UC Davis

UC Davis Previously Published Works

Title

What is the price of tea in China? Goods prices and availability in Chinese cities

Permalink

<https://escholarship.org/uc/item/0027s8tv>

Journal

Economic Journal, 130(632)

ISSN

0013-0133

Authors

Feenstra, RC

Xu, M

Antoniades, A

Publication Date

2020

DOI

10.1093/EJ/UEAA066

Peer reviewed

What is the Price of Tea in China?

Goods Prices and Availability in Chinese Cities *

Robert C. Feenstra

University of California, Davis

and National Bureau of Economic Research

Mingzhi Xu

INSE, Peking University

Alexis Antoniadis

Georgetown University

Abstract

We examine the price and variety of a sample of consumer goods at the barcode level in cities within China. Unlike the United States, in China the prices of goods tend to be lower in larger cities. We explain that difference between the countries by the more uneven spatial distribution of manufacturers' sales and retailers in China, and we confirm the pro-competitive effect of city size on reducing markups there. In both countries, there is a greater variety of goods in larger cities, but that effect is more pronounced in China. Combining the lower prices and greater variety, the price indexes in China for the goods we study fall with city size by around seven times more than in the United States.

Key Words: Retailing prices; Multi-product firms; Pro-competitive effect;

Spatial price difference; Cost-of-living

JEL Classification: E01,F11,L1

*We thank Kyle Bagwell, John Romalis and conference participants at the NBER, Stanford and the North American Summer meetings of the Econometric Society for their helpful comments, along with Qi Liu for research assistance. Financial support from the National Science Foundation and the China Scholarship Council is gratefully acknowledged. Calculations are based on data from the Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at the University of Chicago Booth School of Business, and from the Nielsen Company (China). ©The Nielsen Company. Please direct correspondence to: Robert Feenstra, University of California, Davis, rcfenstra@ucdavis.edu.

Research using prices of consumer goods at the barcode level has shown that these prices do not differ by much across cities in the United States, or across countries with a common currency such as the Euro. Evidence for the United States is provided by Handbury and Weinstein (2015), who find only slight positive effects of population and household income on prices, once they control for barcode, retailer, and city fixed effects, along with household demographics. They confirm, however, that the available product variety of barcode items rises with city size so that the variety-adjusted price index falls. DellaVigna and Gentzkow (2017) investigate the prices of barcode items within retail chains across the United States, and they find that for many chains the prices do not appear to differ at all within states, though there is more price variation between retail chains. They cite similar evidence for the United Kingdom: “roughly half of UK supermarket chains charge uniform prices across stores, too...as do the main UK electronics retailers”.¹ Broadening the coverage of prices and countries to those in the Billion Prices Project, Cavallo, Neiman, and Rigobon (2014) show that eight major retailers in goods comprising about 20% of final consumption expenditure charge nearly uniform prices across the Euro zone.² It is surprising that none of these studies find any evidence for a pro-competitive effect, whereby larger markets lead to lower prices. Prior studies such as Campbell and Hopenhayn (2005) argued that this effect operated, but in the absence of prices at the barcode level, their evidence was actually on the sales of retail firms rather than their prices.

In this paper we present contrasting findings to the above price results using city evidence from China, where we confirm a pro-competitive effect. Unlike the United States, prices at the barcode level are not readily available for China.³ Nevertheless, we were able to obtain the barcode prices and expenditures for four consumer goods across 22 cities in China from Nielsen (China). In addition, we expanded the scope of cities to 60 and the scope of products by adding another 15 goods, including Tea, with all those prices in China collected from a mobile phone application called *Wochacha* (translated as “let me check”). We introduce these data in section 2 and use them to demonstrate a negative relationship between prices at the barcode level and city size for China. We also show that this relationship does not hold consistently using barcode prices for the same goods sold across cities in the United States. We explain that difference between the countries by the more uneven spatial distribution of manufacturers’ sales and retailers in China.

¹DellaVigna and Gentzkow (2017), footnote 9.

²While the dataset for the retailers was for their online prices, Cavallo, Neiman, and Rigobon (2015) confirm that essentially the same prices apply to the offline stores, except for occasional sales within the stores.

³Data at the barcode level for the United States is available by subscription from Nielsen at the Kilts Center for Marketing, <https://research.chicagobooth.edu/nielsen>.

To test more formally for a pro-competitive effect in China, we use a model that explicitly links markups to the market shares of firms. We assume a nested, constant elasticity of substitution (CES) utility function for the representative consumer, and realistically we allow for multi-product firms. This means that firm markups are not fixed by the CES assumption but depend positively on the market shares of these firms, as in Atkeson and Burstein (2008). In other words, these large producers are acting as oligopolists rather than as monopolistically competitive firms. We allow for free entry of these firms, which have varying productivities. In this model, outlined in section 2, larger markets can accommodate more firms and so prices and markups are lower as a result. That result depends on the specification of the sunk costs of entry, however. As in Sutton (1991), if sunk costs are endogenous then the impact of market size on reducing concentration and prices is diminished.

In section 4 we test for the pro-competitive effect of city size on reducing markups in China. For the four main consumer goods, we find support for this hypothesis when using the market share of manufacturing firms in each city to infer their demand elasticities and markups. Still, this approach cannot entirely explain the negative relationship between prices and city size, so we extend our model to also incorporate retail outlets. We find that larger market shares of retailers in smaller cities in China can further explain a portion of higher prices in those cities.

Our model also has implications for product variety that are similar to what Handbury and Weinstein (2015) found for the United States: larger cities have more product variety. Indeed, this effect is more pronounced in China than in the U.S. In section 5 we present these results for our four main goods. For China, the median elasticity of the geometric mean of price with respect to city size (measured by population or GDP) is -0.01 , which is what Handbury and Weinstein (2015) find for the elasticity of *variety-adjusted* price indexes in the United States. The elasticity of the variety-adjusted prices for China shows that these prices fall by as much as seven times more than found by them when we use GDP to measure city size, and fall by at least seven times more when we use city population (as they do for the U.S.), or around seven times more on average.

In section 6, we discuss the implications of our finding for the comparison of prices *between* countries. Comparing barcode prices directly between China and the U.S. is very difficult because the barcode systems are different in the two countries.⁴ Fortunately, we could address this issue by re-

⁴The barcode systems used in China is the European Article Number (EAN-13), whereas that used in the United States is the Universal Product Code (UPC). While it is not difficult to identify similar product *categories*, it is impossible to *exactly* match the items across countries by their barcodes since the barcode systems differ and product descriptions themselves do not match. In online Appendix A3, we provide an example of items that are seemingly identical in China and the U.S. but have different EAN and UPC numbers.

lying on the new price index results in Redding and Weinstein (2018). In the four main goods for which we purchased Nielsen (China) data, there is a significant presence of U.S. brands in the Chinese market, and then the availability of additional Chinese brands leads to greater variety than in the United States and lower Chinese price indexes for that reason. In the other 15 goods for which we only have scraped data, however, there is much less presence of U.S. brands in the Chinese market. In these cases, the observed prices differences between the countries (usually lower prices in China) are partially or fully offset by the variety differences (less variety in China), so that the cost of living in China is not as low as the pure price differences suggest. Tea is an example where the availability of brands with barcodes in China is less than in the U.S., but in that case, we are clearly missing the many Chinese varieties of loose tea sold without barcodes. Section 7 concludes and additional material is gathered in the online Appendix.

1 Prices and Market Shares in China

1.1 Price and Market Share Data

Our price data are obtained from three sources. The first source is the Nielsen (China) Sales database from which we have the value and quantity sold of barcode items in 2013.⁵ We purchased the scanner data for four goods, namely, Laundry Detergent, Personal Wash items, Shampoo, and Toothpaste, covering 22 cities in China, as shown in the red colored regions of Figure 1.⁶ Besides the sales information for each product, we also observe manufacturer information such as brand and sub-brand of each product, and the total quantity of each barcode aggregated over the entire city. In our theory and empirical analysis, we refer to the *brand* as the manufacturing firm (e.g. Tide for laundry detergent or Crest for toothpaste).⁷ So this source of data is limited to only four goods and 22 cities, and the *average price* obtained by dividing the annual sales by the quantity is for the *total* sales of each barcode item in each city, i.e. there is no information provided on the retail outlets.

⁵We convert RMB data to U.S. dollars using the annual average exchange rate.

⁶We use 2011 and 2012 sales information for toothpaste, and 2014 for the other three product categories.

⁷For China, we observe the *brand* (e.g., Crest) and *manufacturer* (e.g., Procter & Gamble) in the Nielsen (China) data. For the United States, we use the first two keywords in the brand description provided by the U.S. Nielsen data to identify brands. For cases with an ambiguous description such as *CLT BR*, we use the first five UPC digits to identify brands (or firms) according to Muth et al. (2016). For example, we treat Crest and Oral-B as two distinct brands, even though both belong to Procter & Gamble.

Figure 1: Regions included in Nielsen Sales Data and Scraped Price Data

The regions covered by Nielsen (China) database are for larger cities (many of them are capital cities). To address this limitation, our analysis relies on our second data source, which are scraped prices collected in 2015 from the mobile phone application *Wochacha* that allows consumers to check for the various prices at supermarkets in each city, i.e. these are offline prices posted in the store but with information provided online.⁸ Details on this second source of data are provided in online Appendix A, and it allows us also to include smaller cities in our sample, with prices for the four main and 15 additional goods. For each barcode item, we obtain a *single* price quote for each retail chain in each city. We are including within our definition of *retail chain* the large multi-outlet retailers (e.g. Walmart or Carrefour, both of which operate in China) and also individual retail outlets that operate in a single city.⁹ Even if there might be more than one outlet of that retail chain in a city, we obtain only a single price quote. That still improves upon the Nielsen (China) dataset that provides for each barcode item and each city only a single price (averaged over all retailers for 2014), for the four main goods. Using the *Wochacha* data, we are able to expand our sample to 60 cities, including the 22 cities provided by Nielsen (China) database, as shown in the yellow colored regions of Figure 1. Our third source of data is the Nielsen Retailing Sales (U.S.) database for 377 metropolitan statistical areas

⁸For studies of online shopping in China see Couture, Faber, Gu, and Liu (2018) and Fan, Tang, Zhu, and Zou (2018).

⁹For China, we directly observe the retailer name in each city from *Wochacha*, e.g., Shanghai-Walmart and Changsha-Walmart (retail chains), or Changsha-Weiran (an individual retail outlet only in Changsha). For the United States, we define a retail chain to be a unique combination of parent code and retailer code, as in DellaVigna and Gentzkow (2017).

(MSAs) of the United States in 2013, at the barcode level for each of the four main and 15 additional goods, with complete information on retailers.

We stress that both our Chinese data (from Nielsen or from *Wochacha*) and our U.S. data do not include all retailers in a city/MSA, but rather, they both focus on large supermarkets (and drug stores for the U.S.). The Nielsen Retailing Sales (U.S.) database reports that it collects: “scanner data from more than 35,000 participating grocery, drug, mass merchandiser, and other stores, covering more than half the total sales volume of US grocery and drug stores”.¹⁰ The comparable fraction for China could be lower due to the prevalence of small, traditional retailers in many parts of that country. Lagakos (2016) argues that in developing countries it may be natural to have more traditional retailers (with low productivity), because consumers do not have access to public transportation that would enable them to benefit from modern supermarkets.¹¹ We do not view the lack of complete data coverage as biasing our price comparisons within or between countries, because we control for the retailer with fixed effects, and so we are interested in the prices of the identical barcode items from the same retailer across regions. The lack of data coverage must be taken into account when constructing the market shares for retailers, however.¹²

In Table 1 we provide summary statistics for the city/MSA characteristics of our U.S. and Chinese sample, including city GDP, population and income per-capita, along with summary statistics for our four main goods, such as the number of retail chains (across all four goods) and the number of brands (for each product) and average price (in US cents). We can see from Table 1 that the sample of 60 Chinese cities includes some smaller cities than in the initial 22 city sample: the ratio of largest to smallest cities by GDP in the 60 city sample is about 32 times, but 7 times for the 22 cities; similarly, in terms of population, the ratio of largest to smallest cities in the 60 city sample is about 18 times, but 8 times for the 22 cities. Still, the smallest population in our Chinese sample of 60 cities is 1.6 million, which obviously excludes many smaller cities in that country. By comparison, the range of MSA GDP and population for the United States is 774 and 338 times, respectively.

¹⁰See: <https://research.chicagobooth.edu/nielsen/datasets#simple2>.

¹¹Evidence to support this idea for China comes from data on car ownership, which in 2014 was 16.9% in China versus 91% in the United States (data from *China Statistical Yearbook* and *American Consumer Survey*). Indeed, Walmart or Carrefour both provide bus service to their stores in China; see <http://www.businessinsider.com/photos-from-shopping-at-chinese-wal-mart-2014-1>.

¹²We will later (in section 4) parameterize the share of retailers covered by our Nielsen (China) data as 50%, like in the United States, but for robustness we also experiment with values as low as 10% as reported in online Appendix D2.

Table 1: Summary Statistics of the United States and China's Cities

Data Sample	Variable	Num. Obs.	Mean	Std	Min	Max
China 22 Cities (Nielsen):						
	GDP (\$ million)	22	129,985	73,884	43,616	312,469
	Per-capita GDP (\$)	22	12,352	3,793	6,025	19,082
	Population (million)	22	10.84	6.27	3.64	29.32
China 60 Cities (Wochacha):						
	GDP (\$ million)	60	75,210	64,716	9,727	312,469
	Per-capita GDP (\$)	60	9,658	4,528	2,680	21,996
	Population (million)	60	7.39	4.92	1.61	29.32
	<i>Four Main Goods</i>					
	Number of Retail Chains	60	5.85	3.63	1.00	16.00
	Num. of Retail Chains / 1 Million Pop.	60	0.85	0.45	0.15	2.47
<u>Laundry Detergent:</u>	Number of Brands	60	28.1	4.7	15.0	36.0
	Average Price (\$)	60	3.52	0.26	3.02	4.22
	Herfindahl Index	60	0.16	0.08	0.10	0.66
	Top 4 Brands Share	60	0.68	0.07	0.54	0.95
<u>Personal Wash Items:</u>	Number of Brands	60	50.8	9.7	31.0	66.0
	Average Price (\$)	60	4.04	0.18	3.61	4.53
	Herfindahl Index	60	0.13	0.03	0.07	0.28
	Top 4 Brands Share	60	0.61	0.06	0.44	0.77
<u>Shampoo:</u>	Number of Brands	60	46.0	5.4	33.0	54.0
	Average Price (\$)	60	6.35	0.26	5.61	7.34
	Herfindahl Index	60	0.11	0.01	0.09	0.12
	Top 4 Brands Share	60	0.56	0.02	0.49	0.60
<u>Toothpaste:</u>	Number of Brands	60	23.9	5.5	15.0	34.0
	Average Price (\$)	60	2.06	0.17	1.75	2.57
	Herfindahl Index	60	0.26	0.10	0.13	0.55
	Top 4 Brands Share	60	0.86	0.09	0.62	0.99
United States MSA:						
	GDP (\$ million)	377	36,754	103,573	1,775	1,374,000
	Per-capita GDP (\$)	377	40,411	12,113	17,533	108,016
	Population (million)	377	0.72	1.65	0.06	20.29
	<i>Four Main Goods</i>					
	Number of Retail Chains	377	7.18	2.26	2.00	20.00
	Num. of Retail Chains / 1 Million Pop.	377	28.74	19.01	0.97	102.70
<u>Laundry Detergent:</u>	Number of Brands	377	44.0	6.0	28.0	69.0
	Average Price (\$)	377	5.83	0.32	4.85	6.86
	Herfindahl Index	377	0.16	0.03	0.10	0.27
	Top 4 Brands Share	377	0.61	0.03	0.52	0.70
<u>Personal Wash Items:</u>	Number of Brands	377	132.0	31.3	32.0	299.0
	Average Price (\$)	377	5.98	0.30	3.60	6.92
	Herfindahl Index	377	0.06	0.01	0.04	0.09
	Top 4 Brands Share	377	0.37	0.03	0.30	0.54
<u>Shampoo:</u>	Number of Brands	377	123.4	25.4	43.0	248.0
	Average Price (\$)	377	6.06	0.60	3.21	8.00
	Herfindahl Index	377	0.06	0.01	0.04	0.09
	Top 4 Brands Share	377	0.43	0.03	0.32	0.52
<u>Toothpaste:</u>	Number of Brands	377	38.4	4.6	17.0	63.0
	Average Price (\$)	377	3.26	0.19	2.68	3.86
	Herfindahl Index	377	0.26	0.02	0.21	0.34
	Top 4 Brands Share	377	0.85	0.02	0.78	0.90

Notes: The table summarizes the city and product statistics for the United States and China. GDP is in millions of dollars; per-capita GDP is in dollars; population is in millions. We convert RMB to US dollars using the annual average exchange rate for 2013, i.e., 6.46 RMB/USD. Average prices are calculated based on the corresponding sampled items. Prices (in US cents) are firstly averaged at barcode-city/MSA level weighted by quantity (for US and China Nilsen; the mean is unweighted for *Wochacha* across retailers within a city), and then are arithmetically averaged across barcode at city/MSA level. The city and MSA macro data are from China Statistical Yearbook (2013) and U.S. Bureau of Economic Analysis (2013).

Figure 2: Number of Retail Chains in the U.S. and China (LOWESS smoothing)

Despite that very wide range of city sizes in the U.S., the number of retailers for the four main goods found in each city/MSA is similar across the two countries, ranging from 2 to 20 in the United States and from 1 to 16 in China. The *distribution* of the number of retail chains across cities by population looks quite different, however, as shown in Figure 2, which is the locally weighted scatterplot smoothing (LOWESS) curve, demonstrating the distribution of retailing chains across cities/MSAs for both China and the United States according to the city percentile by population, with unity indicating the largest city in each country. The figure shows that the number of retail chains is more evenly distributed in the United States, as the rise in number of retailing chains is much flatter when moving from small cities to large ones as compared to China.¹³

Also shown in Table 1, the number of retail chains per 1 million population ranges from about 1 to 100 in the United States, but from 0.15 to 2.5 in China. So the *density* of retailers is significantly lower in China. Consistent with our data sample, the *total* number of retail chain stores per million population is 811 for the U.S. and only 159 for China.¹⁴ Walmart for example, had 432 stores in all of

¹³The more uneven distribution in China would appear more vividly if we used a common scale for the percentile distribution of city sizes by combining the city populations in both countries. Since the smallest Chinese city in our sample of 1.6 million is in fact at the 90th percentile of city size in the United States, the smoothed curve for China would start at about 0.9 on the horizontal axis of the common scale.

¹⁴The number of retail chain stores in the U.S. is from *The Future of Food Retailing 2017* (page 6), and that number in China is from National Bureau of Statistics of China.

China in 2017 (Fung Business Intelligence, 2017), and more than 5,000 stores in the United States.¹⁵ The low density combined with the steeper rise of retail chains with city population in China, as shown in Figure 2, means that the smaller cities in our Chinese sample (which still have populations exceeding 1.6 million) have a distinct lack of large retailers to choose from as compared to larger cities in China or nearly any MSAs in the United States. We therefore expect to find reduced competition between these retailers in smaller cities in China.

Besides the difference in retailer chains, there is also a notable difference in the markets share of brands (e.g. Tide or Crest) between China and the United States. The price data that we scraped from *Wochacha* do not provide us with consumer expenditures, so our only source for such data is from Nielsen (China) data for the four main goods over 22 cities. Recall that the Nielsen's data is limited, however, by having only the total city expenditure for each barcode item, and likewise a single *average* price (over the year) for that barcode item in each of 22 cities. In comparison, the *Wochacha* data provide us with a price quote for each barcode and each retail chain in 60 cities. Conveniently, for the four main goods we can merge the Nielsen's and *Wochacha* data to have both consumer expenditure for each barcode item in each city and a list of the retail chains selling it.

Using these merged data, we estimate a simple demand equation that will be consistent with our constant elasticity of substitution (CES) framework introduced in the next section:

$$\ln(p_{ifc}x_{ifc}) = (1 - \hat{\sigma}) \ln p_{ifc} + \alpha_i + \sum_{r \in R_{ic}} \alpha_r + \beta \ln Y_c \quad (1)$$

where α_i denotes a fixed effect for each barcode, $\sum_{r \in R_{ic}} \alpha_r$ is the sum of retailer fixed effects for the set of retailers R_{ic} selling each barcode in each city, and Y_c denotes city GDP. Both the expenditure data on the left and the average price on the right are from the Nielsen (China) data, i.e. they are annual for the expenditure and price by barcode i , firm f , and city c . Such annual average barcode prices are subject to measurement error due to sales and other promotions, so we do not expect OLS estimation of (1) to provide an unbiased estimate of the demand elasticity σ (the elasticity of substitution in our CES framework) which we allow to differ across the four goods. Accordingly, we implement the generalized method of moments estimator of Hotman, Redding, and Weinstein (2016) to obtain these elasticities for China, which corrects for measurement error and any simultaneous equation bias (see online Appendix B1). Those authors estimate the elasticities from the Nielsen Retailing

¹⁵See: <https://www.statista.com/statistics/269425/total-number-of-walmart-stores-in-the-united-states-by-type/>.

Sales (U.S.) database, while we use the Nielsen (China) data.¹⁶ Their technique double-differences across barcodes and cities, so that σ is estimated but no other coefficients of (1). We then substitute the estimated values of $\hat{\sigma}$ back into (1), and estimate the fixed effects for barcodes and retailers along with the income elasticity using OLS, with the results shown in Table 2. The high R^2 in Table 2 indicates the good performance of (1) for predicting expenditures by barcode within the 22 cities for which we have sales from Nielsen (China). We use this regression to predict sales by barcode *and by* retailer within the 22 cities as $(1 - \hat{\sigma}) \ln p_{ifc} + \hat{\alpha}_i + \hat{\alpha}_r + \hat{\beta} \ln Y_c$, and for the other 38 cities in our 60 city sample by $\hat{\alpha}_i + \hat{\alpha}_r + \hat{\beta} \ln Y_c$.¹⁷

Table 2: Summary of Sales Prediction

	Barcode Elas. (σ)	Income Elas.	Barcode FE	Retailer FE	R^2
(1) Laundry Detergent	5.55	0.83	Y	Y	0.94
(2) Personal Wash	3.85	1.03	Y	Y	0.83
(3) Shampoo	2.23	0.83	Y	Y	0.81
(4) Toothpaste	10.30	0.68	Y	Y	0.95

Notes: Each row reports the regression specified by (1), where the dependent variable is the annual sales of barcode items by city. The barcode elasticities for each good are estimated following the method explained in Appendix B. The income elasticity reports the estimated coefficient of $\ln Y_c$, and barcode and retailer fixed effects are included.

We then compute the Herfindahl indexes and the market share of the top 4 brands, which we report in Table 1 along with the number of brands and the average barcode price (from *Wochacha*) across cities for each of the four main goods. In three of these goods – Laundry Detergent, Personal Wash items and Toothpaste – the maximum values of the Herfindahl indexes and the 4-brand concentration ratio are noticeable larger in China than in the United States. These maximum values tend to occur in the smaller cities in China, indicating the greater concentration there than for U.S. cities.¹⁸ So as we find for retailers, we likewise find that the smaller cities in China have reduced competition between manufacturing firms.

¹⁶Given that our framework is similar to Hottman, Redding, and Weinstein (2016), we follow their method to obtain both product and firm elasticities for China. The estimated product elasticities are shown in Table 2 and the firm elasticities are 4.78, 2.88, 2.21 and 4.29 for Laundry Detergent, Personal Wash items, Shampoo, and Toothpaste, respectively.

¹⁷The inclusion of the firm-level price p_{ifc} in (1), used to estimate the demand elasticity σ as explained in online Appendix B1, is unimportant to the prediction of expenditures which depends almost entirely on the barcode and retailer fixed effects.

¹⁸In the online Appendix, Figure A.3 compares the histograms of brand market shares in the United States and China, which also show greater concentration in China.

1.2 Price Regressions

Table 3: Price Regressions for Consumer Goods, United States and China

	United States				China			
	ln(Income)	ln(Pop)	N	R ²	ln(Income)	ln(Pop)	N	R ²
Four Main Goods								
Laundry detergent	8.22*** (1.43)	0.74 (0.53)	521,963	0.98	-5.00*** (1.65)	-3.96*** (0.80)	103,227	0.94
Personal Wash Items	7.86*** (1.416)	-0.46 (0.917)	744,714	0.93	-5.83 (3.84)	-2.35 (1.86)	151,826	0.92
Shampoo	6.79*** (1.79)	-0.13 (0.64)	1,054,261	0.95	-8.26** (3.15)	-5.48*** (1.81)	136,091	0.92
Toothpaste	5.46*** (1.61)	0.42 (0.54)	515,372	0.91	-5.57*** (1.09)	-4.12*** (1.32)	86,619	0.94
15 Other Goods								
Baby Formula	3.03 (4.49)	0.99 (1.54)	140,095	0.96	-67.93*** (13.220)	-13.17 (11.28)	41,30	0.98
Battery	6.27 (5.14)	-3.69*** (1.05)	345,787	0.85	-0.36 (3.56)	-5.73* (3.21)	14,735	0.98
Biscuit	3.17*** (0.69)	0.23 (0.17)	2,558,449	0.93	-0.99 (2.77)	0.05 (1.62)	145,983	0.94
Cereal	5.95*** (1.16)	0.41 (0.45)	640,753	0.83	-9.48*** (2.86)	-4.03* (2.06)	19,741	0.88
Chips	3.72*** (0.84)	0.37* (0.20)	375,212	0.88	-2.33** (0.90)	-1.70** (0.80)	18,777	0.94
Chocolate	6.03*** (1.23)	-0.32 (0.22)	1,153,028	0.94	-5.73 (4.54)	-3.87 (4.01)	11,997	0.94
Coffee	6.06*** (1.72)	0.00 (0.71)	103,453	0.92	-9.71*** (3.16)	-9.50** (4.34)	22,945	0.95
Diaper	5.58 (3.720)	-1.59* (0.92)	598,382	0.96	-42.25** (16.02)	-11.51 (7.10)	17,683	0.96
Dog Food	4.04 (3.63)	2.01*** (0.75)	273,199	0.96	-11.56*** (2.71)	-11.20*** (2.19)	8,438	0.98
Gum	1.34 (1.02)	0.07 (0.52)	37,296	0.93	-2.19 (1.36)	-0.60 (0.77)	10,975	0.88
Milk	0.37 (1.75)	1.10* (0.64)	90,604	0.90	-4.54*** (1.45)	-1.98 (1.72)	17,29	0.95
Sanitizer	3.82*** (1.24)	0.57* (0.31)	512,976	0.93	-3.64* (2.16)	-1.51 (1.40)	26,436	0.93
Soft Drink	4.19*** (0.77)	0.42 (0.26)	1,192,077	0.96	-3.31*** (0.77)	-1.79*** (0.52)	30,665	0.96
Tea	4.19*** (1.25)	0.07 (0.33)	702,523	0.93	-8.32*** (2.29)	-4.74*** (1.09)	10,445	0.94
Toothbrush	13.18*** (2.30)	-1.74 (1.18)	355,316	0.97	-1.42 (1.35)	-1.28 (2.32)	57,872	0.86

Notes: Each row includes two regressions as shown by equation (2), one for the U.S. and one for China. The dependent variable is per-item price in US cent and each observation is at the barcode-retailer-city level. The independent variables are ln *Income* measured as GDP per capita in each city, ln *Population*, and barcode and retailer fixed effects. Robust standard errors are clustered at the barcode, retailer and city/MSA level and reported in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

In Table 3, we report price regressions for the four main goods and the 15 additional goods for which we have scraped barcode data for 60 cities in China, along with regressions for comparable products in the United States.¹⁹ The dependent variable is the level of the price at the barcode level

¹⁹In online Appendix B3, we report the Chinese price regressions run over the smaller set of 22 cities for which we

(in cents).²⁰ For China, we consistently use the single barcode price quote for each retailer in the 60 cities that was collected from *Wochacha*. For the United States, we use the Nielsen Retailing Sales (U.S.) database to construct a similar (average) price for each barcode and retailer in each of 377 MSAs. The regression is specified as:

$$p_{ifrc} = \alpha_i + \alpha_r + \beta_1 \ln w_c + \beta_2 \ln L_c. \quad (2)$$

On the right of (2) we include fixed effects for barcode and retailer, along with the log income per-capita $\ln w_c$ and the log population $\ln L_c$ of each city, the sum of which are log city GDP, $\ln Y_c = \ln w_c + \ln L_c$. By including fixed effects for barcodes (and hence firms) and also retailers, we are attempting to control for the marginal costs of production for each item.²¹ Of course, income per-capita may be associated with higher marginal costs, too (through higher wages and rents in wealthier cities). For that reason and because firms might charge higher markups to wealthier consumers, we expect to find $\beta_1 \geq 0$. In contrast, city size as measured by population will have a negative impact on prices if there is a pro-competitive effect, so we expect to find $\beta_2 \leq 0$.

The results in Table 3 for the United States, on the left, show that city income per-capita has a persistently positive sign, indicating higher prices in higher income cities, whereas city population has no systematic pattern: its coefficient is either positive or negative, and most often insignificant. The finding that prices are rising with income are consistent with the results of Handbury and Weinstein (2015) and DellaVigna and Gentzkow (2017), though both these studies argue that the positive estimated effect can be biased upwards for various reasons.²² Handbury and Weinstein (2015) find a slight positive impact of city population on prices when pooling across all goods, whereas DellaV-

have price data from *Wochacha* and Nielsen (China). Over this smaller set of Chinese cities, the negative relationship between price and city size observed in Table 3 does not hold significantly, especially when using the (annual average) barcode prices from Nielsen (China). Thus, it was important to expand our sample to 60 cities using the scraped data from *Wochacha*.

²⁰We have run the price regressions in both the level of price and the log of price, with similar qualitative results, and a Box-Cox regression does not prefer one specification over the other. The percentage impact of city size on reducing prices in China is small, however, so the results are more easily seen in levels. Regressions of the log of the city price indexes on log city GDP and population are reported in section 5.

²¹In section 4, we will add provincial fixed effects as well as distance from the nearest factory for each firm to each city to estimate the shipping costs. We have collected that distance information only for the four main goods in China, so it is not used in Table 3.

²²Handbury and Weinstein (2015) show that the effect of income is biased upwards when using the prices of similar but not identical products across cities, or when using identical barcode products but insufficient controls, whereas DellaVigna and Gentzkow (2017) argue that an upward bias occurs even using identical products but using the weekly average price. In Table 3 we are using the annual average price of each barcode item for each retailer and MSA in the United States, consistent with our Chinese dataset where we have a single observation for each barcode-retailer-city. If instead we disaggregate by all the outlets for each retailer in the U.S., then we find stronger positive impacts of MSA income and of city population on prices.

igna and Gentzkow (2017) do not directly investigate the impact of population on prices. Their evidence that prices are nearly uniform within many retail chains across MSAs, however, is consistent with little or no impact of city population. Our results in Table 3 pool across retail chains and non-chain retailers, and we likewise find very little evidence of prices differing systematically with city population in the United States.

The results for China, on the right of Table 3, are quite different. Both city income per-capita and population often appear with negative and significant coefficients. Indeed, in 14 out of the 19 cases, we cannot reject the hypothesis that $\beta_1 = \beta_2 < 0$, so we interpret the negative signs on both coefficients as reflected a *pro-competitive effect of city size (measured by GDP) in reducing prices*. In other words, whether larger city size is due to higher population or higher income per-capita, both variables are associated with reduced prices.²³ We cannot rule out the idea that, *ceteris paribus*, firms in Chinese cities with wealthier consumers would charge higher markups, as we find for the United States, but evidently any such effect is overwhelmed by the negative impact of city size on prices.²⁴ In economic terms, the median coefficient on $\ln w_c$ or $\ln L_c$ computed over the four main goods or the 15 others in Table 3 is about -5 cents, and multiplying that by the natural log of the ratio of highest to lowest GDP among our 60 Chinese cities (which is 3.5), we obtain -17 cents. The average prices for Laundry Detergent, Personal Wash Items and Shampoo are about \$6, so the price reduction is 3 percent of the price, while the average price for Toothpaste is about \$3 so the price reduction is 6 percent. While these pro-competitive effects are not that large, in section 5 we will incorporate the positive impact of city size on the variety of products available, and we will find a more substantial impact on consumers.

Why do we observe such different results for China and the United States? In the remainder of the paper, we will outline a model that exhibits pro-competitive effects, and we shall test this hypothesis formally for China. We will show that the higher prices in small cities in China can be substantially explained by reduced competition between firms and between retailers. The surprising feature of the results in Table 3 is that there is *no systematic evidence* of such a pro-competitive effect

²³This result will be consistent with the nested, CES preferences that we have adopted, depending on the sunk costs of entry as explained in section 3.4, but we stress that it is sensitive to the specification of preferences. Parenti, Ushchev, and Thisse (2017) consider a general class of preferences that allow for differential responses of markups to changes in income per-capita versus population, particularly because rising income per-capita can indicate greater willingness to pay and lead to higher markups.

²⁴To distinguish income per-capita from the inequality of income, we constructed Gini coefficients for each city using the distribution of wages paid by firms. For two products (Laundry Detergent and Shampoo) the Gini has a positive and highly significant impact on price, indicating that greater inequality leads to higher prices. For the other two products the impact of the Gini on price was insignificant (positive for Toothpaste and negative for Personal Wash items).

for the United States. DellaVigna and Gentzkow (2017) argue that such lack of price discrimination is evidence of *managerial decision making costs*, i.e. retailers are not pricing so as to maximize profits. Our contrasting results for China suggest an alternative explanation to account for the lack of price variation in the U.S, and that is the very different spatial distribution of manufacturers and retailers across cities in the two countries. Specifically, we have shown above that there are fewer retailers in the smaller cities of China, and also a greater concentration of producers. In the United States, by contrast, we have shown that cities do not differ by nearly the same extent in the concentration of producers and retailers, even though the MSAs themselves vary greatly in their relative size. In the next section, we attempt to explain these different country outcomes in an oligopoly model with multi-product firms, while also developing a structural test for the pro-competitive effect of city size on markups.

2 Firm Pricing and Choice of Variety

2.1 Nested CES Demand

We now turn to a model of large multi-product firms that can be used to interpret the negative relationship between prices and population for China found in the previous section. We study an economy consisting of $c = 1, \dots, D$ cities or destinations, which differ in population L_c and labor income w_c . Labor is the only factor of production and it is not mobile across cities. We denote spending on the differentiated product by $Y_c = \rho w_c L_c$.²⁵ The preferences of the representative consumer in each city are nested CES. Denoting the set of product varieties sold by firm f in city c by $i \in I_{fc}$, the sub-utility from the products of firm f are given by,

$$X_{fc} = \left(\sum_{i \in I_{fc}} (b_{ifc} x_{ifc})^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \quad \sigma > 1, \quad (3)$$

where σ is the elasticity of substitution across products sold by a firm, and b_{ifc} are the taste parameters for each variety, which we will allow to differ across firms and cities in the general model, but sometimes restrict them to be the same across barcode items i as explained below. Each firm f is a manufacturing firm and not a retailer, which we introduce as another CES nest in section 4.2.

Aggregating across firms $f \in F_c$ that sell in city c , utility of the representative consumer is,

²⁵In our empirical work we let Y_c denote city GDP and w_c denote GDP per-capita, whereas in our model we let Y_c denote spending on the differentiated good in question while w_c denotes the wage.

$$U_c = \left(\sum_{f \in F_c} X_{fc}^{(\eta-1)/\eta} \right)^{\eta/(\eta-1)}, \quad \sigma \geq \eta > 1. \quad (4)$$

As in Hottman, Redding, and Weinstein (2016), we expect that the elasticity of substitution σ across products within a firm is larger than the elasticity η across firms, so we assume that $\sigma \geq \eta > 1$. When the two elasticities are equal, then the nested CES system will collapse to a standard CES utility function.

2.2 Optimal Prices for the Multi-product Firm

Consider a firm producing variety i in city c and delivering it to destination d . The firm chooses the range of products to sell in multiple destinations $d = 1, \dots, D$. The profit-maximization problem for this firm is

$$\max_{p_{ifd}, i \in I_{fd}} \sum_{d=1}^D \left\{ \left[\sum_{i \in I_{fd}} (p_{ifd} - g_{if}(w_c) - T_{cd}) x_{ifd} - k_{ifd} \right] - K_{fd} \right\} \mathbf{1}(I_{fd} \neq \emptyset), \quad (5)$$

where K_{fd} denotes the fixed costs to enter a city and $\mathbf{1}(I_{fd} \neq \emptyset)$ is an indicator variable that takes value of unity if $I_{fd} \neq \emptyset$ and zero otherwise. That is, the firm chooses whether to enter into each city d , and pays the fixed cost of K_{fd} only if it does, along with the fixed costs of k_{ifd} to sell each variety i in city d . In an intertemporal model we could imagine that the fixed costs to enter a city K_{fd} are paid only once, and so we shall refer to these as *sunk costs*, whereas the fixed production costs k_{ifd} are paid every period to produce each variety and lead to increasing returns to scale in production. We let $g_{if}(w_c)$ denote the (constant) marginal costs of producing good i in city c , with factor prices w_c , and selling it in city d with transport costs of T_{cd} . We assume that firms treat the prices of other firms as given under Bertrand competition, and that demand in the various cities is independent.

Focus initially on the choice of optimal prices. If firms sold only a single product i in destination d , so that $s_{ifd} = 1$, then the (positive) elasticity of demand for the CES utility function in (4) is $\eta - (\eta - 1)S_{fd}$, so that the elasticity minus one equals $(\eta - 1)(1 - S_{fd})$. It follows from the usual markup formula that the optimal price is,

$$p_{ifd} = \left[1 + \frac{1}{(\eta - 1)(1 - S_{fd})} \right] [g_{if}(w_c) + T_{cd}]. \quad (6)$$

When the firm sells multiple products, however, then it must take into account how a reduction

in the price of one will decrease demand for its other products: this is the cannibalization effect. Nevertheless, we show in online Appendix C1 that the same pricing formula as in (6) is obtained.²⁶ In other words, when the firm jointly maximizes over all its prices then the markup depends only on the *total* market share S_{fd} of the firm in that city, and this markup is applied to all the products sold by the firm regardless of whether they have different taste parameters b_{ifc} or marginal costs.

2.3 Optimal Variety for the Multi-product Firm

For simplicity, suppose that in each city there is symmetric demand, and also symmetric marginal costs within each firm,²⁷ so that the firm sells N_{fd} varieties in city d . We allow marginal and fixed costs to differ across firms. Then the profit maximization problem (5) is simplified as:

$$\max_{p_{ifd}, N_{fd} \geq 0} \sum_{d=1}^D \{N_{fd} [(p_{fd} - g_f - T_{cd})x_{fd} - k_{fd}] - K_{fd}\} \mathbf{1}(N_{fd} > 0). \quad (7)$$

where we omit the subscript i from price and quantity, p_{fd} and x_{fd} , due to symmetry. The optimal price is still given by (6), though now with symmetry this price is the same across varieties i for each firm in city d . As the firm expands the number of varieties sold, it draws demand away from existing varieties. Taking this cannibalization effect into account, it is shown in online Appendix C2 that the optimal variety is determined by,

$$N_{fd} = \frac{\eta - 1}{\sigma - 1} \left[\frac{S_{fd}(1 - S_{fd})}{\eta - (\eta - 1)S_{fd}} \right] \frac{Y_d}{k_{fd}}, \quad (8)$$

where $Y_d = \rho w_d L_d$ is the total expenditure on the differentiated good in destination d . Substituting this equation into (7) and using the optimal price from (6), it is also shown that the profits from entering a city, *before* deducting the sunk entry costs K_{fd} , are:

$$\pi_{fd} = \left[\frac{(\sigma - \eta) + (\eta - 1)S_{fd}}{\sigma - 1} \right] \left[\frac{S_{fd}}{\eta - (\eta - 1)S_{fd}} \right] Y_d. \quad (9)$$

²⁶This result is also shown by Feenstra (2016), chapter 9, and Hottman, Redding, and Weinstein (2016).

²⁷In other words, we are assuming that $b_{ifd} = b_{jfd}$ and that $g_{fc} = g_{if}(w_c)$ does not depend on i . In online Appendix C3, we generalize the analysis to allow the rising marginal cost of products that are farther from the core-competency of the firm. If we restrict the analysis to iceberg rather than specific trade costs, then we find that the equilibrium condition for the scope of a firm is essentially the same as that shown by (8), but with an extra constant term.

2.4 Firm Entry

In order for the firm to serve city d , we must have $\pi_{fd} \geq K_{fd}$. There are a wide range of possible ways that entry of firms can respond to a rise in city spending Y_d , depending on the specification of the sunk costs of entering a city, K_{fd} . To illustrate these outcomes, we continue to suppose that in each city there is symmetric demand and now we also assume that there is complete symmetry across firms in their marginal costs, fixed cost k_d per variety, and in the sunk costs to enter each city, K_d . It follows that profits π_d appearing in (9) are equal across firms in a city, and they depend on the firm market shares, which equal the inverse of the number of firms as denoted by $S_c = 1/M_c$. So we can rewrite (9) in the form $\pi_d = f(M_d)Y_d$, where $f(M_d)$ is the two bracketed terms appearing on the right of (9); it is readily checked that this function is monotonically decreasing in the number of firms M_d . For generality, we also allow the sunk costs of entering a city to depend on the components of city expenditure $Y_d = \rho w_d L_d$, which we specify as $K_d = K_0 w_d^{\delta_1} L_d^{\delta_2}$. Firm entry occurs until the profits π_d just cover the sunk costs, so the zero-profit condition is written as:

$$\pi_d = f(M_d)Y_d = K_0 w_d^{\delta_1} L_d^{\delta_2}, \quad 0 \leq \delta_1, \delta_2 \leq 1. \quad (10)$$

We consider several specific cases for the sunk costs. In the first case, assume that $\delta_1 = \delta_2 = 0$ so that the sunk costs are just $K_d = K_0$. In this case we are not even allowing the sunk costs to depend on the local wage in city d , but rather, they may reflect the time spent in headquarters to establish a new outlet (though such time is outside our model). Then we invert (10) to solve for the number of firms $M_d = f^{-1}(K_0/Y_d)$. Since f is monotonically decreasing, it follows that the number of firms is increasing in the expenditure $Y_d = \rho w_d L_d$, i.e. increasing in both the wage w_d and population L_d . Then it follows from (6) that the markup of firms is decreasing in *both* these components of expenditure and so is the price charged by firms.²⁸ These results illustrate the pro-competitive effect that we found in Table 3 for China.

As an alternative case, assume that $\delta_1 = 1$ and $\delta_2 > 0$, so that the sunk costs are $K_d = K_0 w_d L_d^{\delta_2}$. We are now supposing that sunk costs use local labor that is paid the prevailing wage, and that sunk costs are increasing in the size of the population. A more sophisticated version of this specification is used by Arkolakis (2010), who refers to it as *endogenous* sunk costs and explains it by the need for advertising to the local population. In this case, the wage cancels out in (10) and we invert it as

²⁸This result is sensitive to the specification of preferences: see footnote 23.

$M_d = f^{-1}(K_0/L_d^{1-\delta_2})$. Then the positive elasticity of firm entry M_d with respect to population will depend on $1 - \delta_2$, and likewise so will the negative elasticity of the markup from (6) with respect to population. Higher values of δ_2 diminish these elasticities, and in the limiting case where $\delta_2 = 1$ then the number of firms is uniquely determined by $M_d = f^{-1}(K_0)$, so there is *no* response of the markup or price to the population. This outcome illustrates our empirical results for the United States in Table 3, where the elasticity of prices with respect to population was either positive or negative and most often insignificant.

These results show that the response of entry and markups to market size depends on the nature of sunk costs. That is what Sutton (1991) argues, using theoretical results from a wide range of oligopoly models.²⁹ Sutton draws together these models by distinguishing *exogenous* versus *endogenous* sunk costs, and like Arkolakis (2010) some years later, he argues that endogenous sunk costs would arise from expenditures on advertising as well as research and development. Under exogenous fixed costs, Sutton predicts that larger markets will have more firms (and lower prices) as a result. Under endogenous fixed costs, however, expenditures on advertising can shift demand towards the firm in question and therefore inhibit entry even in larger markets. In these cases, the link between market size, the number of firms and their prices is diminished.

These observations suggest a reason why the pro-competitive effect in China can be stronger than in the United States: namely, a different structure of sunk costs of entry. Sutton (1991) bases his empirical work on a detailed examination of various industries in countries at a similar stage of development: the United States, United Kingdom, France, Germany, Italy, and Japan. His approach does not really admit the possibility of differing sunk costs between these similar countries, even though he notes that a given industry in one country can evolve from having exogenous sunk costs into having two groups, one with exogenous and another with endogenous sunk costs.³⁰ China has industrialized much more recently than the countries examined by Sutton, so its consumer goods industry may to some degree be in the former group, with exogenous fixed costs. For example, if the major cost to establishing more retail outlets in smaller cities of China is overcoming administrative and logistical hurdles, then that cost can be treated as exogenous from the firms viewpoint, even

²⁹See also the summary in Bagwell (2007).

³⁰Particularly pertinent to the consumer goods we examine is the discussion in Sutton (1991), p. 17, of the frozen food industry in the United States and the United Kingdom: "In each case, it is possible to pinpoint the exact phase at which firms became portioned into two discrete groups: a high advertising group selling primarily to the retail sector, and a non-advertising group selling primarily to the non-retail sector. While the frozen food industries provide an unusually clear-cut illustration of this process, the same process can be seen to operate in a wide number of instances explored in later chapters".

though over time we may well see a greater spread of manufacturers and retailers into smaller cities. In a mature market like the United States, however, it may be that such costs have been substantially paid, so that now manufacturers and retailers are competing by advertising and new product innovation, which are endogenous sunk costs. These ideas from Sutton (1991), extended to allow for different sunk costs across countries, are one explanation for why the spatial distribution of firms differs so much between China and the United States.

For the remainder of the paper we focus mostly on China, and we examine next whether the pro-competitive effect as predicted from our model (with exogenous sunk costs) is borne out in the data for Chinese cities.

3 Estimating the Pricing Equation

We will estimate (6) by moving the markup term to the left, so that the remaining marginal costs plus transport costs on the right will be estimated:

$$\begin{aligned} p_{ifd} \left[1 + \frac{1}{(\eta - 1)(1 - S_{fd})} \right]^{-1} &= g_{if}(w_c) + T_{cd} \\ &= \alpha_i + \alpha_r + \alpha_p + \alpha_{cap} + \beta \ln Dist_{cd} + \epsilon_{ifd}. \end{aligned} \quad (11)$$

The first terms on the right are indicator variables for variety i and retailer r , which together with the error term ϵ_{ifd} reflect the marginal cost of production. We include an indicator variable α_{cap} for the capital city of each province in China, and the distance $Dist_{cd}$ between the production in city c and the destination city d . The term α_p denotes the destination province fixed effects that control for the additional cost incurred that are unexplained by the transportation cost (e.g. higher operation cost due to the higher land rents when a firms sells in a bigger city). We will want to evaluate how prices charged by firms differ across cities of different sizes, and for this purpose we include the variables $\ln w_d$ and $\ln L_d$ on the right of (11), so that we actually estimate is,

$$p_{ifd} \left[1 + \frac{1}{(\eta - 1)(1 - S_{fd})} \right]^{-1} = \alpha_i + \alpha_r + \alpha_p + \alpha_{cap} + \beta_1 \ln w_d + \beta_2 \ln L_d + \beta_3 \ln Dist_{cd} + \epsilon_{ifd}. \quad (12)$$

3.1 Estimating the Pro-Competitive Effect

The presence of a pro-competitive effect is tested with the price equation shown in (12). For each of the four goods, we use both the price (P) and the price divided by the markup (P/MK) as the dependent variables. We use the scraped prices at the barcode level and rely on Nielsen (China) data to calculate or predict firm-destination shares, as in Table 2. Finally, we used company reports to identify the factory locations in China, and therefore compute the distance between the factory and destination markets.³¹

In Table 4, we present the results for the four main goods (Laundry Detergent, Personal Wash items, Shampoo and Toothpaste) controlling for retailer, barcode, and destination province fixed effects. In the first column for each good, we use price-per-item (U.S. cents/item) as the dependent variable. The coefficients of $\ln Population$ and $\ln Income$ are all significantly negative, which implies larger and richer cities will benefit from lower prices for a given product at the barcode level. These initial regressions are very similar to those for the four main goods shown in Table 3, except that we now control for province fixed effects and distance.

The second column for each good in Table 4 shows the results from using P/MK_F as the dependent variable, or the price divided by the markup for the manufacturing firm as shown in (12). Comparing the first two columns for Laundry Detergent, the coefficient on $\ln Population$ falls by about one-third, from -4.45 to -2.86, and the coefficient on $\ln Income$ falls by about one-quarter, from -2.06 to -1.61. Similar reductions in each of these coefficients are found for Toothpaste. For Personal Wash Items and Shampoo, the pro-competitive effect from the manufacturing sector contributes more: the coefficients on both are reduced by roughly one-half as seen by comparing the first and second columns. For all four goods, therefore, we confirm a pro-competitive effect of city size as measured by either the $\ln Population$ and $\ln Income$. Still, we do not fully explain the negative relationship between prices and city size due to controlling for the manufacturing markup. To see whether we can further reduce the magnitude of this negative relationship, we consider next the retail sector.

³¹We have checked via internet sources that the firms in our sample have only one manufacturing location for each good.

Table 4: Price and Price/Markup Regressions for China ($\gamma_d = 0.50$)

	Laundry Detergent				Personal Wash Items			
	Price	P/MK _F	P/MK _R	P/(MK _R MK _M)	Price	P/MK _F	P/MK _R	P/(MK _R MK _M)
<i>ln Income</i>	-2.06* (1.15)	-1.61* (0.85)	-1.46** (0.72)	-1.14** (0.53)	-4.09 (2.81)	-2.53 (1.74)	-1.83 (1.79)	-1.12 (1.13)
<i>ln Population</i>	-4.45** (2.12)	-2.86* (1.48)	-2.78** (1.26)	-1.84** (0.89)	-2.41*** (0.78)	-1.09*** (0.35)	-2.05*** (0.64)	-1.07*** (0.34)
<i>Capital City</i>	-1.94 (2.21)	-1.37 (1.64)	-1.20 (1.20)	-0.85 (0.89)	-2.64 (2.01)	-1.40 (1.11)	-1.94 (1.38)	-1.09 (0.80)
<i>ln Distance</i>	0.57 (0.62)	-0.43 (0.67)	0.28 (0.36)	-0.25 (0.38)	0.68 (0.89)	0.44 (0.59)	0.33 (0.49)	0.21 (0.32)
Observations	103,227	103,227	103,227	103,227	151,826	151,826	151,826	151,826
R-squared	0.94	0.93	0.93	0.933	0.92	0.92	0.92	0.92
	Shampoo				Toothpaste			
	Price	P/MK _F	P/MK _R	P/(MK _R MK _M)	Price	P/MK _F	P/MK _R	P/(MK _R MK _M)
<i>ln Income</i>	-3.95** (1.76)	-2.11** (0.90)	-1.91* (1.12)	-1.02* (0.58)	-4.43** (1.73)	-3.70** (1.42)	-2.77** (1.10)	-2.28*** (0.78)
<i>ln Population</i>	-8.47*** (0.64)	-4.22*** (0.31)	-5.65*** (0.53)	-2.86*** (0.31)	-3.16*** (1.18)	-2.15** (1.06)	-1.00 (1.30)	-0.63 (0.83)
<i>Capital City</i>	-2.82 (2.14)	-1.47 (1.14)	-1.93 (1.25)	-1.01 (0.66)	-0.61 (0.84)	0.07 (0.63)	-0.71 (0.85)	-0.24 (0.61)
<i>ln Distance</i>	1.50 (0.90)	0.98* (0.52)	0.76 (0.51)	0.49* (0.29)	1.18*** (0.36)	1.23** (0.48)	0.64*** (0.19)	0.66** (0.25)
Observations	136,091	136,091	136,091	136,091	86,619	86,619	86,619	86,619
R-squared	0.92	0.92	0.92	0.92	0.94	0.94	0.94	0.94

Notes: Each column is a regression like that shown in (12), with the dependent variable shown at the top of the column: either Price, or Price divided by one or both of the markups denoted by MK_M and MK_R , for manufacturer and retailer, respectively. The dependent variables *ln Income* measured as GDP per capita in each city, *ln Population*, *ln Distance* from the factory to each city, and barcode, retailer and province fixed effects. Price and P/MK are in units of U.S. cents. Robust standard errors are clustered at the barcode, retailer and city level, and are reported in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.2 Extension of the Model to Include Retailers

We extend our model by introducing retailers to the economy, so that we are able to account for a pro-competitive effect in the retailing sector. The preferences of the representative consumer in each city are a three-level nested CES. Denoting the set of product varieties sold by retailer r and produced by firm f in city c by $i \in I_{frc}$, the sub-utility from purchasing the products of firm f in retailer r are,

$$X_{frc} = \left(\sum_{i \in I_{frc}} (b_{ifrc} x_{ifrc})^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \quad \sigma > 1, \quad (13)$$

where σ is the elasticity of substitution across products sold by a firm, and b_{ifrc} are the taste parameters for each variety. Aggregating across firms $f \in F_{rc}$ that sell in retailer r and city c , the sub-utility of the representative consumer purchasing in retailer r is,

$$\mathcal{X}_{rc} = \left(\sum_{f \in F_{rc}} X_{frc}^{(\eta-1)/\eta} \right)^{\eta/(\eta-1)}, \quad \sigma \geq \eta > 1. \quad (14)$$

where η is the elasticity of substitution across firms within retailers. Finally, aggregating across retailers $r \in R_c$ that sell in city c , the utility function can be written as

$$U_c = \left(\sum_{r \in R_c} \mathcal{X}_{rc}^{(\theta-1)/\theta} \right)^{\theta/(\theta-1)}, \quad \eta \geq \theta > 1. \quad (15)$$

where θ is the elasticity of substitution across retailers.

There are two kinds of firms: manufacturers and retailers. Each manufacturer produces multiple products and sells to all the retailers in each city. Both retailers and manufacturers are assumed to play their optimal strategies simultaneously in the subsequent analysis. Manufacturers are not able to sell their products directly to the consumers, and they must sell at the wholesale prices q_{ifrd} to the local retailers. We assume that the local retailers are homogeneous (i.e. symmetric) and maximize their profit by choosing the retail prices, taking the city demand and wholesale prices as given. Consider a retailer r selling product i (produced by firm f) in city d , for price p_{ifrd} . The retailer chooses the prices for all the products provided by each manufacturer, and whether or not to enter each city, with the profit-maximization problem:

$$\max_{p_{ifrd}, \forall i \in I_{f rd} f \in F_{rd}} \left[\sum_{i \in I_{f rd} f \in F_{rd}} (p_{ifrd} - q_{ifrd}) x_{ifrd} - K_d^R \right] \mathbf{1}(r \in R_d) \quad (16)$$

where K_d^R denotes the retailing fixed cost to enter in city c and $\mathbf{1}(r \in R_d)$ is an indicator variable that takes value of unity if retailer r decides to pay the fixed cost and enter city d , and zero otherwise.

Solving the above problem gives the optimal markup of the retailers,

$$p_{ifrd} = \mu_{rc} q_{ifrd}, \quad \mu_{rd} \equiv \left[1 + \frac{1}{(\theta-1)(1-S_{rd})} \right], \quad (17)$$

where μ_{rd} is the markup and S_{rd} is the market share of retailer r in city d . In practice, we have the market shares of retailers from the Nielsen (China) data, extended from 22 to 60 cities using the regressions shown in Table 2. But Nielsen data does not cover the universe of retailers, and so we need to multiply S_{rd} by a parameter γ_d , reflecting the share of retailers covered by the Nielsen

(China) data in city d . In our benchmark results reported in this section we use $\gamma_d = 0.5$, but in online Appendix D2 we experiment with alternative values as low as $\gamma_d = 0.1$, with results as discussed below.

Consider a manufacturing firm producing variety i in city c , and delivering and selling it to *all* the retailers r in destinations $d = 1, \dots, D$. We assume that the firm takes the retailer markup μ_{rd} as given from (17), and then the firm chooses the range of products to sell to all the retailers in each city where it enters, as well as the wholesale prices q_{ifrd} :

$$\max_{q_{ifrd}, i \in I_{frd}} \sum_{d=1}^D \left\{ \sum_{r=1}^{R_d} \sum_{i \in I_{frd}} \left[(q_{ifrd} - g_{if}(w_c) - T_{cd})x_{ifrd} - k_{ifd}^F \right] - K_{fd}^F \right\} \mathbf{1}(I_{fd} \neq \emptyset), \quad (18)$$

where R_d stands for the number of retailers in city d , K_{fd}^F denotes the fixed cost for the manufacturing firm to enter a city, and k_{ifd}^F is the fixed cost to sell a variety.

The final price paid by consumers combines the markups shown in (6) and (17), and is given by,³²

$$p_{ifrd} = \underbrace{\left[1 + \frac{1}{(\theta - 1)(1 - S_{rd})} \right]}_{\text{Retailer Markup}} \underbrace{\left[1 + \frac{1}{(\eta - 1)(1 - \gamma_d S_{fd})} \right]}_{\text{Manufacturer Markup}} \underbrace{[g_{if}(w_c) + T_{cd}]}_{\text{Marginal Cost}}. \quad (19)$$

The first term in (19) denotes the retailer markup, and the second is the manufacturer markup. Larger and richer cities could have both greater number of retailers (lower retailer markups) and more entry of producers (lower manufacturer markups), which will decrease the price in either way.

3.3 Pro-Competitive Effect for Retailers

We re-estimate the price equations allowing for the pro-competitive effect of manufacturing firms and retailers, with the results shown in Table 4. The third column for each good shows the results from using P/MK_R as the dependent variable, or the price divided by the markup for the retailer only, and the fourth column shows the results from using using $P/MK_F MK_R$ as the dependent variable, or the price divided by the markup for both the manufacturing firm and the retailer, as shown in (19). Comparing the first and third columns for Laundry Detergent, the coefficient on $\ln Population$ falls by more than one-third, from -4.45 to -2.78, when we switch from using the price as the dependent variable to using the price divided by retailer markup. That reduction is greater than when we used the price divided by producer markup, as shown in the second column, implying a stronger

³²This solution for the optimal consumer price is derived in online Appendix D1.

pro-competitive effect in the retailing sector. Combining *both* the firm and retailer markups gives us the fourth column, where the coefficient on $\ln Population$ is only -1.84, or less than one-half of the original coefficient of -4.45. The negative coefficient on $\ln Income$ is also reduced when controlling for markups, by roughly 20 percent when using the firm markups and 30 when using the retailer markups, so by nearly one-half in total.

Similar results are found for the other three goods. In all cases, controlling for firm markups and retailer markups (in the fourth column) reduced the coefficients on $\ln Population$ and $\ln Income$ by at least one-half as compared to their original values in the price regression (the first column), and in some cases the coefficients are reduced by considerably more – to one-third or one-quarter of their original size. Controlling for the retailer markup consistently has stronger effect on reducing the $\ln Income$ coefficient than controlling for the manufacturing firm markup. In contrast, for reducing the $\ln Population$ coefficient, controlling for the retailer markup has a greater impact in two goods and a lesser impact in the other two goods. But the finding that the negative coefficients on *both* $\ln Income$ and $\ln Population$ are reduced in all cases indicates that *total* city GDP is the appropriate measure of city size to evaluate the pro-competitive effect.

Still, even with the addition of the retail sector, we do not *fully* explain the impact of city size on prices as due to a pro-competitive effect working through firm and retailer markups: there is some additional negative effect of city size that remains. There are several explanation for this unexplained effect. Empirically, we have already noted that our Nielsen (China) and *Wochacha* data do not cover the universe of retailers, so that the market shares that we are using (which are extended from 22 cities to 60 cities using the predicted sales from the regressions in Table 2) may not be accurate. When we experiment with alternative values for the share of retailers covered by our data as low as $\gamma_d = 0.1$ (see online Appendix D2), the precisely measured coefficients on $\ln Income$ and $\ln Population$ become slightly smaller (in absolute value) than in Table 4 when controlling for the retailer markups, so that we are explaining somewhat more of the pro-competitive effect. Another source of empirical error comes from accepting the administrative definition of city boundaries. In contrast, Dingel, Miscio, and Davis (2018) have recently used nighttime city lights data to construct contiguous city areas for several countries, and they find that the large cities in China should be even larger. Presumably that alternative definition of city boundaries would encompass more entry in the largest cities, and so it would have a stronger pro-competitive effect. This idea is consistent with recent theory in Kokovin, Parenti, Thisse, and Ushchev (2017), whose model predicts the entry of a competitive fringe

of producers in large cities that leads to lower prices. Summing up, a limitation of both our theory and empirical work is that we have not incorporated small producers or retailers, which can lead to lower prices and markups in large cities due to a competitive fringe, and to higher prices in small cities due to traditional but inefficient retailers.

4 Variety of Products in Chinese Cities

4.1 Estimating the Variety Effect

Our model implies that not only prices vary across cities, but also the availability of varieties, so that consumer welfare is affected for both reasons. To estimate the impact of city size on product variety, we return to our baseline model as presented in section 2. Let $I_f \equiv \cap_c I_{fc}$ denote the “common” varieties sold by firm f in all cities within each country, which we assume is non-empty, and let $N_{f,com}$ denote the number of products in this set. When solving for the range of optimal range of product varieties in section 2.3, we assumed symmetric demand and symmetric marginal cost within each firm. Then it follows that the common-goods share of firm f in destination d is:

$$\lambda_{fd} = \frac{N_{f,com} p_{fd} x_{fd}}{N_{fd} p_{fd} x_{fd}} = \frac{N_{f,com}}{N_{fd}} = N_{f,com} \left(\frac{\sigma - 1}{\eta - 1} \right) \left[\frac{\eta - (\eta - 1) S_{fd}}{S_{fd}(1 - S_{fd})} \right] \frac{k_{fd}}{Y_d}, \quad (20)$$

where the final equality follows using (8). Holding fixed the firm share S_{fd} , we see that higher city total expenditure Y_d on the differentiated good leads to reduced common-goods share λ_{fd} for each firm, and hence greater product variety.

We can estimate this equation as a regression by substituting $Y_d = \rho w_d L_d$ and taking logs,

$$\ln \lambda_{fd} = \alpha_{f,prov} + \beta_1 \ln w_d + \beta_2 \ln L_d + \beta_3 \ln \left[\frac{\eta - (\eta - 1) S_{fd}}{S_{fd}(1 - S_{fd})} \right] + \epsilon_{fd}. \quad (21)$$

The first term $\alpha_{f,prov}$ on the right is an indicator variable for the firm-province, which together with the error term ϵ_{fd} reflects variation in $N_{f,com}$ and in the fixed costs k_{fd} of providing each product in city d . The presence of w_d and L_d in (21) reflects the fact that higher expenditure in larger cities *reduces* the common-goods share of firms in that city, λ_{fd} , thereby increasing product variety from each firm. In theory, the coefficients on these variables are both $\beta_1 = \beta_2 = -1$. A negative sign on these estimated coefficients will show how larger cities (measured by population or income per-capita) have a lower common-goods share for firms, and therefore more product variety. The common-

goods share of firms in each city, λ_{fd} , is *inversely* related to their product scope, so the final term before the error on the right of (21) is a U-shaped function of the firm's market share, and should have a coefficient of $\beta_3 = 1$ in theory.

Table 5 exhibits the result of estimating (21). The columns report results for Laundry Detergent, Personal Wash items, Shampoo and Toothpaste, respectively. For the first three products, the regressions results show that larger cities (measured by income per-capita or population) will have smaller expenditure shares on common products for a given firm, though not the expected coefficient of -1 . This result implies that larger cities get access to more product varieties than smaller cities. For Toothpaste, the coefficient on $\ln Population$ is negative as expected, but the coefficient on $\ln Income$ is positive. Consistent with the model, the variable $\ln \frac{\eta - (\eta - 1)S_{fd}}{S_{fd}(1 - S_{fd})}$ contributes to the common product share positively, and is close to its theoretical value of unity.

Table 5: Firm Share Regression for China

	Laundry Detergent	Personal Wash	Shampoo	Toothpaste
$\ln Income$	-0.285*** (0.038)	-0.065*** (0.018)	-0.051*** (0.010)	0.785*** (0.093)
$\ln Population$	-0.255*** (0.057)	-0.116** (0.042)	-0.109*** (0.016)	-0.155** (0.059)
<i>Capital City</i>	0.011 (0.035)	-0.204*** (0.022)	-0.068*** (0.013)	-0.706*** (0.060)
$\ln \frac{\eta - (\eta - 1)S_{fd}}{S_{fd}(1 - S_{fd})}$	0.769*** (0.138)	0.901*** (0.130)	0.519*** (0.086)	0.916*** (0.062)
Observations	420	420	840	660
R-squared	0.820	0.925	0.952	0.888

Notes: Each column is a regression with the dependent variable $\ln \lambda_{fd}$. The dependent variables are $\ln Income$, measured as GDP per capita in each city, $\ln Population$, a dummy variable for *Capital City*, and firm and destination province fixed effects. Robust standard errors are clustered at firm level and reported in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

These empirical results support the inverted U-shaped relation between product scope and market share. Feenstra and Ma (2009) showed that in a model with heterogeneous, multi-product firms, an inverted U-shaped relation would hold between the productivity of the firm and its product scope: more productive firms initially add more products, but then reduce product scope as cannibalization

becomes more important. Raff and Wagner (2013) have shown that such an inverted U-shaped relation holds empirically in a sample of German firms, and Macedoni (2017) obtains this relation in more general theoretical and empirical settings. Our results here are consistent with those authors.

4.2 Variety-Adjusted Price Indexes

We now turn to the general measurement of product variety in the case of non-symmetric demand and prices. For simplicity, we return to the two-level nested CES preferences without distinguishing retailers. Let $\mathbf{p}_{fc}$ denote the vector of prices p_{ifc} for firm f across all product varieties i , with the vector of taste parameters $\mathbf{b}_{fc}$, and let $P_{fc} = e(\mathbf{p}_{fc}, \mathbf{b}_{fc}, I_{fc})$ denote the minimum expenditure needed to obtain one unit of sub-utility, $X_{fc} = 1$. Then $e(\mathbf{p}_{fc}, \mathbf{b}_{fc}, I_{fc})$ takes on the CES form,

$$P_{fc} = e(\mathbf{p}_{fc}, \mathbf{b}_{fc}, I_{fc}) = \left(\sum_{i \in I_{fc}} (p_{ifc}/b_{ifc})^{1-\sigma} \right)^{1/(1-\sigma)}. \quad (22)$$

Feenstra (1994) shows how to measure the effect of new varieties on the expenditure needed to obtain a fixed level of utility, which gives us a variety-adjusted – or exact – price index. Specifically, consider firm f selling to destinations c and d . Suppose that there is a non-empty subset of common products $I_f \subseteq I_{fc} \cap I_{fd}$ sold by firm f in these two cities for which the taste parameters are equal, $b_{ifc} = b_{ifd}$, $i \in I_f$. Then the exact price index between the two cities can be expressed as,

$$\frac{e(\mathbf{p}_{fc}, \mathbf{b}_{fc}, I_{fc})}{e(\mathbf{p}_{fd}, \mathbf{b}_{fd}, I_{fd})} = \left[\prod_{i \in I_f} \left(\frac{p_{ifc}}{p_{ifd}} \right)^{w_{if}(I_f)} \right] \left(\frac{\lambda_{fc}}{\lambda_{fd}} \right)^{\frac{1}{\sigma-1}}. \quad (23)$$

The first term in brackets on the right of (23) is the Sato (1976)-Vartia (1976) price index, where $w_{if}(I_f)$ is the weight defined by:

$$w_{if}(I_f) \equiv \frac{\frac{s_{ifc}(I_f) - s_{ifd}(I_f)}{\ln s_{ifc}(I_f) - \ln s_{ifd}(I_f)}}{\sum_{j \in I_f} \left(\frac{s_{jfc}(I_f) - s_{jfd}(I_f)}{\ln s_{jfc}(I_f) - \ln s_{jfd}(I_f)} \right)}, \quad s_{ifc}(I_f) \equiv \frac{p_{ifc} x_{ifc}}{\sum_{j \in I_f} p_{jfc} x_{jfc}}, \quad (24)$$

and likewise for the shares $s_{ifd}(I_f)$ in city d , also defined over the common set of products. The second term on the right of (23) is the adjustment needed to take into account differing sets of goods available in the two cities, and is defined by:

$$\lambda_{fc} \equiv \frac{\sum_{i \in I_f} p_{ifc} x_{ifc}}{\sum_{i \in I_{fc}} p_{ifc} x_{ifc}} = 1 - \frac{\sum_{i \in I_{fc} \setminus I_f} p_{ifc} x_{ifc}}{\sum_{i \in I_{fc}} p_{ifc} x_{ifc}}, \quad (25)$$

To interpret these formulas, λ_{fc} in (25) denotes the spending in city c on the common products of firm f , sold in both cities, relative to total spending in city c on firm f 's products. Equivalently, it equals one minus the share of expenditure on the unique products sold by firm f only in city c . Having access to more unique varieties in city c implies a smaller expenditure share on common products, λ_{fc} , and a lower exact index in (23). So λ_{fc} is an inverse measure of product variety that indicates a welfare gain.

To extend the exact price index to the nested CES case, let $\mathbf{P}_c$ denote the vector of CES price indexes P_{fc} shown in (22) for all firms $f \in F_c$ in city c , and let $P_c = E(\mathbf{P}_c, F_c) = \left(\sum_{f \in F_c} P_{fc}^{(\eta-1)} \right)^{1/(\eta-1)}$ denote the expenditure needed to obtain utility of one in city c . Let $F \equiv F_c \cap F_d$ denote the non-empty set of common firms selling to both cities c and d . Then using the above results and again from Feenstra (1994), the exact price index between the two cities can be written as,

$$\begin{aligned} \frac{E(\mathbf{P}_c, F_c)}{E(\mathbf{P}_d, F_d)} &= \left[\prod_{f \in F} \left(\frac{P_{fc}}{P_{fd}} \right)^{W_f(F)} \right] \left(\frac{\lambda_c}{\lambda_d} \right)^{\frac{1}{\eta-1}} \\ &= \left[\prod_{f \in F} \prod_{i \in I_f} \left(\frac{p_{ifc}}{p_{ifd}} \right)^{W_f(F)w_{if}(I_f)} \right] \left[\prod_{f \in F} \left(\frac{\lambda_{fc}}{\lambda_{fd}} \right)^{W_f(F)} \right]^{\frac{1}{\sigma-1}} \left(\frac{\lambda_c}{\lambda_d} \right)^{\frac{1}{\eta-1}} \end{aligned} \quad (26)$$

where the Sato-Vartia weights across firms are,

$$W_f(F) \equiv \frac{\frac{S_{fc}(F) - S_{fd}(F)}{\ln S_{fc}(F) - \ln S_{fd}(F)}}{\sum_{g \in F} \left(\frac{S_{gc}(F) - S_{gd}(F)}{\ln S_{gc}(F) - \ln S_{gd}(F)} \right)}, \quad S_{fc}(F) \equiv \frac{\sum_{i \in I_{fc}} p_{ifc} x_{ifc}}{\sum_{g \in F} \sum_{i \in I_{gc}} p_{igc} x_{igc}}, \quad (27)$$

and likewise for the shares $S_{fd}(F)$ for city d defined over the common set of firms. The final term on the right of (26) is defined by:

$$\lambda_c \equiv \frac{\sum_{g \in F} \sum_{i \in I_{gc}} p_{igc} x_{igc}}{\sum_{g \in F_c} \sum_{i \in I_{gc}} p_{igc} x_{igc}} = 1 - \frac{\sum_{g \in F_c \setminus F} \sum_{i \in I_{gc}} p_{igc} x_{igc}}{\sum_{g \in F_c} \sum_{i \in I_{gc}} p_{igc} x_{igc}}, \quad (28)$$

That is, λ_c denotes the spending on the common set of firms F relative to total spending in city c , or one minus the share of spending on firms selling only in city c . The greater the share of spending on unique firms selling only in that city, the lower is the λ_c and the exact price index in (26).

In online Appendix F, we provide graphs of the various components of the exact price index against city population in China and the United States. Here, we report simple regression results of

the components of the exact price index against city GDP and city population in China. Specifically, we calculate first the geometric mean of prices in each city, denoted by P_c^G . We then calculate the Sato-Vartia weighted average of prices in each city, which is the numerator of the Sato-Vartia index in (26), and also the numerator of the variety term in (26):

$$P_c^{SV}(F) \equiv \prod_{f \in F} \prod_{i \in I_f} (p_{ifc})^{W_f(F)w_i(I_f)}, \quad \Lambda_c(F) \equiv \left[\prod_{f \in F} \lambda_{fc}^{W_f(F)} \right]^{\frac{1}{\sigma-1}} \left(\lambda_c \right)^{\frac{1}{\eta-1}}. \quad (29)$$

Multiplying the above two terms we obtain $E_c^F(\mathbf{P}_c, F) \equiv P_c^{SV}(F)\Lambda_c(F)$, which we interpret as the Feenstra exact price index for cities in China. In Table 6 we report the coefficients of regressions of the geometric mean of prices and the exact price index and its components on the log of city GDP or city population. Specifically, we run the regressions:

$$\ln Z_c = \alpha_0 + \alpha_{cap} + \beta_z \ln \text{Size}_c + \epsilon_d, \quad \text{Size}_c = \text{GDP}_c \text{ or } \text{Pop}_c, \quad (30)$$

where Z_c denotes the geometric price index or either of the two components in (29), or their product $E_c^F(\mathbf{P}_c, F)$, α_{cap} is a dummy variable if city c is a provincial capital, and β_z is the estimated elasticity of the index or component Z_c with respect to city size (i.e. GDP or population).

The elasticity of the geometric mean of prices with respect to city size are shown in the first row of Table 6. Remarkably, these elasticities are nearly the same whether we use city GDP or city population as the measure of size, and they range from -0.006 to -0.011 for Laundry Detergent, Personal Wash items and Shampoo. If we multiple those coefficients by the natural log of the ratio of highest to lowest GDP among our 60 Chinese cities (which is 3.5), then we obtain a reduction in prices from 2.1 to 3.9 percent. For Toothpaste the elasticity with respect to GDP is larger at -0.024 , and so the implied reduction in price is 8.4 percent from the smallest to the largest city. These estimates are quite close to the calculation we made in section 2.2 using linear price regressions, where we argued that there was a price reduction of about 3 percent for the first three goods and 6 percent for Toothpaste, from the smallest to the largest cities.

Interestingly, the median elasticity of -0.01 for city size on the geometric mean of prices in China

Table 6: Elasticities of Price Index and its Components with respect to City Size in China

	Laundry Detergent		Personal Wash Items		Shampoo		Toothpaste	
	$\ln GDP_c$	$\ln Pop_c$	$\ln GDP_c$	$\ln Pop_c$	$\ln GDP_c$	$\ln Pop_c$	$\ln GDP_c$	$\ln Pop_c$
$\ln P_c^G$	-0.010*	-0.010	-0.006	-0.007	-0.009*	-0.011	-0.024***	-0.024*
	(0.006)	(0.009)	(0.006)	(0.008)	(0.005)	(0.008)	(0.008)	(0.012)
$\ln P_c^{SV}$	0.002	0.008	-0.008	-0.011	-0.007	-0.007	-0.043***	-0.040*
	(0.008)	(0.012)	(0.008)	(0.011)	(0.006)	(0.009)	(0.015)	(0.024)
$\ln \Lambda_c$	-0.049**	-0.073**	-0.061**	-0.118**	-0.062***	-0.116***	0.026	-0.038
	(0.020)	(0.030)	(0.025)	(0.036)	(0.019)	(0.027)	(0.026)	(0.038)
$\ln E(P_{fc}, F_c)$	-0.047**	-0.065*	-0.069**	-0.129***	-0.070***	-0.123***	-0.017	-0.078
	(0.022)	(0.033)	(0.026)	(0.038)	(0.021)	(0.030)	(0.025)	(0.036)

Notes: Each row reports the results of eight regressions, two for each of the four consumer goods, with the dependent variables: $\ln P_c^G$, which is the simple geometric mean of unit prices; $\ln P_c^{SV}$, which is the Sato-Vartia weighted mean of unit prices; $\ln \Lambda_c$, which is the variety term in (29); and $\ln E(P_{fc}, F_c)$, which is the variety-adjusted, exact price index. The independent variables are either $\ln GDP_c$ which is city GDP or $\ln Pop_c$, whose coefficients are reported in the table, along with a *Capital* dummy variable and constant which are omitted from the table. Standard errors are reported in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

is what Handbury and Weinstein (2015) find for the elasticity of city population on the *variety-adjusted* price index for the United States.³³ In contrast, we find much *larger* impact of city size on the variety-adjusted price index, as shown in the last row of Table 6. When using city GDP as the measure of size, the elasticity on the variety-adjusted price index is as high as -0.07 , or up to *seven times* more than what Handbury and Weinstein (2015) find for the United States. In comparison, when using city population as the measure of size, the elasticity on the variety-adjusted price index is at least -0.07 (rounding up the elasticity of -0.065 for Laundry Detergent to that value), or at least *seven times* more than found by Handbury and Weinstein (2015). This larger fall in prices in China is due primarily to the greater impact of city size on increasing the variety of products, as shown by the elasticities of the (inverse) variety term Λ_c in the third row of Table 6, which nearly equal the elasticity of the exact price index in the final row for Laundry Detergent, Personal Wash items and Shampoo. For Toothpaste, however, the elasticity of product variety with respect to either measure of city size is insignificant, and the dominant impact of city size is on reducing the geometric mean of prices (first row) or the Sato-Vartia weighted mean (second row). So Toothpaste is a counterexample to our finding that most of the reduction in the exact price index in larger cities is due to product variety.³⁴

³³Their preferred estimate for the elasticity of MSA population on the U.S. variety-adjusted price index is -0.011 .

³⁴In online Appendix G, Table A.13, we examine results similar to Table 6 for the 15 other goods in our sample, where the

5 Comparing the Variety-adjusted Price Indexes Between Countries

So far in the paper we have compared the impact of city size on prices and variety in China and the United States, but we have not been able to *directly* compare prices across the two countries because the barcode systems differ.³⁵ In that case, it is impossible to construct the ratio of prices (p_{ifc}/p_{ifd}) that appears in (23) and (26) because the index i is not the same if one city is in China and the other is in the U.S. That difficulty is overcome in Redding and Weinstein (2018) by relying on a procedure to *average* across barcode items (even if they are measured with different systems) and by making an assumption on tastes. We now show how this technique which Redding and Weinstein (2018) propose in a time-series context can be adapted to our cross-country context, to directly compare the variety-adjusted price indexes between the countries.

To briefly review their results, we start with the demand for each product variety which equals $b_{ifc}x_{ifc} = [(p_{ifc}/b_{ifc})/P_{fc}]^{-\sigma}X_{fc}$. Multiplying by (p_{ifc}/b_{ifc}) and dividing by $P_{fc}X_{fc}$, we obtain an equation for the share of each variety within the total sales of firm f , which depends on the CES price index P_{fc} . Inverting that equation to solve for P_{fc} , we readily obtain:

$$P_{fc} = e(\mathbf{p}_{fc}, \mathbf{b}_{fc}, I_{fc}) = s_{ifc}^{1/(\sigma-1)} \left(\frac{p_{ifc}}{b_{ifc}} \right). \quad (31)$$

The term λ_{fc} defined in (25) equals $s_{ifc}/s_{ifc}(I_f)$, so we can replace the share s_{ifc} above by $s_{ifc} = s_{ifc}(I_f)\lambda_{fc}$. Then because (31) holds for *every* product i , we take the unweighted geometric mean across all N_f products in the common set to obtain the formula in Redding and Weinstein (2018) :

$$P_{fc} = \left[\prod_{i \in I_f} s_{ifc}(I_f)^{\frac{1}{N_f(\sigma-1)}} \right] \left[\prod_{i \in I_f} \left(\frac{p_{ifc}}{b_{ifc}} \right)^{\frac{1}{N_f}} \right] \lambda_{fc}^{1/(\sigma-1)}. \quad (32)$$

Notice that this formula gives us the *level* of the CES price index P_{fc} , and not just its ratio as shown in section 5 by equation (23).

We can aggregate over firms in a city using a similar approach. The aggregate demand for each firm's products are $X_{fc} = (P_{fc}/P_c)^{-\eta}(Y_c/P_c)$. Multiplying by P_{fc} and dividing by Y_c , we obtain the share of each firm in city c , which depends on the CES price index P_c . Inverting that equation to solve for P_c , we readily obtain:

expenditure share on each barcode items are estimated from (38). We find that the median elasticity of the variety-adjusted price index with respect to city GDP is again -0.01 , and the highest elasticity of the exact price index with respect to city GDP is -0.07 , which is again seven times more than what Handbury and Weinstein (2015) find for the United States.

³⁵See footnote 4.

$$P_c = E(\mathbf{P}_c, F_c) = S_{fc}^{1/(\eta-1)} P_{fc}.$$

We again replace the share S_{fc} by $S_{fc} = S_{fc}(F)\lambda_c$, and take the unweighted geometric mean over the number of common firms M in the set F selling to *all* cities in each country. Then using that geometric mean along with (32), we obtain:

$$P_c = E(\mathbf{P}_c, F_c) = \left[\prod_{f \in F} \prod_{i \in I_f} \left(\frac{p_{ifc}}{b_{ifc}} \right)^{\frac{1}{MN_f}} S_c(F) \Lambda_c(F), \right. \quad (33)$$

where,

$$S_c(F) \equiv \left[\prod_{f \in F} S_{fc}(F)^{\frac{1}{M(\eta-1)}} \prod_{i \in I_f} s_{ifc}(I_f)^{\frac{1}{MN_f(\sigma-1)}} \right], \quad \Lambda_c(F) \equiv \left[\prod_{f \in F} \lambda_{fc}^{1/M(\sigma-1)} \right] \lambda_c^{1/(\eta-1)}. \quad (34)$$

The first expression on the right of (33) is a simple geometric mean of the prices of common products, but these prices are adjusted for the taste parameters b_{ifc} . We have earlier assumed in our discussion of (23) that the taste parameters b_{ifc} are identical for common goods that are available in every city within each country. We can now follow the approach of Redding and Weinstein (2018) and use the *weaker* assumption that the geometric mean of taste parameters are equal in every city within each country.³⁶

$$\prod_{f \in F} \prod_{i \in I_f} (b_{ifc})^{\frac{1}{MN_f}} = 1, \quad \forall c = 1, \dots, D. \quad (35)$$

With this assumption, the first term on the right of (33) is easily measured as the geometric mean of prices. The second term $S_c(F)$ is a geometric mean of the common firms and common product shares, as shown in (34). We interpret that term as adjusting for differences in tastes for specific products *within* the common set, but subject to the maintained assumption that average tastes are the same as in (35). The third term $\Lambda_c(F)$ is the correction for variety outside the common set as in Feenstra

³⁶Redding and Weinstein (2018) are dealing in a time-series rather than cross-section context, so they rely on the assumption that the geometric mean of CES taste parameters is constant over time. That assumption is obtained as a result when the taste parameters have an i.i.d. random error and the number of common goods is very large, as in equations (2) and (12) of their paper. We do not presume that differences in tastes between cities and countries are random or that the number of common goods is very large, so we impose (35) and (36) as assumptions.

(1994).

To make the comparison between cities in China and in the United States, we need an assumption similar to (35) but applied across countries. To achieve this, let us define the *global brands* by the common brands that sell in both countries, i.e. by the set $G \equiv F^{US} \cap F^{China}$, which we assume is not empty.³⁷ We then pick two large cities in each country, New York (NY) and Shanghai (SH), for the benchmark comparison. Our assumption is that there has been enough convergence of tastes and availability of global brands across these two cities so that the average taste parameters are the same,

$$\prod_{g \in G} \prod_{i \in I_g^{US}} (b_{ig, NY})^{\frac{1}{MN_g}} = \prod_{g \in G} \prod_{j \in I_g^{China}} (b_{jg, SH})^{\frac{1}{MN_g}} = 1. \quad (36)$$

Notice that this assumption can be made even when the barcode systems $i \in I_g^{US}$ and $j \in I_g^{China}$ are different, because we are taking averages over each of these sets.

With this assumption, we can compare the Feenstra price indexes $E_c^F(\mathbf{P}_c, F)$ between *countries*, as follows. We first construct the Redding-Weinstein price indexes in (33) for Shanghai in China and New York in the U.S., with the common international firms G , which are now directly comparable and denoted by $E^{RW}(\mathbf{P}_{SH}, G)$ and $E^{RW}(\mathbf{P}_{NY}, G)$. Next, we renormalize the Feenstra price indexes for the Chinese and U.S. cities/MSAs that are denoted by $E^F(\mathbf{P}_{SH}, F^{China})$ and $E^F(\mathbf{P}_{NY}, F^{US})$, so that the price indexes are the same for Shanghai and New York when using the global brands (G) or the national brands (F^{China} and F^{US}) as the common set. This is achieved by renormalizing as:

$$\begin{aligned} \tilde{E}(\mathbf{P}_c, F^{China}, G) &\equiv \frac{E^F(\mathbf{P}_c, F^{China})}{E^F(\mathbf{P}_{SH}, F^{China})} \times E^{RW}(\mathbf{P}_{SH}, G), \\ \tilde{E}(\mathbf{P}_d, F^{US}, G) &\equiv \frac{E^F(\mathbf{P}_d, F^{US})}{E^F(\mathbf{P}_{NY}, F^{US})} \times E^{RW}(\mathbf{P}_{NY}, G), \end{aligned} \quad (37)$$

where the first equation applies for cities c in China and the second for cities d in the U.S. Since this renormalization is done for all cities in both countries, we obtain variety-adjusted price indexes that are fully comparable between cities and countries.³⁸

³⁷The global brands by goods are listed in Table A.2 in the online Appendix. In practice, G is empty for Chips and we adopt the broader definition $G \equiv [F^{US} \cap (\cup F_c^{China})] \cup [F^{China} \cap (\cup F_c^{US})]$ as the set of global brands that either sell national common products in China or in the U.S.

³⁸In online Appendix H we show that the results of Table 7 are not affected too much by choosing other benchmark pairs of cities in China and the United States, besides Shanghai and New York, to compute the cross-country price indexes.

We apply the method in (37) to our four main goods and the 15 other products. Since we have only scraped price data and not the Nielsen (China) barcode expenditures for the 15 other products besides our four main goods, it is challenging to construct the share term $S_c(F)$ and the variety term $\Lambda_c(F)$ for China that appear in (33). For these other products, we create a proxy for the expenditure on each barcode product by simply counting the number of retail stores that sell each barcode in a city and dividing by the total number of retailers. Then the sales of each barcode item are approximated by:

$$\widehat{\ln Sales_{ic}} = A_c + \frac{\sum_{r \in R_c} \mathbf{1}(i \text{ is observed in retailer } r)}{N_c^{retailers}}, \quad (38)$$

where $\mathbf{1}(i \text{ is observed in retailer } r)$ equals unity if we observe price quotes of barcode product i in retailer r and zero otherwise, R_c denote the set of retailers in city c with the total number $N_c^{retailers}$, and the constant A_c cancels out when taking the exponent and constructing the market shares.³⁹

In Table 7 we summarize each component of the Chinese price index relative to the average value of its counterpart of the United States, where a value greater than one implies the corresponding variable is larger in China. Beginning with the geometric mean of prices, as shown in the first column, these are lower in China for all products except for Shampoo, Baby Formula, Cereal and Milk Powder. These higher prices in China are explained, we believe, by differing quality characteristics of these goods that arise because the barcode systems and therefore the products are not identical across countries: some Shampoo products in China include additives, for example. A different type of quality characteristic was the tainted Milk Powder that occurred in China in 2008, leading to higher prices than in the U.S. for our scraped data in 2015, and similarly for Baby Formula.

In the second column we show the ratio of the terms $[S_{SH}(G)\Lambda_{SH}(G)/S_{NY}(G)\Lambda_{NY}(G)]$ for the benchmark comparison of Shanghai relative to New York. This is the key ratio that adjusts the relative geometric mean of prices to obtain the relative exact price index in these two cities. This ratio is less than unity for our four main goods, indicating that there is greater product variety in China than in the U.S. We interpret that result as saying that the global brands G are sufficiently common in China for these four goods that the presence of *additional* local brands in China gives it greater product variety. For the 15 other goods, however, this ratio is greater than unity for all

³⁹See Antoniadis, Xu, and Feenstra (2019), where it is shown empirically that the market share is exponentially related to the proportion of available retailers carrying the barcode item. It follows that the log of sales is linearly related to this proportion. In the presence for expenditure data for a subset of the sample, then a coefficient can be estimated on the fixed effect for each retailer appearing in the right-most term of (38), as we did in equation (1) to predict the sales in the rest of the sample. We have no such expenditure data for the 15 other goods, however.

products except for Coffee, Gum and Milk Powder. Coffee and Gum are similar to our four main goods, where there are both global brands and (inexpensive) local brands in China that are sufficient popular to result in greater product variety there. Along with Soft Drinks, which also has local brands in China, these three products in addition to the four main goods have price indexes in China relative to the U.S. that are lower than the relative geometric mean of prices shown in the first column.

Table 7: Summary of Exact Price Index by Goods (China relative to the U.S.)
Using Shanghai relative to New York as the Benchmark Comparison

	P_c^G	$\frac{S_{SH}(G)\Lambda_{SH}(G)}{S_{NY}(G)\Lambda_{NY}(G)}$	Other Terms	$\frac{\tilde{E}(\mathbf{P}_c, F^{China}, G)}{\tilde{E}(\mathbf{P}_d, F^{US}, G)}$
Four Main Products				
Laundry Detergent	0.58	0.55	0.98	0.32
Personal Wash Items	0.79	0.86	0.55	0.37
Shampoo	1.60	0.97	0.62	0.98
Toothpaste	0.90	0.29	1.11	0.29
Other Products				
Baby Formula	1.30	1.12	1.02	1.48
Battery	0.24	1.01	1.27	0.31
Biscuits	0.64	2.18	3.05	4.23
Cereal	1.29	1.95	0.70	1.76
Chips	0.31	1.76	1.16	0.64
Chocolate	0.74	1.74	1.37	1.77
Coffee	0.70	0.97	1.00	0.68
Diapers	0.74	1.12	1.04	0.86
Dog Food	0.23	1.46	1.06	0.35
Gum	0.93	0.88	0.82	0.67
Milk Powder	5.29	0.74	1.60	6.20
Sanitizer	0.37	4.28	0.89	1.43
Soft Drink	0.41	1.19	0.67	0.33
Tea	0.81	2.13	1.12	1.93
Toothbrush	0.53	1.57	1.13	0.93

Notes: The table summarizes population weighted average of each variable in relative terms (greater than unity implies the corresponding variable is larger in China.). P_c^G is the geometric mean of prices per ounce for the four main products and per item for other products; $\frac{S_{SH}(G)\Lambda_{SH}(G)}{S_{NY}(G)\Lambda_{NY}(G)}$ denotes the composite share of varieties in Shanghai relative to New York; and Other Terms is the further adjustment needed to obtain $\frac{\tilde{E}(\mathbf{P}_c, F^{China}, G)}{\tilde{E}(\mathbf{P}_d, F^{US}, G)}$, which is averaged across cities in each country to obtain the exact index in China relative to the U.S. Demand elasticities are country specific for the four main products, while we use the U.S. elasticities for other products.

For the other dozen products in Table 7, we find that the variety-adjusted price index in China relative to the U.S., as shown in the final column, is *higher* than the relative geometric mean of prices. The higher price index in China for the remaining products is in most cases due to lower variety

there. In some products (Biscuits, Chocolate, Soap, and Tea), that lower variety is strong enough that it reverses the lower relative price in China, as shown by the geometric mean of prices, to become a higher exact price index in China. It is surprising that there is lower relative variety in China of a product like Tea, which China has specialized in from ancient times. What seems to explain the lower relative variety in our data is that the many loose varieties of Chinese tea are not sold with bar codes, and so we do not capture them in our scraped data. So we miss many local varieties of Chinese tea and end up with only pre-packaged varieties, with Lipton as the unique global brand that is sold in all cities in both countries. The finding of a higher relative price of tea in China illustrates a limitation of our data that likely also applies to fresh fruit, fish, meat, etc., which are goods that are sold without barcodes and so it is especially difficult to measure product variety. This is not the case for most of the other goods we study where most varieties purchased come prepackaged and carry a barcode. For goods sold without a barcode, however, it remains to be seen whether other methods of obtaining data across countries can improve the measurement of prices and especially product variety.

6 Conclusions

We have used barcode data on selected consumer goods to compare prices and the availability of goods across cities in China, and to contrast our findings with the United States. Our main finding is that the prices for consumer goods in China are lower in larger cities and that product variety is higher. The finding that prices are lower in large Chinese cities contrasts with evidence from the United States, where there is not a systematic relationship between prices and city size. We attribute lower prices for larger cities in China to a pro-competitive effect, whereby larger cities attract more brands and retailers which leads to lower markups and prices. Greater variety in larger cities applies to the United States, too, but it is more pronounced in China. We explain both these differences between countries as due to the more uneven spatial distribution of manufacturers' sales and retailers in China as compared to the United States.

We further compare the variety-adjusted price indexes for particular goods between China and the United States. The simple geometric mean of prices is lower in China than in the United States for most goods (but with several exceptions related to product quality). By focusing on "global brands" to construct these geometric means, we can plausibly meet the theoretical restriction in Redding and Weinstein (2018) that the *average tastes* across barcode items in these global brands are the same in the two countries. The main difference between the geometric mean and the exact price index, there-

fore, comes from product variety. The magnitude of these product variety terms in China relative to the U.S. differs across the products in our sample. In the four goods for which we purchased Nielsen (China) scanner data, i.e., Toothpaste, Laundry Detergent, Personal Wash items and Shampoo, the availability of additional Chinese brands leads to greater variety than in the United States, and therefore lower Chinese price indexes for that reason. In the other 15 products for which we only have scraped data, however, there is much less presence of U.S. brands in the Chinese market. In these cases, the observed price differences between the countries (usually lower prices in China) are partially or fully offset by the variety differences (less variety in China), so that the cost-of-living in China is not as low as the price differences suggest, especially for smaller cities.

References

- ANTONIADES, A., M. XU, AND R. C. FEENSTRA (2019): "Distribution as Expenditure," Working paper, Georgetown University.
- ARKOLAKIS, C. (2010): "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 118(6), 1151–1199.
- ATKESON, A., AND A. BURSTEIN (2008): "Pricing-to-market, Trade Costs, and International Relative Prices," *American Economic Review*, 98(5), 1998–2031.
- BAGWELL, K. (2007): "The Economic Analysis of Advertising," *Handbook of Industrial Organization*, 3, 1701–1844.
- CAMPBELL, J. R., AND H. A. HOPENHAYN (2005): "Market Size Matters," *The Journal of Industrial Economics*, 53(1), 1–25.
- 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.
- (2015): "The Price Impact of Joining a Currency Union: Evidence from Latvia," *IMF Economic Review*, 63(2), 281–297.
- COUTURE, V., B. FABER, Y. GU, AND L. LIU (2018): "E-Commerce Integration and Economic Development: Evidence from China," Discussion paper, National Bureau of Economic Research no. 24384.
- DELLAVIGNA, S., AND M. GENTZKOW (2017): "Uniform Pricing in US Retail Chains," Discussion paper, National Bureau of Economic Research no. 23996.
- DINGEL, J., A. MISCIO, AND D. R. DAVIS (2018): "Cities, Lights, and Skills in Developing Economies," Discussion paper, University of Chicago and Columbia University.
- FAN, J., L. TANG, W. ZHU, AND B. ZOU (2018): "The Alibaba Effect: Spatial Consumption Inequality and the Welfare Gains from E-commerce," *Journal of International Economics*, 114, 203–220.
- FEENSTRA, R. C. (1994): "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84(1), 157–177.

- (2016): *Advanced International Trade*, Princeton University Press, 2nd edition.
- FEENSTRA, R. C., AND H. MA (2009): “Optimal choice of Product Scope for Multiproduct Firms under Monopolistic Competition,” in *The Organization of Firms in a Global Economy*, ed. by E. Helpman, D. Marin, and T. Verdier, pp. 173–199. Harvard university Press.
- HANDBURY, J., AND D. E. WEINSTEIN (2015): “Goods Prices and Availability in Cities,” *The Review of Economic Studies*, 82(1), 258–296.
- HOTTMAN, C., S. J. REDDING, AND D. E. WEINSTEIN (2016): “Quantifying the Sources of Firm Heterogeneity,” *Quarterly Journal of Economics*, 131(3), 1291–1364.
- KOKOVIN, S., M. PARENTI, J.-F. THISSE, AND P. USHCHEV (2017): “On the Dilution of Market Power,” Discussion paper, CEPR No. DP12367.
- LAGAKOS, D. (2016): “Explaining Cross-Country Productivity Differences in Retail Trade,” *Journal of Political Economy*, 124(2), 579–620.
- MACEDONI, L. (2017): “Multiproduct Superstars Across Rich and Poor Countries. Theory and Evidence,” Discussion paper, Aarhus University.
- MUTH, M. K., M. SWEITZER, D. BROWN, K. CAPOGROSSI, S. KARNS, D. LEVIN, A. OKRENT, P. SIEGEL, AND C. ZHEN (2016): *Understanding IRI Household-based and Store-based Scanner Data*. United States Department of Agriculture, Economic Research Service.
- PARENTI, M., P. USHCHEV, AND J.-F. THISSE (2017): “Toward a Theory of Monopolistic Competition,” *Journal of Economic Theory*, 167, 86–115.
- RAFF, H., AND J. WAGNER (2013): “Productivity and the Product Scope of Multi-Product Firms: A Test of Feenstra-Ma,” *Economic Bulletin*, 33(1), 173–199.
- REDDING, S. J., AND D. E. WEINSTEIN (2018): “Measuring Aggregate Price Indexes with Demand Shocks: Theory and Evidence for CES Preferences,” Discussion paper, National Bureau of Economic Research no. 22479.
- SATO, K. (1976): “The Ideal log-Change Index Number,” *Review of Economics and Statistics*, 58, 223–228.

SUTTON, J. (1991): *Sunk costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. MIT press.

VARTIA, Y. O. (1976): "Ideal log-Change Index Numbers," *Scandinavian Journal of Statistics*, 3, 121–126.