Diverging Prices of Novel Goods

Nick Huntington-Klein, CSU Fullerton
Alasdair Young, CEO, MTGPrice.com

November, 2017

Diverging Prices of Novel Goods

Nick Huntington-Klein[a,*], Alasdair Young[b]

[a]*California State University, Fullerton*
[b]*CEO, MTGPrice.com*

Abstract

It is well-established that the costs of consumer search lead many markets to exhibit price dispersion in equilibrium. However, relatively little attention is paid to how dispersion changes over time. The presiding assumption is that prices out of equilibrium should converge. We develop a model in which shifts in demand combined with menu costs lead dispersion to naturally rise soon after release for some novel goods. We examine an online market for collectible cards and find that price dispersion increases over time for more than half of the goods in the market.

Keywords: Price dispersion, secondary market, e-commerce, search

JEL: D40, D83, L11, L81

I. INTRODUCTION

Despite its inclusion in many introductory economics courses, the Law of One Price is known to rarely hold empirically (Varian, 1980). Theoretical explanations for the existence of price dispersion include the presence of heterogeneous goods or store characteristics, costly consumer search, imperfectly informed consumers, or bounded rationality among sellers (Salop and Stiglitz, 1977, 1982; Baye and Morgan, 2004; Baye et al., 2006).

The online retail market is a useful setting for the study of price dispersion. Because the internet offers easy search and a number of sites that display prices from many stores at once, one might have ex-ante expected that price dispersion would be nonexistent or at least

---

*Corresponding Author. Email: nhuntington-klein@fullerton.edu. Address: 800 N. State College Blvd., Fullerton, CA, 92831

low online (Smith et al., 2000). However, an extensive literature shows that price dispersion both exists online and persists over time (e.g. Brynjolfsson and Smith, 2000; Brown and Goolsbee, 2002; Chevalier and Goolsbee, 2003; Baye et al., 2004; Ellison and Ellison, 2009; Xing, 2008, 2010; Ghose and Yao, 2011; Einav et al., 2015; Coey et al., 2016; Gorodnichenko and Talavera, 2017; Dinerstein et al., 2017). Even consumers who search multiple stores do not always purchase at the lowest price (De los Santos et al., 2012). Dispersion online cannot be fully explained by differences across stores in reputation or services (Pan et al., 2002; Baylis and Perloff, 2002; Ratchford et al., 2003; Venkatesan et al., 2007).

The existence of price dispersion in equilibrium is well-established both in theory and in empirical research for online markets. There is also theoretical interest in the dynamics of price dispersion, which often frame the dynamics of price dispersion in a differentiated-goods setting (e.g., Rothschild, 1974; Bergemann and Välimäki, 1996; Keller and Rady, 2003), as the result of ongoing experimentation or information shocks among consumers or producers. However, there is relatively little empirical work examining changes in dispersion over time, as opposed to showing that dispersion exists. In a market where price dispersion exists in equilibrium, when would we expect dispersion to rise or fall?

When certain market factors change, we would expect price dispersion to change as well. Verboven and Goldberg (2001) examine how international price dispersion in the European car market changes over time in response to trade policies and exchange rates. Pan et al. (2003) find that price dispersion in online markets dropped from 2000 to 2001, but rebounded by 2003, and was higher after the dot-com bubble burst than before. Baye et al. (2004) find that price dispersion for goods on an online price comparison site varies with market structure, and responds to the entry and exit of firms.

Another potential source of change in price dispersion is the refinement of information about a good. Given that many theoretical explanations of the dynamics of price dispersion rely on uncertainty about product quality, we might expect higher amounts of dispersion for a new good, which has quality of higher uncertainty. We then might expect dispersion

to mellow over time to some baseline equilibrium dispersion as information about product quality or which firms offer low prices is spread, and entry and exit occur (Brynjolfsson and Smith, 2000). There is some evidence that price dispersion decreases as markets mature. Ratchford et al. (2003) find that dispersion in an online price comparison site for consumer goods declined between 2000 and 2001, and attributes this to maturation of the market. Perhaps more tellingly, Brown and Goolsbee (2002) find that the introduction of internet search sites to the online market for life insurance increased price dispersion within demographic groups, but this increase disappeared over time. There is a small literature on price convergence in newly formed markets, as described in Doraszelski et al. (2016), which often finds that prices in new markets tend to converge, but only after long periods of time. However, this related literature focuses on newly formed markets (in the case of Doraszelski et al. 2016, a new electricity market in the United Kingdom formed by deregulation), rather than new goods released into existing market structures, for which the dynamics of price dispersion may be different. In newly formed markets, the market itself must mature, while for new goods in existing market structures, only information about the product itself must mature.

In this paper we focus on the market for novel goods introduced into an existing market. We argue that price dispersion is likely to regularly *increase* for some goods soon after release, even without changes to market structure or other major shocks. While in the long run, price dispersion may diminish to some long-run equilibrium level, in the short run dispersion is likely to increase in response to demand shifts, which should occur as a regular result of quality revelation for new goods. We present a model in which goods are of uncertain quality. Changes in perceived quality shift the demand curve, which in the presence of menu costs leads retailers to change prices in a way that may increase price dispersion.

We find evidence of increasing dispersion shortly after release in the secondary market for a collectible card game, a mature market into which new products of uncertain quality are regularly released and prices are set by individual retailers. While the market for collectible

cards is not of direct interest to most economists, it offers an opportunity to observe shifts in price dispersion for a large number of distinct novel goods that are homogeneous across retailers, sold in meaningful quantities, and are regularly released into a highly competitive market. Simply, this is a clean environment to observe the predicted behavior, which can then be applied to other markets with similar features.

In a data set of 430,521 sales of 537 goods at 909 stores on a price comparison site observed over two years, dispersion in sale prices increases within the first three months after a good's release for more than half of all goods. On average, the coefficient of variation increases by about .009 each week. The existence of price dispersion, and its increase over time, is not accounted for by differences across stores in reputation or shipping policies. The data is consistent with the presence of menu costs for stores and shifting demand curves, which the model suggests are sufficient to produce equilibria with increasing price dispersion.

These empirical results add to the extensive literature on the existence of price dispersion in online markets by describing the short-term dynamics of that price dispersion. Evidence in favor of increasing dispersion as the most common result aligns in some ways with theory (Rothschild, 1974; Bergemann and Välimäki, 1996; Keller and Rady, 2003) but is applied here in a homogeneous-goods setting where differentiated goods are not a likely reason for dispersion dynamics. I suggest menu costs combined with demand shifts as an alternative intuitive explanation, and demonstrate that both of those features are present in the market, and that short-run increasing dispersion is observed.

II. Model

In this section we introduce a simple model of pricing and price dispersion for novel goods, borrowing several elements from the clearinghouse consumer search model of Salop and Stiglitz (1977), in which consumers may pay to access the full distribution of prices. See Baye et al. (2006) for a review of this class of model. Salop and Stiglitz (1977) show that it is

possible for a market with consumer search costs to lead to price dispersion in equilibrium. This occurs because consumers with high search costs would rather pay a high price for the good than work to seek out a low-priced retailer. Some retailers take advantage of this by raising prices, but not all retailers can do so without incentivizing more search among consumers, eliminating sales for high-priced retailers.

The goal of this model is not to provide a comprehensive model of price dispersion, but rather to formalize the intuition that increases in price dispersion can arise from a clearinghouse model in the presence of menu costs as long as demand continues to shift. We suggest that these demand shifts are routine for new products of uncertain quality, and so increases in price dispersion can also be expected to be routine. We will not be showing the uniqueness of an increasing-dispersion equilibrium, but just that the equilibrium exists.

We begin with the clearinghouse model and add demand shifts and menu costs. A qualitative explanation of the mechanism at hand is that retailers with high menu costs must prepare for demand shifts by choosing prices that do not change between periods. As perceived quality changes, the left-behind prices of high menu cost firms allow low menu cost firms to set very high prices (if demand increases), or forces them to set very low prices (if demand decreases), both of which can increase dispersion in the set of prices.

Consider a market of N_c consumers and N_s retailers over two periods. N_s is held fixed since the time frame of the market is short enough, and the stakes for any individual good small enough, that entry is unlikely to occur. Each consumer may buy up to one unit in each period. The good in question is new, and so perceived value of the product may change between periods due to learning, even in the absence of external shocks. Each customer's perceived value is distributed $v_{it} \sim f(x|V_t^*)$ over all consumers in each period.

V_t^* is a shared demand parameter of the good, which may shift from period to period. V_t^* is such that $V_t^* > V_{t'}^* \Rightarrow 1 - F(p|V_t^*) \geq 1 - F(p|V_{t'}^*) \ \forall \ p$, where F is the cumulative distribution function associated with f, and p is the price of the good. In other words, an increase in V_t^* indicates a rightward shift of the demand curve.

Consumers know the distribution of prices but are imperfectly informed about which prices are offered by which retailers for the good. Retailers post prices $P_t = \{p_{1t}, p_{2t}, ..., p_{N_st}\}$ in period t. When shopping each period, a consumer i may choose whether or not to pay a search cost. If they pay the cost, then they observe which stores offer which prices and buy the good from one of the lowest-price retailers at random at a price of $\min_j\{p_{jt}\}$ as long as the minimum price is below their perceived value. If they do not pay the cost, they choose a store at random and pay on average $\bar{p}_t = \sum_j p_{jt}/N_s$, as long as the price they observe is below their perceived value. Search costs vary over consumers. A portion γ have low costs $c_i = c^L$ and $1 - \gamma$ have high costs $c_i = c^H$.

A risk-neutral consumer i will choose to search if and only if $\min_j\{p_{jt}\} + c_i < \bar{p}_t$. Retailers simultaneously set their prices with this search behavior in mind.

A retailer, when setting its price, considers how its effect on the minimum and average prices incentivizes searching, and attempts to maximize its profits. In a given period, a retailer earns expected profits equal to

$$E(\pi_{jt}(p_{jt})) = p_{jt} N_c (1 - F(p_{jt}|V_t^*)) \left[\frac{\alpha_t}{\beta_t N_s} I(p_{jt} = \min_j\{p_{jt}\}) + \frac{1 - \alpha_t}{N_s} \right] \tag{1}$$

where α_t is the share of consumers who choose to search, conditional on $\min_j\{p_{jt}\}$ and $\bar{p}_t$, β_t is the share of retailers setting their price at $\min_j\{p_{jt}\}$, and $I(\cdot)$ is an indicator function. Marginal costs are normalized to zero.

As long as some consumers search ($\alpha_t > 0$), retailers choosing to set their price at the minimum will be in Bertrand competition and will set prices at the market-clearing level. Further, assume that, while retailers may sell multiple units, restocking only occurs between periods,[1] and so the market-clearing price will be set by the lowest consumer valuation rather

---

[1]As described in the next section, for each individual good in the empirical setting of this paper, competitive supply is almost perfectly inelastic, as a result of a production process that generates many outputs and is insensitive to the price of any single output. The restocking assumption stands in here for that inelasticity, allowing the market-clearing price, and differences in market-clearing prices across goods, to be set by demand.

than marginal cost. There is no incentive to sell below the market-clearing price, since all units can be sold at the market-clearing price. Retailers choosing to set their price above the minimum face no direct competitive pressures. High prices are restrained only by the demand curve and the possibility of more searching.

As long as γ and c_H are large enough, there exists a separating equilibrium in which low-search-cost consumers search and high-search-cost consumers do not search ($\alpha_t = \gamma$), while some stores set market-clearing and others set high prices ($0 < \beta_t < 1$). In other words, price dispersion in equilibrium.

The price dispersion equilibrium occurs when $\min_j\{p_{jt}\} + c_H = \bar{p}_t$, at least two retailers are at the minimum price, and no retailer is above its monopoly price. At this equilibrium point, consumers' best response is for those with low search costs to search, and those with high search costs to not search. No retailer will want to raise their price, since that would increase $\bar{p}_t$ just enough to induce consumers with high search costs to search, removing all sales for retailers priced above the minimum (including the defector). Additionally, no retailer will want to lower their price, since at or below the monopoly price, marginally lowering the price would reduce profits.

The condition $\min_j\{p_{jt}\} + c_H = \bar{p}_t$ relies only on the minimum and average prices, and the minimum and average alone cannot identify a single price distribution. There are many possible price distributions that satisfy the equilibrium condition.

This stage game is expanded to a two-period game with menu costs. In the second period, retailer j incurs an additional cost c_j^M if $p_{j2} \neq p_{j1}$. Retailers then choose prices in the first and second period to maximize expected profits

$$E(\Pi_j(p_{j1}, p_{j2})) - c_j^M I(p_{j2} \neq p_{j1}) \equiv E(\pi_{j1}(p_{j1})) + E(\pi_{j2}(p_{j2})) - c_j^M I(p_{j2} \neq p_{j1}). \quad (2)$$

If $V_1^* \neq V_2^*$, the market-clearing and monopoly prices will differ between periods 1 and 2, and so there is an incentive to change prices. Some retailers will have menu costs low enough

7

that they can be effectively ignored. Other retailers, for whom menu costs are high enough that $\max_{p_{j1}, p_{j2}} \{\Pi_j(p_{j1}, p_{j2})\} - \max_{p_j} \{\Pi_j(p_j, p_j)\} < c_j^M$, will choose one price to maximize the expected sum of profits. High menu costs also give retailers the ability to commit to a price in the second period, which other retailers must respond to.

We now consider intuitively how dispersion in prices may change between periods 1 and 2. Consider any set of first-period equilibrium prices with some nonzero amount of dispersion $P_1 = \{p_{11}, p_{21}, ..., p_{N_s1}\}$. P_1 will satisfy $\min_j \{p_{j1}\} + c_H = \bar{p}_1$. Assume that some retailers will not change prices in period 2 due to menu costs.

If $V_2^* > V_1^*$, then retailers that do not change their prices and had prices set at the market-clearing level will sell any remaining inventory briefly before dropping out of the market, since they will then be below the market-clearing level; retailers do not stay in the market below the market-clearing price because of restocking frictions, as previously noted. Retailers who change their prices will set prices either at or above the market-clearing level such that $\min_j \{p_{j2}\} + c_H = \bar{p}_2$ after the sub-market-clearing retailers run out of inventory. If the spread of the high-priced retailers is similar in period 2 to that of period 1, then the overall spread of observed sale prices in period 2 will be greater due to the presence of some sales by the sub-market-clearing retailers. Also, if $V_2^* - V_1^*$ is large, we would expect any remaining retailers with unchanged prices to have relatively low prices, allowing price-changing retailers to have very high prices without incentivizing more search, further increasing the spread of prices. However, since any distribution of high-priced retailers with the same average is an equilibrium, increased price dispersion in the second period is not guaranteed.

If $V_1^* > V_2^*$, the highest menu cost retailers will be left with prices in period 2 that are higher, relative to the market clearing level, than in period 1. These committed high prices will force price-changing retailers to choose market-clearing or low above-market-clearing prices in the period 2 equilibrium. As long as these committed prices are indeed relatively high in period 2, the commitment to high prices will push the distribution of

8

above-market-clearing prices wider as flexible firms choose prices nearer the market-clearing level to maintain $\min_j\{p_{j2}\} + c_H = \bar{p}_2$.

To show that the model can actually produce such equilibria, we present a simple specification of the model in which $f(x|V_t^*) = Unif[V_t^*, V_t^* + 5], V_1^* = .5, V_2^* = 0, c_L = 0, c^H = 1, \gamma = .5$, and $N_s = 4$ such that $c_1^M = c_2^M = c_3^M = 0, c_4^M = \infty$. We then suggest the following equilibrium: $\alpha_1 = \alpha_2 = \gamma, P_1 = \{.5, .5, 2.5, 2.5\}, P_2 = \{0, 0, 1.5, 2.5\}$. Here .5 and 0, respectively, are market-clearing prices, set by demand, keeping in mind the restocking restrictions earlier laid out. Note that the average price is 1.5 in the first period and 1 in the second period, so that $\min_j\{p_{jt}\} + c_H = \bar{p}_t$ in both periods. In period 2, retailer 4 has already committed to a price of 2.5, which is the monopoly price in period 2, and the other retailers know it. This forces any other high-priced retailer to take a relatively low above-market-clearing price in order to maintain the equilibrium condition. The equilibrium described here has increasing price dispersion. P_2 has a larger coefficient of variation than P_1, as well as a higher standard deviation and range.

From the intuition surrounding this model we predict that, in a market in which menu costs are present and demand shifts are routine, we should see price dispersion increasing over time. This is not a unique equilibrium, however, and so increasing price dispersion is unlikely to be a universal result.

In the next section we describe a group of markets that are similar to the model described.

III. SETTING

Our study setting is the online secondary market for Magic: the Gathering (MTG) cards. MTG is a collectible card game introduced in 1993 by Wizards of the Coast, currently a subsidiary of Hasbro Inc.. MTG is played like a standard card game, but each player builds a deck from their own collection of cards. According to Wizards of the Coast, the game has

an estimated 12 million players internationally.[2]

The focus of our study is on the game's secondary market for cards.[3] With few exceptions, Wizards of the Coast only sells cards in randomized "packs." Consumers do not have a choice of which particular cards they will get. In order to obtain a particular card, they must find it on the secondary market or get lucky when opening a pack.

A sizable secondary market has grown around the buying and selling of MTG cards to address player demand for particular cards. Retailers in the market range from informal marketplace tables at local events to comic and card shops to large online stores selling thousands of cards per day. The identical nature of each copy of the cards and the existence of the mature secondary market has made MTG cards amenable to the economic study of online auctions in previous work (Lucking-Reiley, 1999; Reiley, 2006).

The secondary MTG market is of interest to a study of how novel goods are priced because it offers an unusual combination of three properties:

(1) The Magic secondary market is highly competitive. There are a large number of retailers and consumers. We observe 909 online stores in the sample, and there are many more local stores, informal sellers, and small online stores not in the sample. More importantly, products are nearly identical across retailers (assuming the condition of the card is identical; this study only looks at undamaged "Near-mint/Mint" cards), with differences in shipping policy, the ability to keep inventory in stock, and reputation as the only sources of product differentiation. Wizards of the Coast makes no direct attempt to control prices in the secondary market, such as by setting an MSRP. These features mean that findings can be more easily interpreted as a result of the modeled mechanism as opposed to other market imperfections.

(2) New products are regularly released into the market. Wizards of the Coast releases

---

[2]http://magic.wizards.com/en/content/history

[3]In particular, we focus on the secondary market for physical cards. There is also an online version of the game that uses digital cards. The secondary market for digital cards is separate, and cards from one market cannot be shifted into the other within the time frame we examine.

a "set" of new cards roughly every three months. Each set contains 150-250 unique cards, most of which have never been seen before, and which can be purchased in randomized packs. The constant release of new products directly into a competitive market means that it is feasible to observe a large sample of early-stage markets as information is revealed about the products. Additionally, since information about the next set is revealed on a mostly predictable time frame, it is possible to isolate changes in the information set.

(3) Demand is likely to shift soon after release for a few reasons. First, part of the value of a given card is its novelty. Players are generally excited to try out new cards, but novelty naturally wears off.

Additionally, the new products are of uncertain quality. "Quality" here refers to the desirability of the card on the part of players. Cards may be desired because they are powerful in the game, because they fit into a certain theme that players enjoy, because they are novel and interesting, or because they carry attractive art or depict a popular fantastical creature such as a dragon.

Desirability varies widely from one card to the next. Some cards are nearly worthless on the market, and can be had for pennies. Other new cards routinely sell on the secondary market at prices upwards of $50 apiece.[4]

Certain features of desirability, such as whether a dragon is depicted, are observable. Others, especially the strength of a particular card in the game, are somewhat uncertain. Since cards are never played in isolation, the sheer complexity of the game (in which each card interacts with 149 others chosen from a pool of thousands) ensures that the strength of a particular card will be difficult to judge perfectly without experience. The quality of a particular card may be revealed through experience, between players by word of mouth, or by market observations of the price of the card. The fact that card quality can be revealed over time through experience sets the MTG market apart from some other markets with

---

[4]Other cards, such as premium "foil" versions of cards, can go for more than this. Older cards, which are out of print and scarce, can fetch very high prices. A Black Lotus card, last printed in 1993, can sell for more than $15,000 in good condition.

goods of uncertain quality, like the auction market for broadband spectrum, where quality revelation occurred after sales were completed (Cramton, 2013).

Continuing from (3), the MTG market has the convenient quality that, while the total supply of cards increases over time as more product is purchased, the competitive supply for any particular card at any particular time is highly inelastic. Differences in average price between cards at a given time can be interpreted as demand-side differences in perceived quality. In MTG, the standard unit of sale is a "pack," which contains 15 randomized cards from a single set of 150-250 cards. Consumers and retailers cannot buy particular individual cards directly from Wizards of the Coast. This means the production process for each card in the same set is tied together.

As a result of the pack system, the supply of any individual card should be very inelastic. Cards are produced by opening packs, each of which contains exactly one "rare" or "mythic" card. A given particular "rare" card will only appear in about 1.5% of packs, or .8% for a given particular "mythic." So if the price of a given card increases, the effect on the expected value of opening a pack will only be about 1% as large. The price increase will then only have a small impact on the number of packs opened, and thus the quantity supplied of that card.

In summation, while price changes over time may reflect changes in available quantity supplied, as it takes time to actually open and distribute the cards, these changes in available quantity are constant across all cards in a given rarity and set release at a particular time. Accounting for these characteristics, remaining changes can be attributed to shifts in demand.

A final advantage of the MTG market is that there is a large contingent of online stores that make their prices apparent to observers. These prices can be easily observed by the researcher. Data from these online stores make up the sample in this study.

IV. DATA

We use a sample of data on Magic: the Gathering (MTG) card sales from 909 online stores collected by MTGPrice.com between May 10, 2012 and July 27, 2014. The stores we observe are those registered with a large online price comparison portal. Consumers who shop directly through the portal are able to observe comparison prices for a given card, although they may also visit the store websites directly, in which case they would not observe comparison prices. MTGPrice software collects daily records of the sale price and available quantity of each card listed on the portal by each store.

From MTGPrice data on the listed price and available quantity, sales are identified from instances in which the quantity of the card available at a particular store decreases from one day to the next. This approach, rather than relying solely on list prices, ensures that we only observe actual sales, an important distinction to ensure that any dispersion is not simply a result of very-high priced stores that never actually make sales (Ghose and Yao, 2011). However, this will not capture sales in which the store replenished its stock before data was collected again, or sales from stores that censor high available quantities, so some sales are unobserved.

Variables in the data include the name of the card sold, the set it is printed in, the store, the date of sale, the number of copies sold on that day, and the price of sale. Data also include the condition of the card for sale, and we limit all analysis to cards graded as being "Near Mint/Mint" condition. We augment the data with information about which cards are reprints or available as promotional giveaways.

In order to focus solely on cards that are newly released and for which product quality is most uncertain, we track cards only from one week before the cards are first made available by Wizards of the Coast (when information about cards is made available and many stores sell preorders) to four months after release. We study "rare" and "mythic" cards from nine new sets of cards released in the sample window. Cards that are not "rare" or "mythic" typically command prices low enough that transaction costs are likely to swamp actual value. We also

exclude premium "foil" versions of cards, since these are rare enough that the markets for these cards are thin.

In total, we track the sales of 537 different cards, of which 698,319 copies are sold. Most analyses focus on the time window between release and when cards from the next set begin to be revealed, in which 430,521 copies are sold. We limit the sample frame to the time before cards from the next set begin to be revealed because the introduction of new cards that are complements or substitutes of old cards changes the underlying true value of old cards and as such we would expect a structural break in the way cards are priced. The length of time before the next set is revealed is on average 14.4 weeks, but is as long as 19 weeks in one case, longer than the four-month observation window.

The mean price of a sold copy of a card is $3.85. More expensive cards are more heavily traded; taking the mean price for each card over the sample window, the median of those average prices is $1.56. Sales in general are concentrated among the most popular cards but the concentration is not extreme: the most popular 136 of the 537 cards make up half of all sales, and while the mean number of copies sold per card was 802, the median was 682.

10.4% of sales are of reprinted cards, which are already available through a previous printing when re-released in a new set in our sample window and so might be subject to somewhat less uncertainty of quality (although their interactions with new cards would still be untested). A further 17.3% of sales are of new cards that are also available in some limited form outside of packs within the sample window, such as through promotional giveaways or as a part of non-randomized packs sold directly by Wizards of the Coast. These promotional cards may be subject to a less inelastic supply curve than is described in the previous section. 74.2% of sales are of cards that are only available in the new packs. Results are similar whether or not reprints and promotional cards are included, and so all reported analyses include all cards.

Much of the analysis focuses on the dispersion in prices, and when stores decide to change prices. 12.3% of copies are sold at a different price than the previous copy at the

same store, indicating a price change. The average coefficient of variation between stores in prices for a given card in a given week is .258. We use the coefficient of variation as a measure of dispersion because it allows dispersion to be compared across goods of different average prices. Some of the variation in prices may be due to differences across stores in other amenities such as customer support or shipping prices. In the results section I show that results are robust to the use of dispersion net of store effects.

We use this detailed data set of online sales to examine the dynamics of the secondary market for MTG cards. We begin by looking for changes in price dispersion over time.

V. RESULTS

V.i. Price Dispersion

Table 1 depicts the results of regressions predicting the coefficient of variation (standard deviation divided by mean) of a card's price across sales in a given week, starting the week before release and finishing before the cards from the next set start to be revealed. The coefficient of variation is a useful statistic capable of comparing dispersion across goods with different average price levels. Estimates come from the least-squares regression equation

$$CV_{it} = \alpha_i + \beta_1 t + \beta_2 t^2 + \varepsilon_{it} \tag{3}$$

where CV_{it} is the coefficient of variation, α_i is a fixed effect for each card, and t is the number of weeks since the card's release. If some types of cards happen to have higher baseline levels of dispersion, the fixed effect accounts for this, allowing β_1 and β_2 to isolate within-card changes over time.

Model 1 in Table 1 shows the general behavior of price dispersion across stores over time. This model uses one observation per card per week, considering dispersion across all sales in a given week. This reduces small cell size issues if the observation level is daily, although the

15

relationships in this section all hold if observations are at the daily level. Here we establish the empirical result that price dispersion is increasing over time, which we will in the next section link to the model from Section II. The results in Table 1 clearly show increases in the coefficient of variation over time.

Table 1: Price and Price Dispersion Over Time

	(1) Coef. Var.	(2) Coef. Var. (Store-Adjusted)	(3) Coef. Var.
Weeks since release	.009***	.027***	.003*
	(.001)	(.003)	(.002)
(Weeks since release)^2	-.0002**	-.001***	.0001
	(.0001)	(.0002)	(.0004)
Lagged Coef. of Variation			.177***
			(.016)
Card fixed effects	Yes	Yes	Yes
N	4,984	4,984	4,359
R^2	.467	.544	.490

Note: */**/*** indicates statistical significance at the 10%/5%/1% level. Sample size is smaller in model 3 because some observations are missing a lagged coefficient of variation.

The growth in dispersion is not linear, and Models 1 and 2 have negative and significant quadratic terms. However, these quadratic terms are meaningfully small. The sample window extends to only about 15 weeks (varying by set), and so by the end of the sample window, the speed of growth has only dropped by about a third. One interpretation of the concave growth is that this is consistent with smaller or less frequent demand shifts over time as quality is fully revealed, but the data are not sufficient to pin down this interpretation.

The presence and upward trend of price dispersion is not solely a result of store effects, such as differences across stores in reputation, fully stocked inventories, or favorable shipping policies. We calculate the coefficient of variation again, taking into account store effects,[5] and find an average coefficient of variation of .361, actually larger than the base level. In

---

[5]We regress price on a set of store fixed effects and card fixed effects. We subtract the store fixed effects from price and calculate standard deviation using these residuals. To get the coefficient of variation we divide this adjusted standard deviation by the original mean price.
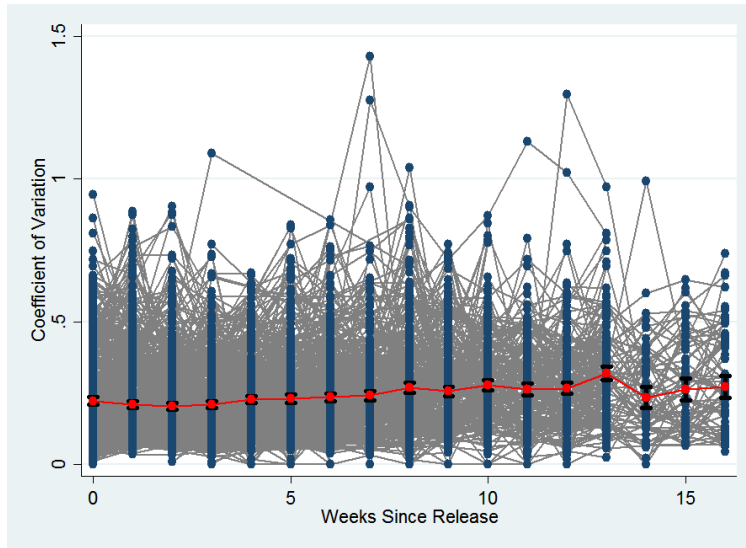
Model 2 we estimate Equation 3 using this adjusted coefficient of variation, finding a much steeper climb over time. Notably, for both the average level of dispersion and the increase over time, if the standard deviation is used instead of the coefficient of variation, there is no meaningful difference between the store-adjusted and unadjusted measures, rather than more dispersion with the store-adjusted measure. We prefer Model 1 to Model 2 for ease of interpretation and because some store fixed effects may be poorly estimated due to small cell sizes, throwing off results. In general, we find that differences in store quality and reputation do not explain price dispersion, and that price dispersion here must to some degree rely on stores with similar average prices across goods charging especially high or low prices for particular goods, or "relative price dispersion" as outlined by Kaplan et al. (2016). In a model of relative price dispersion, dispersion arises because stores take advantage of consumers' preference to purchase multiple goods in the same store.

Prices in this market tend to drop over time, which will be explored further in a later section. This average drop in price does not explain increasing price dispersion. Since the coefficient of variation has the price mean in the denominator, a constant price spread could translate to increasing dispersion if average prices drop. However, results are robust to the use of the standard deviation rather than the coefficient of variation. Additionally, results are still present if the sample is limited only to cards that increased in average price by more than 5% from the beginning to the end of the sample.[6]

We show two less parametric forms of this relationship. In Figure 1 we contrast average price dispersion by week since release with individual-level measures of dispersion by week. There is an overall slight upward trend in average price dispersion, shown by the thick central line, similar to Model 1 in Table 1. The mean level of dispersion is rising, and the mean in any given week is precisely estimated enough, shown by the confidence intervals, that the

---

[6]Growth in the standard deviation (store-adjusted or not) and the store effect-adjusted coefficient of variation is actually faster within cards that increased in price. Growth in the unadjusted coefficient of variation is somewhat smaller, although still positive (the coefficients on week are .004 for the linear term -.0002 on the quadratic). Unadjusted coefficient of variation growth is insignificant at the 5% level, although this is partially due to the smaller sample size; at the original sample size, growth would still be significant.
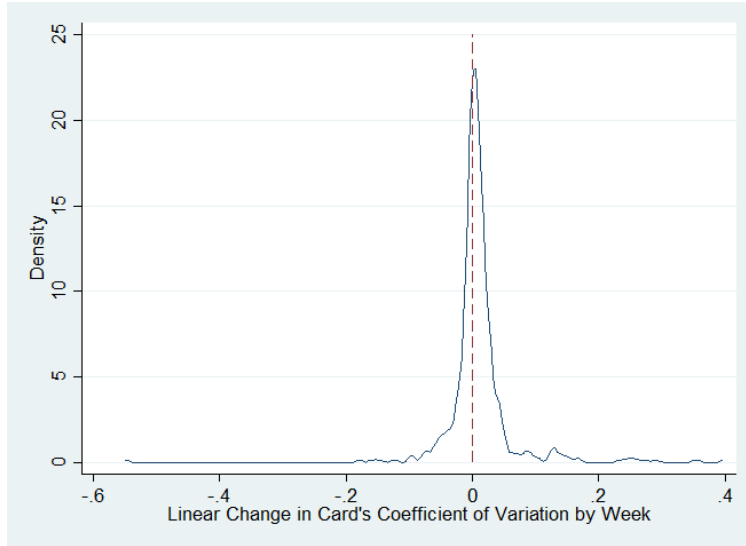
Figure 1: Average Price Dispersion by Week

upward trend is statistically discernible. The figure also shows that price dispersion is not limited to a small number of highly dispersed cards; many cards are above the average price dispersion level.

Changes in dispersion over time can be difficult to see here for individual cards, and so for Figure 2, we run a separate regression of the coefficient of variation on weeks since release separately, specified linearly, for each card. We then show the distribution of this individual linear slope across all cards. A roughly symmetrical distribution emerges, centered above zero. More than half of all cards increase in dispersion over the sample window (63.6%). Among cards that increased in price from the beginning to the end of the sample, it is slightly less than half (49.6%). As predicted by the model, increasing price dispersion over time is not the rule of the market, but rather something that occurs for a significant and non-ignorable portion of the sample.

Figure 2: Distribution of Linear Dispersion/Week Relationship by Card

The presence of temporary spikes in Figure 1 raises the question of how dispersion appears and erodes for given cards. If dispersion of a given card is completely uncorrelated across time, then the overall upward trend in dispersion is dependent entirely on the distribution of dispersion-causing shocks across time (which, when isolated to a single card, is the result of demand shocks, as discussed in III) rather than market structure, and so any market stucture-based explanation is likely to be wrong. On the other hand, if the upward trend disappears once persistence in dispersion is accounted for, this would discount the proposed menu-cost based explanation, since for these new goods demand should continue to gradually shift as information is revealed and novelty is lost, and so dispersion should continue to rise. In Model 3, again in Table 1, we add to Equation 3 a one-week lag in CV_{it}.

Within a given card, if the coefficient of variation is one unit higher than normal in one week, the market for that card is expected to have a coefficient of variation .177 higher than normal the next week. The coefficient is less than one, indicating that markets do engage in self-correction; a very large increase in dispersion for a given card should be expected to be largely corrected by the next week, as we see in the high temporary spikes in Figure 1. This is not too surprising given a menu cost explanation, as large shocks to dispersion would be the time one would expect the benefit of changing price to outweigh the cost, and correction

19

would occur.

Additionally, the coefficient is more than zero, and as such shocks to price dispersion would be predicted to persist for some time. Overall trends in dispersion are not simply driven by the distribution of momentary spikes, but instead there is underlying persistence in dispersion. Finally, even accounting for persistence, a small upward trend in dispersion remains. These results do not isolate a menu costs explanation specifically, but they do fail to produce results inconsistent with menu costs.

The appearance, persistence, and increase over time in price dispersion is established for the secondary MTG market in this section. In the next section we link this result to the menu cost model outlined earlier.

V.ii. Model Conditions and Predictions

Thus far, we have empirical evidence for increasing price dispersion over time on average and for a significant portion of the goods under examination. We also have a model relying on the preconditions of uncertain and changing perceived quality, menu costs, consumer search costs. If this explanation of increasing price dispersion is accurate, we should expect to see evidence that the market setting is as described in Section II.

Here we look for evidence of changing perceived quality and menu costs. Menu costs, if present, should predict that: if prices are in general decreasing (increasing) over time, then retailers with high menu costs should have on average higher (lower) prices. Retailers with high menu costs should be less likely to change their prices, and should change by larger absolute amounts when they do. Finally, as a test of the consumer search model in general, we consider that search costs should be identical across different goods, and as a result the coefficient of variation should be lower for markets with higher average prices.
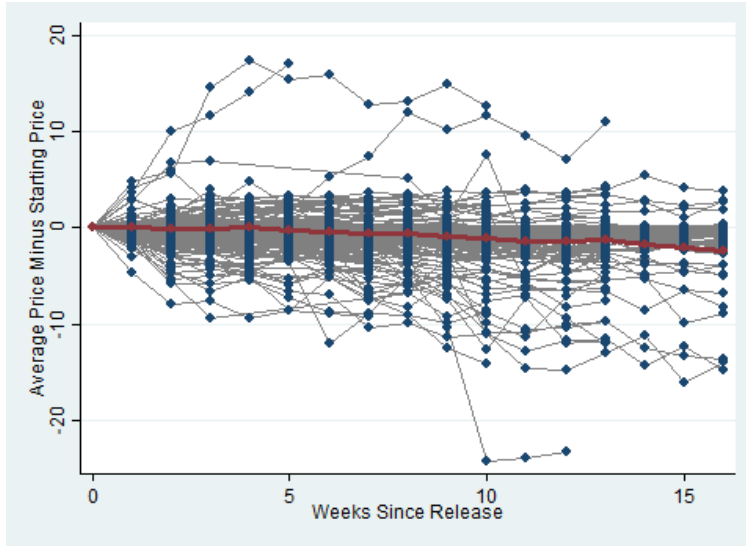
V.ii.1. Shifting Demand

As described in Section III, since the supply of any given card is heavily inelastic with respect to its own price, changes in average price can largely be interpreted as being due to shifts in demand, in addition to a shared supply increase over time for all cards, as retailers have more time to open and sort product. Within a relatively limited time window, changes in price for a given card, relative to other cards experiencing the exact same supply shift, are evidence of shifts in demand rather than supply.

In Figure 3 we plot the average price of each card by the number of weeks since release. Most cards change in average price over the sample window. While some cards climb in price, some significantly so, these are eclipsed by the large number of cards that lose value over the sample window. On average, card prices drop; a linear regression of card price on weeks since release, with card fixed effects included, suggests that each card's price drops by an average of eight cents per week. The overall drop could be attributed to the shared supply increase over all cards or a shared leftward shift in demand, perhaps due to loss of novelty value. However, variation in price changes between cards in the same set cannot be attributed to supply shifts, since supply shifts would affect all cards of the same set, and instead indicate changes in demand (which, as previously described, may include revelation of quality through experience or learning, or a loss of novelty value).

Over all cards, 23.7% increase in price by more than 5% from the beginning of the sample window to the end, 65.2% decrease in price by more than 5%, and the remaining 11.0% stay close to the price they started at. Figure 4 shows the distribution of percentage price changes from the beginning of the sample window to the end. Variation across cards in price changes over time, in a market where supply shifts are shared across cards in the same set, establishes the presence of the card-specific demand shifts necessary to drive increasing dispersion in the model.
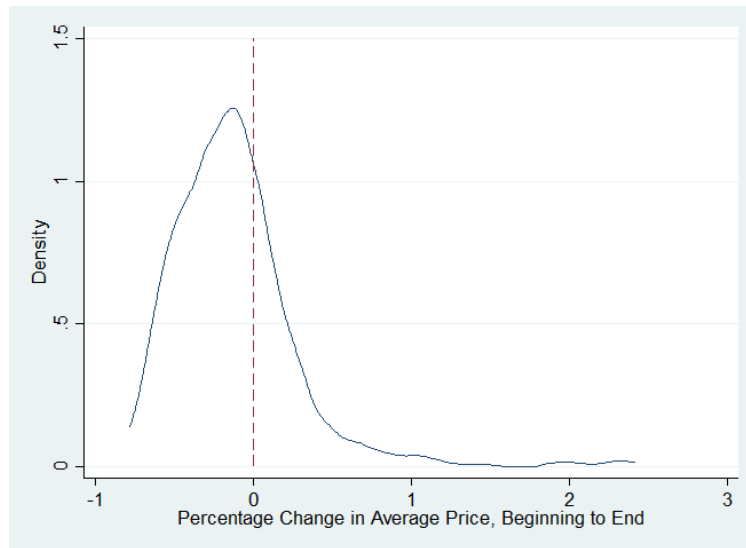
Figure 3: Average Price by Week

V.ii.2. Menu Costs and Store Qualities

Given that average prices change significantly for most cards, in the absence of menu costs we would expect most stores to frequently change their prices to match demand. However, many cards never see price changes at some stores. For each store/card-specific listing we examine whether or not the card ever changes price. Omitting listings that only have sales on one day, fully 55.9% of listings never see a price change.

Stores that change prices less often are not completely unable to do so. Listings that do not change price are more often for cheaper cards, as one would expect given that menu costs should be consistent across cards, but the potential benefits to changing price may be larger for more expensive cards. The average card price among listings that do not change price is $3.11, as compared to $4.38 for those that do change. Listings that do not change price are also on the market for less time: there are on average 22.8 days between the first and last sale for these listings, as compared to 42.8 days for others. The shorter selling spans make sense in the presence of changing average prices - those who have too-low prices will be quickly sold out, and those with too-high prices will cease to see any sales.

If store-specific menu costs are the reason for the lack of price changes, we should expect to consistently see more change at some stores than at others. We regress the no-price-change

22

Figure 4: Distribution of Average Price Changes Over Sample Window

indicator on a set of store fixed effects using a linear probability model. The F statistic for the fixed effects is a significant 8.56.

Testing the above prediction that stores with high menu costs should change prices less often, and change by larger amounts when they do so, we compare stores on the basis of their size. Small stores are more likely to have to change prices by hand or to coordinate their prices with a brick-and-mortar storefront, and should have higher menu costs. Taking the total number of card copies sold at a store over the sample window as the store's size, the average store size for a listing that never saw a price change is 2,665, as compared to 6,339 for listings that changed price. For the ten largest stores, 16.8% of cards are sold at a different price than the previous sale, and the average absolute price change is 43.0 cents, while for stores outside the top 100, 5.8% are sold at a different price than the previous sale, and the absolute average price change is 54.9 cents. We split at the top 10 because there is a minor break in the size of stores at 10. Top 100 is arbitrary, but analysis is robust to the use of other cutoffs.

V.ii.3. Price Differences for Menu Cost Constrained Retailers

The model outlined in Section II suggests that increasing price dispersion is a result of menu cost constrained retailers setting prices that are intended to hold across multiple periods. As the demand curve shifts, high menu cost retailers are left behind and are likely to have prices better suited for old valuations.

The data conforms to expectations. Overall, prices drop over time. We regress card price on an indicator for whether that listing ever changed price and card fixed effects. Controlling for card identity (as opposed to the comparison in the previous section showing that listings that never changed price were more commonly for cheaper cards), listings that never changed price are on average 9.2 cents more expensive than those that changed. We can see whether this effect holds in both directions by splitting the sample by how much the card's average price changed from the beginning of the sample to the end. Among the cards that saw price decreases of 5% or more, listings that never changed price are 14.1 cents more expensive. Among the cards that saw price increases of 5% or more, listings that never changed price are 2.8 cents less expensive. Among the cards that stayed close to their initial price, listings that never changed price were an insignificant .1 cents more expensive.

V.ii.4. Price Dispersion by Average Price

Price dispersion in equilibrium relies on search costs among consumers. If search were truly frictionless, there would be little reason for any consumer to pay anything but the minimum possible price, and dispersion would only be possible through increased service or bundling. However, it is likely that consumer search costs are identical across cards. In an online setting, it costs a consumer as much to find the cheapest supplier of a card that costs on average $1 as it does to find the cheapest supplier of a card that costs on average $50. We should expect less dispersion (measured by the coefficient of variation) among more expensive cards, since a high standard deviation relative to the price indicates larger price differences for expensive cards, and even high search cost consumers would be willing to search if there

24

were such large price differences.

We see evidence of this prediction in the data. We modify Equation 3 by adding the card's average price in that week as an explanatory variable. The coefficient on average price is a negative -.016. In weeks where a card is more expensive, the market for that card exhibits lower levels of variation.

This model, in which coefficient of variation is regressed on price, suffers from endogeneity, since price appears on the left and right hand sides. So, we repeat the analysis using standard deviation alone instead of the coefficient of variation and average price again comes in as negative, with a coefficient of -.041. The addition of mean price does not greatly affect the coefficient on weeks since release.

We are also interested in whether cards that are generally more expensive show lower levels of variation. For this we perform the analysis with standard deviation again, leaving out the individual card fixed effects. The coefficient using standard deviation as the measure of dispersion is .131. Standard deviation is higher for more expensive cards, but still standard deviation rises more slowly than price, in line with what would be expected if consumers faced constant search costs across goods of different average price levels.

VI. CONCLUSION

Consistent with other studies of online markets, we find evidence that price dispersion in the online secondary market for Magic: the Gathering cards exists and is persistent. Further, we show that price dispersion often increases in the months following the release of a product.

Our focus on the dynamics of price dispersion in the market add another layer to studies of dispersion, which commonly focus on static analysis and the presence of dispersion in equilibrium. While we are not the first to approach dispersion as a dynamic phenomenon (Verboven and Goldberg, 2001; Brown and Goolsbee, 2002; Pan et al., 2003; Ratchford et al., 2003; Baye and Morgan, 2004; Doraszelski et al., 2016), we do add emphasis on understanding

how commonly reoccurring changes in markets are likely to predict changes in dispersion. This study proposes a mechanism that leads to the counterintuitive phenomenon of price dispersion increasing over time.

We predict that increasing price dispersion can occur in a market with consumer search costs, shifts in demand, and menu costs. We see evidence that each of these factors are present in the market for Magic cards, and predict several features of the market that are also present, including increasing price dispersion.

The presence of menu costs in the market is an interesting basic result on its own. Static analyses of dispersion do not consider whether it is likely to arise, and basic economic logic as applied to situations like arbitrage suggests that if dispersion is changing, it should be declining towards some equilibrium level. But this is not what we see.

The mechanism of interest, where increases in price dispersion occur as the result of shifting demand interacting with menu costs, should apply whenever meaningful menu costs are in place. Consumer search costs are an nearly unavoidable feature of the shopping experience, even if they are only small. Online stores might be expected to have negligible menu costs, since it is possible to automate price setting. However, the paper offers a number of pieces of descriptive evidence in favor of the existence of menu costs in the market. Given consumer search costs and menu costs, the model requires only the addition of demand shifts, which are likely to occur for new goods of uncertain quality and are observed in the market, in order to produce increasing dispersion.

Unfortunately, the data used in this paper is capable of providing only suggestive evidence of the mechanism of interest, and only so much can be done to rule out alternative explanations. For example, rising dispersion would also be consistent with a price discrimination model like the one proposed for the online book market by Ellison and Ellison (2014). Following their model, stores with both brick-and-mortar and online presence might exhibit behavior that looks like menu costs but is instead price discrimination, setting high online prices hoping to find rare high-value buyers. Adding this insight to the demand shifts in

the present paper, this behavior would, like menu costs, be capable of generating increases in price dispersion. It is unlikely that the Ellison and Ellison (2014) approach applies in the Magic market, since newly released Magic cards are far less numerous in terms of the number of product categories, and are more generally available, than long-tail books, and so the match-quality gains that drive their model are not as relevant.[7] Additionally, when average card prices increase, unchanging listings have lower prices, as shown in Section V.ii.3. This is hard to square with their model. Still, this paper cannot rule out all other potential mechanisms, and some of those other mechanisms may apply better than menu costs in other markets.

While the market for Magic cards is not of direct interest to economists, affirmative empirical results here suggest applying the same analysis to other markets where menu costs likely exist, such as the collectible coin market or the wine market. Looking for similar behavior in a wider range of markets is a clear direction forward, for the purposes of examining both the menu cost mechanism and other potential mechanisms like Ellison and Ellison (2014), and also in general looking more at the empirical short- and long-run dynamics of price dispersion. However common are the types of markets in which the menu-cost mechanism operates, it stands that in this market we observe a rise in price dispersion over time that is consistent and repeated over many goods. This is an unexpected result that speaks directly to the way in which markets function and the ability of a largely competitive market to fulfill one of their vital functions of offering low prices to consumers.

---

[7]A direct test of their model in this data could be performed by comparing stores with brick-and-mortar equivalents to those without. Unfortunately, in the Magic market many very large online stores also have a small brick-and-mortar representation. And so the comparison would need to be based on whether the firm is *primarily* online or brick-and-mortar, rather than simply whether they have a brick-and-mortar store. The data necessary to establish this is unavailable.

VII. ACKNOWLEDGMENTS

VIII. REFERENCES

Baye, M. R., Morgan, J., 2004. Price Dispersion in the Lab and on the Internet: Theory and Evidence. The RAND Journal of Economics 35 (3), 449.

Baye, M. R., Morgan, J., Scholten, P., 2004. Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site. Journal of Industrial Economics 52 (4), 463–496.

Baye, M. R., Morgan, J., Scholten, P., 2006. Information, Search, and Price Dispersion. In: Handbook on Economics and Information Systems. pp. 323–371.

Baylis, K., Perloff, J. M., 2002. Price Dispersion on the Internet: Good Firms and Bad Firms. Review of Industrial Organization 21 (3), 305–324.

Bergemann, D., Välimäki, J., 1996. Learning and Strategic Pricing. Econometrica 64 (5), 1125–1149.

Brown, J. R., Goolsbee, A., 2002. Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry. Journal of Political Economy 110 (3), 481–507.

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

Chevalier, J., Goolsbee, A., 2003. Measuring Prices and Price Competition Online: Amazon.com and BarnesandNoble.com. Quantitative Marketing and Economics 1 (1998), 203–222.

Coey, D., Larsen, B., Platt, B., 2016. A Theory of Bidding Dynamics and Deadlines in Online Retail.

Cramton, P., 2013. Spectrum Auction Design. Review of Industrial Organization 42 (2), 161–190.

De los Santos, B., Hortaçsu, A., Wildenbeest, M. R., 2012. Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior. The American Economic Review 102 (6), 2955–2980.

Dinerstein, M., Einav, L., Levin, J., Sundersan, N., 2017. Consumer Price Search and Platform Design in Internet Commerce.

Doraszelski, U., Lewis, G., Pakes, A., 2016. Just Starting Out: Learning and Equilibrium in a New Market.

Einav, L., Kuchler, T., Levin, J., Sundaresan, N., 2015. Assessing Sale Strategies in Online Markets Using Matched Listings. American Economic Journal: Microeconomics 7 (2), 215–247.

Ellison, G., Ellison, S. F., 2009. Search, Obfuscation, and Price Elasticities on the Internet. Econometrica 77 (2), 427–452.

Ellison, G., Ellison, S. F., 2014. Match Quality, Search, and the Internet Market for Used Books.

Ghose, A., Yao, Y., 2011. Using Transaction Prices to Re-examine Price Dispersion in Electronic Markets. Information Systems Research 22 (2), 269–288.

Gorodnichenko, Y., Talavera, O., 2017. Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration. The American Economic Review 107 (1), 249–282.

Kaplan, G., Menzio, G., Rudanko, L., Trachter, N., 2016. Relative Price Dispersion: Evidence and Theory.

Keller, G., Rady, S., 2003. Price Dispersion and Learning in a Dynamic Differentiated-Goods Duopoly. The RAND Journal of Economics 34 (1), 138–165.

Lucking-Reiley, D., 1999. Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet. The American Economic Review 89 (5), 1063–1080.

Pan, X., Ratchford, B. T., Shankar, V., 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), 433–445.

Pan, X., Ratchford, B. T., Shankar, V., 2003. The Evolution of Price Dispersion in Internet Retail Markets. In: Baye, M. R. (Ed.), Organizing the New Industrial Economy - Advances in Applied Microeconomics. Vol. 12. pp. 85–105.

Ratchford, B. T., Pan, X., Shankar, V., 2003. On the Efficiency of Internet Markets for Consumer Goods. Journal of Public Policy & Marketing 22 (1), 4–16.

Reiley, D. H., 2006. Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet. The RAND Journal of Economics 37 (1), 195–211.

Rothschild, M., 1974. A Two-Armed Bandit Theory of Market Pricing. Journal of Economic Theory 9, 185–202.

Salop, S., Stiglitz, J. E., 1977. Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion. The Review of Economic Studies 44 (3), 493–510.

Salop, S., Stiglitz, J. E., 1982. The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents. The American Economic Review 72 (5), 1121–1130.

Smith, M. D., Bailey, J., Brynjolfsson, E., 2000. Understanding Digital Markets: Review and Assessment. In: Brynjolfsson, E., Kahin, B. (Eds.), Understanding the Digital Economy: Data, Tools, and Research. MIT Press, Cambridge, MA, pp. 99–136.

Varian, H. R., 1980. A Model of Sales. The American Economic Review 70 (4), 651–659.

Venkatesan, R., Mehta, K., Bapna, R., 2007. Do Market Characteristics Impact the Relationship Between Retailer Characteristics and Online Prices? Journal of Retailing 83 (3), 309–324.

Verboven, F., Goldberg, P. K., 2001. The Evolution of Price Dispersion in the European Car Market. The Review of Economic Studies 68 (4), 811–848.

Xing, X., 2008. Does Price Converge on the Internet? Evidence from the Online DVD Market. Applied Economics Letters 15 (1), 11–14.

Xing, X., 2010. Can Price Dispersion be Persistent in the Internet Markets? Applied Economics 42 (15), 1927–1940.