NBER WORKING PAPER SERIES

QUANTUM PRICES

Diego Aparicio
Roberto Rigobon

Working Paper 26646
http://www.nber.org/papers/w26646

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2020, Revised January 2021

Data and codes will be made publicly available. We are grateful to Glenn Ellison for detailed discussions. We also thank Duncan Simester, Alberto Cavallo, Ignacio Puente, Mariana Albano, Pablo Azar, Manuel Bertolotto, Vivek Bhattacharya, Jean-Pierre Dube, Paul Ellickson, Sara Ellison, Miguel Espinosa, Ricardo Fraiman, Renee Gosline, Paul Joskow, Chris Knittel, Carolina Levinton, Juan Cristobal Lopez, Victor Martinez de Albeniz, John McCoy, Mateo Montenegro, Emi Nakamura, Chela Pedreschi, Pilar Posse, Drazen Prelec, Nancy Rose, 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. Nicolas Gomez del Campo and Gina Rey provided outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Quantum Prices
Diego Aparicio and Roberto Rigobon
NBER Working Paper No. 26646
January 2020, Revised January 2021
JEL No. D2,D22,D83,E31,L81,M3

ABSTRACT

Online data was collected for 350,000 products from over 65 fashion retailers in the U.S. and the U.K. Many retailers practice an extreme form of stickiness described as quantum prices: a large number of differentiated products are priced using few sparse prices, with price changes occurring rarely and in large magnitudes. Quantum prices exist within categories (similar products) and across product introductions (over time). Most surprisingly, it also occurs across categories (very different products). Normalized measures indicate substantial price clustering beyond the role of popular prices, assortment size, or digit endings. Quantum prices imply frictions in lumpy price adjustments through product mix, inflation measurement, and in the law-of-one-price.

Diego Aparicio
IESE Business School
Carrer d'Arnus i de Gari, 3-7
Spain
daparicio@iese.edu

Roberto Rigobon
MIT Sloan School of Management
100 Main Street, E62-516
Cambridge, MA 02142
and NBER
rigobon@mit.edu

1 Introduction

The Internet has had a fundamental impact on the way retailers engage with their consumers. First, consumers often observe prices across locations, making geographical price discrimination harder. Old brick-and-mortar retailers were small local monopolies that had the ability to decide the prices and products' offerings. This practice, in turn, implied that uniform prices were rare. Second, and equally important, the internet allowed retailers to organize the data differently. Since the Egyptian land register created around 1400 BC —the first known land ownership record— the data has been organized around two paradigms: geography and socio-economic conditions. This structure is informative because a family's location or economic conditions are good predictors of their customs and purchases. In other words, it is an indicator of preferences. With the introduction of the internet, however, it is now possible to directly observe consumers' behaviors. Every purchase, trip taken, website visited, music listened, movies watched, and searches googled, are breadcrumbs left behind that describe aspects of the personality and preferences. Therefore, clustering the data along the actual behavior constitutes a better estimate of preferences. These two features have changed how retailers engage with customers.

We study the fashion industry and observe three distinct pricing behaviors, which we denote as follows: Traditional, Platform, and Quantum Pricing strategies. Traditional pricing stores —such as Louis Vuitton, Dolce Gabana, Ralph Lauren, Reiss— focus on obfuscating prices from the consumer. These retailers will often make 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, Google— emphasize their business in growing the network and thereby expanding its market share. They tend to have two sources of revenues, namely a mark-up and the revenue from utilizing the consumer information. Their 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; only a few can capture the whole network. These two strategies have received considerable attention in the literature.

Our paper is focused on characterizing a new form of pricing: the *Quantum Pricing* strategy. The quantum pricing is followed by a vast number of firms, including some of the leading and most successful retailers —such as Apple, Zara, H&M, Ikea, Bonobos, and Uniqlo. These retailers tend to cluster prices massively along price points that are far from each other and that are sticky over time. In fact, product introductions tend to occur at those price points, and even in some occasions discounts are set to match those existing lower prices. 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]

[Figure 1]

Figure 1 provides a visual perspective with several striking observations. It shows Uniqlo's distribution of distinct prices and the share of products by category. The distribution is extremely clustered and products from very different categories often have the same identical price. Uniqlo exhibits many remarkable numbers: it is a $21 billion business in sales ($1 billion in the U.S.), has 2,200 stores throughout 22 countries (50 stores in the U.S.), and it is among the 100 most valued brands worldwide. But perhaps most surprising is the price number: its online catalog of 1,500 products has just 13 prices. Additionally, there is not a single different 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.

Taking a step back from Uniqlo, the stylized fact is remarkable when connected to two separate but recent areas in price setting. First, recent studies document the practice of uniform pricing across locations within a chain (DellaVigna and Gentzkow (2019); Hitsch, Hortacsu, and Lin (2019)). Relatedly, there is evidence of price points in the price distribution of supermarket scanner data (Levy, Lee, Chen, Kauffman, and Bergen (2011); Eichenbaum, Jaimovich, and Rebelo (2011); Anderson, Jaimovich, and Simester (2015); Ilut, Valchev, and Vincent (2020); Stevens (2020)). In that sense, quantum prices extend and exacerbate the discrete or uniform pricing: retailers use the same handful of prices not just across locations, across online vs. offline, or across similar products, but also for very different types of products and even across seasons. And second, this behavior stands at odds with the possibility of utilizing advanced price algorithms (Brynjolfsson and McAfee (2014); Calvano, Calzolari, Denicolo, and Pastorello (2020)). Said differently, some of these retailers

[1]The term 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 very sticky, and move in discrete and large quantities.

have state-of-the-art assortment and markdown optimization strategies, e.g. fast-fashion assortment optimization *during* the season (Caro and Gallien (2007); Cachon and Swinney (2011)). These firms could afford to use price algorithms to set prices. Instead, they choose to have a sparse and discrete price grid.

We study pricing strategies in the fashion retail industry in the U.S. and the U.K.[2] We collect online data from 54 retailers in the U.S. and 40 retailers in the U.K, with a total of close to 230,000 and 190,000 distinct products, respectively. This represents a significant share of the apparel, footwear, and accessories industries in both countries, and to the authors' knowledge 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 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 documenting the existence of quantum prices in the fashion industry within retailer-categories (excluding 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 later extend the analysis to across-categories within a retailer and, surprisingly, we find substantial clustering in prices. The average probability that two items from different categories have the same price is about 5%. Furthermore, the tails in the distribution of these probabilities include extreme clustering behaviors. These estimates already provide a model-free motivation for the practice of using few prices for large and diverse assortments.

A number of 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

[2]Fashion retail represents a significant portion of the economy. It accounts for about 4% of the CPI basket (food and beverages, 16%). The combined market of apparel, footwear, bags, and accessories is about $380 billion sales in the U.S., and £55 billion sales in the U.K. The U.S. is the second largest market, next to China, and accounts for about 20% of world sales in the sector. The internet channel accounts for about 17% of sales, but is growing rapidly (20% CAGR).

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, and this fraction increases to 45% 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 variation and between-cluster variation. We find retailers that concentrate prices in a few discrete clusters, and that price clusters have economically meaningful differences between each other. The median price differential between centroids (mid-points) is 30% and 21% in the U.S. and the U.K., respectively.

Moreover, quantum prices are documented in the time-series. Using data collected across collections, we show that the quantum prices are sticky and, in fact, retailers will consistently introduce new products at those quanta. The existence of quantum prices implies that price changes at introduction also exhibit a degenerate distribution: either zero or very large. This suggests a new form of menu cost stickiness, i.e. a cost to opening a new price bucket in the distribution. In fact, although quantum retailers use the same set of prices for new and old 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 distinct prices are the same but the proportion of products in each bucket changes. This type of stickiness is different from the canonical menu cost models (e.g., Golosov and Lucas Jr (2007)), because fashion products tend to have a short duration, often do not experience any price increase, and the main pricing decision takes place at the time of a product introduction. Moreover, it stands in contrast to models that assume firms optimize prices from a relatively unrestricted domain of positive numbers.

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 fashion retail industry in terms of revenue market shares, as well as in terms of the large firm heterogeneity. The data includes department stores, luxury retailers, medium- to low- end pricing retailers, and fast fashion. For example, the data includes Louis Vuitton, Forever 21, GAP, Walmart, Victoria's Secret, and Zara.[3] Moreover, the quantum prices are economically consequential: all the firms sell at the same price across physical stores and the online store, and sales at full price often account for the largest share of revenue in the life cycle of a product.

We conclude the paper discussing the relevance of quantum prices beyond the core empirical findings. We focus on illustrating two implications. First, we include a simulation approach to quantify the extent to which quantum pricing can affect the estimation of the inflation rate. Firms face two frictions that prevent them to flexibly change prices: a sparse price choice set and a standard *à la* Calvo time-dependent rule. Our findings indicate that statistical offices will need to sample a significantly large number of products to capture the average behavior in prices, and moreover the inflation estimates are noisy and dispersed. Intuitively, the price change distribution degenerates into atoms when quantum price strategy is used. Finally, we report large deviations (quantum retailers relative to non-quantum retailers) from the law of one price using matched products collected the same day between the U.S. online store and the U.K. online store. The intuition is, once again, that a sparse price grid adds a friction to price at the LOP.

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 between the exact managerial decision making process. Instead, we build upon the stylized facts and institutional knowledge of price setters to set forward a simple behavioral model in which a restricted price distribution makes price advertising increasingly effective. The novelty lies in modeling the role of price advertising and price recall in the context of large, diverse, and fast-moving assortments that need to be communicated to consumers. We also complement the conceptual discussion with theories of price salience, demand uncertainty, and convenient prices.

[3]The size of some of these retailers is remarkable. For instance, Louis Vuitton had £22 billion sales in fashion and leather goods; Uniqlo had $21 billion sales; Inditex had €28 billion sales and operates over 7,400 stores in over 80 countries; Nike had $37 billion sales. Estimates as per FY2019 annual reports.

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, Neiman, and Rigobon (2014), Gorodnichenko, Sheremirov, and Talavera (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.

The machine learning classifier to categorize products relates to previous efforts to categorize unstructured online data. For example, Cavallo (2012), Cavallo and Rigobon (2016), and Aparicio and Cavallo (2019) categorize online products in order to construct price indices using CPI weights. Our work build upon this literature in that we have a more comprehensive data in fashion retail, and in that we describe a data-based technique which 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)), adverse quality signaling for lower-priced goods (Anderson and Simester (2001)), and the role of variants (Anderson, Jaimovich, and Simester (2015)). See also Nakamura (2008); Levy, Lee, Chen, Kauffman, and Bergen (2011); Eichenbaum, Jaimovich, and Rebelo (2011); Ilut, Valchev, and Vincent (2020); Stevens (2020) using scanner data. Our research provides a rather extreme version of uniform pricing, extending across very different types of products and across seasons, suggesting additional micro- and macro- implications.

Our main framework of price advertising, mainly discussed in the Appendix for space limitations, 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 assortments and fast-turnover products, and where the need (and the strategies) to inform consumers are different from traditional retail. Relatedly, the work of Jung, Kim, Matejka, Sims, et al. (2019); Caplin and Dean (2015); Matějka and McKay (2015) highlights the importance of accounting for limited processing capacity.The model helps conceptualize the effectiveness of the distribution of unique prices and a large menu cost to introducing new prices.

We enhance the conceptual discussion borrowing ideas from demand uncertainty (Ilut, Valchev, and Vincent (2020)), bounded comparability (Hauser and Wernerfelt (1990); Eliaz and Spiegler (2011); Piccione and Spiegler (2012); Bordalo, Gennaioli, and Shleifer (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 surprising number of options in apparel. For example, Forever 21 sells about 2,000 different styles of dresses on a given day. 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 literature on price stickiness and on the managerial costs of updating prices. Some papers in this area include, for instance, Levy, Bergen, Dutta, and Venable (1997); Blinder, Canetti, Lebow, and Rudd (1998); Zbaracki, Ritson, Levy, Dutta, and Bergen (2004); Bhattarai and Schoenle (2014); Anderson, Jaimovich, and Simester (2015). We show that the existence of quantum prices can be thought of as an extreme form of menu cost which restricts the strategies of price adjustments and thereby hinders the computation of inflation.

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, Ridgway, and Netemeyer (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, Bergen, Dutta, and Venable (1997); Knotek (2008, 2011); Levy, Lee, Chen, Kauffman, and Bergen (2011); Ater and Gerlitz (2017); Snir, Levy, and Chen (2017)). Although we study 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 do account for a fraction of the observed quantum pricing. Still, after removing the effect of digit endings, there is substantial price clustering.

2 Online Data

We collect online data from 53 fashion retailers in the U.S. and 40 retailers in the U.K. These retailers are representative of the industry overall, i.e. we cover the largest firms in terms of market share and a wide heterogeneity in terms of style, item composition, branding, and customer base. See Appendix A.1 for the complete list of retailers.

Online data is collected as follows. A script is designed to search the Hyper-Text 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 very similar ID, which we use to keep only one of them. This allows to rule out price clustering that would arise from near perfect substitutes. Additional details about collecting online retail data can be found in Cavallo and Rigobon (2016); Aparicio and Cavallo (2019).

Due to the large scope of retailers and computational limitations, we collect data once per month during 6 months for most of the retailers, during 1 year for a subset of the retailers, and during 2 years for a few of the latter retailers. For the most parts of the analyses, we concentrate on a cross-section of items across retailers from a 6-month period. Note that a cross-section removes any price clustering explained by the same products over time.

The fashion retail sector 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, Neiman, and Rigobon (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 do experience either temporal or permanent discounts towards the end of the season. If discounts are permanent, items have a sale price until it is discontinued, $p^{t=1} > p_s^{t>1}$. And if discounts are temporal, the price will return to the regular price, $p^{t=1} = p^{t=3} > p_s^{t=2}$. In this paper we focus on the regular price, or introduction price, which we consider the most important pricing decision. Sales at the full price account for the largest share of revenue (Ghemawat, Nueno, and Dailey (2003)). Moreover, in the regulatory filings, firms often attribute lower financial returns to excessive markdowns.

[Table 1]

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 4,000 distinct products in each retailer, but there is a large heterogeneity. The 10th and 90th percentile store has 1,279 and 7,463 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 extraordinary 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, that groups items into twelve 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 not feasible to manually assign labels to every single product. Instead we design rules to check random portions of the data or products that exhibit dissimilar characteristics to those in their group (e.g., items that are too expensive in the category), re-train, and re-classify. The final output is a classifier which can consistently categorize products across retailers and across collections. See Appendix A.2 for additional details.

3 Evidence of Quantum Prices

This Section presents formal evidence of quantum prices or price clustering in fashion retail. We show results using a series of clustering measures computed at the retailer or retailer-category level.

3.1 Descriptive evidence

Table 1 indicates that, overall, there are about 200 products per price in the U.S. market. This suggests that many items in the same retailer-category will have the same price. Figure 2 shows the probability that two distinct items in the same retailer-category have the same price. The median probability is close to 10%. Some of these magnitudes are surprisingly large when we consider that some categories are relatively broad, and thus a lower bound to what we could find in more narrowly defined categories (e.g., Jeans instead of Bottoms). Item misclassification should, in any case, push the probabilities downwards. 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.

[Figure 2]

Appendix B.1 shows results of this probability computed at the retailer level. The average probability is 5.2% in the U.S. and 5.3% in the U.K. Although magnitudes are much smaller, there is a fair amount of retailers that use the same prices across categories. The heterogeneity in pricing across retailers is reinforced in Appendix B.2. We find substantial variation of the probabilities (that two items in the retailer-category have the same price) across the number of distinct products. For example, there are many retailer-categories where the probability is 20%, and the number of distinct product varies from less than 100 to over 1,000.

3.2 Normalized measure

Not all prices are equally good in practice. 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 having a lower price differential with respect to \$19.00 than with respect to \$20.99. 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 the number of products or from a range of good prices.

We construct a normalized clustering index that builds on Ellison and Glaeser

(1997).[4] 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 for clustering that occurs at certain prices (e.g., Zara Underwear) which are popular in the category (Underwear). Formally, the index in the retailer-category i is defined as follow:

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

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 number of distinct products in retailer-category i, and the term $\frac{1}{N_i}$ controls for the number of products.

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 B.3.

The normalized measure is substantially larger than those that would be expected if prices were drawn from a Normal or uniform distribution with the same empirical parameters. Prices out of the uniform distribution are restricted to 10 dollar multiples, which increases its clustering and thus provides a stringent benchmark. See Appendix B.4 for additional evidence. 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

[4] 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 (in our case, an industry is a retailer-category).

the data-based index.

Note 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. We discuss three alternatives.

First, we could 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 weighed retailers). Clustering estimates replicate the same patterns and therefore shown in Appendix B.5. The mean and median in the U.S. are 0.106 and 0.083, respectively.

Second, we could 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 run 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 regression in 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.

The results, once again, show that some fraction of firms, after controlling for these practices, have little to no clustering. But there is still a large number of retailer-categories that exhibit substantial price clustering. We view this as the preferred specification that captures the true price clustering in the data. The histogram is shown in Appendix B.5.

The more features we include in a regression like 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 the true 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 $\sum_{j=0}^{8} \beta_{i,\mathrm{E}_j} \mathrm{End}_{i,j}$. Estimates are shown in

Appendix B.5. We now find a larger number of retailer-categories with close to zero price clustering, suggesting that rightmost digits continues to be a popular pricing practice (Kashyap (1995); Stiving and Winer (1997); Anderson and Simester (2003)). However, half of the retailer-categories still exhibit price clustering that cannot be attributed to price endings.

[Table 2]

Table 2 provides summary statistics of the normalized price clustering indices. Note that the indices decrease from (1) to (3) and (4) as we sequentially control for additional 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. Extreme cases like these could be retailer-categories where there is close to one price per product. Column (3) shows that the mean and median of the preferred measure of price clustering is 0.044 and 0.063 in the U.S., and 0.041 and 0.057 in the U.K., respectively. And there is a considerable fraction of cases with clustering greater than 0.05, i.e. 41% in the U.S. and 40% in the U.K. Even in the stringent case that controls for price endings dummies, the mean estimates are 0.038 in the U.S. and 0.029 in the U.K.

The results so far demonstrates that there are sharp differences in pricing strategies. There is a group of retailers that exhibit little if any price clustering. In fact, these might even exhibit less clustering than what we would expect given the price distribution of the category. Many of these retailers fall into what we describe as traditional or platform retailers. Examples are Louis Vuitton and Walmart, respectively. Another set of retailers have price clustering explained by price levels, popular prices, or price endings, and thus controlling for these brings the excess clustering down. A third set of retailers continues to exhibit substantial clustering not driven by these features. These are the retailers which clearly fall into the quantum pricing strategy.

Although price endings accounts for a portion of price clustering practices, there remains a sizable number of retailer-categories which are substantially concentrated. A few reasons might help explain why. First, the data includes all the items in sets of retailers. Price endings can be more prevalent in advertised or popular items, but gathering entire catalogs presumably captures diverse digit endings. This is important because it is unlikely that there is a kink in the demand in Abercrombie at the $8 ending but at the $9 in Uniqlo, and that such kink would not exist at $9 and \$8, respectively. Second, price endings are less commonly thought at the integer level. This would exacerbate price rigidity (from \$29 to \$39 there a 35% difference). In fact, about 51% of the over 230,000 products in the U.S. have a non-zero decimal digit. Lastly, because markdowns are prevalent in apparel, if price endings were important to revenue then one might expect a similar strategy in markdown prices. Still, many fashion retailers set sale prices as percentages or direct price points; and other than a higher frequency of \$9 in markdowns, we find no strong relationship between price endings in regular prices and markdown prices. See Appendix B.6.

3.3 Price clustering across categories and retailers

The clustering measures can also be computed within retailers and across categories, as well as across retailers and within categories. For example, comparing prices between two categories within the same retailer would estimate the degree of price grid overlap. We discuss three results.

First, we compute the probability that two random items, in different categories but in the same retailer, have the same price. This empirical probability is computed sampling two items for every within-retailer category-category pair. In total there are 2,217 pairs in the U.S. and 1,646 in the U.K. 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). For example, in the U.S. there are over 500 retailer category-category pairs (about 25% of total) with more than 5% estimated probability. In the U.S., the mean is 3%, the median is 2%, and the 90th percentile is 10%; in the U.K., the mean is 5%, the median is 2%, and the 90th percentile is 12%. The histogram is shown in Figure 3.

[Figure 3]

Second, we implement the normalized measure in eq. (1) within-retailer and across-categories, as well as across-retailers and within-retailers. Estimates are expected to be lower because items in two different categories cannot in general be more concentrated than they are in their own category. The normalized measure of correlated clustering is based on a modification of Ellison and Glaeser (1997); Ellison, Glaeser, and Kerr (2010), who estimate geographic concentration across manufacturing plants from different industrial sectors. The normalized index is computed as follows:

$$index^c_{i,j} = \frac{\sum_{b=\underline{b}}^{\bar{b}}(s_{i,b} - x_b)(s_{j,b} - x_b)}{1 - \sum_{b=\underline{b}}^{\bar{b}} x_b^2} \tag{3}$$

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 remarkable 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 B.7 shows the histogram of the clustering measure as well as which category pairs tend to be on average more concentrated. The average is 0.018 and the median is 0.016, computed across 74 category pairs in the U.S. (excluding same category pairs). Importantly, this measure will only be large when different 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, in order to examine if retailers use the same prices for similar items. Appendix B.8 indicates no evidence of correlated clustering across retailers. For computational reasons the results are available for a random fraction of the retailers in each country. There are in total 832 and 1,475 within-category retailer-retailer pairs in the U.S. and the U.K., respectively. The estimates are close to 0 en both countries. For example, the mean and median in the U.S. is 0 and 0.002, respectively. The lack of evidence echoes the limited advantage of price endings. If some prices were particularly appealing to consumers then one might expect different retailers concentrating on the same prices.[5]

In summary, 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 B.9 overlaps the clustering measure distributions at the three levels. Table 3 provides selected examples of retailers and their estimated clustering measures.

[5]This evidence does not imply that retailers are not mindful of each others' prices. For instance, retailers can own niche and clustered price buckets to soften price competition. Analyzing these dynamics is beyond the scope of the paper.

[Table 3]

3.4 Robustness: machine learning clustering

In previous analyses we showed evidence of price clustering using predefined buckets of prices. But are these price 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, Hastie, and Tibshirani (2001)) and from the CH index (Caliński and Harabasz (1974)). See Appendix B.10 for methodological details. We define a ratio $\kappa(k) \equiv \frac{\overline{WC}(n_k,k)}{\overline{BC}(k)}$, 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. First, it determines the optimal number of price clusters in the data according to a standard trade-off. And second, we demand that clusters be separated at least 5% from each other.

[Figure 4]

Figure 4 shows the results of k_i^* for all retailer-categories i in the U.S., the corresponding average distance between consecutive centroids (in percentage), and the ratio of k^* to the maximum possible of clusters.[6] Results for the U.K. are shown in Appendix B.10. Overall, the findings are qualitatively consistent with those discussed previously. Moreover, it is reassuring to the stringent normalized measure, defined in eq. (1) and eq. (2), which controls for prices that are too close to each other.

There is a large share of retailer-categories that exhibit medium to substantial price clustering. These are captured by those having a low ratio of k^* relative to the maximum possible k. Then there is another set of cases that are poorly clustered, and tend to exhibit large k^*. And 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 30% in the U.S. and 21% in the U.K., respectively.

[6]Due to the 5% threshold, the maximum number of clusters is not the number of distinct prices, i.e. $\#b$. The maximum k^* is $\lfloor \frac{\log(b^{max}/b^{min})}{\log(1.05)} \rfloor + 1$.

4 Stickiness Dynamics of Quantum Pricing

Section 3 showed evidence of quantum pricing in the cross-section. However, the relevance of the price clusters fades if retailers select new quantum prices in subsequent periods. We now show that quantum prices are very sticky over time. Retailers are not only reluctant to create new prices throughout the life of a product, but also use the same prices for product introductions.

4.1 Product introductions

Fashion products are often characterized by little, if any, upward price changes and a short product life. This raises the question of which prices are retailers choosing throughout collections.[7] In addition, there is substantial seasonality in prices, as measured by the monthly inflation rate in clothing. The average absolute non-seasonally adjusted monthly inflation rate is about 1.9% in the U.S. and 1.8% in the U.K. (see Appendix C.2). Therefore, a priori there is no reason why we should not expect very different, seasonal prices as the assortment composition varies over time. However, we find that the set of prices is remarkably stable.

We take the prices in each retailer-category's initial data collection month, and compare these with the prices observed 1 to 5 months later. This measures the likelihood that retailers use different prices over time in a given category. The top panel in Table 4 shows a significantly large share of common prices, relative to the first month, even though the catalog is changing. The estimates can be interpreted as follows. The share of common prices measures the ratio of common prices (between the first month and m months after) to the distinct prices in month m. For example, in half of the retailer-categories over 75% of the prices observed 5 months later were exactly the same to those in the first month.

[Table 4]

These set of results include both existing products (whose price may change) as well as new products. The bottom panel in Table 4 replicates the analysis but only for product introductions in each of the 1 to 5 following months. Once again,

[7] We document that the mean product life is 3.1 months in the U.S. and 3.2 months in the U.K. In addition, we document asymmetric frequencies of price changes. Only 5.3% of the products have a regular price change, whereas close to 70% experience a sale price. Estimates are comparable to Nakamura and Steinsson (2008) and Cavallo, Neiman, and Rigobon (2014). See Appendix C.1.

the vast majority of products are introduced at the existing prices. For example, the average probability that a new product in month 5 comes in at an existing price is 0.91. Moreover, the 90th percentile probability is 1, which means that there is a fair share of retailer-categories pairs where there is no single product introduction with a new price. Similar results are found in the U.K. (Appendix C.3).

We now explore whether the price stickiness probabilities (line (ii) in Table 4) are related to the clustering measures. We run fixed-effects models at the retailer-category level, where the dependent variable is defined as the average of these m-based probabilities. Regressions are run jointly for all group pairs in the U.S. and the U.K. We find that the clustering index associates with higher price stickiness, as defined by the probabilities that products will be introduced at the existing prices. We also find that measures of price salience and fast-inventory turnover (fast-fashion) are related to a higher price stickiness.[8] Appendix C.4 shows the the complete set of results. The motivation for leveraging price advertising lies in that, for those retailers, price salience and the price distribution is a key strategy of their business (Section 6).

4.2 Adjustments through product mix

Inflation in the U.S. and the U.K. has been hovering around 2%, and input prices often experience large price swings.[9] Although we lack unit costs in the data, it is reasonable to assume that cost shocks are expected, to some degree, to impact final prices. One wonders how retailers might adjust their average prices with a discrete price grid.

We begin using evidence from two stylized retailers: Uniqlo and Ralph Lauren. Panel (a) in Figure 5 compares the prices observed in Uniqlo U.S. between the same categories over two years (same month). 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. In fact, quantum prices are so sticky that

[8]The price salience indicator was constructed as follows. We enrolled in e-mail newsletters from most of the retailers, and recorded whether these included price points or not (promotions in terms of percentages that did not include price points were classified as non price advertisers). We also complemented the newsletters with a simple online exercise: for each retailer we checked whether the landing page included price points or not. These two sources are surprisingly similar. Approximately 36% of the retailers in the data engage in price advertising. Fast-fashion retailers are identified following Caro and Martínez-de Albéniz (2015) and industry reports.

[9]For example, cotton prices increased 20% year-on-year as of April 2018. But then prices decreased 26% year-on-year as of July 2019. See Financial Times (2018); Bloomberg (2018, 2019).

over 90% of the change in the price distribution occurs via modifying the product shares in the old prices. On the other hand, in Panel (b) showing Ralph Lauren U.S., prices are spread out across many points in a price grid. In fact, the price range showed is restricted for comparison purposes; the entire price distribution is significantly wider (Panel (c)). In this example, changes in the price distribution are more evenly split between shares in the same prices and new price buckets (close to 50% each).

[Figure 5]

The observation about the changes in the price distribution can be generalized. We use data on all retailers for which we have one year of data and thus 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 change in the price distribution is decomposed into:

$$\sum_{i=p^{\min}}^{p^{\max}} |w_{i,2} - w_{i,1}| = \underbrace{\sum |w_{i,2} - w_{i,1}|_{w_1 \cap w_2 \neq \emptyset}}_{\text{shares of products in same prices}} + \underbrace{\sum |w_{i,2} - w_{i,1}|_{w_1 \cap w_2 = \emptyset}}_{\text{shares of products in different prices}} \quad (4)$$

The term in the left computes the sum of the absolute differences between the shares of products in the price buckets with observations 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 measure is computed for all retailer-categories.

[Figure 6]

The median proportion of the shares in the same price bins is 0.6, and tends to be larger for price clustered and price advertiser retailers. Panel (a) in Figure 6 shows that the fraction of the change in prices through existing prices is significantly higher for retailers which advertise price salience. The mean share is 0.45 for non price advertiser retailers and 0.80 for price advertisers. Panel (b) shows the histogram of the proportion accounted for the left term in eq. (4).

We propose a second measure of support in the distribution across periods. Similarly, $w_{i,1}$ and $w_{i,2}$ denote the shares of products located in price i (no rounding) in time 1 and time 2, respectively. The minimum support in the price distribution can be computed as:

$$\sum_{i=p^{\min}}^{p^{\max}} \min_{w} \left(w_{i,1} - w_{i,2}\right) \tag{5}$$

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 $w_{12.4,1} = 0.15$ and $w_{12.4,2} = 0$, then we obtain 0. We then sum these minimum shares across all possible prices. The results are similar to Figure 6. The average support for non price advertisers is 44% and for price advertisers it is 78%.

Pricing through product mix can be a source of lumpy price adjustment, i.e. either very small or very large. Appendix C.5 shows larger lumpy price adjustments for price advertising retailers. Lumpy adjustments predict either 0 or economically large changes, which is what the Figure shows. For example, over 50% of the retailer-category pairs with price salience have exactly 0 change in the median price (compared to less than 20% for non-price advertising retailers). The average absolute change, conditional on being different from zero, is 21% and 13% for price and non-price advertisers, respectively.

5 Implications of Quantum Prices

Sections 3 and 4 showed that quantum strategy retailers use a handful of prices for large and diverse assortments, and that such strategy is sticky over time. We now illustrate more general implications of quantum prices.[10] More precisely, we focus on two analyses: inflation measurement and deviations from the law of one price.

5.1 Inflation measurement

We study the inflation measurement implications while being mindful of two limitations, namely assuming the data is representative of the aggregate economy and

[10]Quantum prices can be interpreted as an extreme form of price stickiness. The stickiness is different from canonical time-dependent or state-dependent models. In time-dependent models, the timing of price changes is exogenous, and in state-dependent models, firms choose when and by how much to change prices due to menu costs (Klenow and Kryvtsov (2008)). In this paper, quantum price stickiness is interpreted as closer to state-dependent pricing, because firms decide how to change prices given quantum prices, or how to introduce a new quantum price. But still it suggests a new form of menu cost, namely the cost of introducing a new price.

assuming uniform units sold. The simulation conceptualizes that quantum prices have practical implications to measuring inflation and sampling products. The intuition behind is that quantum prices can be thought of as an extreme form of price rigidity –a large menu cost–, which implies that the distribution of price changes is degenerate. As a consequence, a large sample of distinct products is needed to capture the average price behavior.

We consider the pricing decision of the firm in a partial equilibrium model with a deterministic shock in the spirit of standard inventory models (Nevo (2011)). The complete model setup is presented in Appendix D.1. Prices are not totally flexible due two frictions: *a*) a limited price choice set, and *b*) a limited time-period ability to change prices through product introductions and Calvo pricing. There are two possible reasons for prices to change. First, we allow for *à la* Calvo time-dependent process, which we calibrate to be once a year. And second, we allow for a product to be replaced by a new variety, at which time the price is totally flexible; calibrated to a rate of once every 6 months. The importance of accounting for price changes due to product introductions is reminiscent of Nakamura and Steinsson (2012). These two events are modeled as a Poisson process with expected realizations of 12 and 6 months, respectively. The firms' decision rule is a non-stationary problem between two possible prices, and we solve by backward induction knowing that all firms will start and end at the optimal price.

With these assumptions in hand, we created a simulation of 10,000 items, and solved backwards starting from the steady state. We compute three types of price indexes: a CES index (eq. (12), Appendix D.1), a price index that replicates the CPI (Laspeyres's), and a price index that replicates the consumer expenditure procedure from the Bureau of Economic Analysis (PCE). Formally:

$$\mathbf{P}_{ces} = (\sum p_i^{1-\sigma})^{\frac{1}{1-\sigma}} \tag{6}$$

$$\mathbf{P}_{cpi} = \frac{1}{N} \sum p_i \tag{7}$$

$$\mathbf{P}_{pce} = \sum \omega_i p_i \tag{8}$$

where $\omega_i = q_i / \sum q_i$.

The results are shown in Figure 7. Several remarks are noteworthy. First of all, looking at the left panels, prices do not change at the aggregate level for several periods. That is to say, even though there is a large fraction of firms whose prices are allowed to change –either because of new products or time dependency– firms decide not to do so. The existence of quantum prices resembles a large menu cost. Moreover, the three price indexes produce different inflation paths.

[Figure 7]

It is interesting to observe the dispersion in inflation measurement as a result of taking a sub-sample of the price distribution. We analyze three different sample sizes: 0.1%, 1% and 10% of the items. For example, in the case of 1% we randomly sample 1% of the items and follow them over time. When the item is substituted with a new product, we keep track of the substituted item. We assume no quality change in new products, hence the price change that takes at the time of the substitution is attributed to cost inflation. We create 500 possible sub-samples of 1% items and compute the CES price index.

The left Panels depict the slowest and the fastest inflation paths for each sample size, and the right Panels depict the standard deviation of the price indexes. We begin noting the very large price dispersion when 0.1% of the items are sampled (Panels (a) and (b)). The standard deviation is 2 times larger than the error computed by changing the inflation rate methodology. When 1% of the items are sampled (Panels (c) and (d)), the indexes become closer, and the standard deviation is about the same order of magnitude than changing the methodology from the CPI to the PCE. Finally, when the sub-sample of products is increased to 10% (Panels (e) and (f)) the error in the price indexes becomes negligible.

The parameters in the simulation have been calibrated to match the annual inflation rate in fashion retail (close to 3 percent), with a similar pace of product introductions and price changes, and where a single store like Zara sells about 10,000 items. Additionally, the price clusters in our data are on average separated by close to 30 percent. Overall, these results raise concrete implications for national statistical offices. In order to achieve little sampling error, it will need to survey 10% of the items in each store. For instance, taking 50 fast fashion retailers whose prices are considered clustered, it implies surveying 50,000 items just in a single category; but even if 1% of the items are surveyed, or even if not all stores carry a large assortment as Zara, the human effort required is remarkable. This compares to the Bureau of Labor Statistics's 88,000 items surveyed every two months. In summary, the existence of price clustering and regular price stickiness, as quantified in the simulation analysis, indicates a meaningful bias to measure inflation. A simpler procedure could be to survey average prices weighted by units sold in each category of the store. Although this resembles a PCE-type index at the category level, the bias is small and the collection effort is minimal.

5.2 Law of one price

Additionally, we consider implications in terms of the law of one price (LOP) returning to the case study of the two stylized retailers Uniqlo and Ralph Lauren. Fashion retailers can potentially reduce the pricing frictions from a limited price choice set by designing items to hit specific prices. For instance, compared to retailers which sell stable product lines (e.g, food, electronics, appliances), fashion retailers are known for having short time-to-market and for rotating the assortment within the season (Caro and Gallien (2007)). However, multi-national retailers cannot perfectly produce for a specific quantum price. First, quantum prices are different across countries, and within a country prices are separated by non-trivial price increments. Second, there are exchange rate movements. And third, there are country-level taste differences. Therefore, despite lower price setting costs for new products, quantum prices are expected to generate good-level deviations from the law of one price (LOP) at the events of product introductions.

The data allows to test for product-level LOP because each product has a unique ID which can be utilized to perfectly match the same product in the U.S. online store and the U.K. online store on the same day. The patterns in LOP can be visualized in Figure 8. In particular, we compute the percent of products assigned to each combination of U.S. dollars and U.K. prices. Darker regions indicate a larger share of products assigned to a given bucket. We pool all products throughout the collection period. The heatmaps can be related to the price distributions in Figure 5. Panel (a), which corresponds to Uniqlo, depicts large and discrete price increments between prices. In fact, a handful of buckets are enough to 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.

[Figure 8]

These intuitions can be formally quantified by estimating a fixed-effects model of relative prices and evaluating the residuals. In particular, we regress $Relative_{i,t} = \alpha + \gamma_t + \epsilon_{i,t}$, where $Relative_{i,t}$ denotes p_{UK}/p_{US} (the product-level ratio of U.K.

and U.S. price). Monthly fixed effects are included to control for exchange rate movements. Confirming our visual impressions in Figures 5 and 8, quantum pricing relates to large deviations from the law of one price. The MSE is 67.8% larger in Uniqlo's, compared to Ralph Lauren's case.

Overall, using data from all retailers, we find large good-level deviations from the LOP. We compute the good-level real exchange rate (RER) for each good i, which is defined as $RER_i = \log(p_{i,UK}) - e_{US,UK} - \log(p_{i,US})$; where $e_{US,UK}$ denotes the log of the value of the (average monthly) nominal exchange rate between the US and the UK. Values close to 0 would indicate no deviation to the LOP. The mean and median absolute good-level RER is 0.201 and and 0.196 log points, respectively.[11] We also find that products introduced in the first price bucket in each retailer-category in the U.S. are in general not priced in similar price buckets in the U.K. Additional results are shown in Appendix D.2.

6 Theories of Quantum Prices

We lack the data to empirically test the underlying managerial process. In Appendix E we discuss a variety of models to conceptualize the stylized facts of quantum pricing. We set forward a simple behavioral framework, in which the novelty is to explicitly consider the economies of price advertising and price recall. In particular, fewer prices makes price advertising more effective; and it is increasingly more effective when the same prices are used not only within a category, but also across categories and across seasons. In other words, there is a menu cost to introducing a new price (to the price distribution) for product introductions. Note that, once again, fashion retail is characterized by fast-moving assortments and therefore there is a constant need to inform products and prices to consumers. Several recent papers (Jung, Kim, Matejka, Sims, et al. (2019); Caplin and Dean (2015); Matějka and McKay (2015)) study how optimal discrete actions can arise when information processing is constrained (Shannon's channel capacity).

Additionally, we discuss models featuring convenient prices that generate demand kinks at certain prices (Kashyap (1995); Knotek (2008); Levy, Lee, Chen, Kauffman, and Bergen (2011); Knotek (2011)); price salience and bounded comparability

[11] These are considered relatively large deviations from the law of one price. See recent studies using micro data, e.g. Imbs, Mumtaz, Ravn, and Rey (2005); Gopinath and Rigobon (2008); Cavallo, Neiman, and Rigobon (2014); Gorodnichenko and Talavera (2017).

of attributes (Eliaz and Spiegler (2011); Piccione and Spiegler (2012); Bordalo, Gennaioli, and Shleifer (2013, 2015)); demand uncertainty that makes firms more likely to use prices that were successful in the past (Ilut, Valchev, and Vincent (2020)); and managerial inattention (Fershtman and Kalai (1993); Bloom and Van Reenen (2007); Ellison, Snyder, and Zhang (2018)). These models capture some of the stylized facts and allow for an informed discussion of discrete pricing.

7 Conclusions

This paper used 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 collection combines three pieces that are rarely available together: (i) a large scale cross-section of products that are representative of the diversity of firms in the industry; (ii) a time-series of collections; and (iii) thousands of matched products across countries.

We show evidence that a large fraction of retailers practice what we define as quantum pricing: they set prices that are sparse and separated by large increments. Moreover, and perhaps more surprising, the same handful of prices are used across different categories of products and over time. Quantum pricing represents a remarkable source of price stickiness, suggesting a menu cost to introducing new prices and lumpy price adjustments, which has implications for inflation measurement and international pricing frictions. As the Internet continues to account for a growing share of the economy, it remains an open question whether both the platform and the quantum pricing strategies will intensify.

References

Anderson, E., N. Jaimovich, and D. Simester (2015): "Price stickiness: Empirical evidence of the menu cost channel," *Review of Economics and Statistics*, 97(4), 813–826.

Anderson, E. T., and D. I. Simester (2001): "Price discrimination as an adverse signal: Why an offer to spread payments may hurt demand," *Marketing Science*, 20(3), 315–327.

——— (2003): "Effects of 9 price endings on retail sales: Evidence from field experiments," *Quantitative marketing and Economics*, 1(1), 93–110.

——— (2008): "Research note—does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes," *Marketing Science*, 27(3), 492–500.

Aparicio, D., and A. Cavallo (2019): "Targeted price controls on supermarket products," *The Review of Economics and Statistics*, forthcoming.

Ater, I., and O. Gerlitz (2017): "Round prices and price rigidity: Evidence from outlawing odd prices," *Journal of Economic Behavior & Organization*, 144, 188–203.

Bhattarai, S., and R. Schoenle (2014): "Multiproduct firms and price-setting: Theory and evidence from US producer prices," *Journal of Monetary Economics*, 66, 178–192.

Bils, M., and P. J. Klenow (2004): "Some evidence on the importance of sticky prices," *Journal of political economy*, 112(5), 947–985.

Blinder, A., E. R. Canetti, D. E. Lebow, and J. B. Rudd (1998): *Asking about prices: a new approach to understanding price stickiness*. Russell Sage Foundation.

Bloom, N., and J. Van Reenen (2007): "Measuring and explaining management practices across firms and countries," *The Quarterly Journal of Economics*, 122(4), 1351–1408.

Bloomberg (2018): "Apparel Retailers, Finally Recovering, Eye Cotton Prices Warily," June 1, 2018.

——— (2019): "Cotton Crash Looks Set to Deepen as Stockpiles Keep Growing," July 9, 2019.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013): "Salience and consumer choice," *Journal of Political Economy*, 121(5), 803–843.

——— (2015): "Competition for attention," *The Review of Economic Studies*, 83(2), 481–513.

——— (2020): "Memory, attention, and choice," *The Quarterly Journal of Economics*, forthcoming.

Brown, F. E. (1969): "Price image versus price reality," *Journal of Marketing Research*, pp. 185–191.

Brynjolfsson, E., and A. McAfee (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Cachon, G. P., and R. Swinney (2011): "The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior," *Management science*, 57(4), 778–795.

Caliński, T., and J. Harabasz (1974): "A dendrite method for cluster analysis," *Communications in Statistics-theory and Methods*, 3(1), 1–27.

Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2020): "Artificial intelligence, algorithmic pricing, and collusion," *American Economic Review*, 110(10), 3267–97.

Caplin, A., and M. Dean (2015): "Revealed preference, rational inattention, and costly information acquisition," *American Economic Review*, 105(7), 2183–2203.

Caro, F., and J. Gallien (2007): "Dynamic assortment with demand learning for seasonal consumer goods," *Management Science*, 53(2), 276–292.

——— (2010): "Inventory management of a fast-fashion retail network," *Operations Research*, 58(2), 257–273.

Caro, F., and V. Martínez-de Albéniz (2015): "Fast fashion: business model overview and research opportunities," in *Retail supply chain management*, pp. 237–264. Springer.

Cavallo, A. (2012): "Online and official price indexes: measuring Argentina's inflation," *Journal of Monetary Economics*, 60(2), 152–165.

Cavallo, A. (2018): "Scraped data and sticky prices," *Review of Economics and Statistics*, 100(1).

Cavallo, A., B. Neiman, and R. Rigobon (2014): "Currency unions, product introductions, and the real exchange rate," *The Quarterly Journal of Economics*, 129(2), 529–595.

Cavallo, A., and R. Rigobon (2016): "The Billion Prices Project: Using online prices for measurement and research," *The Journal of Economic Perspectives*, 30(2), 151–178.

DellaVigna, S., and M. Gentzkow (2019): "Uniform pricing in us retail chains," *The Quarterly Journal of Economics*, 134(4), 2011–2084.

Devroye, L. (1981): "Laws of the iterated logarithm for order statistics of uniform spacings," *The Annals of Probability*, pp. 860–867.

Dickson, P. R., and A. G. Sawyer (1990): "The price knowledge and search of supermarket shoppers," *The Journal of Marketing*, pp. 42–53.

Draganska, M., and D. C. Jain (2006): "Consumer preferences and product-line pricing strategies: An empirical analysis," *Marketing science*, 25(2), 164–174.

Dubé, J.-P., G. J. Hitsch, and P. Manchanda (2005): "An empirical model of advertising dynamics," *Quantitative marketing and economics*, 3(2), 107–144.

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011): "Reference prices, costs, and nominal rigidities," *American Economic Review*, 101(1), 234–62.

Eliaz, K., and R. Spiegler (2011): "Consideration sets and competitive marketing," *The Review of Economic Studies*, 78(1), 235–262.

Ellison, G., and E. L. Glaeser (1997): "Geographic concentration in US manufacturing industries: a dartboard approach," *Journal of political economy*, 105(5), 889–927.

Ellison, G., E. L. Glaeser, and W. R. Kerr (2010): "What causes industry agglomeration? Evidence from coagglomeration patterns," *American Economic Review*, 100(3), 1195–1213.

Ellison, S. F., C. M. Snyder, and H. Zhang (2018): "Costs of Managerial Attention and Activity as a Source of Sticky Prices: Structural Estimates from an Online Market," *NBER WP24680*.

Fershtman, C., and E. Kalai (1993): "Complexity considerations and market behavior," *The RAND Journal of Economics*, pp. 224–235.

Financial Times (2018): "Cotton jumps as storm Alberto brings heavy rain to growing areas," May 29, 2018.

Forbes (2018): "15 Revelations from Trader Joe's podcast, from the parking lots to deer stew," June 7, 2018.

Friedman, J., T. Hastie, and R. Tibshirani (2001): *The elements of statistical learning*, vol. 1. Springer series in statistics New York.

Ghemawat, P., J. L. Nueno, and M. Dailey (2003): *ZARA: Fast fashion*, vol. 1. Harvard Business School Boston, MA.

Golosov, M., and R. E. Lucas Jr (2007): "Menu costs and Phillips curves," *Journal of Political Economy*, 115(2), 171–199.

Gopinath, G., and R. Rigobon (2008): "Sticky borders," *The Quarterly Journal of Economics*, 123(2), 531–575.

Gorodnichenko, Y., V. Sheremirov, and O. Talavera (2014): "Price setting in online markets: Does IT click?," *Journal of the European Economic Association*.

Gorodnichenko, Y., and O. Talavera (2017): "Price setting in online markets: Basic facts, international comparisons, and cross-border integration," *American Economic Review*, 107(1), 249–82.

Hauser, J. R., and B. Wernerfelt (1990): "An evaluation cost model of consideration sets," *Journal of consumer research*, 16(4), 393–408.

Hitsch, G. J., A. Hortacsu, and X. Lin (2019): "Prices and promotions in us retail markets: Evidence from big data," *NBER WP 26306*.

Ilut, C., R. Valchev, and N. Vincent (2020): "Paralyzed by fear: Rigid and discrete pricing under demand uncertainty," *Econometrica*, 88(5), 1899–1938.

Imbs, J., H. Mumtaz, M. O. Ravn, and H. Rey (2005): "PPP strikes back: Aggregation and the real exchange rate," *The Quarterly Journal of Economics*, 120(1), 1–43.

Jung, J., J.-H. Kim, F. Matejka, C. A. Sims, et al. (2019): "Discrete actions in information-constrained decision problems," *The Review of Economic Studies*, 86(6), 2643–2667.

Kashyap, A. K. (1995): "Sticky prices: New evidence from retail catalogs," *The Quarterly Journal of Economics*, 110(1), 245–274.

Kaul, A., and D. R. Wittink (1995): "Empirical generalizations about the impact of advertising on price sensitivity and price," *Marketing Science*, 14(3_supplement), G151–G160.

Kehoe, P., and V. Midrigan (2015): "Prices are sticky after all," *Journal of Monetary Economics*, 75, 35–53.

Klenow, P. J., and O. Kryvtsov (2008): "State-dependent or time-dependent pricing: Does it matter for recent US inflation?," *The Quarterly Journal of Economics*, 123(3), 863–904.

Knotek, E. S. (2008): "Convenient prices, currency, and nominal rigidity: Theory with evidence from newspaper prices," *Journal of Monetary Economics*, 55(7), 1303–1316.

——— (2011): "Convenient prices and price rigidity: cross-sectional evidence," *Review of Economics and Statistics*, 93(3), 1076–1086.

Lal, R., and C. Matutes (1994): "Retail pricing and advertising strategies," *Journal of Business*, pp. 345–370.

Levy, D., M. Bergen, S. Dutta, and R. Venable (1997): "The magnitude of menu costs: direct evidence from large US supermarket chains," *The Quarterly Journal of Economics*, 112(3), 791–824.

Levy, D., D. Lee, H. Chen, R. J. Kauffman, and M. Bergen (2011): "Price points and price rigidity," *Review of Economics and Statistics*, 93(4), 1417–1431.

Lichtenstein, D. R., N. M. Ridgway, and R. G. Netemeyer (1993): "Price perceptions and consumer shopping behavior: a field study," *Journal of marketing research*, pp. 234–245.

Matějka, F., and A. McKay (2015): "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 105(1), 272–98.

Monroe, K. B. (1973): "Buyers' subjective perceptions of price," *Journal of marketing research*, pp. 70–80.

Monroe, K. B., and A. Y. Lee (1999): "Remembering versus knowing: Issues in buyers processing of price information," *Journal of the Academy of Marketing Science*, 27(2), 207–225.

Nakamura, E. (2008): "Pass-through in retail and wholesale," *American Economic Review*, 98(2), 430–37.

Nakamura, E., and J. Steinsson (2008): "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 123(4), 1415–1464.

Nakamura, E., and J. Steinsson (2012): "Lost in transit: product replacement bias and pricing to market," *American Economic Review*, 102(7), 3277–3316.

Nevo, A. (2011): "Empirical models of consumer behavior," *Annu. Rev. Econ.*, 3(1), 51–75.

Orbach, B. Y., and L. Einav (2007): "Uniform prices for differentiated goods: The case of the movie-theater industry," *International Review of Law and Economics*, 27(2), 129–153.

Piccione, M., and R. Spiegler (2012): "Price competition under limited comparability," *The quarterly journal of economics*, 127(1), 97–135.

Prelec, D., and G. Loewenstein (1998): "The red and the black: Mental accounting of savings and debt," *Marketing science*, 17(1), 4–28.

Rhodes, A. (2014): "Multiproduct retailing," *The Review of Economic Studies*, 82(1), 360–390.

Schindler, R. M., and P. N. Kirby (1997): "Patterns of rightmost digits used in advertised prices: implications for nine-ending effects," *Journal of Consumer Research*, 24(2), 192–201.

Shin, J. (2005): "The role of selling costs in signaling price image," *Journal of Marketing Research*, 42(3), 302–312.

Simester, D. (1995): "Signalling price image using advertised prices," *Marketing Science*, 14(2), 166–188.

Snir, A., D. Levy, and H. A. Chen (2017): "End of 9-endings, price recall, and price perceptions," *Economics Letters*, 155, 157–163.

Stevens, L. (2020): "Coarse pricing policies," *The Review of Economic Studies*, 87(1), 420–453.

Stiving, M., and R. S. Winer (1997): "An empirical analysis of price endings with scanner data," *Journal of Consumer Research*, 24(1), 57–67.

Thomas, M., and V. Morwitz (2005): "Penny wise and pound foolish: the left-digit effect in price cognition," *Journal of Consumer Research*, 32(1), 54–64.

Tibshirani, R., G. Walther, and T. Hastie (2001): "Estimating the number of clusters in a data set via the gap statistic," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2), 411–423.

Voss, G. B., A. Parasuraman, and D. Grewal (1998): "The roles of price, performance, and expectations in determining satisfaction in service exchanges," *The Journal of Marketing*, pp. 46–61.

Wall Street Journal (2017): "Fashion Ads, a Last Bastion of Print, Are Going Digital," January 25, 2017.

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004): "Managerial and customer costs of price adjustment: direct evidence from industrial markets," *The review of Economics and Statistics*, 86(2), 514–533.

8 Figures

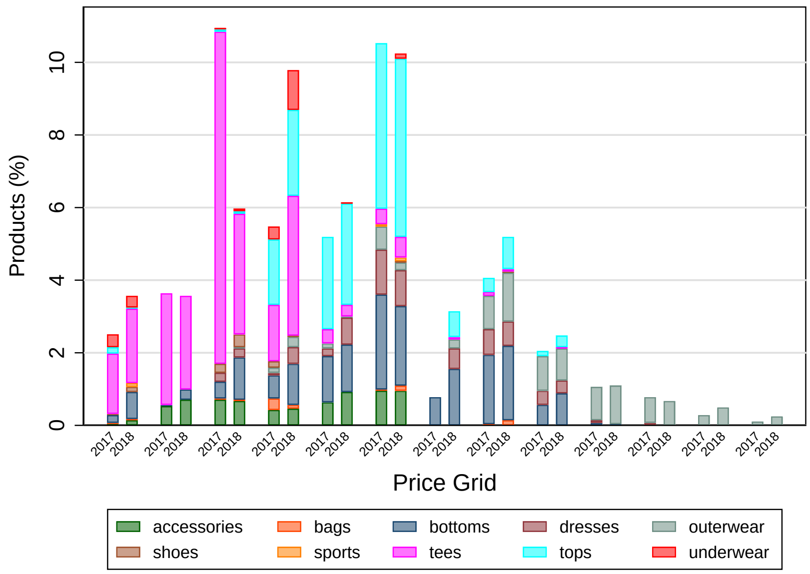

Figure 1: Uniqlo's Price Grid and Product Mix

Notes: Uniqlo's (U.S.) price distribution (without rounding) and the percent of products in across each price point. Data includes prices between $9 and $95, and Spring 2017 and Spring 2018.

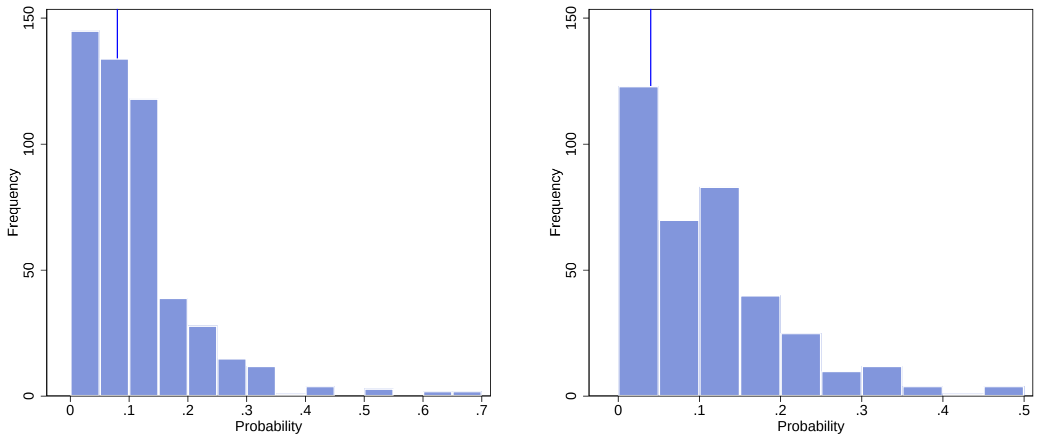

(a) US

(b) UK

Figure 2: Probability that two items have the same price

Notes: Histogram of the probability that two different items in the same retailer-category have the same price. Probability calculated at the retailer-category level.

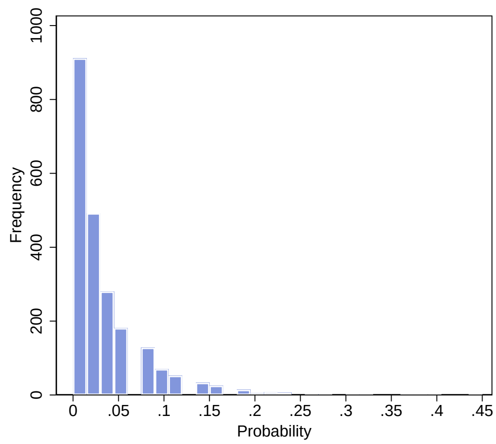

(a) US

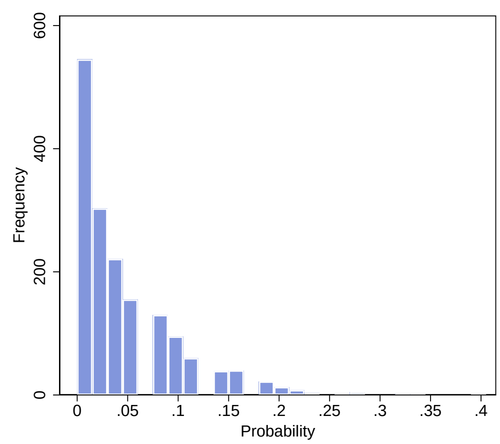

(b) UK

Figure 3: Probability of same price within-retailer and across-category

Notes: Histogram of the probability that two items from two different categories in the same retailer have the same price. Probability computed from each category pair within a retailer.

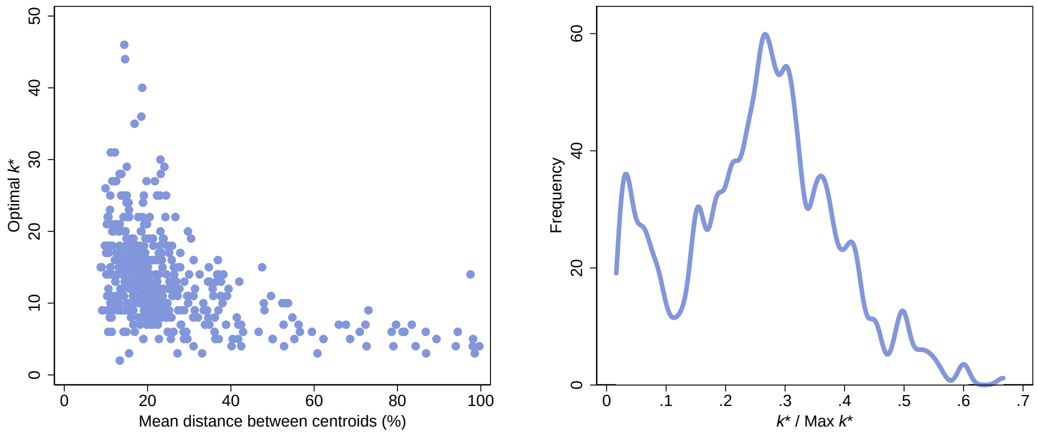

(a) US

(b) US

Figure 4: Optimal number of price clusters in the data

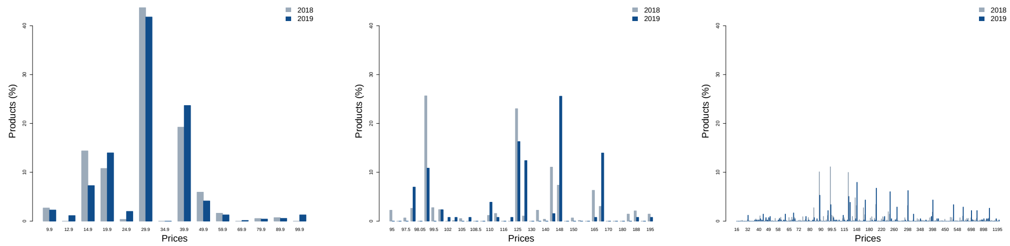

(a) Uniqlo US　　(b) Ralph Lauren US　　(c) Ralph Lauren US (Full)

Figure 5: Price adjustment

Notes: Panels (a) and (b) show the price distribution of the same categories and time periods in Uniqlo U.S. and Ralph Lauren U.S., respectively. The bars sum up to 100% each year. Data from the same month is used across the two years to account for seasonality. Results are similar using different months or years.

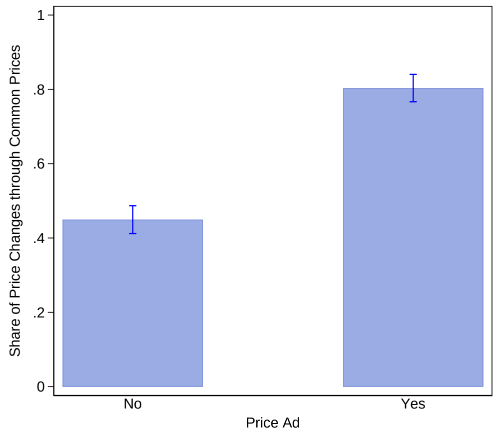

(a) Price salience

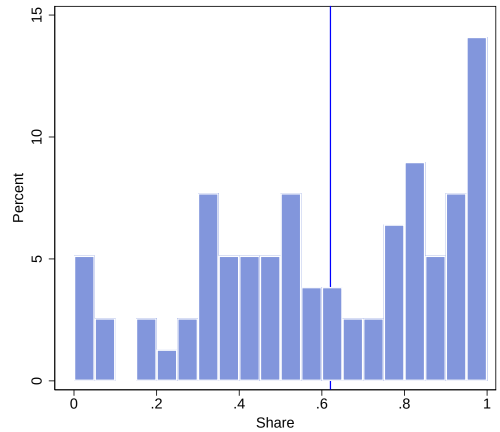

(b) Product shares in price buckets

Figure 6: Price changes through product mix

Notes: Panel (a) shows the average change in the price distribution that comes through modifying shares of products in the old prices. The measure is first calculated for all retailer-categories and then the average is reported across price- and non price- advertiser retailers. Error bars are SEMs. Panel (b) shows the histogram of the portion of the change in the price distribution that takes place via adjusting shares of products in the existing price buckets. Vertical line depicts the median share.

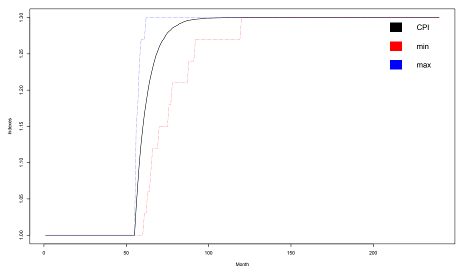

(a) Price Index: 0.1% of items

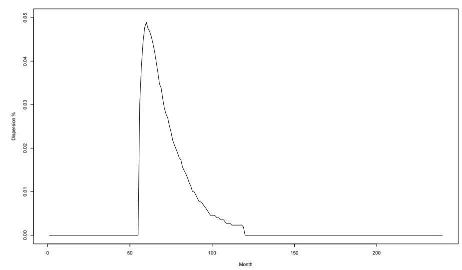

(b) Price Dispersion: 0.1% of items

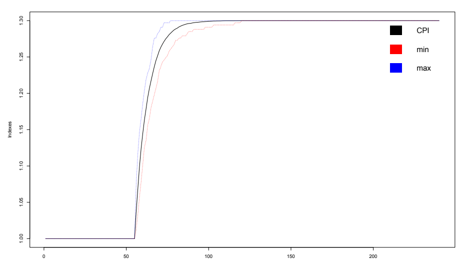

(c) Price Index: 1% of items

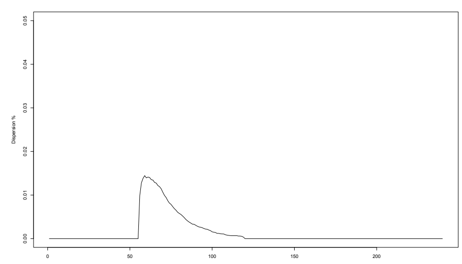

(d) Price Dispersion: 1% of items

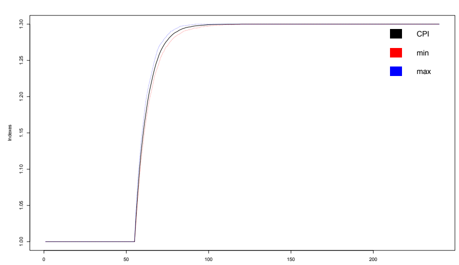

(e) Price Index: 10% of items

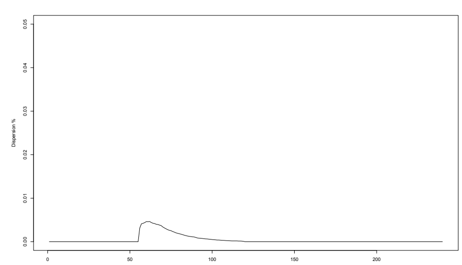

(f) Price Dispersion: 10% of items

Figure 7: Dispersion in Inflation Rate for Different Sample Sizes

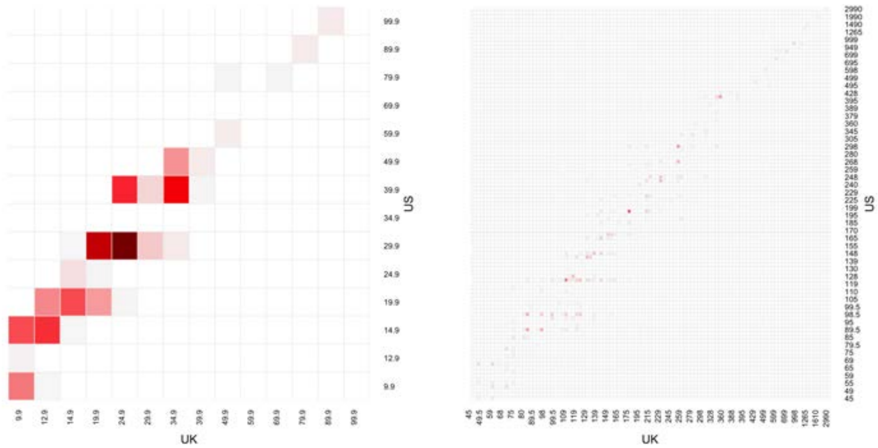

(a) Uniqlo (b) Ralph Lauren

Figure 8: Law of One Price

Notes: Panels (a) and (b) show a heatmap that represents the proportion of matched products allocated in each combination of U.S. dollars and U.K. pounds. In both panels the color key goes from the same range of 0% to 20% of products. Darker regions indicate larger proportion of products in a price bucket.

9 Tables

Table 1: Summary Statistics

		US	UK
(i)	Time period	March 2017 to May 2018	March 2017 to May 2018
(ii)	Average months per retailer[a]	4.6	4.5
(iii)	Observations[b]	241,928	199,619
(iv)	Retailers	54	40
(v)	Distinct goods	230,717	188,558
(vi)	Average distinct goods[c]	4,278	4,718
(vii)	Distinct goods (10%pct.)	1,279	1,122
(viii)	Distinct goods (90%pct.)	7,463	13,800
(ix)	Distance between prices (%)[d]	13	13
(x)	Average Items / Prices	49	87
(xi)	Items / Prices (10%pct.)	11	9
(xii)	Items / Prices (90%pct.)	92	192

Notes: [a]Equal weight average across retailers. [b]Excludes duplicates in terms of product, category, country, retailer. For example, the same product collected in two different months will appear only once. [c]Equal weight average across retailers. [d]Average distance across consecutive prices. Computed as equal weight average across retailers.

Table 2: Summary statistics on the normalized clustering index

		Index (1)	Index (2)	Index (3)	Index (4)
US					
(i)	Mean	0.075	0.083	0.044	0.025
(ii)	Median	0.098	0.106	0.063	0.038
(iii)	Percent of cases where:				
	< 0	0.2	0	0.8	3
	$[0, 0.025)$	5.4	7.9	23.6	46.8
	$[0.025, 0.05)$	23.6	19.4	34.5	27.8
	> 0.05	70.8	72.6	41.1	22.4
UK					
(i)	Mean	0.08	0.09	0.041	0.019
(ii)	Median	0.097	0.107	0.057	0.029
(iii)	Percent of cases where:				
	< 0	0	0	0	5.6
	$[0, 0.025)$	5.9	8.1	26.9	54.3
	$[0.025, 0.05)$	19.4	18.5	33.3	24.7
	> 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 preferred normalized index that controls for prices, range of prices, popular prices. Index (4) is similar to (3) and adds price ending dummies. Indices are defined in the main text.

Table 3: Examples of price clustering

		Retailer	$\frac{\text{Items}}{\text{Prices}}$ [a]	Prob.[b]	Prob.[c]	Index(1)[d]
	US					
(i)	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
	UK					
(iv)	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 among those that exhibit low, medium, and high price clustering. The following measures are within-retailer across-category averages. [a]The number of distinct items per distinct price. [b]Average probability that the price of two distinct items in the same category is the same. [c]Average probability that the price of two distinct items in different categories is the same. [d]Normalized clustering measure as defined in equation (1).

Table 4: Price stickiness

		m months after				
		$m = 1$	$m = 2$	$m = 3$	$m = 4$	$m = 5$
(i)	**Share of common prices**[a]					
	$p10\%$	0.73	0.63	0.57	0.48	0.41
	Median	0.91	0.89	0.86	0.81	0.75
	Average	0.91	0.87	0.86	0.83	0.76
	$p90\%$	1.0	1.0	1.0	0.99	0.92
(ii)	**Prob introducing a new good at existing prices**[b]					
	$p10\%$	0.61	0.53	0.48	0.55	0.60
	Median	0.96	0.96	0.93	0.93	0.92
	Average	0.96	0.93	0.91	0.93	0.91
	$p90\%$	1	1	1	1	1

Notes: Results are averages across retailer-categories in the U.S. [a]Ratio of common prices (between the prices observed in the first month and month m) to the prices in month m. [b]Probability that the price of a product introduction in month m was among the observed prices in the first month.