



Full Length Articles

Quantum prices[☆]Diego Aparicio^{a,*}, Roberto Rigobon^b^a IESE Business School, University of Navarra, Spain^b MIT and NBER, USA

ARTICLE INFO

Article history:

Received 4 January 2022

Received in revised form 3 May 2023

Accepted 4 May 2023

Available online 9 May 2023

Repository link: <https://data.mendeley.com/datasets/gvn68hz8sg/1>

JEL codes:

D22

M3

L81

E7

Keywords:

Behavioral pricing

Price frictions

Saliency

Recall

Cross-country retail

ABSTRACT

Many retailers practice an extreme form of discrete pricing defined as *quantum prices*: differentiated products are priced using few and sparse price buckets. To show this, online data was collected for 350,000 products from over 65 fashion retailers in the U.S. and the U.K. This pricing strategy is observed within categories and across categories (i.e., similar or even disparately distinct products like jeans and bags have an identical price), as well as in product introductions, where new products come in at previous price buckets. Normalized indices indicate substantial price clustering after controlling for popular prices, convenient prices, assortment size, or digit endings. Quantum prices have implications for price adjustments through product shares, markdown prices, and for the law-of-one-price. A behavioral model of price saliency and recall is discussed.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

How firms set prices remains a core question of interest in economics, operations, marketing, and behavioral sciences. It is known from theoretical and empirical studies that firms do not choose prices from an unrestricted price grid. The seminal work of Maskin and Tirole (1988) writes: “The price space is discrete, i.e., firms cannot set prices in units smaller than, say, a penny.” In the past decades, the extant literature has documented that firms and consumers exhibit biases or heuristics leading to realized prices that are significantly more discrete.

[☆] The paper is based on one chapter of Diego Aparicio's PhD dissertation—Who gratefully acknowledges the guidance of, and discussions with, Glenn Ellison. The authors thank Ignacio Puente, Alberto Cavallo, Mariana Albano, Pablo Azar, Manuel Bertolotto, Fernando Broner, Vivek Bhattacharya, Jean-Pierre Dube, Pierre Dubois, Paul Ellickson, Sara Ellison, Miguel Espinosa, Ricardo Fraiman, Renee Gosline, Paul Joskow, Chris Knittel, Carolina Levinton, Jura Liaukonyte, Juan Cristobal Lopez, Victor Martinez de Albeniz, John McCoy, Mateo Montenegro, Emi Nakamura, Augusto Ospital, Chela Pedreschi, Pilar Posse, Drazen Prelec, Nancy Rose, Duncan Simester, Luminita Stevens, Nader Tavassoli, Julian Villanueva, Birger Wernerfelt, Mike Whinston, seminar participants at MIT, the Golub Center for Finance and Policy, and the VI Financial Information Forum (CEMLA), and pricing managers and industry consultants, for helpful comments. We also thank the Editor Martin Uribe and two anonymous referees for critical feedback. Nicolas Gomez del Campo, Gina Rey, Nestor Santiago-Perez, and Sophie Reynolds provided excellent research assistance.

* Corresponding author.

E-mail address: daparicio@iese.edu (D. Aparicio).

We study the fashion and clothing retail industry and observe three distinct pricing behaviors: Traditional, Platform, and what will be defined as *Quantum* pricing strategies. Two of these have received considerable attention in the literature. Traditional pricing stores—such as Louis Vuitton, Dolce Gabbana, Ralph Lauren, Reiss—tend to obfuscate prices from the consumer, often making prices the least salient attribute. In many cases, products have no tags attached, or prices are carefully hidden from the shopping experience. Platform pricing retailers—such as Amazon, Walmart, Wayfair—emphasize growing their network and expanding market share. They tend to have two sources of revenues, namely a mark-up and the gains from utilizing customer information. Prices are often decided by advanced algorithms, and their distribution is dense, almost continuous on a price line. The nature of the competition implies that few retailers can follow this approach because only a few can capture the whole network.

Our paper focuses on a novel form of pricing: “Quantum Pricing.” It is followed by a vast number of firms, including some of the leading and most successful firms, such as Zara, H&M, Ikea, Apple, Bonobos, and Uniqlo. These firms tend to massively cluster prices along price points that are distant from each other and that are sticky over time. In fact, product introductions tend to target precisely those price buckets. We therefore define *quantum prices* as the property of a firm’s pricing strategy of using sparse, clustered, and sticky values, to price large and diverse product lines.¹ Said differently, the key properties are: (a) few and sparse price buckets for a large product line, (b) the same prices for different types of products, and (c) prices that are sticky across product introductions.

Panel (a) in Fig. 1 shows Uniqlo’s distribution of distinct prices and the share of products by category. Several striking observations can be made. The distribution is extremely clustered and products from highly dissimilar categories often have an identical price. Uniqlo exhibits many remarkable numbers: \$21 billion in sales (\$1 billion in the U.S.), operates 2200 stores in 22 countries (50 in the U.S.), and it is among the 100 most valued brands worldwide. But perhaps most surprising is the price number: Uniqlo’s online catalog of 1500 products has just 13 prices. Additionally, there is not a single new price point across years. That is to say, the set of 13 prices remains constant; what changes is the share (mix) of products allocated in each price bucket.

One may wonder “This discreteness could originate simply because various products have a similar unit cost.” That is, firms apply a constant markup and then round prices. Panel (b) provides a compelling visual argument against this case. In collaboration with a department store retailer, we obtained regular prices and unit costs for all dresses sold in one year. Surprisingly, there is no uniform markup: for a given price bucket there is a wide range of unit costs. Again, a substantial price clustering is seen across differentiated products.

Taking a step back from these stylized examples, quantum pricing is a novel contribution to two recent bodies of literature. First, quantum pricing is at odds with the possibility of utilizing advanced price algorithms (Brynjolfsson and McAfee, 2014). As an illustration, we collected UberX prices for a ride between Boston’s Museum of Fine Arts and the Celtics Stadium (3.6 miles, see Appendix A). In a single week, there were 156 prices (in many cases, prices varied by as little as 1 cent). While rides in various hours are not identical “products,” it is still notable to observe 156 prices in one week vs. only 13 prices for 1500 products in two years. Why are some firms not using algorithms to set prices? Presumably, dynamic markdowns may be difficult if the brand maintains uniform pricing policies across stores and online/offline (Zara, 2023). However, Uniqlo or Zara can afford to use algorithms to set regular prices (e.g., the studies of Cachon and Swinney (2011); Ban et al. (2019) develop and test assortment planning optimization models in collaboration with Zara). Nevertheless, they choose to have a sparse and discrete price grid.


Second, our findings contribute to the literature studying various aspects of discrete pricing. While we review the literature later, it is helpful to note the seminal role of price endings (Monroe, 1973; Schindler and Kirby, 1997; Thomas and Morwitz, 2005). There is also evidence of price points in supermarket scanner data (Levy et al., 2011; Eichenbaum et al., 2011; Ilut et al., 2020; Stevens, 2020), as well as uniform prices across variants of the same product (Anderson et al., 2015; Draganska and Jain, 2006). More recently, there is work showing the role of uniform prices for the same product, across physical stores within a chain (DellaVigna and Gentzkow, 2019; Hitsch et al., 2021; Aparicio et al., 2021). However, our analysis separates from these studies as follows.

To examine quantum pricing, we consider *exact* price points. That is, price coarseness is not potentially magnified by rounding weighted-average weekly scanner prices. This is important because it allows to carefully show that quantum prices exist after controlling for price endings. Additionally, we consider the pricing problem of cross-country retailers with control over the final price. This is rarely the case in supermarkets, for example, where stores and chains deal with multiple brands, minimum price requirements, and recommended prices. Similarly, the discrete pricing that we examine is not focused on the prices for the same SKU across stores (e.g., DellaVigna and Gentzkow (2019)) or the same price across color variants for the same SKU. Instead, we examine the dynamics of a firm using a handful of prices for disparately dissimilar products (a pair of shoes and a sweater). Moreover, we show that the same prices are used across seasons—Even though the cost of printing price tags for new products does not depend on which digit appears on the tag.


With these ideas in mind, quantum prices exhibit features previously not considered together: (a) sparse prices for a large assortment, (b) clustering across differentiated products, and (c) price stickiness at introduction.

We collected online data from 54 retailers in the U.S. and 40 retailers in the U.K, with a total of nearly 230,000 and 190,000 distinct products, respectively. This data represents a significant share of the apparel, footwear, and accessories industries in both countries, and to the authors’ knowledge it is the largest data collection effort in this market. Online data lacks a structured hierarchy that groups products by similarity (as opposed to Nielsen or IRI scanner data, for example). To address this problem, we

¹ The term “quantum pricing” is reminiscent of quantum mechanics, where the electrons jump from one level of energy to another in discrete, non-divisible quanta. Prices in fashion retail can be quite sticky, and move in discrete and large quantities.



(a) Uniqlo's Price Grid



(b) A Retailer's Price vs. Cost

Fig. 1. Discrete Price Distributions.

Notes: Panel (a) shows Uniqlo's (U.S.) price distribution (without rounding) and the percent of products in each price bucket. Data includes prices between \$9 and \$95, in Spring 2017 and Spring 2018. Panel (b) shows the distribution of the regular price and unit cost (without rounding), vs. the price-sorted product index. Data includes a cross-section of over 13,000 distinct SKUs in the dresses category during one year. Prices and costs are multiplied by a constant to mask the retailer identity, but the mark-up relationship is preserved.

design a semi-supervised machine learning classifier that, reading HTML code and product-level descriptions, classifies each item into twelve categories that are consistent across retailers and over time. The categories are: accessories, bags, bottoms, dresses, jewelry, outerwear, shoes, sports, suits, tees, tops, and underwear. Broadly speaking, we could interpret a retailer–category (e.g., Zara Dresses) as the relevant choice space of a consumer buying a specific product.

We begin the paper by documenting the existence of quantum prices in the fashion industry within retailer–categories (excluding duplicate variants in colors or sizes). Note that because some categories are already relatively broad (e.g., sweaters

and shirts are “Tops”), the estimates are lower bounds to using more narrowly defined categories. The average probability that two items in a retailer–category have the same price is close to 10%. We also extend the analysis to across categories within a retailer and, surprisingly, report substantial clustering in prices.

Several statistical techniques are used to test the robustness of the observed price clustering. For example, clustering could be explained by a corporate policy to use \$9 dollar endings. Thus, we use a normalized price clustering index that removes the effects from popular prices, price endings, ranges of prices, or number of products. The index is a modification of the geographical concentration index introduced in Ellison and Glaeser (1997). This measure takes the price distribution of each retailer–category and controls for shares of products concentrated on, for example, certain digit endings or popular prices in the overall category. We find that about 25% of the retailer–categories have little to no clustering, which increases to 50% when we control for price digits. Therefore, digit endings remains a popular strategy in the industry. However, about half of the remaining retailer–categories still exhibit statistically large clustering, including 22% of them which are found to be remarkably concentrated. Price clustering is also observed using an unsupervised machine learning approach. The method estimates the price clusters in the data through a trade-off in the within–cluster and between–cluster variations. We find that firms concentrate prices in a few discrete clusters, and that price clusters have economically meaningful distances. The median price differential between centroids (mid–points) is roughly 21%.

Importantly, quantum prices are documented in the time–series. We show that (a) retailers will consistently introduce new products at the same quantum prices, and (b) this concentration exists in sale prices. This implies that price changes at introduction also exhibit a degenerate distribution: either zero or very large. This behavior suggests that firms face a form of “menu cost.” In the spirit of Shugan (1980)’s “cost of thinking,” there is a “cost of opening” a new price bucket in the distribution. Although retailers use a fixed set of prices for old and new products, they will change the product mix across price buckets. That is to say, a retailer will put a larger (lower) fraction of products into the more expensive (cheaper) prices. Critically, the set of prices is identical but the proportion of products in each bucket changes. The practice of repeatedly using quantum prices amplifies deviations from the law of one price (LOP)—based on *matched* products between the U.S. and U.K. online stores. Intuitively, the prices that would hit the LOP are not available in a sparse quantum ladder.

These pricing observations are not oddities associated with isolated firms. Quantum prices are observed in a wide range of retailers, and the data is representative of the industry in terms of revenue market shares and firm heterogeneity. Fashion retail is a significant portion of the economy, accounting for about 4% of the CPI basket (food and beverages, 16%). The combined market of apparel, footwear, bags, and accessories is about \$380 billion in sales in the U.S., and £55 billion in the U.K. The data includes department stores, from luxury to low–end, and fast–fashion retailers. Moreover, the scale of some firms is remarkable; to illustrate, in 2022, Louis Vuitton had £39 billion in sales, Inditex 33 billion, and Nike \$47 billion.

Given the striking patterns that are documented, it is useful to conceptualize in which circumstances quantum pricing can be an optimal strategy. In doing so, we are mindful that it is beyond the scope of the paper to empirically disentangle the exact managerial decision–making process. Instead, we build upon stylized facts and institutional knowledge of price–setters to offer a behavioral model in which a restricted price distribution makes *advertising price salience* increasingly effective. The novelty lies in modeling the role of price salience and price recall in the context of large, diverse, and fast–moving assortments that need to be communicated to consumers. Overall, we find robust evidence that price salience explains the decision to cluster the price distribution. We also complement the conceptual discussion with theories of bounded comparability, convenient prices, demand uncertainty, and managerial inattention.

As mentioned, our work examines pricing and implications of a large, international market. While it captures nuances and dynamics in more than 65 firms, we realize that from a macro standpoint it is helpful to incorporate many industries of the economy. At the end of the paper we report simple evidence showing quantum prices in furniture, books, music, and restaurants. Future research should more systematically examine discrete pricing throughout end–consumer markets.

1.1. Related literature

This paper relates to several strands of literature. Methodologically, using online product–level data to study price frictions is similar to that in Cavallo et al. (2014), Gorodnichenko et al. (2014), and Cavallo (2018). Our work contributes to these papers in that we describe novel patterns of price setting, product introductions, and extreme forms of price rigidity in the fashion retail industry. Using a machine learning classifier to categorize products relates to previous efforts to categorize unstructured online data. For example, Cavallo (2012); Cavallo and Rigobon (2016); Aparicio and Cavallo (2019) categorize online products in order to construct price indices using CPI weights. Our work builds on this literature as we have a more comprehensive data in fashion retail, and we offer a data–based technique that can be used in more general settings to identify sub–categories that are consistent across retailers, time, and countries.

A body of literature has documented forms of price points in the distribution of prices. There is work highlighting uniform prices in retail chains due to managerial inertia (DellaVigna and Gentzkow (2019)), managerial costs or consumer fairness (Orbach and Einav (2007)), consumer fairness for different sizes (Anderson and Simester (2008)), homogenous preferences for flavors (Draganska and Jain (2006)), the role of reference prices (Eichenbaum et al. (2011); Klenow and Malin (2010)), adverse quality signaling for lower–priced goods (Anderson and Simester (2001)), and the role of variants (Anderson et al. (2015)). See also Nakamura (2008); Levy et al. (2011); Ilut et al., 2020; Stevens (2020) using scanner data. Our research provides a rather extreme version of uniform pricing, that is, across different types of products and across seasons.

Our guiding framework of price salience advertising and recall relates to prior work in marketing (Simester (1995); Shin (2005); Rhodes (2014)). However, these papers do not explore the fashion retail industry, which is characterized by large,

diverse, and fast-moving assortments that must be advertised quite frequently. Our model helps conceptualize the effectiveness of the distribution of unique prices and a large menu cost to introducing new prices. The related studies of Jung et al. (2019); Caplin and Dean (2015); Matějka and McKay (2015) also highlight the importance of accounting for limited processing capacity. We enhance the conceptual discussion borrowing ideas from demand uncertainty (Ilut et al. (2020)), bounded comparability (Hauser and Wernerfelt (1990); Eliaz and Spiegler (2011); Piccione and Spiegler (2012); Bordalo et al. (2013, 2020)), and convenient prices (Knotek (2008)). That firms limit the number of prices within a category due to overload or bounded comparability concerns is not unreasonable given the surprisingly large number of options in apparel. For example, Forever 21 sells about 2000 different styles of dresses on a given day; and even in a narrower sub-category like casual dresses there are 450 options. Consumers cannot possibly compare all product attributes; in fact, a consumer cannot even take more than a few dresses to the fitting room. Therefore, we expect consumers to rely on decision heuristics to screen out products.

Our evidence of price stickiness relates to a large empirical literature on price stickiness and on the managerial costs of updating prices. Some papers in this area include, for instance, Levy et al. (1997); Blinder et al. (1998); Zbaracki et al. (2004); Bhattarai and Schoenle (2014); Anderson et al. (2015). We show that the practice of quantum pricing restricts the strategies of price adjustments, spillovers to discount prices, and amplifies deviations from the LOP (a matched product between two countries).

Finally, price points in the distribution can be related to a literature that associates price endings with consumer demand, price recall, or value inferences. See Monroe (1973); Dickson and Sawyer (1990); Lichtenstein et al. (1993); Schindler and Kirby (1997); Thomas and Morwitz (2005). Our work also relates to a growing literature that connects price endings in supermarket products with price rigidity (Kashyap (1995); Levy et al. (1997); Knotek (2008, 2011); Levy et al. (2011); Ater and Gerlitz (2017); Snir et al. (2017)). Although we examine a different market, our findings clearly demonstrate that, despite the rise of algorithmic pricing in the online channel, digit endings continue to be popular and account for a fraction of the observed quantum pricing. Nevertheless, after removing the effect of digit endings, substantial price clustering is observed.

2. Online data

We collected online data from 53 fashion retailers in the U.S. and 40 retailers in the U.K. These retailers are representative of the industry—we cover a wide range of firms in terms of market share, style, item composition, branding, and customer base. See Appendix B.1 for the complete list of retailers.

A script is designed to search the HyperText Markup Language (HTML) public code of a retailer's website. The program automatically stores the data of each item, including product description, ID, price, sale price, promotion description, new arrivals indicator, and sales indicator. The ID is an item-specific identifier assigned by the retailer. Products that come in different colors will often have the same or a very similar ID, which we use to keep only one of them. This process rules out price clustering from nearly perfect substitutes.

Due to the large scope of retailers and computational limitations, we collect monthly data during six months for most of the retailers, and during one year or two years for a subset of the retailers. Note that a cross-section removes any price clustering explained by the same products over time.

Fashion retail is characterized by some distinctive features. Products have a short duration, which ranges from a few weeks to several months (Caro and Gallien (2010); Cavallo et al. (2014)). This likely produces an asymmetric price stickiness during the life of a good. Products are introduced at a certain price $p^{t=1}$ and often do not experience any price increase. However, products experience either temporal or permanent discounts toward the end of the season. If discounts are permanent, items have a sale price until it is discontinued, $p^{t=1} > p_s^{t=1}$; if discounts are temporal, the price will return to the regular price, $p^{t=1} = p^{t=3} > p_s^{t=2}$. We primarily focus on the regular price (introduction price), which we consider the crucial pricing decision; later, we report some analyses using markdown prices. Sales at the full price account for the largest share of revenue (Ghemawat et al. (2003)). Moreover, in the regulatory filings, firms often attribute lower financial returns to excessive markdowns.

Table 1 provides summary statistics of the data coverage. We have a cross-section of over 230,000 and 188,000 distinct products in the U.S. and the U.K., respectively. In total, there are over 350,00 products (some exact matched products are collected in both countries). There are on average over 4000 distinct products in each retailer, though there is a large heterogeneity. The 10th and 90th percentile store has 1279 and 7463 distinct products in the U.S., respectively. This illustrates the diversity in the retailers covered, since some fashion retailers sell very few items (e.g., Hermes) while others sell extraordinarily large assortments (e.g., Zara). There is also heterogeneity in the relationship between prices and products. For example, the 10th and 90th percentile store in the U.K. sells 9 and 192 products per price, respectively.

An initial requirement to study price setting is to identify what classes of products have certain prices. However, scraped online data is not structured this way. Data is collected without labels, product names are often inconsistent across retailers, and therefore we need classification rules that can group similar items together. These classifications are necessary to study cross-section and time-series pricing. For example, clustering that takes place at a popular price $\$x$ should receive little weight; but *what price* is a popular price must be learned from the overall category price distribution across retailers.

We construct a semi-supervised machine learning classifier based on decision trees to group items into 12 categories: accessories, bags, bottoms, dresses, jewelry, outerwear, shoes, sports, suits, tees, tops, and underwear. The approach is semi-supervised because there is no unequivocal procedure to validate the classification. We rely on the retailers' webpage categories and our interpretation of the product description to create these classification rules. Moreover, due to the large quantity of data, it is infeasible to manually assign labels to each product. Instead, we designed rules to check random portions of the data or products that

Table 1
Summary Statistics.

	U.S.	U.K.
Time period	March 2017 to May 2018	March 2017 to May 2018
Avg. months	4.6	4.5
Observations ^a	241,932	199,619
Retailers	54	40
Distinct prices	1126	827
Distinct goods	230,720	188,558
Avg. distinct goods	4278	4718
Distinct goods (10%pct.)	1279	1122
Distinct goods (90%pct.)	7463	13,800
Distance between prices (%) ^b	13	13
Avg. Items / Prices	49	87
Items / Prices (10%pct.)	11	9
Items / Prices (90%pct.)	92	192

Notes: Equal-weight averages across retailers. ^aExcludes duplicates in terms of country x retailer x category x product (e.g., the same product collected in two months will appear only once). ^bAverage distance between consecutive prices.

exhibit dissimilar characteristics to those in their group (e.g., items that are too expensive in the category), and then re-train and re-classify. The final output is a classifier which can consistently categorize products across retailers and across collections. See Appendix B.2 for additional details.

3. Evidence of quantum pricing: across products

This Section presents formal evidence of quantum prices or price clustering. Throughout, both terms are used interchangeably. We propose a series of clustering measures computed at the retailer or retailer–category level.

3.1. Descriptive evidence

Table 1 indicates that, overall, there are 204 and 228 prices products per price in the U.S. and U.K., respectively. This already signals that many items in the same retailer–category have an identical price. To quantify this, we randomly and repeatedly sample two distinct items from a retailer–category to estimate the probability of an identical price. Fig. 2 shows the average probability across retailer–categories. The median probability is close to 10%. Some of these magnitudes are surprisingly large (e.g., the 90th percentile is 22%), considering that some categories are relatively broad (such as Jeans vs. Bottoms). Therefore, for this probability to be this large it needs to be the case that even different items like sweaters and shirts (or jeans and chinos), which belong to the same category, have the same price.

This pricing heterogeneity across retailers is reinforced in Appendix C.1, showing a substantial variation of the probabilities (that two items in the retailer–category have an identical price) vs. assortment size. For instance, there are many retailer–




Fig. 2. Probability 2 Items Have an Identical Price.

categories where the probability exceeds 20%, for varying assortment sizes from less than 100 to over 1000. Additionally, Appendix C.2 shows that price clustering tends to decrease with brand expensiveness.

3.2. Normalized measure

Not all prices are equally good. A body of literature argues that consumers practice left-to-right processing for multiple-digit prices, and due to processing costs and lower returns to rightmost digits, consumers either drop off rightmost digits or overweight the left ones (e.g., Schindler and Kirby (1997); Thomas and Morwitz (2005)). For these reasons, \$19.99 might be perceived as closer to \$19.00 than to \$20.00. Therefore, we would like to measure price clustering after controlling for prices that are popular in the category or the retailer, or that may arise mechanically from assortment size or from a range of good prices.

We construct a normalized clustering index that builds on Ellison and Glaeser, 1997. (Ellison and Glaeser (1997) measure geographic concentration across manufacturing firms in the U.S., and would like to control for regions that are naturally better for certain industries, which in our case can be translated as a retailer–category.) The core of this index lies in comparing the observed price frequencies in a given retailer–category against the observed price frequency in the overall category. Intuitively, we want to penalize clustering that occurs at certain prices (e.g., Zara Underwear) which are popular in the category (Underwear). Formally, the index in retailer–category i is defined as follows:

$$index_i = \frac{\sum_{b=\underline{b}}^{\bar{b}} (s_{i,b} - x_{c,b})^2 - \left(1 - \sum_{b=\underline{b}}^{\bar{b}} x_{c,b}^2\right) 1/N_i}{\left(1 - \sum_{b=\underline{b}}^{\bar{b}} x_{c,b}^2\right) (1 - 1/N_i)} \tag{1}$$

We bin the distribution of prices into buckets of 1 dollar (or 1 pound), i.e. prices are rounded to the nearest integer. This is a very conservative control for price endings because prices like \$19.90 and \$19.50 will be treated as the same, and therefore penalized according to the greater overall frequency of \$20. In Eq. (1), $s_{i,b}$ is the share of items in retailer–category i at bucket b , and $x_{c,b}$ is the share of items in category c at bucket b . The sum goes from the minimum price in category c (\underline{b}_c), to the maximum price in category c (\bar{b}_c). Finally, N_i is the number of distinct products in retailer–category i , and the term $\frac{1}{N_i}$ controls for assortment size.

The index can be interpreted as the excess probability that two items in the same group will have the same price, given the size of the retailer–category N_i , and the empirical distribution of prices in the category. Because the index is normalized to be between 0 and 1, values close to 0 should be interpreted as retailer–categories not exhibiting excess price clustering. Values above 0.025 indicate statistically large price clustering (Ellison and Glaeser (1997)). We find that a fraction of retailer–categories with no more than the expected clustering, as well as a fair share of cases with medium to large clustering. The mean and median in the U.S. are 0.098 and 0.075, respectively, both of which are considered large estimates of price concentration. The histogram is shown in Appendix C.3. In general, we find that a retailer–category exhibits similar price concentration in the U.S. and the U.K. (Appendix C.4).

To further build intuition, the normalized measure shows substantial concentration compared to prices drawn from a Normal or uniform distribution with the same empirical parameters. Consider prices drawn from a uniform distribution, restricted to 10-dollar multiples—which increases clustering and thus yields a stringent benchmark. The mean values are 0.010 and 0.052 for the Normal and uniform indices, respectively; these compare to a value of 0.098 for the data-based index. See Appendix C.5 for additional evidence.

Next, we consider three alternative specifications to Eq. (1). The results, summarized in Table 2, reinforce the main findings. It is helpful to observe that Eq. (1) compares each retailer–category price bin share against the category price bin share (the first term in the numerator). We could, however, replace the category price bin share ($x_{c,b}$) with a series of alternative price market shares, and evaluate changes in the clustering index.

First, we use the price shares that would be observed under a Normal distribution. We define $x_{c,b} = \Phi(b) - \Phi(b - 1)$, where $\Phi(\cdot)$ refers to the CDF from a Normal distribution with μ as the average price in the category and σ as the standard deviation of the prices in the category (equally-weighted retailers). Clustering estimates replicate the same patterns and therefore shown in Appendix C.5. The mean and median in the U.S. are 0.106 and 0.083, respectively.

Second, we can more severely control for prices and price levels (ranges of good prices, or cheap and expensive products relative to others in the category) as follows. We estimate retailer–category Poisson regressions of price counts on price, price squared, and category shares:

$$S_{i,b} = \alpha_i + \beta_{i,1} bin_{i,b} + \beta_{i,2} bin_{i,b}^2 + X_{c,b} + e_i \tag{2}$$

$S_{i,b}$ denotes the count of items in retailer–category i priced at bin b (instead of $s_{i,b}$, which are shares), $bin_{i,b}$ is the price bin b , $bin_{i,b}^2$ is the price bin squared, and $X_{c,b}$ are count of items in category c priced at bin b . Once we estimate Eq. (2), we obtain predicted counts, $\hat{S}_{i,b}$, and convert these to predicted shares, $\hat{s}_{i,b}$, using the sum of predicted counts, \hat{N}_i . These predicted shares are used in Eq. (1) instead of the price shares $x_{c,b}$, and the normalized index is recalculated. Counts of 0 items are ignored in the regressions and forced to predict a share equal to 0.

Table 2
Normalized Price Clustering Index.

		Index (1)	Index (2)	Index (3)	Index (4)
US					
(i)	Mean	0.098	0.106	0.063	0.025
(ii)	Median	0.075	0.083	0.045	0.038
(iii)	% of retailer-categories:				
	<0	0	0	0.6	2.8
	[0,0.025)	5.4	7.9	23.6	47.2
	[0.025,0.05)	23.6	19.4	34.7	27.6
	>0.05	71.0	72.6	41.1	22.4
UK					
(i)	Mean	0.097	0.107	0.057	0.029
(ii)	Median	0.080	0.090	0.041	0.019
(iii)	% of retailer-categories:				
	<0	0	0	0	5.6
	[0,0.025)	5.9	8.1	26.9	53.8
	[0.025,0.05)	19.4	18.5	33.3	25.3
	>0.05	74.7	73.4	39.8	15.3

Notes: Index (1) is the baseline normalized index in Eq. (1). Index (2) is the normalized index using price bin shares from a Normal distribution. Index (3) is the normalized index in Eq. (2) that controls for prices, range of prices, popular prices. Index (4) adds price ending dummies.

Once again, after controlling for these practices, some firms exhibit little to no price clustering. However, many retailer-categories still exhibit substantial price clustering. The mean and median are 0.063 and 0.045 in the U.S., and 0.057 and 0.041 in the U.K., respectively. Moreover, a considerable fraction of cases has clustering greater than 0.05 (41% in the U.S. and 40% in the U.K.). The histogram is shown in Appendix C.5.

The more features included in Eq. (2), the more stringent the clustering measure, and therefore the smaller the excess price clustering that will be estimated in the data. For instance, we can investigate the portion of price clustering that is due to a firm having firm-category specific price endings policies. The third variant does exactly this. We included price ending (integer) dummy variables in the regressions, i.e. the terms $\sum_{j=0}^8 \beta_{i,E} \text{End}_{ij}$. We now find a larger number of retailer-categories with close to zero price clustering, suggesting that rightmost digits continues to be a popular heuristic (Kashyap (1995); Stiving and Winer (1997); Anderson and Simester (2003)). However, half of the retailer-categories still exhibit a notable price concentration that cannot be captured by price endings. The median estimates are 0.038 in the U.S. and 0.019 in the U.K. The histogram is shown in Appendix C.5.

Table 2 provides summary statistics of the normalized price clustering indices, which decrease from (1) to (3) and (4) as we sequentially control for more price features. In some cases, the clustering index can be slightly less than 0 if a retailer-category exhibits less price clustering than what is expected.

The results so far demonstrate sharp pricing heterogeneity. A group of firms exhibit little if any price clustering. Other firms have price clustering that can be attributed to price levels, popular prices, convenient prices, or price endings. However, there is considerable price concentration not driven by these features.² In summary, quantum pricing is a more general phenomena than convenient prices or digit endings, both of which are found to be popular practices.

3.3. Price clustering: across categories or across retailers

The measures of price concentration can also be computed within retailers and across categories, as well as across retailers and within categories. Next we discuss three results. First, we compute the probability that two random items, in different categories of the same retailer, have the same price. This empirical probability is computed sampling two items for every within-retailer category-category pair. We find a surprising amount of correlated clustering, especially if we consider that some categories are relatively broad (e.g., sweaters and shirts in Tops). In the U.S., for example, there are over 530 retailer category-category pairs (about 25% of total) with more than 5% estimated probability, the mean is 3%, and the 90th percentile is 10%. The histogram is shown in Appendix C.7.

Second, we implement the normalized measure in Eq. (1) within-retailer and across-categories, as well as across-retailers and within-category. Estimates are expected to be lower because items from two categories cannot in general be more concentrated

² Convenient prices tend to be operationalized as multiples of 25¢ and mental processing heuristics in digit endings tend to examine 99¢ not integers; whereas quantum ladders are considerably more discrete. This is consistent with Knotek (2011)'s prediction that *convenient prices should be less relevant in fashion*, who writes: "convenience is unlikely to affect traditional apparel stores, where there are fewer customers, each of whom spends a considerable amount of time shopping and transactions are often executed with credit or debit cards." Additionally, our data reveals the practice of many digit endings. If there are kinks in demand at key digits, it is puzzling that there is a kink in Bonobos at \$8 but at \$9 in Uniqlo, but the kink would not exist at \$9 and \$8, respectively. A similar analogy applies to markdown prices. Many fashion retailers set sale prices as percentages or direct price points, and there is weak relationship between price endings in regular prices and markdown prices. See Appendix C.6.

than in their own category. The normalized measure of correlated clustering is based on a modification of Ellison and Glaeser (1997) and Ellison et al. (2010), who estimate geographic concentration across manufacturing plants from different industrial sectors. The normalized index is computed as follows:

$$\text{index}_{i,j}^c = \frac{\sum_{b=\underline{b}}^{\bar{b}} (s_{i,b} - x_b)(s_{j,b} - x_b)}{1 - \sum_{b=\underline{b}}^{\bar{b}} x_b^2} \quad (3)$$

Where i and j denote two categories within the same retailer, and x_b is the average price bin share between the two categories at price bucket b . The sums go from the minimum to the maximum prices observed in either category.

We estimate meaningful price clustering across categories. For example, there are 325 category pairs (about 15%) with clustering above 0.05 in the U.S. The mean is 0.023 in the U.S and 0.031 in the U.K. This reinforces the evidence that some retailers use the same prices for very different types of products. Appendix C.7 shows the histogram of the clustering measure as well as which category pairs tend to be on average more concentrated. Importantly, this measure will only be large when distinct categories use the exact same price, not a similar price.

Finally, we calculate the normalized correlated clustering index in Eq. (3) within-category and across-retailers to examine whether retailers use the same prices for similar items. Appendix C.8 indicates no evidence of correlated clustering across retailers. For computational reasons the results are computed for a random fraction of the retailers in each country. The estimates are close to 0 in both countries. For example, the mean and median in the U.S. is 0.0 and 0.0, respectively. The lack of evidence echoes the constrained edge of price endings: if some prices were particularly appealing to consumers then one might expect different retailers concentrating on the same prices.

In sum, we estimate a statistically large degree of price clustering within a retailer-category, a significant but smaller degree across categories within the same retailer, and no clustering across retailers within the same category. Appendix C.9 overlaps the clustering measure distributions at the three levels. Table 3 provides some examples of retailers and their clustering measures.

3.4. Machine learning price clustering

In previous analyses we showed evidence of price clustering using all price buckets. But are all those buckets economically meaningful? For instance, some of these distinct prices can be too close from each other, and we might want to consider those as belonging to a same price cluster. We use an unsupervised machine learning approach to address this question.

We define a clustering index borrowing ideas from the popular k -means literature (for a review see Friedman et al. (2001)) and from the CH index (Caliński and Harabasz (1974)). See Appendix C.10 for methodological details. We define a ratio $\kappa(k) \equiv \frac{WC(n_k, k)}{BC(k)}$, which relates the within-cluster variation (a series of price buckets within cluster k) to the between-cluster variation (centroids k and $k - 1$). This method accomplishes two objectives: (a) it determines the optimal number of price clusters in the data according to a standard trade-off, and (b) clusters are separated by at least 5% from each other.

Fig. 3 shows the results for the U.S. data (see Appendix C.10 for U.K.). More specifically, we show k_i^* for all retailer-categories i , the corresponding average distance between consecutive centroids (in percentage), and the ratio of k^* to the maximum possible of clusters. Note: due to the 5% threshold, the maximum number of clusters is not the number of distinct prices (i.e., $\#b$); the maximum k^* is $\lfloor \frac{\log(b^{\max}/b^{\min})}{\log(1.05)} \rfloor + 1$. Overall, there is a sizable share of retailer-categories that exhibit medium to high price clustering. These are captured by those having a low ratio of k^* relative to the maximum possible k . Intuitively, there is more coarseness than the set of distinct prices. Then there is another set of cases that are poorly clustered and tend to exhibit large k^* . Importantly, for the vast majority of the cases, the price clusters tend to be meaningfully separated from each other. Panel (a) shows many cases where the average distance between centroids is between 10% and 30%. The median average distance between centroids is approximately 20%. Note that this separation in prices is a lot more extreme than that predicted by price endings or convenient prices, i.e. increments of 25 or 99 or even \$9.

4. Evidence of quantum pricing: introductions, sale prices, and LOP

Section 3 discussed the existence of quantum pricing in the cross-section of either similar or very dissimilar products. Does quantum pricing emerge in the time-series? One might argue that the relevance of price clusters fades if retailers simply start over every season, i.e. they are flexible to choose new price buckets for new products. However, we demonstrate a third feature of quantum prices: prices are very *sticky* over time.³

³ We provide a caveat that we use the term “stickiness” somewhat loosely to capture the behavior of a firm that prices products across seasons using a sticky menu of “old” prices. This is a different dimension from the canonical price rigidity in the New Keynesian literature, in which “individual good prices change rarely” (Kehoe and Midrigan, 2015). In this domain, see for example Kashyap (1995); Blinder et al. (1998); Bils and Klenow (2004); Klenow and Kryvtsov (2008); Nakamura and Steinsson (2008); Eichenbaum et al. (2011); Nakamura and Steinsson (2013).

Table 3
Examples of Price Clustering.

		Retailer	Items ^a Prices	Prob. ^b	Prob. ^c	Index(1) ^d
(i)	US Low clustering	Louis Vuitton	2.8	0.02	0.01	0.04
		Aritzia	8.5	0.05	0.02	0.05
(ii)	Medium clustering	Gap	20.5	0.16	0.06	0.11
		Victoria Secret	7.6	0.08	0.02	0.08
(iii)	High clustering	Bonobos	18.5	0.21	0.07	0.22
		Uniqlo	30.7	0.34	0.14	0.31
(iv)	UK Low clustering	Gucci	4.5	0.02	0.01	0.06
		Ralph Lauren	5.1	0.04	0.01	0.06
(v)	Medium clustering	Burberry	7.3	0.08	0.02	0.09
		HM	75.1	0.15	0.08	0.08
(vi)	High clustering	Zara	85.7	0.2	0.08	0.13
		Uterque	13.2	0.17	0.08	0.19

Notes: Selected retailers to illustrate low, medium, and high price clustering. The following measures are within-retailer across-category averages. ^aThe number of distinct items per distinct price. ^bAverage probability that two distinct items in the same category have an identical price. ^cAverage probability that two distinct items in different categories have an identical price. ^dNormalized clustering measure as per Eq. (1).

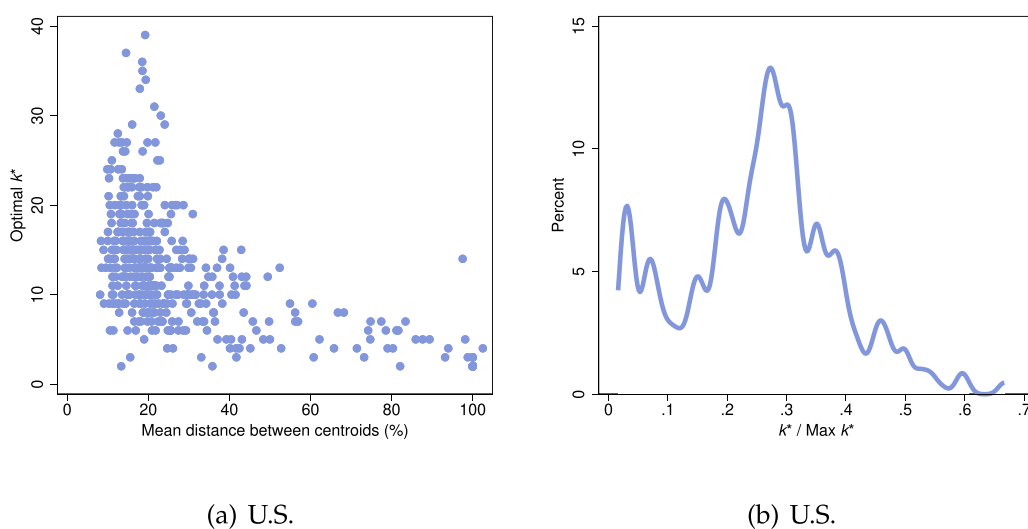


Fig. 3. Optimal Number of Price Clusters.

4.1. Pricing new products

Fashion products are often characterized by little, if any, upward price changes and a short product life. (The mean product life is 3.1 months in the U.S. and 3.2 months in the U.K. See Appendix D.1.) As a result, retailers are constantly deciding *what prices* to set for new products. This is a difficult problem for obvious reasons: uncertain demand, seasonality, cost changes, fabrics, trends, etc. Besides difficult, it is economically consequential—the introduction price accounts for a product's core revenues.

We begin by characterizing a fast pace of product introductions. We identify the products in each retailer–category's initial month, and then compute the share of the persistent catalog 1 to 5 months later. Line (i) in Table 4 shows that the average share of persistent catalog in the second month is 0.55 and in the sixth month is 0.17. A lower magnitude indicates that most of the products in that month are new.

We now ask: Given that new products are introduced very frequently (and old products are discontinued), what prices do firms set for the new products? Actually, firms tend to consistently use the same price buckets. Line (ii) in Table 4 shows a

Table 4
Price Stickiness at Introduction.

	m-th Month				
	m = 2	m = 3	m = 4	m = 5	m = 6
Panel A. Share of persisting catalog ^a					
p25%	0.42	0.28	0.21	0.02	0.00
Average	0.55	0.44	0.38	0.25	0.17
p75%	0.74	0.62	0.56	0.41	0.25
Panel B. Prob introduction at existing prices ^b					
p25%	0.83	0.81	0.78	0.77	0.80
Average	0.87	0.84	0.84	0.83	0.85
p75%	1.00	1.00	1.00	0.99	0.96

Notes: Equal-weight averages across retailer–categories in the U.S. and the U.K. ^aRatio of catalog in the first month to the catalog in month *m*. ^bProbability that the price of a product introduction in month *m* equals a price observed in the first month.

large share of common prices, relative to the first month, even though the assortment is changing. In other words, the vast majority of products are introduced at the existing prices. To illustrate, the average probability that the price of a new product in the sixth month coincides with an “old price” is 85%. Note that there are retailer–categories where the estimate is 0, i.e. there is not a single “new price.”

Does quantum pricing exacerbate this price stickiness? Although the “menu cost” for printing a tag with \$24 or \$26 are the same, if the old price grid had \$24 (and not \$26), then a new product is more likely to be priced at \$24. We estimate the following fixed-effects model:

$$Prob.Sticky_{r,c} = \alpha + \beta Index_{r,c} + \gamma_c + \delta_u + e_{i,t} \tag{4}$$

Where *Prob. Sticky_{i,t}* is the average of the *m*-based probabilities (line (ii) in Table 4) and *Index_{r,c}* denotes the baseline measure of price concentration at the retailer–category level (Eq. (1)). Additionally, γ_c and δ_u capture category and country fixed effects, respectively. We include U.S. and U.K. observations.

Column (1) of Table 5 indicates strong evidence that quantum pricing amplifies price stickiness at introduction. Said differently, retailers that use sparse and clustered prices across differentiated products, are more likely to use those identical prices over time—Even though the effort of printing a tag for new products is the same regardless of the digit on the paper.

4.2. Re-pricing through product shares in price buckets

Prices of raw materials often experience large price swings. To illustrate, cotton prices increased 20% year-on-year as of April 2018; and then prices decreased 26% year-on-year as of July 2019 (Financial Times (2018); Bloomberg (2018, 2019)). Although we lack unit costs in the data, it is reasonable to assume that cost shocks partly impact retail prices. How might retailers adjust average prices with a discrete price grid?

We motivate the discussion using evidence from two stylized retailers: Uniqlo and Ralph Lauren. Panel (a) in Fig. 4 compares the prices in Uniqlo U.S. for the same categories over two years (same month to account for seasonality, and bars sum up to 100% each year). Uniqlo, which is characterized by strong measures of price clustering (Table 3), appears to adjust prices by *changing shares of products in the existing prices*. There is so much price stickiness at introduction that over 90% of the change in the price distribution occurs via modifying the product shares in the old price buckets. On the other hand, Ralph Lauren U.S. in Panel (b), prices are spread out across many points in a grid. In this case, changes in the price distribution are more evenly split between shares in the same prices and new price buckets (close to 50% each).⁴

The observation about changes in the price distribution can be generalized. We use all retailers with one year of data, allowing for a comparable assortment year-on-year. The change in the price distributions in each retailer–category is decomposed as follows. Let $w_{i,1}$ and $w_{i,2}$ denote the shares of products located in price *i* (no rounding) in time 1 and time 2, respectively. The price distribution is decomposed as:


$$\sum_{i=p^{\min}}^{p^{\max}} |w_{i,2} - w_{i,1}| = \underbrace{\sum |w_{i,2} - w_{i,1}|_{w_1 \cap w_2 \neq \emptyset}}_{\text{shares in overlap price ladder}} + \underbrace{\sum |w_{i,2} - w_{i,1}|_{w_1 \cap w_2 = \emptyset}}_{\text{shares in non-overlap price ladder}} \tag{5}$$

⁴ Overall, that price changes in fashion are observed during events of product introductions (vs. during product life) echoes the patterns explained by the seminal work of Nakamura and Steinsson (2008): “Price changes due to product substitutions are a second class of price changes that we argue is fundamentally different from the regular price changes typically emphasized by macroeconomists. This source of price flexibility is particularly important for durable goods. For example, the spring and fall clothing seasons in apparel and the new model year for cars are associated with a large number of price changes due to the introduction of new products.”

Table 5
Quantum Pricing and Stickiness at Introduction.

	(1)	(2)	(3)
Price Index ^a	0.301*** (0.077)	0.891*** (0.293)	0.778** (0.361)
Observations	802	73	73
R ²	0.043	0.371	0.294
Country FE	YES	YES	YES
Category FE	YES	YES	YES

Notes: ^aNormalized clustering measure as per Eq. (1). Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.



(a) Uniqlo U.S.

(b) Ralph Lauren U.S.

Fig. 4. Re-Pricing Through “Product Shares in Price Buckets”.

The term in the left computes the sum of the absolute differences between the shares of products in the overlap price ladder, i.e. the set of price buckets that exist in both periods. The term in the right computes the sum of the absolute differences located in prices that are observed in only one of the periods. We then compute the fraction that each term represents in the price change distribution. The median proportion of shares in the same price bins is 60%.

Mirroring the previous analysis, we now examine whether quantum pricing exacerbates this price stickiness for pricing new products. Therefore, we reestimate Eq. (4) with the outcome $OverlapPrices_{r,c}$, which denotes the share of overlap prices in retailer r and category c . A larger value captures that a set of identical prices are used throughout products introductions. Column (2) of Table 5 shows, once again, a strong relationship between quantum pricing and price stickiness at introduction.

As robustness, we propose a second measure of support in the price distribution across periods. Similarly, $w_{i,1}$ and $w_{i,2}$ denote the shares of products located in price i in time 1 and time 2, respectively. The minimum support in the price distribution can be computed as:

$$\sum_{i=p^{min}}^{p^{max}} \min_w (w_{i,1} - w_{i,2}) \tag{6}$$

For every single price in the distribution of the retailer–category, we identify the minimum share of products located at that price (between the first and second period). For instance, if $w_{12,4,1} = 15\%$ and $w_{12,4,2} = 5\%$, then we keep 5%. We then sum these minimum shares across all possible prices.

We reestimate Eq. (4) with the outcome $MinSupport_{r,c}$, which denotes the share of overlap prices in retailer r and category c . As with the previous measure, column (3) of Table 5 shows a strong relationship between quantum pricing and the support in the price distribution over time. In summary, quantum pricing presents a novel source of price stickiness at introduction: retailers are reluctant to “open” price buckets for new products, instead they will allocate products into the existing buckets. This evidence can microfound a behavioral pricing model whereby a firm, unable to set an optimal price, chooses the closest from the existing price menu. Later in the paper we discuss conceptual frameworks.

Table 6
Quantum Pricing in Markdown Prices.

	(1)	(2)
Price Index ^a	0.728*** (0.091)	
Prob. Equal Price ^b		0.655*** (0.071)
Observations	637	637
R ²	0.317	0.348
Country FE	YES	YES
Category FE	YES	YES
Assortment Size	YES	YES

Notes: ^aNormalized price clustering measure as per Eq. (1); ^bAverage probability that two random items in a retailer–category have identical price. Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

4.3. Markdown pricing

The most important pricing decision in fashion retail is setting the introduction (regular) price. The second is the discount price. Thus far, we have studied *regular prices*—But does quantum pricing exist in *markdown prices*? In the analyses that follow, we shift attention to the discount price. The data is well-suited to evaluate clustering in sale or markdown prices because fashion is characterized by declining prices throughout the life cycle. In our data, 74% (U.S.) and 58% (U.K.) of items exhibit at least one markdown price (see Appendix D.1).

We find substantial concentration in markdown prices. To show this, we first compute an analogous measure of concentration in prices to that in Section 3. In particular, the probability that two random items in the same retailer–category have an identical discount price (conditional on observing a discount). The average probability is 11% (and the 90th percentile is 25%). The histogram is shown in Appendix D.2.

We then estimate the following fixed-effects model:

$$\text{Prob. Equal Price}_{r,c} = \alpha + \beta \text{Index}_{r,c} + \gamma_c + \delta_u + e_{r,c} \quad (7)$$

Where $\text{Prob. Equal Price}_{r,c}$ is the probability that two random items in retailer r and category c have an identical sale price; $\text{Index}_{r,c}$ denotes the baseline measure of price concentration (Eq. (1)); and γ_c and δ_u capture category and country fixed effects, respectively.

Column (1) in Table 6 shows that retailer–categories with high price concentration in regular prices also exhibit high concentration in sales prices. Column (2) shows that the effect is similar if we use the empirical probability that two random items have the same (regular) price. For visual evidence plotting the price clustering patterns in regular prices vs. discount prices, see Appendix D.3. One may also interpret these findings as indicating that the more sparse and clustered the price grid, the more it spills over to other managerial decisions, such as distinct categories, new products, or discount prices.

4.4. Deviations from the law of one price

Fashion retailers can potentially reduce international pricing frictions from a discrete price grid by designing items to hit specific prices. Still, this time-to-market cannot be perfect: there are exchange rate movements, tariffs, taxes, and distribution costs. And as a result, we expect to observe deviations from the law of one price. These factors, however, are industry-specific. The question, therefore, is whether quantum pricing exacerbates deviations from the LOP. The visual examples in Figs. 1, 3, and 4 suggest that sparse price buckets might inflate the delta-LOP. Intuitively, the prices that would be close to the LOP are simply not available in the menu of quantum prices, compared to a platform's continuous price line.

The data allows to test for product-level LOP because each product has a unique ID which can be utilized to perfectly match the identical product in the U.S. and U.K. online stores, on the same day. To motivate the discussion, we return to the stylized case of Uniqlo and Ralph Lauren.

Uniqlo's and Ralph Lauren's LOPs are shown in Fig. 5. In particular, we compute the percent of matched products assigned to each combination of U.S. dollars and U.K. prices. In both panels the color key goes from the same range of 0% to 20% of products. Darker regions indicate a larger share of products assigned to a given bucket. Panel (a), which corresponds to Uniqlo, depicts large and discrete price increments between prices. In fact, a handful of buckets characterize Uniqlo's pricing across countries. In contrast, Ralph Lauren in Panel (b) exhibits a richer range of prices which can more flexibly accommodate exchange rate movements or local taste differences.

We formalize these intuitions by extending the analysis to all retailers in the data and estimating a fixed-effects model:

$$\text{RER}_{i,t} = \alpha + \gamma_{r,t} + \delta_c + N_{r,c} + \text{Sales}_{r,c} + \varepsilon_{i,t} \quad (8)$$

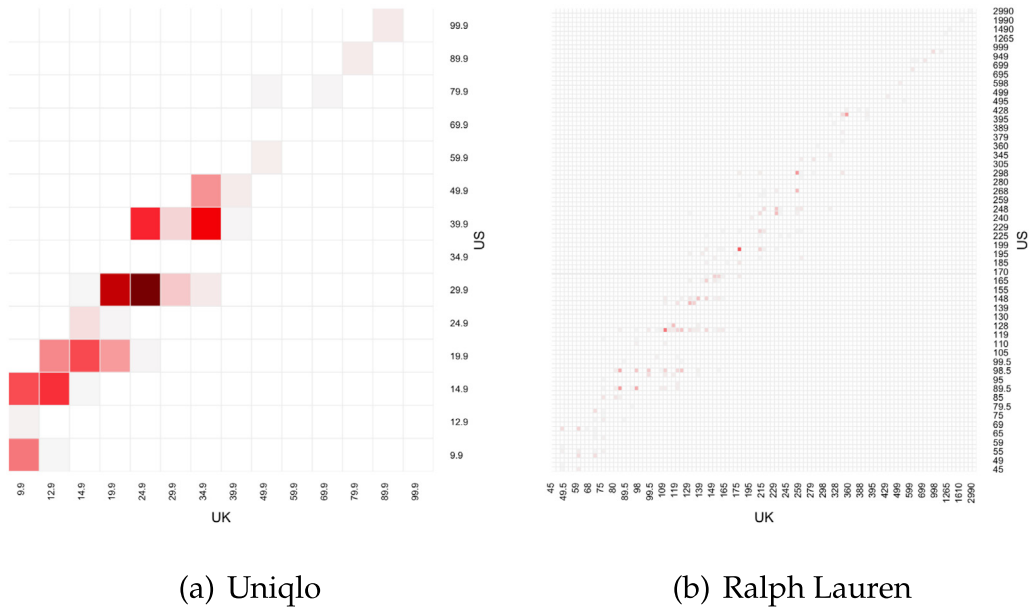


Fig. 5. Law of One Price.

Where $RER_{i,t}$, defined as $\log(p_{i,UK}) - e_{US,UK} - \log(p_{i,US})$, denotes the good-level real exchange rate (RER) for item i in month t in retailer r and category c , and $e_{US,UK}$ denotes the log of the value of the (average monthly) nominal exchange rate between the US and the UK. Additionally, $\gamma_{r,t}$ and δ_c control for retailer-month and category fixed effects, respectively, and $N_{r,c}$ and $Sales_{r,c}$ account for the (log) average assortment size and average markdown probability in the retailer-category, respectively. We consider the first month of each item i , i.e. the introduction period. When $RER_{i,t}$ takes values close to 0, it indicates no deviation from the LOP.⁵ The results are similar if we estimate Eq. (8) with the relative price, defined as the product-level ratio p_{UK}/p_{US} .

Next we obtain the absolute residuals, namely $\hat{\varepsilon}_{i,t} \equiv |RER_{i,t} - \hat{RER}_{i,t}|$, and evaluate the extent to which quantum pricing exacerbates these deviations:

$$\hat{\varepsilon}_{i,t} = \alpha + \beta Index_{r,c} + e_{i,t} \tag{9}$$

Our variable of interest is $Index_{r,c}$, which stands for the baseline measure of price concentration in the retailer r and category c (return to Eq. (1)), averaged across U.S. and U.K. Note that Eq. (8) already captures FX movements and retailer-specific effects. (As shown in Appendix C.4, the retailer-category price concentration measure is similar in the U.S. vs. the U.K.)

The results are shown in Table 7. The estimates in column (2), which include detailed covariates, indicate that an additional 1-point in the clustering index increases the absolute LOP deviations by roughly 8%. As per column (3), the results are similar if we use the empirical probability that two items have an identical price. In summary, we find evidence that firms with greater price concentration exhibit significantly wider deviations from the LOP. We emphasize that our analysis is silent about whether/how these deviations are costly for the firm. We encourage scholars to further examine how a firm's reluctance to expand the price menu enhances our understanding of monetary and international economics using good-level data.

5. Theories of quantum pricing

We lack the data to empirically test the underlying managerial process. Here we set forward a guiding behavioral framework, where its novelty is to explicitly consider the economies of advertising price salience and of price recall. While it captures the core empirical findings, it is not the only one. In the Appendix we discuss complementary models to conceptualize the stylized facts of quantum pricing: convenient prices that generate demand kinks at certain prices (Kashyap (1995); Knotek (2008); Levy et al. (2011); Knotek (2011)); bounded comparability of attributes (Eliaz and Spiegel (2011); Piccione and Spiegel (2012); Bordalo et al. (2013, 2015)); demand uncertainty that makes firms more likely to use prices that were successful in the past (Ilut et al. (2020)); and managerial inattention and mark-up rules (Fershtman and Kalai (1993); Bloom and Van Reenen (2007); Ellison et al. (2018)).

⁵ Overall, we find large good-level deviations from the LOP. The mean and median absolute good-level RER is 0.20 and 0.19 log points, respectively. See recent studies using micro data, e.g. Imbs et al. (2005); Gopinath and Rigobon (2008); Cavallo et al. (2014); Gorodnichenko and Talavera (2017). Prices in the U.K. are inclusive of VAT (value added tax). Thus, we adjust prices in the U.S. by the average state sales tax of 5.1%. State-level sales taxes were obtained from the Tax Foundation. Throughout, we exclude products with no match in both countries.

Table 7
Deviations from LOP.

	(1)	(2)	(3)
Price Index ^a	2.868*** (0.895)	8.290*** (0.581)	
Prob. Equal Price ^b			3.454*** (0.406)
Observations	52,541	52,541	52,541
R ²	0.001	0.025	0.027
Retailer x Month FE	NO	YES	YES
Category FE	NO	YES	YES
Assortment Size	NO	YES	YES
Markdown	NO	YES	YES

Notes: ^aNormalized price clustering measure as per Eq. (1); ^bAverage probability that two random items in a retailer–category have identical price. Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

5.1. Advertising effectiveness and recall

We propose a model in partial equilibrium where quantum pricing makes advertising price salience more effective; and it is increasingly more so when the same prices are used across categories and across seasons. In other words, there is a friction to opening a new price (to the price distribution) for product introductions.


Fashion companies invest substantially on media advertisement (Ghemawat et al. (2003); Wall Street Journal (2017)). However, only some retailers are salient about prices. Consider two opposite examples. In Uniqlo, prices have a central role in advertising and signage. This is true online as well as in physical stores. Prices are prominently displayed in the ads, promotions, and signs across the store. In fact, there are employees—runners—who go around the store rearranging items such that prices, sizes, and colors are in the right location and visible to the consumer. Louis Vuitton, on the other hand, makes prices the least salient attribute. Price labels are carefully obfuscated, you often need to ask for help to a merchant, there are no markdown racks, and prices are never displayed in the ads. See Fig. 6 for visual examples.

Although prior literature examines whether and how advertising should include prices (Lal and Matutes (1994); Simester (1995); Shin (2005)), the interaction between the distribution of unique prices, price rigidity, and consumer memory has received little attention. A relevant exception is Rhodes (2014), which introduces a model where consumers update beliefs on prices based on the advertised prices. Although not in pricing, several recent papers (Jung et al. (2019); Caplin and Dean (2015); Matějka and McKay (2015)) study optimal discrete actions when information processing is constrained (Shannon’s channel capacity).

Consider a retailer that sells a multi-product assortment C. Let C be defined as follows.

$$C \equiv \left\{ (x_{1,1}, \dots, x_{1,N_1}), \dots, (x_{M,1}, \dots, x_{M,N_M}) \right\}$$

$x_{i,m}$ is product i in category m and N_m is the number of products in this category. Then $N = N_1 + \dots + N_M$ denotes the number of products in the catalog. The cost of advertising is $g(A)$, where A denotes the number of distinct prices and products included in the advertisement. We can think of A as “bits of information” (Jacoby (1984)) competing for a space in the ad. As illustrated in Fig. 6, either with or without prices, newsletters only advertise a handful of items.



(a) Uniqlo (b) Abercrombie (c) Louis Vuitton (d) Gucci

Fig. 6. Advertising Examples.

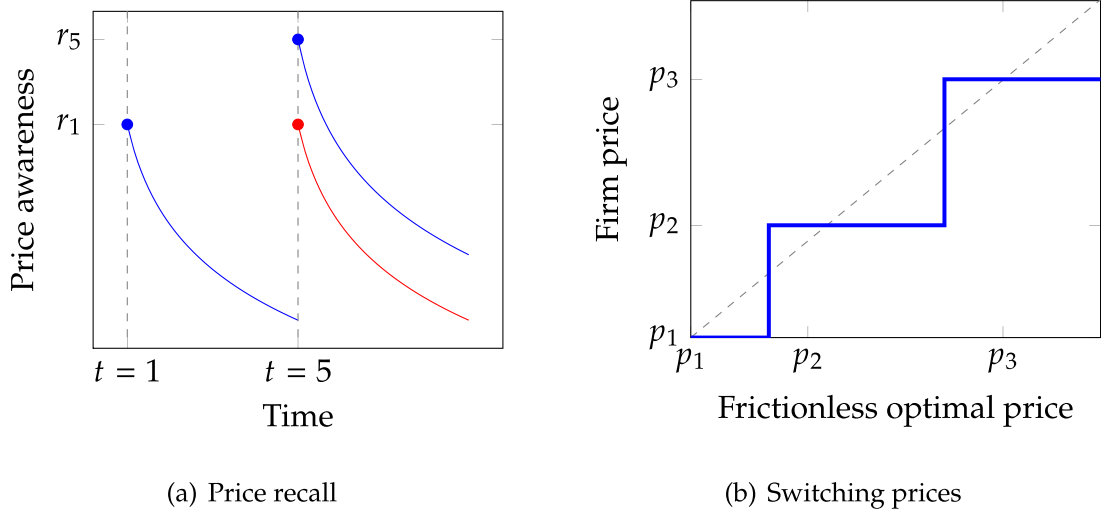


Fig. 7. Price Recall and Advertising. Notes: Panel (a) shows the relationship between price awareness and advertising. At time $t = 5$, the retailer can advertise a different price (red) and achieve r_1 awareness, or advertise the same price (blue) and increase awareness to r_5 . Fewer prices allows to more effectively advertise different products together, and, if done over time, to more consistently build a price positioning. Panel (b) shows three advertised prices p_1, p_2, p_3 . The retailer will stick to the advertised price p_1 and, if switching to a different price, it will jump discretely to p_2 which has been used before.

There are two types of consumers. A fraction λ are price-sensitive, of which γ_{ps} are informed and $(1 - \gamma_{ps})$ are only aware through advertising. And a fraction $(1 - \lambda)$ are quality-sensitive (γ_{qs} are informed). Consumers' product or price memory, or brand positioning awareness, fades over time but is updated with new advertisements. For evidence on frictions to price recall, see Monroe (1973); Dickson and Sawyer (1990); Lichtenstein et al. (1993); Monroe and Lee (1999). Recall (r) decays exponentially. An initial ad reaches $r_1 = r_0$ of the $(1 - \gamma)$ uninformed. In the next period, if no new ad is seen, then $r_t = r_0\alpha^{t-1}$; and when the same ad is seen, then $r_t = r_0\alpha^{t-1}(1 - \beta) + \beta$; where $\alpha, \beta < 1$.

Panel (a) in Fig. 7 provides visual guidance. An initial ad of x, p reaches r_1 price awareness. Price recall starts fading afterwards until the retailer sends out a new advertisement. If the retailer advertises a new price, p' , awareness jumps back to the r_1 level; while advertising the old price, p , it increases to r_5 . This can be thought of as a menu cost to introducing a new price; it is the cost of giving away the price positioning that was previously achieved through signage and advertising. Note: it might be tempting to jam or saturate the consumer with advertising to boost recall, but this is likely to increase dissatisfaction, as shown by empirical evidence on *pulsing*, i.e. advertising that oscillates from high to low levels (Dubé et al. (2005)).

The problem of the firm is to choose prices and advertising as follows:

$$\max_{A, p} \sum_m \sum_i^{N_m} (p_{i_m} - c_{i_m}) \left\{ \lambda \left[\gamma_{ps} + (1 - \gamma_{ps}) 1(i_m \in A) \right] \right\} D_{ps}(p_{i_m}) + (p_{i_m} - c_{i_m}) \left\{ (1 - \lambda) \left[\gamma_{qs} + (1 - \gamma_{qs}) 1(i_m \in A) \right] \right\} D_{qs}(p_{i_m}) - g(A) \tag{10}$$

where p_{i_m} and c_{i_m} are the price and cost of product i in category m ; $1(i_m \in A)$ is an indicator variable if product (or price) i_m was included in the advertisement. For tractability $D_j(p_{i_m})$ is a linear demand, $1 - b_j p_{i_m}$, for $j = ps, qs$.

The model allows for simple predictions that reflect institutional details of price setters in the industry. For example, when consumers are price-sensitive and informed, retailers do not cluster prices. Platform strategy retailers (e.g., Walmart's everyday low prices), focus on getting the lowest price, and therefore have no incentives to stick to price points over time. When consumers are price-sensitive, uninformed, and advertising is costly, firms cluster prices. Retailers like Uniqlo and Zara will use few prices for advertising efficiencies, and it will be costly to use different prices for new products. When retailers have both price- and quality- sensitive consumers, firms will not cluster prices. Retailers like Ralph Lauren and Armani have low/high labels and a wide price distribution to attract a broad customer base. We formalize some of these cases as follows. Proofs for Results (1)–(4) are shown in Appendix E.1.

Result 1. In a single-period two-product problem, for a set of parameters where consumers are price-sensitive and advertising is costly, it is optimal to cluster prices: $p_1^* = p_2^*$.

In Result 1, the retailer clusters prices but also restricts the product space. The retailer could produce a high-quality product; but its core demand is price-sensitive, and therefore it is better off selling low-price products at the same price. Price clustering allows to more efficiently inform consumers about a cross-section of products.

Result 2. In a single-period two-product problem, for a set of parameters where there are price- and quality- sensitive consumers, it is not optimal to cluster prices: $p_1^* < p_2^*$.

In Result 2, the retailer produces high-end and low-end products to appeal to both types of consumers. This resembles the standard vertical differentiation solution.

Result 3. In a two-period two-product problem, for a set of parameters where consumers are price-sensitive and uninformed, and there are advertising efficiencies over time, it is optimal to use sticky prices: $p_1^t = p_2^t = p^*$.

In Result 3, the retailer overcomes uninformed consumers using sticky prices. A new product is introduced in each period at the same price. Sticky prices allow to build price awareness. The effect is expected to be larger with more products (cross-section information gains).

Result 4. In a two-period problem, two products in t_1 and one product in t_2 , for a set of parameters where consumers are price-sensitive and uninformed, it is optimal to switch price to advertised prices: $p_3^t = p_2^t$.

In Result 4, the retailer is first selling two products at different prices (costs are such that $p_1^{t*} < p_2^{t*}$). In the second period, the retailer faces unanticipated increased cost for a new product. Instead of setting the frictionless $p_3^{t*} < p_2^{t*}$, it is optimal to jump to the advertised p_2^{t*} . Intuitively, it picks the closest price from the existing price menu. This is illustrated in Panel (b) in Fig. 7. The retailer jumps between advertised prices instead of introducing new prices.

A behavioral model with price memory and advertising efficiencies can conceptualize price clustering within- and across- categories. It also predicts a reluctance to introduce new prices over time, even for product introductions—which obviously require a price tag. Interestingly, price adjustments should be lumpy: if switching a portion of new products to a different price, these will jump to an existing bucket in the price ladder. Next, we provide some field evidence that price salience explains higher measures of price concentration.

5.2. Evidence of price salience

We make empirical progress in supporting the role of price salience. We define price salience as the behavior of a retailer that explicitly advertises prices in their communications and signage. To operationalize it, we enrolled in e-mail newsletters from all retailers, and then identified whether these included prices or not (promotions in terms of percentages without prices were classified as non-price salient). We complemented the newsletters with a simple online exercise: we checked whether the landing page included price points or not. Again, see Fig. 6 for some examples. Approximately 36% of the retailers in the data are price salient.

We next gauge the degree to which price salience magnifies price clustering:

$$Index_{r,c} = \alpha + \beta Salience_r + e_{r,c} \tag{11}$$

Where $Index_{r,c}$ denotes the baseline measure of price concentration for retailer r and category c (Eq. (1)) and $Salience_r$ is an indicator variable that takes value 1 if retailer r is classified as price salient.

The results are shown in Table 8. We find that price salience lifts up the normalized price clustering measure by 13%. We obtain a similar pattern of results when we reestimate Eq. (11) with: the probability that two random items have an identical price (Fig. 2), the probability that two random items have an identical markdown price (Table 6), the average probability of price stickiness at introduction (Table 5), the proportion of shares in the same price buckets (Eq. (5)), and the min support in the price distribution (Eq. (6)). These estimates are shown in columns (2) to (6), respectively, in Table 8. That is to say, firms that are price salient are significantly more concentrated in regular and markdown prices, and significantly more price sticky at introduction.

5.3. Summary

Models of convenient prices, demand uncertainty, salience and bounded rationality, and inattention (Appendix E.2) predict forms of price clustering that are consistent with the data, but generally do not account for across-categories or for sticky product introductions. See Appendix E.6 for a summary table. For example, digit-ending pricing is abundant in the data—as predicted by models with demand kinks. However, after controlling for price endings, we still observe substantial price concentration. Moreover, why would we observe identical kinks across retailer-categories but not across retailers? The models listed above do not attend to the effects of price signage. International fashion retailers often sell large and diverse assortments with short life cycle, and have complete control over the final price (vs. supermarkets, etc.). As a result, firms must constantly inform consumers

Table 8
Quantum Pricing and Price Salience.

	(1)	(2)	(3)	(4)	(5)	(6)
Price Salience	0.013** (0.006)	0.036*** (0.007)	0.023** (0.010)	0.073*** (0.011)	0.339*** (0.044)	0.397*** (0.053)
Constant	0.093*** (0.003)	0.095*** (0.004)	0.092*** (0.006)	0.849*** (0.009)	0.437*** (0.036)	0.449*** (0.038)
Observations	804	804	804	782	65	65
R ²	0.005	0.029	0.006	0.035	0.436	0.453

Notes: Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.


about new products—and presumably, this is costly. When firms are salient about prices, one strategy is to pick a niche set of prices and stick to those for various products and maintain them over time. In a sense, there is a “menu cost” to opening a new price bucket. The behavioral role of price salience finds empirical support (Table 8), though it obviously cannot model all managerial nuances and we hope it sparks further debate.

6. Other markets: music, books, furniture, and restaurant menus


This paper examined the fashion retail market. However, assessing whether quantum pricing is a more general phenomenon is essential from a macro standpoint. Here, we use data from four additional markets to show the existence of quantum prices. More precisely, we collected online data from the following firms:

- Apple iTunes store: top 197 music albums.
- Barnes and Noble: top 2111 best-seller books (hardcover).
- Ikea: 7193 furniture, decoration, and kitchen items.
- Cheesecake Factory: 307 items from DoorDash's restaurant menu.


Following our methods in Section 3, we estimate that the average probability that two distinct items of the same firm have an identical price is about 7%. This estimate is four to five times higher than that predicted by a uniform distribution (in 1-dollar increments). Thus, there is strong evidence of price clustering.




(a) Apple iTunes



(b) Barnes and Noble



(c) Ikea



(d) Cheesecake Factory

Fig. 8. Discrete and Clustered Price Distributions in Other Markets.

It is perhaps revealing to directly observe the price distributions. As captured by the bar spikes (bins of \$0.5), Fig. 8 shows severely clustered and discrete price distributions. For instance, 16% of Apple's albums cost \$12, 14% of BN's hardcover books cost \$30, 9% (28 cakes) of Cheesecake Factory's cost \$59, and 7% (485 items) of Ikea's cost \$15. We provide a caveat that these datasets are not intended to capture representative patterns of the economy; instead, it delivers evidence that quantum pricing transcends beyond fashion retail.

7. Conclusions

This paper built a novel dataset with over 350,000 different products from over 65 retailers in the U.S. and the U.K. to study pricing strategies in the fashion retail industry. The data combines three pieces that are rarely available together: (i) a large-scale cross-section of products that are representative of the diverse firms in the industry; (ii) a time-series of catalogs; and (iii) thousands of matched products across countries.

Many firms practice what we define as *quantum pricing*: (a) setting prices that are sparse and separated by large increments, (b) the same handful of prices are used across different categories of products, and (c) the same prices are used for product introductions over time. Moreover, firms that “bunch” regular prices do so in discount prices. Quantum pricing can be thought of as a novel source of stickiness—firms are reluctant to “open a new price”—suggesting a friction in the pricing decision, with implications for price adjustments and LOP deviations across countries. Finally, we discussed a behavioral model of price salience and recall.

Data availability

QP (Original data) (Mendeley Data).

Declaration of Competing Interest

None.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jinteco.2023.103770>.

References

- Anderson, E.T., Simester, D.I., 2001. Price discrimination as an adverse signal: why an offer to spread payments may hurt demand. *Mark. Sci.* 20 (3), 315–327.
- Anderson, E.T., Simester, D.I., 2003. Effects of 9 price endings on retail sales: evidence from field experiments. *Quant. Mark. Econ.* 1 (1), 93–110.
- Anderson, E.T., Simester, D.I., 2008. Research note—does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes. *Mark. Sci.* 27 (3), 492–500.
- Anderson, E., Jaimovich, N., Simester, D., 2015. Price stickiness: empirical evidence of the menu cost channel. *Rev. Econ. Stat.* 97 (4), 813–826.
- Aparicio, D., Cavallo, A., 2019. Targeted price controls on supermarket products. *Rev. Econ. Stat.* 103 (1), 60–71.
- Aparicio, D., Metzman, Z., Rigobon, R., 2021. The pricing strategies of online grocery retailers NBER WP 28639.
- Ater, I., Gerlitz, O., 2017. Round prices and price rigidity: evidence from outlawing odd prices. *J. Econ. Behav. Organ.* 144, 188–203.
- Ban, G.-Y., Gallien, J., Mersereau, A.J., 2019. Dynamic procurement of new products with covariate information: the residual tree method. *Manuf. Serv. Oper. Manag.* 21 (4), 798–815.
- Bhattarai, S., Schoenle, R., 2014. Multiproduct firms and price-setting: theory and evidence from US producer prices. *J. Monet. Econ.* 66, 178–192.
- Bils, M., Klenow, P.J., 2004. Some evidence on the importance of sticky prices. *J. Polit. Econ.* 112 (5), 947–985.
- Blinder, A., Canetti, E.R., Lebow, D.E., Rudd, J.B., 1998. Asking about Prices: A New Approach to Understanding Price Stickiness. Russell Sage Foundation.
- Bloom, N., Van Reenen, J., 2007. Measuring and explaining management practices across firms and countries. *Q. J. Econ.* 122 (4), 1351–1408.
- Bloomberg, 2018. Apparel Retailers, Finally Recovering, Eye Cotton Prices Warily June 1, 2018.
- Bloomberg, 2019. Cotton Crash Looks Set to Deepen as Stockpiles Keep Growing July 9, 2019.
- Bordalo, P., Gennaioli, N., Shleifer, A., 2013. Salience and consumer choice. *J. Polit. Econ.* 121 (5), 803–843.
- Bordalo, P., Gennaioli, N., Shleifer, A., 2015. Competition for attention. *Rev. Econ. Stud.* 83 (2), 481–513.
- Bordalo, P., Gennaioli, N., Shleifer, A., 2020. Memory, attention, and choice. *Q. J. Econ.* 135 (3), 1399–1442.
- Brynjolfsson, E., McAfee, A., 2014. The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. WW Norton & Company.
- Cachon, G.P., Swinney, R., 2011. The value of fast fashion: quick response, enhanced design, and strategic consumer behavior. *Manag. Sci.* 57 (4), 778–795.
- Caliński, T., Harabasz, J., 1974. A dendrite method for cluster analysis. *Commun. Stat. Theor. Methods* 3 (1), 1–27.
- Caplin, A., Dean, M., 2015. Revealed preference, rational inattention, and costly information acquisition. *Am. Econ. Rev.* 105 (7), 2183–2203.
- Caro, F., Gallien, J., 2010. Inventory management of a fast-fashion retail network. *Oper. Res.* 58 (2), 257–273.
- Cavallo, A., 2012. Online and official price indexes: measuring Argentina's inflation. *J. Monet. Econ.* 60 (2), 152–165.
- Cavallo, A., 2018. Scraped data and sticky prices. *Rev. Econ. Stat.* 100 (1).
- Cavallo, A., Rigobon, R., 2016. The billion prices project: using online prices for measurement and research. *J. Econ. Perspect.* 30 (2), 151–178.
- Cavallo, A., Neiman, B., Rigobon, R., 2014. Currency unions, product introductions, and the real exchange rate. *Q. J. Econ.* 129 (2), 529–595.
- DellaVigna, S., Gentzkow, M., 2019. Uniform pricing in us retail chains. *Q. J. Econ.* 134 (4), 2011–2084.
- Dickson, P.R., Sawyer, A.G., 1990. The price knowledge and search of supermarket shoppers. *J. Mark.* 42–53.
- Draganska, M., Jain, D.C., 2006. Consumer preferences and product-line pricing strategies: an empirical analysis. *Mark. Sci.* 25 (2), 164–174.
- Dubé, J.-P., Hitsch, G.J., Manchanda, P., 2005. An empirical model of advertising dynamics. *Quant. Mark. Econ.* 3 (2), 107–144.
- Eichenbaum, M., Jaimovich, N., Rebelo, S., 2011. Reference prices, costs, and nominal rigidities. *Am. Econ. Rev.* 101 (1), 234–262.
- Eliaz, K., Spiegler, R., 2011. Consideration sets and competitive marketing. *Rev. Econ. Stud.* 78 (1), 235–262.
- Ellison, G., Glaeser, E.L., 1997. Geographic concentration in US manufacturing industries: a dartboard approach. *J. Polit. Econ.* 105 (5), 889–927.
- Ellison, G., Glaeser, E.L., Kerr, W.R., 2010. What causes industry agglomeration? Evidence from coagglomeration patterns. *Am. Econ. Rev.* 100 (3), 1195–1213.

- Ellison, S.F., Snyder, C.M., Zhang, H., 2018. Costs of Managerial Attention and Activity as a Source of Sticky Prices: Structural Estimates from an Online Market NBER WP24680.
- Fershtman, C., Kalai, E., 1993. Complexity considerations and market behavior. *RAND J. Econ.* 224–235.
- Financial Times, 2018. Cotton Jumps as Storm Alberto brings Heavy Rain to Growing Areas May 29, 2018.
- Friedman, J., Hastie, T., Tibshirani, R., 2001. *The Elements of Statistical Learning*, vol. 1. Springer series in statistics, New York.
- Ghemawat, P., Nueno, J.L., Dailey, M., 2003. ZARA: Fast Fashion. vol. 1. Harvard Business School Boston, MA.
- Gopinath, G., Rigobon, R., 2008. Sticky borders. *Q. J. Econ.* 123 (2), 531–575.
- Gorodnichenko, Y., Talavera, O., 2017. Price setting in online markets: basic facts, international comparisons, and cross-border integration. *Am. Econ. Rev.* 107 (1), 249–282.
- Gorodnichenko, Y., Sheremirov, V., Talavera, O., 2014. Price setting in online markets: does IT click? *J. Eur. Econ. Assoc.* 16 (6), 1764–1811.
- Hauser, J.R., Wernerfelt, B., 1990. An evaluation cost model of consideration sets. *J. Consum. Res.* 16 (4), 393–408.
- Hitsch, G.J., Hortacsu, A., Lin, X., 2021. Prices and promotions in us retail markets: evidence from big data. *Quant. Mark. Econ.* 49, 289–368.
- Ilut, C., Valchev, R., Vincent, N., 2020. Paralyzed by fear: rigid and discrete pricing under demand uncertainty. *Econometrica* 88 (5), 1899–1938.
- Imbs, J., Mumtaz, H., Ravn, M.O., Rey, H., 2005. PPP strikes back: aggregation and the real exchange rate. *Q. J. Econ.* 120 (1), 1–43.
- Jacoby, J., 1984. Perspectives on information overload. *J. Consum. Res.* 10 (4), 432–435.
- Jung, J., Kim, J.-H., Matejka, F., Sims, C.A., et al., 2019. Discrete actions in information-constrained decision problems. *Rev. Econ. Stud.* 86 (6), 2643–2667.
- Kashyap, A.K., 1995. Sticky prices: new evidence from retail catalogs. *Q. J. Econ.* 110 (1), 245–274.
- Kehoe, P., Midrigan, V., 2015. Prices are sticky after all. *J. Monet. Econ.* 75, 35–53.
- Klenow, P.J., Kryvtsov, O., 2008. State-dependent or time-dependent pricing: does it matter for recent US inflation? *Q. J. Econ.* 123 (3), 863–904.
- Klenow, P.J., Malin, B.A., 2010. Microeconomic evidence on price-setting. *Handbook of Monetary Economics*, vol. 3. Elsevier, pp. 231–284.
- Knotek, E.S., 2008. Convenient prices, currency, and nominal rigidity: theory with evidence from newspaper prices. *J. Monet. Econ.* 55 (7), 1303–1316.
- Knotek, E.S., 2011. Convenient prices and price rigidity: cross-sectional evidence. *Rev. Econ. Stat.* 93 (3), 1076–1086.
- Lal, R., Matutes, C., 1994. Retail pricing and advertising strategies. *J. Bus.* 345–370.
- Levy, D., Bergen, M., Dutta, S., Venable, R., 1997. The magnitude of menu costs: direct evidence from large US supermarket chains. *Q. J. Econ.* 112 (3), 791–824.
- Levy, D., Lee, D., Chen, H., Kauffman, R.J., Bergen, M., 2011. Price points and price rigidity. *Rev. Econ. Stat.* 93 (4), 1417–1431.
- Lichtenstein, D.R., Ridgway, N.M., Netemeyer, R.G., 1993. Price perceptions and consumer shopping behavior: a field study. *J. Mark. Res.* 234–245.
- Maskin, E., Tirole, J., 1988. A theory of dynamic oligopoly, II: price competition, kinked demand curves, and Edgeworth cycles. *Econometrica* 571–599.
- Matějka, F., McKay, A., 2015. Rational inattention to discrete choices: a new foundation for the multinomial logit model. *Am. Econ. Rev.* 105 (1), 272–298.
- Monroe, K.B., 1973. Buyers' subjective perceptions of price. *J. Mark. Res.* 70–80.
- Monroe, K.B., Lee, A.Y., 1999. Remembering versus knowing: issues in buyers processing of price information. *J. Acad. Mark. Sci.* 27 (2), 207–225.
- Nakamura, E., 2008. Pass-through in retail and wholesale. *Am. Econ. Rev.* 98 (2), 430–437.
- Nakamura, E., Steinsson, J., 2008. Five facts about prices: a reevaluation of menu cost models. *Q. J. Econ.* 123 (4), 1415–1464.
- Nakamura, E., Steinsson, J., 2013. Price rigidity: microeconomic evidence and macroeconomic implications. *Ann. Rev. Econ.* 5 (1), 133–163.
- Orbach, B.Y., Einav, L., 2007. Uniform prices for differentiated goods: the case of the movie-theater industry. *Int. Rev. Law Econ.* 27 (2), 129–153.
- Piccione, M., Spiegler, R., 2012. Price competition under limited comparability. *Q. J. Econ.* 127 (1), 97–135.
- Rhodes, A., 2014. Multiproduct retailing. *Rev. Econ. Stud.* 82 (1), 360–390.
- Schindler, R.M., Kirby, P.N., 1997. Patterns of rightmost digits used in advertised prices: implications for nine-ending effects. *J. Consum. Res.* 24 (2), 192–201.
- Shin, J., 2005. The role of selling costs in signaling price image. *J. Mark. Res.* 42 (3), 302–312.
- Shugan, S.M., 1980. The cost of thinking. *J. Consum. Res.* 7 (2), 99–111.
- Simester, D., 1995. Signalling price image using advertised prices. *Mark. Sci.* 14 (2), 166–188.
- Snir, A., Levy, D., Chen, H.A., 2017. End of 9-endings, price recall, and price perceptions. *Econ. Lett.* 155, 157–163.
- Stevens, L., 2020. Coarse pricing policies. *Rev. Econ. Stud.* 87 (1), 420–453.
- Stiving, M., Winer, R.S., 1997. An empirical analysis of price endings with scanner data. *J. Consum. Res.* 24 (1), 57–67.
- Thomas, M., Morwitz, V., 2005. Penny wise and pound foolish: the left-digit effect in price cognition. *J. Consum. Res.* 32 (1), 54–64.
- Wall Street Journal, 2017. Fashion Ads, a Last Bastion of Print, Are Going Digital January 25, 2017.
- Zara, 2023. Help-Shop-Price Changes March 2023.
- Zbaracki, M.J., Ritson, M., Levy, D., Dutta, S., Bergen, M., 2004. Managerial and customer costs of price adjustment: direct evidence from industrial markets. *Rev. Econ. Stat.* 86 (2), 514–533.