THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume LII December 2004 No. 4

PRICE DISPERSION IN THE SMALL AND IN THE LARGE: EVIDENCE FROM AN INTERNET PRICE COMPARISON SITE*

MICHAEL R. BAYE†
JOHN MORGAN‡
PATRICK SCHOLTEN§

This paper examines four million daily price observations for more than 1,000 consumer electronics products on the price comparison site Shopper.com. We find little support for the notion that prices on the Internet are converging to the 'law of one price.' In addition, observed levels of price dispersion vary systematically with the number of firms listing prices. The difference between the two lowest prices (the 'gap') averages 23 per cent when two firms list prices, and falls to 3.5 per cent in markets where 17 firms list prices. These empirical results are an implication of a general 'clearinghouse' model of equilibrium price dispersion.

I. INTRODUCTION

A number of recent studies provide conflicting pictures of the competitiveness of Internet markets.[1] For example, Brynjolfsson and Smith [2000] find that E-commerce markets for books and CDs are far from frictionless, with price ranges of around 30 per cent. In contrast, Ellison and Ellison [2004] report dispersion of about 5 per cent for computer memory.

What accounts for the differences in the levels of price dispersion observed in different online markets?

One potential explanation is that price dispersion is a *disequilibrium* phenomenon that is being corrected over time. The Ellison and Ellison data

*We gratefully acknowledge the comments of the editor and two anonymous referees, and seminar participants at Indiana University, the University of California at Berkeley and at Bentley College.

† Author's affiliations: Department of Business Economics, Kelley School of Business, 1309 East Tenth Street, Indiana University, Bloomington, IN 47405-1701, U.S.A.
e-mail: mbaye@indiana.edu.

‡Haas School of Business, 545 Student Services Building #1900, University of California at Berkeley, Berkeley, CA 94720-1900, U.S.A.
e-mail: morgan@hass.berkeley.edu.

§Department of Economics, 175 Forest Street, Bentley College, Waltham, MA 02452-4705, U.S.A.
e-mail: Pscholten@bentley.edu.

[1] See Bakos [2000] and Smith, Bailey and Brynjolfsson [2000] for excellent surveys of this work.

was collected several years after that of Brynjolfsson and Smith, and the lower price dispersion might reflect the fact that prices have moved toward perfectly competitive equilibrium as consumers became more skillful in comparison shopping in online markets. This explanation is consistent with the view that the Internet will ultimately lead to a perfectly competitive equilibrium:

> 'The explosive growth of the Internet promises a new age of perfectly competitive markets. With perfect information about prices and products at their fingertips, consumers can quickly and easily find the best deals. In this brave new world, retailers' profit margins will be competed away, as they are all forced to price at cost.' *The Economist*, November 20, 1999, p. 112.

An alternative explanation is that price dispersion is an *equilibrium* phenomenon and that the differences in price dispersion in the two studies stem from differences in market structure. Competing sellers in the markets studied by Ellison and Ellison number in the hundreds whereas fewer than twenty sellers compete in the markets studied by Brynjolfsson and Smith, and the lower dispersion might stem from these differences in the number of sellers.

To examine these competing explanations, we assembled a dataset containing four million price observations in an online market for consumer electronics products. These data are daily price quotes from merchants selling the top 1,000 products covered by Shopper.com—a leading price comparison site on the Internet. The data span the time horizon from August 2, 2000, through March 31, 2001. The number of firms listing prices for these products varies a great deal—both cross sectionally and over time—thus permitting us to examine the impact of variations in the number of listing firms (and hence market structure) on price dispersion. Data from price comparison sites, such as the one analyzed in this paper, offer a unique opportunity to (1) quantify the role that the number of firms play in explaining differences in levels of dispersion for different products, and (2) differentiate among alternative theoretical models of price dispersion. To the best of our knowledge, there have been no empirical studies of price dispersion on the Internet that examine how price dispersion varies with market structure.

We find little evidence to support the view that price dispersion is a disequilibrium phenomenon that is being corrected over time. Instead, we find persistent price dispersion that depends on market structure. Specifically, despite the fact that consumer usage of price comparison sites increased by 12.9 per cent during the eight month period we study,[2] we find

[2] This figure is based on comparing price comparison site usage in the period 2000–01 to the period 2001–02 using figures provided in Brynjolfsson, Montgomery and Smith [2003].

no statistical evidence of any decline in levels of price dispersion. Further, we find systematic differences in price dispersion depending on the number of firms listing prices for a given product: The level of price dispersion is greater when small numbers of firms list prices than when large numbers do. For example, for products where only two firms list prices, the gap between their prices (which is also the range of prices) averages 23 per cent. In contrast, for products where 17 firms list prices (the average in our sample), the gap between the two lowest prices falls to about 3.5 per cent.

The remainder of the paper proceeds as follows: Section II discusses several equilibrium explanations of price dispersion, and offers a general clearinghouse model that matches many of the institutional characteristics of price comparison sites such as Shopper.com. Our model subsumes a number of existing models, including Baye and Morgan [2001], Narasimhan [1988], Rosenthal [1980], Shilony [1977] and Varian [1980], as special cases. We show that the general model predicts that price dispersion is an equilibrium phenomenon and that, for the leading special cases, price dispersion is predicted to be greater in the small than in the large. Section III summarizes our data and collection methodology and highlights differences between the Shopper.com site and competing services (such as shopbots) available on the web, as well as the strengths and limitations of our dataset. Empirical results are presented in Section IV, while Section V attempts to discriminate among special cases of the general model based on the findings in Section IV. Finally, an appendix contains formal proofs of various assertions made in the text.

II. THEORETICAL CONSIDERATIONS

A number of papers in the economics literature predict that price dispersion will persist in the Internet age. For example, Reinganum [1979], Burdett and Judd [1982], Gatti [2001], and many others show that equilibrium price dispersion can arise if there is a *positive* marginal cost of obtaining each price quote. This provides an appealing rationale for price dispersion documented in (1) conventional retail markets where consumers must incur the incremental costs of searching for prices at firms' brick-and-mortar stores,[3] and (2) those electronic markets where consumers incur the incremental costs of searching for prices at different firms' online stores.[4]

While these models are relevant when it is costly to obtain each and every price quote, price comparison sites such as Shopper.com, mySimon.com, and EvenBetter.com now make it possible for consumers to obtain a list of prices for a given product for what is close to a *zero* marginal cost of obtaining each price quote. A product search at any one of these sites will

[3] See, for instance, Pratt, *et al.* [1979], Carlson and Pescatrice [1980], and Sorensen [2000].
[4] See, for instance, Smith, Bailey and Brynjolfsson [2000] as well as Bakos [2000].

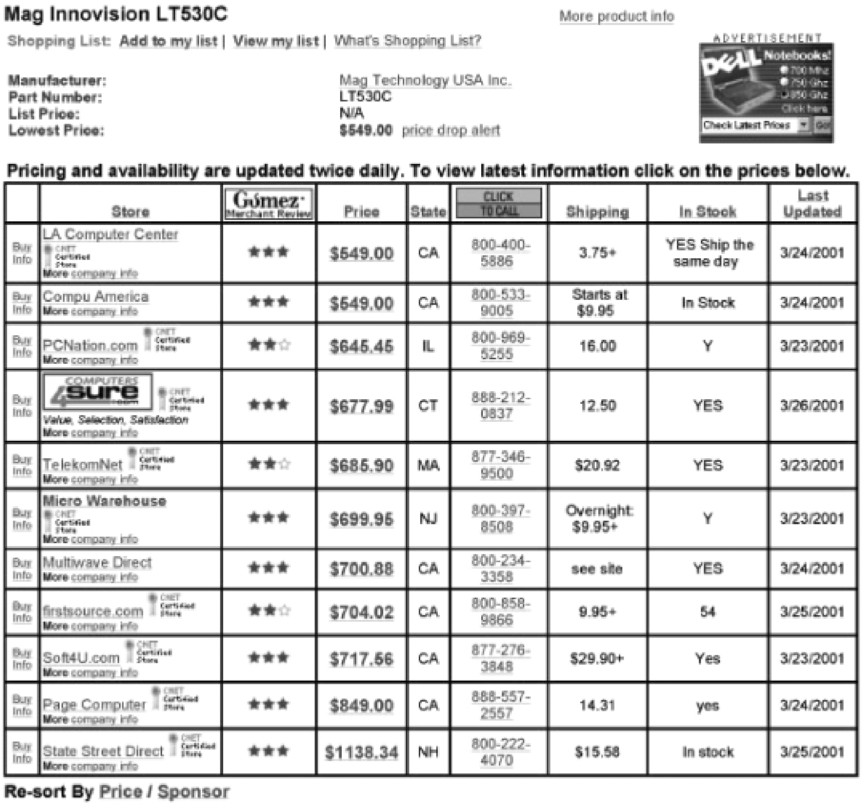

Mag Innovision LT530C

More product info

Shopping List: Add to my list | View my list | What's Shopping List?

ADVERTISEMENT

Manufacturer:	Mag Technology USA Inc.
Part Number:	LT530C
List Price:	N/A
Lowest Price:	$549.00 price drop alert

Pricing and availability are updated twice daily. To view latest information click on the prices below.

	Store	Gómez Merchant Review	Price	State	CLICK TO CALL	Shipping	In Stock	Last Updated
Buy Info	LA Computer Center CNET Certified Store More company info	★★★	$649.00	CA	800-400-5886	3.75+	YES Ship the same day	3/24/2001
Buy Info	Compu America More company info	★★★	$649.00	CA	800-533-9005	Starts at $9.95	In Stock	3/24/2001
Buy Info	PCNation.com CNET Certified Store More company info	★★☆	$645.45	IL	800-969-5255	16.00	Y	3/23/2001
Buy Info	COMPUTERS 4sure.com CNET Certified Store Value, Selection, Satisfaction More company info	★★★	$677.99	CT	888-212-0837	12.50	YES	3/26/2001
Buy Info	TelekomNet CNET Certified Store More company info	★★☆	$685.90	MA	877-346-9500	$20.92	YES	3/23/2001
Buy Info	Micro Warehouse CNET Certified Store More company info	★★★	$699.95	NJ	800-397-8508	Overnight $9.95+	Y	3/23/2001
Buy Info	Multiwave Direct More company info	★★★	$700.88	CA	800-234-3358	see site	YES	3/24/2001
Buy Info	firstsource.com CNET Certified Store More company info	★★☆	$704.02	CA	800-858-9866	9.95+	54	3/25/2001
Buy Info	Soft4U.com CNET Certified Store More company info	★★★	$717.56	CA	877-276-3848	$29.90+	Yes	3/23/2001
Buy Info	Page Computer CNET Certified Store More company info	★★★	$849.00	CA	888-557-2557	14.31	yes	3/24/2001
Buy Info	State Street Direct CNET Certified Store More company info	★★★	$1138.34	NH	800-222-4070	$15.58	In stock	3/25/2001

Re-sort By Price / Sponsor

Figure 1
Screenshot from Shopper.com

return a listing of prices that different merchants charge for the same product.[5] For example, consider a consumer who wants to purchase a Mag Innovision LT530C flat panel monitor using Shopper.com. One mouse click on March 26, 2001, brought up the list of prices displayed in Figure 1. Notice that these prices are dispersed, ranging from a low of $549 to a high of $1,138.34.

Clearinghouse models, which we discuss in detail below, more closely match the environment consumers encounter at price comparison sites. These models assume that information about prices is available through a clearinghouse, such as Shopper.com, and that some or all consumers access the list of prices to identify the 'best' price. As we will show, these models predict equilibrium price dispersion that varies with market structure.

[5] Products with identical manufacturer part numbers.

II(i). *Measuring Price Dispersion*

Before describing clearinghouse models in more detail, it is useful to discuss briefly some theoretical issues that arise when tracking price dispersion at sites such as Shopper.com. Traditionally, economists have used the coefficient of variation (cf. Sorensen [2000] and Carlson-Pescatrice [1980]) or the price range (*cf.* Brynjolfsson and Smith [2000]) to measure price dispersion in homogeneous product markets. When the law of one price holds, all firms in the market charge the same price and these measures of price dispersion are all zero. Thus, it would seem natural to examine the coefficient of variation or range in prices over time to examine whether price dispersion is a disequilibrium phenomenon that is being corrected over time. There is, however, a theoretical difficulty with this approach: The coefficient of variation and range can indicate significant price dispersion even when the underlying data are consistent with a competitive equilibrium.

Consider the list of prices displayed in Figure 1. One can hardly imagine a more dramatic departure from the law of one price. Based on the range, price dispersion is over 107 per cent of the lowest price; based on the coefficient of variation, it is 22.4 per cent. Yet one can argue that these data are also consistent with a situation where products are identical and all consumers purchase at one of the two firms charging the lowest price, which happens to be marginal cost. That is, these data are consistent with a competitive equilibrium in which no firm can gain by adjusting its price. To see this, suppose the 11 firms listing prices in Figure 1 are price-setting oligopolists and each has a marginal cost of \$549. Given this list of prices, price-sensitive consumers will naturally buy from one of the two firms offering the lowest price of \$549. While firms charging prices above \$549 do not have sales, they have no incentive to gain consumers by pricing at or below their costs of \$549. Likewise, since two firms are charging the lowest price in the market, neither can gain by unilaterally raising or lowering its price. Thus, the apparent price dispersion is arguably a fiction: The list of prices comprises an equilibrium in which all transactions take place at the perfectly competitive price (\$549).

We focus on a measure of price dispersion that alleviates this problem. Suppose the prices charged by $n \geqslant 2$ firms for a given product are ordered from lowest to highest, so that $p_1 \leqslant p_2 \leqslant \ldots \leqslant p_n$. We define 'the gap,' $G = p_2 - p_1$, to be the difference between the two lowest prices. Clearly, the classical Bertrand model implies that the gap between the two lowest prices is zero in any equilibrium (symmetric or otherwise). Thus, in any competitive equilibrium, price dispersion measured by G is zero (and therefore independent of the number of firms). The gap measure also has the empirically desirable property that it gives greater weight to low prices, which are presumably more likely to lead to sales than high prices. Thus, it is a proxy for a quantity-weighted measure of price dispersion if, as seems likely, consumers visiting the site tend to be price-sensitive.

II(ii). *Clearinghouse Models*

While $G=0$ in a competitive equilibrium, we show in this section that positive gaps between the two lowest prices always arise as equilibrium behavior in clearinghouse models. The distinguishing feature of clearinghouse models is that identical firms sell to two types of consumers: Those who buy at the lowest price listed at the clearinghouse, and those who do not. Consumers who do not buy at the lowest listed price may be loyal to a particular firm (as in Narasimhan, Rosenthal, or Shilony) or may be unwilling or unable to access the site (as in Baye-Morgan and Varian).[6] These models all predict that the list of prices obtained at the site will exhibit price dispersion despite quite different assumptions regarding the number of firms, product homogeneity, firms' decisions to list prices at the clearinghouse, consumers' decisions to utilize the clearinghouse, and the fees charged by the clearinghouse to consumers and firms using its services to acquire or transmit price information. Furthermore, these models predict that the level of price dispersion depends on the number of firms that list prices. In particular, all of these models predict that the expected difference between the lowest two prices is greater in the small than in the large.[7]

To establish this, we develop a general clearinghouse model which includes each of these models as a special case.

Suppose that there are $n>1$ firms with constant marginal cost $m\geqslant 0$ competing in a market by offering some identical product to consumers. This market is served by a price clearinghouse. Firms must decide what price to charge for the product and whether to list this price at the clearinghouse. Let p_i denote the price charged by firm i. It costs a firm an amount $\phi\geqslant 0$ if it chooses to list its price. All consumers have unit demand and a maximal willingness to pay of $r>m$[8]. Of these, $L\geqslant 0$ consumers per firm are price-insensitive 'loyal' consumers and will purchase from the firm to which they are loyal if its price does not exceed r. Otherwise, they do not buy the product at all.[9] A number, $S>0$, of the consumers are price sensitive 'shoppers.' These consumers first consult the clearinghouse and buy at the lowest price

[6] See also Salop and Stiglitz [1977], Spulber [1995], Stahl [1989], Stahl [2000] and Janssen and Moraga [2000]. These models also share the property that some fraction of consumers observe the complete list of prices offered by firms.

[7] The predictions of models differ in other dimensions, however, and we will explore some of these in Section V.

[8] It is straightforward to modify the model to allow for downward sloping demand.

[9] An alternative interpretation of 'loyal' consumers is as follows: 'Loyal' consumers are fully price sensitive, but do not access the clearinghouse site. Their search technology is such that the optimal strategy is to choose a single firm at random and buy from it if its price is at or below r. This interpretation is the one offered in Varian and Baye-Morgan.

listed there provided this price does not exceed r. If no prices are advertised at the clearinghouse or all listed prices exceed r, then a 'shopper' visits one of the firms at random and purchases if its price does not exceed r.

Several well-known clearinghouse models emerge as special cases of the general model. For instance, letting $M \geqslant 0$ denote a constant, the general model reduces to:

- The Baye-Morgan [2001] model when $\phi > 0$ and $L = \frac{M}{n}$;
- The Varian [1980] model when $\phi = 0$ and $L = \frac{M}{n} > 0$;
- The Narasimhan [1988] model when $\phi = 0$, $L = \frac{M}{n} > 0$, and $n = 2$;
- The Rosenthal [1980] and Shilony [1977] models when $\phi = 0$ and $L > 0$ is constant.

Proposition 1 establishes that if it is not too costly for firms to list prices at the clearinghouse, price dispersion *always* arises in the general clearinghouse model.

Proposition 1: *Suppose* $0 \leqslant \phi < \frac{n}{n-1}(r-m)S$. *Then in a symmetric equilibrium to the general clearinghouse model:*

1. *The expected gap between the two lowest listed prices is strictly positive.*
2. *The distribution of prices listed at the clearinghouse is:*

$$F(p) = \frac{1}{\alpha}\left(1 - \left(\frac{\frac{n}{n-1}\phi + (r-p)L}{(p-m)S}\right)^{\frac{1}{n-1}}\right) \text{ on } [p_0, r],$$

where

$$p_0 = \frac{\frac{n}{n-1}\phi + Lr + Sm}{L+S}$$

and

$$\alpha = 1 - \left(\frac{n\phi}{(n-1)(r-m)S}\right)^{\frac{1}{n-1}}.$$

Proof: The proof of part 1 follows from the fact that the distribution of listed prices is atomless and has non-degenerate support. The proof of part 2 follows as a consequence of a more general existence result (Theorem 1), which is stated and proved in the appendix.

Thus, under a wide variety of modeling approaches, the general clearinghouse model predicts equilibrium price dispersion, the level of which depends on market structure. Furthermore, under mild assumptions discussed in the appendix, the leading special cases of this model all share the following property relating price dispersion to market structure.

Proposition 2: *In the Baye-Morgan, Varian, Rosenthal, and Shilony models, price dispersion (as measured by the expected gap) is greater in the small (when 2 firms list prices) than in the large (when an arbitrarily large number of firms list prices).*

Proof: The proof is contained in the appendix.

III. DATA

We base our empirical analysis on four million daily price listings by different merchants selling the most popular 1,000 products at Shopper.com for the eight month period August 2, 2000 – March 31, 2001.[10] As noted above, Shopper.com is an 'information clearinghouse' that specializes in price comparisons for identical consumer electronics products sold by different firms. It touts the most comprehensive price catalog for these items on the Internet, with more than 100,000 products. Moreover, there is considerable firm participation on the site—at any given time, there are more than one million price quotes listed there. Shopper.com generates over 175,000 qualified leads per day to merchants listing prices on its site.[11] Thus, there is also considerable consumer traffic on the site. Shopper.com is owned and operated by Cnet.com, which is consistently among the most viewed sites on the Internet. Each month over 9 million unique consumers access Cnet.[12] In addition to price information, users of Shopper.com have one-click access to Cnet's extensive database of technical specifications and reviews. The Cnet site is ranked first among consumer electronics shopping sites and tenth among all web sites on the Internet.[13]

We gathered information from the site once per day by running a program written in the PERL programming language (known hereafter as 'the spider'), which downloaded this data. For each of the top 1,000 products listed at the site on a given date, the spider collected the product rank for each product and the prices listed by all firms selling that product. The product rank variable consists of a number from 1 to 1,000 indicating each product's

[10] With four million observations, one might expect firms to occasionally make errors in posting their prices. We sometimes observed prices that appeared to reflect a misplaced decimal, such as a merchant quoting a price of \$1000 or \$1 instead of \$100. While the results presented below are based on the cleaned dataset with outliers omitted, the qualitative results presented below are not affected by the inclusion or exclusion of outliers.

[11] A qualified lead occurs when a consumer 'clicks-through' from the Shopper.com site to a merchant's site.

[12] According to a June, 2000, study by Media Metrix.

[13] Based on 100hot.com rankings as of January 18, 2001.

relative popularity measured by the number of qualified leads for that product in the recent past. The information posted at Shopper.com (including prices) is updated twice each day.[14] Consequently, the products included in our sample as well as their rank changes over time. Items in our sample include the Palm III and Palm V personal digital assistants, Canon G1 digital camera, Office 2000 software, and the HP Deskjet 930C inkjet printer.

Table I provides various summary statistics for our data, including the number of competing firms, price levels, and three different measures of price dispersion (the range, coefficient of variation, and the percentage gap between the lowest two prices). Notice that the percentage gap measure of price dispersion (defined as difference in the lowest two prices relative to the lowest price) is the unit-free analog of the Gap measure defined above. Since all of these measures of dispersion are zero for products sold by a single firm, we distinguish between observations where only a single firm lists a price for a product on a given day (denoted as 'Single-Price Listings' in Table I), and those where two or more firms list prices (denoted as 'Multi-Price Listings'). Various measures of price dispersion summarize the set of prices offered for a given product on a given date. Thus, the relevant unit of observation for these measures is what we term a 'product date.' With daily price observations for 1,000 products over an eight month period, there are about a quarter-million product dates. As shown in Table I, our analysis of price dispersion consists of 214,337 product dates with multi-price listings and 13,743 with single-price listings.

Compared to existing studies, the products in our dataset tend to be fairly expensive, with an average price of \$513 across all products and dates.[15] The average minimum price is \$458, or about 12 per cent lower than the average price. Notice that both the average price and average minimum price tend to be higher for less popular products (those with higher ranks). Products with multiple price listings have a lower average price and average minimum price than those with single price listings. Of course, since the mix of products

[14] Merchants have the opportunity to update price quotes twice daily – once at 1:00am and again at 2:00pm (Pacific time). Thus, between each price observation that we collect, each firm had at least one opportunity to change its price in response to rivals' behavior. An audit of prices on April 27, 2001, revealed that over three-fourths of firms update their listing information at least once every twenty-four hours.

[15] More formally, the averages referred to in the table are constructed as follows. Let J_{it} denote the set of firms listing a price for product rank i at time t. Let I_t denote the set of product ranks for which 1 or more prices are listed in period t. Let T be the set of time periods. Finally, let p_{jit} denote the price charged by firm j for product rank i at time t Then the average price in all listings is

$$\frac{1}{\sum_{t\in T}|I_t|}\sum_{t\in T}\sum_{i\in I_t}\left(\frac{\sum_{j\in J_{it}}p_{jit}}{|J_{it}|}\right).$$

Similar methodology was used to construct the other averages.

TABLE I
SUMMARY STATISTICS

	All Product Ranks	Product Ranks 1 – 250	Product Ranks 251 – 500	Product Ranks 501 – 750	Product Ranks 751 – 1000
Total Number of Prices					
Multi-Price Listings	3,925,947	1,202,912	960,709	904,256	858,070
Single-Price Listings	13,743	2,846	3,416	3,785	3,696
Average Price in					
All Listings	$513.23	$472.73	$494.91	$529.60	$555.64
	(882.8)	(665.2)	(838.3)	(1,039.6)	(941.7)
Multi-Price Listings	$491.64	$461.07	$476.41	$486.56	$543.08
	(760.8)	(590.7)	(706.1)	(820.0)	(892.0)
Average Minimum Price in					
All Listings	$457.62	$417.94	$442.78	$475.77	$493.93
	(818.7)	(611.9)	(781.3)	(980.0)	(855.4)
Multi-Price Listings	$432.47	$403.40	$420.97	$428.91	$477.09
	(678.2)	(525.1)	(630.9)	(733.7)	(792.4)
Average Number of Firms in					
All Listings	17.27	21.17	16.90	15.91	15.12
	(11.7)	(14.1)	(10.8)	(10.4)	(10.0)
Multi-Price Listings	18.32	22.23	17.91	16.97	16.10
	(11.3)	(13.7)	(10.3)	(9.9)	(9.6)
Price Dispersion Measures					
Product dates with					
Multi-Price Listings	214,337	54,108	53,633	53,299	53,297
Single-Price Listings	13,743	2,846	3,416	3,785	3,696
Average Range of Prices in					
All Listings	$123.43	$123.88	$117.21	$118.78	$133.87
	(239.5)	(202.5)	(220.5)	(249.3)	(278.3)
Multi-Price Listings	$131.35	$130.40	$124.67	$127.22	$143.15
	(244.9)	(205.7)	(225.3)	(256.0)	(285.5)
Average Coefficient of Variation in					
All Listings	9.10%	9.06%	9.15%	9.10%	9.10%
	(8.0)	(7.2)	(7.9)	(8.4)	(8.6)
Multi-Price Listings	9.69%	9.54%	9.73%	9.75%	9.74%
	(7.9)	(7.1)	(7.8)	(8.3)	(8.5)
Average Gap in Low Prices in					
All Listings	4.39%	3.79%	4.03%	4.71%	5.03%
	(16.2)	(20.4)	(9.9)	(15.4)	(17.3)
Multi-Price Listings	4.67%	3.99%	4.29%	5.04%	5.38%
	(16.7)	(20.9)	(10.2)	(15.9)	(17.8)

Note: Standard deviations are in parentheses.

being offered might differ between single price and multiple price listings, these differences in the levels of prices must be interpreted with caution.

On average, about 17 firms list prices for each product in our sample. Products ranking in the top 250 tend to attract more firms than products not ranked in the top 250. The average range in prices is between $123 and $131, depending upon whether one includes or excludes single-price listings. Levels of price dispersion differ a great deal depending on the measure used. The average range in prices is about thirty per cent, while the average gap between the two lowest prices is about 4.5 per cent. The coefficient of variation lies between these two measures of dispersion, averaging about 9.5 per cent. Interestingly, while the average coefficient of variation is invariant to product rank, the average percentage gap between the lowest two prices is smaller for more popular products. One might therefore speculate that product popularity is a key determinant of price dispersion. However, notice that the more popular products also tend to have more price listings, on average. As we shall see below, differences in the number of firms—not product ranks—are the key to explaining differences in price dispersion across products.

There is considerable variation in the number of firms listing prices for products in our data. Table II shows that single-firm markets accounted for 13,743, or 6.03 per cent, of product dates. Over 80 per cent of all product

TABLE II
FREQUENCY DISTRIBUTION OF THE NUMBER OF FIRMS LISTING PRICES

Number of Firms	Frequency	Percent
1	13743	6.03
2	8791	3.85
3	8615	3.78
4	7363	3.23
5	7325	3.21
6	6972	3.06
7	6649	2.92
8	6708	2.94
9	5723	2.51
10	5924	2.60
11	5949	2.61
12	5967	2.62
13	6085	2.67
14	5814	2.55
15	5898	2.59
16	5751	2.52
17	6185	2.71
18	6044	2.65
19	6154	2.70
20	6441	2.82
21	6408	2.81
22	6426	2.82
23	6834	3.00
24	6877	3.02
25	6265	2.75
26	6404	2.81
27	6231	2.73
28	5853	2.57
29	5292	2.32
30	4655	2.04
31	4132	1.81
32	3379	1.48
33	3046	1.34
34	2721	1.19
35	2341	1.03
36	1879	0.82
37	1592	0.70
38	1391	0.61
39	1074	0.47
40	831	0.36
41	687	0.30
42	548	0.24
43	375	0.16
44	294	0.13
45	263	0.12

TABLE II. (Contd.)

Number of Firms	Frequency	Percent
46	224	0.10
47	268	0.12
48	296	0.13
49	298	0.13
50	309	0.14
51	332	0.15
52	334	0.15
53	328	0.14
54	309	0.14
55	296	0.13
56	237	0.10
57	236	0.10
58	189	0.08
59	141	0.06
60	132	0.06
61	72	0.03
62	67	0.03
63	31	0.01
64	39	0.02
65	26	0.01
66	8	0.00
67	2	0.00
68	3	0.00
69	0	0.00
70	0	0.00
71	0	0.00
72	0	0.00
73	0	0.00
74	0	0.00
75	0	0.00
76	1	0.00
77	0	0.00
78	1	0.00
79	0	0.00
80 or more	2	0.00

dates have between 2 and 30 prices listed, with the number of listings roughly uniformly distributed over this range. Observations where 31 to 40 firms list prices are more rare, accounting for less than 10 per cent of all product dates. Product dates where more than 40 firms list prices account for less than 3 per cent of our data.

Before turning to the analysis of the data, it is useful to highlight some of the strengths and limitations of our study. Key strengths of the dataset used in our study are its duration (eight months), its size (four million price observations), and its composition (1,000 different consumer electronics items). The average *low* price for a product in our dataset is about $460. In contrast, previous studies of price dispersion on the Internet have focused on price dispersion at an instant in time, and have documented price ranges of up to 30 per cent for products such as books and CDs, which typically cost around $15. One might argue that price differences of $4.50 on a $15 item reflect the willingness of some consumers to pay a premium to use a merchant with whom they have an ongoing relationship. It seems less plausible that the price ranges observed in our dataset ($123 on a $513 consumer electronics item) are primarily due to such factors. Another possible explanation for the price dispersion documented in previous studies is that there are economies of scale in shipping these products: it may be optimal for consumers to pay above the lowest price for a single item in order to purchase a low-priced *bundle* from a single merchant. This explanation of price dispersion seems less plausible for the products in our dataset: Shipping costs are small compared to the average price in our sample, and electronics products (such as digital cameras or personal digital assistants) would seem to be less likely to be purchased in bundles than books or CDs.

An important consideration when analyzing data from price comparison services is the veracity and 'seriousness' of the offers listed there. The Shopper.com site has a number of advantages in this regard. First, in contrast to sites relying on shopbot technology,[16] the prices listed at Shopper.com are directly input by the firms themselves. Moreover, price listing on Shopper.com are not free. Specifically, a merchant wishing to list its product pays a one-time, non-refundable fee of $1,000. In addition, at the beginning of each month, it pays an additional fee of $100. Merchants who receive over 250 qualified leads in a given month must pay $0.50 per lead for the first 50,000 leads, and $0.60 for each additional lead. In light of Shopper.com's fee structure and the fact that the site generates over 175,000 qualified leads per day, merchants would seem to have a sharp incentive to post serious prices. A firm attempting a bait and switch strategy – listing a low price with no intention of honoring it – is exposed to considerable downside risk in the form of generating numerous qualified leads (costing at least 50 cents each) while generating few sales and presumably alienating potential customers. On the other hand, firms listing artificially high prices are unlikely to generate enough sales from the site to justify the associated fixed fees of listing.

Second, we conducted an audit of prices listed at Shopper.com for ten randomly selected products among the top 1,000. Since Cnet updates the prices listed on Shopper.com twice per day while firms are free to update prices at their own sites continuously, one would expect some differences in

[16] A shopbot is an automated search engine that visits multiple E-retailers' sites to collect information about prices and other attributes of consumer goods and services. Early shopbots suffered from the defect that information listed there was at times irrelevant and inaccurate. When we began our study, we considered using the price listing site mySimon.com, which is based on shopbot technology. We rejected this approach because search results tended to include a great deal of 'noise.' For example, a product search using the search term 'Palm V' returned a list of products including not only our target item, but also a Deluxe Leather Carrying Case, a Palm V HotSync Cradle, a Palm V Travel Charger, and a Palm V modem. For this reason, we began collecting data from the Shopper.com site rather than from shopbots. We note that the technology used by shopbots has dramatically improved in recent months, and it now appears possible to collect accurate price information through mySimon.com and many other shopbots.

prices to arise even if, at the time of the listing, all prices listed were 100 per cent accurate. In fact, we found that 96 per cent of the 171 prices audited were accurate to within \$1. Moreover, 100 per cent of the low prices were accurate.[17]

Third, there is evidence that consumers can indeed purchase products listed on Shopper.com at the prices listed on the site. We purchased over 30 items (ranging in price from a \$30 headset to a \$600 flat panel monitor) from a number of different merchants listing prices at Shopper.com. In all cases, the prices we paid and the goods received corresponded to the information posted at the site.[18] This is not surprising, since Shopper.com uses a variety of reputational mechanisms that punish vendors who might otherwise be tempted to post erroneous information. For these reasons, we think there is strong evidence to suggest that the price quotes contained in our dataset are serious.

In addition, firms listing prices on Shopper.com make sales almost exclusively online and thus are highly dependent on maintaining their reputations in online channels. At the time of our study, of the three largest 'big box' consumer electronics retailers (BestBuy, Circuit City and CompUSA), only BestBuy had online presence (and only near the end of the study). Thus, the typical merchant selling products at Shopper.com had no brick and mortar presence whatsoever.

The primary limitation of our data is that we were unable to obtain data on the actual quantities of goods purchased at the observed prices.[19] Classical Bertrand models predict that all consumers will purchase from the low-priced firm while clearinghouse models predict that a positive fraction of customers will purchase only at the lowest price while other consumers who are brand loyal or uninformed will purchase at higher prices. Lacking quantity data, we cannot assess whether the predicted sensitivity of consumer behavior more closely matches the Bertrand or clearinghouse predictions. In particular, the

[17] The clearinghouse models discussed in Section 2 operate under the assumption that firms cannot or do not price discriminate. To examine whether this is the case at Shopper.com, we also conducted an audit of ten randomly selected products and compared the price listed on Shopper.com with that obtained by eschewing Shopper.com and going directly to each merchant's site. For the 132 price listings sampled, there were only three cases where prices at the merchant's site were higher than those listed at Shopper.com. In these cases, prices at the three merchants' sites were higher by only \$1.17, \$1.83, and \$0.11. The lowest prices for these items were, respectively, \$214.99, \$185, and \$40.

[18] Our personal experience, as well as more than two years of data on the top 37 products, suggests that shipping costs are fairly constant across firms; see Baye, Morgan, and Scholten (forthcoming).

[19] Other limitations of our data stem from tradeoffs made due to the sheer volume of data being collected. We initially downloaded all of the information listed at the Shopper.com site for a subset of the products, and results were robust to incorporating shipping costs, inventory, reputational ratings, and a variety of other variables. We thus opted to collect the most relevant information on a larger number of products rather than more extensive information on a smaller number. This approach substantially reduced file sizes (enabling us to more thoroughly analyze the data) and reduced the Spider's demand for bandwidth at Shopper.com's site (reducing the probability of Cnet.com taking action to block us from their site).

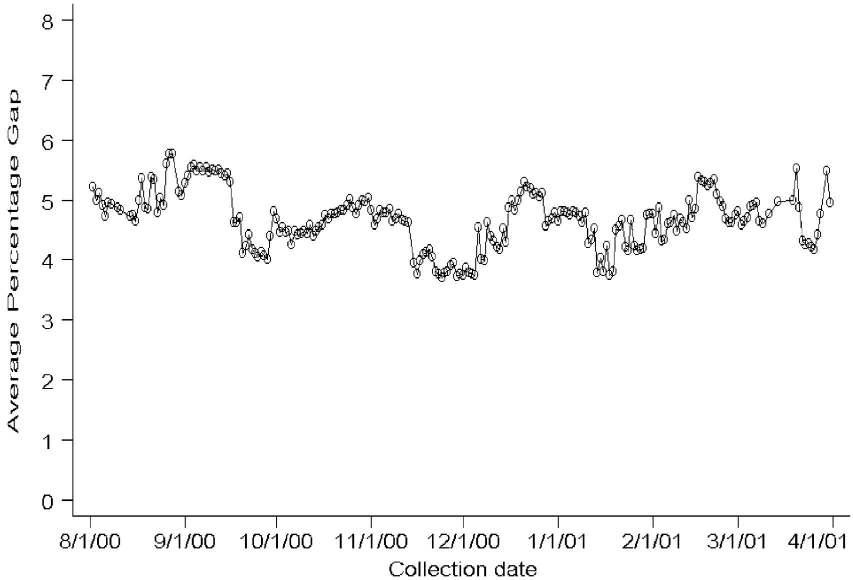

Figure 2
Average Percentage Gap Over Time

classical Bertrand model predicts that a consumer's demand for an individual firm's product is perfectly elastic, while clearinghouse models predict that the demand for an individual firm's product is highly elastic, but not perfectly elastic.[20] Some evidence on this issue is contained in Ellison and Ellison [2004], who examine price and quantity data on computer memory chips sold over the Internet. Their data consists of prices and quantities from a single vendor that lists its price on Pricewatch.com. They find that consumer's are very price sensitive with an estimated elasticity of demand for the firm's product that exceeds 25. This is consistent with what one would expect based on clearinghouse models.

IV. RESULTS

Figure 2 presents a time series graph of the average percentage gap for the period surveyed.[21] As this figure reveals, there is no discernible trend in price

[20] To see this, notice that by raising its price slightly above marginal cost, a firm in a clearinghouse model does not lose demand from uninformed or brand-loyal customers. Furthermore, it only loses informed or price-conscious customers if the price increase results in another firm charging the lowest price.

[21] In examining price dispersion at Shopper.com, we restrict attention to product-dates where two or more firms list prices since dispersion is trivially zero when only a single firm lists a price.

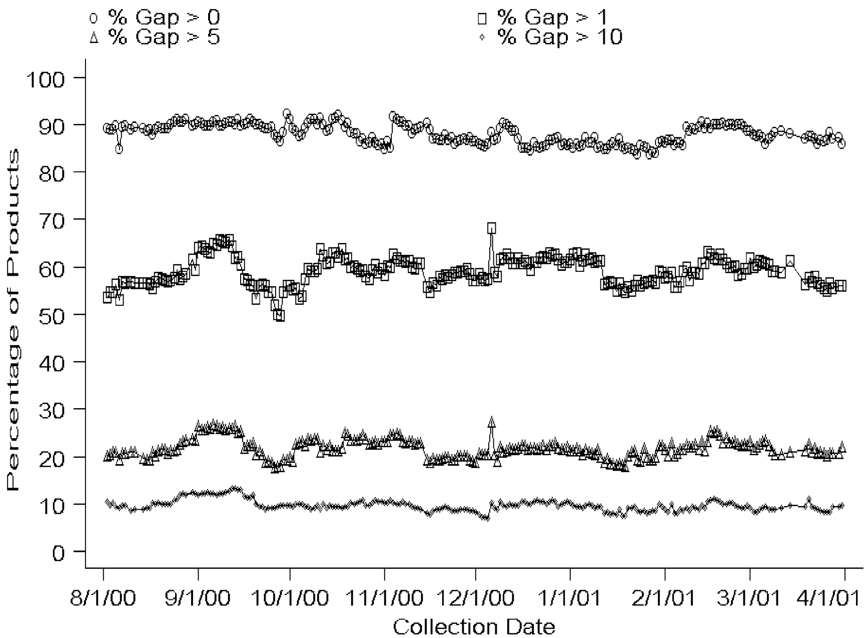

Figure 3
Percentage of Products with Various Percentage Gaps

dispersion over the survey period. Similar analyses for the other measures of dispersion, such as the coefficient of variation and range, likewise show little in the way of a trend. The average coefficient of variation is about 10 per cent in both August 2000 and March 2001. The average percentage range declines by about three per cent.

Figure 3 presents a time series of the fraction of products for which the percentage gap exceeds 0, 1, 5, and 10 per cent. As the figure shows, price dispersion over this period is indeed a pervasive and stable phenomenon. On virtually any date in our sample, there is a strictly positive gap between the lowest two prices for more than 90 per cent of the 1,000 products sampled. About half of all products have a gap of 1 per cent or more, about 20 per cent of the products have a gap of over 5 per cent, and about 10 per cent of the products have gaps exceeding 10 per cent. Thus, a considerable number of products have economically significant gaps between the two lowest prices, and the distribution of gaps has remained relatively unchanged during the survey period.

In short, there is little evidence that price dispersion is a disequilibrium phenomenon that is being corrected over time. If price dispersion is an equilibrium phenomenon, then levels of dispersion should vary with market structure. Figure 4 plots the average percentage gap across all product dates

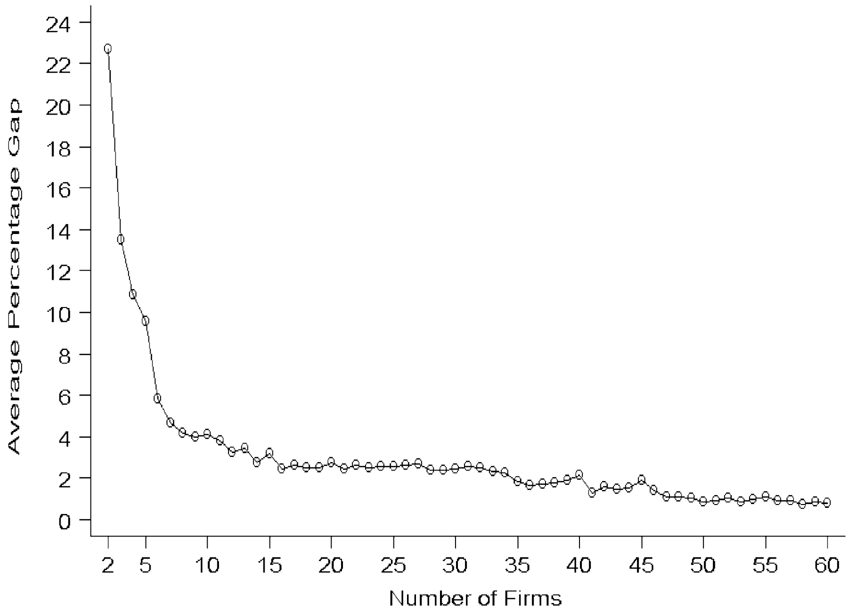

Figure 4
Average Percentage Gap by Number of Firms

against the number of firms listing prices for that product. Notice that the average percentage gap declines sharply as the number of firms listing prices increases. For products where only two firms list a price, the percentage gap averages about 23 per cent. As the number of firms listing prices increases, the percentage gap falls dramatically. It is around 4 per cent for products where ten firms list prices. When fifteen or more prices are listed, the gap is less than 3 per cent.[22]

Figure 4 suggests that price dispersion might vary systematically in the small and in the large. However, this graph fails to account for systematic variation in the number of firms over time as well as across product ranks. In particular, as we saw in Table I, the percentage gap is smaller for more popular products, but more popular products tend to have more firms listing prices. To further confound these effects, over the survey period, there was a substantial decline in the number of firms listing prices on Shopper.com (and by E-retailers generally). The average number of firms listing prices declined about 25 per cent, from 20 to 16 firms.

[22] The average range also depends on market structure. Specifically, the range is significantly higher when many firms list prices than when few firms list prices. For products where only two firms list a price, the range averages about 23 per cent. When five or more firms list prices, the range increases to a neighborhood of 40 per cent.

To help disentangle these effects, we use a simple econometric model to examine the relationship between price dispersion and market structure. We report results based on the gap measure (which, as noted above, provides a more accurate measure of price dispersion in some environments).[23] We regress price dispersion for a particular product date against a number of dummy variables that capture the effects of differences in market structure across products and across time. These controls are potentially important, since the level of price dispersion in the general clearinghouse model depends on the relative size of the market. We use dummy variables for product rank to proxy for these cross-sectional effects (since product rank is a rough measure of the popularity of a product) and 229 time dummies (one for each date) to account for potential dynamic effects.

These results are summarized in Table III, and include a variety of specifications that demonstrate a robust relationship between numbers of firms listing prices for a given product and price dispersion.[24] Model 1 presents a very simple specification of the relationship between price dispersion and numbers of price listings with no controls and where numbers of firms listing prices are pooled into three bins. Model 2 uses this same specification but adds product rank dummies. Model 3 uses individual dummies for numbers of firms listing prices, while Model 4 uses this same specification and adds controls for product rank. Finally, Model 5 is the most general specification, since it controls for both product rank and time fixed effects. In all cases, reported t-statistics are based on robust standard errors to control for potential heteroskedasticity.[25]

The results in Table III are supportive of the view portrayed in Figure 4 that the percentage gap is lower when a large number of firms list prices than when a small number of firms do. Models 1 and 2 indicate that, compared to the case where more than 20 firms list prices, the gap is about 13.5 per cent higher when fewer than five firms list prices, and about 3.2 per cent higher when 5 to 10 firms list prices. Beyond 10 firms, there is little difference in the percentage gaps. Models 3 through 5 show that the results are robust to the bins used to categorize numbers of firms, controls for product rank effects (in Model 4), and potential date effects (in Model 5).

[23] We also ran regressions using the coefficient of variation and range measures of price dispersion. These regressions also reject the hypothesis that price dispersion is invariant to the number of firms at the one per cent significance level.

[24] The results reported here treat the number of firms as exogenous. To control for potential endogeneity, we also ran a variety of 2SLS regressions which instrumented for numbers of firms using product ranks and obtained qualitatively similar results. Further, based on a Hausman test, we failed to reject the hypothesis that the OLS and 2SLS estimates are identical.

[25] The findings reported here are also robust to the use of cluster analysis to control for a potential lack of independence across time.

Model 5 of Table III permits us to test the hypothesis that price dispersion is diminishing over time against the null hypothesis that the coefficients on the date fixed effects are jointly zero (as would be the case if price dispersion is stable over time). As Table III shows, the p-value for this test is 0.24. Thus, based on the gap measure of price dispersion, we find no evidence for diminishing price dispersion, which suggests that price dispersion may be an equilibrium phenomenon. Note that, while the results indicate that price dispersion is lower for the most popular products (those ranked in the top 100), the economic magnitude of these effects are very small compared to the impact on price dispersion of the number of firms listing prices.

One might speculate that the stable price dispersion documented above stems from the fact that new products are entering the sample. In other words, price dispersion for existing products might be falling over time, but convergence is masked by new products entering the dataset with highly dispersed prices. While data limitations prevent us from controlling for such effects with the entire dataset, we collected additional information on the top 100 products that permits us to show that the qualitative results contained in Table III continue to hold with a fixed set of products. In particular, we collected product information for the top 100 products for which price quotes existed on the first date in our sample, and followed this fixed set of products over the eight month period. This permits us not only to examine whether our finding (that the gap is greater in the small than in the large) is robust to controls for product-specific effects, but also to examine whether there is any evidence for convergence based on a fixed set of products. These results are summarized in Table IV.

The first thing to notice is that including product-specific fixed effects dramatically increases the R^2 from less than 10 per cent to about 50 per cent in all of the specifications in Table IV. Furthermore, even with controls for product-specific effects and data containing a fixed set of products, the gap remains greater in the small than in the large. Looking at Model 1, the gap is 12.97 per cent when four or fewer firms list prices, and falls to essentially zero when five or more firms list prices. Models 2, 3, and 4 show that these findings are robust to controls for product rank effects, the bins used to count the number of firms, and date fixed effects.

We can also use this subset of the data to re-examine whether dispersion is decreasing over time. The test of this hypothesis against the null hypothesis that the date fixed effects are jointly zero in Table IV fails to replicate the corresponding test in Table III. Specifically, holding fixed the set of products, we reject the null hypothesis that all the date fixed effects are jointly zero at the 1% level. To examine whether this stems from a downward trend in dispersion, Model 5 includes a linear trend variable, Date Trend. The results here show that, contrary to the notion that price dispersion is decreasing over the lifespans of a fixed set of products, the

Table III
Impact of the Number of Firms Listing Prices on the Percentage Gap

Dependent variable: Percentage Gap. The sample is drawn from Shopper.com for the period August 2, 2000 to March 31, 2001. Each model estimates an OLS regression of the dependent variable on market and product variables obtained from Shopper.com. Coefficients on the date fixed effects are suppressed. Robust t-statistics are reported in parentheses to the right.

	Model 1		Model 2		Model 3		Model 4		Model 5	
Dummy Variable for:	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
Number of Firms Listing Prices										
Between 2 and 4 Firms	0.1362	(49.9)	0.1352	(48.8)						
Between 5 and 10 Firms	0.0316	(45.8)	0.0308	(44.8)						
Between 11 and 20 Firms	0.0058	(22.5)	0.0051	(18.8)						
2 Firms					0.2074	(33.2)	0.2052	(32.6)	0.2063	(104.2)
3 Firms					0.1151	(34.5)	0.1126	(33.5)	0.1142	(57.1)
4 Firms					0.0892	(25.4)	0.0871	(25.0)	0.0887	(41.8)
5 Firms					0.0760	(25.1)	0.0736	(24.5)	0.0752	(35.3)
6 Firms					0.0389	(27.6)	0.0366	(25.8)	0.0381	(17.6)
7 Firms					0.0268	(29.4)	0.0249	(26.8)	0.0264	(12.0)
8 Firms					0.0223	(27.3)	0.0204	(24.6)	0.0220	(10.0)
9 Firms					0.0203	(24.2)	0.0183	(21.5)	0.0200	(8.6)
10 Firms					0.0212	(24.8)	0.0190	(22.0)	0.0206	(8.9)
11 Firms					0.0187	(22.8)	0.0166	(20.1)	0.0182	(7.9)
12 Firms					0.0131	(18.4)	0.0114	(15.6)	0.0131	(5.7)
13 Firms					0.0145	(16.0)	0.0128	(13.8)	0.0145	(6.3)
14 Firms					0.0080	(12.7)	0.0064	(9.5)	0.0077	(3.3)
15 Firms					0.0122	(11.5)	0.0103	(9.6)	0.0115	(5.0)
16 Firms					0.0048	(7.9)	0.0031	(4.8)	0.0044	(1.9)
17 Firms					0.0065	(11.0)	0.0045	(7.1)	0.0060	(2.6)
18 Firms					0.0058	(10.1)	0.0040	(6.4)	0.0057	(2.5)
19 Firms					0.0058	(10.6)	0.0036	(6.2)	0.0054	(2.4)
20 Firms					0.0079	(13.2)	0.0056	(8.9)	0.0074	(3.3)
21 Firms					0.0046	(9.5)	0.0025	(4.7)	0.0040	(1.8)
22 Firms					0.0066	(11.1)	0.0042	(6.7)	0.0057	(2.5)
23 Firms					0.0055	(10.0)	0.0032	(5.6)	0.0046	(2.1)
24 Firms					0.0064	(10.7)	0.0042	(6.8)	0.0056	(2.6)
25 Firms					0.0063	(10.3)	0.0042	(6.8)	0.0051	(2.3)
26 Firms					0.0066	(11.5)	0.0046	(7.8)	0.0059	(2.6)
27 Firms					0.0073	(12.7)	0.0056	(9.4)	0.0063	(2.8)

TABLE III. (Contd.)

Dependent variable: Percentage Gap. The sample is drawn from Shopper.com for the period August 2, 2000 to March 31, 2001. Each model estimates an OLS regression of the dependent variable on market and product variables obtained from Shopper.com. Coefficients on the date fixed effects are suppressed. Robust t-statistics are reported in parentheses to the right.

	Model 1		Model 2		Model 3		Model 4		Model 5	
Dummy Variable for:	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
28 Firms					0.0045	(8.8)	0.0029	(5.5)	0.0036	(1.5)
29 Firms					0.0046	(8.8)	0.0030	(5.5)	0.0038	(1.6)
30 Firms					0.0052	(9.3)	0.0032	(5.5)	0.0037	(1.4)
Product Rank Categories										
Product Ranks 101 – 200			0.0235	(11.0)			0.0231	(10.8)	0.0228	(14.7)
Product Ranks 201 – 300			0.0084	(12.5)			0.0083	(12.2)	0.0079	(5.1)
Product Ranks 301 – 400			0.0081	(11.7)			0.0080	(11.2)	0.0076	(4.9)
Product Ranks 401 – 500			0.0096	(11.7)			0.0089	(10.8)	0.0086	(5.5)
Product Ranks 501 – 600			0.0114	(11.8)			0.0108	(11.2)	0.0104	(6.7)
Product Ranks 601 – 700			0.0129	(11.5)			0.0121	(10.8)	0.0117	(7.5)
Product Ranks 701 – 800			0.0189	(13.1)			0.0175	(12.1)	0.0171	(10.9)
Product Ranks 801 – 900			0.0144	(12.5)			0.0135	(11.6)	0.0130	(8.3)
Product Ranks 901 – 1000			0.0121	(11.9)			0.0110	(10.7)	0.0106	(6.7)
Other Controls										
Intercept	0.0236	(180.2)	0.0121	(29.2)	0.0196	(98.2)	0.0101	(23.6)	0.0092	(7.2)
Date Fixed Effects	No		No		No		No		Yes	
Number of Observations	214,337		214,337		214,337		214,337		214,337	
R^2	0.06		0.07		0.08		0.08		0.08	
Null Hypotheses:										
All Date Fixed Effects are Zero										
p-value									0.24	
All Number of Firm Effects are Zero										
p-value	0.00		0.00		0.00		0.00		0.00	

TABLE IV
IMPACT OF THE NUMBER OF FIRMS LISTING PRICES ON THE PERCENTAGE GAP (FIXED SAMPLE)

Dependent variable: Percentage Gap. The fixed sample of 88 products is drawn from Shopper.com for the period August 2, 2000 to March 31, 2001. Each model estimates an OLS regression of the dependent variable on market and product variables obtained from Shopper.com. Coefficients on the product fixed effects and date fixed effects are suppressed. Robust t-statistics are reported in parentheses to the right.

	Model 1		Model 2		Model 3		Model 4		Model 5	
Dummy Variable for:	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
Number of Firms Listing Prices										
Between 2 and 4 Firms	0.1297	(12.2)	0.1268	(12.0)						
Between 5 and 10 Firms	– 0.0006	(0.2)	0.0019	(0.5)						
Between 11 and 20 Firms	0.0041	(1.8)	0.0074	(3.2)						
2 Firms					0.1632	(11.9)	0.1462	(10.7)	0.1459	(10.7)
3 Firms					0.0992	(8.7)	0.0836	(7.4)	0.0836	(7.4)
4 Firms					0.1511	(9.9)	0.1357	(9.4)	0.1318	(9.0)
5 Firms					0.0263	(3.5)	0.0170	(2.2)	0.0175	(2.2)
6 Firms					0.0140	(2.1)	0.0034	(0.5)	0.0063	(0.9)
7 Firms					– 0.0092	(1.5)	– 0.0206	(3.8)	– 0.0171	(3.1)
8 Firms					– .0002	(0.0)	– 0.0103	(1.9)	– 0.0082	(1.6)
9 Firms					– 0.0041	(0.8)	– 0.0112	(2.1)	– 0.0094	(1.8)
10 Firms					0.0176	(2.8)	0.0119	(1.8)	0.0132	(2.0)
11 Firms					0.0254	(5.6)	0.0185	(3.9)	0.0216	(4.7)
12 Firms					0.0221	(4.6)	0.0165	(3.5)	0.0210	(4.4)
13 Firms					0.0101	(2.5)	0.0066	(1.6)	0.0092	(2.2)
14 Firms					0.0058	(1.6)	0.0055	(1.5)	0.0068	(1.9)
15 Firms					– 0.0007	(0.2)	– 0.0001	(0.0)	– 0.0010	(0.2)
16 Firms					– 0.0025	(0.7)	0.0004	(0.1)	– 0.0001	(0.0)
17 Firms					0.0030	(0.8)	0.0056	(1.4)	0.0052	(1.3)
18 Firms					0.0004	(0.1)	0.0043	(1.2)	0.0035	(1.0)
19 Firms					– 0.0018	(0.6)	0.0010	(0.3)	– 0.0006	(0.2)
20 Firms					0.0061	(1.7)	0.0057	(1.5)	0.0048	(1.3)
21 Firms					0.0003	(0.1)	0.0002	(0.1)	– 0.0009	(0.2)
22 Firms					0.0010	(0.3)	– 0.0027	(0.8)	– .0034	(1.0)
23 Firms					0.0056	(1.7)	0.0062	(1.8)	0.0063	(1.9)
24 Firms					0.0036	(0.8)	0.0037	(0.8)	0.0024	(0.5)
25 Firms					0.0082	(2.1)	0.0068	(1.7)	0.0076	(1.9)
26 Firms					0.0098	(2.6)	0.0061	(1.6)	0.0077	(2.0)
27 Firms					0.0062	(2.4)	0.0043	(1.6)	0.0042	(1.6)

TABLE IV. (Contd.)

Dependent variable: Percentage Gap. The fixed sample of 88 products is drawn from Shopper.com for the period August 2, 2000 to March 31, 2001. Each model estimates an OLS regression of the dependent variable on market and product variables obtained from Shopper.com. Coefficients on the product fixed effects and date fixed effects are suppressed. Robust t-statistics are reported in parentheses to the right.

	Model 1		Model 2		Model 3		Model 4		Model 5	
Dummy Variable for:	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
28 Firms					0.0039	(1.6)	0.0045	(1.8)	0.0022	(0.9)
29 Firms					0.0056	(2.5)	0.0056	(2.3)	0.0041	(1.8)
30 Firms					– 0.0033	(1.9)	– 0.0072	(3.4)	– 0.0058	(3.0)
Product Rank Categories										
Product Ranks 11 – 20			– 0.0001	(0.0)			0.0006	(0.3)	0.0007	(0.4)
Product Ranks 21 – 30			0.0079	(3.0)			0.0078	(3.1)	0.0073	(3.0)
Product Ranks 31 – 40			0.0133	(4.3)			0.0130	(4.5)	0.0132	(4.6)
Product Ranks 41 – 50			0.0151	(4.2)			0.0127	(3.7)	0.0113	(3.3)
Product Ranks 51 – 60			0.0157	(4.2)			0.0133	(3.8)	0.0130	(3.7)
Product Ranks 61 – 70			0.0135	(3.4)			0.0105	(2.7)	0.0101	(2.6)
Product Ranks 71 – 80			0.0011	(0.3)			– 0.0030	(0.8)	– 0.0035	(0.9)
Product Ranks 81 – 90			– 0.0097	(2.3)			– 0.0163	(3.9)	– 0.0162	(3.9)
Product Ranks 91 – 100			0.0003	(0.1)					– 0.0073	(1.4)
Other Controls										
Intercept	0.0186	(16.8)	0.0125	(5.4)	0.0161	(11.8)	– 0.0060	(1.1)	– 1.5894	(10.4)
Date Trend									0.0001	(10.4)
Date Fixed Effects	No		No		No		Yes		No	
Product Fixed Effects	Yes		Yes		Yes		Yes		Yes	
Number of Observations	9,457		9,457		9,457		9,457		9,457	
R^2	0.47		0.48		0.49		0.52		0.51	
Null Hypotheses:										
All Date Fixed Effects are Zero										
p-value							0.00			
All Number of Firm Effects are Zero										
p-value	0.00		0.00		0.00		0.00		0.00	

coefficient on Date Trend is positive and statistically significant. This finding, that price dispersion is increasing over time, is inconsistent with the view that dispersion is a disequilibrium phenomenon that is being corrected over time.

V. DISCRIMINATING AMONG CLEARINGHOUSE MODELS

The findings reported above – that price dispersion depends on market structure and is not diminishing over time – is broadly consistent with the notion that dispersion is an equilibrium phenomenon. We conclude by taking a closer look at special cases of the general clearinghouse model in an attempt to discriminate among them.

Before doing so, we note that the broad findings reported above are also consistent with 'naïve' pricing by firms, whereby 'zero-intelligence' sellers simply post prices at random that range from marginal cost to the monopoly price. In this case, the distribution of prices is independent of n, but nonetheless the expected gap declines as the number of sellers increases due purely to order-statistic effects. This is in contrast to clearinghouse models, where order-statistic effects are confounded by strategic responses by firms that lead to changes in the equilibrium distribution of prices as the number of firms changes.

To compare these models, we perform the following calibration. We set consumers maximal willingnesses to pay, r, equal to the average maximum price observed in the data, which is \$563. We normalize the number of consumers to be unity and set $S = 0.13$, which is based on estimates by Brynjolfsson, Montgomery and Smith [2003] for the percentage of Internet users using price comparison sites over the 2000–2002 period. The number of loyal customers per firm is simply $1 - S$ divided by the average number of firms in our sample. In the case of the Varian and Baye-Morgan models, the total number of loyals, M, is simply $1 - S$. Marginal cost is calibrated based on the U.S. Census Bureaus estimate of the average margin for Electronic Shopping and Mail Order Retailers (NAICS 4541), which is 38.5 per cent.[26] Since clearinghouse models predict that the average transactions price is a weighted-average of the average minimum price and the average list price, one can use these data to estimate marginal cost by $m = \$371$.[27] This completely calibrates all of the models but Baye-Morgan, which also requires a calibrated value for ϕ. We set $\phi = \$3.33$, which is the average cost per day of listing a price at Shopper.com during the period of our study.

[26] Table 6: *Estimated Gross Margin as Per Cent of Sales by Kind of Business*, U.S. Census Bureau, Revised June 1, 2001.

[27] The average transaction price is $T = SE[p_1] + (1 - S)E[p]$. Based on a gross margin of 38.5 per cent, marginal cost is $m = .615T$. The average minimum price and the average listed price in our data, coupled with the estimate of S discussed above, yields our estimate of m.

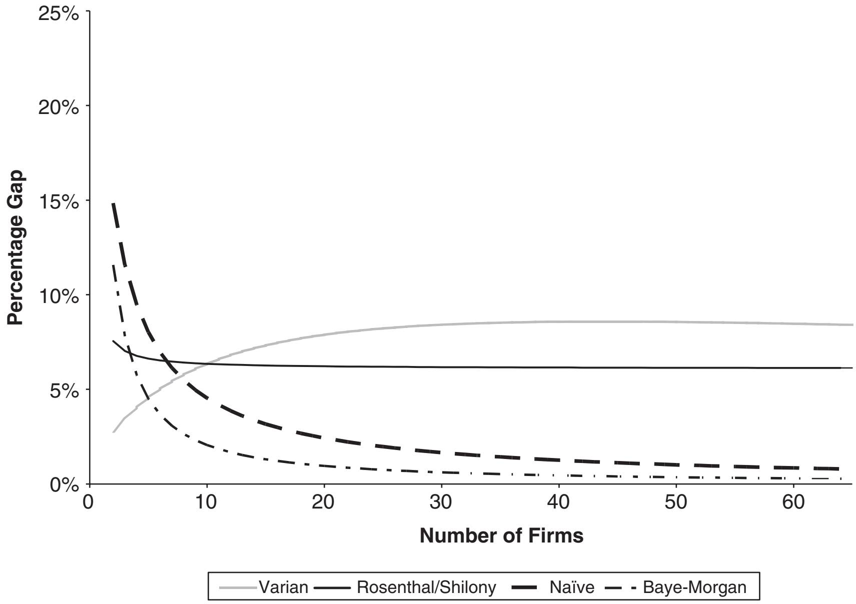

Figure 5
Calibrated Percentage Gaps

Based on this somewhat crude calibration, Figure 5 displays the theoretical relationship between the gap and the number of firms listing prices, which are the theoretical analogues of Figure 4. While Proposition 2 showed that all of the models predict price dispersion which is greater in the small than in the large, Figure 5 illustrates that their predictions are dramatically different when the number of listing firms is in the range occurring in the data. In particular, notice that the Varian model initially predicts an *increasing* relationship between the percentage gap and the number of firms listing prices. This stems from the fact that the strategic effect of an increase in the number of firms on the equilibrium price distribution initially overwhelms the order-statistic effect, thereby leading to a predicted gap (for $n < 40$) that is the opposite of the pattern observed in the data. The Rosenthal/Shilony models more closely match the data, but the strategic effect again pushes the distribution of prices in a direction opposite to the order statistic effect, thus leading to only a modest decline in price dispersion. The Baye-Morgan model and the Naïve model both predict a pattern similar to that observed in Figure 4.

While Figure 5 reveals sharp differences in the predictions of special cases of the general clearinghouse model, it does not permit one to distinguish between equilibrium behavior (where changes in n have strategic effects on price distributions) and Naïve behavior (where they

TABLE V
IMPACT OF THE NUMBER OF FIRMS LISTING PRICES ON THE AVERAGE PRICE (FIXED SAMPLE)*

Dependent variables: Average Price and natural log of Average Price. The fixed sample of 88 products is drawn from Shopper.com for the period 2 August, 2000 to March 31, 2001. Models 1 and 2 estimate OLS regressions of the Average Price and the natural log of Average Price, respectively, on market and product variables obtained from Shopper.com. Models 3 and 4 estimate two-stage least squares regressions of the dependent variables Average Price and the natural log of Average Price, respectively, on market and product variables obtained from Shopper.com. Coefficient estimates on the product fixed effects and date fixed effects are suppressed. Robust t-statistics are reported in parentheses to the right.

Dependent Variable	Model 1 Average Price		Model 2 ln(Average Price)		Model 3 ln(Average Price)		Model 4 ln(Average Price)	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
Independent Variables								
Number of Firms	– 0.3695	(2.7)			– 1.0026	(3.5)		
ln(Number of Firms)			– 0.0137	(2.5)			– 0.0245	(1.9)
Intercept	1167.40	(110.1)	7.05	(393.6)	204.45	(19.1)	5.35	(169.1)
Date Fixed Effects	Yes		Yes		Yes		Yes	
Product Fixed Effects	Yes		Yes		Yes		Yes	
Number of Observations	9741		9741		9741		9741	
R^2	0.98		0.99		0.98		0.99	

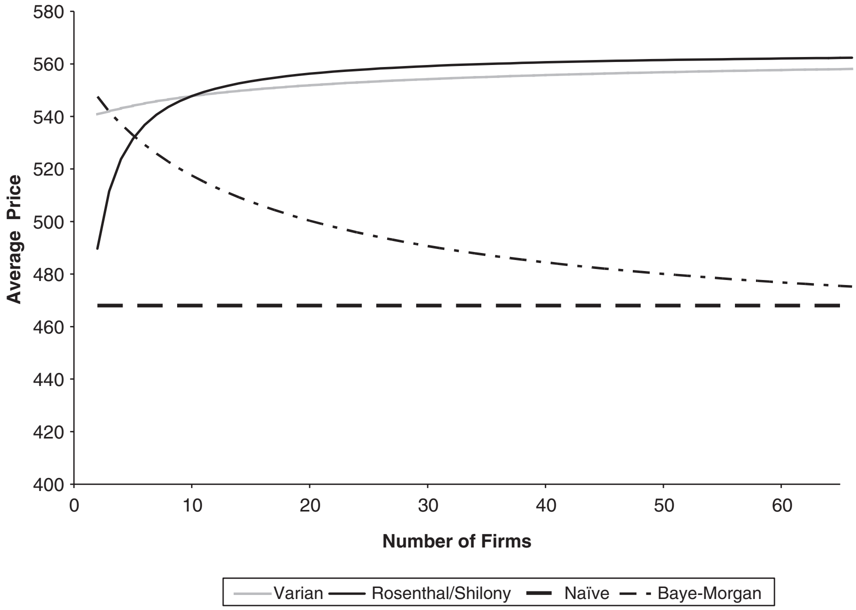

Figure 6
Calibrated Average Prices

do not). A simple way of testing whether strategic effects are present is to examine whether the first moment of the empirical distribution of prices varies with n.[28] Accordingly, we regress average prices on numbers of firms along with controls for product and date fixed effects. These results are presented in Table V, and are based on the same fixed set of products used in the analysis contained in Table IV in order to permit controls for product-specific effects. Specifically, Table V reports results of both linear and a log-linear specifications of the dependent variable and uses both OLS and 2SLS (where we instrument for the number of firms using product rank). In all cases, the results indicate a negative relationship between average price and the number of competing firms. We can reject at conventional significance levels the null hypothesis (implied by the Naïve model) that there is no relationship between average prices and numbers of firms listing prices in favor of the one-sided alternative that average prices decline with the number of competing firms. In short, we find evidence of strategic effects in the data.

Are the observed strategic effects consistent with the predictions of the general clearinghouse model? To examine this question, we again use the

[28] We are grateful to the editor for suggesting this test.

calibrated models to examine the theoretical relationship between average price and the number of firms. As Figure 6 reveals, the Varian and Rosenthal/Shilony models predict that greater competition leads to higher average prices. This is not only counter-intuitive, but also contrary to the empirical findings reported in Table V. In contrast, the Baye-Morgan model predicts that greater competition leads to lower average prices, which is consistent with the data.

VI. CONCLUSION

Our analysis indicates that the levels of price dispersion observed at Shopper.com were stable over a period in which consumer usage of price comparison sites increased by about 13 per cent. Our empirical finding that both the level of dispersion and average prices are greater in the small than in the large can be rationalized by a clearinghouse framework, although some special cases of the general model more closely match the observed data than others. Unfortunately, our dataset is not detailed enough to permit the structural estimation required to discriminate further among competing clearinghouse models. The findings reported here suggest that future research along these lines might prove useful.

At a more general level, our results suggest that it is useful to control for market structure in online markets when comparing levels of dispersion across products or over time. Such controls permit one to disentangle dynamic effects (such as learning on the part of firms or the increased usage of price comparison sites by consumers) from market structure effects. Viewed in this context, our results compliment recent work by Brown and Goolsbee [2002], who provide convincing evidence that the emergence of the Internet led to sharp declines in insurance premiums during the 1990s. Our results suggest that part of this decline might have stemmed from increases in the number of insurance companies choosing to list their rates at price comparison sites.

APPENDIX

The proof of part 2 of Proposition 1 relies on the following theorem, which establishes the existence of equilibrium in the general clearinghouse model.

Theorem 1: *Let* $0 \leqslant \phi < \frac{n-1}{n}(r-m)S$. *Then, a symmetric equilibrium of the general clearinghouse model:*

1. *Each firm lists its price at the clearinghouse with probability*

$$\alpha = 1 - \left(\frac{\frac{n}{n-1}\phi}{(r-m)S}\right)^{\frac{1}{n-1}}.$$

2. *When a firm lists its price at the clearinghouse it charges a price drawn from the distribution*

$$F(p) = \frac{1}{\alpha}\left(1 - \left(\frac{\frac{n}{n-1}\phi + (r-p)L}{(p-m)S}\right)^{\frac{1}{n-1}}\right) \text{ on } [p_0, r],$$

where

$$p_0 = \frac{\frac{n}{n-1}\phi + Lr + Sm}{L+S}.$$

3. *When a firm does not list its price, it charges a price equal to r.*

Proof

First, observe that if a firm does not list its price at the clearinghouse, it is a dominant strategy to charge a price of r.

Next, notice that $\alpha \in (0, 1]$ whenever

$$\frac{n\phi}{(n-1)(r-m)S} < 1.$$

This condition holds, since $\phi < \frac{n-1}{n}(r-m)S$.

We next show that F is a well-defined cdf with $m < p_0 < r$. First,

$$\begin{aligned} p_0 &= \frac{\frac{n}{n-1}\phi + (Lr + Sm)}{(L+S)} \\ &< \frac{(r-m)S + (Lr+Sm)}{(L+S)} \\ &= r, \end{aligned}$$

where the inequality follows from the fact that $\phi < \frac{n}{n-1}(r-m)S$. Furthermore,

$$\begin{aligned} p_0 &= \frac{\frac{n}{n-1}\phi + (Lr + Sm)}{(L+S)} \\ &\geq \frac{(Lr+Sm)}{(L+S)} \\ &> \frac{(L+S)m}{(L+S)} \\ &= m, \end{aligned}$$

where the weak inequality follows from the fact that $\phi \geqslant 0$ and the strict inequality follows since $r > m$.

By construction, $F(p_0) = 0$. To see that $F(r) = 1$, we compute

$$\begin{aligned} F(r) &= \frac{1}{\alpha}\left(1 - \left(\frac{\frac{n}{(n-1)}\phi}{(r-m)S}\right)^{\frac{1}{n-1}}\right) \\ &= \frac{1}{\alpha}\alpha. \end{aligned}$$

It remains to show that F is strictly increasing in p :

$$\begin{aligned}\frac{\partial F(p)}{\partial p} &= \frac{Z^{\frac{1}{n-1}-1}}{(n-1)\alpha}\frac{(r-m)L+\left(\frac{n}{n-1}\phi\right)}{(p-m)^2 S} \\ &> 0,\end{aligned}$$

where

$$Z = \frac{\frac{n}{n-1}\phi + (r-p)L}{(p-m)S} > 0.$$

Next, we show that, conditional on listing a price, a firm can do no better than pricing according to F. It is obvious that choosing a price above or below the support of F is dominated by choosing a price in the support of F. A firm choosing a price p in the support of F earns expected profits of

$$E\pi(p) = (p-m)\left(L + \left(\sum_{i=0}^{n-1}\binom{n-1}{i}\alpha^i(1-\alpha)^{n-1-i}(1-F(p))^i\right)S\right) - \phi.$$

Using the binomial theorem, we can rewrite this as:

$$\begin{aligned}E\pi(p) &= (p-m)\left(L + \left((1-\alpha F(p))^{n-1}\right)S\right) - \phi \\ &= (p-m)\left(L + \left(\frac{\frac{n}{n-1}\phi + (r-p)L}{(p-m)S}\right)S\right) - \phi \\ &= (r-m)L + \frac{\phi}{n-1}.\end{aligned}$$

Since this is independent of p, it follows that F is a best response to the other $n-1$ firms' pricing based on F.

When $\phi = 0$, it is a weakly dominant strategy to list. It remains to show that when $\phi > 0$ and $\alpha \in (0,1)$, a firm earns the same expected profits regardless of whether it lists its price. But a firm that does not list earns expected profits of

$$\begin{aligned}E\pi &= (r-m)\left(L + \frac{S}{n}(1-\alpha)^{n-1}\right) \\ &= (r-m)L + \frac{\phi}{n-1},\end{aligned}$$

which equals the expected profits earned by listing any price $p \in [p_0, r]$. *Q.E.D.*

Proof of Proposition 2

1. Baye-Morgan Model

Note that in the Baye-Morgan model, there is a distinction between the number of competing firms (n) and the number of firms listing prices at the clearinghouse, which we denote as k. To evaluate the case where the number of *listings* becomes arbitrarily large requires one to first let the number of competing firms go to infinity and then evaluate the order statistics of the limit distribution as $k \to \infty$ to obtain the expected gap as the number of listings on the clearinghouse grows large.

First, notice that as $n \to \infty$, the limit distribution of listed prices is

$$\lim_{n\to\infty} F(p) = \lim_{n\to\infty} \frac{1}{\alpha}\left(1 - \left(\frac{\frac{n\phi}{n-1} + (r-p)L}{(p-m)S}\right)^{\frac{1}{n-1}}\right)$$
$$= \frac{\ln\left(\frac{\phi+(r-p)L}{(p-m)S}\right)}{\ln\left(\frac{\phi}{(r-m)S}\right)}$$
$$= F^*(p) \text{ on } \left[p_0^*, r\right],$$

where $p_0^* = \frac{\phi+Lr+Sm}{L+S}$.

Since $F^*(p)$ is atomless with positive support, it is clear that $E(G) > 0$ for finite k. To show that the gap is zero in the limit, it is sufficient to establish that the expectation of the 2nd lowest of k draws from F^*, $E\left[p_2^{(k)}\right]$, converges to the lower support of the distribution as $k \to \infty$. That is

$$\lim_{k\to\infty} E\left[p_2^{(k)}\right] = p_0^*.$$

To establish this result, denote the cumulative distribution of the 2nd lowest of k draws by $H(t)$. It is well-known that for any cdf F,

$$H(p) = \left(1 - (1-F(p))^n - nF(p)\,(1-F(p))^{n-1}\right),$$

with corresponding density $h(p)$.

Hence,

$$E\left[p_2^{(k)}\right] = \int_{p_0^*}^{r} th(t)dt,$$

where $h(\cdot)$ is evaluated using $F(\cdot) = F^*(\cdot)$ and with the corresponding density $f^*(\cdot)$. Now, fix $\varepsilon > 0$. Then

$$E\left[p_2^{(k)}\right] = \int_{p_0^*}^{p_0^*+\varepsilon} th(t)dt + \int_{p_0^*+\varepsilon}^{r} th(t)dt$$
$$< \left(p_0^*+\varepsilon\right)H\left(p_0^*+\varepsilon\right) + r\left(1 - H\left(p_0^*+\varepsilon\right)\right).$$

Thus, to prove the result requires only that we show that $lim_{k\to\infty} H\left(p_0^*+\varepsilon\right) = 1$. To see this, notice that

$$\lim_{k\to\infty} H\left(p_0^*+\varepsilon\right)$$
$$= \lim_{k\to\infty}\left(1 - \left(1 - F^*\left(p_0^*+\varepsilon\right)\right)^k - kF^*\left(p_0^*+\varepsilon\right)\left(1 - F^*\left(p_0^*+\varepsilon\right)\right)^{k-1}\right)$$
$$= \left(1 - \lim_{k\to\infty}\left(1 - F^*\left(p_0^*+\varepsilon\right)\right)^k - \lim_{k\to\infty} kF^*\left(p_0^*+\varepsilon\right)\left(1 - F^*\left(p_0^*+\varepsilon\right)\right)^{k-1}\right)$$
$$= 1 - F^*\left(p_0^*+\varepsilon\right)\lim_{k\to\infty} k\left(1 - F^*\left(p_0^*+\varepsilon\right)\right)^{k-1}.$$

Since $\left(1 - F^*\left(p_0^* + \varepsilon\right)\right) \in (0, 1)$, it then follows (from L'Hopital's rule) that

$$\lim_{k\to\infty} k\left(1 - F^*\left(p_0^* + \varepsilon\right)\right)^{k-1} = 0,$$

which establishes the result.

Finally,

$$\lim_{k\to\infty} E[G] = \lim_{k\to\infty}\left(E\left[p_2^{(k)}\right] - E\left[p_1^{(k)}\right]\right) = 0$$

follows from the fact that $p_0^* \leqslant lim_{k\to\infty} E\left[p_1^{(k)}\right] < lim_{k\to\infty} E\left[p_2^{(k)}\right] = p_0^* + \varepsilon$ for small $\varepsilon > 0$ *Q.E.D.*

2. Varian Model

In the Varian model, $F(p)$ is atomless with positive support, so it is clear that $E[G] > 0$ for finite n. To show that $lim_{n\to\infty} E[G] = 0$, it is sufficient to show that $lim_{n\to\infty} E\left[p_2^{(n)}\right] = 0$. Fix $\varepsilon > 0$. Now

$$\begin{aligned}\lim_{n\to\infty} E\left[p_2^{(n)}\right] &= \lim_{n\to\infty} \int_{p_0}^{r} th(t)dt \\ &= \lim_{n\to\infty} \int_{p_0}^{p_0+\varepsilon} th(t)dt + \int_{p_0+\varepsilon}^{r} th(t)dt \\ &< \lim_{n\to\infty} ((p_0 + \varepsilon)H(p_0 + \varepsilon) + r(1 - H(p_0 + \varepsilon))).\end{aligned}$$

In the Varian model,

$$F(p) = \left(1 - \left(\frac{\frac{M}{n}(r - p)}{(p - m)S}\right)^{\frac{1}{n-1}}\right)$$

and

$$p_0 = \frac{\frac{M}{n}r + Sm}{\frac{M}{n} + S}.$$

Taking limits yields

$$\lim_{n\to\infty} H(p_0 + \varepsilon) = 0$$

and

$$\lim_{n\to\infty} p_0 = m.$$

Hence

$$\lim_{n\to\infty} E\left[p_2^{(n)}\right] = m + \varepsilon.$$

Finally,

$$\lim_{n\to\infty} E[G] = \lim_{n\to\infty}\left(E\left[p_2^{(n)}\right] - E\left[p_1^{(n)}\right]\right) = 0$$

follows from the fact that $m \leqslant lim_{n\to\infty} E\left[p_1^{(n)}\right] < lim_{n\to\infty} E\left[p_2^{(n)}\right] = m + \varepsilon$ for small $\varepsilon > 0$.

Q.E.D.

3. *Rosenthal/Shilony Model*

Since each firm in the Rosenthal/Shilony model earns expected profits of $E\pi_i = (r - m)L$, the assumption that L is a positive constant implies that expected industry profit, $\sum_{i=1}^{n} E\pi_i = n(r - m)L$, tends to infinity as the number of firms increases without bound. To mitigate this shortcoming while allowing initial entrants to bring additional loyal consumers, it is necessary to slightly modify the model by assuming

$$L = min\left(L^*, \frac{\Pi}{(r - m)n}\right),$$

where $\Pi < \infty$ and $L^* < \infty$ are positive constants. In this case, it is easy to show that industry profits are bounded from above by Π. In fact, for any finite n, expected industry profits are

$$\sum_{i=1}^{n} E\pi_i = n(r - m)L = min(n(r - m)L^*, \Pi)$$

and furthermore,

$$\lim_{n\to\infty} \sum_{i=1} E\pi_i = \Pi < \infty.$$

There exists a fixed $\Pi < \infty$ such that the modified model is identical to the original models for small and intermediate values of n (say, $n \leqslant 65$). Since all of the empirical results discussed in the text are based on markets where $n \leqslant 65$, none of the quantitative results discussed in the text depend on this modification.

Since $F(p)$ is atomless with positive support, it is clear that $E[G] > 0$ for finite n. To show that $lim_{n\to\infty} E[G] = 0$, notice that the limit version of the model is identical to the Varian model; therefore

$$\lim_{n\to\infty} E[G] = 0.$$

Q.E.D.

REFERENCES

Bakos, Y., 2000, 'The Emerging Landscape of Retail E-Commerce,' *Journal of Economic Perspectives*, 19(1), pp. 63–82.

Baye, M. R. and Morgan, J., 2001, 'Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,' *American Economic Review*, 91(3), pp. 454–474.

Baye, M. R., Morgan, J. and Scholten, P., forthcoming, 'Persistent Price Dispersion in Online Markets,' in Jansen, D. (ed.), *The New Economy*, (University of Chicago Press).

Brown, J. R. and Goolsbee, A., 2002, 'Does the Internet Make Markets More Competitive? Evidence from the Insurance Industry,' *Journal of Political Economy*, 110(4), pp. 481–507.

Brynjolfsson, E. and Smith, M. D., 2001, 'Consumer Decision-Making at an Internet Shopbot: Brand Still Matters,' *Journal of Industrial Economics*, 49(4), pp. 541–558.

Brynjolfsson, E. and Smith, M. D., 2000, 'Frictionless Commerce? A Comparison of Internet and Conventional Retailers,' *Management Science*, 46(4), pp. 563–585.

Brynjolfsson, E., Montgomery, A. and Smith, M. D., 2003, 'The Great Equalizer: The Role of Shopbots in Electronic Markets', *Carnegie Mellon University Working Paper*.

Burdett, K. and Judd, K. L., 1983, 'Equilibrium Price Dispersion,' *Econometrica*, 51, pp. 955–969.

Carlson, J. A. and Pescatrice, D. R., 1980, 'Persistent Price Distributions,' *Journal of Economics and Business*, 33(1), pp. 21–27.

Ellison, G. and Ellison, S. F., 2004, 'Search, Obfuscation, and Price Elasticities on the Internet,' *mimeo*.

Gatti, J. R. J., 2001, 'Equilibrium Price Dispersion with Sequential Search,' *mimeo*.

Janssen, M. and Moraga, J. L., 2000, 'Pricing, Consumer Search and the Size of Internet Markets,' *Tinbergen Institute Discussion Paper* TI200-0042/1.

Narasimhan, C., 1988, 'Competitive Promotional Strategies,' *Journal of Business*, 61(4), pp. 427–449.

Pratt, J. W., Wise, D. A. and Zeckhauser, R., 1979, 'Price Differences in Almost Competitive Markets,' *Quarterly Journal of Economics*, 93(2), pp. 189–211.

Reinganum, J. F., 1979, 'A Simple Model of Equilibrium Price Dispersion,' *Journal of Political Economy*, 87(4), pp. 851–858.

Rosenthal, R. W., 1980, 'A Model in Which an Increase in the Number of Sellers Leads to a Higher Price,' *Econometrica*, 48(6), pp. 1575–1579.

Salop, S. C. and Stiglitz, J. E., 1977, 'Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion,' *Review of Economic Studies*, 44(9), pp. 493–510.

Shilony, Y., 1977, 'Mixed Pricing in Oligopoly,' *Journal of Economic Theory*, 14(2), pp. 373–388.

Smith, M. D., Bailey, J. and Brynjolfsson, E., 2000, 'Understanding Digital Markets: Review and Assessment,' in Brynjolfsson, E. and Kahin, B. (eds.), *Understanding the Digital Economy: Data, Tools, and Research*, ((MIT Press).

Sorensen, A., 2000, 'Equilibrium Price Dispersion in Retail Markets for Prescription Drugs,' *Journal Political Economy*, 108(4), pp. 833–850.

Spulber, D. F., 1995, 'Bertrand Competition when Rivals' Costs are Unknown,' *Journal of Industrial Economics*, 43(1), pp. 1–11.

Stahl, D. O. II., 1989, 'Oligopolistic Pricing with Sequential Consumer Search,' *American Economic Review*, 79(4), pp. 700–712.

Stahl, D. O., 2000, 'Strategic Advertising and Pricing in E-Commerce,' in Baye, M. R. (ed.) *Volume 9: Advances in Applied Microeconomics*, ((Elsevier Science).

Varian, H., 1980, 'A Model of Sales,' *American Economic Review*, 70(4), pp. 651–659.