

Estimating US Consumer Gains from Chinese Imports*

Liang Bai
University of Edinburgh

Sebastian Stumpner[†]
Université de Montréal

June 5, 2017

Abstract

This paper investigates the size of US consumer gains from the growth in Chinese imports during 2004-12. Drawing on the insight from [Arkolakis et al. \(2012\)](#), we use the change in the US's domestic share of expenditure as a summary statistic for the welfare consequences of foreign trade shocks. We compute this measure for 222 consumer good categories, and estimate its effect on US prices using Chinese exports to Europe as an instrument. Utilizing barcode-level price and expenditure data from AC Nielsen, we construct category-level inflation rates under CES preferences, and find significant effects of Chinese imports on US prices. Comparing the median category in terms of import penetration to a category with no exposure, we find that prices in the median category grew by 0.45 percentage points less per year. The effect is driven by both changes in the prices of existing goods and the entry of new goods. In contrast to the gains from final-good imports, results for intermediate-good imports are inconclusive. We also find no evidence of heterogeneous effects across consumer groups by income or region. A simple benchmarking exercise suggests that Chinese imports led to a 0.29 percentage point per year reduction in the ideal price index.

*We are grateful to David Dorn, Swati Dhingra, Jason Garred, Steve Redding, Andres Rodriguez-Clare, Thomas Sampson, Catherine Thomas, Mike Waugh, Robert Zymek, and seminar/conference participants at CESI HKUST, DEGIT Nottingham, DIW Berlin, Edinburgh, ETSG Helsinki, Glasgow, LSE, Université de Montréal, the Montréal Micro-Macro workshop, RAND, and the World Bank for their helpful comments. Liucija Latanauskaite, Matus Luptak and Tomaz Norbutas provided excellent research assistance.

[†]Bai: University of Edinburgh, School of Economics (liang.bai@ed.ac.uk). Stumpner: Université de Montréal, Department of Economics (sebastian.stumpner@umontreal.ca).

1 Introduction

Recent years have seen a surge in US imports from China. What are the welfare consequences of this import growth for the US? Previous research has shown that it may have accounted for up to one-quarter of the contemporaneous decline in US manufacturing employment, and reduced lifetime earnings of affected workers.¹ Standard economic theory suggests these negative effects should be accompanied by significant consumer gains. In this paper, we estimate the size of these gains using highly disaggregated micro-data, and investigate the channels through which they are realized.

Focusing on the years 2004-2012, we first estimate the effect of Chinese import growth on consumer prices across different product categories, assuming CES preferences and using barcode-level price and expenditure data from AC Nielsen. 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 0.45 ppt less per year. A simple benchmarking exercise implies that the ideal price index declined by 0.29 ppt per year as a result.

We next investigate the channels underlying this effect. Results show that it is driven by both existing goods having lower inflation and the introduction of new goods, suggesting the presence of pro-competitive effects as well as variety gains (Broda and Weinstein (2006), Feenstra and Weinstein (2010)). Moreover, Chinese import penetration leads to an increase in the number of consumed varieties, as well as higher rates of entry and exit. Variety growth is 1 ppt higher in a category with median China trade shock compared to one experiencing no shock. In contrast to the gains arising from final-good imports, results for intermediate-good imports are inconclusive. We also find no evidence for heterogeneous effects across consumer groups by income or region.

Our main estimation equation is consistent with a large class of standard 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. Our identifying variation in the change of domestic expenditure share comes from an instrumental variable, closely related to Autor et al. (2013), namely Chinese import penetration in the five largest European economies.²

We construct constant elasticity of substitution (CES) price indices at the level of 222 con-

¹See Autor et al. (2013) and Autor et al. (2014).

²We use data for Germany, France, the United Kingdom, Italy and Spain, instead of all European countries, due to their availability of sectoral-level output data for the years in our sample.

sumer product categories using data from AC Nielsen’s Homescan Panel, following the methodologies developed by [Sato \(1976\)](#), [Vartia \(1976\)](#) and [Feenstra \(1994\)](#). The data contain information on consumer-goods purchases (both prices and expenditures) from a sample of approximately 60,000 US households at the barcode-level. Defining a product at this level ensures that we do not confound price changes of individual products with changes in the composition of varieties if products were more coarsely defined. Moreover, tracking the prices of individual barcodes allows us to hold product quality constant. While the data do not cover all consumer expenditure on tradable goods (non-barcode products are excluded), we show that Chinese import penetration was even larger in product groups that were absent in the Nielsen data. Hence our results are likely to represent a lower bound of the actual effect.

The data on consumption is combined with data on international trade flows at the HS 6-digit level from UN Comtrade. In particular, we construct a concordance between 1100 Nielsen product modules and 5226 HS 6-digit categories. Finally, we rely on US and European sector-level production data to compute changes in the domestic share of US expenditure as well as Chinese import penetration in Europe.

The main threat to identification is that both prices in the US and imports from China may be driven by demand or supply shocks in the US instead. For instance, a positive US demand shock (or negative US supply shock) could lead to higher US prices and an increase in imports from China. The resulting bias could therefore lead OLS to under-estimate the actual effect. In order to address this concern, we follow [Autor et al. \(2013\)](#) and instrument for the change in the domestic share of US expenditure using Chinese import penetration in Europe. This allows us to isolate the variation due to changes in Chinese supply (e.g. productivity growth, falling trade costs) and not from shocks originating in the US.³

A second threat to identification is that supply shocks in the rest of the world (ROW) may be correlated with those in China. There are two reasons why this is unlikely to be a major concern. First, the rise of China is much more prominent than that of any other country exporting to the US during our study period. Second, we observe that our instrument has a positive effect on the China share of US expenditure, but a negative effect on the ROW share. That is, Chinese imports are displacing both US and ROW products in the US market. If correlated supply shocks between China and the ROW were a concern, we would instead expect

³If demand shocks across the US and Europe were correlated, this would lead to a bias in our estimates. However, we would only overestimate the true effect, if demand shocks across the US and Europe were negatively correlated, which we regard as very unlikely.

a positive relationship between our instrument and the ROW share in US expenditure.

Our results hold in a series of robustness checks. Among others, we control explicitly for US productivity growth at the level of individual product categories, and we construct an alternative instrument to deal with the possibility of correlated supply shocks between China and other countries. We also show that our main results are robust to the exclusion of all food and drinks categories.

This paper fits into a growing empirical literature on the global welfare implications of China’s rapid growth and re-integration into the world economy. In the US context, much of the focus has been on labor market adjustment. The well-known studies of [Autor et al. \(2013\)](#), [Autor et al. \(2014\)](#), and [Pierce and Schott \(2016\)](#) estimate the effects of rising Chinese imports on manufacturing employment and workers’ labor market outcomes. In particular, [Autor et al. \(2013\)](#) find rising imports cause higher unemployment, lower labor force participation, and reduced wages in local labor markets that house import-competing industries. Their baseline estimates suggest that the China trade shock can explain up to one-quarter of the contemporaneous decline in US manufacturing employment between 1990 and 2007. Using individual-level, longitudinal data covering 1992 to 2007, [Autor et al. \(2014\)](#) find that workers more exposed to trade with China exhibit lower cumulative earnings and employment. Consistent with these findings, [Pierce and Schott \(2016\)](#) link the sharp drop in US manufacturing employment after 2000 to the US granting permanent normal trade relations (PNTR) to China.^{4 5}

Our work contributes to this literature by providing the first micro-estimates of US consumer gains, using barcode-level data, from the recent growth in Chinese imports. In doing so, it is closely related to the work of [Amiti et al. \(2016\)](#), who estimate the effect of China’s WTO entry on the US manufacturing price index. Using Chinese firm-product-destination level export data between 2000 and 2006, the authors find a cumulative reduction in the price index of 7.3 percent. They then attribute this effect to two distinct policy changes, one where the US grants PNTR status to China, and the other where China reduces its own input tariffs. In particular, they

⁴Beyond the negative labor market consequences in the US, [Bloom et al. \(2016\)](#) find that Chinese import competition led to both increased technical change within firms and a reallocation of employment between firms, in twelve European countries between 1996 and 2007. Meanwhile, using census data for Brazil, [Costa et al. \(2016\)](#) document faster wage growth in locations benefiting from rising Chinese commodity demand between 2000 and 2010.

⁵Using a quantitative trade model, [Hsieh and Ossa \(2016\)](#) find small spillover effects of China’s productivity growth on welfare around the world between 1995 and 2007 (within 0.2% of cumulative real income growth). Addressing the same question, [Di Giovanni et al. \(2014\)](#) simulate two alternative growth scenarios: a “balanced” one in which China’s productivity grows at the same rate in each sector, and an “unbalanced” one in which China’s comparative advantage sectors catch up disproportionately faster to the world productivity frontier. They find much larger welfare gains from the latter scenario in a sample of 74 countries.

find the latter policy accounting for two thirds of the overall effect, operating through increased imported inputs and higher firm productivity. While their work sheds light on the origins of China’s recent export surge, and the policies ultimately responsible, our approach focuses more on the various channels of adjustment in the US.⁶

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 imports to Europe as an instrument for the change in domestic share of US expenditure.

This paper also fits into the much broader literature on the effects of trade integration on prices and consumer surplus (e.g. [Broda and Weinstein \(2006\)](#), [Feenstra and Weinstein \(2010\)](#), and [De Loecker et al. \(2016\)](#)).⁷ The seminal paper by [Broda and Weinstein \(2006\)](#) was the first to structurally estimate the gains from variety for the US. They consider the entry of foreign varieties, while ignoring the potential exit of previously consumed varieties, something that we explicitly account for in this paper. Another point of departure is the definition of a variety. While it is defined as a combination of an HS 10-digit product category and an origin country in [Broda and Weinstein \(2006\)](#), it is defined at the level of a barcode in our study. In a related paper, [Feenstra and Weinstein \(2010\)](#) analyze both the variety and pro-competitive effects of globalization using firm-level sales data and a translog expenditure function. They find an important role for both of these effects, and that exit of US firms from the domestic market reduces the gains from variety. All of these results are consistent with our findings.

The rest of this paper is organized as follows. Section 2 will describe in detail our data sources and the construction of key variables. Section 3 will outline our empirical strategy and discuss summary statistics. Section 4 will present our main results as well as robustness checks. Section 5 offers some concluding remarks.

⁶Using Italian manufacturing firm-level data on output prices covering the period 1990-2006, [Bugamelli et al. \(2010\)](#) find that import competition from China reduces both price growth and markups. Meanwhile, using a panel of 325 manufacturing industries from 1997 to 2006, [Auer and Fischer \(2010\)](#) show that imports from nine low-wage countries are associated with strong downward pressure on US producer prices.

⁷See [Costinot and Rodriguez-Clare \(2013\)](#) for a recent survey.

2 Data Sources and Methodology

This section first describes the various data sources we employ, before discussing the construction of key variables such as product category-level inflation rates, changes in the US domestic share of expenditure, and the relationship between them.

2.1 Data Sources

We use data from several sources. The two main datasets we employ are international trade flows at the HS 6-digit level, and household purchases and product prices at the barcode (or Universal Product Code, henceforth UPC) level.

The international trade data (i.e. bilateral imports and exports) come from BACI (Base pour l'Analyse du Commerce International), which takes as input the UN's Commodities Trade Statistics database (COMTRADE), at the HS 6-digit level (2002). This is a widely used dataset, covering over 5,200 commodities and 240 countries/regions.

Data on the prices of consumer goods and volume of purchases come from a proprietary dataset from AC Nielsen. This longitudinal dataset contains a sample of around 60,000 US households who continually provide information to Nielsen about their household demographics, what products they buy, as well as when and where the products were bought. A key advantage of the Nielsen data is the availability of price and expenditure information at the barcode level. For our study period, over 1.5 million UPC codes are present. These are grouped into 1,138 different product modules by Nielsen. Some examples include olive oil, pasta-spaghetti, dental accessories, cameras, batteries, and printers.

Nielsen recruits households via mail or online, and offers incentives (such as monthly prize drawings and gift points) to join and to remain active in reporting transactions. Households that do not regularly report their transactions are removed from the sample and new households are added. In doing so, Nielsen aims to make the sample nationally representative in terms of demographics. The panelists were provided with in-home scanners to record all of their purchases (price and quantity) at the UPC level. 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, 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 high (88%).

In order to merge the two aforementioned datasets, we constructed a concordance between the 5,226 6-digit HS codes and the 1,138 Nielsen product modules. A large part of the HS codes are accounted for by intermediate goods (e.g. industrial chemicals, minerals, copper wires, semiconductor devices, etc.), and therefore do not match directly to Nielsen product modules, which only contain final consumer goods. Often times, complex merges are required in which several different HS codes and Nielsen product modules are combined to form one product category.⁸ The merge was carried out with the help of online tools such as the US Census Bureau’s Schedule B Search Engine⁹ and the Canadian Importers Database¹⁰, which can be used to 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 match. The resulting concordance contains 319 distinct categories, spanning 1,138 Nielsen product modules and 846 HS 6-digit commodities. Our main analytical sample is a subset of 222 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 (26 categories).¹¹ Examples of analyzed categories include coffee, bottled water, computer software, vacuum cleaners, and creams and cosmetics.

As mentioned in the introduction, our identification strategy requires the computation of Chinese import penetration in Europe. This, in turn, requires category-level expenditure and output data in Europe. We focus on the five largest European economies - Germany, France, the United Kingdom, Italy, and Spain - both because of their relative size and their availability of such output data for our study period.

In particular, output data for Europe come from the UNIDO Industrial Statistics Database at the 4-digit level of ISIC (INDSTAT 4). It contains information on more than 150 manufacturing sectors and sub-sectors for the period 1990 onwards. Here we follow a two-step procedure to compute expenditure at the category level. We first allocate output from each ISIC code to its corresponding HS codes, using a concordance obtained from WITS.¹² Here we make a

⁸For example, the Nielsen products “Milk chocolate” and “Dark chocolate” were combined with the HS 6-digit codes “Chocolate in blocks weighing less than 2kg” and “Chocolate in blocks weighing more than 2kg” to form one category named “Chocolate.”

⁹<https://uscensus.prod.3ceonline.com/>

¹⁰<https://www.ic.gc.ca/app/scr/ic/sbms/cid/searchProduct.html?lang=eng>

¹¹This is defined as either less than the 1st, or greater than the 99th, percentile values.

¹²<http://wits.worldbank.org>

proportionality assumption where the output-to-export ratio is assumed to be constant within each ISIC code. We then compute expenditure for each of our product categories using our own concordance between HS codes and Nielsen product modules. A similar procedure is used to compute category-level expenditure in the US, using output data at the NAICS 6-digit level from the Bureau of Economic Analysis (BEA).

Finally, in order to study the indirect effects of Chinese imports on US consumer prices via intermediate inputs, we use the detailed direct requirements table from the BEA to measure input-output linkages across the roughly 430 BEA commodity codes.

2.2 Computing Category-Level Inflation Rates

To compute category-level inflation rates, we start with the following non-symmetric CES consumption function. Consumption in product category i at time t is given by an aggregate over different varieties 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:

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

When the set of goods being consumed remains constant over time, inflation can be computed following the formula derived by [Sato \(1976\)](#) and [Vartia \(1976\)](#):

$$\frac{P_{it}}{P_{it-1}} = \prod_k \left(\frac{p_{it}^k}{p_{it-1}^k} \right)^{\omega_{it}^k}$$

That is, inflation is a weighted geometric sum of individual price changes where the weights ω_{it}^k sum to one and can be expressed as a function of expenditure shares $s_{it}^k = \frac{p_{it}^k c_{it}^k}{\sum_k p_{it}^k c_{it}^k}$:

$$\omega_{it}^k = \frac{\frac{s_{it}^k - s_{it-1}^k}{\log(s_{it}^k) - \log(s_{it-1}^k)}}{\sum_k \frac{s_{it}^k - s_{it-1}^k}{\log(s_{it}^k) - \log(s_{it-1}^k)}}$$

As noted in [Feenstra \(2010\)](#), the numerator is a logarithmic mean of the market shares s_{it}^k and

s_{it-1}^k and is therefore always situated between them.

With a changing set of goods, [Feenstra \(1994\)](#) shows that inflation can be written as:

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

where I_{it}^E denotes the set of “staying” goods that are present in both $t-1$ and t . The full inflation rate is now the product of the Sato-Vartia formula and a correction factor that measures the contribution of the extensive margin (i.e. the effect of new and disappearing varieties).

The term λ_{it} is defined as the fraction of expenditure at time t that goes towards staying goods. Intuitively, when the share of expenditure going to staying goods is declining, this must mean that entering varieties are more competitive than exiting varieties. This in turn contributes to a reduction in the cost of living. The extent to which the introduction of new goods affects the inflation rate depends on the elasticity of substitution σ . In particular, if goods are highly substitutable, the effect on inflation will be more muted.

Given our sample, there are two alternative ways of computing the intensive margin inflation rate. First, we can restrict the sample to include only those goods that appear in every year between 2004 and 2012. We call these goods “long stayers.” Second, we can restrict the sample to include all goods that appear in any consecutive pair of years. We call these goods “short stayers.” In other words, the second way of computing the intensive margin involves updating the definition of “old” and “existing” goods each year.

It should be noted that the price of each good p_{it}^k is the average unit price for that UPC code among all Nielsen purchases in a given year, while the expenditure on each good is the total among all such purchases. The expenditure shares of each good in a given year s_{it}^k is, in turn, computed as the ratio of total expenditure on UPC code k in year t over total expenditure on the product category i it belongs to in year t .

Finally, to compute inflation rates, we assume an elasticity of substitution between product varieties equal to 5 in our baseline specification, which is roughly the mean of the elasticities estimated by [Broda and Weinstein \(2006\)](#) for product groups at a similar level of aggregation. In robustness exercises, we also consider different values of σ , and find the results to be qualitatively unchanged.

2.3 Domestic Share of Expenditure (DSE) and Inflation

How can we assess the effects of a positive supply shock in China on prices in the US? [Arkolakis et al. \(2012\)](#) show that in a large class of trade models, including [Anderson \(1979\)](#), [Eaton and Kortum \(2002\)](#), [Krugman \(1980\)](#), and [Melitz \(2003\)](#), 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 in these models 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:

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

where DSE denotes the domestic share of expenditure and θ is the trade elasticity, that is, the elasticity of trade flows with respect to trade costs.¹³ Since the change in domestic share of expenditure 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.

To compute the domestic share of expenditure for each product category i in year t , we first compute total US expenditure on product category i as total production plus imports minus exports.¹⁴ We then measure DSE as the fraction of expenditure (X_{it}) that does not go towards imports (M_{it}):

$$DSE_{it} = \frac{X_{it} - M_{it}}{X_{it}}$$

3 Empirical Strategy and Summary Statistics

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 between 2004 and 2012. Following equation 1, we relate the log change of the price index to the log

¹³In the model by [Eaton and Kortum \(2002\)](#) and in a Melitz-model with a Pareto distribution, the trade elasticity is related to the parameter governing the dispersion of productivities.

¹⁴We compute production at the category-level using production data for 6-digit NAICS industries, and using a crosswalk from NAICS to HS.

change in the domestic share of expenditure in each product category:

$$\Delta \log(P_i) = \alpha + \beta \Delta \log(DSE_i) + \epsilon_i \tag{2}$$

In subsequent specifications we decompose the effect on inflation into intensive and extensive margins. We then turn our attention to the effect on category-level expenditure growth, the change in number of goods purchased, as well as product entry and exit.

3.1 Identifying Trade Shocks

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 supply shocks originating from China. For instance, a positive US demand shock should lead to higher US prices and may also affect the US domestic share of expenditure. Likewise, a positive US supply shock would tend to lower US prices and increase the domestic share of expenditure, as US products become more competitive. This type of bias would therefore lead OLS to underestimate the true effects. To address this concern, we follow [Autor et al. \(2013\)](#) and use Chinese import penetration in Europe as an instrument, defined as:

$$IP_i = \frac{\Delta M_i^{Europe,Chn}}{X_i^{Europe}},$$

that is, as the change in European imports from China during our sample period, normalized by initial European expenditure. Intuitively, if the growth in US imports from China during 2004-2012 is driven either by productivity growth in China or a reduction in trade barriers as a result of China’s accession to the WTO, we should observe an increase in Chinese exports to other developed countries, such as those in Europe.

One potential concern with such a strategy is that demand or supply shocks may be correlated between the US and Europe. However, this would only be problematic if shocks were *negatively* correlated. If, for instance, demand shocks between the US and Europe were positively correlated, as we would expect, one should observe an increase in Chinese import penetration in Europe to be associated with an increase in US prices. This is the opposite of what we find in the data.

A second threat to identification is that supply shocks in the rest of the world (ROW) may be correlated with those in China. This could potentially lead us to overestimate the role of Chinese imports in determining US prices. We believe this is unlikely to be a major concern.

First, the rise of China in world trade and among US imports is much more prominent than that of any other country exporting to the US during the time period under study. Figure 1 shows the evolution of US import shares by origin region. While the shares of EU countries, Canada and Mexico, and other emerging economies has either remained stable or declined over time, China’s share in US imports has increased by almost 10 percentage points during 2003-12. Second, as discussed in the next section, we observe that while our instrument has a positive effect on the China share of US expenditure, it has a negative effect on the ROW share. That is, Chinese imports are displacing both US and ROW products. If correlated supply shocks between China and the ROW were significant, we would instead expect a positive relationship between our instrument and the US expenditure share on ROW products.

With these points in mind, we run the following first-stage regression:

$$\Delta \log(DSE_i) = \gamma + \delta \cdot IP_i + \nu_i \tag{3}$$

where IP_i is the China import penetration measure for Europe in category i .

3.2 Summary Statistics

Table 1 shows summary statistics for some of our key dependent and explanatory variables. Assuming σ takes on a value of 5, the average cumulative inflation rate across product categories during 2004-2012 was 1.8%, or 0.2% per annum. Using only goods that are present in consecutive years, and therefore disregarding the contribution of new varieties, the average cumulative inflation rate becomes 13.4% during 2004-2012, or 1.5% per annum. There is a large amount of variation across product categories for each of these measures, with standard deviations of 26 and 20 percentage points respectively. Product categories with the five lowest inflation rates were electro-thermic appliances, babies’ clothing, flashlights, heating appliances, and cooking appliances; while those with the five highest were meslin flour, cigarettes, frozen orange juice, shortening, and sugar.

At the same time, the US domestic share of expenditure declined in the majority of consumer goods categories. The average decline equals 0.16 log points. The median category saw a decline of 0.042 log points. There is also a large amount of variation in the change of domestic share across categories, with a standard deviation of 0.39 log points. The distribution of the log change in domestic share of expenditure is illustrated in Figure 4. Product categories with the

five largest declines in domestic share were female contraceptives, ammonia, coffee/tea makers, toasters, and vacuum cleaners; while those with the five largest increases were dairy spreads, frozen beef, babies' clothing, sparkling wine, and canned beef.

Finally, Chinese imports into Europe also grew rapidly. For the average category, the increase in European imports from China from 2004-12 equaled 4.9% of total expenditure in 2004, with a standard deviation of 8.9 percentage points. Product categories with the five highest import penetrations were pet accessories, flashlights, other household appliances, food processors, and toasters; while those with the five lowest import penetrations were grated cheese, fresh cheese, other cheese, cider and beer, and cornmeal.

Given the Nielsen Homescan panel's specific coverage of consumer goods, it is important to highlight some similarities and differences between the HS codes that are included in our analysis and those that are not. First, the HS codes that are included account for 18.3% of total US imports, and 19.2% of US imports from China, in 2004. These fractions remain largely unchanged during our sample period, reaching 20.0% and 17.2% in 2012 respectively.

In terms of the US domestic share of expenditure, the HS codes that are included in our analysis experienced a decrease from 82.0% to 77.8% (Figure 2), while the entire set of HS codes (including intermediate goods) experienced a decrease from 73.0% to 69.8% (Figure A.1). While the products included in our analysis have a larger domestic share of expenditure, the magnitude of the decline is very similar across the two groups. In terms of the China trade shock, the HS codes that are included in our analysis saw the China share of expenditure increase from 2.7% in 2004 to 4.6% in 2012 (Figure 2), while the entire set of HS codes saw the China share increase from 3.8% to 7.2% over the same period (Figure A.1). This suggests that the products in our sample had slightly *less* exposure to the increase in Chinese imports, compared to the products that are outside of our sample.

4 Results

In this section we discuss the estimated effects of import growth from China on inflation and product variety in the United States between 2004 and 2012. We will examine both the intensive and extensive margins of inflation, separately identifying the contribution of new varieties. We will test whether any effect on inflation differs across income groups as well as regions, before investigating the role of trade in intermediate goods. Finally, we perform several robustness

exercises, and report the findings from a simple benchmarking exercise.

Before discussing the results on inflation, Table 2 (Panel B) presents the first-stage of the IV regression. Columns 3 and 4 are unweighted and weighted specifications respectively. The main coefficient on China’s import penetration in the five-largest European economies is strongly negative and highly statistically significant across both specifications. Intuitively, the product categories that saw a larger increase in Chinese import growth to Europe also saw a larger decrease in the share of US domestic expenditure. The first stage also yields a sufficiently strong F-statistic of 24.

Table 3 decomposes this first-stage relationship further by looking separately at the change of the relative China share in US expenditure, as well as the rest-of-the-world share. In particular, the log change in the domestic share of expenditure can be approximated as follows:

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

where *CSE* and *RSE* denote the China share of US expenditure and the ROW share of US expenditure, respectively. That is, the log change in the domestic share of expenditure equals the relative decline in the China share plus the relative decline in the ROW share. Table 3 looks at the effect of our instrumental variable on the two terms on the right-hand side. Here we observe that a rise in Chinese import growth to Europe is positively correlated with the change in the relative China share of US expenditure, and negatively correlated with the change in the ROW share. Taken together, this implies Chinese imports are displacing both US and ROW products.

4.1 Effect on Inflation Rates

Table 2 (Panel A) presents our estimated effects of changes in the domestic share of expenditure on US inflation. Columns 1 and 2 are simple OLS estimates, while columns 3 and 4 are results when the US domestic share of expenditure is instrumented by Chinese import penetration in Europe. As discussed in the previous section, if demand or supply shocks in the US affect both prices and imports from China, then least squares would likely underestimate the true effect. We see some suggestive evidence of this in our results. While the OLS and IV coefficient estimates are both positive and highly statistically significant, the IV estimates are larger in magnitude. For instance, in the unweighted IV specification (Column 3), comparing a product category with

median change in the domestic share of expenditure (decline by 4.2%) to one with no change, prices in the median category grew by 3.66 ppt less cumulatively over our study period, or 0.45 ppt less per year. Using estimates from the weighted regression yields a slightly larger effect.

4.2 Intensive vs. Extensive Margins of Inflation Rates

We next analyze different sources of this reduction in the cost of living. In particular, we are interested in the role of intensive margin effects, i.e. the effect of Chinese imports on the prices of previously existing goods, and potential gains from variety as in [Broda and Weinstein \(2006\)](#), i.e. consumer gains arising from changes in the set of available varieties.

4.2.1 Intensive Margin Effects

We conduct three separate tests to evaluate the role of intensive margin effects. In the first two tests, we focus directly on the intensive margin of our cost-of-living inflation formula. In particular, we consider two different ways of computing this margin, depending on the definition of “existing” goods. First, we compute the year-on-year inflation rate, using only the set of goods that are consumed in both periods. We then update the set of “existing” goods for each consecutive pair of years, before compounding these annual inflation rates into a cumulative value for the entire sample period. This is the procedure followed in constructing our “short-stay” inflation measure. This has the advantage that it includes all goods that are observed for consecutive periods, but the disadvantage that it also includes price changes of new (and potentially foreign) goods that only entered at a later time. As a second test, we re-compute the intensive margin using only the set of goods which are consumed in all sample years (called the “long-stay” inflation measure). Roughly 40% of the 500,000 goods sold in 2004 fall into this category. This approach likely identifies more cleanly the intensive margin effects, since a substantial share of the entering products might be imported from China.

Finally, competitive pressure might not only result in lower prices of previously consumed varieties, but also in lower sales and exit. In order to estimate these effects, we shift our attention to the level of individual barcodes. In particular, we focus on the set of all goods available in 2004, and track their prices, expenditure, market share, and potential exit through time.

Results for the first two tests can be found in Columns 1-4 of Table 4 (Panel A). All results show that intensive-margin price growth was lower in product categories facing stronger import competition from China. While the results on the long-stay inflation are smaller in magnitude

than those for the short-stay measure, they remain highly significant. Comparing the median category in terms of the reduction in domestic share of expenditure to a category without change, prices of existing goods declined by roughly 1 ppt more during our study period.

We next shift our attention to the level of individual products. Table 5 presents results from the third test discussed above, using barcode-level regressions. In particular, we estimate the following specification:

$$y_{ik} = \alpha + \beta \Delta \log(DSE_i) + \epsilon_{ik} \quad (4)$$

for category i and barcode k . y_{ik} is the annualized log change for a particular outcome (price, market share, expenditure) from 2004 until the last date in which the product is observed.¹⁵ 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 as outcome a dummy variable that equals one if the last year in which the product is observed is before 2012.

Columns 1-4 of Panel A (unweighted) and Panel B (expenditure-weighted) present results for roughly half a million products that were available in 2004. The results suggest that products in markets that were subject to more intense import competition from China (i) reduced their prices by more, (ii) experienced a drop in sales and market share, and (iii) were more likely to exit the market.¹⁶ Columns 5-7 of both panels show that the same results on prices, sales, and market share also hold for the subset of just over 200,000 goods that are still available in 2012. The magnitude of the price decline is more muted for the long stayers, indicating that products that eventually exit exhibit a particularly strong price decline in response to import penetration. In terms of magnitude, the results are also in line with the ones obtained at the category-level. Column 5 of Table 5 and columns 3 and 4 of Table 4 all show results for the same sample of goods, namely long stayers, with very similar coefficients.

How can these intensive margin effects be interpreted? Since we do not directly observe markups, but only output prices, our results are consistent with two alternative interpretations: The existence of pro-competitive effects (i.e., reductions in markups) or a differential change in factor costs across consumer good categories. In a model with labor as only factor of production and perfect labor mobility across sectors, wages would equalize, and the differential decline in

¹⁵That is, for a product that exits at year T and for outcome x , we compute $y_{ik} = \frac{1}{T-2004} (\log(x_{ik,T}) - \log(x_{ik,2004}))$. For a product that does not exit, we have $T = 2012$.

¹⁶Given our previous result of no differential change in category-level expenditure, a reduction in expenditure on old goods must also mean that their market share declines. Columns 3 and 7 show that this is indeed the case.

prices may be interpreted as evidence for pro-competitive effects. However, since we cannot rule out that factor costs changed differentially across product categories, our results are only suggestive of pro-competitive effects.¹⁷

4.2.2 Extensive Margin Effects

Next, we test for consumer gains arising from changes in the set of available varieties. To do so, we estimate the effect of import penetration on the Feenstra (1994) correction factor across categories. The results are shown in Columns 5-6 of Table 4, and suggest substantial variety gains. The extensive margin makes up roughly 40% of the overall adjustment, while the intensive margin is responsible for the remaining 60%.¹⁸

These results are very closely related to those from the seminal paper by [Broda and Weinstein \(2006\)](#), who look at the contribution of variety growth through international trade for the reduction in the cost of living in the US. They define a product as a ten-digit HTS category and a variety as a product-country combination, and show that from 1972 to 2001, US consumers gained from a surge in available varieties of consumer goods.

Our extensive margin analysis differs from their work in two respects. First, we define a variety at the finest possible level, namely that of a barcode. This more detailed definition allows us to better disentangle intensive margin from extensive margin effects. In particular, by defining a variety as a combination of a ten-digit HTS category and an origin country, [Broda and Weinstein \(2006\)](#) might understate the role of variety growth for welfare.¹⁹ Second, [Broda and Weinstein \(2006\)](#) only observe the increase in imported varieties, but not the potential exit of domestic varieties that might result from more import competition. As shown in the next section, we find that exit of previously consumed varieties is an important margin of adjustment to an increase in foreign competition. Crucially, however, variety gains still play a key role even after accounting for the exit of previously consumed varieties.

Given heterogeneous product quality, variety gains as defined above can arise with or without

¹⁷Our results are also consistent, however, with the findings in [Feenstra and Weinstein \(2010\)](#), who provide evidence for the existence of pro-competitive effects in the U.S. from international competition by structurally estimating the effect of globalization on markups using firm-level data.

¹⁸This decomposition is computed as the ratio between the coefficient estimate from Column 5 of Table 4 (0.328), corresponding to the extensive margin, and that from Column 3 of Table 2, Panel A (0.836), corresponding to the overall adjustment.

¹⁹For instance, observing an increase in US imports of French red wine might be due to more varieties of French red wine becoming available. Unless the initial level of imports was zero, [Broda and Weinstein \(2006\)](#) would not count this increase as a variety gain.

an increase in the total number of consumed varieties. We therefore perform additional tests on whether Chinese imports have led to a more crowded product space. Columns 3 and 4 of Table 4 (Panel B) report results for the number of different product varieties consumed. Here we find that variety growth is 1 ppt higher in a category with median change in the domestic share of expenditure compared to one experiencing no change. However, the effect on the number of varieties does not seem to go through market size. When we estimate the effect of the China trade shock on category-level expenditure, we do not find any effect (Columns 1 and 2). This suggests that the entry of foreign goods leads to a decline in sales per product. The absence of a differential effect on expenditure together with the evidence on different inflation rates also suggests that the elasticity of substitution across categories is close to one.

Finally, we study the channels behind the variety gains by looking at average entry and exit rates across product categories.²⁰ Results are shown in Columns 5-8 of Table 4 (Panel B). Unsurprisingly, categories that experience stronger Chinese import exposure (and thereby lower domestic share of expenditure) are characterized by higher entry rates. At the category-level, the entry of new Chinese products leads to more exit of other varieties, which confirms our previous results at the barcode-level.

4.3 Heterogeneous Effects by Income Group and Region

Another useful feature of the Nielsen Homescan data is the availability of household characteristics such as annual income and area of residence, allowing us to estimate heterogeneous treatment effects for different income groups and regions within the US.²¹ To do so, we first divide individuals into five similarly-sized groups according to their reported annual household income. These correspond to: (i) less than \$30,000, (ii) between \$30,000 and \$50,000, (iii) between \$50,000 and \$70,000, (iv) between \$70,000 and \$100,000, and (v) above \$100,000, respectively. We then compute category-level inflation rates during 2004-12 separately for each of these five groups. Doing so allows the set of goods consumed within each category, as well as relative weights across categories, to vary by income group.

To investigate heterogeneous effects across regions, we follow a similar procedure, whereby we

²⁰For each category and year, we define the fraction of entrants as the share of currently consumed goods that were not consumed in the previous period. Similarly, we define the fraction of exiters as the share of all currently consumed goods that are not consumed in the following year.

²¹In a similar exercise, [Faber \(2014\)](#) finds that high-income households benefitted more from trade liberalization than low-income households during Mexico's NAFTA accession.

first group individuals according to their place of residence, into one of four US census regions. These are: (i) the Northeast, (ii) the Midwest, (iii) the South, and (iv) the West, respectively. The choice of these groupings is made in order to strike a balance between the representativeness of our sample within each group on the one hand, and the degree of heterogeneity on the other.²² It is worth noting that there are potentially important dimensions of heterogeneity that we are unable to study. For instance, households may differ in terms of the fraction of their expenditure on tradables vs. nontradables, as in [Fajgelbaum and Khandelwal \(2016\)](#). In such scenarios, poorer households may benefit more from growing imports, if more of their spending is concentrated on tradable goods.

With that in mind, the results are shown in Table 6. The first notable finding is that our earlier aggregate result on inflation holds consistently across income groups (Panel A) and regions (Panel B). The second main finding is that these effects do not differ significantly across these groups. In terms of the point estimates in Panel A, there is some suggestive evidence that the effect size decreases monotonically with income, indicating that poorer households may have benefited more from the growth in Chinese imports than richer households. However, this difference is not statistically significant (p-value of 0.44 when comparing the highest-income group to the lowest-income group). In Panel B, the coefficient estimate for the South is larger than that for the West, with the Northeast and Midwest taking on values in between. These differences, however, are not statistically significant at conventional levels, suggesting that the different sub-groups analyzed here have benefited to largely the same extent.

We can further decompose the result for income groups as follows. First, aggregate inflation might differ for poor and rich households because of differences in weights on the various consumer good categories. While for each income group, the category-level expenditure share (out of all expenditure) is mildly negatively correlated with our instrument (that is, categories with high Chinese import penetration tend to be low-expenditure categories), this is not differentially the case between poor and rich households.

The second reason is that there may be differences in category-level inflation rates across the two groups. However, while the previous results already indicate no significantly different effect on category level inflation rates among the groups, we nevertheless further decompose this effect here. In particular, differences in category-level inflation rates can in turn arise for one of

²²The Nielsen Homescan data is designed to be a nationally representative sample. Therefore we risk losing representativeness when the number of sub-groups, either by income or region, becomes large.

three reasons: Differences in the composition of products consumed, differences in price growth for the same product, and extensive margin differences. To this end, we derive the following exact decomposition of differences in category-level inflation rates across income groups g and h (dropping the subscript i for a category for simplicity):

$$\begin{aligned} \Delta \log(P_g) - \Delta \log(P_h) &= \sum_{k \in I_t^E} (\omega_{gt}^k - \omega_{ht}^k) \Delta \log(p_t^k) + \sum_{k \in I_t^E} \omega_t^k (\Delta \log(p_{gt}^k) - \Delta \log(p_{ht}^k)) \\ &\quad + \frac{1}{\sigma - 1} (\Delta \log(\lambda_{gt}) - \Delta \log(\lambda_{ht})) \end{aligned}$$

where $\Delta \log(p_t^k)$ is average price growth of item k and ω_t^k is the average share of item k across the two groups.²³ The set I_t^E is the union of the sets of "staying" goods for the two groups.²⁴

To test for the role of these components, we divide households into two groups by income (high and low) with the cutoff being an annual income of \$70,000. Results for this decomposition are in table 7. While the aggregate difference of the effect is not significantly different between the two consumer groups, the slightly larger effect for low-income households (albeit not significant) seems to be associated to the composition term. That is, within categories, a larger share of low-income household expenditure is concentrated on goods with a relatively strong price decline. In contrast, there is no evidence for a differential change in prices for the same product across low and high income groups, or for a differential effect of the extensive margin.

4.4 Effect of Intermediate Inputs

So far in the analysis we have only been concerned with the role of import penetration in final consumption goods. In practice, a great deal of international trade is in intermediate goods, and such flows can in theory be an additional channel through which imports can affect price and variety growth. We now investigate this channel by constructing a measure, again at the product category level, that aims to capture import penetration in the intermediate inputs (or IPII) associated with that category.

To do so, we make use of input shares from the BEA's input-output tables, map them into the product categories used in our analysis, and compute the measure of IPII as the weighted average of input-level changes in the domestic share of expenditure, with the weights γ_{ij} given

²³We have $\Delta \log(p_t^k) = \frac{\Delta \log(p_{gt}^k) + \Delta \log(p_{ht}^k)}{2}$ and $\omega_t^k = \frac{\omega_{gt}^k + \omega_{ht}^k}{2}$.

²⁴If a product is only consumed in both periods by group h , but not by group g , then $\omega_{gt}^k = 0$.

by input shares:

$$IPII_i = \sum_{j=1}^S \gamma_{ij} \Delta \log(DSE_j)$$

We then instrument for $IPII$ using an analogous measure that replaces industry j 's log DSE change by the corresponding Chinese import penetration into Europe. Therefore, we have two measures of import penetration, one for final goods, and one for intermediate inputs, for each product category, and we can analyze their effects on inflation separately.

Table 8 presents the IV results, where we include $IPII$ as an additional explanatory variable. Compared to the results from Table 2 (Panel A), the effect of $\Delta \log(DSE)$ in final goods, induced by the China trade shock, remains entirely unchanged. Meanwhile, an increase in $IPII$, again induced by the China trade shock, does not have a statistically significant effect on US inflation. The coefficient estimates for $IPII$ are similar in size to those of $\Delta \log(DSE)$ in final goods, but much less precisely estimated, despite of a strong first-stage F-statistic of 12.1.

4.5 Robustness

In this subsection we present results from several robustness checks. First, we directly control for US productivity growth by including the log change in US real output per worker as an additional explanatory variable in our inflation regressions. Data for the growth in real output at the six-digit industry level are taken from the BEA, while data for employment are obtained from County Business Patterns.²⁵ The results are in shown in Columns 1 and 2 of Table 9. Here we see the expected negative coefficient on the US productivity variable (although statistically insignificant), but more importantly controlling for US productivity growth does not change our estimates of the effect of import penetration on prices.

In a second robustness exercise, we deal explicitly with potential supply shocks from third-party countries that are also exporting to the US. In particular, we focus on supply shocks from Canada and Mexico, the two countries with the highest growth in exports to the US after China during our sample period. The concern is that Chinese supply shocks may be correlated across categories with supply shocks in Canada and Mexico, so that our estimates pick up their joint effect. While the decline in the share of US expenditure on rest-of-the-world products caused by Chinese import penetration (Table 3) should alleviate such concerns, we nevertheless deal

²⁵We focus on real output per worker, instead of TFP, given the absence of good capital stock data at the six-digit industry level.

explicitly with this issue here. In particular, we construct an alternative instrumental variable by isolating variation in Chinese supply shocks that is orthogonal to potential supply shocks in Canada and Mexico. To do so, we define this alternative IV as the residual from a regression of Chinese import penetration into Europe on Mexican and Canadian import penetration into Europe, and then re-run our estimation. Results are presented in Columns 3 and 4 of Table 9, and show the coefficient on $\Delta \log(DSE)$ to be virtually unchanged.

We further check for the sensitivity of our results to the exclusion of certain product categories from the sample. In Columns (5) and (6) of Table 9, we re-estimate the IV regression for cumulative inflation, this time excluding all categories that are either food or drinks. The resulting coefficient in the weighted specification falls by 25% compared to the full sample, but remain statistically significant at the 5% level.

Table A.1 presents results from another robustness check, using alternative values of σ in computing cumulative inflation rates. Specifically, we experimented with values of 2, 3 and 10 (while the main regressions used a value of 5). In all cases our results are essentially unchanged, although the point estimates are necessarily smaller for larger values of σ , as the contribution of new goods becomes more muted.

4.6 Benchmarking the Impact of China Trade on US Consumer Prices

Using our cross-sectional estimates, we can carry out a simple aggregation exercise, where we compute the effect of changes in the domestic share of expenditure, due to Chinese import growth, on an overall consumer price index.

Given Cobb-Douglas preferences across categories (consistent with the result that expenditure shares remain unaffected), the change in aggregate price index is given by:

$$\Delta \log(P) = \sum_i \omega_i \Delta \log(P_i)$$

where ω_i is the relative weight of category i in the overall consumption bundle. The predicted change in the aggregate price index induced by supply shocks in China is then the weighted average of the category-level changes:

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

Consistent with theory, we 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.

In line with the first stage equation, we have $\widehat{\Delta \log(DSE_i)}^{Chn} = c + \hat{\delta} IP_i$, where the constant c is unknown, $\hat{\delta}$ is the first-stage coefficient estimate, and IP_i the instrumental variable.²⁶ To solve for c , we assume that in the aggregate time series, the China shock is responsible for the same decline in the domestic share of expenditure as in the cross-section among product categories. That is, we assume that in the aggregate,

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

where κ is the share of the cross-sectional variance in $\Delta \log(DSE_i)$ that is due to supply shocks in China:

$$\kappa = \frac{Var(\widehat{\Delta \log(DSE_i)}^{Chn})}{Var(\Delta(\log(DSE_i)))} = \frac{Var(\hat{\delta} IP_i)}{Var(\Delta(\log(DSE_i)))}$$

In our sample, κ takes on a value of 0.22, same as the R-squared of the first-stage regression, and $\Delta \log(DSE) = -0.047$. This implies that $\widehat{\Delta \log(DSE)}^{Chn} = -1.06\%$. That is, roughly a quarter of the aggregate decline in the U.S. domestic share of expenditure is attributed to supply shocks in China. Since supply shocks in China tend to reduce the domestic share of US expenditure by less than they increase the China share of US expenditure (as reported in section 4), the increase of the China share in US expenditure should provide an upper bound for the decline of $\widehat{\Delta \log(DSE)}^{Chn}$. Since the China share of US expenditure has increased by roughly 2 ppt (figure 2), a value of -1.06ppt for $\widehat{\Delta \log(DSE)}^{Chn}$ seems reasonable.

We then approximate the aggregate decline in the domestic share of expenditure due to the China shock as follows:²⁷

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

²⁶Since the US domestic share of expenditure can change for reasons that are unrelated to supply shocks in China, the constant c potentially differs from the estimated coefficient $\hat{\gamma}$.

²⁷This comes from a first-order approximation of the aggregate log decline in the domestic share of expenditure.

In our data, $\hat{\delta} = -1.98$, and $\sum_i \frac{\omega_i DSE_i}{DSE} IP_i = 0.023$. This delivers a value for the constant term of $c = 0.035$. We can then compute the aggregate price decline due to the China shock as

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

Doing so, we find an aggregate decline of the ideal price index of 2.27 percentage points for the period 2004-2012, or 0.29 percentage points per year.

5 Concluding Remarks

In this paper we have tried to estimate the size of US consumer gains from the recent growth in Chinese imports by adopting a micro-econometric approach. Using barcode-level price data from AC Nielsen and exploiting cross-product variation in import penetration between 2004 and 2012, we find that prices declined by more in product categories with higher Chinese import penetration: comparing a category with median China trade shock to one with no change, prices in the median category grew by 0.45 ppt less per year. A simple benchmarking exercise suggests that inflation was 0.29 ppt lower per year due to Chinese import penetration.

We addressed the potential endogeneity of imports by using Chinese exports to European countries as an instrument. We find that the effect of import penetration on prices is driven by both the intensive and extensive margins, suggesting the presence of pro-competitive effects for old goods as well as variety gains from new goods. While Chinese import penetration leads to more entry of new goods, thus benefitting consumers (as in [Broda and Weinstein \(2006\)](#)), this effect is somewhat muted due to the exit of previously consumed varieties. Still, we find evidence for net gains in the number of consumed varieties and an overall positive contribution of the extensive margin. In contrast, results for the effect of increased intermediate goods imports on final goods prices are inconclusive. We also test for heterogeneous effects of Chinese import penetration across different consumer groups by income and region, but do not find any differential effect of Chinese import growth.

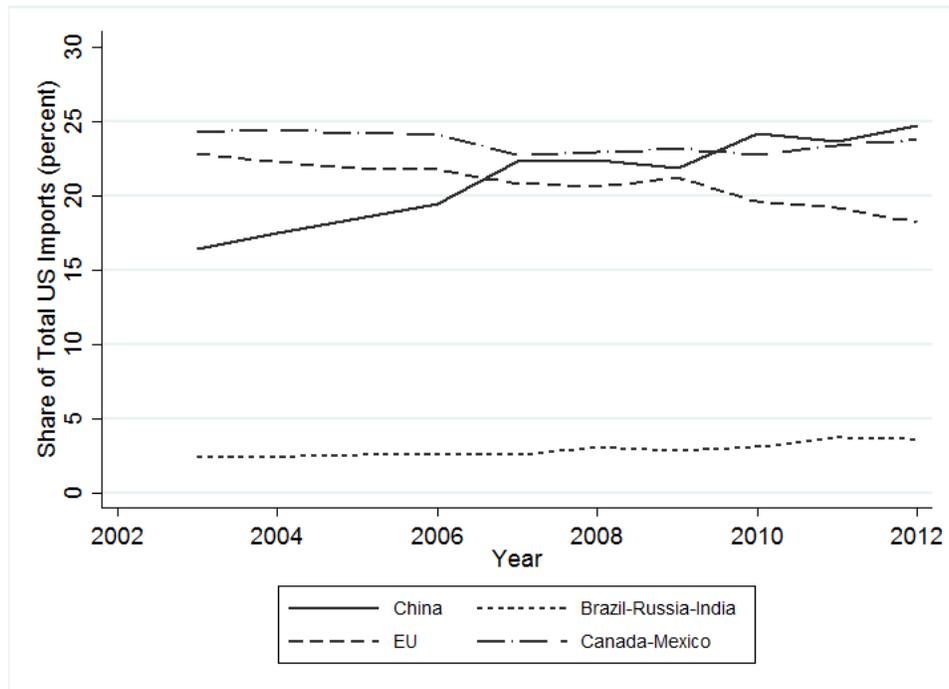
Taken together, these results suggest substantial gains to US consumers from the recent growth in trade with China. These have the potential to offset at least some of the negative labor market consequences due to import competition.

References

- Amiti, Mary, Mi Dai, Robert Feenstra, and John Romalis**, “How did China’s WTO entry benefit U.S. consumers?,” *mimeo, FRBNY*, 2016.
- 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.
- Broda, Christian 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.
- Bugamelli, Matteo, Silvia Fabiani, and Enrico Sette**, “The pro-competitive effect of imports from China: an analysis of firm-level price data,” 2010.
- Costa, Francisco, Jason Garred, and João Paulo Pessoa**, “Winners and losers from a commodities-for-manufactures trade boom,” *Journal of International Economics*, 2016, *102*, 50–69.
- Costinot, Arnaud and Andres Rodriguez-Clare**, “Trade Theory with Numbers: Quantifying the Consequences of Globalization,” *Handbook of International Economics*, 2013, *4*.
- 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.

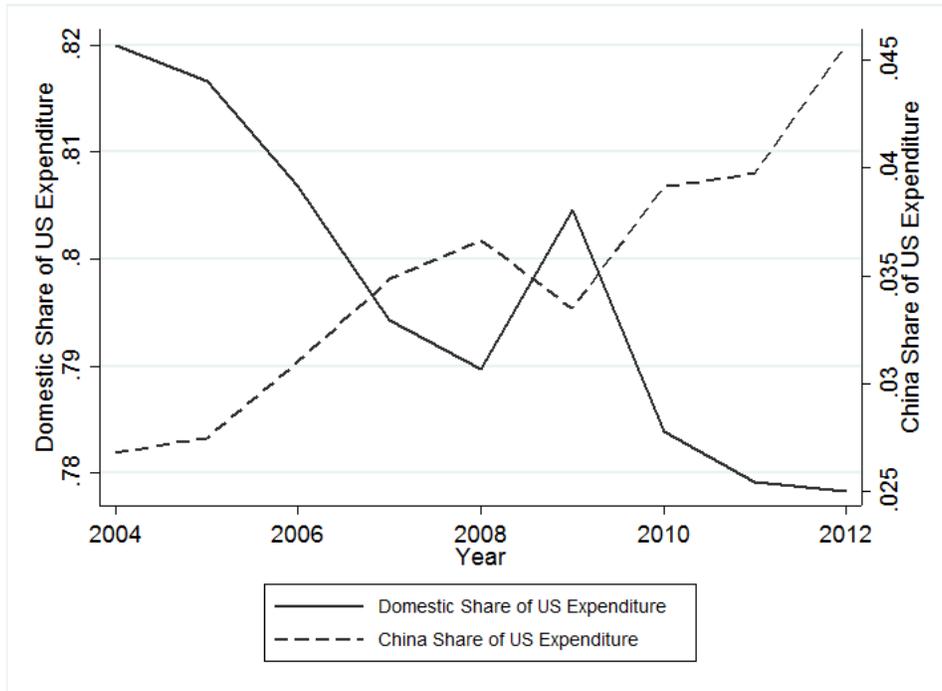
- Faber, Benjamin**, “Trade Liberalization, The Price of Quality, and Inequality: Evidence from Mexican Store Prices,” *mimeo, UC Berkeley*, 2014.
- Fajgelbaum, Pablo D. and Amit K. Khandelwal**, “Measuring the Unequal Gains from Trade,” *Quarterly Journal of Economics*, 2016, *131* (3), 1113–1180.
- Feenstra, Robert C.**, “New Product Varieties and the Measurement of International Prices,” *American Economic Review*, 1994, *84* (1), 157–177.
- , “Measuring the gains from trade under monopolistic competition,” *Canadian Journal of Economics*, 2010, *43* (1), 1–28.
- **and David E. Weinstein**, “Globalization, Markups, and U.S. Welfare,” *NBER Working Paper No. 15749*, 2010, *43*.
- 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.
- Hsieh, Chang-Tai and Ralph Ossa**, “A global view of productivity growth in China,” *Journal of International Economics*, 2016, *102*, 209–224.
- Krugman, Paul**, “Scale Economies, Product Differentiation, and the Pattern of Trade,” *American Economic Review*, 1980, *70* (5).
- Loecker, Jan De, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik**, “Prices, markups, and trade reform,” *Econometrica*, 2016, *84* (2), 445–510.
- 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.
- Vartia, Y.O.**, “Ideal Log-Change Index Numbers,” *Scandinavian Journal of Statistics*, 1976, *3*, 121–126.

Figure 1: Changing Composition of US Imports



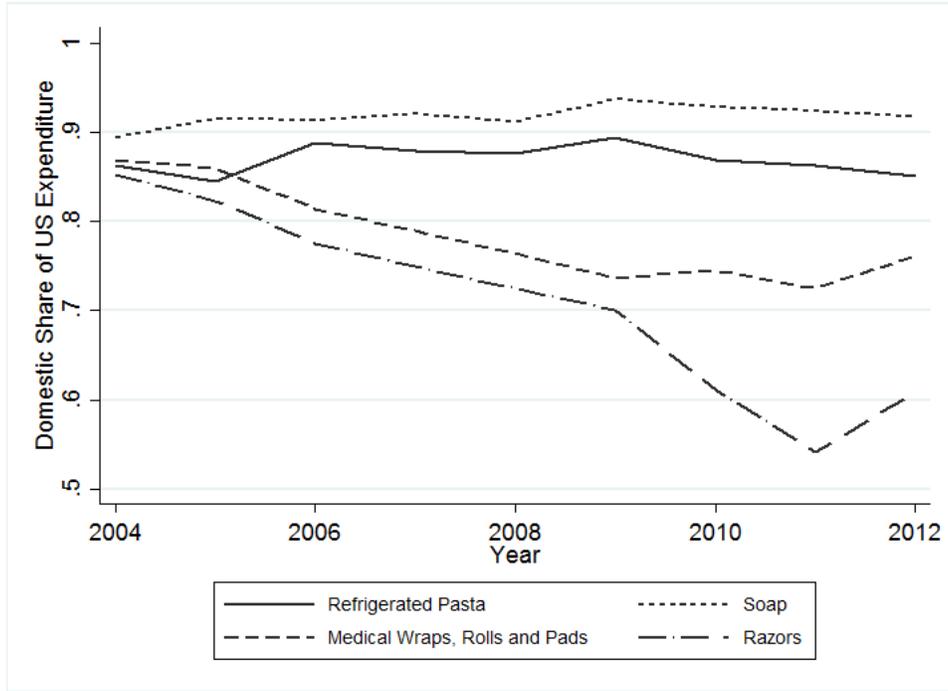
Notes: Using UN Comtrade data at the HS-6 digit level, this figure plots the share of total US imports accounted for by: (i) China, (ii) the European Union, (iii) Canada and Mexico, and (iv) the other BRIC countries (i.e. Brazil, Russia, and India) respectively during 2003-2012.

Figure 2: Changing Composition of US Expenditure



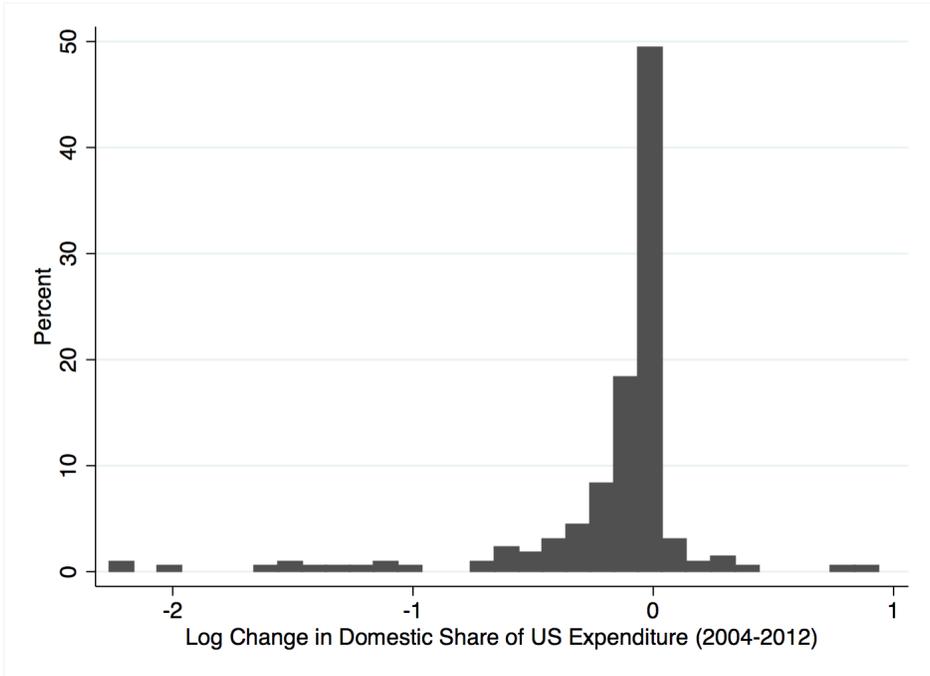
Notes: This figure shows the evolution of both the domestic and China shares of US expenditure (left and right axes respectively) during 2004-2012, using the HS 6-digit codes in our sample. Total US expenditure is computed as US production + imports - exports. Output data for the US come from the Bureau of Economic Analysis, while imports and exports data come from UN Comtrade.

Figure 3: Product Heterogeneity



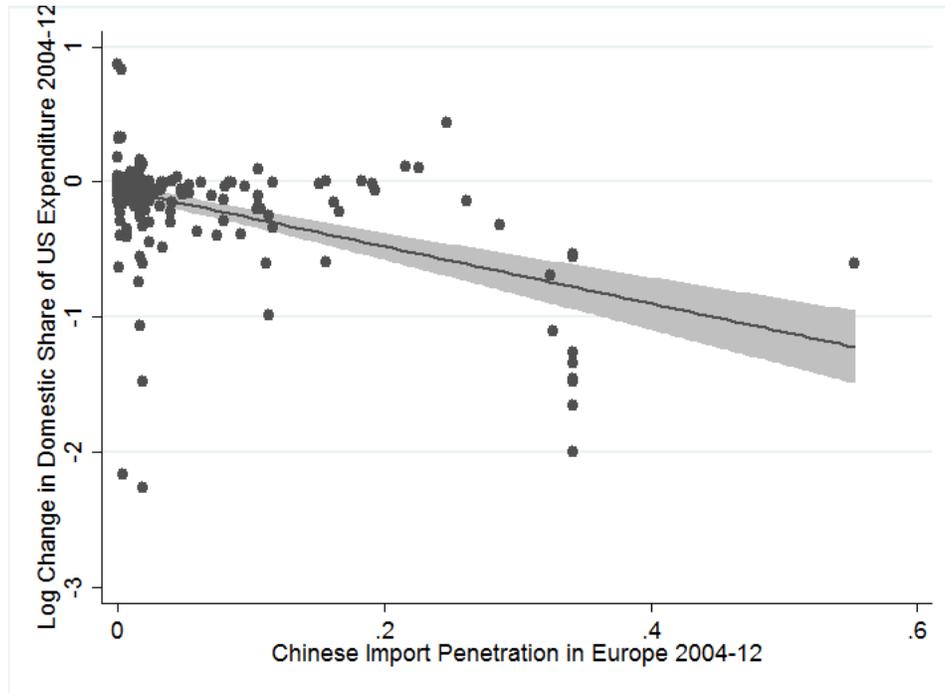
Notes: This figure illustrates the variation in domestic share of US expenditure across different product categories during 2004-2012 in our sample, using: (i) refrigerated pasta, (ii) soap, (iii) razors, and (iv) medical wraps, rolls and pads as examples. For a more systematic illustration of the variation across all product categories, please see Figure 4.

Figure 4: Distribution of Log Change in Domestic Share of US Expenditure



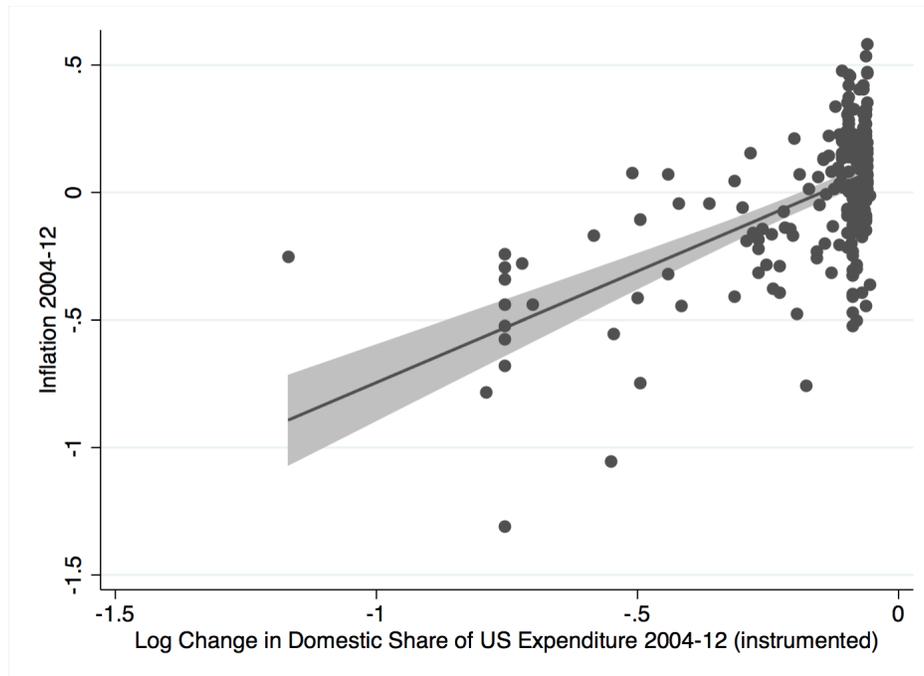
Notes: This figure is a histogram of the log change in domestic share of US expenditure, the main explanatory variable in our analysis, during 2004-2012 across the product categories in our sample.

Figure 5: First Stage



Notes: This figure plots the first-stage relationship at the product category level, between Chinese import penetration in Europe and the log change in domestic share of US expenditure during 2004-2012. Each dot represents a product category. The gray solid line is the OLS linear fit, while the shaded area corresponds to the 95% confidence interval.

Figure 6: Effect on Inflation



Notes: This figure plots the relationship between cumulative inflation and the log change in domestic share of US expenditure (instrumented by Chinese import penetration in Europe) during 2004-2012. Cumulative inflation is computed assuming an elasticity of substitution of 5 across different barcodes within a product category. Each dot represents a product category. The gray solid line is the 2SLS linear fit, while the shaded area corresponds to the 95% confidence interval.

Table 1: Summary statistics

Variable	Mean	Std. Dev.	Min	Max	Obs
Panel A: Inflation (category level, 2004-2012)					
Inflation ($\sigma = 3$)	0.934	0.311	0.141	1.749	222
Inflation ($\sigma = 5$)	1.027	0.262	0.269	1.785	222
Inflation ($\sigma = 10$)	1.087	0.231	0.386	1.805	222
Panel B: Inflation (various margins, 2004-2012)					
Inflation (using only goods present in consecutive years)	1.141	0.207	0.515	1.822	222
Inflation (using only goods present in all years)	1.204	0.177	0.759	1.854	222
Inflation (Feenstra (1994) Correction Factor, $\sigma = 5$)	0.889	0.1	0.523	1.015	222
Panel C: Additional Outcomes, 2004-2012					
Log Change of Expenditure	0.460	0.312	-0.419	1.778	222
Log Change in the Number of Barcodes	0.025	0.240	-0.677	0.804	222
Average Rate of Entry	0.228	0.091	0.080	0.491	222
Average Rate of Exit	0.227	0.085	0.083	0.489	222
Panel D: China Trade Shock, 2004-2012					
Log Change in Domestic Share of Expenditure (US)	-0.158	0.389	-2.259	0.873	222
China Import Penetration (Europe)	0.051	0.093	0	0.554	222

Table 2: Import Penetration and Inflation (2004-2012)

	(1)	(2)	(3)	(4)
Panel A: Cumulative Inflation	OLS		IV	
Domestic Share of Expenditure (US)	0.165** (0.074)	0.378*** (0.064)	0.871*** (0.202)	0.662*** (0.111)
Weights	No	Yes	No	Yes
R^2	0.051	0.122	n/a	n/a
N	222	222	222	222
Panel B: First Stage				
China Import Penetration (Europe)			-1.976*** (0.406)	-1.795*** (0.460)
Weights			No	Yes
R^2			0.22	0.37
N			222	222
F-stat			23.7	15.2

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Domestic share of expenditure is the log change in the same variable for the US during 2004-2012. The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). In constructing the inflation rate reported here, we have used a within-category elasticity of substitution equal to 5. The weights used in columns 2 and 4 are total US expenditure in a given category in 2004. The sample includes all product categories that fall within the 1st and 99th percentiles of China import penetration, cumulative inflation and domestic share of expenditure in the US.

Table 3: Decomposing the First Stage

	(1)	(2)	(3)	(4)
	Relative China Share 04-12		Relative ROW Share 04-12	
China Import Penetration (Europe)	2.430*** (0.545)	1.946*** (0.496)	-1.277*** (0.418)	-0.801** (0.342)
Observations	222	222	222	222
R-squared	0.397	0.470	0.149	0.094
Weights	No	Yes	No	Yes

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Share of Chinese (rest of the world) expenditure is the simple change in the share of US expenditure towards goods from China (the rest of the world). The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). The weights used in columns 2 and 4 are total US expenditure in a given category in 2004. The sample includes all product categories that fall within the 1st and 99th percentiles of China import penetration, cumulative inflation and domestic share of expenditure in the US.

Table 4: Import Penetration, Inflation, Expenditure Growth, and Product Scope (2004-2012)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Inflation (Various Margins)	Short-Stay		Long-Stay		New Varieties			
Domestic Share of Expenditure (US)	0.533*** (0.122)	0.405*** (0.069)	0.267*** (0.060)	0.191*** (0.051)	0.338*** (0.084)	0.257*** (0.056)		
Panel B: Additional Outcomes	Expdt Growth		No. of Products		Rate of Entry		Rate of Exit	
Domestic Share of Expenditure (US)	0.044 (0.126)	0.151 (0.113)	-0.234** (0.101)	-0.209 (0.137)	-0.221*** (0.055)	-0.175*** (0.040)	-0.201*** (0.051)	-0.160*** (0.036)
Weights	No	Yes	No	Yes	No	Yes	No	Yes
N	222	222	222	222	222	222	222	222
1st-stage F-stat	23.7	15.2	23.7	15.2	23.7	15.2	23.7	15.2

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. In Panel A, columns 1-4 present results on the intensive margin of the inflation effect, while columns 5-6 the extensive margin. Specifically, columns 1-2 use an inflation rate constructed using the same set of goods in consecutive years, while columns 3-4 study the effect on an inflation rate constructed using the same set of goods throughout the sample period. In other words, we have excluded the effect of new goods on inflation in columns 1-4. The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). The weights used are total US expenditure in a given category in 2004. The sample includes all product categories that fall within the 1st and 99th percentiles of China import penetration, cumulative inflation and domestic share of expenditure in the US.

Table 5: Barcode level: Average effect on products

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Unweighted regressions	Price	Expenditure	Market Share	Exit	Price	Expenditure	Market Share
Domestic Share of Expenditure (US)	0.786*** (0.276)	1.190* (0.718)	1.166 (0.751)	-2.932 (2.090)	0.389*** (0.102)	1.357*** (0.378)	1.244*** (0.323)
Observations	496,602	496,602	496,602	496,602	209,002	209,002	209,002
Weights	No	No	No	No	No	No	No
Sample	All	All	All	All	Stayer	Stayer	Stayer
Panel B: Weighted regressions	Price	Expenditure	Market Share	Exit	Price	Expenditure	Market Share
Domestic Share of Expenditure (US)	0.618*** (0.162)	2.066** (0.975)	1.957** (0.962)	-3.675*** (1.343)	0.298*** (0.076)	1.315* (0.733)	1.177* (0.651)
Observations	496,602	496,602	496,602	496,602	209,002	209,002	209,002
Weights	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sample	All	All	All	All	Stayer	Stayer	Stayer

Notes: Standard errors in parentheses are clustered at the product category level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The instrument used is China import penetration in the five largest EU economies (Germany, France, United Kingdom, Italy and Spain). The sample includes all products observed in 2004. The weights used are total expenditure on a given product in 2004.

Table 6: Heterogeneous Effects of Import Penetration on Inflation (2004-2012)

	(1)	(2)	(3)	(4)	(5)
Panel A: By Income Group	< 30k	(30k – 50k)	(50k – 70k)	(70k – 100k)	> 100k
Domestic Share of Expenditure (US)	0.609*** (0.119)	0.566*** (0.101)	0.538*** (0.087)	0.521*** (0.092)	0.499*** (0.079)
Weights	Yes	Yes	Yes	Yes	Yes
<i>N</i>	220	220	220	220	220
1st-stage F-stat	24.5	24.5	24.5	24.5	24.5
Panel B: By Region	(1) Northeast	(2) Midwest	(3) South	(4) West	
Domestic Share of Expenditure (US)	0.569*** (0.106)	0.559*** (0.100)	0.621*** (0.106)	0.511*** (0.085)	
Weights	Yes	Yes	Yes	Yes	
<i>N</i>	220	220	220	220	
1st-stage F-stat	24.5	24.5	24.5	24.5	

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. In Panel A, . The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). The weights used are total US expenditure by the relevant group in a given category in 2004. The sample includes all product categories that fall within the 1st and 99th percentiles of China import penetration, cumulative inflation and domestic share of expenditure in the US.

Table 7: Decomposing Differences in Inflation Rates

	(1)	(2)	(3)	(4)	(5)
	Inflation 04-12	Inflation 04-12	Composition	Price Diff	Ext Margin
Domestic Share of Expenditure (US)	0.652*** (0.106)	0.617*** (0.102)	-0.022** (0.010)	-0.008 (0.005)	-0.006 (0.011)
Observations	222	222	222	222	222
R-squared	0.065	0.058	0.027		0.003
Weights	Yes	Yes	Yes	Yes	Yes
Group	Low Income	High Income	High - Low	High - Low	High - Low

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). The weights used are total US expenditure in a given category in 2004. The *composition*, *price difference*, and *extensive margin* terms refer to differences in the inflation rates between rich and poor households arising from differences in the within-category composition of consumed products, differences in price growth for the same product, and difference in the extensive margin contribution, respectively. The sample includes all product categories that fall within the 1st and 99th percentiles of China import penetration, cumulative inflation and domestic share of expenditure in the US.

Table 8: Import Growth in Intermediate Inputs and Inflation (2004-2012)

Dependent Variable: Cumulative Inflation	(1)	(2)
	IV	
Domestic Share of Expenditure in Final Goods (US)	0.861*** (0.195)	0.664*** (0.108)
Domestic Share of Expenditure in Intermediate Inputs (US)	0.472 (1.738)	-0.150 (2.454)
Weights	No	Yes
<i>N</i>	222	222
1st-Stage F-stat	12.3, 13.8	8.56, 4.10

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Share of domestic expenditure is the log change in the same variable for the US during 2004-2012. The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). In constructing the inflation rate reported here, we have used a within-category elasticity of substitution equal to 5. The weights used in columns 2 and 4 are total US expenditure in a given category in 2004. The sample includes all product categories that fall within the 1st and 99th percentiles of China import penetration, cumulative inflation and domestic share of expenditure in the US.

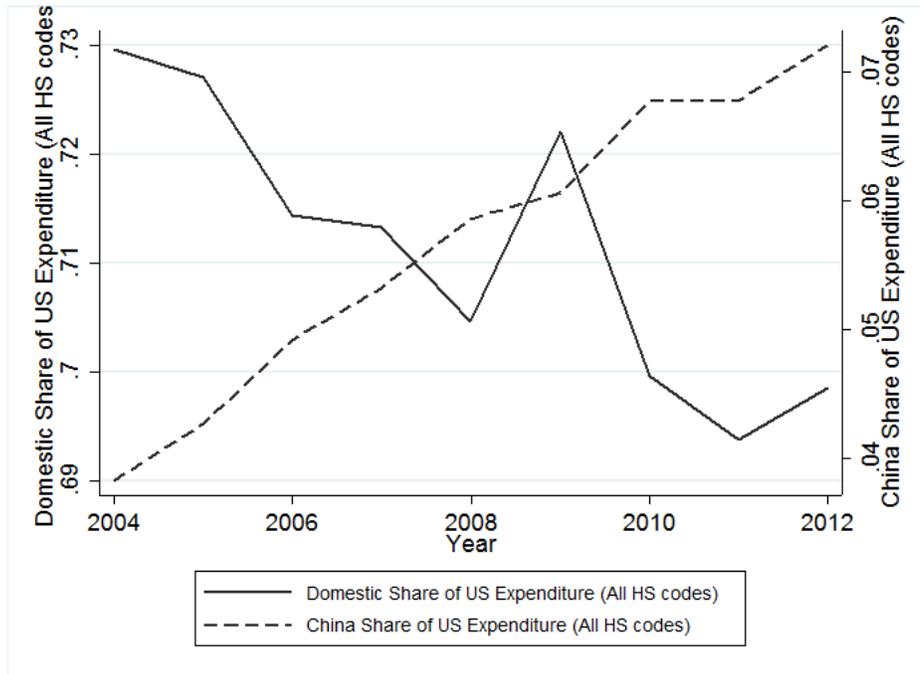
Table 9: Robustness Checks: Effect on Cumulative Inflation

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var: Cumulative Inflation	Controlling for productivity growth	Alternative Instrument	Alternative Instrument	Alternative Instrument	Excluding Food and Drinks	Excluding Food and Drinks
Domestic Share of Expenditure (US)	0.807*** (0.180)	0.653*** (0.112)	0.877*** (0.205)	0.660*** (0.111)	0.563*** (0.185)	0.232** (0.102)
$\Delta \log(\text{Real Output per Worker})$	-0.251 (0.158)	-0.038 (0.026)				
<i>N</i>	220	220	222	222	104	104
Weights	No	Yes	No	Yes	No	Yes
1st-Stage F-stat	28.2	15.5	23.2	15.2	15.0	12.1

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). The weights used are total US expenditure in a given category in 2004. The sample includes all product categories that fall within the 1st and 99th percentiles of both China import penetration and cumulative inflation. The alternative instrument is the residual from a regression of Chinese import penetration into Europe on Mexican and Canadian import penetration into Europe.

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 2004-2012, using all HS 6-digit codes. Total US expenditure is computed as US production + imports - exports. Output data for the US come from the Bureau of Economic Analysis, while imports and exports data come from UN Comtrade.

Table A.1: Robustness Check: Alternative Values for the Elasticity of Substitution

Dependent Variable:	(1) Inflation: $\sigma = 2$	(2) Inflation: $\sigma = 2$	(3) Inflation: $\sigma = 3$	(4) Inflation: $\sigma = 3$	(5) Inflation: $\sigma = 10$	(6) Inflation: $\sigma = 10$
Domestic Share of Expenditure (US)	1.885*** (0.452)	1.432*** (0.269)	1.209*** (0.285)	0.919*** (0.162)	0.683*** (0.157)	0.519*** (0.085)
Observations	222	222	222	222	222	222
Weights	No	Yes	No	Yes	No	Yes
1st-stage F-stat	23.7	15.2	23.7	15.2	23.7	15.2

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The instrument used is China import penetration in the five largest European economies (Germany, France, United Kingdom, Italy and Spain). The weights used are total US expenditure in a given category in 2004. The sample includes all product categories that fall within the 1st and 99th percentiles of both China import penetration and cumulative inflation.