Accepted Manuscript

Title: Location Still Matters: Evidence from an Online Shopping Field Experiment

Authors: John Morgan, David Ong, Zemin Zachary Zhong

PII: S0167-2681(17)30325-6
DOI: https://doi.org/10.1016/j.jebo.2017.11.021
Reference: JEBO 4203

To appear in: *Journal of Economic Behavior & Organization*

Received date: 7-1-2017
Revised date: 20-10-2017
Accepted date: 23-11-2017

Please cite this article as: Morgan, John, Ong, David, Zhong, Zemin Zachary, Location Still Matters: Evidence from an Online Shopping Field Experiment.Journal of Economic Behavior and Organization https://doi.org/10.1016/j.jebo.2017.11.021

Location Still Matters: Evidence from an Online Shopping Field Experiment

By JOHN MORGAN, DAVID ONG AND ZEMIN (ZACHARY) ZHONG[1]

[1] Morgan: Haas School of Business, UC Berkeley, Berkeley, CA 94720-1900, (email: morgan@haas.berkeley.edu); Ong: HSBC Business School, Peking University, PHBS Building 708, University Town, Shenzhen, P.R.C, 518055 (e-mail: dvdong@gmail.com); Zhong: Haas School of Business, UC Berkeley, Berkeley, CA 94720-1900, (email: zachary_zhong@haas.berkeley.edu). This research was funded by the Dean's Fund No.2009039, Peking University Shenzhen Graduate School. The authors appreciate comments from David Reily, Fang-Fang Tang, Young-Joon Park, Yuk-fai Fong, Fuhai Hong, Jan Potters, Yan Chen, Anthony Dukes, the audience in the 2010 WISE International Workshop of Experimental Economics, 2011 Econometric Society North American Meeting, 2011 INFORMS Marketing Science Conference, 2011 Economic Science Association International Meeting, and seminar participants at Peking University, National Taiwan University, and Academic Sinica. The authors gratefully acknowledge Yuan (Yvette) Qiu and Wei (Vivian) Yang for excellent research assistance work.

Highlights:

- Many empirical studies show that sellers post different prices for homogeneous goods.
- However, seller heterogeneity is difficult to control.
- We sell identical simple goods from sellers that were identical except in name.
- Out of 514 sales, 73 were of the higher priced item.
- Price sensitive consumers do not necessarily buy the lowest priced item.
- We distinguish between search cost and limited attention based price dispersion.

ABSTRACT: Many empirical studies of online price dispersion show that sellers post different prices for homogeneous goods. However, seller heterogeneity is difficult to control for and posted prices may not reflect price dispersion in actual transactions. We contribute to this literature by selling identical simple goods (cell phone credits) at different prices from sellers that were identical except in name and with minimal ratings. The only way consumers could find us in this extremely thick market is to rank by price from lowest to highest. Out of 514 sales, 73 were of the higher priced item, for which we had non-negligible demand even when the price gap was 2.5%. Thus, even this selected sample of price-sensitive consumers do not necessarily buy the lowest priced item, all else being equal. Using *independent variation in screen location and price, we are able to distinguish for the first time between search cost and limited attention based price dispersion.*

Keywords: Price dispersion, field experiment, Internet, electronic market
JEL Codes: C93, D4, D8, M3, L11

I. Introduction

A visit to Amazon, eBay, or a host of price comparison sites will quickly confirm Hal Varian's famous quip (1980) that "the law of one price is no law at all." Sellers offer apparently identical items for wildly varying prices, with ranges often exceeding 100 percent. At first blush, this seems puzzling—why would some firms charge higher prices than others? The practical (though facile) answer is that firms charge more because (some) consumers will pay more. Yet this is hardly satisfying. Why do consumers pay more? How many consumers pay more? If consumers do pay more, then why don't all firms charge the higher price?

Before the rise of online markets, economists argued that such price dispersion was the product of search frictions. Some consumers were simply unaware of the lower prices and, at least in expectation, it was not worth their while to find out. Stigler (1961) was first to formalize this idea, and such approaches dominated economists' thinking until the late 1990s.

Arguably, the internet revolutionized consumer search. When a consumer could simply scan down a price list, or even sort the list with the click of a mouse, appeals to search frictions became less plausible. While a number of models, starting with Baye and Morgan (2001) offered rational explanations of price dispersion through sophisticated strategic behavior, simpler and more direct explanations ran along two main lines. The first line asserted that, although products appear identical, there are, in fact, myriad differences across sellers. An online consumer is not simply buying the product, but consuming the entire purchase experience, from ease of use of the seller's website and design of the shop webpage, to after sales service features like warranties, return policies, and restocking fees. From this perspective, price dispersion merely reflects the differing costs and willingness to pay for various types of purchase experiences. Sellers charging high prices simply offer a better experience than those charging lower prices. Moreover, some consumers value these aspects and are willing to pay for them.

A second line of argument reflected the growing incorporation of behavioral factors into economic models—in this case, limited attention. By this rationale, scarce consumer attention causes lower prices to sometimes go unnoticed. While it might seem unlikely that a consumer will overlook a low price in a setting where sorting by prices is simple, in practice, seemingly small costs, like clicking a mouse or scanning a page, prove large. In a structural model of search costs, Hong and Shum (2006) estimate that the (mainly psychological) cost of evaluating each price listing on an internet price comparison site exceeds $3. Market forces bear this out—competition is fierce to be the top listing following search queries on Google. The same is true at the price comparison site Nextag, in spite of the fact that the default search order is easily sortable (by price among other things) with a single mouse click. Likewise, featured listings on eBay constitute an important revenue stream despite a variety of easily used search and sorting tools that remove their privileged position. In effect, limited attention models represent a return to traditional search frictions explanations of price dispersion, albeit with psychological microfoundations now added. Moreover, as we show, even seemingly price sensitive consumers who rely on sorting tools do not always transact at the lowest price.

Distinguishing between these two competing lines proves difficult. Both predict price dispersion in posted offers and transactions. Both predict that higher prices will lead to fewer sales than lower offers. And both predict that price dispersion, both posted and transactional, will persist. Thus, without extraordinary insight into the details of the transaction experience, traditional field data is of limited use.

We turn to field experiments to study the impact of limited attention on consumer choice. The key advantage of this approach is that, by taking on the identities of sellers, we can ensure that the "real" good—the combination of the item on offer and the purchase experience—is exactly identical save for price. To be precise, we created two storefronts at the large online sales platform Taobao. We offered goods (mobile phone top-up cards) that were identical in all respects except price. Such goods are simple enough that unique store capabilities, such as assistance in use, technical advice, and so on, are of little consequence. Initially, our stores had no reputations. We maintained this absence of reputation throughout the experiment by not incorporating buyer feedback and through other means. We randomized the identity of the higher priced store for each of our items, so that neither store had a consistent pricing pattern. Our main treatment was to vary the gap between the low and high price offers for an identical item.

Because our storefronts lacked both reputation and sales history, they appear deep within the default search results at Taobao.[1] However, by choosing extremely low prices, we ensured that our listings appeared at or near the top of results obtained when consumers sorted by price. This creates a selection effect—our potential shoppers were only those willing to click to sort by price. For those who actually purchased from our shops, our lack of reputation and sales history also selects for low valuation for seller reputation. But, by limiting our sample to such consumers, we likely *understate* the magnitude of the true effect because such individuals are more conscientious searchers and more price sensitive.

Our findings suggest that limited attention plays a critical role in consumer choice. Even among the price sensitive consumers we study, out of 514 sales, 73 (or roughly 14%) were for the higher priced offerings. When the price gap was small (0.05%), fully 25% of our customers purchased at the higher price. Raising the percentage price gap produced a roughly linear fall in the percentage of total sales (by volume) occurring at the higher price. Nonetheless, even with a 2% price gap, 6% of our sales were at the higher price.

This would seem to imply that, as the price gap increases, price becomes more salient, so sales at the higher price fall. However, the nature of listings at Taobao is such that, following a sort by price, the *position* of the listing on the screen also changes. As the price gap increases, the higher priced offer appears lower down the page. Fortunately, the distance down the page depends on the number of competitors offering similar prices, which varies. We exploit this variation to separate price and location effects. Consistent with the limited attention hypothesis, regressions indicate that location alone drives choice—when the listings are close to one another, regardless of price, for higher priced offer enjoys more sales at locations higher on the screen.

Apart from our findings, we also contribute by augmenting existing studies of online pricing, most of which focus on Western firms and customers. China is the fastest growing retail market in the world. Moreover, its vast physical distances offer a compelling reason for a flourishing e-retail market. We are the first transaction level e-retail field experiment there.

The paper proceeds as follows: We place our work in the context of the extant literature in section II. Section III describes our field experiments and the institutional framework of Taobao. Section IV presents our main results with ancillary results offered in the Appendix. Finally, section V concludes.

II. Literature Review

The phenomenon of online price dispersion is well documented. Substantial price dispersion was found for books (Brynjolfsson & Smith, 2000; Clay, Krishnan, & Wolff, 2001), DVDs (Xing, 2007), consumer electronics (Baye, Morgan, & Scholten, 2004), travel services (Clemons, Hann, & Hitt, 2002), and even for commodity items (Ellison & Ellison, 2005; Pan, Ratchford, & Shankar, 2004). This price dispersion cannot easily be attributed to search costs (Levin, 2011), especially as price differences are often larger online than offline where search costs are presumably higher (Pan et al., 2004).

Such price dispersion might be explained by seller differences such as reputation, after-sales service, ease of returns, restocking fees, and so on. Baye et al. (2004), Baylis and Perloff (2002), and Pan, Ratchford, and Shankar (2002) try to control for this using proxy variables, but such controls are crude at best, leaving open the possibility that seller heterogeneity is the main driver of price dispersion.

Even if price dispersion arises with truly identical offers (save for price), such dispersion may be largely absent in terms of transaction prices, as several researchers point out. (See, e.g., Baye et al., 2004; and Pan et al., 2004). Worse yet, posted prices may not constitute real offers—low priced items are sometimes “stalking horses” designed to entice the consumer to the firm’s website where the consumer will be upsold the “true” product. Ellison and Ellison (2009) convincingly show the prevalence of such obfuscation strategies in consumer markets for DRAM chips. Thus, the connection between observed price dispersion and consumer choice behavior is tenuous at best.

Outside of consumer e-retail markets, researchers find evidence that higher priced vendors receive non-negligible transaction volume, all else equal. For instance, Ghose and Yao (2011) document choices by government agencies in the US Federal Supply Service system opting for higher price offers. The incentives of such agencies differ, sometimes markedly, from those of consumers.

Our contribution is to study the degree of transactional price dispersion in e-retail arising when sellers are truly identical and consumers are revealed to be price sensitive. Furthermore, we augment this literature by distinguishing between price and screen location as factors determining consumer choice. By doing so, we are able to distinguish for the first time between search cost based reasons for price dispersion for homogenous goods sold by homogenous retailers and limited attention based price dispersion

Others, such as Hossain and Morgan (2006) and Brown, Hossain, and Morgan (2010), have posed as online sellers to investigate behavioral aspects of consumer choice. Their main concerns center on the impact of observability of various elements of price, such as shipping and handling, on consumer choice. A key novelty of our approach is that we pose as multiple *competing* sellers. Our results come via comparison of the attributes of a competing offer in the presence of a fixed baseline offer. Also unlike these papers, we wish to make all aspects of offers as transparent as possible.

Pricing and price dispersion have also been studied in laboratory settings. For instance, Dufwenberg and Gneezy (2000) have subjects play the role of identical sellers and compare prices with different numbers of competitors. Consumer choice in their setting is pre-programmed (buyers always purchase at the lowest listed price). Despite this, they observe considerable price dispersion. In this case, dispersion derives not from cognitive limitations of consumers, but of competitors. Because the usual Bertrand result entails play of weakly dominated strategies, any relaxation of full rationality produces equilibria with “beauty contest” like features—in essence, the game becomes one of guessing the offer made by a rival seller and then undercutting it. By contrast, our concerns center on limited cognition/attention of consumers.

Distortions in consumer’s choice due to limited attention have been observed in many settings.[2] For instance, Lacetera, Pope, and Sydnor (2012) observe how similarly trivial search costs produce significant price effects in the used car market. Specifically, consumers appear to round the odometer readings of used cars to only the first digit, leading to a price discontinuity. An analogous phenomenon is well-known to retailers, as the widespread practice of selling items at $19.99 suggests. A key novelty in our study concerns selection. The extant literature mainly documents the behavior of average consumers. In contrast, our study concerns the choices made by a selected sample of consumers opting to order offers by price. Thus, these consumers have revealed acute price awareness. That these consumers should also make a significant number of "mistakes" in their purchasing represents a new finding.

III. Experimental Design and Rationale

An ideal experiment separating transaction volume from “strategic” price dispersion compared with effects driven by seller-product differences would occur in a market where the consumer’s assessment of the product’s use value is mostly or entirely divorced from after-sales service, returns policy, and other seller-specific issues. It would involve sellers providing users with an identical experience in completing

[2] Examples include: Hossain and Morgan (2006) and Brown et al. (2010).

the sale following a click on the comparison site. The sellers themselves should be identical in all reputational characteristics, such as feedback received and sales history. And it would provide an identical experience to all buyers landing on the comparison site and reviewing the available offers. Finally, the market should be sufficiently liquid to obtain transaction volume necessary to make statistically meaningful comparisons.

We sought to approximate these conditions as closely as possible. First, we chose the price comparison site Taobao as our platform. Taobao is the largest e-retail platform in the world by gross merchandise value (GMV), which amounted to 170 billion USD in 2012. While selling can occur by auction or posted price, nearly all transactions on Taobao use the latter selling mechanism, as we did in our experiments.

We sold mobile phone top-up cards in our experiments, which account for between 3 and 4 million transactions per day. These cards provide users with credits directly added to their mobile accounts that may be used for voice or text. The product is simple in that all cards work in the same way and have extremely low failure rates. It is a familiar product to users, so seller guidance as to the use of the product is unnecessary. Finally, in the unlikely event of failure, remediation occurs through the card provider rather than the seller.

Top-up cards are delineated by amount of credit, denominated in CNY, and province, of which there are 31 in China. In addition, a card covering all of China is available. A card for a given province may not be used to top-up a phone whose SIM card's designation is for some other province. Cards come in 10, 20, 30, 50, and 100 CNY denominations.

In our main treatment, we opened two, nearly identical, online storefronts whose inventory consisted of 20 CNY top-up cards for each province as well as all of China. We selected this denomination for two reasons. First, it provided an affordable means to maintain sufficient inventory so that stockouts were never a problem. Second, there are slightly fewer sellers in this market (although sellers still numbered over one hundred thousand), and this made it slightly easier to ensure that our offers would be noticed by consumers.

Top-up card sellers on Taobao must obtain a special software in order to operate. Fortunately, there is an active market for such selling software, and we purchased ours from the jieyitong.net platform for 200 CNY each. The name of shops, their brand so to speak, consists mainly of a rough description of the items in which they specialize along with an abbreviation of the seller's username. Following this industry's practice, we named our shops "Top-up XYZ" where XYZ represent a random sequence of letters associated with the first three letters of a meaningless username which varied between the two shops. For instance, in some treatments, we named one store "Top-up EDH," associated with username edhubikk, and the other "Top-up WIG," associated with the username wigkupa. We intentionally chose meaningless usernames to minimize any associations with existing brands and to ensure that both storefronts were on an equal footing. Because our storefronts were newly created, they shared identical (nil) sales histories and reputation. Thus, to the degree possible, we eliminated perceptions of seller quality from buyer choice between our competing storefronts. Figure A1 and A2 in the Appendix provide screenshots of the image our storefronts presented to would-be buyers. Notice that there is virtually nothing to differentiate the two. Indeed, we even opted not to include any logo or profile photo whatsoever to ensure an identical user experience.

Our main treatment variation were the prices (including shipping and handling) charged at each shop. The custom in this market is to offer free shipping, and we followed this practice. Thus, the price the user observed on Taobao was also the price they paid, as was the case with most of our competitors. We also opted to price extremely low, even below the cost of the card itself. Such a practice is not unusual in this market because many sellers intentionally make losses on cards for the benefit of accumulating both history and reputation for being reliable. Indeed, the practice is so widespread that Taobao imposes minimum prices for top-up cards. In our case, the permitted minimum was 19.10 CNY at the time of our experiments. Our low price offers in all treatments were .01 above the minimum while our higher price varied by treatment, ranging from 0.05% to 2.51% above this price. Finally, to ensure that neither store was seen as having an advantage in pricing, we randomly assigned the high offer between stores for each of the 32 provinces plus all of China. Table 1 summarizes our treatments.

We kept the offers associated with each treatment in place for one day, after which we changed the higher offer and re-randomized the identity of the store offering the lower price for each of the 32 different locations associated with a card of a given denomination. While this might seem to be a short duration, volume on Taobao is such that we obtained considerable sales volume even within this short time horizon. Our design also has the advantage of minimizing demand fluctuations within a treatment as intra-day demand variation is likely to be minimal.

We selected the baseline price point of 19.11 CNY to ensure that the list of prices (the landing page in the parlance of e-commerce) a buyer faced when considering our offer would not vary across buyers. Specifically, when buyers search by item, Taobao presents them with a default landing page not ordered by price. While the precise algorithm used to delineate the ordering is secret, casual investigation quickly reveals that sales history and reputation play a large role. Our shops never appeared on this page since they lacked both history and reputation. Buyers are, however, able to sort sellers produced from a given search query by price, from lowest to highest or, if they wish, highest to lowest. We selected our prices so that, in the great majority of cases, our listings would appear on the landing page when a buyer sorted by price. Given the depth of the market, it seems highly unlikely that a buyer engaged in any other search strategy would uncover our offers. Thus, we can say, with reasonable confidence, that when our listing appeared on the buyer's landing page, the process resulted from a search by price. Figure A3 in the Appendix shows a screenshot of Treatment 2.

When choosing among offers, most buyers scan the page from top to bottom; thus most buyers would encounter our low price listing first and only later see the higher offer. However, this viewing strategy may not be universal, and it is quite possible that some buyers, perhaps those attracted to the middle of the page, might encounter the listings in the opposite order. As this is a field experiment, we have no practical means of ocular measurement to track the viewing process used by each potential buyer.

One difficulty is that, as each storefront experiences success, its sales history and possibly reputation will increase: thereby creating differences between stores over time. We sought to minimize these possible heterogeneities in several ways. First, by randomizing the identity of the firm making the lower price offer, both firms should, in principle, gain sales and reputation at the same rate. Second, we opted to delay the posting of buyer feedback (ratings) by 15 days from the time that they rated the store. This ensured that our stores maintained an equal (zero) reputation over much of the experiment. Third, Taobao records a store's sales history only over the preceding 30 days from the time a buyer lands on a page listing our offer. We exploited this by creating a gap between batches of treatments exceeding this interval. Specifically, treatments 1 and 2 occurred during the first two weeks of November 2010 treatments 3-6 occurred throughout January 2011, and treatment 0 occurred in April 2011.

We introduced this gap for other reasons as well. First, we worried about a demand shift occurring during the (Western) holiday buying season, which traditionally occurs starting in the last week in November and continuing throughout most of December. Because our results are comparative, changes in levels of demand should, in principle have no impact on the difference in the likelihood of choosing the higher offer, but would change the volume of transactions occurring at each price point. However, the degree of limited attention, our prior hypothesis as an important price dispersion driver, could plausibly increase as harried holiday shoppers might be systematically less careful than at other times.

In our "restart," we reset the competitive position of both shops by closing our earlier shops and creating two new store names and associated usernames using the same protocol as in the earlier treatments. Here again the goal was to approximate, as closely as possible, identical firms in the ideal experiment.

We limited the amount of the price difference across all treatments mainly to ensure that both listings would appear on a consumer's landing page following a price search. Were the price gap too large, the higher price firm would appear on a later page, introducing an additional confound, a discrete search friction, into the design.

To summarize, we designed our experiment to study transactional price dispersion when two identical firms compete. Under the purely rational perspective (as in models such as Baye and Morgan, 2001), listed price dispersion should produce no transactional price dispersion as rational consumers should never purchase a homogeneous product at the higher price when shown on the same search result as the lower price offer.

In contrast, consumers suffering from limited attention will buy non-negligible amounts at the higher price. To see this formally, suppose that, following a search that returns n listings, a limited attention consumer only notices one of these. Listing i is noticed with probability p_i, strictly decreasing in i when firms are indexed in order of appearance from the top of the results page, i.e., listing 1 appears at the top and listing n at the bottom. Such a listing pattern would predict a decreasing market share for the higher-priced lower-ranked shop, as the price and therefore distance gap between our two shops increased. Price, in this model, is irrelevant—a consumer purchases so long as the price of the noticed firm does not exceed the consumer's willingness to pay. Thus, not controlling for position, higher prices produce fewer sales but, controlling for position relative to the lower price offer, price is independent of the volume of sales.[3]

IV. Results

Because the premise of the experiment is to create a neutral setting with respect to competing storefronts, we test this by pooling all treatments, controlling for price and screen position (non-parametrically), and using dummy variables for firm identity. Table 2 displays the results of this analysis.

[3] The above model represents a polar case of limited attention as would-be buyers notice only a single listing. It may be readily generalized to allow buyers to notice a subset $k < n$ of the listings. This would not qualitatively change the results, but would mean that, even after controlling for location, price level would remain a significant determinant of sales volume.

As expected, the coefficient associated with shop identity is neither statistically or economically significant. Thus, consumer choice behavior appears unaffected by the particulars of each shop apart from experimentally induced variation in price and naturally occurring variation in screen position, which we describe in more detail below.

It is well known that screen position drives consumer attention and, as a consequence, choice behavior. Firms bid considerable amounts to obtain the top sponsored search position for keyword searches at sponsored search sites like Google. While we care mainly about the effect of price on choice, a price ranked search will necessarily produce differences in screen position—the lower priced listing will be listed higher among the set of search results. The amount by which the offers appear on the results screen also depends on the magnitude of the price difference. Larger price differences will admit more competing offers lying between the two listings of the experimental firms. Thus, when the offers of the two firms differ by a relatively small amount, their listings will appear closer together, adjacent in cases where the price difference is extremely small. When the offers differ by a lot, the listings will appear farther apart, with non-experimental listings appearing in between. The number of intervening listings varies with the competitive environment for each of the top-up card markets, i.e., the provinces where the cards are valid. This naturally occurring variation, together with a treatment intervention described below allows us to identify location effects separately from price effects on choice.

As a baseline, we consider a situation where the firms are absolutely identical in all respects, including price. We refer to this baseline as treatment 0. Even though the firms are identical, they will still differ in the order in which they appear on the screen—the listing posted slightly earlier appears above that posted slightly later. In all cases, the two listings appeared in adjacent positions within the returned search. If this small difference in location were irrelevant, we would expect the two firms to achieve roughly equal sales. By contrast, the listing located at the higher position achieves roughly double the sales volume of the listing at the lower position.

To formally test for the existence of a location effect, we perform a Poisson regression of sales volume on a dummy for being the listing located lower on the screen. We also control for store identity and card region. Table 3, column (1) reports the results of this analysis. As that table shows, even this small location difference is highly significant. The magnitude of the effect is Exp(-0.704) or approximately about a 49% drop in sales volume, all else equal. Treating the number of buyers as the dependent variable produces a similar finding as seen in column (4) of Table 3.

We can only conjecture as to possible reasons why consumers are more likely to choose the higher listing. However, with nothing else to distinguish the offers, it is intuitive that consumers might break the tie by choosing the higher-up listing. Where consumers sorted by price; nonetheless, higher listings are, in general, lower priced listings, so consumers following this heuristic strategy will, more often than not, make good choices. In any case, the consistency between our treatment 0, where there was no price difference, and treatment 1, where there as small price difference between our items, further suggests that location, but not small price differences drove differences in sales[4].

Treatment 1 introduces small price variation between the two listings. Indeed, the difference is so small that, in every case, the two listings appear in adjacent locations. Since price and location perfectly co-vary in this treatment, one cannot separately estimate different effects. Instead, we pool with the baseline treatment under the hypothesis (verified later) that such small price differences have no effect on choice. Columns (2) and (5) of Table 3 report the results of a Poisson specification for sales volume and unique buyers, respectively. The estimated coefficient is now larger (in absolute value) than in the baseline treatment, which implies that the impact of being in a disadvantageous screen location is now slightly larger than under the baseline. This is intuitive since there is now a (small) price difference reinforcing the location difference. But, as we discuss in reference to Table 4 below, one should not make much of this small difference.

Treatment 2 increases the magnitude of the price difference between the two listings. As a consequence, listings are no longer always adjacent. In some cases, one or two additional listings by competitors appear between the listings of the two experimental firms. We code this as "distance", where distance equal to 1 indicates adjacency and distance equal to x indicates that x - 1 competing listings appear between the two experimental firms. Ignoring price variation for the moment and pooling treatments 0-2 reveals that screen location consistently impacts sales (Columns (3) and (6) of Table 3). The coefficient estimates show that screen position is a central choice driver—moving down one position halves sales, moving down two drops them by a whopping 72% (1 – Exp(-1.75)). Beyond two positions, however, we lack sufficient information to make precise estimates. Under the full rationality benchmark, we would expect no systematic effect of screen position on sales; after all, the competing offers may be easily inspected, appearing close to one another on the landing page. The limited attention model, however, suggests exactly this type of effect, though it is silent as to the magnitude of lost attention as screen position gets lower.

4 We observed that our two existing shops in our prior treatments had slightly different ratings, the higher rated store with the later posting time was lower down on the screen when items were sorted by price. Given that the stores were otherwise identical, this difference in location can only be explained by the earlier posting time. However, consumers cannot be expected to know this.

The worry, of course, is that the specification in Table 3 suffers from omitted variable bias, most notably, the omission of price. We performed the same analysis including dummy variables for each price level. This flexible specification allows for the possibility that demand increases as well as decreases with price or that demand varies non-monotonically. Regardless, the results are unchanged by this addition—price is economically and statistically insignificant while the coefficient estimates of location remain quantitatively and statistically similar. In short, small differences in price compared to the baseline of 19.11 CNY seem not to matter to consumer choice.

This is not to say that consumers universally ignore price. Table 4 extends the specification of Table 3 with the addition of dummy variables for the higher price points associated with treatments 3-6. To gain sufficient power to identify screen location effects at longer distances, we pool screen locations of distance 3 or greater.[5] Columns (1) – (4) use sales volume as the dependent variable; columns (5) – (8) use the number of unique buyers. For sufficiently large price differences, price becomes an important consideration.

Notice that, regardless of the sample, coefficient estimates on adjacent and distance 2 listings are remarkably consistent. Distances of 3 or more are marginally significant, but distance seems to matter less, and price more, as listings are further from one another. Price becomes a significant consideration once it is more than 2% higher than the baseline of 19.11. But, even when the high price firm is charging 2.4% more, the magnitude of the coefficient is similar to the impact of being two positions below and charging the same price. Treating the number of buyers as the dependent variable, price becomes significant only for the highest price treatment.[6]

The key observation arising from Table 4 is that, even for consumers electing to sort by price, price itself forms only a secondary consideration so long as it is "competitive" with a lower priced offer. Given the nature of the selection in our sample of buyers, we suspect that our estimates *overstate* the importance of price to the populace at large. In short, even in a setting where consumers *have* sorted by price, and

[5] We chose a number of different topcodes for pooling larger distances in screen position. The results are little affected by this choice, so we opted to report the simplest specification.

[6] To our knowledge, most Taobao transactions happened on PC screens at the time of the experiment in 2011. Our screenshots (available on request) show that Taobao had 40 listings per page. When the higher price is 19.59, our higher priced listings were often on the second page. But our listings were never on the 2nd page when our price was 19.39 or 19.49. For the treatment with the price of 19.49, we add a dummy variable for page break in Table A1 in the Appendix. Although the page break dummy has a large and significant coefficient, adding it barely changes the price coefficient, though it does make it less significant.

have demonstrated a low valuation for quality, a firm offering the higher price still attracts a non-trivial number of consumers.

Our results are closely related to the literature on rank-ordered choice (Hefti & Heinke, 2015). The distance effect we find in our experiment is qualitatively consistent with the finding of the literature within the first page (Pan et al., 2007). Also, similar to the literature, we find that the first page dominates (Jansen, Spink, and Saracevic, 2000; Smith and Brynjolfsson, 2001). Both the distance and the first page effects are consistent with our limited attention hypothesis (Hefti and Heinke, 2015). We contribute to this literature with a field experiment, which allows us to cleanly identify the causal effect of distance on sales. Furthermore, we extend the finding to a new setting (the Taobao platform).

V. Discussion

Our findings suggest that behavioral factors, such as limited attention, produce transactional price dispersion, even in settings where the purchase experience is exactly identical. Even though our consumers elected to sort offers by price, they appear more sensitive to the *location* of a listing relative to rival offers than to the price itself. Only at price differences of 2% or greater do consumers start to respond to price differences.

This should not be taken to mean that charging higher prices represents an effective business strategy. In exchange for a small increase in price, sellers suffer massive decreases in volume owing to their disadvantageous screen location. Depending on costs, it might be an effective strategy to price high and then to pay the necessary fees to appear at the top of the default list of search results a la Google. On a platform like Taobao, limited attention generates considerable value. Because location matters more than price, firms are willing to pay for privileged positions. They can profitably do so by charging premium prices.

While we do not claim that strategic seller behavior that leads to price dispersion does not exist, we find little evidence in favor of the “strategic” price dispersion of the sort envisaged by Baye and Morgan (2001). Competition for price-sensitive “shoppers” drives dispersion, at least theoretically, yet real-world “shoppers” are not terribly price-sensitive, and thus, defeat the rationale for producing price uncertainty suggested by the theory. Transactional price dispersion displays patterns much more consistent with behavioral models than with the *homo economicus* players inhabiting most theory models on the subject. Our findings can be explained by a menu effect where consumers exhibit errors in execution of a choice among adjacent or similarly coded items on menus in other domains such as voting (Shue and Luttmer, 2007) and investments (Rashes 2001).

We stress that our results represent a useful step, rather than the last word, on transactional price dispersion. Our analysis is limited in scope in so far as our experiment was with one set of price-sensitive/reputation insensitive consumers for one homogenous product in one large online market. However, as we discussed, these consumers should be the least susceptible to limited attention to price, and therefore, our results would strongly suggest the generality of our findings to other less price-sensitive consumers for other products in other markets. Establishing this implication is for future work. An important follow-up would investigate the behavior of the "non-shoppers" visiting the platform—consumers who do not search by price. Such an investigation would also exogenously vary screen position through things like bidding for prime internet real estate atop the default landing page. Finally, a follow-up study would pay closer attention to variation in competition. We know only the distance between listings, but clearly one would like to know more about the universe of competitors than this. The likely presence of higher quality competitive stores with prices intermediate to our low and high priced stores would lower the share of our sales at our higher price store. Gaining this information requires additional scraping or direct access to the platform that was beyond the scope of this research.

REFERENCES

Baye, M. R., & Morgan, J. (2001). Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review*, 454-474.

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.

Baylis, K., & Perloff, J. M. (2002). Price dispersion on the internet: good firms and bad firms. *Review of Industrial Organization, 21*(3), 305-324.

Brown, J., Hossain, T., & Morgan, J. (2010). Shrouded attributes and information suppression: Evidence from the field. *Quarterly Journal of Economics, 125*(2), 859-876.

Brynjolfsson, E., & Smith, M. D. (2000). Frictionless commerce? A comparison of Internet and conventional retailers. *Management Science, 46*(4), 563-585.

Clay, K., Krishnan, R., & Wolff, E. (2001). Prices and price dispersion on the web: evidence from the online book industry. *Journal of Industrial Economics, 49*(4), 521-539.

Clemons, E. K., Hann, I.-H., & Hitt, L. M. (2002). Price dispersion and differentiation in online travel: An empirical investigation. *Management Science, 48*(4), 534-549.

Dufwenberg, M., & Gneezy, U. (2000). Price competition and market concentration: an experimental study. *International Journal of Industrial Organization, 18*(1), 7-22.

Ellison, G., & Ellison, S. F. (2005). Lessons about markets from the Internet. *Journal of Economic Perspectives*, *19*(2), 139-158.

Ellison, G., & Ellison, S. F. (2009). Search, obfuscation, and price elasticities on the Internet. *Econometrica, 77*(2), 427-452.

Ghose, A., & Yao, Y. (2011). Using transaction prices to re-examine price dispersion in electronic markets. *Information Systems Research, 22*(2), 269-288.

Hefti, A., & Heinke, S. (2015). On the economics of superabundant information and scarce attention. *Œconomia. History, Methodology, Philosophy,* (5-1), 37-76.

Hong, H., & Shum, M. (2006). Using price distributions to estimate search costs. *RAND Journal of Economics, 37*(2), 257-275.

Hossain, T., & Morgan, J. (2006). ...plus shipping and handling: revenue (non) equivalence in field experiments on eBay. *Advances in Economic Analysis & Policy, 5*(2).

Jansen, B. J., Spink, A., & Saracevic, T. (2000). Real life, real users, and real needs: a study and analysis of user queries on the web. *Information Processing & Management,* 36(2), 207-227.

Lacetera, N., Pope, D. G., & Sydnor, J. R. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review*, *102*(5), 2206-2236.

Levin, J. D. (2011). The economics of internet markets: National Bureau of Economic Research.

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. (2004). Price dispersion on the internet: a review and directions for future research. *Journal of Interactive Marketing*, *18*(4), 116-135.

Pan, B., Hembrooke, H., Joachims, T., Lorigo, L., Gay, G., & Granka, L. (2007). In google we trust: Users' decisions on rank, position, and relevance. *Journal of Computer‐Mediated Communication*, *12*(3), 801-823.

Rashes, M. S. (2001). Massively confused investors making conspicuously ignorant choices (mci–mcic). *Journal of Finance*, *56*(5), 1911-1927.

Shue, K., & Luttmer, E. F. (2009). Who misvotes? The effect of differential cognition costs on election outcomes. *American Economic Journal: Economic Policy*, *1*(1), 229-257.

Smith, M. D., & Brynjolfsson, E. (2001). Consumer decision‐making at an Internet shopbot: Brand still matters. *The Journal of Industrial Economics*, *49*(4), 541-558.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, *69*(3), 213-225.

Varian, H. R. (1980). A model of sales. *American Economic Review*, *70*(4), 651-659.

Xing, X. (2007). Does price converge on the internet? Evidence from the online DVD market. *Applied Economics Letters*, *15*(1), 11-14.

TABLE 1: SUMMARY OF TREATMENTS

Treatment	Date	Low Price	High Price	Average distance between listings*	Total sales	Percentage of sales at higher price
1	Nov 26-27, 2011	19.11	19.12	1.0 (0.0)	129	33%
2	Nov 30, 2011	19.11	19.19	2.2 (0.85)	78	24%
3	Jan 2, 2011	19.11	19.29	4.3 (2.2)	88	18%
4	Jan 6, 2012	19.11	19.39	8.3 (3.3)	166	14%
5	Jan 5, 2012	19.11	19.49	21 (5.8)	78	5%
6	Jan 4, 2012	19.11	19.59	38 (14)	79	1%
0	Apr 10-12, 2012	19.11	19.11	1.0 (0.0)	105	N/A

Notes: We ran two sessions of Treatment 1, four sessions of Treatment 0, and one session each for the rest.

*: Standard deviation in parentheses. We define adjacent listings to have distance=1, two listings with just one other listing in between have distance=2, and so on.

TABLE 2: NO SHOP EFFECTS

Dependent variable: sales								
Regressions	Poisson					Negative Binomial		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shop dummy for Treatment 1&2	-0.109	-0.109	-0.0739	-0.109	-0.155	-0.106	-0.0320	0.0319
	(0.396)	(0.389)	(0.303)	(0.325)	(0.233)	(0.360)	(0.244)	(0.255)
Shop dummy for Treatment 3-6	0.0286	0.0804	0.117	0.0286	0.223	0.0158	-0.0263	0.371
	(0.250)	(0.236)	(0.244)	(0.242)	(0.152)	(0.259)	(0.147)	(0.212)
Shop dummy for Treatment 0	0.0953	0.0953	-0.280	0.0953	0.0953	0.0726	0.168	0.168
	(0.307)	(0.307)	(0.201)	(0.240)	(0.240)	(0.300)	(0.244)	(0.244)
Session F.E.	Y	Y				Y		
Price dummies		Y						
Session -by-Region R.E.			Y		Y			Y
Session -by-Region F.E.				Y	Y		Y	Y
N	704	704	704	360	360	704	360	360

Notes: Each observation is sales of a 20 CNY card for each region (32), shop (2) and sessions (7 treatments, and we have 2 sessions for treatment 1 and 4 sessions for treatment 0, adding up to 11 sessions in total). Therefore we have N=32×2×11=704. For the most flexible Treatment-by-Region fixed effect models (4), (5), (7), and (8), we allow for each region in each session to have a separate fixed effect, and consequently we have 172 groups (344 observations) dropped because of all zero sales.

Standard errors in parentheses. For (1), (2), and (6), S.E. are robust standard errors. For (3) and (4), S.E. are clustered by session-region. For (5), (7), and (8), S.E. are standard asymptotic standard errors, and robust or bootstrap standard errors lead to similarly insignificance.

* p<0.05, ** p<0.01, *** p<0.001.

TABLE 3: DISTANCE EFFECTS

Conditional fixed effect Poisson Regression. Dependent variable: sales

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Treatment 0, all sales	Treatment 0&1, all sales	Treatment 0~2, all sales	Treatment 0, unique buyers	Treatment 0&1, unique buyers	Treatment 0~2, unique buyers
Distance=1	-0.704**	-0.724***	-0.697***	-0.644**	-0.762***	-0.743***
	(0.235)	(0.167)	(0.158)	(0.232)	(0.169)	(0.159)
Distance=2			-1.735***			-2.110***
			(0.259)			(0.421)
Distance≥3			-1.386			-1.386
			(1.099)			(1.099)
Shop dummies	Y	Y	Y	Y	Y	Y
Session-by-Region F.E.	Y	Y	Y	Y	Y	Y
N	88	160	198	88	160	196
Incidence rate ratios (IRR)						
Distance=1	0.494**	0.485***	0.498**	0.525**	0.467***	0.476**
	(0.116)	(0.081)	(0.078)	(0.122)	(0.079)	(0.076)
Distance=2			0.176***			0.121***
			(0.046)			(0.051)
Distance>=3			0.25			0.25
			(0.27)			(0.27)

Notes: Each observation is sales of a 20 CNY card for each region and each shop. We define distance for the lower price (19.11) listing in Treatment 1-6 and the listing on the top in Treatment 0 to be zero. The adjacent listing have distance=1, the listing with just one other listing in between have distance=2, and so on. IRRs are obtained by exponentiating the Poisson regression coefficient. It shows the effect of RHS variable on the expected sales. For instance, having a listing with distance=1, or just below the baseline (distance=0), will reduce sales to 49.4% in (1). Adding price dummies for Treatment 1 (p=19.12) and 2 (p=19.19) does not change the result. Using non-clustered standard errors or negative binomial regressions does not change the result.

Standard errors (clustered for session-region) in parentheses.

* p<0.05, ** p<0.01, *** p<0.001.

TABLE 4: DISTANCE AND PRICE EFFECTS

Conditional fixed effect Poisson Regression. Dependent variable: sales

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sample	Treatment 0~3, all sales	Treatment 0~4, all sales	Treatment 0~5, all sales	Treatment 0~6, all sales	Treatment 0~3, unique buyers	Treatment 0~4, unique buyers	Treatment 0~5, unique buyers	Treatment 0~6, unique buyers
Distance=1	-0.696***	-0.697***	-0.697***	-0.697***	-0.745***	-0.746***	-0.745***	-0.746***
	(0.156)	(0.156)	(0.156)	(0.156)	(0.157)	(0.157)	(0.157)	(0.157)
Distance=2	-2.087***	-2.067***	-2.070***	-2.051***	-2.386***	-2.348***	-2.361***	-2.342***
	(0.325)	(0.307)	(0.305)	(0.304)	(0.472)	(0.458)	(0.457)	(0.456)
Distance≥3	-1.208*	-1.193*	-1.196*	-1.181*	-1.185*	-1.158	-1.167	-1.153
	(0.592)	(0.596)	(0.597)	(0.594)	(0.601)	(0.604)	(0.605)	(0.604)
Price= 19.29	-0.343	-0.351	-0.350	-0.358	-0.322	-0.338	-0.333	-0.341
	(0.540)	(0.538)	(0.539)	(0.536)	(0.535)	(0.533)	(0.535)	(0.533)
Price= 19.39		-0.692	-0.690	-0.697		-0.799	-0.794	-0.801
		(0.658)	(0.658)	(0.656)		(0.659)	(0.659)	(0.659)
Price= 19.49			-1.783*	-1.786*			-1.611	-1.616
			(0.861)	(0.857)			(0.834)	(0.832)
Price= 19.59				-3.198**				-3.071**
				(1.181)				(1.182)
Shop dummies	Y	Y	Y	Y	Y	Y	Y	Y
Session-by-Region F.E.	Y	Y	Y	Y	Y	Y	Y	Y
N	236	286	316	360	230	278	308	350
Incidence rate ratios (IRR)								
Distance=1	0.498***	0.498***	0.498***	0.498***	0.475***	0.474***	0.474***	0.474***
	(0.078)	(0.078)	(0.078)	(0.078)	(0.074)	(0.074)	(0.074)	(0.074)
Distance=2	0.124***	0.127***	0.126***	0.129***	0.092***	0.096***	0.094***	0.096***
	(0.040)	(0.039)	(0.039)	(0.039)	(0.043)	(0.044)	(0.043)	(0.044)
Distance≥3	0.298*	0.303*	0.303*	0.307*	0.306*	0.314	0.311	0.316
	(0.177)	(0.181)	(0.180)	(0.182)	(0.184)	(0.190)	(0.188)	(0.191)
Price= 19.29	0.710	0.704	0.705	0.699	0.724	0.713	0.717	0.711
	(0.384)	(0.378)	(0.380)	(0.374)	(0.388)	(0.380)	(0.383)	(0.379)
Price= 19.39		0.501	0.501	0.498		0.450	0.452	0.449
		(0.329)	(0.330)	(0.327)		(0.296)	(0.298)	(0.296)
Price= 19.49			0.168*	0.168*			0.200	0.199
			(0.145)	(0.144)			(0.167)	(0.165)
Price= 19.59				0.041**				0.046**
				(0.048)				(0.055)

Notes: Each observation is sales of a 20 CNY card for each region and each shop. We define distance for the lower price (19.11) listing in Treatment 1-6 and the listing on the top in Treatment 0 to be zero. The adjacent listing have distance=1, the listing with just one other listing in between have distance=2, and so on. IRRs are obtained by exponentiating the Poisson regression coefficient. It shows the effect of RHS variable on the expected sales. Adding price dummies for Treatment 1 (p=19.12) and 2 (p=19.19) does not change the result. Using non-clustered standard errors, using negative binomial regressions or adding additional control of distances does not change the result.

Standard errors (clustered for session-region) in parentheses. * p<0.05, ** p<0.01, *** p<0.001.

Appendix

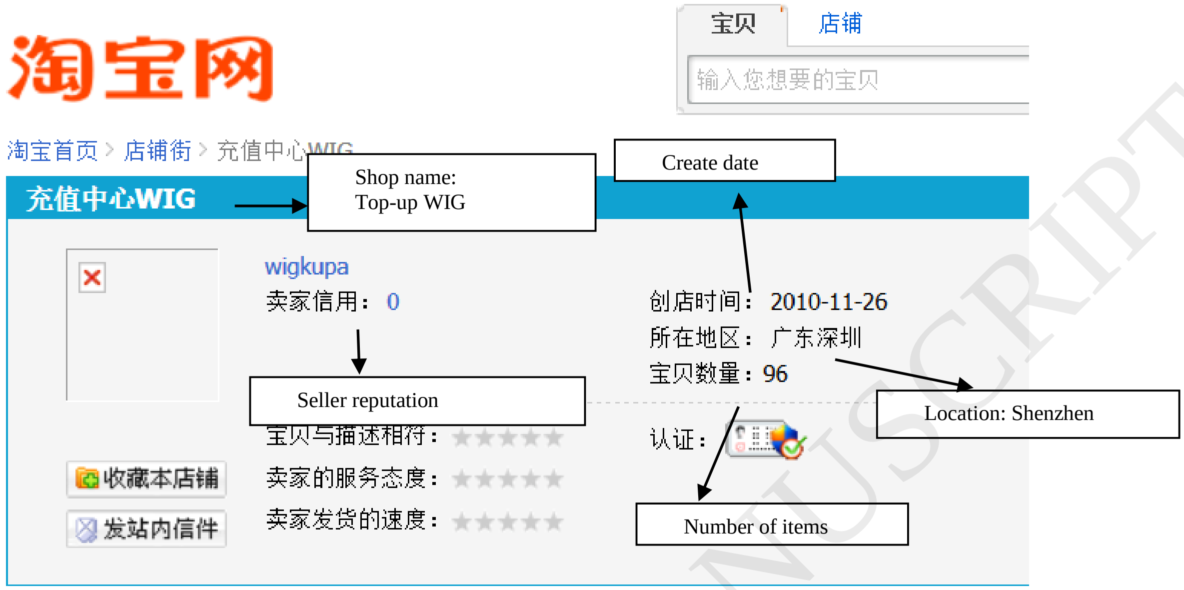

FIGURE A1: SHOP 1 SCREENSHOT

Notes: We did not use any picture for our sellers. Hence, it is shown as a broken link.

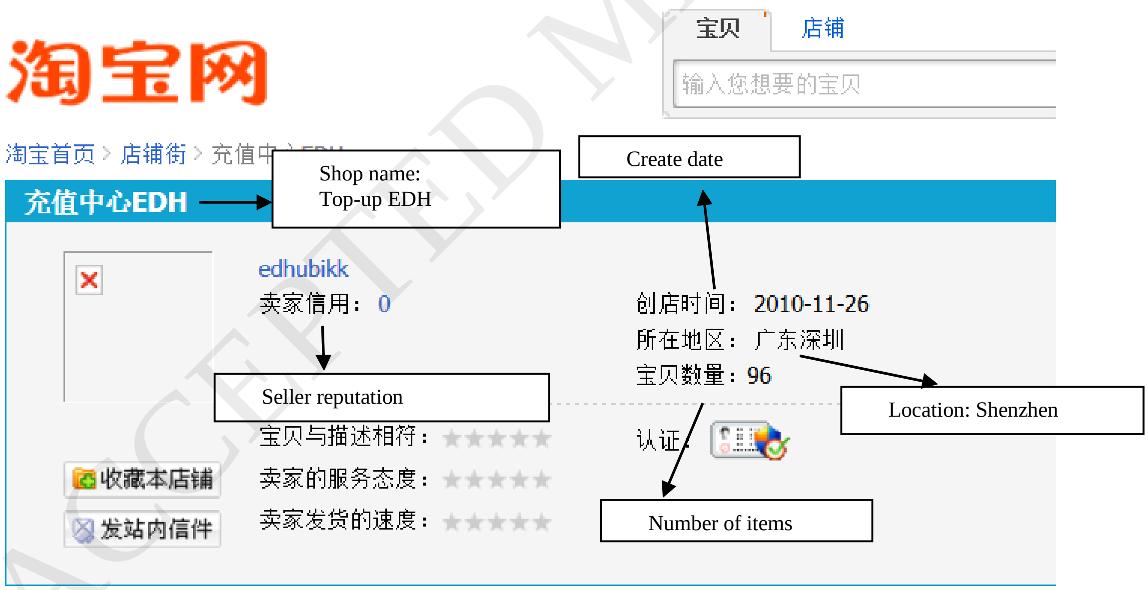

FIGURE A2: SHOP 2 SCREENSHOT

Notes: We did not use any profile picture for our sellers. Hence, it is shown as a broken link.

FIGURE A3: LISTINGS SCREENSHOT

TABLE A1: DISTANCE AND PRICE EFFECTS

Conditional fixed effect Poisson Regression. Dependent variable: sales

	(1)	(2)
Sample	Treatment 0~6, all sales	Treatment 0~6, unique buyers
Distance=1	-0.697***	-0.746***
	(0.156)	(0.157)
Distance=2	-2.056***	-2.347***
	(0.304)	(0.456)
Distance≥3	-1.185*	-1.157
	(0.595)	(0.604)
Page break	-12.22***	-12.28***
	(1.175)	(1.189)
Price=19.29	-0.356	-0.339
	(0.537)	(0.534)
Price=19.39	-0.695	-0.799
	(0.657)	(0.659)
Price=19.49	-1.785*	-1.614
	(0.858)	(0.833)
Price=19.59	-2.732*	-2.648*
	(1.182)	(1.182)
Shop dummies	Y	Y
Session-by-Region F.E.	Y	Y
N	360	350
Distance=1	0.498***	0.474***
	(0.0777)	(0.0744)
Distance=2	0.128***	0.0957***
	(0.0388)	(0.0436)
Distance≥3	0.306*	0.314
	(0.182)	(0.190)
Page break	0.00000491***	0.00000462***
	(0.00000577)	(0.00000550)
Price=19.29	0.700	0.713
	(0.376)	(0.380)
Price=19.39	0.499	0.450
	(0.328)	(0.296)
Price=19.49	0.168*	0.199
	(0.144)	(0.166)
Price=19.59	0.0651*	0.0708*
	(0.0770)	(0.0837)

Notes: Each observation is sales of a 20 CNY card for each region and each shop. We define distance for the lower price (19.11) listing in Treatment 1-6 and the listing on the top in Treatment 0 to be zero. The adjacent listing have distance=1, the listing with just one other listing in between have distance=2, and so on. IRRs are obtained by exponentiating the Poisson regression coefficient. It shows the effect of RHS variable on the expected sales. Adding price dummies for Treatment 1 (p=19.12) and 2 (p=19.19) does not change the result. Using non-clustered standard errors, using negative binomial regressions or adding additional control of distances does not change the result.

Standard errors (clustered for session-region) in parentheses. * p<0.05, ** p<0.01, *** p<0.001.