 percent of loyal consumers, a firm reaps the highest possible profit by charging their reservation prices. The result of those two forces is the price dispersion we observe at the price comparison site. Alternatively the discrepancy may indicate the theoretical model is restrictive. A measure that is of interest to retailers and the gatekeeper is the net maximal profit margin, about 20 percent of the maximal price.

Table 13: NLS Parameter Estimates Using Nonparametric CDF

	Canon PowerShot A560	Canon PowerShot SD800	Sony CyberShot DSC-T200	Pentax K10D SLR
Updating Cost	1.30 (1.08)	2.36 (3.23)	2.32 (3.06)	3.00 (5.79)
Proportion of Shoppers	0.16 (0.0424)	0.21 (0.0534)	0.26 (0.0663)	0.21 (0.0529)
Net Maximal Profit Margin	72.17 (5.99)	88.91 (13.06)	81.38 (13.00)	205.03 (28.32)
Number of Competing Firms	3.59 (0.38)	5.15 (0.83)	3.70 (0.56)	4.71 (0.76)
Observations	214	100	115	126
Minimized Square Value	0.80	0.28	0.36	0.57
k / X	1.80%	2.65%	2.85%	1.46%
X / p_{max}	21.87%	21.17%	16.31%	19.11%

Note: standard deviations in parentheses are obtained using 500 bootstrap samples.

[23] It is calculated by *(1-S)/N* from the theoretical model. We use the average of proportion of shoppers in the four markets, 21 percent.

From Table 13, we see the minimized square value is small, implying a fine model fit. For visual interpretation, we also plot the smooth nonparametric distribution (the dotted lines are the 95 percent confidence interval of the estimated distribution line) and the theoretical distribution on the same graph for the four digital cameras in Figure 14 below.

Figure 14: Nonparametric CDF and Estimated Analytical CDF

Part I. Canon PowerShot A560

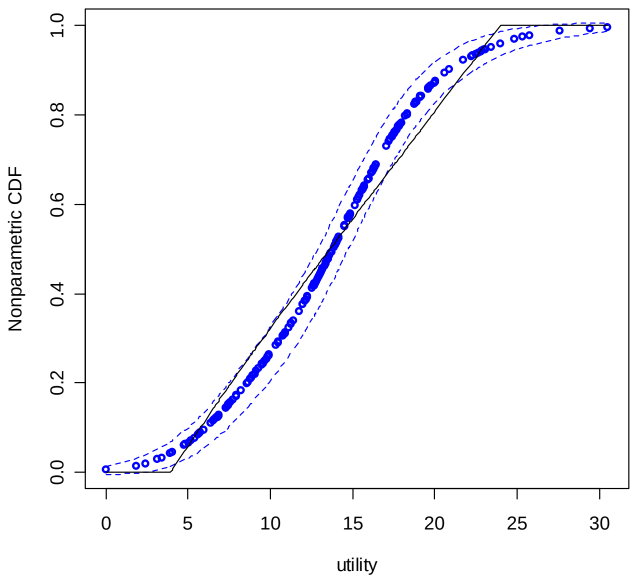

Part II. Canon PowerShot SD800

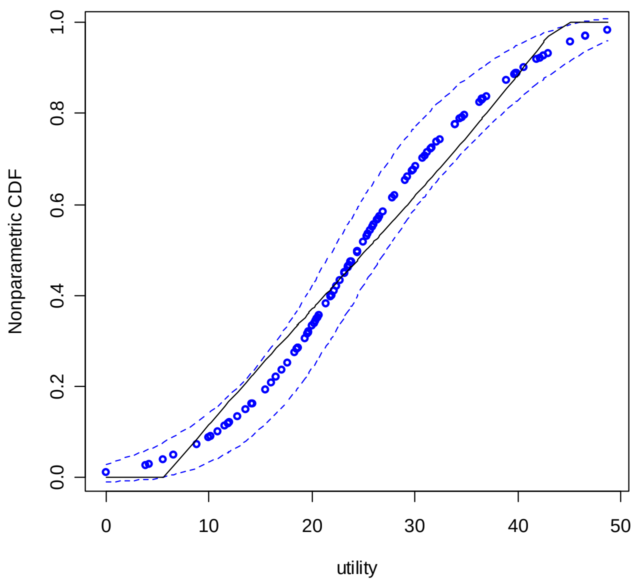

Part III. Sony CyberShot DSC-T200

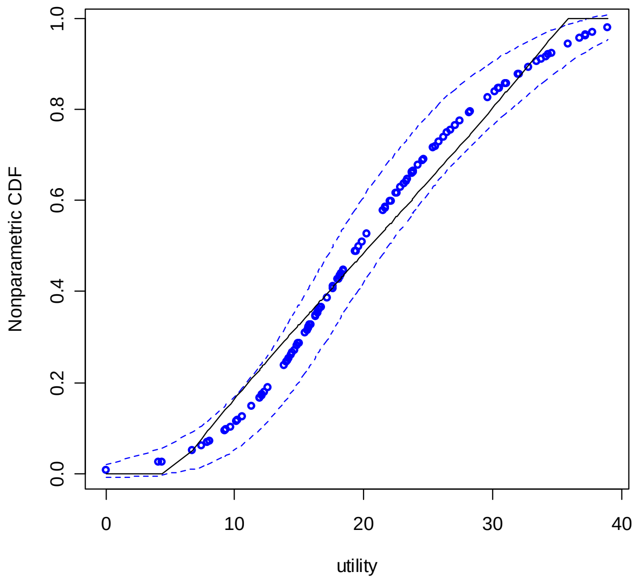

Part IV. Pentax K10D SLR

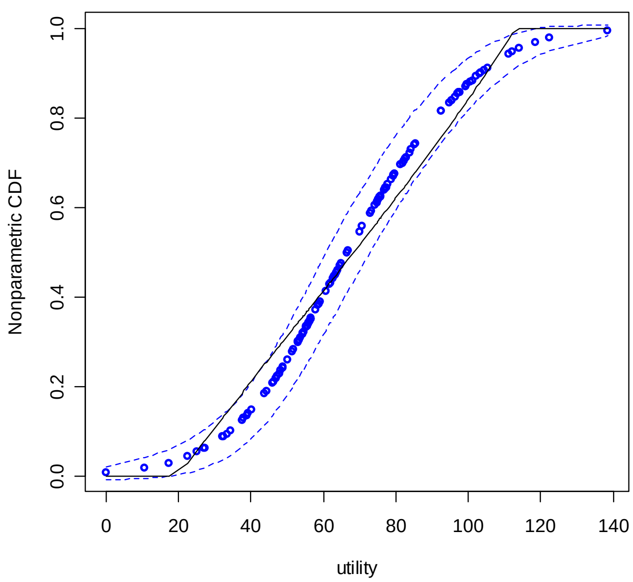

The theoretical distribution is contained inside the 95 percent confidence interval of the nonparametric distribution except for some high values of utility Canon PowerShot A560, which has the largest minimized square value in estimation. The kinks at zero and one occur because of the restriction imposed by the cumulative distribution function. Overall, the theoretical model fits the data reasonably well considering the few parameters we use in estimation and strong symmetric conditions of the model.

6. Conclusion

In this paper, we empirically estimate the theoretical model of our companion paper (see Gao, 2008) which encompasses Baye and Morgan (2001) and Varian (1980). Our estimates suggest that factors that generate price dispersion are similar across the four digital cameras. The maximal profit margin ranges from 16 to 22 percent and the estimated proportion of consumers who visit the price comparison site is 16 to 26 percent. When store managers believe to be competing with four to five stores, the estimated proportion of consumers who do not compare prices and buy from their local store is about 18 percent. The cost of adjusting prices at the comparison site is small across the four products. We also compare our model that allows for vertical differentiation and the general clearing house model that do not in terms of the consistent model implications with the real world data. Intra-distributional movement is shown to be more frequent and more thorough in adjusted prices which control for store and time effects than posted prices. Consumers are unlikely to predict the identity of the lowest-priced firm over time as a result. Our analysis shows that firms keep price relatively constant for three weeks.

Although our theoretical model which simplifies dramatically the observed real world price dispersion down to only four parameters explains the data reasonably well, we recognize the limitations of our assumption and some possible improvements. First, we can get rid of the symmetric condition in the model. Arnold, Li, Saliba and Zhang (2008) develop a model of asymmetric loyal consumers in a duopoly market. Second, measurement errors of firms or decision errors of consumers can be introduced into the model by assuming a distribution that generates the error term.

7. Appendix

Table I: Price Dispersion Measures in 50 Digital Camera Markets

Price Dispersion Measures	Mean (Std. Dev.)	Min, Max
Price Range ($)	168.69 (142.66)	23.09, 853.43
Value of Information ($)	58.11 (53.96)	9.18, 281.58
Price Gap ($)	14.02 (38.89)	1.13, 278.38
Coefficient of Variation	9.51 (3.16)	2.93, 16.47
Value of Information over Minimum Price	0.12 (0.0364)	0.03, 0.19

Figure II: Price Dispersion Measures

Figure A: Three Price Dispersion Measure in Dollars

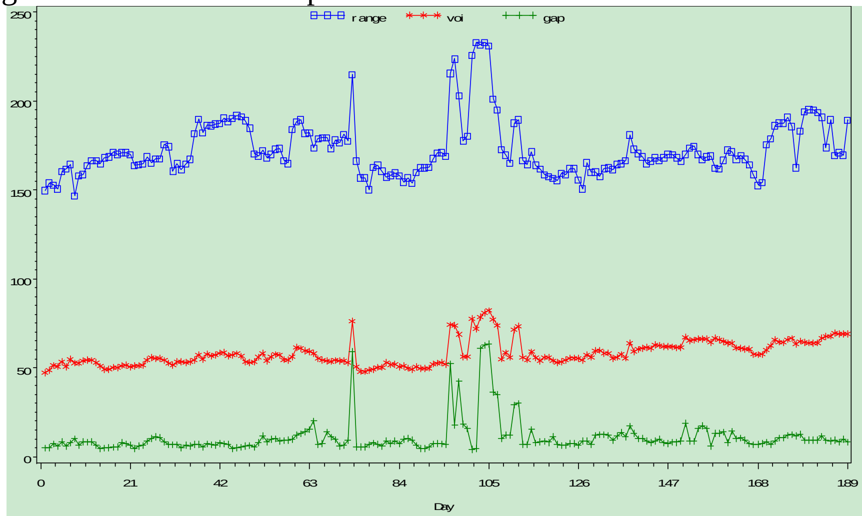

Figure B: Unit-free Price Dispersion Measures: Coefficient of Variation

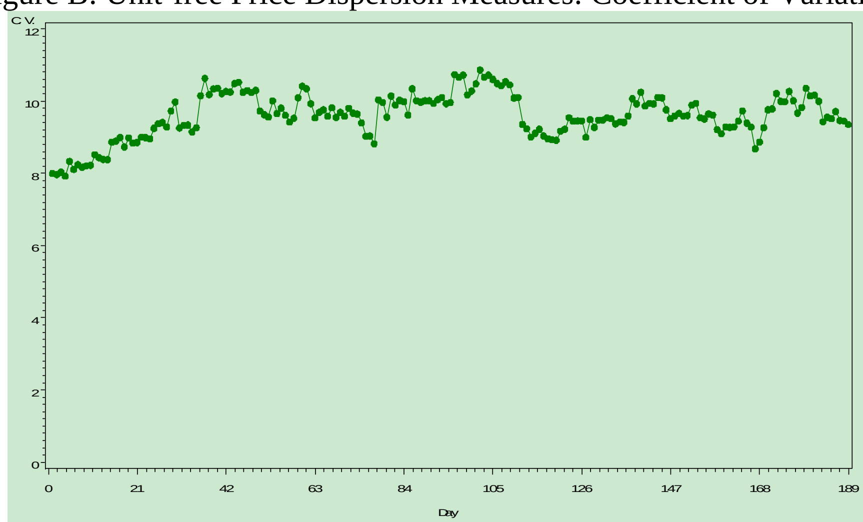

Figure C: Unit-free Price Dispersion Measures: VOI as % of Min. Price

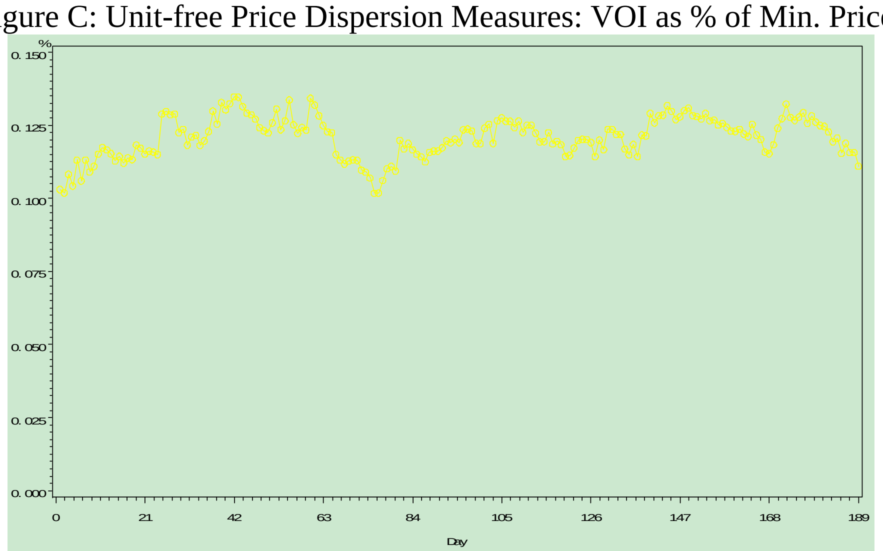

Table III: Descriptive Statistics

Part I: Canon PowerShot A560 Digital Camera

	Mean (Std. Dev.)	Min, Max	Number of Observations
Product Price	\$144.53 (\$26.21)	\$119.99, \$329.99	4595
Review Rating	4.31 (0.76)	2, 5	3629
Featured Store	0.15 (0.3604)	0, 1	4595
Promotion	0.01 (0.1006)	0, 1	4595
Hacker-safe Feature	0.36 (0.4800)	0, 1	4595
Shipping Fee	\$4.53 (\$4.87)	\$0, \$19.52	4206
Tax	\$1.22 (3.06)	\$0, \$19.80	4264
Number of Stores	24.31 (11.84)	3, 43	189 days

Part II: Canon PowerShot SD800 Digital Camera

	Mean (Std. Dev.)	Min, Max	Number of Observations
Product Price	\$295.00 (\$38.42)	\$229.00, \$419.97	2209
Review Rating	4.40 (0.71)	2, 5	1958
Featured Store	0.34 (0.4732)	0, 1	2209
Promotion	0.02 (0.1334)	0, 1	2209
Hacker-safe Feature	0.54 (0.4989)	0, 1	2209
Shipping Fee	\$3.91 (\$4.66)	\$0, \$14.49	2031
Tax	\$1.40 (4.94)	\$0, \$27.92	2038
Number of Stores	11.69 (6.29)	4, 25	189 days

Part III: Sony CyberShot DSC-T200 Digital Camera

	Mean (Std. Dev.)	Min, Max	Number of Observations
Product Price	\$370.85 (\$24.65)	\$301.00, \$499.00	2513
Review Rating	4.19 (0.81)	2, 5	2245
Featured Store	0.29 (0.4533)	0, 1	2513
Promotion	0.11 (0.3750)	0, 1	2513
Hacker-safe Feature	0.47 (0.4993)	0, 1	2513
Shipping Fee	\$1.92 (\$4.53)	\$0, \$34.20	2297
Tax	\$4.99 (\$9.50)	\$0, \$37.00	2306
Number of Stores	13.30 (4.51)	1, 20	189 days

Part IV. Pentax K10D SLR Digital Camera

	Mean (Std. Dev.)	Min, Max	Number of Observations
Product Price	\$827.96 (\$83.76)	\$649.95, \$1072.95	2745
Review Rating	4.39 (0.77)	2, 5	2295
Featured Store	0.26 (0.4388)	0, 1	2745
Promotion	0.08 (0.3340)	0, 1	2745
Hacker-safe Feature	0.36 (0.4790)	0, 1	2745
Shipping Fee	\$3.59 (\$5.58)	\$0, \$22.47	2509
Tax	\$4.98 (\$15.27)	\$0, \$81	2528
Number of Stores	14.61 (5.50)	4, 24	189 days

References

Arnold, Michael A., Chenguang Li, Christine Saliba and Lan Zhang (2008). “Asymmetric Market Shares, Advertising, and Pricing: Equilibrium with an Information Gatekeeper.” Working Paper.

Baye, Michael R. and John Morgan (2001). "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." *American Economic Review* 91, pp.454-474.

Baye, Michael R., John Morgan, and Patrick Scholten (2004a.). "Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site." *Jounal of Industrial Economics* 52, pp.463-496.

Baye, Michael R., John Morgan, and Patrick Scholten (2004b.). "Persistent Price Dispersion in Online Markets." In D. Jansen, ed., *The New Economy.* Chicago: University of Chicago Press.

Baye, Michael R., John Morgan, and Patrick Scholten (2004). "Temporal Price Dispersion: Evidence from an Online Consumer Electronics Market." *Journal of Interactive Marketing* 18, pp.101-115.

Baye, Michael R., John Morgan, and Patrick Scholten (2006). "Information, Search, and Price Dispersion." *Handbook of Economics and Information Systems.* T.Hendershott, ed., North Holland: Elsevier.

Baye, Michael R., Rupert J. Gatti, Paul Kattuman and John Morgan (2006). "Did the Euro Foster Online Price Competition? Evidence from an International Price Comparison Site." *Economic Inquiry* 44(2), pp. 265-279.

Baylis, Kathy and Jeffrey M. Perloff (2002). "Price Dispersion on the Internet: Good Firms and Bad Firms." *Review of Industrial Organization* 21(3), pp.305-324.

Burdett, Kenneth and Kenneth L. Judd (1983). "Equilibrium Price Dispersion." *Econometrica* 51(4), pp.955-969.

Clay, Karen, Ramayya Krishnan, Eric Wolff and Danny Fernandes (2002). “Retail Strategies on the Web: Price and Non-price Competition in the Online Book Industry.” *Journal of Industrial Economics* 50(3), pp.351-367.

Hong, Han and Matthew Shum (2006). "Using Price Distributions to Estimate Search Costs." *The Rand Journal of Economics* 37(2), pp. 257-275.

Iyer, Ganesh and Amit Pazgal (2003). “Internet Shopping Agents: Virtual Colocation and Competition.” *Marketing Sciences* 22(1), pp.85-106.

Lach, Saul (2002). “Existence and Persistence of Price Dispersion: An Empirical Analysis.” *Review of Economics and Statistics* 84(3), pp. 433-444.

Lewis, Matthew (2008). “Price Dispersion and Competition with Differentiated Sellers.” *Journal of Industrial Economics* forthcoming.

Moraga-González, José Luis and Matthijs Wildenbeest (2008). "Maximum Likelihood Estimation of Search Costs." *European Economic Review* 52, pp.820-848.

Pan, Xing, Brian Ratchford and Venkatesh Shankar (2002). "Can Price Dispersion in Online Markets be Explained by Differences in E-Tailer Service Quality?" *Journal of the Academy of Marketing Science* 30(4), pp. 433-445.

Pan, Xing, Brian Ratchford and Venkatesh Shankar (2004). "Price dispersion on the internet: A review and directions for future research." *Journal of Interactive Marketing* 18(4), pp.116-135.

Reinganum, Jennifer F. (1979). "A Simple Model of Equilibrium Price Dispersion." *Jounal of Political Economy* 87(4), pp.851-858.

Salop, Steven and Joseph E. Stiglitz (1977). "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion." *Review of Economic Studies* 44(3), pp.493-510.

Sorensen, Alan T. (2000). “Equilibrium Price Dispersion in Retail Markets for Prescription Drugs.” *Journal of Political Economy* 108(4), pp.833-850.

Stahl, Dale (1989). “Oligopolistic Pricing with Sequential Consumer Search.” *American Economic Review* 79(4), pp. 700-712.

Varian, Hal R. (1980). "A Model of Sales." *American Economic Review* 70(4), pp.651-659.

Varian, Hal R. (2000). “Market Structure in the Network Age,” in *Understanding the Digital Economy*, Erik. Brynjolfsson and Brian Kahin eds. Cambridge, MA: MIT Press.

Venkatesan, Rajkumar, Kumar Mehta and Ravi Bapna (2007). “Do Market Characteristics Impact the Relationship Between Retailer Characteristics and Online Prices?” *Journal of Retailing* 83(4), pp.309-324.

Wildenbeest, Matthijs R. (2008). "An Empirical Model of Search with Vertically Differentiated Products." Working Paper.

Yang, Guoning (2008). “Firm Heterogeneities, Click-through Fees and Pricing in Oligopoly: Theory and Evidence.” Job Market Paper.