THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Estimating US consumer gains from Chinese imports

OPEN ACCESS

Estimating US Consumer Gains from Chinese Imports[*]

Liang Bai
University of Edinburgh

Sebastian Stumpner[†]
Université de Montréal &
Banque de France

First Draft: April 29, 2016
This Draft: October 15, 2018

Abstract

We estimate the size of US consumer gains from Chinese imports during 2004-2015. Using barcode-level price and expenditure data, we construct inflation rates under CES preferences, and use Chinese exports to Europe as an instrument. We find significant negative effects of Chinese imports on US prices. This effect is driven by both changes in the prices of existing goods and the entry of new goods and it is similar across consumer groups by income or region. A simple benchmarking exercise suggests that Chinese imports led to a 0.19 ppt annual reduction in the price index for consumer tradables.

1

1 Introduction

Recent years have seen a significant shift in US expenditure towards imports from China (Figure 1). What are the welfare consequences of this change for the US? Previous research has shown that imports from China may have accounted for up to one-quarter of the contemporaneous decline in US manufacturing employment (Autor et al. (2013)), and reduced lifetime earnings of affected workers (Autor et al. (2014)). Standard economic theory suggests these negative effects should be accompanied by significant consumer gains. The magnitude and distribution of these gains, however, have yet to be systematically estimated. In this paper, we provide the first micro-estimates of these effects using barcode-level price and expenditure data, decompose them into various margins of adjustment, and investigate heterogeneous effects across consumer groups by income and region.

Using variation across product categories in both cost-of-living inflation rates and Chinese import growth during 2004-15, our results show sizable gains for US consumers. Comparing a product category with median China trade shock to one with no change, prices in the median category grew by 0.17 percentage points less per year. A simple benchmarking exercise implies that the ideal price index for the set of tradable goods we analyze declined by 0.19 percentage points per year.

We next decompose the effect of Chinese import penetration on US prices along various margins of adjustment. Results show that two thirds of the effect is driven by lower inflation among existing goods (intensive margin) and one third by the introduction of new goods and disappearance of old goods (extensive margin). This suggests the presence of both pro-competitive effects (Feenstra and Weinstein (2017)) and variety gains (Broda and Weinstein (2006)). Chinese imports lead to higher rates of product entry and exit, but no net change in the number of consumed varieties. In contrast to the gains from final-good imports, results for intermediate-goods are inconclusive. We also find no evidence of heterogeneous effects across consumer groups by income or region.

The consumption data we use come from Nielsen's Homescan Panel, which contains information on purchases and demographic characteristics for a nationally-representative sample of around 60,000 US households between 2004 and 2015. The data on purchases are highly detailed, giving us information about prices paid by households and their expenditure for over 1.5 million barcoded goods. To link this dataset with international trade and production data, we build a new concordance from Nielsen product modules to HS 6-digit commodities, and by extension to NAICS 6-digit and ISIC 4-digit industries. Our resulting dataset covers roughly half of expenditure on consumer tradables. Despite this limited coverage, we find evidence suggesting the results generalize to the entire manufacturing sector. A more detailed discussion of both the benefits and limitations of the data is included in the following section.

Our main estimation equation is consistent with a large class of trade models. In particular, we use the insight from Arkolakis et al. (2012) that the domestic welfare effects of a foreign

shock can be summarized by the change in the share of expenditure on domestically produced goods. We then compare the evolution of prices in product categories with differential change in the domestic share of expenditure (DSE), induced by supply shocks in China.

Estimating the causal effect of Chinese imports on US inflation has a number of empirical challenges. The main threat to identification is that both prices and imports from China may be driven by demand or supply shocks in the US instead. To deal with this, we use a strategy similar to Autor et al. (2013) and instrument for the change in the domestic share of US expenditure using Chinese import penetration in Europe. We discuss this and other potential threats to identification in the next section.

This paper fits into a growing empirical literature on the global welfare implications of China's rapid growth: see Autor et al. (2013), Autor et al. (2014), and Pierce and Schott (2016) on manufacturing employment in the US, and Bloom et al. (2016) on technical change in Europe.[1] Our work contributes to this literature by providing the first micro-estimates of US consumer gains, using barcode-level data. It is closely related to Amiti et al. (2017), who estimate the effect of China's WTO entry on the US manufacturing price index. While their work sheds light on the origins of China's recent export surge and the policies responsible, our approach constructs directly cost-of-living inflation rates from micro data, investigates the various margins of adjustment, and estimates heterogeneous effects of the China trade shock for different consumer groups.[2] [3]

Another related paper is Broda and Romalis (2008), which investigate the distributional consequences of US-China trade. There are a number of important differences between our two studies. First, we study in detail the channels through which Chinese imports affect domestic prices, such as intensive margin price growth, variety effects, and the role of imported intermediate goods. Second, our main empirical equation is consistent with a large class of trade models, which allows for a tighter link between import growth and welfare changes. Finally, we adopt a different empirical strategy, using Chinese exports to Europe as an instrument for the change in the domestic share of US expenditure.

---

[1]For recent studies using quantitative trade models, see Di Giovanni et al. (2014) and Hsieh and Ossa (2016).

[2]Other empirical studies on the effect of international trade on producer prices include Bugamelli et al. (2010) and Auer and Fischer (2010).

[3]In more recent, ongoing work, Jaravel and Sager (2018) use price data from the CPI to estimate the effect of Chinese imports on US consumer prices during 2000-07. While the CPI data has broader coverage, the advantage of the Homescan data is that we can work with price and expenditure information based on actual purchases made by households. This allows us to analyze the effect of Chinese imports on effective prices paid by households, instead of retailers' posted prices. It also allows us to explicitly deal with heterogeneous product quality which may be particularly important for new goods entering from China. Finally, the ability to link each individual purchase to household characteristics enables us to compute category-by-group inflation rates to investigate heterogeneous effects.

2 Empirical Approach

How can we assess the effects of a positive supply shock in China on prices in the US? Arkolakis et al. (2012) analyze domestic welfare effects of foreign shocks in a large class of trade models characterized by Dixit-Stiglitz preferences, one factor of production, linear cost functions, no external economies of scale, and either perfect or monopolistic competition.[4] In this class of models, the domestic welfare effects of any foreign shock can be summarized by the change in the share of expenditure on domestically produced goods. Taking the domestic wage as numeraire, welfare equals the inverse of the domestic price index. That is, US welfare increases through lower prices if the Chinese supply shock results in a lower share of US expenditure on domestically produced goods:

$$\Delta \log(P) = \frac{1}{\theta} \Delta \log(DSE), \tag{1}$$

where DSE denotes the domestic share of expenditure and θ is the trade elasticity. Intuitively, welfare changes in the home country depend on changes in the terms of trade. When the terms of trade improve, the country imports more ($\Delta \log(DSE) < 0$) and the local price index goes down. Since the change in the DSE serves as a sufficient statistic for the price change, it includes all relevant general equilibrium effects that may arise as result of a Chinese supply shock, such as the indirect effects through the bilateral terms of trade between the US and third countries, and any resulting changes in those countries' market shares.

In our empirical application, we use a version of this equation at the level of individual product categories. In a multi-sector context, equation 1 extends to the sector level if one assumes perfect competition, no fixed costs, and perfect factor mobility across sectors. In that case, we can write

$$\Delta \log(P_i) = \frac{1}{\theta} \Delta \log(DSE_i) \tag{2}$$

for each sector i.

2.1 Data Sources

We employ two main datasets: i) international trade flows from UN Comtrade at the HS 6-digit level, obtained from CEPII (Gaulier and Signago (2010)), and ii) household purchases and product prices from AC Nielsen's Homescan Panel at the barcode level.

The Homescan dataset has been used widely in empirical economic studies (e.g. Broda and Weinstein (2010), Bronnenberg et al. (2012), Handbury and Weinstein (2014), Hottman et al. (2016)). It comes from a sample of around 60,000 US households who continually provide information about their demographic characteristics and product purchases. Information on

---

[4]This class of models includes Anderson (1979), Eaton and Kortum (2002), Krugman (1980), and Melitz (2003) with a Pareto distribution.

product purchases include price and expenditure at the level of individual barcodes and are reported by households using a handheld scanner. For our study period, we observe over 1.5 million barcodes, grouped by Nielsen into 1,147 product modules.[5] A more detailed description of the data is given in Appendix B.

There are numerous advantages to using this data in our context. First, defining a product at a very fine level is important for distinguishing price changes of the same product from changes in product composition. It also allows us to hold within-product quality constant, since quality of the same barcoded good should not change over time (product updates usually result in a new barcode). Second, we can observe effective prices paid by households, rather than retailers' posted prices. This is important as households may change retailers to benefit from lower prices.[6] Third, observing a larger share of household purchases for our set of categories (instead of a smaller sample of prices) helps in estimating the importance of variety gains. Finally, linking purchases to household characteristics gives us an opportunity to study heterogeneous effects across income groups and regions.

A potential drawback of the Nielsen data is that it comes with its own classification of consumer goods into "product modules", for which no existing concordance to trade and production data is available. To overcome this, we build such a concordance ourselves, to HS 6-digit trade, and by extension to US (NAICS) and international (ISIC) production data. We leave a detailed description of this procedure for Appendix B.

Another drawback of the Nielsen data is its limited sectoral coverage. Average expenditure in 2015 amounted to $5,113, or 12.6% of reported household income. According to data from the Consumer Expenditure Survey, spending on tradable goods in 2015 accounted for 26.6% of average household income.[7] This suggests our data represent around half of total expenditure on tradables. On the trade side, Nielsen categories represent 18.3% of total US imports, and 23.9% of US imports from China, in 2004. Figure A.1 compares the evolution of the DSE and the China share for Nielsen categories vs. all HS codes. While the products in Nielsen are on average more closed (initial DSE of 82% vs. 73% of all HS codes), the magnitude of decline in DSE is very similar across the two groups. Furthermore, while the increase in the China share is sizable among Nielsen categories (3.1 ppt), it is even larger among all traded goods. This suggests our results may provide a lower bound for the effect of Chinese imports. Despite these similarities, we provide a more formal test of external validity in Section 3.3, using the producer price index.

An additional concern with the Nielsen data is related to within-sample coverage. For instance, households may report purchases of food more reliably than purchases of electronics. Table A.1 shows the distribution of spending and number of goods across broad product classes

---

[5]Unfortunately, the barcode cannot be used to identify a product's country of origin.

[6]See, e.g., Coibion et al. (2015)

[7]To arrive at this number, we add up expenditure on food at home, apparel and services, vehicle purchases, gasoline and motor oil, other vehicle expenses, and all other expenditures.

for the categories in our sample. While food is the largest product group, we still observe a substantial amount of purchases in groups such as household/office/school supplies, health and beauty, and also electrical appliances. Figure A.2 compares the distribution of spending across product groups for our sample of Nielsen categories to that implied by BEA expenditure ($= production + imports - exports$). Even though we do observe some over-representation of food and drinks in Nielsen, the distributions look reasonably similar. Despite this, we may still be concerned about measurement error within product categories if households do not report all purchases, *and* if this under-reporting varies across product categories. We address this concern by directly controlling for the quality of coverage and its evolution in Section 3.3.

Finally, we also make use of sectoral output data for Europe and the US to compute category-level expenditures. For Europe, we focus on Germany, France, Italy, Spain, and the UK, both because of their size and data availability. Details on the concordance between industry-level production data and HS-level trade data are left for Appendix B.

2.2 Category-Level Inflation Rates

Consumption in category i at time t is given by a non-symmetric CES aggregate over different varieties (i.e. barcodes) k:

$$C_{it} = \left(\sum_k a_i^{k\frac{1}{\sigma}} c_{it}^{k\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

The terms a_i^k denote unobserved product quality, which is assumed to be constant for a given barcode over our sample period. The ideal price index for this consumption bundle is given by:

$$P_{it} = \left(\sum_k a_i^k p_{it}^{k\,1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

For a changing basket of goods, inflation can be written as:[8]

$$\frac{P_{it}}{P_{it-1}} = \underbrace{\left(\prod_{k \in I_{it}} \left(\frac{p_{it}^k}{p_{it-1}^k} \right)^{\omega_{it}^k} \right)}_{\text{intensive margin}} \underbrace{\left(\frac{\lambda_{it}}{\lambda_{it-1}} \right)^{\frac{1}{\sigma-1}}}_{\text{extensive margin}},$$

where I_{it} denotes the set of "staying" goods that are present in both $t-1$ and t, and the weights ω_{it}^k sum to one and depend on market shares s_{it-1}^k and s_{it}^k.[9]

The term λ_{it} equals the fraction of expenditure at time t that goes towards staying goods.

---

[8]See Sato (1976), Vartia (1976), and Feenstra (1994).

[9]The weight equals $\omega_{it}^k = \dfrac{\frac{s_{it}^k - s_{it-1}^k}{\log(s_{it}^k) - \log(s_{it-1}^k)}}{\Sigma_k \frac{s_{it}^k - s_{it-1}^k}{\log(s_{it}^k) - \log(s_{it-1}^k)}}.$

Intuitively, when the share of expenditure on staying goods is declining, this must mean that entering varieties are more competitive than exiting varieties, which reduces the cost of living. The effect on inflation is larger the more complementary are the varieties, i.e., the lower is the elasticity of substitution σ.

We compute the price of each good p_{it}^k as the quantity-weighted average unit price for that barcode among all Nielsen purchases in a given year. Finally, we assume $\sigma = 5$ in our baseline estimates, but also consider alternative values in robustness exercises. Summary statistics for our inflation measure and all other relevant variables can be found in Appendix Table A.3.

2.3 Identification

Our main empirical estimation exploits cross-product variation in import penetration to identify the effect of the China trade shock on consumer prices and product varieties in the US. Following equation (2), we relate the log change of the price index to the log change in the domestic share of expenditure in each product category:

$$\Delta \log(P_i) = \alpha + \beta \Delta \log(DSE_i) + Z_i' \gamma + \varepsilon_i \tag{3}$$

To compute DSE_{it}, we first calculate total US expenditure on product category i as production + imports - exports. We then measure DSE_{it} as the fraction of expenditure that does not go towards imports.

The main threat to identification is that both prices and imports from China may be driven by demand or supply shocks in the US rather than Chinese supply shocks. For instance, a positive US demand shock should raise US prices and may also affect the US domestic share of expenditure. Similarly, a positive US supply shock would tend to lower US prices and increase the domestic share, as US products become more competitive. These types of biases could therefore lead OLS to either under- or over-estimate the true effects. As a first step to addressing these concerns, we include a number of control variables (the Z_is) to capture supply and demand changes in the US: growth in US labor productivity and average income, age, and income growth of relevant consumers at the product category level. Adding these controls should also alleviate concerns that our results are driven by a potential correlation between foreign and US specific shocks (e.g. a correlation between Chinese and US productivity shocks).

To the extent these controls do not fully capture the effects of US demand and supply shocks, we follow Autor et al. (2013) and use Chinese import penetration in Europe as an instrument. This is computed as the 2004-15 change in EU imports from China, divided by 2004 EU expenditure. The idea of the instrument is to isolate the part of the variation in ΔDSE that is due to supply changes in China, rather than supply or demand changes in the US. Intuitively, if the growth in US imports from China during 2004-2015 is driven either by productivity growth in China or a reduction in trade barriers, we should observe a corresponding

increase in Chinese exports to other developed countries, such as those in Europe.

A concern with this identification strategy is that supply shocks in the rest of the world may be correlated with those in China. This could lead us to overestimate the role of Chinese imports. However, this is unlikely to be a major concern. First, the rise of China in both world trade and US imports is much more prominent than any other country during our study period. Second, while our instrument has a positive effect on the China share of US expenditure, it has a negative effect on the rest-of-the-world (ROW) share. Third, our results are unchanged when we modify the instrument by using the residual from a regression of Chinese import penetration into the EU on Mexican, Canadian, and other Asian imports into the EU ("orthogonalized IV").[10]

A second concern with this identification strategy is that demand or supply shocks may be correlated between the US and Europe. In particular, this would be problematic if such shocks were *negatively* correlated, as this would lead us to overestimate the effect. We regard this as very unlikely. If they were positively correlated instead, one should observe an increase in Chinese import penetration in Europe to be associated with an increase in US prices, thereby leading us to underestimate the true effect. The orthogonalized IV, by using only the component of Chinese import penetration in Europe that is orthogonal to other countries' import penetration in Europe, should remove the influence of EU specific shocks on the instrument. If correlated shocks between the EU and US were important, we would therefore expect a significant increase in the coefficient when we use this IV. The fact that the estimated coefficient is unchanged suggests that, conditional on our set of controls, the role of correlated shocks across the US and EU is small.

3 Main Results

3.1 Inflation

We first estimate the effect of Chinese import penetration in Europe on the U.S. domestic share of expenditure for our sample of 232 product categories. Results are shown in Columns 1-2 of Table 1 (Panel A). All estimations are weighted with 2004 category-level household expenditure. As expected, categories with higher growth in European imports from China saw larger declines in the U.S. DSE. This first-stage relationship yields a strong F-statistic of 78.1.

Panel B of Table 1 decomposes this relationship. We approximate the change in the DSE as

$$\Delta \log(DSE) \approx -\frac{\Delta CSE}{DSE_{2004}} - \frac{\Delta RSE}{DSE_{2004}},$$

where *CSE* and *RSE* denote the China and ROW shares of US expenditure, respectively. That is, the log change in *DSE* equals the relative decline in the China share plus that in the ROW

---

[10]The other Asian countries include Japan, Korea, India, Taiwan, Vietnam, Malaysia, Thailand, and Indonesia.

share. Results show that European imports from China is positively correlated with the relative China share of US expenditure, and negatively correlated with the ROW share. This implies that Chinese imports are displacing both US and ROW products.

The remainder of Table 1 presents our estimated effects of $\Delta\log(DSE)$ on inflation. Both the least squares (Panel C) and IV estimates (Panel D) are positive and highly statistically significant, with the IV estimates slightly larger in magnitude. Introducing our demand and supply-side controls reduces the coefficient from 0.5 to 0.36, but it remains significantly different from zero. Our preferred specification in Column 4 implies that, comparing a product category with median change in *DSE* (decline by 5.0%) to one with no change, prices in the median category grew by 1.82 ppt less cumulatively during 2004-2015, or 0.17 ppt less per year.

3.2 Intensive vs. Extensive Margins

We next analyze the sources of this reduction in the cost of living. Specifically, we are interested in the effect of Chinese imports on the prices of pre-existing varieties (the intensive margin), and the set of available varieties (the extensive margin, a la Broda and Weinstein (2006)).

We define two measures of intensive-margin gains. Our "short-stay" measure compounds the year-on-year intensive-margin part of the inflation formula. This includes all goods that are observed for at least two consecutive years, including new (and potentially foreign) goods that only entered after 2004. In comparison, our "long-stay" measure uses only those goods that are consumed in all years. Around 30% of all goods sold in 2004 fall into this category. Panel A of Table 2 show that import penetration leads to a decline in both measures. The effect is somewhat stronger for the "short-stay" measure, which could reflect price cutting for soon-to-exit products.

Table A.4 presents further intensive-margin results from barcode-level regressions. We restrict attention to products that were on the market in 2004, and track their prices and sales over time. We then estimate the following specification:

$$y_{ik} = \alpha + \beta\Delta\log(DSE_i) + Z_i'\gamma + \varepsilon_{ik} \tag{4}$$

for category *i* and barcode *k*. y_{ik} is the annualized log change for a particular outcome (e.g. price, expenditure) from 2004 until the last year in which the product is observed. For this estimation, $\Delta\log(DSE_i)$ is defined as the annualized log change in the domestic share of expenditure over the same period, and is instrumented by a similarly-defined measure of Chinese import penetration in Europe. To measure exit, we define a dummy variable that equals one if the last year in which the product is observed is before 2015. The results suggest that products with more exposure to Chinese imports (i) reduced their prices by more, (ii) experienced a drop in sales, and (iii) were more likely to exit the market. Columns 4-5 show similar findings for

long stayers.

Since we don't observe mark-ups, only output prices, these intensive-margin results are consistent with two alternative interpretations: i) the existence of pro-competitive effects (i.e. reductions in mark-ups), or ii) a differential change in factor costs across product categories. In a model with labor as the only factor of production and perfect labor mobility across sectors, wages would equalize, and the differential decline in prices may be interpreted as evidence for pro-competitive effects. However, since we cannot rule out that factor costs changed differentially across products, our results are only suggestive of such effects.[11]

In Panel B of Table 2, we turn to extensive margin gains, defined as those due to the Feenstra adjustment factor. Results show that entry and exit of products contributed to lower inflation. In total, these account for roughly one third ($= 0.127/0.364$) of the decline.[12] These results are closely related to those in the seminal paper by Broda and Weinstein (2006), who find important variety gains from international trade that reduce cost of living in the US. Compared to their work, which defines a product as a ten-digit HTS category and a variety as a product-country pair, our definition of a variety as an individual barcode allows for a cleaner separation between intensive- and extensive-margin effects. Another notable difference between our two papers is the potential exit of domestic varieties due to import competition, which is unaccounted for in Broda and Weinstein (2006). By including this channel, our results extend their earlier work by showing that variety gains play a key role even after accounting for the exit of previously-consumed varieties.

Extensive margin gains can arise with or without a net increase in the number of consumed varieties. Panel C of Table 2 tests whether Chinese imports have led to a more crowded product space. Here we don't find any evidence of differential growth effects in either variety or expenditure. The absence of an expenditure effect, together with the findings on inflation, suggests an elasticity of substitution across categories close to one. Panel D study the channels behind variety gains by looking at average entry and exit rates.[13] Unsurprisingly, categories that experience stronger Chinese import growth have higher entry rates. Meanwhile, imports also lead to more exit of other varieties. Figure A.3 further splits up the effect by price quintile in which entry or exit occurs. While Chinese imports lead to more entry of new varieties at all quintiles of the price distribution, exit is concentrated at the bottom of the distribution.

Finally, Panel A of Appendix Table 3 shows results for import penetration in intermediate inputs. We compute the weighted average of intermediate goods' change in DSE, $\Delta \log(DSE_i)^{IG} = \sum_j s_{ij} \Delta \log(DSE_j)$, where the weights come from the BEA direct requirements table. The in-

---

[11]They are therefore consistent with Feenstra and Weinstein (2017), who provide evidence for pro-competitive effects in the U.S. from international competition.

[12]In Column 2 of Panel B, we construct the adjustment factor by defining a new variety as one that has never appeared in the data before, rather than one that didn't appear in the previous year. The definition of an exiting variety is changed in an analogous way.

[13]For each category-year, we define the fraction of entrants as the share of currently-consumed goods that were not consumed in the previous period. Exiters are defined analogously. We then average entry and exit rates over time.

strument for this variable is analogously defined. Unfortunately, the coefficient is very imprecisely estimated, leaving us unable to draw any firm conclusions regarding the effect of intermediate goods imports from China.

3.3 Robustness

We carry out a number of robustness exercises to investigate: i) whether our inflation results are driven by the Nielsen data's specific coverage of goods for the products within our sample, and ii) the extent to which our results generalize to the entire manufacturing sector.

For the first issue, we may be concerned that the quality of coverage (i.e. the ratio of reported to actual expenditure) differs across product categories. This could happen, for instance, if households only scan purchases from larger shopping trips (such as for food), but don't report purchases from one-off trips, such as for electrical appliances. Our estimates may then be biased if (i) the quality of coverage is correlated with the degree of Chinese import penetration and (ii) low coverage leads to a systematic mismeasurement of inflation. To address this, we first compute the coverage ratio for each product category, as the ratio of projected Nielsen purchases to imputed BEA expenditures,[14] and its change during 2004-2015. We find no correlation between our instrument and the change in the coverage ratio ($\rho = -0.02$), and only a mild negative correlation with the 2004 coverage ratio ($\rho = -0.08$). Very importantly, when we add quartile dummies of both the baseline coverage ratio and its change to our regression, the coefficient on ΔDSE becomes slightly larger, and remains highly significant (Column 1, Panel C of Table 3). Furthermore, the results are unchanged when we restrict the sample to the top half of categories in terms of initial coverage (Column 2, Panel C).

Next, we evaluate how representative our inflation result is for the entire manufacturing sector. We re-run our estimation using PPI data at the six-digit industry level, once for the 128 industries that map into our set of product categories, and once for the entire manufacturing sector (251 industries).[15] Results are shown in Panel D of Table 3. The point estimate for the subset of industries equals 0.27, which is slightly lower than what we find using the Nielsen data. This is likely due to the absence of an extensive margin when constructing the PPI.[16] Most importantly, however, the point estimates for the 128 and 251 industries are remarkably similar (0.27 and 0.33 respectively), suggesting that our main inflation result largely generalizes to the entire manufacturing sector.

Panels B and E present further robustness tests. First, we use the orthogonalized IV in Panel B, and find the coefficient to be unchanged. As discussed in section 2.3, this suggests that neither correlated supply shocks between China and other countries nor correlated shocks

---

[14]The median for this ratio is 43%.

[15]A BEA industry maps into one of our categories if the industry maps into at least one HS code that is covered by our sample.

[16]Other potentially important dimensions of difference between the PPI and our intensive-margin measures include: i) product coverage, ii) sampling frequency, and iii) price formula.

between the US and Europe are driving our results. Panel E tests whether any potential pre-trend in inflation is influencing our estimates. Here we control directly for PPI inflation between 1997 and 2000, the period immediately before China's accession to the WTO. Compared to Panel D, the coefficient estimates are around 20% lower, but remain highly significant.

Some additional robustness tests on the inflation results are shown in Appendix Table A.5. Here we i) exclude all food and drinks categories, all electrical appliances, and all categories not within the 10th-90th percentile values for inflation and Δ DSE (Panel A), ii) use alternative values for the elasticity of substitution (Panel B), iii) perform regressions without weights and with BEA-based expenditure weights (Panel C), and iv) employ a gravity IV, similar to Autor et al. (2013), using the difference between Chinese and U.S. export growth to Europe (Panel D).

3.4 Heterogeneous Effects by Income Group and Region

Since the Nielsen data include households' income and area of residence, we can estimate heterogeneous treatment effects within the US.[17] To do so, we first divide individuals into five similarly-sized groups based on their reported annual household income. We then compute category-level inflation rates separately for each group, allowing the set of goods and good-specific weights within each category to vary. Similarly, we can group individuals into one of four US census regions (Northeast, Midwest, South, West) based on their residence.

The results are shown in Table 4. Firstly, our earlier aggregate result on inflation holds consistently across income groups (Panel A) and regions (Panel B). Secondly, these effects are similar across groups. In Panel A, the coefficient is increasing in income, but the differences are too small to be statistically significant.

There are potentially important dimensions of heterogeneity we are unable to study. For instance, households may differ in the fraction of their expenditure allocated to tradables vs. nontradables, as in Fajgelbaum and Khandelwal (2016).[18] In such cases, poorer households may benefit more from growing imports, if more of their spending is concentrated on tradables.[19]

3.5 Aggregation

We now use our cross-sectional estimates to carry out a simple aggregation exercise to estimate the effect of Chinese import growth on the overall consumer price index for our product cate-

---

[17]In a similar exercise, Faber (2014) finds that high-income households benefited more from trade liberalization than low-income households during Mexico's NAFTA accession.

[18]These results are also consistent with Borusyak and Jaravel (2017) who, using a quantitative model, find that the distributional effects from trade liberalization for the US mostly operate through the earnings channel.

[19]Indeed, we see some evidence consistent with this in our data, as Nielsen purchases in 2015 constituted 27% of household income for those in the lowest income quintile, compared to 6% for those in the highest income quintile.

gories. We assume Cobb-Douglas preferences across categories (consistent with the result that category expenditure shares remain unaffected), and write the predicted change in the aggregate price induced by Chinese supply shocks as follows:

$$\widehat{\Delta \log(P)}^{Chn} = \sum_i \omega_i \widehat{\Delta \log(P_i)}^{Chn},$$

where ω_i is the weight of category i in aggregate consumption.

We next assume that a foreign trade shock only affects prices if there is a change in the DSE. In other words, the level effect in the second stage is zero:

$$\widehat{\Delta \log(P_i)}^{Chn} = \hat{\beta}_{IV} \widehat{\Delta \log(DSE_i)}^{Chn}$$

where $\hat{\beta}_{IV}$ is the second-stage coefficient estimate, and $\widehat{\Delta \log(DSE_i)}^{Chn}$ is the predicted change in $log(DSE)$ that is caused by the China trade shock.

In order to derive $\widehat{\Delta \log(DSE_i)}^{Chn}$, we make use of our first-stage equation together with the assumption that in the aggregate time series, the China shock is responsible for the same share of the decline in *DSE* as in the cross-section (Appendix C gives more detail on this step of the aggregation). This procedure implies that roughly half of the aggregate 5.7ppt decline in the U.S. DSE is attributed to supply shocks in China.

With the values for $\widehat{\Delta \log(DSE_i)}^{Chn}$ in hand, we then use the expressions above to compute the aggregate change in the price index for consumer tradables. Doing so, we find a decline in the ideal price index of 2.1 ppts for the period 2004-2015, or 0.19 ppts per year.

4 Conclusion

This paper provides the first micro-estimates of US consumer gains from Chinese import growth during 2004-2015, using barcode-level data on prices and expenditure. We build exact price indices under CES preferences, and estimate the effect of Chinese imports on the constructed cost-of-living inflation. Our results indicate that Chinese import growth brought sizeable gains to US consumers: comparing a category with median China trade shock to one with no change, prices in the median category grew by 0.17 percentage point less per year. A simple benchmarking exercise suggests that inflation for our set of tradable consumer goods was 0.19 percentage point lower per year as a result.

Roughly two thirds of these gains can be attributed to price changes of existing goods, while the remaining third is due to variety gains. Though Chinese imports lead to more entry of new goods, this effect is somewhat muted due to the exit of previously-consumed varieties. Nonetheless, we find an overall positive contribution from the extensive margin. Unlike final goods, the effects of imported intermediate goods remains unclear. Finally, we don't find evi-

dence for differential gains across consumers by income or region. That said, poor households might still benefit more because they spend a larger share of their income on tradables.

While previous research has highlighted the negative labor market consequences due to import competition, our results suggest substantial gains to US consumers from the recent growth in trade with China. These ought to be taken into account in both the debate around, and the design of, US trade policy.

References

Amiti, Mary, Mi Dai, Robert Feenstra, and John Romalis, "How did China's WTO entry benefit U.S. consumers?," *NBER Working Paper No. 23487*, 2017.

Anderson, James E., "A Theoretical Foundation for the Gravity Equation," *American Economic Review*, 1979, *69* (1), 106–116.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New trade models, same old gains?," *The American Economic Review*, 2012, *102* (1), 94–130.

Auer, Raphael and Andreas M Fischer, "The effect of low-wage import competition on US inflationary pressure," *Journal of Monetary Economics*, 2010, *57* (4), 491–503.

Autor, David, David Dorn, and Gordon H Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *The American Economic Review*, 2013, *103* (6), 2121–2168.

_ , _ , _ , **and Jae Song**, "Trade Adjustment: Worker-Level Evidence," *The Quarterly Journal of Economics*, 2014, *129* (4), 1799–1860.

Bloom, Nicholas, Mirko Draca, and John Van Reenen, "Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity," *The Review of Economic Studies*, 2016, *83* (1), 87–117.

Borusyak, Kirill and Xavier Jaravel, "The Distributional Effects of Trade: Theory and Evidence from the United States," 2017.

Broda, Christian and David E Weinstein, "Product creation and destruction: Evidence and price implications," *American Economic Review*, 2010, *100* (3), 691–723.

_ **and David Weinstein**, "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 2006, *121*.

_ **and John Romalis**, "Inequality and Prices: Does China Benefit the Poor in America?," *mimeo, University of Chicago*, 2008.

Bronnenberg, Bart J, Jean-Pierre H Dubé, and Matthew Gentzkow, "The evolution of brand preferences: Evidence from consumer migration," *American Economic Review*, 2012, *102* (6), 2472–2508.

Bugamelli, Matteo, Silvia Fabiani, and Enrico Sette, "The pro-competitive effect of imports from China: an analysis of firm-level price data," 2010.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implication," *The American Economic Review*, 2015, *105* (3), 993–1029.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo, "Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?," *Quantitative Marketing and Economics*, 2010, *8* (2), 207–239.

Faber, Benjamin, "Trade Liberalization, The Price of Quality, and Inequality: Evidence from Mexican Store Prices," *mimeo, UC Berkeley*, 2014.

Fajgelbaum, Pablo D. and Amit K. Khandelwal, "Measuring the Unequal Gains from Trade," *Quarterly Journal of Economics*, 2016, *131* (3), 1113–1180.

Feenstra, Robert C., "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 1994, *84* (1), 157–177.

— **and David E. Weinstein**, "Globalization, Markups, and U.S. Welfare," *Journal of Political Economy*, 2017, *125*.

Gaulier, Guillaume and Soledad Signago, "BACI: International Trade Database at the Product Level. The 1994-2007 Version," *CEPII Working Paper No. 2010-23*, 2010.

Giovanni, Julian Di, Andrei A Levchenko, and Jing Zhang, "The global welfare impact of China: Trade integration and technological change," *American Economic Journal: Macroeconomics*, 2014, *6* (3), 153–183.

Handbury, Jessie and David E Weinstein, "Goods prices and availability in cities," *The Review of Economic Studies*, 2014, *82* (1), 258–296.

Hottman, Colin J, Stephen J Redding, and David E Weinstein, "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 2016, *131* (3), 1291–1364.

Hsieh, Chang-Tai and Ralph Ossa, "A global view of productivity growth in China," *Journal of International Economics*, 2016, *102*, 209–224.

Jaravel, Xavier and Erick Sager, "What Are the Price Effects of Trade? Evidence From the US and Implications for Quantitative Trade Models," *mimeo, London School of Economics*, 2018.

Krugman, Paul, "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 1980, *70* (5).

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6).

Pierce, Justin R and Peter K Schott, "The surprisingly swift decline of US manufacturing employment," *The American Economic Review*, 2016, *106* (7), 1632–1662.

Sato, Kazuo, "The Ideal Log-Change Index Number," *Review of Economics and Statistics*, May 1976, *58*, 223–228.

Soderbery, Anson, "Estimating Import Supply and Demand Elasticities: Analysis and Implications," *Journal of International Economics*, 2015, *96* (1), 1–17.

Vartia, Y.O., "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics*, 1976, *3*, 121–126.

Figure 1: Changing Composition of US Expenditure
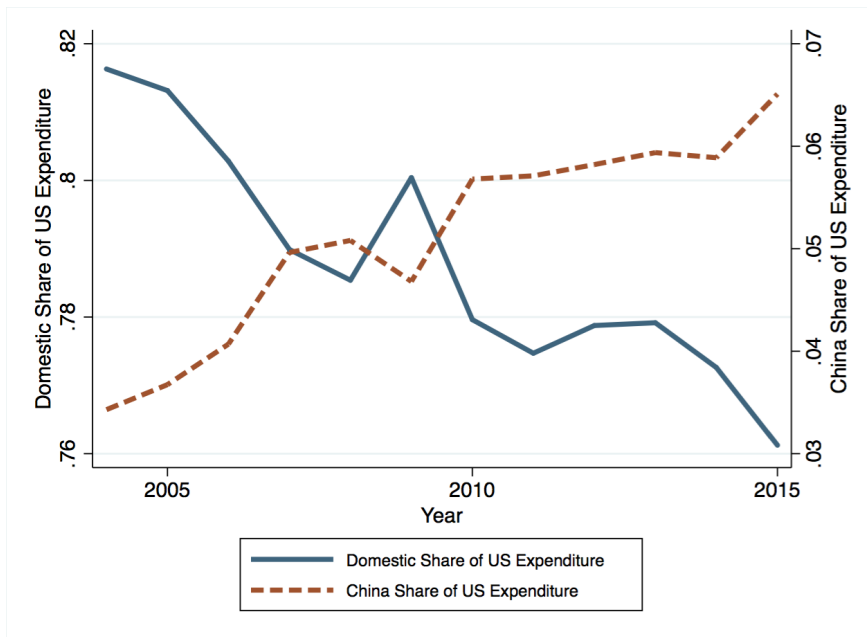
Notes: This figure shows the evolution of both the domestic and China shares of US expenditure (left and right axes respectively) during 2004-2015, using the set of HS 6-digit codes in our sample. Total US expenditure is computed as production + imports - exports.

Table 1: Import Penetration and Inflation

Dependent Variable:	A. First Stage		B. Decomposition of First Stage	
	Δ DSE		Δ China Share	Δ ROW Share
	(1)	(2)	(1)	(2)
China IP (Europe)	-2.630***	-2.677***	2.054***	-0.703**
	(0.315)	(0.303)	(0.381)	(0.328)
Controls	No	Yes	Yes	Yes
1st-Stage F-stat	69.6	78.1	N/A	N/A

Dependent Variable:	C. Least Squares		D. Instrumental Variables	
	Inflation		Inflation	
	(1)	(2)	(1)	(2)
Δ DSE	0.407***	0.290***	0.500***	0.364***
	(0.060)	(0.059)	(0.122)	(0.114)
Controls	No	Yes	No	Yes

Notes: $N = 232$. Robust standard errors in parentheses. *** p<0.01, ** p<0.05,* p<0.10. Δ DSE is the log change in the domestic share of expenditure for the US during 2004-2015. The dependent variable in Column 1 (2) of Panel B is the change in the Chinese (ROW) share in US expenditure, divided by the initial domestic share. The instrument used is China import penetration in the five largest European economies (Germany, France, UK, Italy and Spain). Unless otherwise noted, all regressions include the following controls: US productivity growth and quartile dummies for average income, age and income growth among the households with purchases in each category. All regressions are weighted by 2004 projected Nielsen expenditure.

Table 2: Margins of Adjustment

Dependent Variable:	A. Intensive Margin		B. Extensive Margin	
	Short-Stay Inflation	Long-Stay Inflation	New Varieties	New Varieties (robust)
	(1)	(2)	(1)	(2)
Δ DSE	0.237***	0.137***	0.127***	0.140***
	(0.074)	(0.038)	(0.044)	(0.046)

Dependent Variable:	C. Expenditure and Product Space		D. Entry and Exit Rates	
	Expenditure	No. of Products	Entry	Exit
	(1)	(2)	(1)	(2)
Δ DSE	0.228	-0.086	-0.077***	-0.075***
	(0.196)	(0.179)	(0.018)	(0.022)

Notes: $N = 232$. Robust standard errors in parentheses. *** p<0.01, ** p<0.05,* p<0.10. Panels A and B decompose the inflation result into intensive and extensive margins. The inflation rate in Panel A Column 1 uses all goods available in any consecutive years, while that in Column 2 is constructed only from goods that are consumed in every year. Panel B studies the role of new varieties (i.e. extensive margin), where the robust measure (Column 2) excludes goods that drop in and out of the sample. Δ DSE is the log change in the domestic share of expenditure for the US during 2004-2015. The instrument used is China import penetration in the five largest European economies (Germany, France, UK, Italy and Spain). All regressions include the following controls: US productivity growth and quartile dummies for average income, age and income growth among the households with purchases in each category. All regressions are weighted by 2004 projected Nielsen expenditure.

Table 3: Intermediate Goods and Robustness

Dependent Variable:	A. Interm. Goods	B. Orthogonalized IV	C. Internal Coverage	
	Inflation	Inflation	Inflation	
	(1)	(1)	(1)	(2)
Δ DSE	0.395***	0.365***	0.428***	0.352***
	(0.122)	(0.113)	(0.116)	(0.098)
Δ DSE, Intermediate Goods	-1.005			
	(1.178)			
Observations	232	232	232	118
Coverage Controls	No	No	Yes	Yes
Sample	All Cat.	All Cat.	All Cat.	High Coverage Cat.

Dependent Variable:	D. External Validity		E. Pre-Trends	
	PPI Inflation		PPI Inflation	
	(1)	(2)	(1)	(2)
Δ DSE	0.271**	0.333***	0.210**	0.249***
	(0.119)	(0.101)	(0.099)	(0.083)
Observations	128	251	128	251
Controls	No	No	PPI Trend	PPI Trend
Sample	Nielsen Ind.	All Ind.	Nielsen Ind.	All Ind.

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05,* p<0.10. Δ DSE is the log change in the domestic share of expenditure for the US during 2004-2015. Δ DSE, Intermediate Goods (US) is the weighted average of the log change in the DSE among the goods that an industry uses as inputs. The instrument used is China import penetration in the five largest European economies (Germany, France, UK, Italy and Spain), and weighted Chinese import penetration for the case of intermediate goods. Panel B presents results using an orthogonalized instrument, which is the residual from a regression of China import penetration in Europe on import penetration in Europe by other important trading partners of the US (i.e. Canada, Mexico, and numerous Asian countries). Panel C presents our inflation results with additional coverage controls: the coverage ratio is computed as that between Nielsen and BEA expenditure, and both its level in 2004 and change during 2004-15 are included as quartile dummies. Column 2 reports results using a sub-sample of 118 categories with above-median values in the coverage ratio. Panel D reports results from re-doing our analysis using PPI inflation, first on the subset of 128 BEA industries that map into our sample (Column 1), and then on the full sample of 250 BEA industries (Column 2). Estimations in Panel D include no controls, because we cannot compute the demand side controls for the full set of 250 industries. Panel E includes the 1997-2000 growth rate of the PPI as control. Unless otherwise noted, all regressions include the following controls: US productivity growth and quartile dummies for average income, age and income growth among the households with purchases in each category. All regressions are weighted by 2004 projected Nielsen expenditure.

Table 4: Heterogeneity

	A. Estimations by Income Group				
Dependent Variable: Inflation	$< 30k$ (1)	$(30k - 50k)$ (2)	$(50k - 70k)$ (3)	$(70k - 100k)$ (4)	$> 100k$ (5)
Δ DSE	0.255*** (0.088)	0.258*** (0.084)	0.276*** (0.080)	0.288*** (0.082)	0.312*** (0.092)

	B. Estimations by Region			
Dependent Variable: Inflation	Northeast (1)	Midwest (2)	South (3)	West (4)
Δ DSE	0.259*** (0.078)	0.269*** (0.091)	0.313*** (0.094)	0.259*** (0.085)

Notes: $N = 232$. Robust standard errors in parentheses. *** p<0.01, ** p<0.05,* p<0.10. Panel A shows results using income group-specific inflation rates, while Panel B shows results using region-specific inflation rates. Δ DSE is the log change in the domestic share of expenditure for the US during 2004-2015. The instrument used is China import penetration in the five largest European economies (Germany, France, UK, Italy and Spain). Unless otherwise noted, all regressions include the following controls: US productivity growth and quartile dummies for average income, age and income growth among the households with purchases in each category. All regressions are weighted by 2004 projected Nielsen expenditure.

A   Additional Figures and Tables

Figure A.1: Changing Composition of US Expenditure
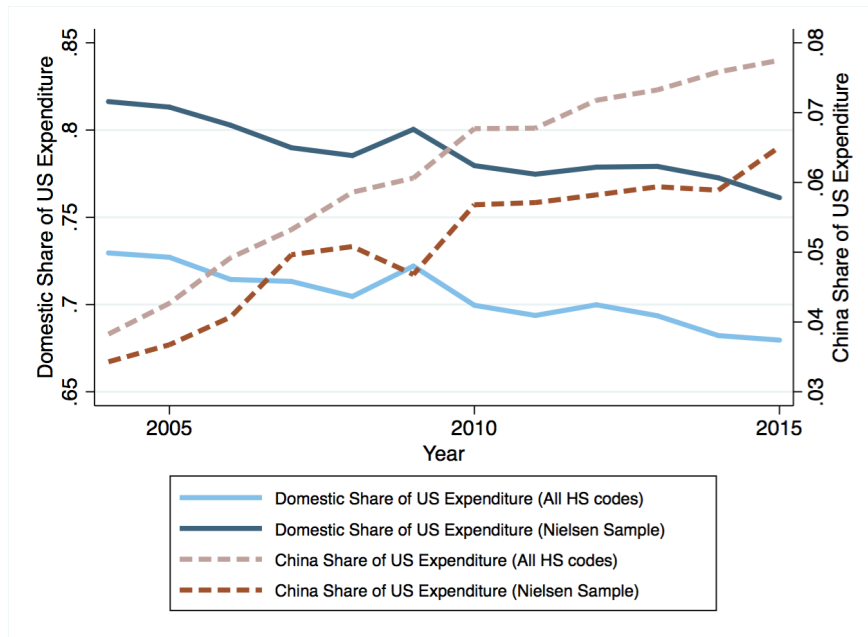
Notes: This figure shows the evolution of both the domestic and China shares of US expenditure (left and right axes respectively) during 2004-2015, comparing all HS 6-digit codes with those in the Nielsen sample. Total US expenditure is computed as production + imports - exports.

Figure A.2: Nielsen Product Coverage
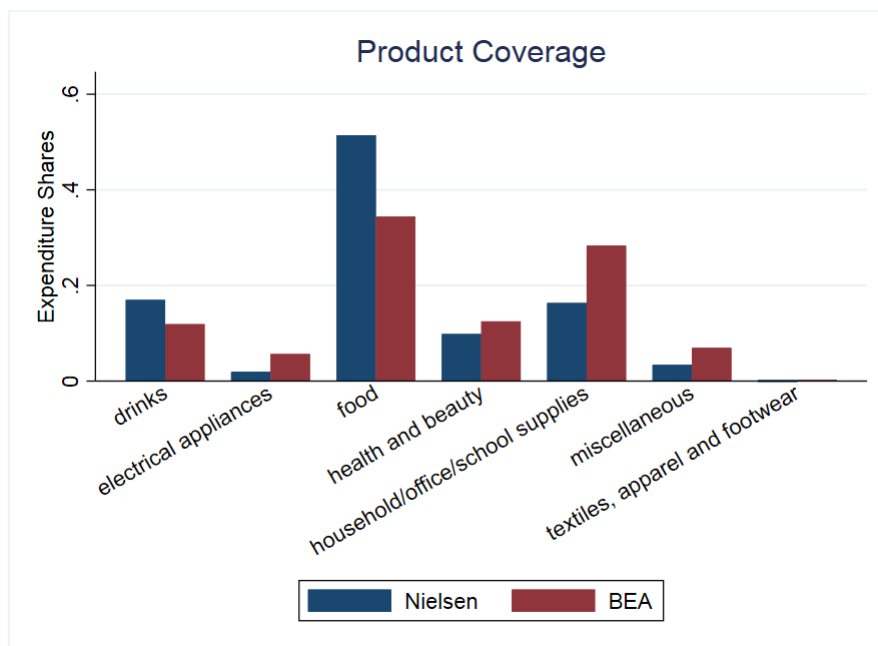
Notes: This figure shows the expenditure shares by broad product groups, averaged during 2004-2015, for both projected Nielsen and imputed BEA expenditures. Imputed BEA expenditure is computed as production + imports - exports. Note that this captures only the part of imputed BEA expenditure that map into the Nielsen data, and not necessarily all BEA-implied expenditure on a particular broad product group.

Figure A.3: Effect on Entry and Exit Rates by Price Quintile
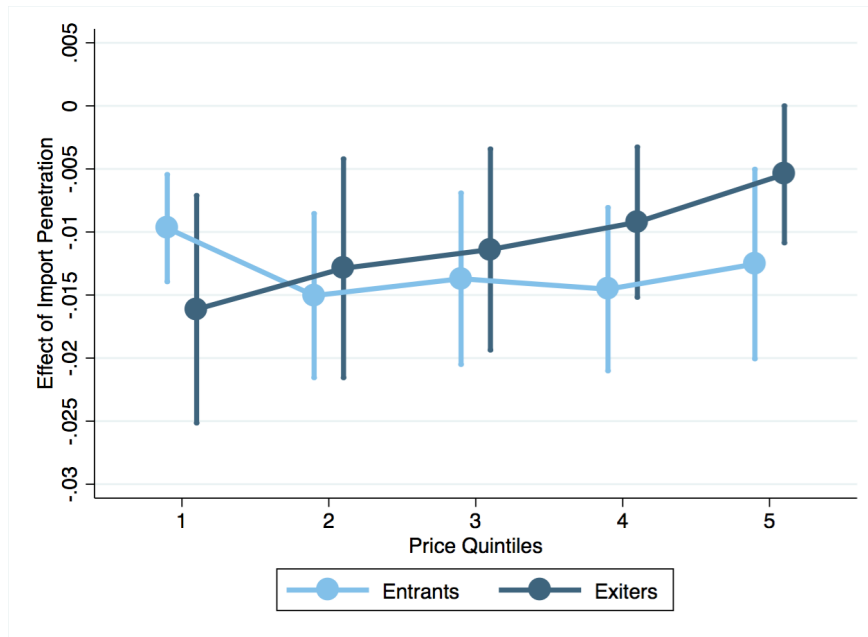
Notes: This figure shows coefficient estimates (and 95% confidence intervals) of quintile-specific entry and exit rates on Δ DSE, instrumented by Chinese import penetration in Europe. For each category and year, we first compute quintile-specific entry and exit rates. The entry rate is computed as $EntryRate_{c,t}^q = \frac{M_{ct}^q}{N_{ct}}$, where M_{ct}^q denotes the number of entrants in category c at year t in price quintile q, and N_{ct} is the number of all varieties observed in c and t. We then average entry rates over time. Exit rates are computed analogously. All estimations include our main set of controls: US productivity growth and quartile dummies for average income, age and income growth among the households with purchases in each category. All regressions are weighted by 2004 projected Nielsen expenditure.

Table A.1: Nielsen Expenditures - by broad product groups

Product Group	(1) Number of Categories	(2) Expenditure ($ bn) 2004	(3) Expenditure ($ bn) 2015	(4) Number of Barcodes 2004	(5) Number of Barcodes 2015
Drinks	27	58.7	69.9	65187	81011
Electrical Appliances	19	7.6	7.3	12225	14989
Food	93	168	226	281771	320785
Health and Beauty	31	33.2	42.9	86102	98222
Household/Office/School Supplies	52	58.3	65.1	144734	175500
Miscellaneous	8	13.0	11.3	8945	6803
Textiles, Apparel and Footwear	2	0.2	0.2	986	918
Total	232	339	423	599950	698228

Notes: Based on projected Nielsen expenditures for our analytical sample of 232 product categories.

Table A.2: Nielsen - HS Concordance Merge Types

Merge Type	(1) Number of Categories	(2) Number of Nielsen Product Modules	(3) Number of HS 6-digit Codes
1:1	125	125	125
1:n	51	51	284
m:1	87	564	87
m:n	61	407	382
Total	324	1147	878

Notes: A 1:1 merge refers to a case where a single Nielsen product module is matched to a single HS 6-digit code. A 1:n merge refers to a case where a single Nielsen product module is matched to multiple HS 6-digit codes. A m:1 merge refers to a case where multiple Nielsen product modules are matched to a single HS 6-digit code. A m:n merge refers to a case where multiple Nielsen product modules are matched to multiple HS 6-digit codes.

Table A.3: Summary Statistics

Variable	Mean	Std. Dev.	P10	P50	P90
A. Inflation (category level)					
Inflation (σ = 3)	0.91	0.37	0.40	0.94	1.38
Inflation (σ = 5)	1.02	0.32	0.60	1.04	1.42
Inflation (σ = Broda Weinstein Elast.)	0.96	0.40	0.41	1.02	1.43
B. Inflation (various margins)					
Intensive Margin (short stayers)	1.17	0.25	0.86	1.17	1.48
Intensive Margin (long stayers)	1.26	0.22	1.06	1.24	1.49
Extensive Margin (σ = 5)	0.86	0.13	0.67	0.89	0.98
C. Additional Outcomes					
Log Change of Expenditure	0.04	0.48	-0.48	0.08	0.48
Log Change in the No. of Barcodes	0.02	0.37	-0.34	0.04	0.45
Average Entry Rate	0.24	0.09	0.13	0.23	0.36
Average Exit Rate	0.24	0.09	0.13	0.22	0.35
D. China Trade Shock					
Δ DSE	-0.18	0.39	-0.57	-0.05	0.02
China IP (Europe)	0.05	0.11	0.00	0.02	0.13

Notes: $N = 232$. All variables are computed for the time period 2004-15. P10, P50, and P90 refer to the 10th, 50th, and 90th percentile values, respectively. The short stayer intensive margin inflation measure includes all goods that are observed in any two consecutive years, while the long stay measure only includes goods that are observed in every year. For the average entry (exit) rate, we first compute the category-by-year entry (exit) rate as the number of entering (exiting) barcodes, divided by all barcodes in that category and year, and then average these rates over time. Δ DSE is the log change in the domestic share of expenditure for the US during 2004-2015. China IP is China import penetration in the five largest European economies (Germany, France, UK, Italy and Spain).

Table A.4: Barcode-level Regression Results

Dependent Variable:	Price	Expenditure	Exit	Price	Expenditure
	(1)	(2)	(3)	(4)	(5)
Δ DSE	0.649***	3.898**	-6.409***	0.140	2.198**
	(0.171)	(1.574)	(1.505)	(0.114)	(1.024)
Observations	472,335	472,335	472,335	142,993	142,993
Sample	All	All	All	Stayer	Stayer

Notes: Category-level clustered standard errors in parentheses. *** p<0.01, ** p<0.05,* p<0.10. Columns 1-3 use as sample all barcodes that were present in 2004, while Columns 4-5 use as sample barcodes that are present in every year. Δ DSE is the log change in the domestic share of expenditure for the US during 2004-2015. The instrument used is China import penetration in the five largest European economies (Germany, France, UK, Italy and Spain). Unless otherwise noted, all regressions include the following controls: US productivity growth and quartile dummies for average income, age and income growth among the households with purchases in each category. All estimations are weighted with initial expenditure for a given barcode.

Table A.5: Further Robustness

	A. Sub-Sample Analysis			B. Inflation w/ Alternative Elasticities		
	No Food & Drinks	No Electr. Appl.	10th-90th			
Dependent Variable:	Inflation	Inflation	Inflation	$\sigma = 3$	$\sigma = 10$	$\sigma = BW$
	(1)	(2)	(3)	(1)	(2)	(3)
Δ DSE	0.212**	0.208**	0.437**	0.491***	0.293***	0.925***
	(0.093)	(0.090)	(0.214)	(0.156)	(0.091)	(0.193)
Observations	112	213	167	232	232	232

	C. Alternative Weights				D. Gravity IV
	First Stage		IV		
Dependent Variable:	Δ DSE	Δ DSE	Inflation	Inflation	Inflation
	(1)	(2)	(3)	(4)	(1)
China IP (Europe)	-2.498***	-2.050***			
	(0.351)	(0.436)			
Δ DSE			0.668***	0.745**	0.387***
			(0.159)	(0.294)	(0.146)
Weights	No	BEA	No	BEA	Nielsen

Notes: $N = 232$ (unless otherwise indicated). Robust standard errors in parentheses. *** p<0.01, ** p<0.05,* p<0.10. Δ DSE is the log change in the domestic share of expenditure for the US during 2004-2015. The instrument used is China import penetration in the five largest European economies (Germany, France, UK, Italy and Spain). Panel A reports inflation results when we separately exclude food and drinks (Column 1), electrical appliances (Column 2), and categories outside of the 10th and 90th percentile values in inflation and *ΔDSE* (Column 3). Columns 1 and 2 of Panel B show results when the inflation rate is computed using an elasticity of substitution equal to 3 and 10 respectively, while Column 3 shows results for category-specific elasticities, estimated using the methodology by Broda and Weinstein (2006), and obtained from Soderbery (2015). Panel D uses a gravity-model based instrument, constructed as the difference between Chinese and US import growth in the EU, weighted by the initial China share in EU expenditure. Unless otherwise noted, all regressions include the following controls: US productivity growth and quartile dummies for average income, age and income growth among the households with purchases in each category. The weights used in Panels A, B. and D are total projected Nielsen expenditure in a given category in 2004. BEA weights in Panel C are total US expenditure in a given category in 2004, computed as production - exports + imports.

B Data Appendix

B.1 Description of AC Nielsen Homescan Data

The Homescan data come from a private vendor, AC Nielsen, and are made available through the Kilts Center at the University of Chicago. It is a panel of roughly 60,000 households that provide information about their product purchases (price and expenditure) at the level of individual barcodes.[20] For our study period, we observe over 1.5 million barcodes, grouped by Nielsen into 1,147 product modules. Some examples of product modules include olive oil, pasta-spaghetti, dental accessories, cameras, batteries, and printers.

Nielsen recruits households by mail or online, and provides incentives to join and to remain active in reporting transactions. Examples of incentives include monthly prize drawings and gift points. Households that do not regularly report their transactions are removed from the sample and new households are introduced. In all of our results, we use the household projection weights provided by Nielsen to make the sample demographics representative of national demographics.[21] The panelists were provided with in-home scanners to record all of their purchases at the universal product code (UPC) level. These are the prices effectively paid by households, and include discounts (e.g., getting the second item at 50% off). Prices of products are collected from one of two sources. If the store in which the product was bought also reports to Nielsen's store-level survey Scantrack, then the price reported from the store is taken directly. If not, then the household's reported price is used. Einav et al. (2010) test the accuracy of the price data by using a sample of transactions for which they observe both the retailer's price and the household's recorded price. They find that, even though mistakes in price entry do occur, the correlation between the two is reasonably high (88%).

Household income in the Nielsen data is reported in 16 brackets. To compute the share of expenditure in income, we impute income using the average value within each bracket. For the two extreme brackets ($> \$100,000$ and $< \$5,000$), we impute income as \$100,000 and \$5,000, respectively. Household age is computed as the average age between the male and female household head.

We use quartile dummies for average income, growth in average income and average age of consumers at the product-category level as control variables. We compute these variables from the Nielsen data, as expenditure-weighted average of the variable among all households with purchases in a partcular product category.

---

[20]From 2004-2006, sample size is limited to roughly 40,000 households.

[21]We also ran all regressions without using household projection weights to compute category-level inflation rates, and find very similar results.

B.2   Description of Production Data

European production data are obtained from UNIDO at the 4-digit ISIC industry level. Data on U.S. production are obtained at the 6-digit industry level from the BEA. In order to map these production data to 6-digit HS codes, we use a concordance between ISIC and HS obtained from WITS (http://wits.worldbank.org) and a concordance between BEA-industries and HS codes made available by Thibault Fally. To map production data into HS codes, we make a proportionality assumption whereby the export-to-output ratio is assumed to be constant within each ISIC code. We then compute output for each product category using our concordance between HS codes and Nielsen product modules. A similar procedure is used for the US.

B.3   Protocol for Merging Nielsen and Comtrade Datasets

Given the differential classifications used in Nielsen's Consumer Panel and the COMTRADE database, we need to construct a concordance between the two datasets. This was done using a complete list of Nielsen's 1,147 product modules and 5,226 HS 6-digit commodities. A majority of HS codes correspond to intermediate goods and non-barcoded consumer goods, and therefore do not match to Nielsen product modules. The merge was carried out using online tools such as the US Census Bureau's Schedule B Search Engine[22] and the Canadian Importers Database[23], which can identify relevant HS codes for a given product. We aimed to produce the largest number of merged categories possible, while ensuring all relevant Nielsen modules and HS commodities are included within each category. The resulting concordance contains 324 distinct categories, spanning 1,147 Nielsen product modules and 878 HS 6-digit commodities. Table A2 lists the number of categories by merge type (i.e. 1:1, 1:n, m:1, m:n). Our main analytical sample is a subset of 232 categories, due to missing values for Chinese import penetration in Europe (32 categories, almost all fresh foods), missing values for inflation (40 categories that had zero expenditure in one or more years), and finally excluding those with extreme values in our dependent and explanatory variables (20 categories).[24] To illustrate the protocol used for carrying out this exercise, we hereby discuss some examples for each type of merge.

B.3.1   1:1 Merges

This is the simplest type, where a Nielsen product module fits exactly into an HS code, and there is no other product module that fits into the same HS code. For example, to find the HS codes corresponding to the product module "Fresh Apples" (4010), we type in "apples" in the Canadian Importers Database Search Engine, which reveals four items: "080430 - Pineapples

---

[22]https://uscensus.prod.3ceonline.com/

[23]https://www.ic.gc.ca/app/scr/ic/sbms/cid/searchProduct.html?lang=eng

[24]This is defined as either less than the 1st, or greater than the 99th, percentile values.

- Fresh Or Dried"; "080810 - Apples - Fresh"; "081330 - Apples - Dried"; and "200820 - Pineapples Nes - Prepared, Whether Or Not Sugared, Sweetened or Spirited". This yields a 1:1 merge with the HS 6-digit commodity "080810 - Apples - Fresh".

Similar merges are found for a number of other fresh fruits and vegetables, such as oranges, strawberries, carrots, potatoes, and so on. Furthermore, there are also some 1:1 merges for products that are neither food nor drink. For instance, the product module "shelf paper and wall coverings" (7325) and the HS code "wallpaper and similar wall coverings" (481490) results in the category "wallpaper"; while the product module "toaster and toaster oven appliance" (7756) and the HS code "toasters, electrical" (851672) yields the category "toasters". Similarly, the product module "vacuum and carpet cleaner appliance" (7772) and the HS code "vacuum cleaners, including dry and wet vacuum cleaners, with self-contained electric motor" (850910) merge into a single category "vacuum cleaners".

B.3.2   1:n Merges

This is the case where a Nielsen product module has more than one HS 6-digit counterpart. We follow a similar procedure to the one above. For instance, to find the HS codes corresponding to the product module "batteries" (7870), we search through the entire set of HS 6-digit commodities, using both Excel and the Search Engines mentioned above. This reveals the relevant HS codes to be as follows: "Primary cells and primary batteries, manganese dioxide (850610)"; "Primary cells and primary batteries, mercuric oxide (850630)"; "Primary cells and primary batteries, silver oxide (850640)"; "Primary cells and primary batteries, lithium (850650)"; "Primary cells and primary batteries, air-zinc (850660)"; "Primary cells and primary batteries, n.e.s. in 85.06 (850680)". Other similar cases of this type include "printers", "heating appliances", and "fridges and freezers".

B.3.3   m:1 Merges

This is the case where multiple Nielsen product modules map into a single HS 6-digit commodity. It is relatively rare, compared to the other types. One such example is "food processors", where five Nielsen product modules ("RBC food processor and grinder appliance" (6063), "blender appliance" (7757), "mixer appliance" (7758), "juicer appliance" (7760) and "food processor and grinder appliance" (7763)) map into the HS commodity "Food grinders and mixers; fruit/veg. juice extractors, dom., with self-contained elec. motor" (850940). Another example is "creams and cosmetics", where a number of Nielsen product modules ("hand cream"; "hand and body lotions"; "baby care products - lotions", etc.) map into the HS commodity "beauty/make-up preps and preps for the care of the skin, including sunscreen/sun tan preps" (330499).

32

B.3.4　m:n Merges

These are the remaining cases where certain products are classified along different dimensions by Nielsen and COMTRADE. For instance, "tea" is classified into "tea - herbal - instant", " tea - herbal bags", "tea - packaged", "tea - bags", "tea - mixes", "tea - instant" and "tea - herbal packaged" in the Nielsen data, while it is classified into "tea, green (not fermented), whether or not flavoured, in immediate packings of a content not >3kg", "tea, green (not fermented), whether or not flavoured, in immediate packings of a content >3kg", "tea, black (fermented), whether or not flavoured, in immediate packings of a content not >3kg", "tea, black (fermented), whether or not flavoured, in immediate packings of a content >3kg" in the trade data. Similarly, "coffee" is classified into "ground and whole bean coffee", "coffee - soluble flavored", "coffee - soluble" in the Nielsen data, while it is classified into "coffee - not roasted, not decaffeinated", "coffee - not roasted, decaffeinated", "coffee - roasted, not decaffeinated", and "coffee - roasted, decaffeinated" in the trade data.

C　Aggregation

This section gives more details on the aggregation performed in the main part of the paper. We start from the predicted change in the aggregate price index induced by the China trade shock:

$$\widehat{\Delta \log(P)}^{Chn} = \sum_i \omega_i \widehat{\Delta \log(P_i)}^{Chn}$$

We next assume that a foreign trade shock only affects prices if there is a change in the domestic share of expenditure. In other words, the level effect in the second stage is zero:

$$\widehat{\Delta \log(P_i)}^{Chn} = \hat{\beta}_{IV} \widehat{\Delta \log(DSE_i)}^{Chn}$$

where $\hat{\beta}_{IV}$ is the second-stage coefficient estimate, and $\widehat{\Delta \log(DSE_i)}^{Chn}$ is the change in log domestic share caused by the China trade shock.

From the first-stage equation, the predicted decline in the DSE takes the form $\widehat{\Delta \log(DSE_i)} = \hat{\mu} + \hat{\delta} IP_i + \hat{\rho} Z_i$. While the second component ($\hat{\delta} \times IP_i$) captures the degree of a differential change in the DSE across categories due to the China trade shock, the estimated constant $\hat{\mu}$ captures the common change in the DSE across categories, which may or may not be related to supply shocks in China. We thus express the decline in the DSE of category i that is a result of the China trade shock as

$$\widehat{\Delta \log(DSE_i)}^{Chn} = c + \hat{\delta} \times IP_i,$$

where the constant c is unknown. This is assuming that the controls Z_i (U.S. productivity growth, average income and age of U.S. consumers) do not change as a response to Chinese supply changes.

To solve for c, we assume that in the aggregate time series, the China shock is responsible for the same share of the decline in the domestic share of expenditure as in the cross-section among product categories. That is,

$$\widehat{\Delta \log(DSE)}^{Chn} = \kappa \Delta \log(DSE),$$

where κ is the share of the cross-sectional variance in $\Delta \log(DSE_i)$ that is due to supply shocks in China.[25] In our sample, κ takes on a value of 0.535. The decline of the aggregate domestic share of expenditure is $\Delta \log(DSE) = -0.057$. This implies that $\widehat{\Delta \log(DSE)}^{Chn} = -0.030$. That is, roughly half of the aggregate decline in the US domestic share of expenditure is attributed to supply shocks in China.

We then approximate the aggregate decline in the domestic share of expenditure due to the China shock as follows:[26]

$$\widehat{\Delta \log(DSE)}^{Chn} = \sum_i \frac{\omega_i DSE_i}{DSE} \widehat{\Delta \log(DSE_i)}^{Chn} = c + \hat{\delta} \sum_i \frac{\omega_i DSE_i}{DSE} IP_i$$

The expression $\sum_i \frac{\omega_i DSE_i}{DSE} IP_i$ can be computed directly from the data and equals 0.025. Using the previous result $\widehat{\Delta \log(DSE)}^{Chn} = -0.030$ and the estimate $\hat{\delta} = -2.677$, we then obtain a value for the constant term of $c = 0.031$. This allows us to compute the aggregate price decline due to the China shock as

$$\widehat{\Delta \log(P)}^{Chn} = \hat{\beta}_{IV} \sum_i \omega_i \widehat{\Delta \log(DSE_i)}^{Chn}$$

Doing so, we find an aggregate decline in the ideal price index of 2.1 percentage points for the period 2004-2015, or 0.19 percentage points per year.

---

[25]That is, $\kappa = \frac{Var(\widehat{\Delta \log(DSE_i)}^{Chn})}{Var(\Delta(\log(DSE_i)))} = \frac{Var(\hat{\delta} IP_i)}{Var(\Delta(\log(DSE_i)))}$.

[26]The aggregate domestic share of expenditure can be written as $DSE = \sum_i \omega_i DSE_i$. Taking a first-order approximation gives $\Delta \log(DSE) = \sum_i \frac{\omega_i DSE_i}{DSE} \Delta \log(DSE_i)$