What are the Price Effects of Trade? Evidence from the U.S. and Implications for Quantitative Trade Models*

Xavier Jaravel, London School of Economics
Erick Sager, Federal Reserve Board

April 2024

Abstract

This paper finds that U.S. consumer prices fell substantially due to increased trade with China. With comprehensive price micro-data and two complementary identification strategies isolating supply shocks in China, we estimate that a 1pp increase in import penetration in a product category causes a 1.9% decline in consumer prices, relative to less exposed product categories. This change in relative prices is one order of magnitude larger than in standard trade models that abstract from strategic price-setting. We find a large fall in domestic prices, driven by intensified competition and declining markups. Product categories catering to low-income consumers experienced larger price declines.

JEL codes: F10, F13, F14

*For thoughtful comments, we thank our discussants Vanessa Alviarez, Teresa Fort, Colin Hottman, Amit Khandelwal and Paul Pivetteau, as well as Pol Antras, Mary Amiti, David Atkin, David Autor, Arnaud Costinot, Kirill Borusyak, Thomas Chaney, David Dorn, Robert Feenstra, Andrew Figura, Thesia Garner, Giammario Impullitti, Oleg Itskhoki, Thierry Mayer, Marc Melitz, Daniel Sturm, Felix Tintelnot, Jon Vogel, and numerous seminar participants at the ASSA Meetings, Atlanta Fed, Barcelona GSE Summer Forum, Berkeley, CEA, College de France, Dartmouth, FTC, Harvard, HBS, ISES, JHU, LBS, LSE, Mannheim, MIT, PIIE, the 2018 NBER Summer Institute, San Francisco Fed, Sciences-Po, SED, UCLA and WAITS. We also thank Fergal Hanks and Gabriel Leite Mariante for excellent research assistance. This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data and all results have been reviewed to ensure that no confidential data are disclosed. The views expressed in this paper are those of the authors alone and do not necessarily reflect those of the Bureau of Labor Statistics, Federal Reserve Board of Governors, the Federal Reserve System or the U.S. government.

I Introduction

What is the impact of trade on consumer prices? Canonical trade models predict that trade benefits consumers through lower prices but may hurt some workers through reduced earnings (e.g., Stolper and Samuelson (1941)). While recent reduced-form evidence indicates that increased trade with China had a large adverse impact on U.S. labor markets (e.g., Autor et al. (2013), Autor et al. (2014), Pierce and Schott (2016), Bloom et al. (2019)), much less is known about the potential benefits to U.S. consumers through lower prices. The magnitude of the price response is an empirical question, because various mechanisms could be at play. As trade with China increases, to what extent do retailers adjust prices facing U.S. consumers? Are price changes driven by products imported from China, or is there a broader impact on prices of domestically-produced goods? To the extent that prices fall, which consumers benefit most and how do the gains in purchasing power for consumers compare to the losses for workers through job disruptions? Data limitations explain the relative scarcity of evidence on these questions, which can only be answered with comprehensive price data.

In this paper, we use micro data from the U.S. Bureau of Labor Statistics (BLS) to obtain comprehensive coverage of price dynamics over a long panel, going back to the 1980s with both consumer prices and producer prices. We estimate the response of prices to the rise in trade induced by supply shocks in China. Estimating the causal response of U.S. consumer prices is challenging because of potential omitted variable biases and reverse causality. For example, China has a comparative advantage in specific product categories that may be on different inflation trends, such as consumer electronics or apparel.

To overcome this challenge, we combine two complementary research designs borrowed from recent work studying the consequences of trade with China on employment across U.S. industries. Pierce and Schott (2016) leverage a change in U.S. trade policy passed by Congress in October 2000, which eliminated potential tariff increases on Chinese imports.[1] This research design uses transparent policy variation and lends itself to year-by-year tests for pre-trends. But the effects of trade induced by changes in policy uncertainty may differ from those of more common permanent changes in foreign productivity. To gauge the generality of our main estimates, we also use the empirical strategy of Autor et al. (2014), who instrument for changes in import penetration from China in the U.S. with contemporaneous changes in eight comparable economies. Finally, we demonstrate the plausibility of a causal interpretation of our estimates by implementing a series of stringent falsification and robustness tests. We find no evidence for pre-trends. Moreover, we implement a stringent triple-difference test using price data from France and the instrument from Pierce and Schott (2016), and we find that there is no similar reaction of prices in France where

[1] A similar research design was developed independently and concurrently by Handley and Limão (2017).

there was no policy change.

Our IV estimates indicate that the price effects of increasing trade with China are large. With the instrument from Pierce and Schott (2016), a one percentage point increase in the import penetration rate from China causes a fall in inflation of 2.2 percentage points (s.e. 0.47).[2] With the instrument from Autor et al. (2014), the corresponding fall in consumer prices is 1.4 percentage points (s.e. 0.45). With both instruments jointly, the IV coefficient is -1.9 percentage points (s.e. 0.38).[3]

We investigate several potential mechanisms that could account for the magnitude of our IV estimate. To do so, we must overcome another challenge: prior work does not provide guidance about the extent to which cross-industry regression estimates can be related to quantitative trade models, which make welfare statements in terms of economy-wide variables, such as the theoretical price index of a representative agent (see, e.g., Arkolakis et al. (2012)). We first show how to interpret our IV coefficient within the structure of standard quantitative trade models: even though we cannot estimate the "missing intercept", the estimated relative price effect across product categories is a useful identified moment that can serve as a diagnostic tool to distinguish between classes of trade models (e.g., *à la* Nakamura and Steinsson (2018)). In the set of trade models characterized by Arkolakis et al. (2012), the price response across product categories is predicted to be equal to the inverse of the trade elasticity. We find that our estimated price effect is about one order of magnitude larger than predicted by these models. To account for the observed effect, this class of models requires an implausibly small trade elasticity, around 0.3, while common estimates are around 4.25 (Simonovska and Waugh (2014)).

We focus on linking our cross-industry IV specification to the price response predicted by the Melitz (2003)-Chaney (2008) model, which we take as our preferred member of the Arkolakis et al. (2012) class.[4] We study the predicted relative price response in this model, in general equilibrium, and highlight that it differs from our IV estimate by one order of magnitude, even after accounting for the potential divergence between the measured CPI and the model-based exact price index. We also show formally that other members of the Arkolakis et al. (2012) class make very similar quantitative predictions, including Armington (1969), Eaton and Kortum (2002), and models with

[2]Notably, the magnitude of our estimate is quite large, even relative to existing work. Compared to Bai and Stumpner (2019), the only other study to consider consumer price responses to the China shock, we recover a much larger price response when considering the entire non-shelter component of the CPI as opposed to their focus on grocery items. Amiti et al. (2020) study producer prices, and thus we can make complementary inferences about consumer prices and we find different mechanisms contribute to the price effect.

[3]While we sometimes use the terms "inflation" and "price effects" for convenience in the remainder of the paper, all of our results must be interpreted in terms of changes in relative prices across product categories that have different exposure to supply shocks in China, as we make clear in the theoretical propositions derived below. We show that the relative price changes obtained from the empirical estimates are directly comparable to the structural parameters of quantitative trade models.

[4]Specifically, we study a trade model with CES preferences and heterogeneous firms engaged in monopolistic competition (Melitz (2003)) whose productivity follows a Pareto distribution (Chaney (2008)).

intermediate inputs as in Caliendo and Parro (2015) and Ossa (2015); so do other leading trade models outside this class, such as Arkolakis et al. (2019).

To uncover the relevant mechanisms, we document which products drive the price response. Using a statistical decomposition, we isolate the roles of continued products (as opposed to new products) and domestically-produced goods (as opposed to foreign products). We find that continued products account for approximately 70% of the overall price effects. To isolate the role of U.S.-produced goods, we identify U.S. goods in the Consumer Price Index (CPI) sample using specification checklists. We find that domestic prices account for a substantial fraction of the overall price effects, between 44% and 85% depending on the specification. We confirm the role of the domestic price response using the Producer Price Index (PPI) sample, which covers domestic manufacturers only.

The domestic price response for continued products could result from two potential effects of increased trade with China on U.S. manufacturers: changes in production costs, or changes in markups. We first assess the role of changes in domestic production costs, which we decompose into several potential sources: wages, intermediate inputs and offshoring, and returns to scale and productivity.

Although changes in domestic production costs are theoretically plausible, in practice we find that they can account for only a small fraction of the estimated price response.[5] Wages fall in response to trade with China, but both public data and administrative data used in prior studies (e.g., Autor et al. (2014)) indicate that the wage effects are much smaller than would be needed to explain the observed change in domestic relative prices across product categories.[6] To assess the role of intermediate inputs, we use the BEA's input-output table and measure upstream and downstream changes in trade with China for each product category. We find that upstream and downstream trade does not help explain the estimated price effects. Finally, by displacing domestic goods and reducing the scale of domestic production, increased import competition with China could affect domestic production costs through (decreasing) returns to scale. In fact, empirical studies find that tradable U.S. industries have *increasing* returns to scale; explaining the estimated price response would require an elasticity of the opposite sign and about five times larger in magnitude than benchmark estimates of returns to scale.

Next, we turn to the potential relevance of domestic markups. Using the setting of Edmond et al. (2015), we show theoretically that the large price effect can be explained by models that

[5] We focus on the importance of changes in production costs in explaining the large cross-sectional relationship between changes in import penetration rates and consumer prices; we find it to be small empirically. However, this empirical result does not imply that changes in production costs are unimportant for the aggregate welfare effects of trade.

[6] Because the labor share in total domestic output for the relevant product categories is very low, explaining the estimated 2% fall in domestic prices (due to a 1pp increase in import penetration from China) requires a very large wage response, on the order of 20%, which we can reject empirically.

feature strategic interactions in pricing. Intuitively, as Chinese producers become more productive they reduce their prices, which leads U.S. producers to reduce their markups through strategic interactions. Because of the fall in U.S. prices, U.S. consumers do not substitute as much toward the products from China. Therefore, the equilibrium change in import penetration rate from China is lower than it would be without the price response for U.S.-produced goods. As a result, the model yields a large reduced-form relationship between changes in import penetration and price changes across product categories.

We conduct several empirical tests of the markup channel, to assess its relevance to our baseline price response estimate. First, we examine the response of estimated markups for publicly-listed firms in Compustat, following the methodology of De Loecker et al. (2020) to estimate markups. We observe a fall in estimated markups: when the import penetration rate from China increases by one percentage point, domestic markups fall by 1.8 percentage points (s.e. 0.85). This estimate is large in magnitude and statistically indistinguishable from the IV coefficient for the response of domestic prices. Moreover, the observed changes in the *distribution* of markups are consistent with the predictions of the model: as trade increases, markups fall primarily at the top of the markup distribution (e.g., there is no effect at the 10th percentile but a large effect at the 90th percentile). Second, we address the incomplete coverage of the Compustat sample and potential limitations of markup estimation by returning to our main sample and assessing whether heterogeneity in the estimated price effects is consistent with key predictions of the markup channel. We document that the price effects are significantly larger in industries where domestic market concentration is higher and where China's initial market share is lower. These patterns are in line with the model: there is more domestic market power to be disrupted by China when the domestic market is more concentrated.

Finally, we discuss how our estimates shed light on the distributional effects of the "China shock." We first benchmark our estimates of the price response, which benefits consumers, to the employment effects estimated in prior work. Using the IV estimates for the price and employment effects, our baseline specification indicates that falling prices in product categories that are more exposed to trade with China create \$411,464 in consumer surplus for each displaced job. The estimates vary from \$288,147 to \$477,555 across specifications.[7] These large magnitudes suggest that it may be possible to compensate those who suffer from the labor market impacts of trade shocks, and that there are large distributional effects from producer surplus toward consumer

[7] These calculations reflect partial-equilibrium differences across industries with different levels of exposure to rising import penetration from China. General equilibrium effects induced by the China shock could affect all industries simultaneously. If displaced manufacturing jobs lead to more job creation in other industries, then the increase in consumer surplus per "destroyed" job at the aggregate level (rather than per "displaced" jobs across industries) would be larger. For example, using a quantitative general equilibrium model, Galle et al. (2022) estimate that overall employment increases because, in their model, the China shock leads to an increase in the average real wage across US commuting zones. Instead, we take a conservative approach and compare price effects to job displacements across industries.

surplus in each product category. In contrast, in the class of standard trade models nested by Arkolakis et al. (2012), the predicted increase in consumer surplus per displaced job would be attenuated by a factor of ten and would be on the order of $40,000 per displaced job, similar to average annual labor earnings in the sample.

Lastly, we investigate distributional effects across consumers and find that the price response is larger in product categories that cater to lower-income households. For example, for product categories with a share of sales to college graduates *below* median, the magnitude of the price response is about five times larger than for the categories with a share above median. The patterns are similar with other proxies for consumer income. These results indicate that distributional effects can arise because of differences in the price responses to trade shocks. This channel appears to be quantitatively important and is novel relative to other mechanisms investigated in prior work (e.g., Fajgelbaum and Khandelwal (2016), Carroll and Hur (2020) and Borusyak and Jaravel (2021) examine differences in spending shares on imports, and Hottman and Monarch (2020) document differences in import price inflation across income groups).

Prior work. This paper relates and contributes to several strands of the literature. First, our estimates of the benefits of trade with China for consumers through lower prices complement a large literature that has documented adverse effects for employment (e.g., Autor et al. (2013), Autor et al. (2014), Pierce and Schott (2016) and Bloom et al. (2019)), mortality (Pierce and Schott (2020)), marriage, fertility and children's living circumstances (Autor et al. (2019)), domestic innovation and investment (Pierce and Schott (2018) and Autor et al. (2020a)), and political polarization (Autor et al. (2020b)).

Second, a growing literature examines the reduced-form impact of changes in trade on domestic producer prices or import prices measured at the border, but no paper uses comprehensive data on consumer prices as we do. Our findings advance the literature by leveraging a comprehensive data set representative of the market basket of U.S. consumers (which includes both domestically produced and imported product varieties) and the prices they actually pay, which allows for an in-depth investigation of the identifying assumptions for causal identification (e.g., with pre-trend tests). We estimate a much larger price effect of the China shock than in the only available study of consumer prices by Bai and Stumpner (2019), which covers consumer packaged goods.[8] Our work is

[8]For our purposes, scanner data such as those used by Bai and Stumpner (2019) suffer from three drawbacks. First, the sample covers fast-moving consumer goods and is not representative of several important product categories for trade with China (e.g., apparel, consumer electronics, appliances, and other slow-moving consumer goods). Second, the sample starts in 2004, making it impossible to test for pre-trends prior to the "China shock" or to study the period 2000-2007, which has been the focus of the literature studying the labor market effects of trade with China. Third, it is not possible to isolate domestic goods or track a set of products that were already available before China joined the WTO. Bai and Stumpner (2019)'s point estimates for the price response are about 60% smaller than ours; in Appendix Table A1, we implement our IV specification in a sub-sample of products that approximates the sample of Bai and Stumpner (2019): we obtain a point estimate statistically indistinguishable from theirs. This result highlights the importance of considering the full consumption basket, as we do. A complete description of how the

also closely related to prior work by Amiti et al. (2020), who provide evidence on the impact of trade with China on U.S. producer prices in manufacturing.[9] Relative to them, (i) we measure consumer prices in the U.S., rather than producer prices, (ii) we focus on a complementary mechanism, i.e. the importance of the price response for domestically-produced goods through strategic pricing, and (iii) we show how to interpret the cross-industry estimates through the lens of standard trade models despite the missing intercept. Finally, Amiti et al. (2019b), Cavallo et al. (2021), Fajgelbaum et al. (2020) and Flaaen et al. (2020) estimate the effects of the 2018 "trade war" on import and producer prices over a one-year horizon; our work predates these studies and complements them by estimating the response of consumer prices to the historical "China shock" over a long horizon, close to a decade.[10] Given our focus on consumer prices, we do not analyze import prices at the border, which have been studied in other work; instead, we directly study consumer prices for imported goods.[11]

Third, we formally link our specification and data to the predictions of leading quantitative trade models that are most commonly used for policy analysis, accounting for general equilibrium effects, sectoral heterogeneity, the "missing intercept", and potential deviations between the theoretical price index and the measured CPI. In particular, our results indicate that the Melitz (2003)-Chaney (2008) model, and more broadly the members of the Arkolakis et al. (2012) class, do not offer a good approximation to observed relative price changes and their distributional effects. By showing the importance of the "pro-competitive effects of trade" to explain the observed relative price changes, our paper is part of a large literature that has investigated the relationship between international trade and markups (e.g., Brander and Krugman (1983), Levinsohn (1993), Krishna and Mitra (1998), Atkeson and Burstein (2008), Epifani and Gancia (2011), Edmond et al. (2015), Feenstra and Weinstein (2017), Feenstra (2018), Impullitti and Licandro (2018), and Arkolakis et al. (2019)). Complementary to the structural approaches in these papers, we use reduced-form identification strategies to provide direct causal evidence.[12] We find that the estimated price effects are consistent

analysis of Bai and Stumpner (2019) relates to ours is provided in Appendix A.

[9] Amiti et al. (2020) study the impact of China's WTO entry on prices using the U.S. PPI as well as data on the unit value of imports. They find that a key mechanism is China lowering its own import tariffs on intermediate inputs. A complete description of how their analysis relates to ours is provided in Appendix A.

[10] Moreover, focusing on the historical "China shock" allows us to draw a comparison between the gains to consumers through lower prices and the losses to workers through job disruptions in a unified empirical setting.

[11] Given space constraints, a more detailed analysis of the literature is reported in Appendix A. We show that there is a growing empirical consensus about the large effects of foreign supply shocks on producer prices. Relative to this prior work, our contribution is threefold: (i) to study consumer prices rather than producer prices (only Bai and Stumpner (2019) studied consumer price, focusing on a selected sample of products where the price response turns out to be much smaller than in our comprehensive sample); (ii) to show that the large magnitude of the price effects is robust to multiple potential concerns about causal identification, which could not be addressed in prior work due to data limitations; (iii) to show formally that these cross-industry estimates are inconsistent with leading trade models used for policy analysis, which was not noted in prior work; we do so by providing a structural interpretation of the reduced-form regression coefficients. Appendix A also reviews prior work on import prices at the border.

[12] Nakamura and Zerom (2010) and Auer et al. (2021) also take a reduced-form approach but examine exchange-rate shocks, instead of foreign competition shocks as we do. De Loecker et al. (2016) study India's trade liberalization and domestic markups using producer price data; relative to them, we study a high-income economy and examine

with models featuring strategic price setting (e.g., Atkeson and Burstein (2008) and Edmond et al. (2015)) but that they are too large to be rationalized by models featuring monopolistic competition and variable elasticities of substitution (with conventional parameter values, e.g. Arkolakis et al. (2019)).

The paper is organized as follows: Section II present the data and variable definitions, Section III estimates the reduced-form effect of increased trade with China on U.S. consumer prices, Section IV distinguishes between potential mechanisms, and Section V estimates the distributional effects.

II Data

In this section, we describe the data sources, define the samples and key variables we use in the analysis, and present summary statistics.

II.A Data Sources, Samples and Variable Definitions

Consumer Price Index. Our main outcome variable is inflation faced by U.S. consumers across product categories. We measure this outcome using the micro data underlying the Consumer Price Index, available from the Bureau of Labor Statistics' internal CPI Research Database (CPI-RDB). The CPI-RDB contains all product-level prices on goods and services collected by the BLS for use in the CPI since January 1988, excluding shelter. A product is defined as a specific item available in a specific store, such as a 500 ml bottle of Coca-Cola on the shelf of a specific Whole Foods Market store in Washington, DC. The BLS data collectors track prices monthly or bi-monthly, depending on the product category, and they identify products using bar codes whenever possible.

Our goal is to estimate the price effects of trade shocks defined at the level of a product category, therefore we aggregate the product-level price changes into category-level price changes. We do so following the BLS' procedure to compute official aggregate inflation statistics, which is described in Appendix B.A. We obtain 222 product categories spanning the full range of final consumption goods and services, with the exception of shelter. These categories, called Entry Level Item (ELI) categories, are the most detailed categories in the BLS' product classification. They are ideal for our purposes because they offer a comprehensive coverage of consumption and are sufficiently detailed such that we expect product substitution to occur primarily within, rather than across categories. For example, a bottle of Coca-Cola belongs to the "Carbonated Drinks" ELI; other examples of ELIs include "Washers & Dryers," "Woman's Outerwear," or "Funeral Expenses."

We leverage the price micro data to build alternative category-level price indices, which we use for various robustness tests and extensions. Alternative category-level price indices help us address potential measurement issues. For example, the baseline CPI index uses quality adjustments when the BLS data collector is unable to find the exact same product in the exact same store from

consumer prices, inclusive of potential changes in retail markups.

one period to the next (e.g., the 500 ml bottle of Coca-Cola might no longer be on the shelf at Whole Foods and might have been replaced with a 500 ml bottle of Pepsi). Given that BLS quality adjustments may not perfectly account for potential changes in underlying product characteristics in such cases, we build an alternative price index based solely on price changes for "continued products" (i.e., those instances when the same item in the same store is observed from one period to the next). We also leverage the micro data to build alternative price indices that help decompose the sources of the price effects we document. For example, we can isolate the role of the price response of products made in the United States.[13]

In addition to its flexibility for inflation measurement, the CPI price data set features other noteworthy advantages. The CPI data set is available over a long panel and covers the representative consumer's market basket comprehensively. This allows us to implement stringent tests for "pre-trends" and assess the plausibility of a causal interpretation of our IV estimates.[14] Although the main data set extends back to 1988, to conduct a more complete analysis of pre-trends we build a similar data set going back to 1977, following Nakamura et al. (2018). Online Appendix B.B describes the construction of this extended sample. Moreover, the CPI measures prices inclusive of retail margins, which is the relevant price for consumers.[15]

A limitation is that the sample frame keeps a fixed number of items in each product category, which makes it impossible to study changes in product variety over time. The available evidence to date suggest that increased trade with China may lead to an increase in product variety, which lowers consumers' effective price index through love of variety.[16]

Trade data. Our main independent variable is the import penetration rate from China over time and across product categories. For product category i, the import penetration rate from China at

[13]See Section IV for a complete discussion. Note that such robustness tests and statistical decompositions would not be possible by using the publicly-available inflation series from the BLS. Another downside of the public data from the BLS, relative to the CPI-RDB data set, is that the publicly available product categories are coarser than ELIs and their definitions change over time; as a result it is difficult to build a balanced panel of detailed product categories over a long time horizon in this data set.

[14]In contrast, scanner data is restricted to consumer packaged goods and is only available after 2000, making it impossible to appropriately assess the validity of the research design. For example, the Nielsen scanner data is available from 2004 onward and offers limited coverage of several product categories in which trade with China is particularly important, such as consumer electronics, household appliances and apparel (for a discussion of expenditures coverage in Nielsen scanner data, see for example Jaravel (2019)).

[15]Therefore we use CPI inflation as our preferred outcome, rather than import or producer price indices.

[16]The impact of trade on product variety remains debated, because domestic exit may offset the increase in foreign varieties available to consumers. Broda and Weinstein (2006) show that the number of foreign varieties increases, but they do not observe domestic varieties in the trade data. Hsieh et al. (2020) suggest that import variety gains are counteracted by exactly analogous domestic variety losses, but they observe plants instead of products. Bai and Stumpner (2019) directly measure changes in product variety using barcode data and estimate that increased trade with China led to a change in product variety that lowered the cost of living for U.S. consumers, as measured by the Feenstra (1994) adjustment factor in a CES framework.

time t is defined as:

$$ChinaIP_{it} = \frac{Imports_{it}^{China}}{DomesticProduction_{it} + TotalImports_{it} - TotalExports_{it}}, \quad (1)$$

where the denominator corresponds to domestic absorption. To make our results comparable with prior work examining the impact of increased import competition with China on employment, we use the measures of import penetration from China built by Acemoglu et al. (2016) at the level of Standard Industrial Classification (SIC) industries. Following their approach, we consider long differences, i.e. the change in the China import penetration rate over two relatively long periods, 1991-2000 and 2000-2007.

Although our baseline trade exposure measure facilitates the comparison with prior work, it could suffer from potential limitations, which we relax in various extensions. In standard trade models, the price effects of trade are related to changes in the overall import penetration rate (not just from China), sometimes with a specific functional form (e.g., the log change in the domestic expenditure share in Arkolakis et al. (2012)). We compute this alternative measures by matching trade data recorded under Harmonized System (HS) codes (from China and from the rest of the world) to domestic production data from the NBER-CES Manufacturing database, using the concordance of Pierce and Schott (2012). In another robustness test, we adjust the denominator in equation (1) for distribution margins. We estimate these margins using the BEA's input-output (IO) table: for each industry the ratio of total output in *purchaser* prices to total output in *producer* prices gives the distribution and transportation margins.

Main analysis sample. Our main analysis sample brings together the CPI inflation data (by ELI categories), the trade data (by SIC industries) and the instruments (by SIC industries for Autor et al. (2014) and NAICS industries for Pierce and Schott (2016)). The ELI categories are more aggregated than SIC and NAICS industries, therefore we build many-to-one crosswalks from SIC and NAICS industries to ELIs and aggregate all variables accordingly.[17]

Input-Output sample. To investigate robustness to aggregation choices and test specific mechanisms, we also build a linked data set at the level of 6-digit IO industries. We use the BEA's 2007 IO table because it is the most disaggregated during our sample. We build a many-to-one match from ELI categories to these industries and then aggregate the data. The variables we build based on input-output linkages are discussed in Section IV and the data construction is described in Online Appendix B.D.

Producer Price Index sample. To assess the role of domestic prices in the overall price effects, we use data from the BLS's Producer Price Index (PPI) data set, which tracks producer prices for

[17]The crosswalks are described in Online Appendix B.C.

products manufactured in the United States. Appendix B provides more information on all data construction steps.

Additional data sets and variables. Finally, we supplement our analysis sample with several ELI-level variables to assess the robustness of our main estimates and study heterogeneity in the treatment effect. We use a product hierarchy from the BLS that classifies ELIs in various groups (e.g., to assess the role of apparel or high-tech goods), expenditure shares from the public-use Consumer Expenditure Survey,[18] along with trade elasticities, average wages, capital intensity, total factor productivity, and market concentration from Broda and Weinstein (2006), the NBER-CES Manufacturing Database and the U.S. Census. We also use data from the French CPI to implement placebo tests and Compustat data to measure markups following De Loecker et al. (2020). These variables are introduced when relevant in subsequent sections.

II.B Summary Statistics

Table 1 reports the summary statistics for our main analysis sample, from 1991 to 2007.

The first three rows describe the CPI inflation data. Across non-shelter ELI categories, inflation was on average 1.15% per year, but with a large standard deviation of 6.75 percentage points across industry-years. The share of continued products corresponds to the share of product-level price changes for which the exact same item is priced by the data collector from one month to the next. Continued products account for over 80% of all observations on average, which makes it possible to build a price index based on these observations only. The third row reports the share of unavailable products, which corresponds to instances when the data collector was unable to find the same item from one month to the next.

Rows four and five of Table 1 describe the changes in import penetration rates from China. The average (annualized) change in import penetration from China in the United States is 66 basis points in our sample. There is large variation across ELIs and periods, with a standard deviation of 1.62 percentages points. The change in import penetration from China in developed economies comparable to the United States has similar properties.

The remainder of Table 1 reports summary statistics for several variables defined at the ELI level. The NTR gap is on average 21% and exhibits large variation across ELIs. The table also reports various indicators for product categories, showing the fraction of goods, apparel products, high-tech products, and the set of durable goods defined in Bils (2009), reported in Appendix Table A2.

[18] We use the Consumer Expenditure Survey data set as processed by Borusyak and Jaravel (2021). This data set provides information on the characteristics of consumers across about 600 very detailed product categories, called UCC. We implement a many-to-one match of UCCs to ELIs, by hand.

III Trade with China and U.S. Consumer Prices

In this section, we estimate the effect of trade with China on U.S. consumer prices using two complementary identification strategies.

III.A Research Design

Several challenges arise when estimating the causal effect of increased trade with China, stemming from supply shocks, on U.S. consumer prices. To understand the main threats to identification, suppose we were to estimate a regression of the change in U.S. CPI inflation on the change in import penetration from China across U.S. product categories over time. A causal interpretation of the OLS estimate from this specification could be misleading, because there may be many unobserved supply and demand shocks affecting U.S. industries that may correlate with trade with China and have a direct effect on U.S. consumer prices.

For example, China may decide to enter product categories where U.S. suppliers are easy to out-compete due to low TFP growth, implying higher U.S. inflation in these product categories and an upward bias of the OLS estimate. Moreover, omitted variable biases may stem from the fact that China has a comparative advantage in specific product categories, which may be on different inflation trends compared with other product categories. For instance, trade with China is large for computers, consumer electronics and apparel. Because of high rates of innovation for computers and consumer electronics, and because of the "fashion cycle" for apparel, these categories are characterized by low inflation, implying a downward bias of the OLS estimate. There are thus multiple potential sources of bias with offsetting effects, such that it is not possible to sign the potential bias.

Given these identification challenges, we use two complementary research designs borrowed from recent work.

Variation in the NTR gap. Pierce and Schott (2016) and Handley and Limão (2017) focus on a specific change in U.S. trade policy passed by Congress in October 2000, which eliminated potential tariff increases on Chinese imports and became effective when China joined the WTO at the end of 2001. This policy change is known as the granting of "Permanent Normal Trade Relations" (PNTR) to China. Although it did not change the import tariff rates the U.S. actually applied to Chinese goods, it reduced the uncertainty over these tariffs. Before China was granted PNTR, U.S. import tariffs on Chinese goods needed to be renewed by Congress. Without renewal U.S. import tariffs on Chinese goods would have jumped back to high non-NTR tariffs rates assigned to non-market economies, which were originally established under the Smoot-Hawley Tariff Act of 1930. The "NTR gap" is the difference between the actual import tariffs on Chinese goods and non-NTR tariffs. We treat the NTR gap instrument as our benchmark, because it allows for stringent falsification tests, discussed below.

Variation in import penetration from China in other countries. We also use the empirical strategy of Autor et al. (2014), who instrument for the change in import penetration from China across U.S. industries with changes in import penetration from China across industries in eight comparable developed economies. This research design addresses threats to identification that stem from U.S.-specific supply or demand patterns, i.e. changes in U.S. supply or U.S. demand across industries that are not correlated with supply and demand changes in the group of eight comparable economies.

III.B Pre-Trends Analysis

To assess the plausibility of the exclusion restrictions, we implement pre-trend tests.

First, we want to assess whether the NTR gap becomes related to CPI inflation only after the policy change is passed, i.e. after 2000. We use the CPI-RDB database to measure inflation in the pre-period, going back to 1988. We then run a flexible event-study panel specification:

$$\pi_{it} = \sum_{k=1988}^{2007} \beta_k NTR\,Gap_i \cdot 1_{\{k=t\}} + \lambda_i + \delta_t + \varepsilon_{it}, \tag{2}$$

where t indexes year, i indexes ELI categories, π_{it} is the CPI inflation rate, $1_{\{k=t\}}$ is an indicator variable for year t, λ_i is ELI fixed effects, and δ_t is year fixed effects.[19] The path of the year-specific reduced-form coefficients $\{\beta_k\}_{k=1988}^{2007}$ is informative about the plausibility of the identification condition. The exclusion restriction, $\mathbb{E}\left[NTR\,Gap_i \cdot \varepsilon_{it}|i,t\right] = 0$, cannot be tested directly, but if it is valid then there should be no relationship between the treatment and inflation prior to the policy change, and we would expect to find $\beta_k = 0$ for any year prior to 2000.

Panel A of Figure 1 reports the set of reduced-form coefficients from equation (2), along with their 95% confidence intervals (standard errors are clustered by ELIs). This figure shows a striking pattern. From 1988 until 2000, the estimated reduced-form coefficients are small and hover around zero and a F-test cannot reject the null of no effect. But after 2000, the coefficients become negative and statistically significant. This pattern supports the plausibility of a causal interpretation of the relationship between the NTR gaps and inflation outcomes. However, it does not rule out the possibility that other shocks, correlated with both the NTR gap and affecting CPI inflation, may have occurred specifically after 2000. We return to this hypothesis later, with a placebo test using French CPI data.

With the instrument from Autor et al. (2014), testing for pre-trends is more challenging. The instrument is the change in import penetration from China in other developed economies, which does not have a precise start date. Trade with China starts increasing in the late 1980s, therefore it

[19] ELI fixed effects introduce collinearity with the NTR gap, therefore we normalize $\beta_{1988} = 0$.

is instructive to examine whether there is a relationship between the increase in trade with China in our main analysis sample and inflation in the 1980s.

Panel B of Figure 1 presents the placebo reduced-form specifications in the extended CPI sample (1977-1986), across ELIs. We regress the average inflation rate over the sample on the instrument from Autor et al. (2014) (sub-figure (a)) and from Pierce and Schott (2016) (sub-figure (b)).[20] With both instruments, there is no relationship with inflation.

III.C Baseline Estimates

The previous reduced-form specifications support the plausibility of the research design by documenting the absence of pre-trends, but they do not yield properly scaled estimates of the impact of trade with China on U.S. consumer prices. We now turn to IV specifications.

Instrumental Variables framework. We implement a difference-in-differences IV design after aggregating the data over two long periods, 1991-1999 and 2000-2007. Because the effect of a change in import penetration from China on consumer prices may occur with some delay, an IV specification allowing only for contemporaneous effects (i.e., within the same year) may be misspecified. Following prior work (e.g., Autor et al. (2013), Autor et al. (2014) and Acemoglu et al. (2016)), we implement specifications with contemporaneous effects within periods spanning several years. This approach is natural for both of the instruments we use: the policy change making the NTR gap relevant occurs in the early 2000s, and the increase in the import penetration from China in the other developed economies becomes more pronounced in the early 2000s as China joins the WTO.[21]

The baseline IV specification uses ELI fixed effects, as in the analysis of pre-trends in Subsection III.B. ELI fixed effects allow for each product category to be on its own inflation trend over time. Intuitively, we examine whether ELIs that were relatively more exposed to import competition from China in the 2000s (relative to the 1990s) also have lower inflation rates in the 2000s (relative to the 1990s), using the two instruments.

Our IV specification is:

$$\begin{aligned} \pi_{it} &= \beta \Delta ChinaIP_{it} + \nu X_{it} + \delta_i + \delta_t + \varepsilon_{it}, \qquad (3) \\ \Delta ChinaIP_{it} &= \gamma \mathbf{Z_{it}} + \widetilde{\nu} X_{it} + \tilde{\delta}_i + \tilde{\delta}_t + \eta_{it}, \end{aligned}$$

[20]The specifications are $\bar{\pi}_i = \Delta ChinaOther_i + \nu X_i + \varepsilon_{it}$ and $\bar{\pi}_i = NTR\,Gap_i + \nu X_i + \varepsilon_{it}$, where $\Delta ChinaOther_i$ is the annualized change in import penetration from China in other developed economies from 1991 to 2007, $\bar{\pi}_i$ is average annual inflation for ELI i from 1977 to 1986, and X_i is a vector of fixed effects for apparel and durable goods.

[21]Permanent Normal Trade Relationship were granted to China by the U.S. Congress in October 2000 and became effective upon China's accession to the WTO at the end of 2001. U.S. prices may react during 2001 in anticipation of increased import competition at the end of the year, or the response might materialize only after 2001. Empirically, we find that the IV estimates are similar in our baseline panel specification using long differences and in an alternative panel specification considering shorter periods after 2001 (Appendix Table A3).

where i indexes ELIs, t indexes periods (1991-1999 and 2000-2007), π_{it} is the average annual CPI inflation rate over the period, $\Delta China_{it}$ is the average annual change in import penetration rate from China, X_{it} is a set of time-varying controls, δ_i ELI fixed effects, and δ_t period fixed effects. $\mathbf{Z_{it}}$ is a vector of instruments, which varies across specifications. Under the identification condition $\mathbb{E}\left[\mathbf{Z_{it}} \cdot \varepsilon_{it} | X_{it}, i, t\right] = 0$ and relevance condition $\mathbb{E}\left[\mathbf{Z_{it}} \cdot \Delta ChinaIP_{it} | X_{it}, i, t\right] \neq 0$, the coefficient β gives the relationship, causally induced by a supply shock in China, between a 1 percentage point increase in the import penetration rate from China and the level of inflation faced by U.S. consumers.[22]

We start with just-identified IV specifications with a single instrument, using in turn the NTR gap and the change in import penetration from China in the other developed economies. Since the NTR gap is relevant only after 2000 (after the policy change), we set $\mathbf{Z_{it,1}} = (NTR\,Gap_i \cdot PostPNTR_t)$, with $PostPNTR_t = 1$ for the period 2000-2007. The change in import penetration from China in the other developed economies offers variation in both periods: $\mathbf{Z_{it,2}} = \Delta ChinaIP\,Other_{it}$. After using the two instruments separately, we use them jointly.

Results. Figure 2 reports binned scatter plots depicting the first-stage and reduced-form specifications using the NTR gap as the instrument. Panel A shows the first stage: the larger the NTR gap, the larger the increase in import penetration from China. Panel B depicts the reduced-form relationship: the CPI inflation rate is significantly lower in ELIs with a higher NTR gap.

Panel A of Table 2 reports the baseline IV estimates, using the NTR gap as the source of identifying variation in equation (3). The OLS coefficients for the first-stage and reduced-form relationships are reported in Columns (1) and (2). We find that a 10 percentage point increase in the NTR gap leads to an increase in the import penetration rate from China of 33.3 basis points and to a fall in the inflation rate of 74.3 basis points. These findings indicate that the policy change had a large impact of trade with China and on consumer prices.

The IV estimate in Column (3) indicates that a one percentage point increase in the import penetration rate from China leads to a fall in inflation of 2.23 percentage points. This coefficient is precisely estimated, with a 95% confidence interval ranging from -3.15 to -1.30. The F-statistic indicates that the instrument is strong. In Column (4), we run the same specification with OLS. The relationship between trade and prices remains large, but not as strong as with the instrument: the coefficient is -0.90, which suggests that omitted variables biases or reverse causality attenuate the estimated relationship between trade and consumer prices by over 50%. Finally, since the policy change was implemented in 2000, it is instructive to restrict the sample to the second period only (2000-2007). With only one period, ELI fixed effects would absorb the entire variation in the data, therefore we only include fixed effects for durable goods and apparel. Column (5) reports

[22] All specifications use consumption weights. Results without weights are similar, see Appendix Table A4.

the results: the estimated IV coefficient of -2.15 remains large, statistically significant, and is very similar to the baseline in Column (3). The standard errors increase by over 60% in Column (5) relative to Column (3), which shows it is useful to leverage the full sample with ELI fixed effects to increase power.

Panel B of Table 2 present the results using as an instrument the change in trade with China in other developed economies. Column (1) reports the corresponding IV coefficient: when import penetration from China increases by one percentage point, consumer prices fall by 1.44%. Column (2) repeats the specification after restricting the sample to the second period only: the IV coefficient remains similar, equal to -1.27. These coefficients are precisely estimated and the F statistics are strong.

Column (3) reports the IV estimate when using both instruments jointly. The IV coefficient is -1.91 and is precisely estimated, with a standard error of 0.38. Because we now have an over-identified equation, we can run the test of over-identifying restrictions of Hansen (1982). With a p-value for the J statistic of 0.21, we cannot reject that the over-identification restrictions are valid.[23]

III.D Robustness

We now implement several robustness tests. In Columns (1) through (4) of Panel A of Table 3, we examine whether the estimates remain stable as we change the set of fixed effects and time varying controls, and we examine their sensitivity to the exclusion of outlier categories with particularly low inflation rates. The NTR gap is used as the instrument. Column (1) repeats the IV specification after replacing the 222 ELI fixed effects with a set of fixed effects for ten broad product categories defined by the BLS (called "major categories", they are defined by the first 2 digits of each ELI). The IV coefficient falls slightly to -2.75. In Column (2), we re-introduce ELI fixed effects as in the baseline but also include period-specific fixed effects for apparel and durable goods (i.e., inflation can vary systematically across periods). The IV coefficient remains similar, increasing slightly to -1.78. In Column (3), we exclude ELIs in the bottom 10% of the inflation distribution over our sample. Doing so leaves the IV coefficient virtually unchanged compared to the baseline, at -2.26. In Column (4), we control for inflation in the 1990s, interacted with period fixed effects. The results remain unchanged. Column (5) repeats the specification after including the period-specific controls used by Pierce and Schott (2016) (a fixed effect for high-tech sectors and controls for contract intensity and union membership), which again leaves the IV coefficient almost unaffected, at -2.10. We also assess the sensitivity of our results by aggregating the data to the level of coarser industries,

[23] The Hansen test could reject due to heterogeneous treatment effects, even when the exclusion restrictions hold. Therefore, rejecting would not necessarily be a sign that either of the instrument is invalid. Although our two instruments leverage different sources of variation, we can hypothesize that there is a common value of β. The J test shows that we cannot reject this hypothesis.

the 6-digit IO industries defined by the BEA's 2007 input-output table, with 170 industry-by-period observations. In Column (6), the IV coefficient is -2.94 and is significant at the 5% level.

The estimated effects so far do not account for the possibility of correlated changes in the patterns of overall trade across ELIs. Trade with other countries may be a source of omitted variable bias. If other countries tend to *increase* their import penetration in the U.S. at the same time as China, then we might conflate the price effects of China with those of other trading partners. Column (1) of Panel B of Table 3 shows that, in fact, China tends to *displace* other trading partners of the US: when import penetration from China increases by 1 percentage point, overall import penetration increases by only 78 basis points. The IV coefficient based on the overall change in import penetration is larger than when considering trade with China alone, because overall trade increases by *less* than the change in trade with China alone suggests. Columns (2) and (3) report the results when instrumenting the change in overall import penetration with the NTR gap. The IV coefficients fall further, to approximately -3.70.

Although we view the estimation of causal effects of trade on consumer prices as a key contributions of this paper, due to space constraints we report several additional robustness checks in Appendix C.A.

III.E Falsification Tests using French CPI Data

One potential confounding factor remains unaddressed so far: unobserved time-varying shocks could bias our estimates. With this in mind, we implement a placebo test using data from the French CPI, which is publicly available across 132 detailed product categories called COICOP. We link our main analysis sample to COICOP categories, aggregate all variables at that level, and repeat the IV strategy from equation (3) with the French CPI as the outcome, using the NTR gap as the instrument.

Panel A of Figure 3 reports the placebo reduced-form. There is no relationship between the NTR gap and inflation across product categories in France. Panel A of Table 4 reports the results: the first-stage in Column (1) and the IV estimate for the U.S. in Column (2) are similar to the preceding analysis, except that we now run the regression across COICOP categories rather than ELIs. Column (3) shows that the reduced form with the French CPI is not significant, and Column (4) reports a precisely estimated null IV coefficient with the French CPI, at -0.074 (s.e. 0.38). The coefficient remains small and insignificant with the alternative specification in Column (5).

We also estimate a triple-difference IV specification, reporting the estimates in Panel B of Table 4. The specification is the same as in equation (3), except that the outcome is now the *difference* between U.S. CPI inflation and French CPI inflation. With this differenced outcome, we effectively control for category-by-period fixed effects for inflation; the estimated IV coefficient only depends on inflation differences between the United States and France for the same product category. We still allow for COICOP fixed effects, i.e. for permanent differences in inflation rates between the U.S. and France for each product category. Panel B of Figure 3 depicts the clear negative reduced-form relationship with this differenced outcome. Column (2) of Panel B of Table 4 reports the corresponding coefficient. The IV coefficients in Columns (3) and (4) are similar to the baseline results, ranging from -2.08 to -2.52, and are statistically significant at the 5% level. These facts alleviate the remaining concerns over time-varying unobserved shocks.

IV Mechanisms

We now investigate a series of potential mechanisms that could account for the estimated price effects of trade with China across product categories. We show how to interpret the estimated price effects in light of standard quantitative trade models. Using statistical decompositions, we demonstrate the important contributions of continued and domestic products to the overall price effects. Finally, we study heterogeneity in the estimated price effects across product categories to distinguish between different potential channels that could explain the response of domestic products, including intermediate inputs, offshoring, changes in wages and TFP, and markups. We find empirical support primarily for the markup channel. Using a simple model of strategic price setting, we also establish that the markup channel is plausible quantitatively.

IV.A Connecting the IV Estimate to the Melitz-Chaney Model

Our IV estimate is a useful identified moment that can serve as a diagnostic tool to distinguish between classes of trade models, *à la* Nakamura and Steinsson (2018). We connect our IV estimate to the standard quantitative trade model of Melitz (2003) and Chaney (2008), with N countries engaging in international trade. By perturbing the model with sector-specific productivity shocks in China and estimating equation (3) within the model, we derive the structural expression for the reduced-form cross-sector regression coefficient, $\hat{\beta}$. Connecting our regression specification to theory makes clear that the estimate for $\hat{\beta}$ explicitly embeds changes in general equilibrium, including endogenous changes in wages and product variety, as well as potential heterogeneous effects across sectors.[24]

The representative consumer in country j has Cobb-Douglas preferences across S sectors indexed by s, $U_j = \Pi_{s=1}^{S} \left(Y_j^s\right)^{\mu_s}$ with $\sum_s \mu_s = 1$. Each sector consists of differentiated varieties over which the representative consumer has CES preferences with an elasticity of substitution $\sigma > 1$,

[24] Our approach thus demonstrates formally how to connect our regression specification to the theoretical prediction of the model, accounting for (i) the fact that the cross-sector regression cannot account for the "missing intercept", i.e. we cannot use standard expressions derived in prior work relating changes in the domestic expenditure share and welfare at the macro level (e.g., Arkolakis et al. (2012), Ossa (2015)); (ii) all GE effects, some of which affect the cross-industry slope because they differ across sectors (e.g., changes in product variety) while we demonstrate that others are differenced out in the cross-sector specification (e.g., changes in domestic wages); (iii) the fact that the measured price index may differ from the theoretical, model-consistent price index.

$Y_j^s = \left(\sum_{i=1}^N \int_{\Omega_{ij}^s} y_{ij}^s(\omega)^{\frac{\sigma-1}{\sigma}} d\omega\right)^{\frac{\sigma}{\sigma-1}}$, where Ω_{ij}^s is the set of varieties from sector s available to the consumer in country j and produced in country i, and $y_{ij}^s(\omega)$ is the quantity of each such variety $\omega \in \Omega_{ij}^s$.

Firms in country i have aggregate productivity A_i^s and idiosyncratic productivity z that is distributed according to $G_i^s(z)$. They are monopolistic competitors within a sector and produce varieties according to a linear production technology that takes labor as an input, $y_{ij}^s(z) = A_i^s z l_{ij}^s$. Firms take the economy-wide wage w_i as given, face an iceberg cost τ_{ij}^s and pay a labor-denominated fixed cost $w_i f_{ij}^s$ to operate in the market, which they pay as long as their profits are positive. In this setting, firms' markups are constant and there is an endogenous entry cutoff. The endogenous domestic wage changes induced by the foreign supply shocks is denoted $d\log(w_d)$. Finally, we denote the trade elasticity by θ, which is the Pareto shape parameter when the idiosyncratic firm productivity distribution is Pareto. For completeness, we also consider a setting with heterogeneous Pareto shape parameters across sectors, denoted θ^s. Appendix D.A.1 provides additional detail about this standard setting.

We now perturb the equilibrium with shocks to foreign country f's marginal costs of production, which are heterogeneous across sectors, and we analyze the general equilibrium cross-sectoral relationship between changes in the domestic expenditure share and changes in the price index for domestic consumers. Proposition 1 provides a structural interpretation for the estimated regression coefficient, $\hat{\beta}$, as the sectoral average relationship between the exact consumer price index in each sector s, denoted P_s, which includes both domestic and imported items, and the expenditure share on domestically-produced items in each sector s, denoted S_{sd}.

Proposition 1 ***[changes in spending shares and exact price indices in the Melitz-Chaney model].*** *Sector-specific supply shocks to firms from trade partner f induce a sector-by-sector relationship in the domestic economy between overall price indices and domestic expenditure shares,*

$$d\log(P_s) = \frac{1}{\theta} d\log(S_{sd}) + d\log(w_d), \tag{4}$$

holding in general equilibrium with endogenous wages and firm entry-and-exit. The reduced-form empirical specification $\Delta\log(P_s) = \alpha + \beta\Delta\log(S_{sd}) + \varepsilon_{sd}$ yields the estimated regression coefficient $\hat{\beta}$ with the following structural interpretation:

$$\hat{\beta} = \frac{1}{\theta}. \tag{5}$$

Moreover, if trade elasticities are heterogeneous across sectors, then the estimated regression coefficient is $\hat{\beta} = \sum_s \omega_s \frac{1}{\theta^s}$, where the sector weights satisfy $\sum_s \omega_s = 1$, $\omega_s \geq 0$, and are provided in Appendix D.A.3.

Proof: see Appendix D.A.2-D.A.3.

Proposition 1 shows that, with the log change in the domestic expenditure share as the endogenous variable, the Melitz-Chaney model predicts that our IV specification (3) should recover an estimate related to the inverse of the trade elasticity: $\widehat{\beta} = \frac{1}{\theta}$.[25] In Appendix D.A.4, we show that this prediction holds more broadly in the set of models considered by Arkolakis et al. (2012). Given that the standard estimate for the trade elasticity is $\theta \approx 4.25$ (Simonovska and Waugh (2014)),[26]using the log change in the domestic expenditure share as the endogenous variable, standard trade models predict an IV coefficient of $\widehat{\beta} = \frac{1}{\theta} \approx 0.23$.

To test this prediction, we implement our IV specification with the log change in the domestic expenditure share as the endogenous variable.[27] The estimates are reported in Table 5 and are large in magnitude: 2.57 (s.e. 0.96) with the instrument from Pierce and Schott (2016), 3.46 (s.e. 1.41) with the instrument from Autor et al. (2014), and 3.10 (s.e. 0.96) with both instruments.[28] These estimates are much larger than the predicted IV coefficient of 0.25. The trade elasticity that would be necessary to match the point estimate (with both instruments) is $\hat{\theta} = \frac{1}{3.10} \approx 0.32$. This trade elasticity is implausibly small: benchmark estimates are typically above 1 and generally close to 4.[29] These results show that standard quantitative trade models do not match the estimated price response. In the presence of heterogeneous trade elasticities across industries, Proposition 1 shows that the regression coefficient is in the convex hull of the inverse of the sector-specific trade elasticities. By Jensen's inequality, $\widehat{\beta} > 1/\mathbb{E}[\theta^s]$: in Appendix Table A5, we compute the weighted average in Proposition 1 and obtain $\widehat{\beta} = \sum_s \omega_s \frac{1}{\theta^s} = 0.39$. Thus, the heterogeneity in trade elasticities leads to a substantial increase in the coefficient but remains insufficient to account for our large IV estimate.[30]

Proposition 1 shows that the IV estimate is about ten to five times too large relative to what would be expected from the change in trade flows, according to standard trade models. However, because of endogenous entry and exit, the changes in the exact CES price index characterized in Proposition 1 could differ from changes in the "measured" Consumer Price Index we use in

[25]When the trade elasticity varies across product categories, our IV estimator recovers a weighted average of the trade elasticities.

[26]As noted by Arkolakis et al. (2011), the trade literature and the international macro literature do not agree on the value of the trade elasticity. Macro models, which focus on short-run fluctuations, generally set a low value for this parameter (e.g., an Armington elasticity of 1.5 in Backus et al. (1993)). In contrast with the international macro studies, our empirical analysis focuses on medium-run responses, for which the elasticities from the trade literature are the natural benchmark.

[27]The specification is given by 3, except that $\Delta ChinaIP_{it}$ is replaced by the log change in the domestic expenditure share.

[28]Our baseline specification estimates the relationship between the change in the overall price index and the change in import shares from China across sectors. The effect is larger when we run the theory-consistent specification, estimating the price effect relative to the change in the domestic spending shares as in Proposition 1. This finding implies that the baseline specification is, if anything, attenuated relative to the standard theory-based specification we pursue here.

[29]In recent work, Boehm et al. (2022) estimate smaller trade elasticities. At a one-year horizon, the trade elasticity is 0.76, which is still over twice as large as the required elasticity to match our estimates. Their estimates increase to about 1.75-2.25 after seven years.

[30]The results are similar when we adjust the domestic expenditure share by retail margins (unreported).

our empirical analysis, which does not account for changes in product variety. For example, if there is significant product exit, the exact price index will be smaller than the measured CPI. We characterize this difference in Corollary 1.1, denoting by $\widetilde{P}_s$ the measured price index for U.S. consumers for sector s, by $\widetilde{P}_{sd}$ the measured price index over domestically-produced goods, and by S_{si} the expenditure share on country i.

Corollary 1.1 ***[changes in spending shares and measured CPI in the Melitz-Chaney model].*** *The sectoral exact price index can be decomposed into the measured CPI and an unobserved product variety correction, such that every sector s satisfies:*

$$\underbrace{\frac{d\log(P_s)}{dS_{sd}}}_{exact\,price\,index} = \underbrace{\frac{d\log(\widetilde{P}_s)}{dS_{sd}}}_{measured\,CPI} + \underbrace{\frac{\sum_i S_{si}\cdot d\log(E_{si})}{dS_{sd}}}_{product\,variety},$$

with $d\log(E_{si})$ *the correction for entry-exit defined in Appendix D.A.5, equation (D11).*

Moreover, denoting the measured CPI over domestic products by $\widetilde{P}_{sd}$*, the estimated regression coefficient in the reduced-form specification,* $\Delta\log(\widetilde{P}_{sd}) = \alpha + \beta\Delta\log(S_{sd}) + \varepsilon_{sd}$*, is* $\hat{\beta} = 0$.

Proof: see Appendix D.A.5.

Available evidence suggest that overall product variety (weakly) increases in response to foreign supply shocks.[31] In that case, the exact price index should fall more than the measured CPI. Thus, the first part of Corollary 1.1 shows that the divergence between the theoretical and measured price index cannot help standard models match our IV estimate.[32] Furthermore, the second part of Corollary 1.1 derives a sharp prediction: in the Melitz-Chaney model, there is no cross-sector relationship between the prices of domestically-produced goods and foreign supply shocks, which we proceed to test in the next subsection.

IV.B The Roles of Continued and Domestic Products

In order to understand the discrepancy in magnitudes between our IV estimate and the theoretical relationships above, next we empirically document that U.S. prices did in fact respond to increased import penetration from China. We begin our empirical investigation of potential mechanisms with a simple statistical decomposition by product characteristics, documenting the extent to which new products and products that were made in the U.S. generated the estimated price effect.

[31] Bai and Stumpner (2019) measure changes in product variety using barcode data and estimate that increased trade with China led to a change in product variety that lowered the cost of living for U.S. consumers, using the Feenstra (1994) adjustment factor in a CES framework.

[32] Furthermore, in Appendix D.A.5 we use the Melitz-Chaney model to derive a bound showing that, with unmeasured entry-exit, the IV estimate using the measured CPI is predicted to be inferior to 0.25. We conclude that unobserved changes in product variety cannot reconcile the predictions of the Melitz-Chaney model with our IV estimate.

Statistical decompositions. We will begin our decomposition of the estimated price effect by denoting the subset of interest by A, and let s_i^A be the share of items within product category i that belongs to subset A (which, as we define below, will correspond to continued products and domestic products). Omitting time subscripts and letting B denote the complementary set, we obtain an exact decomposition for the CPI inflation rate for each product category as

$$\pi_i = \underbrace{s_i^A \pi_i^A}_{\equiv \widetilde{\pi}_i^A} + \underbrace{s_i^B \pi_i^B}_{\equiv \widetilde{\pi}_i^B}, \tag{6}$$

where π_i is the inflation rate for product category i as in Section III, π_i^A is the inflation rate for products within subset A and s_i^A is the spending share on A. Finally, $\widetilde{\pi}_i^A$ is the contribution of subset A to overall inflation in category i, which depends both on the inflation rate within A and on how much spending is devoted to A.

In the remainder of this section, we examine the contributions of continued goods and domestic goods to overall inflation. We first run our IV specification (3) with continued goods inflation or domestic goods inflation as the outcome (π_i^A), which is directly informative about the price response for these sets of goods. These results do not provide a proper decomposition because they ignore the share of spending on the relevant set of products. If a set of products accounts for a small share of spending, its overall impact on category-level inflation may be small even if it has a large inflation response to trade. Therefore we repeat the IV specification with the share-adjusted inflation rate ($\widetilde{\pi}_i^A$) as the dependent variable. By linearity of OLS, the ratio of the IV coefficient with the share-adjusted inflation rate ($\widetilde{\pi}_i^A$) to the baseline IV coefficient (with outcome π_i) gives the share of the overall effect accounted for by products within subset A.

In the remainder of this Section we focus on the NTR gap instrument. The results with the change in import penetration from China in other countries are similar (Appendix Table A6).

The role of continued products. Panel A of Table 6 documents the impact of trade with China on "inflation for continued products", which is defined as inflation for the set of products which are available across consecutive periods. Continued products inflation excludes new products (termed "product substitutions" by the BLS) from the computation of inflation. This decomposition allows us to test whether the overall price response to trade stems from declining prices for new products (i.e., inflation would fall via product substitutions) or from declining prices for pre-existing products (continued products inflation). Across all specifications, we find a robust pattern of lower inflation for continued products in response to increased trade with China. Columns (1) and (2) indicate that inflation for continued products falls by 3 percentage points for each 1 percentage point increase in import penetration from China. Using the decomposition in equation (6), Columns (3) and (4) show that continued products account for approximately 70% of the overall price effects from Table

2.

The role of domestic products. We now examine the contribution of domestic products. We first continue working with the CPI data set, before presenting additional evidence from the PPI data set.

Evidence from the CPI sample. To assess whether the price effects are driven by U.S. goods as opposed to foreign (Chinese) goods, we identify U.S. goods in the CPI using specification checklists. For each product in the CPI, characteristics are recorded in specification checklist files. We use the specification checklists to gather information on the country of origin for each product and then repeat each estimation exercise on subsamples of U.S. products. While checklists for some categories of items have explicit flags for country of origin information (e.g., "Was the product made in the United States; Yes or No?"), others have entries that the data collectors populate with text (e.g., "Write in the country in which the product was made.").[33]

Panel B of Table 6 reports the response of prices to trade with China when only taking into account U.S. goods in the CPI sample. Columns (1) and (2) show that prices for domestic goods experience a large fall, similar to the full sample, with point estimates ranging from -1.94 to -2.73 across specifications. Using the statistical decomposition, Columns (3) and (4) show that domestic prices change account for a substantial fraction of the overall price effects, between 44% and 85% depending on the specification.

Evidence from the PPI sample. We assess the robustness of our results using the producer prices from the Producer Price Index. The PPI sample only takes into account price changes for products manufactured in the U.S. We run an IV specification identical to (3), except that the outcome variable is now the PPI inflation rate, and that the level of observation is a 6-digit NAICS code.

Panel C of Table 6 reports the results in the PPI sample. Columns (1) and (2) show that the prices of domestic U.S. manufacturers fall in response to trade with China. The point estimates are very similar to the CPI sample, ranging from -2.50 to -1.86 across specifications, and are statistically significant at the 1% level. Columns (3) and (4) show that the point estimates remain similar in magnitude, falling by about 20 to 50 basis points only, when we repeat the estimate with PPI inflation for continued products as the outcome. These results confirm the importance of continued and domestic products in accounting for the overall price effects.

[33] Appendix B.A describes the specification checklists and the parsing algorithms we use to retrieve countries of origin from text entries. Appendix Table A7 reports summary statistics on the number of product categories with explicit flags for country of origin.

IV.C The Role of Changes in Domestic Production Cost

The response of domestic prices could result from changes in production cost for U.S. manufacturers, or changes in markups. With a standard production function, cost minimization yields the change in the domestic production cost,

$$\Delta \log(c_i) = -\Delta \log(A_i) + \alpha_i^K \Delta \log(r_i) + \alpha_i^L \Delta \log(w_i) + \alpha_i^I \Delta \log(p_{I,i}), \tag{7}$$

where A_i is total factor productivity, the factor shares $\alpha_i^K, \alpha_i^L, \alpha_i^I < 1$ sum to one. r_i is industry i's rental rate for capital, $w_i$is industry i's wage, and $p_{I,i}$ is the price of a composite bundle of intermediate inputs to industry i. We now investigate whether changes in domestic production costs across industries can account for the price effects across industries, using various proxies for the terms in equation (7).

We test each cost channel one-by-one in order to present the evidence in the simplest manner, and find that the point estimates are much too small relative to what would be necessary to account for the baseline estimate of the price effect. Thus, second-order terms (e.g., joint tests that include covariation between costs) are unlikely to generate larger effects and better account for the baseline estimate.

Imported intermediate inputs. The measure of Chinese import penetration we have used so far is meant to reflect exposure to import competition, not to imported intermediate inputs. But it could be the case that an industry's change in import penetration from China happens to be correlated with changes in Chinese import competition faced by that industry's domestic suppliers. Similarly, if the industry sells to other domestic producers, then the Chinese import penetration measure could be correlated with import competition faced by downstream industries. Conceptually, exposure to rising import penetration from China via buyer-supplier linkages could be a source of omitted variable bias across product categories.[34]

To examine whether buyer-supplier linkages affect our results, we first compute the correlations between our baseline measure of import competition and indirect exposure via domestic suppliers or domestic buyers. The results are reported in Figure 4 and Appendix Table A8.[35] We find that the correlations are positive but small: when the import penetration rate from China increases by 1 percentage point in industry j, the share of intermediate inputs from China in industry j's total output increases by only ten basis points, and there is only a 2 basis point increase in import competition via domestic buyer industries.

In Panel A of Table 7, we directly establish that the price effects are not driven by I-O linkages

[34] Note that our focus is to investigate the observed relationship between changes in trade and changes in prices *across industries*. This exercise is conceptually different from an assessment of the role of intermediate inputs in the gains from trade (for example, see Ossa (2015)).

[35] Appendix C.B discusses the data construction steps.

by repeating our IV specification from equation (3) while controlling for indirect exposure to trade with China via suppliers or buyers. Controlling for supplier and buyer effects in turn (Columns (2) and (3)) or jointly (Column (4)) yields stable point estimates hovering between -2.89 and -3.24, which are very close to the baseline result of -2.94 in Column (1).[36]

Next, we highlight standard quantitative models with intersectoral linkages predict much smaller price effects than our IV estimates. To introduce intersectoral linkages into the Melitz-Chaney model of Section IV.A, we follow Caliendo and Parro (2015) and assume that firms use a linear production technology using labor and intermediate inputs: $y_{ij}^{s}(z) = A_i^s z(\ell_{ij}^s)^{1-\alpha_k}\left[\Pi_{s'=1}^{S}\left(m_{ij}^{s,s'}(z)\right)^{\zeta_{s,s'}}\right]^{\alpha_s}$, where $m_{ij}^{s,s'}$ denotes the composite intermediate good from sector s' used in production by firms from country i selling in j in sector k. The share of intermediate inputs in value added is denoted by α_k, and intermediate inputs have shares that sum to one, i.e. $\sum_{k'}\zeta_{k,k'} = 1$. Given our empirical finding that industries that are more exposed to import competition are only slightly more exposed to the imported intermediate inputs channel (Figure 4), we assume that the change in the price index for intermediates experienced by a given sector is smaller than the price change of the sector itself. This assumption delivers the bound on the IV estimate reported in Corollary 1.2.

Corollary 1.2 ***[bound for changes in import shares and exact price indices with intersectoral linkages].*** *Assume that $\Sigma_{k'\neq k}\zeta_{k,k'}d\log(P_j^{k'}) < d\log(P_j^k)$ and that the distribution of firms' productivity is Pareto with shape parameter θ. Then supply shocks for trading partner f induce a cross-sector relationship in the domestic economy between import shares and consumer price indices in general equilibrium with endogenous wages and entry-exit, such that the estimated regression coefficient in the reduced-form specification, $\Delta\log(P_j^k) = \alpha + \beta\Delta\log(S_{jj}^k) + \varepsilon_{kj}$, satisfies:*

$$\hat{\beta} < \sum_k \omega_k \frac{1}{\theta(1-\alpha_k)},$$

with weights ω_k satisfying $\sum_k \omega_k = 1$.

Proof: see Appendix D.A.5.

Corollary 1.2 shows that, with intermediate inputs, equation (**??**) needs to be adjusted by a factor $(1-\alpha_k)$. According the the BEA input-output table, in our sample $\alpha_k = 56.4\%$, implying $\hat{\beta} < 0.53$, which is much smaller than the IV estimate. This result formalizes the idea that intermediate inputs cannot explain the magnitude of the IV coefficient.

[36] These results should not be interpreted as demonstrating that imported intermediate inputs play no role (see, e.g., De Loecker et al. (2016)), but they show that it is implausible for the intermediate inputs channel to explain our large IV estimates, for the simple reason that our main measure of import penetration captures import competition and is not strongly correlated with imports of intermediate inputs.

Offshoring. Intermediate inputs may play a role independent of I-O linkages. For example, a U.S. manufacturer of water bottles could use plastic imported from China, in which case imported intermediate inputs would be accounted for by our I-O analysis above, because "plastic" and "water bottles" are distinct product categories. But if the U.S. producer offshores production to China and re-imports the finished product (i.e., the water bottle, not plastic), then the I-O analysis would not accurately account for trade-induced changes in production cost. The potential concern is that increased trade with China in an industry does not correspond to intensified import competition, but rather to an increase in "offshoring" trade between related parties.

We examine the importance of this potential channel using the related-party trade database of the U.S. Census Bureau, as in Antràs and Chor (2013). Related-party trade includes trade by U.S. companies with their subsidiaries abroad as well as trade by U.S. subsidiaries of foreign companies with their parent companies. If offshoring drives the price effects, we expect to find larger effects (for a given increase in trade with China) in product categories where related-party trade accounts for a larger fraction of trade with China.

The share of trade with China occurring between related parties is very low during the period we study, with a median of 4% (Appendix Table A9). Although these summary statistics suggest that offshoring may not drive our results for the average category, related-party trade is important for a small fraction of product categories: the 90th percentile of the distribution of related-party shares is 38%. In Column (1) of Panel B of Table 7, we repeat our IV specification after interacting the endogenous variable with an indicator for categories with a share of related-party trade with China above the 90th percentile; the instrument is also interacted with this indicator. We find that the estimated price effects remains stable and that the interaction term is not significant.

Returns to scale and productivity. Increased import competition with China could affect domestic production costs by displacing domestic goods and reducing the scale of domestic production. To rationalize the evidence in Table 6 through decreasing returns to scale, the marginal cost of production should fall by 2% as domestic production falls by 1% (due to displacement by China). Recent empirical studies have estimated this elasticity (e.g., Costinot et al. (2019), Jaravel (2019) and Faber and Fally (2017)). They find that for tradable U.S. industries returns to scale are *increasing*, with elasticities of prices to quantities ranging between -0.1 and -0.4; we would need an elasticity of the opposite sign and five times larger in magnitude.

In our context, two mechanisms could potentially yield an elastic marginal cost of production: industry-specific factors and endogenous changes in technology. If an industry relies on industry-specific factors, then a fall in production could lead to a substantial fall in production costs, because the supply curve is inelastic for these factors. For example, capital investments may be irreversible, in which case the industry-specific rental rate of capital may fall substantially as quantities fall. According to equation (7), this effect should be particularly important for capital-intensive industries (through the term $\alpha_i^K \Delta \log(r_i)$ in equation (7)).

In Column (2) of Panel B of Table 7, we examine whether the magnitude of the price effect varies across sectors depending on their capital intensity. Using an indicator variable for industries above the median capital intensity, we find no heterogeneity in the effect.

Another possibility is that import competition may affect productivity through endogenous technology. If increased competition spurs domestic firms to adopt or invent cost-reducing technologies (e.g., Bustos (2011), Bloom et al. (2016), Aghion et al. (2018)), then change in productivity could rationalize our results, through the term $\Delta \log(A_i)$ in equation (7). However, recent evidence about the China shock in the United-States suggests that innovation by domestic firms fell in response to the shock (Autor et al. (2020a)).

To further examine the potential productivity channel, we use our IV framework to examine the response of Total Factor Productivity, as measured in the NBER-CES database for manufacturing industries.[37] In Panel C of Table 7, Columns (1) and (2) report that both TFP measures fall in response to increased trade with China, which is consistent with the evidence from Autor et al. (2020a) using patent data. We caution that the evidence on TFP should only be viewed as suggestive, because we do not have access to the underlying micro data and cannot investigate the sensitivity of the estimates to alternative measures of TFP.

Wages. Changes in wages across industries could be another reason for changes in domestic production cost. Although this channel is theoretically plausible, we find that in practice it can explain little of the evidence on domestic prices.

The first piece of evidence is that industries exposed to trade with China are not very labor intensive: the labor share of total cost is small ($\alpha_{L,i}$ in equation (7)). The NBER-CES Manufacturing database linked to our sample indicates that the share of labor in total value added for product categories within manufacturing was about 27% in our sample period. Furthermore, the share of labor in total domestic output is only 10.9%, because these industries use intermediate inputs intensively. To explain a 2% fall in domestic prices due to increased import penetration from China, the wage response should be very large, on the order of 20%.

Using worker-level administrative data provided by Autor et al. (2014), we find that a one percentage point increase in the import penetration rate from China leads to a 39.3 basis point fall in wages (Column (3) of their Table III). In Panel C of Table 7, we use public wage data from the NBER-CES and County Business Patterns databases. We find no significant wage effects, either in County Business Patterns data for all workers (Column (3)), or in NBER-CES data for production

[37] The 5-factor TFP measures uses non-production workers, production workers, energy, materials and capital. The 4-factor TFP measure is calculated similarly, but using total materials cost spending rather than separating it into energy and non-energy materials. These measures attempt to capture TFPQ, as defined in Hsieh and Klenow (2009).

or non-production workers (Columns (4) and (5)). Although the public data is imperfect and likely fails to capture the negative effects that Autor et al. (2014) were able to estimate precisely, we can confidently rule out the large wage changes that would be required to meaningfully affect domestic production costs.

IV.D The Role of Changes in Markups

Having established that changes in domestic production costs are unlikely to drive the price effects, we now examine the potential relevance of domestic markups.

Connecting the IV specification to oligopolistic competition models. We start with a simple theoretical exercise: could changes in markups plausibly explain the observed domestic price response, or are the observed price effects too large? We first show that, in a flexible model of oligopolistic competition, the domestic price response is predicted to be of the same magnitude as the overall price response, consistent with the empirical evidence from prior sections. We then highlight that, in a stylized Cournot competition model, conventional parameter choices are consistent with the magnitude of the price response observed in the data, as well as with auxiliary evidence on the response of domestic markups.

We consider a standard setting following Edmond et al. (2015), where the economy consists of two countries, Domestic and Foreign, with a single factor of production, labor, that is in inelastic supply and immobile between countries.[38] We focus on the domestic country throughout. A representative consumer has nested-CES preferences in which oligopolistic competition arises from the existence of only a finite number of competitors within nests. Specifically, there is a continuum of sectors indexed by $s \in [0,1]$ such that consumers' utility is given by $Y = \left(\int_0^1 y(s)^{\frac{\epsilon-1}{\epsilon}} ds\right)^{\frac{\epsilon}{\epsilon-1}}$, where $\epsilon > 1$ is the elasticity of substitution across sectors. Each sector s consists of a finite number of domestic and foreign intermediate producers, such that consumers' consumption aggregator in each sector combines n_{sd} domestic and n_{sf} imported products, with $y(s) = \left(\sum_{j=1}^{n_{sd}} y_j^d(s)^{\frac{\gamma-1}{\gamma}} + \sum_{j=1}^{n_{sf}} y_j^f(s)^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}}$,where $\gamma > \epsilon$ is the elasticity of substitution across goods j within a sector s. In our data, sectors correspond to ELIs.

Within a sector, firms produce using a linear production technology taking labor as an input, $y_j^i(s) = a_{sij} l_{sij}$, where producer-level productivity a_{sij} is drawn from a sector-specific distribution in each country $i \in \{d, f\}$. Firms take as given their country's wage w_i, face an iceberg trade cost τ_{si}, and pay a labor-denominator fixed cost $w_i f_{si}$ to operate in the market.

Firms compete oligopolistically within a sector. We first derive a proposition based on a non-parametric representation of oligopolistic competition using markup elasticities, as in Amiti et al. (2019a). Let p_{sij} and μ_{sij} denote respectively the price and markup of firm j from country i in

[38] This setting is also very close to that in Atkeson and Burstein (2008).

sector s. A change in a firm's marginal cost passes-through into its own price at a rate given by $\frac{1}{1+\Gamma_{sij}}$, where $\Gamma_{sij} \equiv -\frac{\partial \log(\mu_{sij}(.))}{\partial \log(p_{sij})}$ is the firm's "own-price markup elasticity". Under monopolistic or perfect competition, $\Gamma_{sij} = 0$.

We now perturb the equilibrium with a change in Foreign's productivity. Proposition 2 provides a decomposition isolating the role of domestic prices, and gives a bound for the change in the domestic price index, relative to the overall price index, in response to these supply shocks. We use the following notation: P_s denotes the overall consumer price index in sector s, P_{sd} the consumer price index for domestically-produced goods in sector s, S_{sd} the expenditure share on domestically-produced goods, and S_{sf} on foreign products, S_{sij} the expenditure share on firm j from country i's output and, finally, η_s denotes the sectoral trade elasticity with respect to a change in productivity (see equation (D25) of Appendix D.C.2).

Proposition 2 ***[decomposition and bound for the price response in an oligopolistic competition model].*** *Perturbing the equilibrium with a change in Foreign's productivity, for each sector s the relationship between price changes for domestic consumers and changes in the import penetration rate from Foreign is:*

$$d\log(P_s) = \frac{1}{(\gamma-1)} d\log(S_{sd}) + d\log(P_{sd}),$$

with a domestic price response of

$$d\log(P_{sd}) = \frac{1}{(\gamma-1)} \frac{1-\Omega_{sd}}{\Omega_{sd}} d\log(S_{sd}) + d\log(w_d),$$

where $\Omega_{sd} = \sum_{j=1}^{n_{sd}} \frac{S_{sdj}}{S_{sd}} \frac{1-S_{sdj}}{1-S_{sdj}+\Gamma_{sdj}}$ and the domestic markup response is given by

$$\frac{1-\Omega_{sd}}{\Omega_{sd}} = \sum_{j=1}^{n_{sd}} \left(\frac{\frac{S_{sdj}}{S_{sd}} \frac{1-S_{sdj}}{1-S_{sdj}+\Gamma_{sdj}}}{\Omega_{sd}} \right) \frac{\Gamma_{sdj}}{1-S_{sdj}}.$$

Moreover, the difference between the change in the overall sectoral price index and in the sectoral price index for domestic products satisfies the following bounds:

$$0 \leq -\left(\frac{d\log(P_s)}{dS_{sf}} - \frac{d\log(P_{sd})}{dS_{sf}} \right) \leq \frac{1}{\eta_s \cdot S_{sd}}.$$

Proof: see Appendix D.C.2.

The first part of Proposition 2 shows that the overall price response can be decomposed into two terms: while the first term is analogous to the Melitz-Chaney model, capturing the substitution between domestic and foreign goods as in Proposition 1, the second term reflects the domestic price response due to changes in domestic markups, which is equal to zero when markups are constant, i.e. markup elasticities Γ_{sdj} are null.

Corollary 2.1 below shows that the insights from Proposition 2 carry over to our regression analysis, providing a structural interpretation for the estimated coefficient $\hat{\beta}$ in a general oligopolistic competition model.[39]

Corollary 2.1 ***[estimated regression coefficients in an oligopolistic competition model].*** *Sector-specific supply shocks for trading partner f induce a cross-sector relationship in the domestic economy between import shares and consumer price indices, with the reduced-form specification $\Delta \log(P_s) = \alpha + \beta \Delta S_{sd} + \varepsilon_{sd}$, such the estimated regression coefficient $\hat{\beta}$ has the following structural interpretation:*

$$\hat{\beta} = \int_0^1 \omega_s \frac{1}{(\gamma - 1)\Omega_{sd} S_{sd}} ds,$$

and the reduced-form specification with the change in the domestic price index as a dependent variable, $\Delta \log(P_{sd}) = \alpha_{dom} + \beta_{dom} \Delta S_{sd} + \epsilon_{sd}$, yields coefficient:

$$\hat{\beta}_{dom} = \int_0^1 \omega_s \frac{1 - \Omega_{sd}}{(\gamma - 1)\Omega_{sd} S_{sd}} ds,$$

where the sector weights ω_s satisfy $\int_0^1 \omega_s ds = 1$ and $\omega_s \geq 0$, and are provided in Appendix D.A.2.

Moreover, the estimated regression coefficients satisfy the bounds:

$$0 \leq \hat{\beta} - \hat{\beta}_{dom} \leq \frac{1}{\eta \cdot S_d},$$

where η is a weighted average of sectoral elasticities defined in Appendix D.C.3.

Proof: see Appendix D.C.3.

The second part of Corollary 2.1 delivers a specific testable quantitative prediction, holding in a general model of oligopolistic competition with unrestricted firm heterogeneity. Indeed, the estimated regression coefficient for the domestic price index is predicted to be smaller than (in absolute value) but very close to the change in the overall price index, with an upper bound for the difference of $\frac{1}{\eta \cdot S_d}$.[40] With $\eta = 4.25$ (Simonovska and Waugh (2014)) and $S_d = 1 - 0.0452$ (Acemoglu et al. (2016), for 1999), we obtain that the response of the domestic price index must be within 0.25pp of the overall price response, which is in line with our results in Tables 2 and 6.[41] This prediction stands in contrast with models where the domestic price response stems from intermediate inputs, in which the predicted domestic price response is much smaller, as shown in Corollary 1.2.

Next, to assess whether conventional parameter values in oligopolistic price-setting models are

[39] Note that under monopolistic competition, $\Gamma_{sdj} = 0$ for all firms j and the predicted coefficient is identical to Proposition 1, since $\Omega_{sd} = 1$.

[40] Given our assumption, in line with Edmond et al. (2015)'s benchmark model, that there is no entry or exit, the theoretical price indices in Proposition 1 are equal to the measured price indices in our data. For a discussion of the potential divergence between theoretical and measured price index, see Sections IV.A and D.A.5.

[41] The domestic CPI price response is -1.94 (Col. 2 of Panel B of Table 6) and the overall CPI response is -2.23 (Col. 3 of Panel A of Table 2). These point estimates are statistically indistinguishable and the difference between them is 0.29pp, very close to the bound derived in Proposition 2.

consistent with the magnitude of our empirical estimates, we consider a stylized, tractable setting in which one U.S. producer and one Chinese producer compete in each sector. The magnitude of the price response, in terms of observables and elasticities, is given by Corollary 2.2.

Corollary 2.2 ***[price effects in an illustrative head-to-head Cournot competition model].*** *Assuming head-to-head Cournot competition between Domestic and Foreign within each sector, the estimated regression coefficient $\hat{\beta}$ has the following structural interpretation:*

$$\hat{\beta} = \int_0^1 \omega_s \frac{1 + \Gamma_{sd}/S_{sf}}{(\gamma - 1)S_{sd}} ds, \tag{8}$$

with Cournot markup elasticity given by,

$$\Gamma_{sd} = \frac{(\gamma - \epsilon)S_{sf}}{\gamma(\epsilon - 1) + (\gamma - \epsilon)S_{sf}} (\gamma - 1)S_{sd}.$$

Proof: see Appendix D.C.4.

As in Atkeson and Burstein (2008) and Edmond et al. (2015), Cournot competition implies that markups depend on firms' market shares. Domestic markups fall as the market share of the foreign producer increases. The magnitude of the fall in markups is governed by Γ_{sd}, which depends on expenditure shares and the elasticities γ and ϵ. We calibrate γ, ϵ and the relative productivity of Domestic and Foreign in order to match (i) our IV estimate of 1.91 using equation (8), (ii) the trade elasticity $\eta = 4.25$ (Simonovska and Waugh (2014)) using equation (D25) in Appendix D.C.2, and (iii) the foreign expenditure share, set to $S_{sf} = 0.0452$ (Acemoglu et al. (2016), for 1999). We obtain $\gamma = 8.72$ and $\epsilon = 1.43$, which is close to conventional parameters in the literature.[42] Moreover, these estimates imply a value for the markup elasticity $\Gamma_{sd} = 0.59$, which is very close to the untargeted benchmark empirical estimate $\widehat{\Gamma_{sd}} = 0.62$ of Amiti et al. (2019a).[43] These illustrative results show that, with standard parameter choices in oligopolistic price-setting models and with a level of markup elasticities consistent with auxiliary empirical evidence, the magnitude of the reduced-form relationship can be one order of magnitude larger than in the Melitz-Chaney model.[44]

Intuitively, Chinese producers reduce their prices when they experience a positive productivity shock, which leads U.S. producers to also reduce their prices due to strategic interactions. Because of the U.S. price response, the equilibrium change in the spending share on the product from China is lower than it would be absent this price response. As a result, the relationship between changes

[42] For example, in a rich structural model with firm heterogeneity, Edmond et al. (2015) obtain $\gamma = 10.5$ and $\epsilon = 1.24$. Estimating a full structural model with oligopolistic competition and firm-level heterogeneity is beyond the scope of our paper.

[43] We obtain the benchmark estimate from Table 1, Column 5 of Amiti et al. (2019a): their coefficient for the passthrough of changes in marginal cost is $\hat{\alpha} = 0.616$, implying $\hat{\Gamma} \equiv 1/\hat{\alpha} - 1 = 0.6234$.

[44] The results with Betrand competition instead of Cournot are similar (Appendix D.C.4). Although a full-fledged quantitative analysis is beyond the scope of this paper, the calibration results from the stylized Cournot and Bertrand competition models suggest that oligopolistic competition model can deliver an IV estimate of the correct order of magnitude.

in import penetration from China and price changes can be very large.[45]

Having established that the markup channel can plausibly explain large price effects, we now test several implications of theory to assess the empirical relevance of the markup channel.[46] These tests corroborate that the trade shock disrupted market power and consequently reduced prices through markup reductions.

The response of estimated markups. First, we examine whether estimated markups for domestic producers fall in response to increased trade with China. We follow the methodology of De Loecker et al. (2020) to estimate markups for publicly-listed firms in Compustat.[47] In this sample, indexing firms by i and years by t, the gross markup can be written $\mu_{it} = \theta^v \cdot \frac{SALES_{it}}{COGS_{it}}$, where θ^v is the elasticity of output to variable inputs, which multiplies the ratio of sales to the cost of goods sold. Intuitively, the gross markup corresponds to the ratio of the consumer price to the producer's shadow value of an additional unit of output.

Although the production approach to markup estimation has well-known limitations (e.g., Raval (2019)), it provides an instructive test for our purposes. Estimated markups from the Compustat sample can be used to test two predictions from the theoretical framework introduced above. First, do we observe a fall in estimated markups as trade with China increases in a product category? Second, do we see a larger response at the top of the markup distribution? In recent work, De Ridder et al. (2022) validate the approach of De Loecker et al. (2020) for studying *changes* in markups, which is precisely our focus.[48]

Panel A of Table 8 present the results of the analysis with estimated markups as the outcome. We repeat our IV specification (3) over two periods (1991-1999 and 2000-2007), but the level of aggregation is now a 6-digit NAICS code (instead of an ELI) and the outcome is the annualized change in the net markup (instead of the annualized inflation rate). Expressed in percentage points, the net markup is defined as $\tilde{\mu}_{it} = (\mu_{it} - 1) \cdot 100$. Column (3) reports the IV coefficient: when the import penetration rate from China increases by one percentage point, domestic markups fall by 1.75 percentage points (s.e. 0.848). This estimate is statistically indistinguishable from the IV coefficients for the response of domestic prices (from Table 6).

Panel (a) of Figure 5 reports this relationship for the average markup. Furthermore, Panels (b), (c) and (d) of Figure 5 document changes in the distribution of markups across industries with

[45] In Appendix D.C.4, we show that the regression coefficient can be unboundedly large, i.e. $\hat{\beta} \to \infty$, in certain limit cases.

[46] In Appendix D.D, we show that models featuring endogenous markups through Variable Elasticity of Substitution (VES) preferences, as Arkolakis et al. (2019), predict a domestic price response that is one order of magnitude too small relative to our IV estimates.

[47] We compute gross markups over time for each firm in the Compustat sample and we aggregate the firm-level data to 6-digit NAICS codes, using sales weights.

[48] Using firm-level administrative data, De Ridder et al. (2022) show that markup estimates from revenue data are biased for an analysis in *levels*, but not in *changes* (see e.g. their Figure 5, panel (b)).

heterogeneous exposure to increased trade with China. Panel (b) shows that there is no change at the bottom of the markup distribution: the reduced-form is flat for the 10th percentile of markups. In contrast, there is a negative relationship for the 50th percentile (Panel (c)), and the relationship becomes steeper for the 90th percentile (Panel (d)). Consistent with the predictions of oligopolistic competition models, the response of markups is much stronger at the top of the markup distribution. Finally, Panels (e) and (f) of Figure 5 show that profitability ratios deteriorate in industries that are more exposed to trade with China.[49] Appendix C.C discusses additional empirical results.

Heterogeneity by market structure. Second, to collect evidence beyond the sample of publicly-listed Compustat firms, we now assess whether heterogeneity in the estimated price effects *across* product categories is consistent with the predictions of the markup channel. We can (indirectly) test for the relevance of the markup channel by studying heterogeneity in the IV estimates across subsamples.

Per Proposition 2 and Corollary 2, the predicted price effect is increasing in the domestic markup elasticity. The markup elasticity depends on market structure and can therefore vary across industries. In a model with firm heterogeneity, it can be shown that is larger when the domestic market is more concentrated. This prediction is intuitive: when the domestic market is more concentrated, an increase in import competition from China disrupts domestic market power relatively more, therefore we expect to estimate larger price effects. Second, the expression for the predicted effect in Corollary 2 shows that the magnitude of the effect is decreasing in China's initial market share (starting from an equilibrium with a small spending share on Chinese goods, as in the data). Intuitively, there is less room for China to disrupt market power (at the margin) in an industry where it already has a high market share.

To measure domestic market concentration, we work with the PPI sample. The PPI sample frame provides weights for "value of shipments" for each establishment, therefore we are able to construct a Herfindahl index directly from the sample, instead of linking external data on market concentration. We create two indicator variables, one for product categories with a Herfindahl index above the median and one for product categories with an initial import penetration rate from China above the 75th percentile. We then implement our IV specification in subsamples and with interaction terms, interacting the indicator variables with the endogenous variable and the instrument.

Panel B of Table 8 presents the results. The specifications with interactions in Column (1) and the sub-sample specifications in Columns (2), (3) and (4) show that, in response to increased trade

[49] These results are consistent with the findings of Autor et al. (2020a), who document a negative cross-industry relationship between rising import penetration from China and firms' book values and stock market values (their Table 1). The fact that profitability deteriorates in sectors more exposed to rising trade with China is an additional piece of evidence suggesting that falling production costs do not drive the domestic price response.

with China, PPI inflation falls more in product categories that are more concentrated and falls less in categories that were initially more exposed to trade with China. The interaction terms in Column (1) are precisely estimated and significant at the 1% level. In the sub-sample of categories with domestic concentration below median (Column (3)), the point estimate is close to the prediction of Arkolakis et al. (2012). Appendix C.C shows that the results are similar with CPI data.

Overall, the observed heterogeneity in price effects across product categories suggest that markup responses are an important explanatory mechanism.

V The Distributional Effects of the China Shock

In this section, we discuss how our estimates help shed light on the distributional effects of the China shock.

V.A Displaced Jobs vs. Consumer Surplus

While we have documented that prices decline in U.S. industries with rising import penetration, numerous studies have shown that increasing import penetration rates from China have disrupted the U.S. labor market (e.g, Autor et al. (2013), Autor et al. (2014), Acemoglu et al. (2016), and Pierce and Schott (2016)). Employment declines in U.S. industries more exposed to rising import competition, which is detrimental to displaced U.S. workers who are not able to transition without cost to another industry.

Using our IV estimates, we can characterize the tradeoff between rising consumer surplus and displaced jobs across industries.[50] If the import penetration rate from China increases by one percentage point more in industry A than in industry B, what is the impact on (relative) consumer surplus and jobs in these two industries? We can answer this question using IV estimates for the price effects (denoted β_{price}) and the employment effects (denoted β_{emp}), provided that they are scaled properly.

A first-order approximation to the change in consumer surplus (in dollars) from the trade shock for industry j is given by $\Delta CS_j = \left(\frac{-\beta_{price}}{100}\Delta ChinaIP_j\right) \cdot Cons_j$, where "$Cons_j$" is total consumption (or "domestic absorption") for industry j and "$\frac{-\beta_{price}}{100}\Delta ChinaIP_j$" is the fall in

[50] As indicated earlier, the calculations carried out in this section are partial-equilibrium differences across industries with different levels of exposure to rising import penetration from China, i.e. we focus on differences in consumer surplus across industries rather than making statement about overall consumer surplus, which would necessarily entail an extrapolation since general equilibrium effects induced by the China shock could affect all industries simultaneously. As is well-known (e.g., Adao et al. (2019)), GE effects affecting all industries are not reflected in our cross-industry IV estimates as they are absorbed by fixed effects. The tradeoff between increasing consumer surplus and displaced jobs could therefore be different at the aggregate level, once GE effects are accounted for. For example, if displaced manufacturing jobs lead to more job creation in other industries (e.g., Bloom et al. (2019), Galle et al. (2022)), then the increase in consumer surplus per "destroyed" job would be larger than the increase in consumer surplus per "displaced" job documented in Table 9. In this sense, not adjusting for job reallocation is conservative for our purposes and provides a lower bound for policy analysis.

prices induced by industry j's trade shock. Similarly, the number of displaced jobs is $\Delta Jobs_j = \left(\frac{-\beta_{emp}}{100}\Delta ChinaIP_j\right) \cdot Emp_j$, where Emp_j is total employment in industry j.

Assuming that industries A and B initially have the same levels of total consumption and employment, if import penetration increases by one percentage point more in A than in B, the tradeoff between rising consumer surplus and displaced jobs in A relative to B is given by

$$\frac{\Delta CS_j}{\Delta Jobs_j} = \frac{\beta_{price}}{\beta_{emp}} \cdot \frac{Cons_j}{Emp_j}. \tag{9}$$

If industry j employs few workers but accounts for large share of aggregate consumption, then a large amount of consumer surplus can be created per displaced job, as long as β_{price} and β_{emp} are similar.

Panel A of Table 9 reports informative summary statistics about the ratio $\frac{Cons_j}{Emp_j}$, focusing on the set of ELIs within goods only (for which we obtain data on domestic absorption and employment from the NBER-CES Manufacturing database). The summary statistics are reported for 2000, at the outset of the China shock. The first row shows that average annual labor earnings in the sample are about \$33,000 on average. Because the labor share is low, the average value-added of domestic producers per job is higher, around \$120,000 (row 2). And because these industries use a lot of intermediate inputs, total domestic sales per job is much higher, about \$305,000 on average (row 3). Finally, since trade is important for consumption in these industries, total domestic absorption per job is even higher, approximately \$390,000 per job (row 4).

For comparability with the estimated price effects, we run IV specifications for employment in our sample. We repeat specification (3) with the log change in employment as the outcome. We estimate that a one percentage point increase in import penetration from China leads to a fall in employment of 1.83% with the NTR gap instrument, 1.77% with the change in import penetration in other developed economies, and 1.82% with both instruments (Appendix Table A10). The estimates are similar whether we consider all employment, or production workers and non-production workers separately; the magnitudes are in line with prior work (e.g., Table 2 of Acemoglu et al. (2016)).

Panel B of Table 9 characterizes the tradeoffs between rising consumer surplus and displaced jobs across industries, using our IV estimates and equation (9). Because the ratio $\frac{Cons_j}{Emp_j}$ varies across industries, the tradeoff depends on which industry is affected by rising import competition.

We first consider a counterfactual increase in the import penetration rate from China of one percentage point for a representative industry with the average ratio of total consumption to employment (in our sample of goods). We compute $\frac{\Delta CS}{\Delta Jobs} = \frac{\beta_{price}}{\beta_{emp}} \cdot \frac{\sum_j Cons_j}{\sum_j Emp_j}$. In Column (1), with the NTR gap instrument, consumer surplus increases by \$477,555 for each job displaced by trade with China. The estimate remains large, at \$317,383, when using trade with China in other developed economies as the instrument (Column (2)). With both instruments, the estimate yields \$411,464 in consumer surplus per displaced job (Column (3)).

Next, we repeat these calculations by focusing on the industries that were affected by the rise in import penetration from China between 2000 and 2007. If the ratio $\frac{Cons_j}{Emp_j}$ is systematically higher or lower for affected industries, the tradeoff between consumer surplus and employment for the historical China shock could differ from what the previous analysis suggests. We compute $\frac{\Delta CS}{\Delta Jobs} = \frac{\beta_{price}}{\beta_{emp}} \cdot \frac{\sum_j \Delta ChinaIP_j \cdot Cons_j}{\sum_j \Delta ChinaIP_j \cdot Emp_j}$, i.e. the consumption-to-employment ratio is computed with rising import penetration from China as weights. The results are reported in Columns (4) to (6) of Panel B of Table 9. They are slightly attenuated compared to the baseline but remain very large in magnitude, ranging from \$288,147 to \$433,565 across specifications.[51]

Thus, our estimates imply that industries that are more exposed to trade with China create hundreds of thousands of dollars in consumer surplus for each displaced job. Using the predicted price effects from the class of standard trade models nested by Arkolakis et al. (2012), the increase in consumer surplus would be attenuated by a factor of ten and would be on the order of \$40,000 per displaced job, which is similar to average annual labor earnings in this sample.[52]

V.B Distributional Effects via the Expenditure Channel

Finally, we examine whether the price response differs across product categories that cater to households of different income levels.

A growing literature characterizes the distributional effects of trade through the expenditure channel, focusing on differences in spending shares on imports across consumer groups (e.g., Fajgelbaum and Khandelwal (2016), Borusyak and Jaravel (2021), He (2018), and Hottman and Monarch (2020)). We investigate a distinct mechanism: does the rate of pass-through of trade shocks into consumer prices vary systematically with consumer income?

We proceed in two steps. First, we repeat our IV strategy in subsamples of product categories catering to different income groups; second, we use these new estimates to quantify whether this

[51] The implications of declining markups for U.S. producer surplus are ambiguous: if markups reflect market power and economic profit, then producer surplus may have fallen; but if markups merely offset fixed costs such that a zero-profit condition holds, then there is no change in producer surplus.

[52] While we focused on consumer surplus and the tradeoff with displaced jobs, which allows us to make statements that are both informative for policy and robust to general equilibrium effects and the missing intercept (see footnote (50)), other studied have extrapolated from cross-industry results to aggregate effects (e.g., Amiti et al. (2020) and Bai and Stumpner (2019)). Following this line of work, in Online Appendix Table A11 we compute the equivalent variation for increased trade with China from 2000 to 2007, for the average U.S. household, assuming that there are no GE effects affecting prices in all industries. Under this simplifying assumption, the cross-industry IV estimates accurately capture the price effects. Using our baseline IV estimate, we find that in 2007 the (annual) purchasing power of the representative U.S. household was about \$1,500 higher, thanks to lower prices induced by increased trade with China from 2000 to 2007. The estimates range from \$1,105 to \$1,711 across specifications. Assuming that prices do not revert back in the longer run, this result indicates that the China shock increased the purchasing power of U.S. households by about 2% (in the CEX, average annual expenditures were \$49,638 in 2007). In GE, increasing import penetration may induce an overall fall in domestic prices to restore trade balance; therefore the increase in purchasing power for domestic consumers could be larger after accounting for GE effects.

mechanism has a substantial impact on distributional effects across income groups.[53]

We start by running our IV specification (3) in subsamples of product categories whose expenditure shares vary across groups of consumers. For robustness, we split the sample around the median using three alternative variables reflecting consumer income: the share of sales to college graduates, the expenditure elasticity, and the shares of sales to households with an annual income above \$60,000.[54]

The results are reported in Panel A of Table 10. The price effects are large and significant in all subsamples, but they are much larger in product categories that sell to lower-income households. Columns (1) and (2) show that the point estimate for product categories with a share of sales to college graduates above median is only 21% (= 0.91/4.28) of the point estimate for the categories below median. The difference is similar when splitting by expenditure elasticity (0.83/4.62 = 18.3%, in Columns (3) and (4)), while it is attenuated when splitting by the share of sales to households with income above \$60,000 (1.18/2.93 = 40.2%, in Columns (5) and (6)).[55]

Next, we examine whether the estimated heterogeneity in price effects implies substantial distributional effects across income groups. We compute a first-order approximation to the equivalent variation from a change in prices for each consumer group i, expressed as a percentage of initial expenditures for each group, $EV_i = \sum_j s_j^i \hat{p}_j$, where s_j^i is the expenditure share by consumer group i on product category j, and $\hat{p}_j$ is the percentage change in product category j's price index that is induced by the trade shock. We compute this price change as $\hat{p}_j = \frac{\beta_j}{100} \Delta ChinaIP_j$, where β_j is our IV estimate for j (which can vary across product categories as in Panel A of Table 10) and $\Delta ChinaIP_j$ is the increase in import penetration rate from China in j between 2000 and 2007.

We compute the difference in the equivalent variation for high-income and low-income groups, standardized by the average equivalent variation across groups, given by $\Delta EV \equiv \frac{EV^{HI} - EV^{LI}}{EV^{All}} = \frac{\sum_j \left(s_j^{HI} - s_j^{LI}\right) \beta_j \Delta ChinaIP_j}{\sum_j s_j^{All} \beta_j \Delta ChinaIP_j}$.

Intuitively, income group i benefits more if it spends more on categories that are more exposed to rising trade with China ($\Delta ChinaIP_j$) and that feature a larger price response to the shock (β_j).[56] As shown by Fajgelbaum and Khandelwal (2016) and Borusyak and Jaravel (2021), the

[53] In a sample of consumer packaged goods, Bai and Stumpner (2019) examine whether price responses to trade shock differ across income groups *within* the same detailed product category (e.g., between different varieties of beer) and find no difference. In contrast, we document substantial heterogeneity in price responses *across* product categories that tend to target different income groups (e.g., between beer and wine).

[54] We use the spending shares from the CEX for the year 2000, as processed by Borusyak and Jaravel (2021). We match the CEX consumption categories (UCCs) to ELI as explained in Online Appendix B.C.

[55] Because we split the sample, the first-stage F statistics fall. The table reports the results with LIML, which yields very similar point estimates and alleviates concerns about weak instruments. To maximize power, we use both instruments jointly. The Hansen J statistics indicate that we cannot reject the over-identifying restrictions. We obtain similar results with interaction terms in a single specification, instead of repeating the analysis in subsamples.

[56] As shown in these formulas, when we compare the effects across consumer groups we difference out any GE effect affecting all product categories. For this reason, our cross-industry IV estimates are well-suited for the estimation of distributional effects: although they cannot recover aggregate GE effects without additional assumptions, they characterize cross-industry effects accurately.

distributional effects arising from the expenditure channel can be expressed as ΔEV in a general equilibrium model, i.e. our estimates are directly informative about heterogeneity in the expenditure channel despite the fact that they do not recover the "missing intercept".[57]

Panel B of Table 10 reports the results. Column (1) imposes a homogeneous price response to trade shocks, using our baseline estimate for β for all categories (-1.91%, from Table 2). We find that higher-income groups benefit proportionally more from increased trade with China: 6.19% more for college-educated households relative to those without a college degree; 8.39% more for households earnings above $60,000 a year relative to those earning less; and 14.53% more for households earning above $100,000 relative to below $30,000.

These differences result from the fact that, between 2000 and 2007, import penetration from China increased faster in product categories that sell relatively more to higher-income groups (e.g., in consumer electronics rather than in food products). This finding is confirmed in Column (4) in a sample restricted to goods only (including services tends to attenuate the differences, because higher-income groups spend more on services and services are not exposed to trade with China). These patterns are consistent with prior work by Borusyak and Jaravel (2021).

In Column (2), we allow the price response to vary across product categories, depending on the share of sales to households earning above $60,000 a year, as in Panel A of Table 10. In this case, the patterns are *reversed* and higher-income groups now benefit proportionally *less*: 9.64% less for college relative to non-college; 19.54% less for those earnings above $60,000 relative to those below; and 23.13% less for those above $100,000 relative to below $30,000. In Column (3), these differences are magnified, ranging from 13.94% to 36.29%, when we specify heterogeneous price effects using the estimates based on expenditure elasticities (because with these estimates, the price effects are even larger for low-income groups, as shown in Panel A of Table 10).

Columns (5) and (6) confirm these findings in the sample of goods: with heterogeneous pass-through of the trade shocks from China, higher-income groups benefit relatively less, while they benefit relatively more with homogeneous pass-through. The patterns are similar when using heterogeneous pass-through rates by the share of sales to college-educated households (not reported).

Taken together, these findings indicate that accounting for heterogeneous price responses across product categories can be important to accurately characterize the distributional effects of trade via the expenditure channel.

[57]See, e.g., equation (6) in Fajgelbaum and Khandelwal (2016). Intuitively, the "missing intercept" is differenced out when computing the expenditure channel, which is defined as the difference in the equivalent variation across groups, ΔEV. The missing intercept does matter for the overall welfare gains from trade, which we do not claim to speak to; instead, we focus on the specific distributional effects that arise from the expenditure channel.

VI Conclusion

This paper has presented new evidence on the price effects of trade by leveraging a comprehensive price data set from the Bureau of Labor Statistics. Most previous work on the "China shock" emphasized its detrimental consequences for U.S. employment. Our findings convey a different message: the price effects of trade with China were large and beneficial to U.S. consumers. We estimate that falling prices in product categories that were more exposed to trade with China created hundreds of thousands of dollars in consumer surplus for each displaced job. These prices effects are particularly large in product categories selling to low-income consumers.

Our estimates of the impact of rising import penetration on consumer prices are much larger than predicted by standard quantitative trade models such as Melitz (2003)-Chaney (2008) and other members of the Arkolakis et al. (2012) class. We showed that there is a large fall in *domestic* prices, driven by intensified competition and declining markups. By disrupting domestic market power, trade can have substantial price effects that benefit consumers, potentially at the expense of domestic producer surplus. These findings highlight the importance of including endogenous markups and strategic pricing into quantitative trade models used for policy analysis. In a period of rising concentration and rising markups in the United States (Autor et al. (2020c), De Loecker et al. (2020)), the pro-competitive effects of trade may be particularly valuable to U.S. consumers.

While the costs of free trade are disproportionately borne by particular workers, industries, and regions, the large magnitude of the price effects suggest that it may be possible to compensate those who suffer from the labor market impacts of trade shocks. Developing and testing such redistribution schemes is a particularly promising direction for research and policy going forward.

References

Acemoglu, Daron, David Autor, David Dorn, Gordon H Hanson, and Brendan Price, "Import competition and the great US employment sag of the 2000s," *Journal of Labor Economics*, 2016, *34* (S1), S141–S198.

Adao, Rodrigo, Costas Arkolakis, and Federico Esposito, "Spatial linkages, global shocks, and local labor markets: Theory and evidence," *National Bureau of Economic Research*, 2019.

Aghion, Philippe, Antonin Bergeaud, Matthieu Lequien, and Marc J Melitz, "The impact of exports on innovation: Theory and evidence," *Working Paper*, 2018.

Amiti, Mary, Mi Dai, Robert C Feenstra, and John Romalis, "How did China's WTO entry affect US prices?," *Journal of International Economics*, 2020, *126*, 103339.

_ , Oleg Itskhoki, and Jozef Konings, "International Shocks, Variable Markups and Domestic Prices," *The Review of Economic Studies*, 2019, *86* (6), 2356–2402.

_ , Stephen J Redding, and David E Weinstein, "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, November 2019, *33* (4), 187–210.

Angrist, Joshua D and Alan B Krueger, "Empirical strategies in labor economics," in "Handbook of labor economics," Vol. 3, Elsevier, 1999, pp. 1277–1366.

Antràs, Pol and Davin Chor, "Organizing the global value chain," *Econometrica*, 2013, *81* (6), 2127–2204.

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

_ , _ , Dave Donaldson, Andrés Rodríguez-Clare et al., "The elusive pro-competitive effects of trade," Technical Report 1 2019.

_ , Jonathan Eaton, and Samuel Kortum, "Staggered adjustment and trade dynamics," *Working Paper*, 2011.

Armington, Paul S, "A theory of demand for products distinguished by place of production," *Staff Papers*, 1969, *16* (1), 159–178.

Atkeson, Andrew and Ariel Burstein, "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 2008, *98* (5), 1998–2031.

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.

_ , Ariel T Burstein, and Sarah M Lein, "Exchange rates and prices: evidence from the 2015 Swiss franc appreciation," *American Economic Review*, 2021, *111* (2), 652–86.

Autor, David, David Dorn, and Gordon Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

_ , _ , Gordon H Hanson, Pian Shu, and Gary Pisano, "Foreign competition and domestic innovation: Evidence from US patents," *American Economic Review: Insights*, 2020, *2* (3), 357–74.

_ , _ , Gordon Hanson, and Jae Song, "Trade adjustment: Worker-level evidence," *The Quarterly Journal of Economics*, 2014, *129* (4), 1799–1860.

_ , _ , _ , and Kaveh Majlesi, "Importing political polarization? The electoral consequences of rising trade exposure," *American Economic Review*, 2020, *110* (10), 3139–83.

_ , _ , _ et al., "When work disappears: Manufacturing decline and the falling marriage-market value of young men," *American Economic Review: Insights*, September 2019, *1* (2), 161–78.

_ , _ , Lawrence F Katz, Christina Patterson, and John Van Reenen, "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Backus, David, Patrick J Kehoe, and Finn E Kydland, "International business cycles: theory and evidence," *NBER Working Paper*, 1993.

Bai, Liang and Sebastian Stumpner, "Estimating US consumer gains from Chinese imports," *American Economic Review: Insights*, 2019, *1* (2), 209–24.

Bai, Xue, Kala Krishna, and Hong Ma, "How you export matters: Export mode, learning and productivity in China," *Journal of International Economics*, 2017, *104*, 122–137.

Bernard, Andrew B, Jonathan Eaton, J Bradford Jensen, and Samuel Kortum, "Plants and productivity in international trade," *American economic review*, 2003, *93* (4), 1268–1290.

Bils, Mark, "Do higher prices for new goods reflect quality growth or inflation?," *The Quarterly Journal of Economics*, 2009, *124* (2), 637–675.

_ and Peter J Klenow, "Some evidence on the importance of sticky prices," *Journal of Political Economy*, 2004, *112* (5), 947–985.

— , — , **and Benjamin A Malin**, "Reset price inflation and the impact of monetary policy shocks," *American Economic Review*, 2012, *102* (6), 2798–2825.

Bloom, Nicholas, Kyle Handley, Andre Kurmann, and Philip Luck, "The Impact of Chinese Trade on US Employment: The Good, The Bad, and The Apocryphal," *Working Paper*, 2019.

— **, 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.

Boehm, Christoph E, Andrei A Levchenko, and Nitya Pandalai-Nayar, "The long and short (run) of trade elasticities," *National Bureau of Economic Research Working Paper*, 2022.

Borusyak, Kirill and Xavier Jaravel, "The Distributional Effects of Trade: Theory and Evidence from the United States," *Available at SSRN 3269579*, 2021.

Brander, James and Paul Krugman, "A reciprocal dumping model of international trade," *Journal of international economics*, 1983, *15* (3-4), 313–321.

Brandt, Loren and Peter M Morrow, "Tariffs and the Organization of Trade in China," *Journal of International Economics*, 2017, *104*, 85–103.

Broda, Christian and David E Weinstein, "Globalization and the Gains from Variety," *The Quarterly journal of economics*, 2006, *121* (2), 541–585.

Burstein, Ariel and Gita Gopinath, "International prices and exchange rates," *Handbook of international economics*, 2014, *4*, 391–451.

Bustos, Paula, "Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms," *American economic review*, 2011, *101* (1), 304–40.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 2015, *82* (1), 1–44.

Carluccio, Juan, Erwan Gautier, and Sophie Guilloux-Nefussi, "Dissecting the Impact of Imports from Low-Wage Countries on French Consumer Prices," *Working Paper*, 2018.

Carroll, Daniel and Sewon Hur, "On the Heterogeneous Welfare Gains and Losses from Trade," *Journal of Monetary Economics*, 2020, *109*, 1–16.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang, "Tariff Passthrough at the Border and at the Store: Evidence from US Trade Policy," *American Economic Review: Insights*, March 2021, *3* (1), 19–34.

Chaney, Thomas, "Distorted gravity: the intensive and extensive margins of international trade," *American Economic Review*, 2008, *98* (4), 1707–21.

Costinot, Arnaud, Dave Donaldson, Margaret Kyle, and Heidi Williams, "The more we die, the more we sell? a simple test of the home-market effect," *The quarterly journal of economics*, 2019, *134* (2), 843–894.

Deaton, Angus, "Rice prices and income distribution in Thailand: a non-parametric analysis," *The Economic Journal*, 1989, *99* (395), 1–37.

Dixit, Avinash K and Robert S Pindyck, *Investment under uncertainty*, Princeton university press, 1994.

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

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu, "Competition, markups, and the gains from international trade," *American Economic Review*, 2015, *105* (10), 3183–3221.

Epifani, Paolo and Gino Gancia, "Trade, markup heterogeneity and misallocations," *Journal of International Economics*, 2011, *83* (1), 1–13.

Faber, Benjamin and Thibault Fally, "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data," *Working Paper*, 2017.

Fajgelbaum, Pablo D and Amit K Khandelwal, "Measuring the unequal gains from trade," *The Quarterly Journal of Economics*, 2016, *131* (3), 1113–1180.

_ , Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, "The return to protectionism," *The Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Feenstra, Robert C, "New product varieties and the measurement of international prices," *The American Economic Review*, 1994, pp. 157–177.

_ , "Restoring the product variety and pro-competitive gains from trade with heterogeneous firms and bounded productivity," *Journal of International Economics*, 2018, *110*, 16–27.

_ and David E Weinstein, "Globalization, markups, and US welfare," *Journal of Political Economy*, 2017, *125* (4), 1040–1074.

Flaaen, Aaron B, Ali Hortacsu, and Felix Tintelnot, "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines," Technical Report 7 2020.

Gagnon, Etienne, David López-Salido, and Nicolas Vincent, "Individual price adjustment along the extensive margin," *NBER Macroeconomics Annual*, 2013, *27* (1), 235–281.

Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi, "Slicing the pie: Quantifying the aggregate and distributional effects of trade," *The Review of Economic Studies*, 2022.

Goldberg, Pinelopi K. and Rebecca Hellerstein, "How Rigid are Producer Prices?," *Working Paper*, 2009.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely, "Are US industries becoming more concentrated?," *Forthcoming, Review of Finance*, 2018.

Handley, Kyle and Nuno Limão, "Policy uncertainty, trade, and welfare: Theory and evidence for china and the united states," *American Economic Review*, 2017, *107* (9), 2731–83.

Hansen, Lars Peter, "Large sample properties of generalized method of moments estimators," *Econometrica: Journal of the Econometric Society*, 1982, pp. 1029–1054.

He, Zheli, "Trade and Real Wages of the Rich and Poor," *Working Paper*, 2018.

Head, Keith and Thierry Mayer, "Gravity equations: Workhorse, toolkit, and cookbook," in "Handbook of international economics," Vol. 4, Elsevier, 2014, pp. 131–195.

Hottman, Colin J and Ryan Monarch, "A matter of taste: Estimating import price inflation across US income groups," *Journal of International Economics*, 2020, *127*, 103382.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

_ , Nicholas Li, Ralph Ossa, and Mu-Jeung Yang, "Accounting for the New Gains from Trade Liberalization," *Journal of International Economics*, 2020, *127*, 103370.

Impullitti, Giammario and Omar Licandro, "Trade, firm selection and innovation: The competition channel," *The Economic Journal*, 2018, *128* (608), 189–229.

Jaravel, Xavier, "The unequal gains from product innovations: Evidence from the us retail sector," *Quarterly Journal of Economics*, 2019.

Kamin, Steven B, Mario Marazzi, and John W Schindler, "The impact of Chinese exports on global import prices," *Review of International Economics*, 2006, *14* (2), 179–201.

Konny, Crystal G., Brendan K. Williams, and David M. Friedman, "Big Data in the U.S. Consumer Price Index: Experiences and Plans," in Katharine G. Abraham, Ron S. Jarmin, Brian Moyer, and Matthew D. Shapiro, eds., *Big Data for 21st Century Economic Statistics*, Chicago, IL: University of Chicago Press, 2019.

Konüs, Alexander A, "The problem of the true index of the cost of living," *Econometrica*, 1939, pp. 10–29.

Konus, Alexander A and SS Byushgens, "On the problem of the purchasing power of money," *Economic Bulletin of the Conjuncture Institute, Supplement*, 1926, pp. 151–72.

Krishna, Pravin and Devashish Mitra, "Trade liberalization, market discipline and productivity growth: new evidence from India," *Journal of Development Economics*, 1998, *56* (2), 447–462.

Levinsohn, James, "Testing the imports-as-market-discipline hypothesis," *Journal of International Economics*, 1993, *35* (1-2), 1–22.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The rise of market power and the macroeconomic implications," *The Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

_ , 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), 1695–1725.

Moulton, Brent R. and Karin E. Moses, "Addressing the Quality Change Issue in the Consumer Price Index," *Brookings Papers on Economic Activity*, 1997, *28* (1), 305–366.

Nakamura, Emi and Dawit Zerom, "Accounting for incomplete pass-through," *The Review of Economic Studies*, 2010, *77* (3), 1192–1230.

_ and Jon Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

_ and Jón Steinsson, "Identification in macroeconomics," *Journal of Economic Perspectives*, 2018, *32* (3), 59–86.

_ , _ , Patrick Sun, and Daniel Villar, "The elusive costs of inflation: Price dispersion during the US great inflation," *The Quarterly Journal of Economics*, 2018, *133* (4), 1933–1980.

Ossa, Ralph, "Why trade matters after all," *Journal of International Economics*, 2015, *97* (2), 266–277.

Pierce, Justin R and Peter K Schott, "A concordance between ten-digit US Harmonized System Codes and SIC/NAICS product classes and industries," *Journal of Economic and Social Measurement*, 2012, *37* (1, 2), 61–96.

_ and _ , "The surprisingly swift decline of US manufacturing employment," *American Economic Review*, 2016, *106* (7), 1632–62.

_ and _ , "Investment responses to trade liberalization: Evidence from US industries and establishments," *Journal of International Economics*, 2018, *115*, 203–222.

_ and _ , "Trade liberalization and mortality: Evidence from US counties," *American Economic Review: Insights*, 2020, *2* (1), 47–64.

Raval, Devesh, "Testing the Production Approach to Markup Estimation," *Working Paper*, 2019.

Ridder, Maarten De, Basile Grassi, and Giovanni Morzenti, "The hitchiker guide to markup estimation," *Working paper*, 2022.

Simonovska, Ina and Michael E Waugh, "Trade models, trade elasticities, and the gains from trade," *National Bureau of Economic Research Working Paper*, 2014.

Stolper, Wolfgang F and Paul A Samuelson, "Protection and real wages," *The Review of Economic Studies*, 1941, *9* (1), 58–73.

U.S. Bureau of Labor Statistics, *BLS Handbook of Methods*, Washington, DC: Bulletin of the United States Bureau of Labor Statistics, U.S. Dept. of Labor, 2018.

Yitzhaki, Shlomo, "On using linear regressions in welfare economics," *Journal of Business & Economic Statistics*, 1996, *14* (4), 478–486.

Figure 1: Testing for Pre-trends

Panel A: Event Study for NTR Gap in Main Analysis Sample

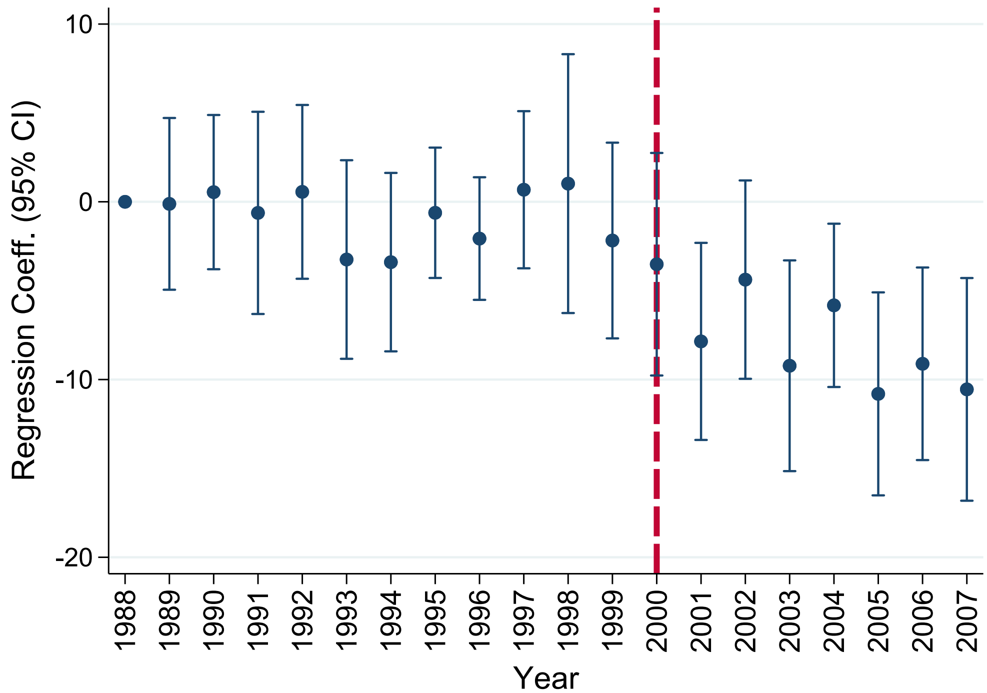

Panel B: Placebo Reduced-Forms in Extended CPI Sample

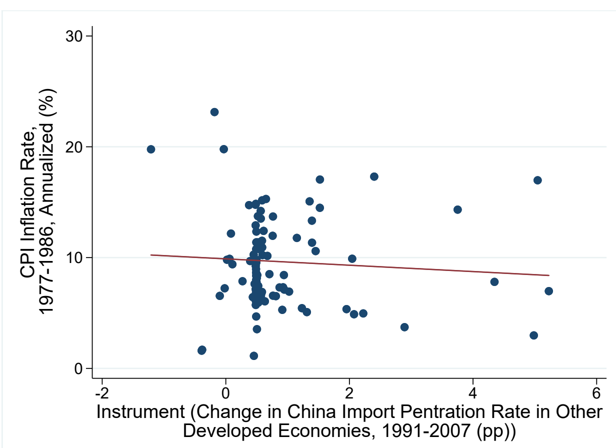

(a) China IP in other developed economies

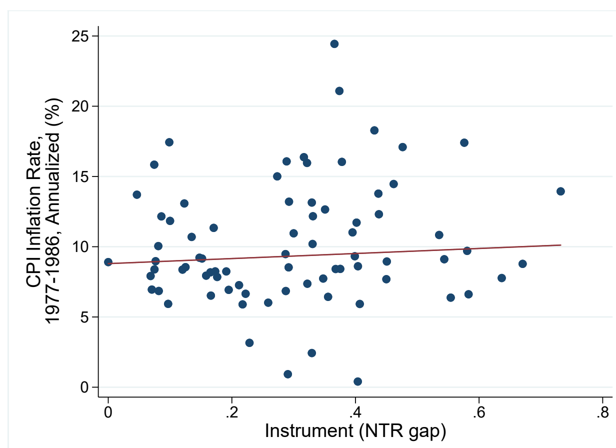

(b) NTR gap

Notes: Panel A reports the estimates from the specification described in equation (2). Panel B reports the binned scatter plots for the reduced-form specifications in the extended CPI sample. Each dot represents 1% of the data and the OLS best-fit line is reported in red. The level of observation is an ELI.

Figure 2: IV Estimates with the NTR Gap

Panel A: First Stage

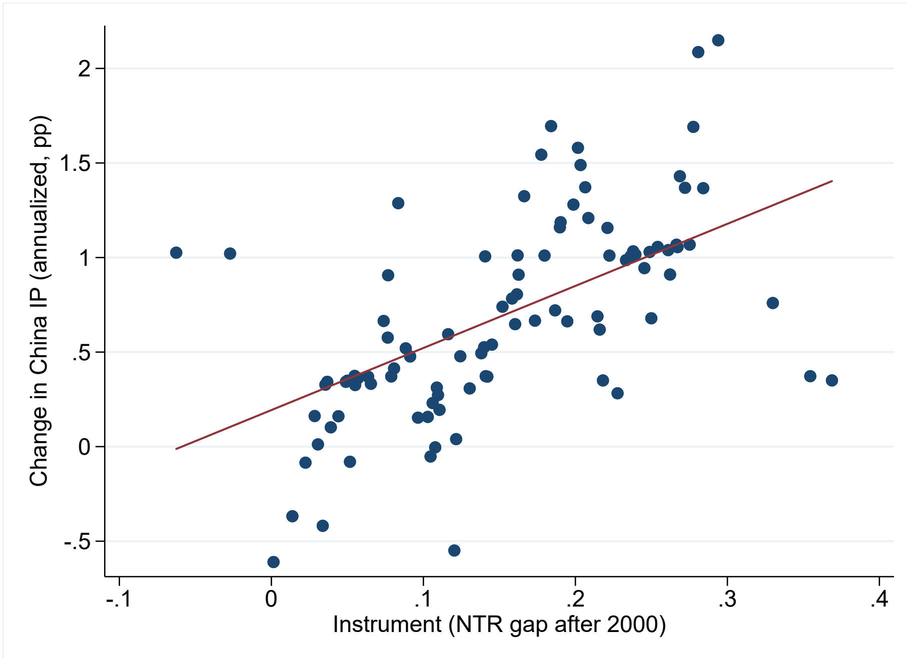

Panel B: Reduced-Form

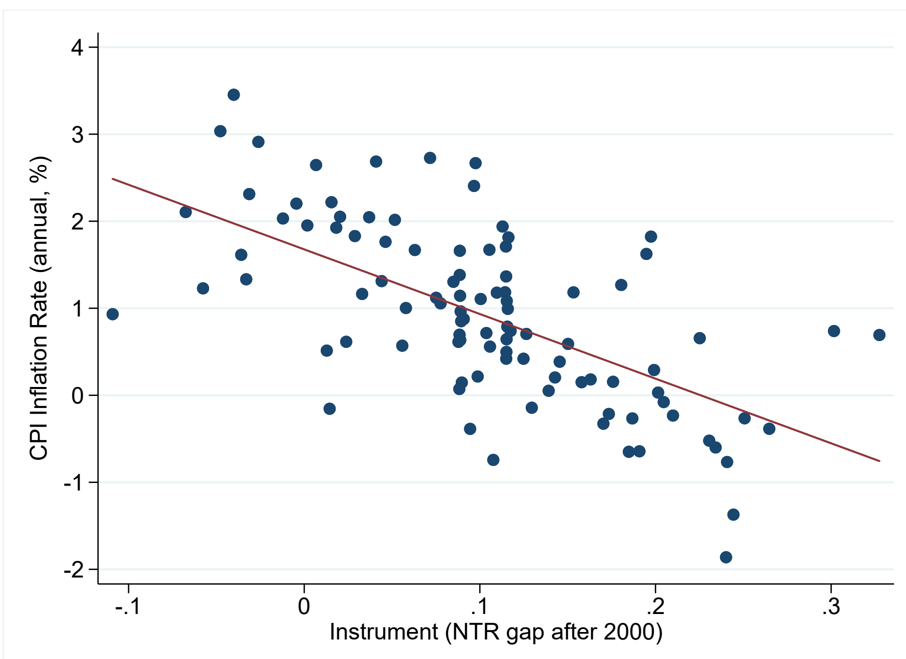

Notes: This figure reports the binned scatter plots for the first-stage (Panel A) and reduced-form (Panel B) relationships of the IV strategy using the NTR gap as an instrument. Each dot represents 1% of the data and the OLS best-fit line are reported in red. The level of observation is an ELI-by-period cell. Consumption weights are used for the OLS best-fit line.

Figure 3: Falsification Tests with the French CPI Data

Panel A: Placebo Reduced-Form

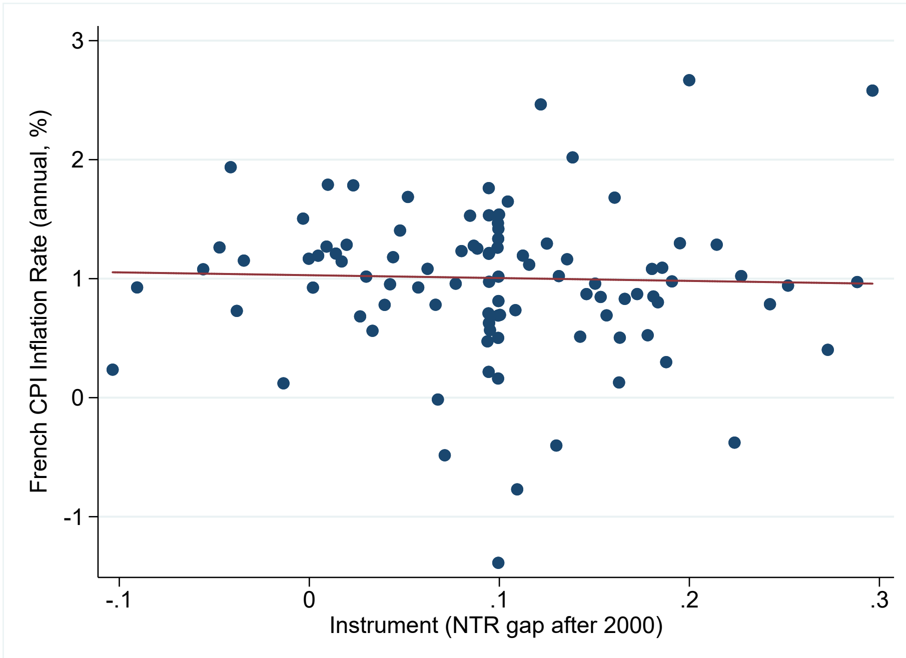

Panel B: Reduced-Form for Triple-Difference Specification

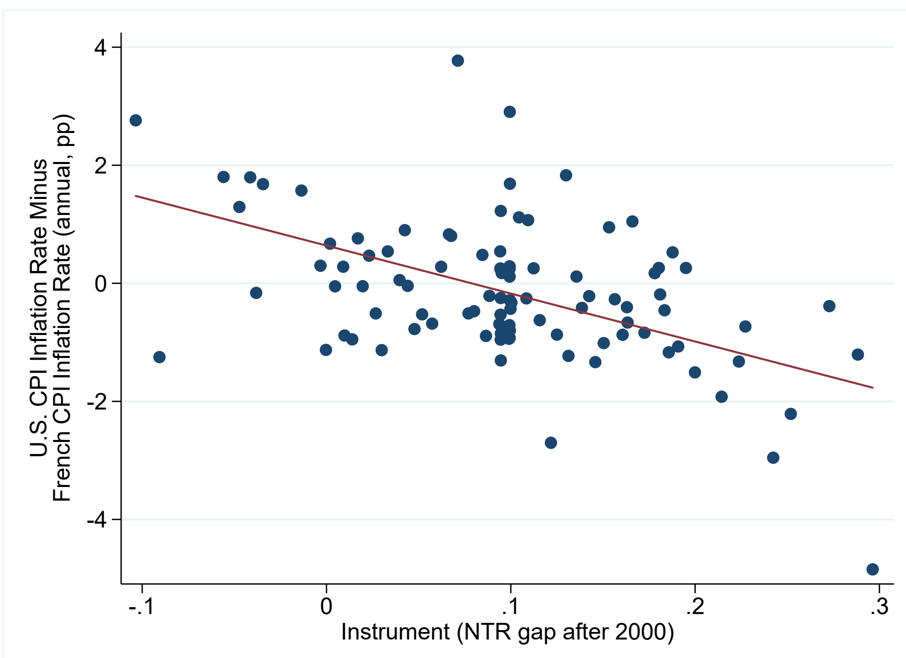

Notes: This figure reports the binned scatter plots for reduced-form relationships for the placebo test and triple difference, using the NTR gap as an instrument in the linked US-French CPI sample. Each dot represents 1% of the data, using consumption weights, and the OLS best-fit line is shown in red.. The level of observation is a COICOP-by-period cell.

Figure 4: The Role of Input-Output Linkages for Exposure to Trade with China

Panel A: Relationship between Direct Import Competition and Exposure via Domestic Suppliers

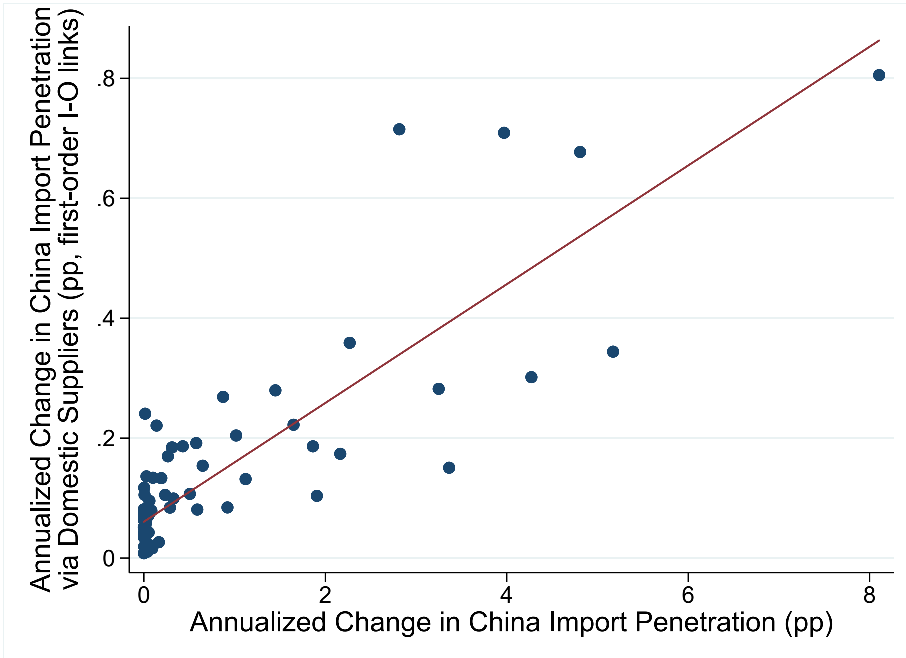

Panel B: Relationship between Direct and Indirect Import Competition

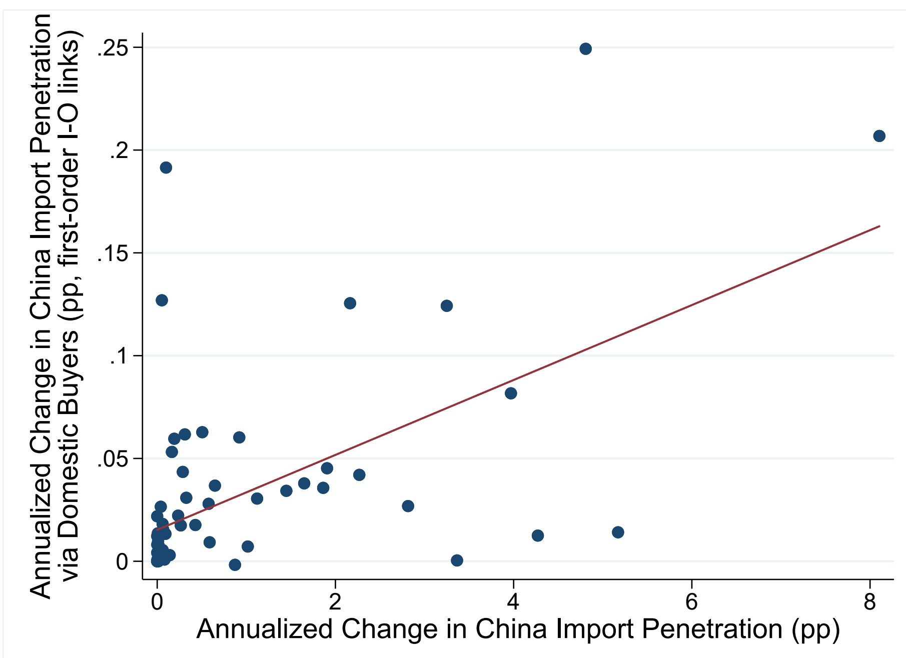

Notes: This figure shows the relationship between direct and indirect exposure to trade with China via domestic suppliers (Panel A) and buyers (Panel B). The level of observation is a 6-digit IO industry-by-period. Each dot represents 1% of the data, using consumption weights, and the OLS best-fit line is shown in red.

Figure 5: The Role of Markups

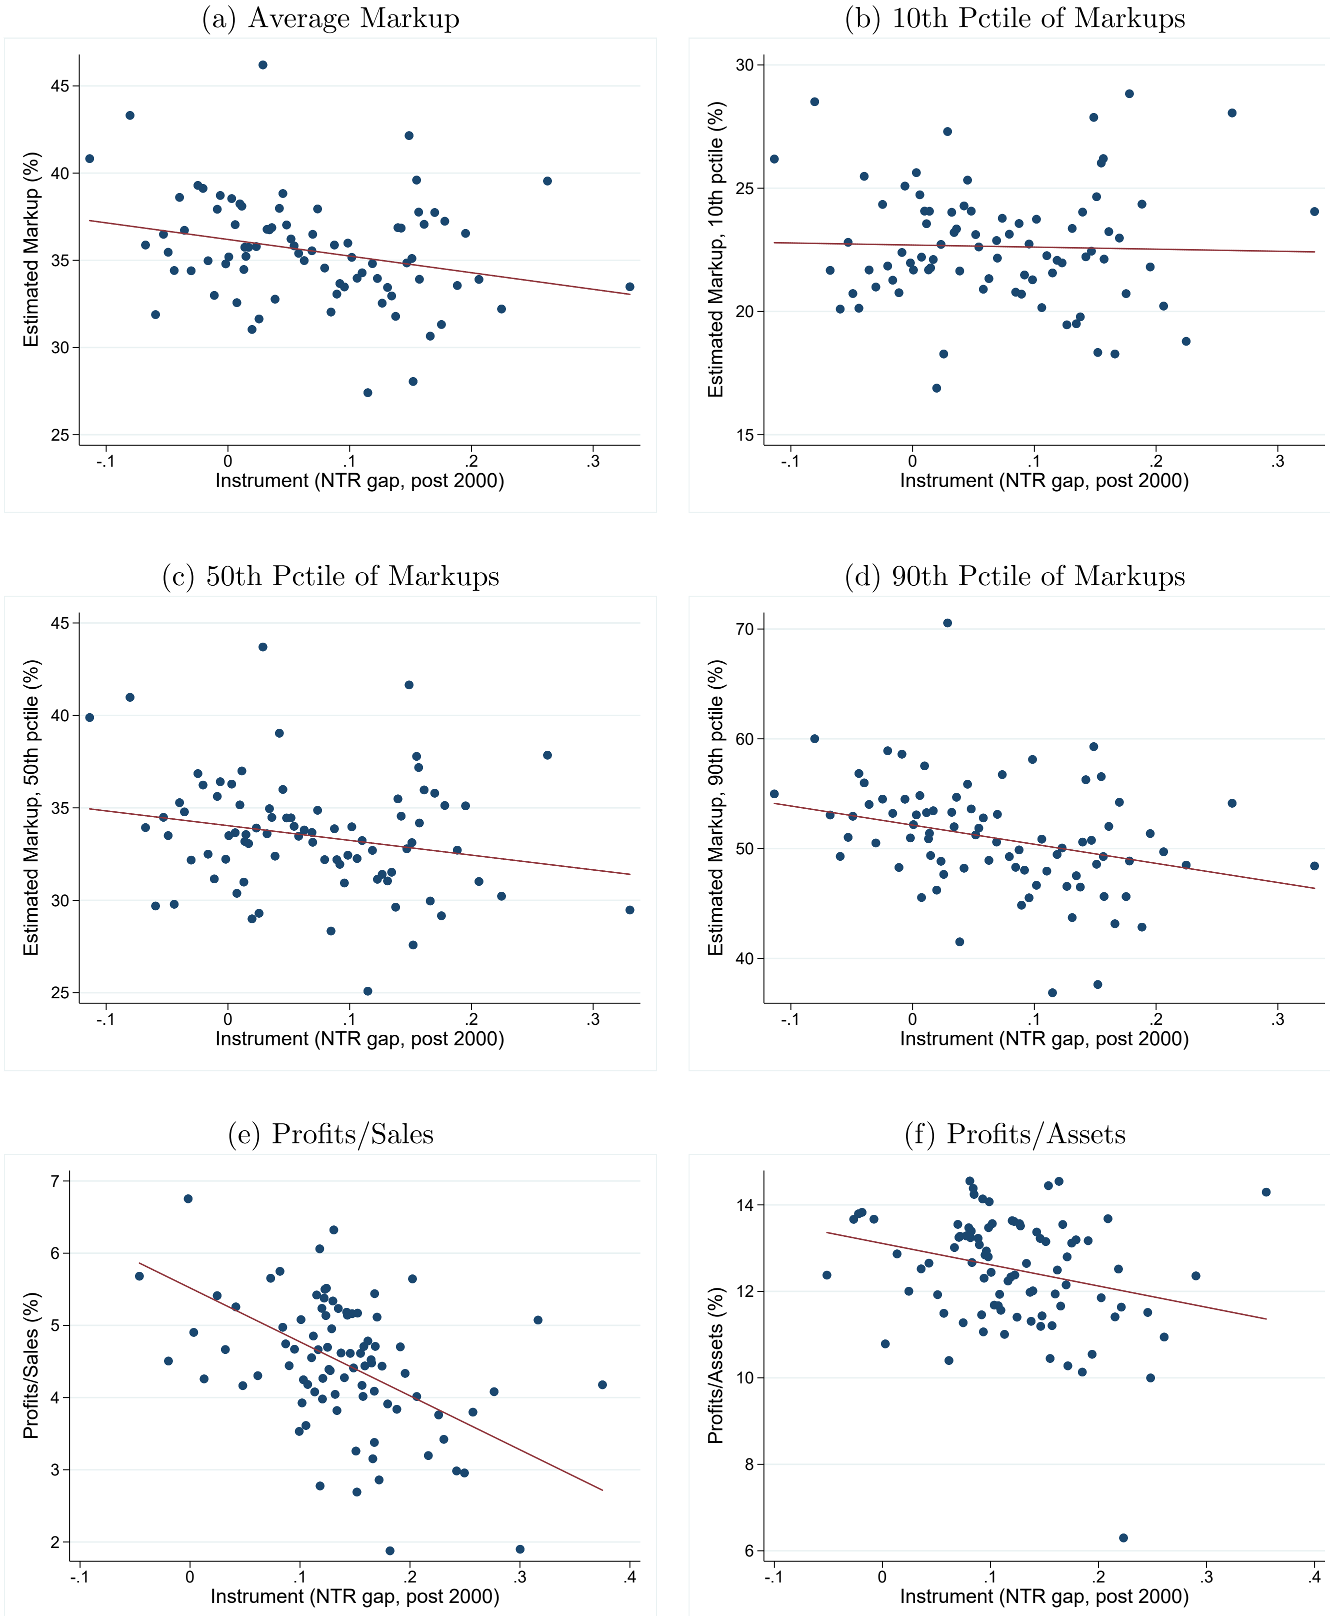

Notes: This figure reports the reduced-form relationships in the Compustat sample, described in Section IV.D. Each dot represents 1% of the data, with consumption weights. The level of observation is a NAICS industry-by-period.

Table 1: Summary Statistics

	Mean	S.D.	p10	p50	p90	Observations	
						N	Aggreg. Level
Inflation, all (%)	1.15	6.75	-7.48	2.21	7.57		
Share of continued products (%)	80.75					3,774	ELI-by-year
Share of unavailable products	4.92						
ΔChina IP in U.S. (pp, annualized)	0.66	1.62	0.00	0.00	2.63	444	ELI-by-period
ΔChina IP in developed economies	0.47	1.07	0.00	0.01	1.63		
NTR Gap	0.21	0.21	0.00	0.13	0.55		
Share of Goods	0.78						
Share of Durables	0.18					222	ELI
Share of Apparel	0.11						
Share of High Tech	0.08						

Notes: This table presents summary statistics for the main variables used in the analysis, which are described in Sections II.A and II.B. The sample covers years 1991 to 2007, which are divided into two periods: 1991-1999 and 2000-2007. Depending on the variable, the level of observations is an ELI-by-year, an ELI-by-period, or an ELI.

Table 2: Baseline Instrumental Variable Estimates

Panel A: With the NTR Gap

	Δ China IP (pp)	U.S. CPI Inflation (pp)			
	OLS	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)
NTR Gap	3.33^{***}	-7.43^{***}			
	(0.54)	(2.21)			
Δ China IP (pp)			-2.23^{***}	-0.90^{***}	-2.15^{***}
			(0.47)	(0.32)	(0.77)
First-stage F			38.14		23.13
ELI F.E.	✓	✓	✓	✓	
Period-specific Goods F.E.	✓	✓	✓	✓	
2000-2007 only					✓
Goods, Durables & Apparel F.E.					✓
N	444	444	444	444	222

Panel B: With the Change in Import Penetration from China in Other Developed Economies

	U.S. CPI Inflation (pp)		
	IV	IV	IV
	(1)	(2)	(3)
Δ China IP (pp)	-1.44^{***}	-1.27^{***}	-1.91^{***}
	(0.45)	(0.28)	(0.38)
First-stage F	26.23	405.69	27.34
Hansen J			0.21
ELI F.E.	✓		✓
Period-specific Goods F.E.	✓		✓
2000-2007 only		✓	
Goods, Durables & Apparel F.E.		✓	
Instruments:			
Δ China IP Other	✓	✓	
NTR Gap & Δ China IP Other			✓
N	444	222	444

Notes: The specifications are described in Section III.C. The level of observation is an ELI-by-period cell. The sample includes all ELIs from 1991 to 2007, with variables averaged over two periods, 1991-1999 and 2000-2007. Consumption weights are used. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level.

Table 3: Robustness of IV Estimates

Panel A: Specifications with Alternative Sets of Controls

	U.S. CPI Inflation (pp)					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ China IP (pp)	-2.75^{***}	-1.78^{***}	-2.26^{***}	-2.49^{***}	-2.10^{***}	-2.94^{**}
	(0.79)	(0.59)	(0.48)	(0.61)	(0.62)	(1.43)
First-stage F	30.23	24.01	37.50	26.59	23.19	9.541
Major Category F.E.	✓					
ELI F.E.		✓	✓	✓	✓	
Period-specific Goods F.E.	✓	✓	✓	✓	✓	✓
Durables & Apparel Time-Varying F.E.		✓				
Excluding Deflationary ELIs			✓			
Time-varying controls for High-tech, Contract intensity and Union membership					✓	
1990s inflation by period F.E.				✓		
6-digit IO Fixed Effects						✓
Instrument: NTR Gap	✓	✓	✓	✓	✓	✓
N	444	444	400	444	444	170

Panel B: Specifications with the Overall Change in Import Penetration

	Δ All IP (pp)	U.S. CPI Inflation (pp)	
	OLS	IV	IV
	(1)	(2)	(3)
Δ China IP (pp)	0.78^{***}		
	(0.15)		
Δ All IP (pp)		-3.68^{**}	-3.67^{**}
		(1.60)	(1.36)
First-stage F		27.59	18.55
ELI F.E.	✓	✓	
Period-specific Goods F.E.	✓	✓	✓
Durables & Apparel Time-Varying F.E.	✓	✓	✓
Instrument: NTR Gap		✓	✓
N	444	444	444

Notes: The level of observation is an ELI-by-period cell and the sample includes all ELIs from 1991 to 2007, with variables averaged over two periods, 1991-1999 and 2000-2007. Column (5) of Panel A is an exception: the data is aggregated from ELIs to 6-digit industries defined in the BEA's IO table. Consumption weights are used. Standard errors are clustered by ELIs or 6-digit IO industries. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table 4: Falsification Tests with the French CPI Data

Panel A: Placebo IV in France

	Δ China IP (in U.S., pp)	U.S. CPI Inflation (pp)	French CPI Inflation (pp)		
	OLS	IV	OLS	IV	IV
	(1)	(2)	(3)	(4)	(5)
NTR Gap	3.22***	−2.59***	−0.24		
	(0.71)	(0.904)	(1.82)		
Δ China IP (in U.S., pp)				−0.074	−0.27
				(0.38)	(0.91)
First-stage F		20.71		20.71	15.73
COICOP F.E.	✓	✓	✓	✓	
Period-specific Goods F.E.	✓	✓	✓	✓	
2001-2007 only					✓
Goods, Durables & Apparel F.E.					✓
N	264	264	264	264	132

Panel B: Triple-Difference IV

	Δ China IP (in U.S., pp)	U.S. Infl. *Minus* French Infl. (pp)		
	OLS	OLS	IV	IV
	(1)	(2)	(3)	(4)
NTR Gap	3.22***	−8.12**		
	(0.71)	(3.27)		
Δ China IP (in U.S., pp)			−2.52**	−2.08**
			(1.09)	(0.93)
First-stage F			20.71	15.73
COICOP F.E.	✓	✓	✓	
Period-specific Goods F.E.	✓	✓	✓	
2000-2007 only				✓
Goods, Durables & Apparel F.E.				✓
N	264	264	264	132

Notes: The specifications are described in Section III.D. The level of observation is a COICOP-by-period cell, with variables averaged over two periods, 1991-1999 and 2000-2007. Consumption weights are used. Standard errors are clustered by COICOPs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table 5: IV Estimates with the Log Domestic Expenditure Share

	U.S. CPI Inflation		
	IV	IV	IV
	(1)	(2)	(3)
Δ Log Domestic Expenditure Share	2.57***	3.46**	3.10***
	(0.9601)	(1.411)	(0.961)
First-stage F	13.211	17.197	11.599
Hansen J			0.568
ELI F.E.	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓
Instruments:			
NTR Gap	✓		✓
Δ China IP Other		✓	✓
N	444	444	444

Notes: This table reports the IV estimates with the log change in the domestic expenditure share as the endogenous variable (the choice of the endogenous variable is the only difference with equation (3) in the main text). As described in Section II.A, the trade data is measured at the level of HS codes, while the domestic production data comes from the NBER-CES Manufacturing database. Column (1) uses the NTR gap instrument, Column (2) uses the change in the import penetration rate from China in other developed economies, and Column (3) uses both instruments jointly. The Hansen J statistic in Column (3) indicates that we cannot reject the overidentification restriction. Consumption weights are used. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table 6: The Roles of Continued and Domestic Goods

Panel A: IV Estimates for Continued Goods in Main Sample (CPI)

	U.S. CPI Inflation, Continued Products (pp)		Contribution to U.S. CPI Inflation (pp) *[%]*	
	(1)	(2)	(3)	(4)
Δ China IP (pp)	-3.00^{***}	-3.23^{***}	-1.54^{***} *[69%]*	-1.54^{***} *[72%]*
	(0.79)	(1.62)	(0.46)	(0.74)
ELI F.E.	✓		✓	
Period-specific Goods F.E.	✓		✓	
2000-2007 only		✓		✓
Goods, Durables & Apparel F.E.		✓		✓
N	444	222	444	222

Panel B: IV Estimates for Domestic Goods in Main Sample (CPI)

	U.S. CPI Inflation, Domestic Products (pp)		Contribution to U.S. CPI Inflation (pp) *[%]*	
	(1)	(2)	(3)	(4)
Δ China IP (pp)	-1.94^{***}	-2.73^{***}	-0.98^{**} *[44%]*	-1.82^{***} *[85%]*
	(0.59)	(0.96)	(0.42)	(0.63)
ELI F.E.	✓		✓	
Period-specific Goods F.E.	✓		✓	
2000-2007 only		✓		✓
Goods, Durables & Apparel F.E.		✓		✓
N	444	222	444	222

Panel C: IV Estimates for Continued and Domestic Goods in PPI Sample

	U.S. PPI Inflation (pp)		U.S. PPI Infl., Continued Products (pp)	
	(1)	(2)	(3)	(4)
Δ China IP (pp)	-2.50^{**}	-1.86^{**}	-2.02^{**}	-1.66^{**}
	(1.01)	(0.78)	(0.93)	(0.81)
First-stage F	19.22	20.07	19.22	20.07
NAICS F.E.	✓		✓	
Period-specific Goods F.E.	✓		✓	
2000-2007 only		✓		✓
Goods, Durables & Apparel F.E.		✓		✓
N	550	275	550	275

Notes: Panel A and B use the main analysis sample, while panel C uses the PPI sample. The specifications are described in Section IV.B. In all panels, the instrument is the NTR gap and the level of observation is an industry-by-period cell. First-stage F statistics in Panel A and B are the same as in Table 2. Standard errors are clustered by industries. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table 7: Mechanisms

Panel A: Indirect Exposure to Trade with China

	U.S. CPI Inflation (pp)			
	(1)	(2)	(3)	(4)
Δ China IP (pp)	-2.943^{**}	-3.214^{**}	-2.892^{**}	-3.241^{**}
	(1.435)	(1.539)	(1.427)	(1.582)
First-stage F	9.541	6.691	8.727	6.270
Controls:				
Δ China IP Supplier (pp)		✓		✓
Δ China IP Buyer (pp)			✓	✓
6-digit IO F.E.	✓	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓	✓
N	170	170	170	170

Panel B: Offshoring and Returns to Scale

	U.S. CPI Inflation (pp)	
	(1)	(2)
Δ China IP (pp)	-1.685^{***}	-2.01^{***}
	(0.403)	(0.70)
Δ China IP ×Interaction	−0.074	0.43
	(1.274)	(0.71)
First-stage F	9.403	8.05
Interacted indicators:		
Related Trade > p90	✓	
Capital Intensity > Median		✓
Period-specific Goods F.E.	✓	✓
Durables & Apparel Time-Varying F.E.	✓	✓
N	306	306

Panel C: Wages and Total Factor Productivity

	TFP Growth (pp)		Wage Growth (pp)		
	4-factor TFP	5-factor TFP	All	Production	Non-production
	(1)	(2)	(3)	(4)	(5)
Δ China IP (pp)	-0.629^{**}	-0.632^{**}	0.0806	0.0856	0.4105
	(0.290)	(0.292)	(0.1249)	(0.1726)	(0.2994)
First-stage F	24.155	24.155	26.328	26.328	26.328
Period-specific Goods F.E.	✓	✓	✓	✓	✓
Durables & Apparel Time-Varying F.E.	✓	✓	✓	✓	✓
N	300	300	306	306	306

Notes: In Panel A, the level of observation is a 6-digit IO industry-by-period cell. In Panels B and C, the level of observation is an ELI-by-period cell. The sample is restricted to ELIs that can be matched to the NBER-CES Manufacturing database.In all panels, the instrument is the NTR gap.*** denotes statistical significance at the 1% level, ** at the 5% level.

Table 8: The Role of Markups

Panel A: The Response of Estimated Markup in Compustat Sample

	Δ China IP (pp)	U.S. Markups (pp)	
	OLS	OLS	IV
	(1)	(2)	(3)
NTR Gap	5.414^{***}	-9.52^{***}	
	(1.051)	(4.34)	
Δ China IP (pp)			-1.75^{**}
			(0.848)
First-stage F			26.49
NAICS F.E.	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓
N	796	796	796

Panel B: Heterogeneity by Market Structure in PPI sample

	U.S. PPI Inflation (pp)				
	Interacted Specs.	Subsample Specs.			
	(1)	(2)	(3)	(4)	(5)
Δ China IP (pp)	-0.47^{***}	-3.47^{***}	-0.31	0.18	-1.88^{**}
	(0.26)	(1.72)	(0.40)	(0.57)	(0.82)
Δ China IP×High Concentration	-1.70^{**}				
	(0.96)				
Δ China IP×High China IP	2.31^{**}				
	(1.08)				
First-stage F	111.85	209.20	174.75	26.10	296.58
NAICS F.E.	✓	✓	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓	✓	✓
Subsample	All	High Conc.	Low Conc.	High China IP	Low China IP

Notes: In both panels, the level of observation is a 6-digit NAICS-by-period cell. Panel A uses the Compustat sample while Panel B uses the PPI sample. In Panel B, "High Concentration" product categories have a level of domestic market concentration above median in 1997 (resp. below for "Low Concentration"). "High China IP" product categories have an import penetration rate from China above the 75th percentile in 1999 (resp. below for "Low China IP"). Standard errors are clustered by 6-digit NAICS industries.*** denotes statistical significance at the 1% level, ** at the 5% level, * at the 10% level.

Table 9: Consumer Surplus per Displaced Job across Product Categories

Panel A: Summary Statistics across Product Categories (year 2000, goods only)

	Total	S.D.	p10	p50	p90
Average Labor Earnings ($)	33,305	10,276	21,318	28,875	43,321
Value-Added of Domestic Producers ($) / Job	121,897	179,052	55,531	98,172	268,606
Total Sales of Domestic Producers ($) / Job	305,250	262,578	103,795	225,279	545,720
Domestic Absorption ($) / Job	390,998	376,320	150,591	357,974	625,686
$N = 174$					

Panel B: Estimates of Consumer Surplus per Displaced Job

	Uniform 1pp Increase in Import Penetration from China			Observed Change in Import Penetration from China, 2000-2007		
	(1)	(2)	(3)	(4)	(5)	(6)
Consumer Surplus per Displaced Job, $	477,555	317,383	411,464	433,565	288,147	373,562
IV Estimates:						
- NTR gap: $\beta_{price} = -2.23$ $\beta_{emp} = -1.834$	✓			✓		
- ΔChina IP Other: $\beta_{price} = -1.44$ $\beta_{emp} = -1.774$		✓			✓	
- Both: $\beta_{price} = -1.91$ $\beta_{emp} = -1.815$			✓			✓

Notes: Panel A presents summary statistics for ELIs matched to the NBER-CES Manufacturing database. Panel B presents estimates of consumer surplus per displaced jobs in this sample, 2000 dollars. Columns (1) to (3) use $\frac{\beta_{price}}{\beta_{emp}} \cdot \frac{\sum_j Cons_j}{\sum_j Emp_j}$, and columns (4) to (6) use $\frac{\beta_{price}}{\beta_{emp}} \cdot \frac{\sum_j \Delta ChinaIP_j \cdot Cons_j}{\sum_j \Delta ChinaIP_j \cdot Emp_j}$.

Table 10: Distributional Effects via the Expenditure Channel

Panel A: IV Estimates across Subsamples

	U.S. CPI Inflation (pp)					
	(1)	(2)	(3)	(4)	(5)	(6)
2SLS: Δ China IP (pp)	-4.28^{***}	-0.91^{***}	-4.62^{**}	-0.83^{**}	-2.93^{***}	-1.18^{***}
	(1.59)	(0.35)	(1.94)	(0.40)	(1.00)	(0.41)
LIML: Δ China IP (pp)	-4.57^{***}	-0.93^{**}	-5.42^{**}	-0.84^{**}	-3.18^{***}	-1.22^{***}
	(1.80)	(0.37)	(2.61)	(0.42)	(1.17)	(0.43)
First-stage F	6.80	8.64	3.07	12.13	10.01	7.39
Hansen J	0.31	0.40	0.64	0.23	0.08	0.36
Above/below median?	$<$	$>$	$<$	$>$	$<$	$>$
Splitting variable	Sales Share to College Educ.		Expenditure Elasticity		Sales Share to Inc. $>$ \$60k	
ELI F.E. & period F.E.	✓		✓		✓	
N	166	166	166	166	166	166

Panel B: Estimates of Distributional Effects

	All Product Categories			Goods Only		
	(1)	(2)	(3)	(4)	(5)	(6)
Distributional Effects, $\frac{EV^{HI}-EV^{LI}}{EV^{All}}$:						
- College vs. non-college	6.19%	-9.64%	-13.94%	19.46%	3.20%	-1.22%
- Income above vs. below \$60k	8.39%	-19.54%	-26.60%	17.47%	-10.03%	-16.98%
- Income above \$100k vs. below \$30k	14.53%	-23.13%	-36.29%	26.97%	-9.35%	-22.04%
IV Estimates:						
- Homogeneous	✓			✓		
- Heterogeneous by sales share to inc. $>$ \$60k		✓			✓	
- Heterogeneous by expenditure elasticity			✓			✓

Notes: Panel A reports the results of IV specifications across subsamples, which are described in Section V.B. The level of observation is an ELI-by-period. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level. Panel B reports the estimates of distributional effects across groups, using the formula $\frac{EV^{HI}-EV^{LI}}{EV^{All}} = \frac{\sum_j \left(s_j^{HI}-s_j^{LI}\right)\beta_j \Delta ChinaIP_j}{\sum_j s_j^{All}\beta_j \Delta ChinaIP_j}$, described in Section V.B.

For Online Publication

Appendix to "What Are the Price Effects of Trade? Evidence from the U.S. and Implications for Quantitative Trade Models"

Xavier Jaravel, London School of Economics

Erick Sager, Federal Reserve Board

March 2024

Contents

A Comparison to Other Estimates in the Literature

In this appendix, we compare the magnitude of our estimates to other studies that have examined the relationship between trade, prices, and markups, in the United States and abroad. We first review the estimates from prior work, and then highlight the four contributions we make relative to this literature.

Literature review. Kamin et al. (2006) study the OLS relationship between Chinese exports and both US import prices and US producer prices. They document that a one percentage point rise in the Chinese import share of a given sector during 1997–2002 was associated with a 1.097 percentage point lower annual import inflation in that sector, which is similar to our OLS estimate for consumer prices (while the IV is larger). They do not find evidence of a significant impact on U.S. producer prices, in contrast with our results and with the other papers described below. Because their analysis with producer prices stops in 2001, the difference in time coverage could potentially explain this discrepancy with our results.

Auer and Fischer (2010) study the impact of U.S. imports from nine low-wage countries between 1997 and 2006. They use an instrument leveraging differences in labor-intensity across sectors, leveraging the idea that low-wage countries have a comparative advantage in labor-intensive industries. They find that when the nine low-wage countries capture a 1% share in a sector, U.S. producer prices decrease by 2.35%. This point estimate is close to ours.

Carluccio et al. (2018) assess the impact of imports from low-wage countries on CPI and PPI inflation in France during 1994-2014. For the PPI estimates, they use an IV strategy based on differences in labor-intensity across sectors, similar to Auer and Fischer (2010). The find that when the nine low-wage countries capture a 1% share in a sector, French producer prices decrease by 1.208%. The corresponding estimate when they focus on imports from China is 1.84%, which is close to our estimate for consumer prices in the U.S.[1]

Amiti et al. (2020) study the impact of China's WTO entry on U.S. producer prices. They measure prices in the U.S. using the U.S. PPI as well as data on the unit value of imports. They find that a key mechanism is China lowering its own import tariffs on intermediate inputs. Using an IV strategy exploiting changes in tariffs, they estimate that the fall in Chinese import tariffs upon WTO entry reduced Chinese firms' costs and in turn lowered their production cost and their export prices, which benefited the U.S. market. Relative to them, (i) we measure consumer prices in the U.S., rather than producer prices, (ii) we focus on a complementary mechanism, i.e. the importance

[1] To assess the impact on consumer prices, rather than producer prices, they conduct an accounting exercise at the level of 3-digit COICOP categories, which is not directly comparable to our analysis.

of the price response for domestically-produced goods through strategic pricing, and (iii) we show how to interpret the cross-industry estimates through the lens of standard trade models despite the missing intercept. Although Amiti et al. (2020) do not focus on estimating the relationship between changes in import penetration and prices in the U.S., the corroborative evidence reported in their Table 9 imply a large decline in U.S. producer prices for small changes in import penetration from China, with magnitudes close to our estimates.[2] [3]

Bai and Stumpner (2019) estimate the impact of trade with China in U.S. consumer prices using Nielsen scanner data, from 2004 to 2015. The scanner data has the key advantage of measuring changes in product variety over time and heterogeneity across household groups within industries. They find that trade with China led to lower consumer prices for continued products and an increase in product variety in the U.S. While this study focuses on the product category covered in the Nielsen data, much of trade with China occurs in other categories, including consumer electronics, apparel, and slow-moving consumer goods. Our estimates complement this study by examining a broader sample over a different time period. In particular, we are able to test for pre-trends and conduct several test for the identification strategy using the instrument of Pierce and Schott (2016). Moreover, Bai and Stumpner (2019)'s point estimates are lower than ours: implementing a regression across product categories, they estimate that a 1pp increase in the import penetration rate from China is associated with a 0.78% fall in the price index (see their Table 1[4]). In Appendix Table A1, we implement our IV specification in a sub-sample of products that approximates the sample of Bai and Stumpner (2019), focusing on consumer packaged goods. We estimate a smaller

[2]Amiti et al. (2020) report a negative relationship between the change in the Chinese export price index and the change in the import penetration rate from China in Column (1) of their Table 9, with a point estimate of -0.376 (s.e. 0.194), and a positive relationship with the price index for domestically-produced goods, measured with U.S. shipment deflators in Column (5) of their Table 9, with a point estimate of 0.635 (s.e. 0.164). These two point estimates imply that, when the import penetration rate from China increases by 1pp, the producer price index for domestically-produced goods falls by 1.69pp ($= 0.635/(-0.376)$). This implied relationship is statistically indistinguishable from the point estimates we obtain when focusing on domestically-produced goods in Panels B and C of our Table 6. For example, we obtain a point estimate of -1.94 (s.e. 0.59) in Column (1) of Panel B of Table 6.

[3]Amiti et al. (2020) obtain a large pass-through of Chinese input-tariff reductions on the price of Chinese exports to the U.S. and on the price of firms from non-China countries selling in the U.S. Their estimated pass-through is greater than unity, which cannot be rationalized by standard oligopolistic competition models (e.g., Atkeson and Burstein (2008), Burstein and Gopinath (2014)). In contrast, in Section IV.D we use a standard oligopolistic competition model to account for our estimates of the price response to increased import competition. Given prior evidence on cost passthrough rates below one (e.g., Amiti et al. (2019a)), our assessment is that the cost passthrough above unity found in Amiti et al. (2020) is an interesting feature that is relevant for the market for intermediate inputs in China, but that it does not preclude the use of standard oligopolistic competition models (with cost passthrough below one) to rationalize our estimates.

[4]Regressing the change in the price index on the log change in the domestic expenditure share, Bai and Stumpner (2019) obtain a coefficient of 0.500 (s.e. 0.122) in Column 1, Panel D of their Table 1. To compare this point estimate with our specification, using the change in the import penetration rate from China as the independent variable, we rescale their estimates by the initial domestic expenditure share, which is equal to 82% for all HS codes in their sample, as mentioned on their page 5. In addition, we adjust for the divergence between the change in the domestic expenditure share and the change in the import penetration rate from China, using the point estimates from Column 1 of Panel A and Column (1) of Panel B of their Table 1. Thus, the relevant coefficient for comparison with our estimates in Tables 2 and A1 is: $-0.78 = 0.500 \times (-2.630/2.054) \times (1/0.82)$.

price response, with a point estimate of 0.68 (s.e. 0.39), statistically indistinguishable from theirs. These results highlight the importance of considering the full consumption basket.

De Loecker et al. (2016) study trade liberalization in India, examining declines in input and output tariffs between 1989 and 1997. They find that the change in input tariffs is the dominant force and leads to an increase in the markup of domestic producers. Input tariff liberalization leads to a fall in input costs, with incomplete cost pass-through to prices and increased domestic markups. De Loecker et al. (2016) also provide evidence that output tariff liberalization exerted pro-competitive pressure on markups, but they find that this effect is smaller that the upward pressure on markups from falling inputs costs.[5] The overall effect, taking into account the average declines in input and output tariffs between 1989 and 1997, is that markups, on average, increased by 12.6 percent following trade liberalization. Accordingly, the authors conclude that, because of markup adjustments, "producers benefited relative to consumers", and that "the short-run gains to consumers appear small, especially considering that we observe factory-gate prices rather than retail prices". In contrast, we study a developed country setting and focus on the import competition shock from China, studying retail prices faced by consumers. We find large price effects benefiting consumers. Most of increased trade with China during the period we study is in final goods, and we do not attempt to study the input cost channel in detail, as we do not have the required firm-level import data.

Contributions relative to prior work. Although estimates vary across studies and are not always comparable, there is a growing empirical consensus that foreign supply shocks have substantial effects on producer prices and markups. Nonetheless, prior work has not examined whether the observed price response can help discriminate between different trade models. Indeed, the predictions of quantitative trade models are usually expressed in terms of economy-wide variables, such as the theoretical price index of a representative agent (see, e.g., Arkolakis et al. (2012)). Instead, we show how to provide a structural interpretation to the cross-sector regression coefficient, accounting for general equilibrium effects, and allowing for sectoral heterogeneity, the "missing intercept", as well as potential deviations between the theoretical price index and the measured CPI. Our approach to formally linking our regression specification to standard quantitative trade models, in particular the Melitz-Chaney model, is new to the literature[6] and provides a principled methodology to distinguish between classes of quantitative trade models based only on cross-sector estimates. In addition, we make another three contributions to the literature:

First, we provide an analysis of the reduced-form effects of trade using micro-data on the

[5]Their estimate is not directly comparable to ours because they do not estimate the equilibrium relationship between changes in domestic markups and changes in import penetration induced by a foreign supply shock.

[6]Indeed, prior work did not point out that the relationship between price changes and changes in expenditure share is about one order of magnitude larger than predicted by standard quantitative trade models that abstract from strategic price-setting.

prices that *consumers* actually face at the store. This is not the case in almost any of the studies mentioned above. These studies use micro data on import/export prices or producer prices, rather than retail prices, i.e. they do not take into account potential changes in retail margins, which are typically large for traded goods and should matter for consumers. To the best of our knowledge, only Bai and Stumpner (2019) study the impact of trade using data on retail prices. Although Bai and Stumpner (2019) have micro data on consumer prices linked to trade shocks, as previously discussed their sample is subject to limitations that we can relax, in terms of coverage of product categories and time periods.

Second, we provide a simple comparative analysis of job losses and price changes in response to increased trade with China for the United States. Our sample frame allows us to compare the job displacement estimates, using the same research design as in prior studies, to the price effects in a unified empirical framework.

Third, our setting lends itself to an in-depth analysis of the identification assumption, leveraging several instruments and a series of falsification and robustness tests. We show that the magnitude of the price effects is robust to multiple potential concerns about causal identification that could not be addressed in prior work due to data limitations; we can thus establish that the large estimated magnitude of the price response is not due to confounders and is an important identified moment for quantitative trade models to match.

Related literature on passthrough estimation. A large literature has estimated the passthrough of a firm's production costs into its own prices (e.g., Amiti et al. (2020)), or the passthrough of competitor prices into a focal firm's prices (e.g., Amiti et al. (2019a)).[7] In our setting, it is challenging to use import prices from China, computed as unit values, to estimate a passthrough rate of prices of imports into retail consumers prices. Indeed, the unit values of imports must be computed within "products (e.g., HS10) by country of origin" cells, raising standard concerns about composition effects due to a potential change in the underlying set of products within each cell. Our focus on the relationship between import shares and consumer prices across industries addresses the concerns about composition effects for unit values, without loss in terms of interpretability since our reduced-form specification retains a structural interpretation in standard quantitative trade models.

[7] See also, among many others, Burstein and Gopinath (2014) and Cavallo et al. (2021).

B Data Appendix

B.A Consumer Price Index Data

This section contains information about the Consumer Price Index (CPI). For additional information, we refer the reader to chapter 17 of the BLS Handbook of Methods (U.S. Bureau of Labor Statistics (2018)).

Overview. Our price data set is known as the CPI Research Database (CPI-RDB), which is maintained by the Division of Price and Number Research at the Bureau of Labor Statistics. This is a restricted access data set that contains the micro data underlying the non-shelter component of Consumer Price Index (CPI). The CPI-RDB contains all product-level prices on goods and services collected by the BLS for use in the CPI since January 1988.[8] Although the number of individual prices used to construct the CPI has changed over time, the BLS currently collects data on approximately 80,000 core products and 130,000 total products per month from about 27,000 retail outlets across 87 geographical areas in the United States. The sampling frame for the non-shelter component of the CPI represents about 70% of consumer expenditures. Descriptions of and summary statistics for the CPI-RDB in prior years of data can be found in Bils and Klenow (2004), Nakamura and Steinsson (2008), and Bils et al. (2012) who base their research on the same data set.

The CPI-RDB contains prices and sampling weights for each individual item in the non-shelter component of the CPI. We use the CPI-RDB to construct inflation by disaggregated categories called Entry Level Items (ELIs). The BLS defines ELIs for the practical construction of the CPI. There are nearly 360 ELIs between 1988-1998 and 270 ELIs after a 1998 revision of definitions. We collapse the number of ELIs to 222 in order to maintain a consistent definition before and after a 1998 revision to the ELI structure. Examples of ELIs are "Carbonated Drinks," "Washers & Dryers," "Woman's Outerwear," and "Funeral Expenses".

The item structure of the CPI is grouped from broadest to most narrow product category: Major Groups, Item Strata, and ELIs. The Appendix Table of Konny et al. (2019) provides the list of ELIs, item structure, item weights and number of quotes contained in the CPI as of August 2018 (note that the sample used in our paper does not draw upon post-2007 data; the full list corresponding to our sample is available upon request).

Index Construction. The BLS constructs a *matched-model* price index, which means that the BLS selects a set of products and then collects the prices for those products over time. This enables the BLS to construct price changes for the same product each month. These price changes are then aggregated to construct elementary price indexes for each product category.

[8]The CPI-RDB was extended to cover 1977-1987 by Nakamura et al. (2018).

Sampling. The sampling frame for the non-shelter component of the CPI represents about 70% of consumer expenditures. Although the number of individual prices used to construct the CPI has changed over time, the BLS currently collects data on approximately 80,000 products per month from about 23,000 retail outlets across 87 geographical areas in the United States.

The BLS uses a Point-of-Purchase survey to identify the probability that consumers visit a particular outlet within a defined geographical region. Upon selecting outlets that are representative of consumer points-of-purchase, the BLS uses data provided by individual outlets on sales by product (within the specified ELI) in unison with the Consumer Expenditure Survey to construct probabilities that each product is purchased. Therefore, products within an ELI and geographical region are chosen to be representative by outlet and by product within the outlet. In the last step of sample construction, the BLS constructs multiple representative samples of products and chooses the one that minimizes sampling error through a sample variance reduction algorithm.

Aggregation. Aggregation to the ELI proceeds as follows. Let $p_{i,t}$ be a price quote within a given product category (ELI) in a month t, and let $\omega_{i,t-1}$ be its accompanying sampling weight. Following BLS procedure, we aggregate individual price quotes to the product category level using a *Geometric Laspeyres Index* (or, as it is alternatively called, the "Geometric Mean Index"), in which quantity information is incorporated through the share of expenditures in the base period,

$$I_t \equiv \exp\left(\sum_{i=1}^{N} \omega_{i,t-1} \log\left(\frac{p_{i,t}}{p_{i,t-1}}\right)\right)$$

where N is the supposed number of price quotes collected between times $t-1$ and t, and the sampling weight $\omega_{i,t-1}$ measures product-level expenditures,[9]

$$\omega_{i,t-1} \equiv \sum_{i=1}^{N} \frac{p_{i,t-1}q_{i,t-1}}{\sum_{j=1}^{N} p_{j,t-1}q_{j,t-1}} \ .$$

Relationship between the measured CPI and consumer welfare. To relate the Geometric Laspeyres Index to the change in consumer welfare, we follow Konüs (1939), Deaton (1989) and Fajgelbaum and Khandelwal (2016) and define the change in welfare for the representative consumer as the equivalent variation $EV_{t-1,t}$ divided by initial expenditures X_{t-1}, which we denote $d\log U_t$. Considering a fixed set of products and small shocks to prices, denoted $d\log p_t \equiv log(p_{i,t}) - log(p_{i,t-1})$, the envelope theorem (Roy's identity) implies that price changes affect consumer welfare in proportion to the initial spending shares across products indexed by i, regardless of the demand system: to the first order, the change in consumer welfare is given by $d\log U_t = -\sum_{i=1}^{N} \omega_{i,t-1} d\log p_t =$

[9] The BLS weighting procedure for aggregation to the product category level has two components. First, the main product-level weighting is performed by BLS through probability sampling, i.e. through the selection of retail outlets and individual products within those outlets. Second, the CPI-RDB provides additional weights for each product-level price that correct for sampling error to ensure weights reflect expenditure shares.

$1-\sum_{i=1}^{N}\omega_{i,t-1}\left(\frac{p_{i,t}}{p_{i,t-1}}\right)$.[10] The Geometric Laspeyres Index thus provides a first-order approximation to the change in consumer welfare (see also Konus and Byushgens (1926)).[11] Changes in product variety may introduce first-order changes in consumer welfare that are not captured by the measured CPI, which we discuss in Appendix D.A.[12]

Item Rotation. New products are phased into the CPI once every four years after initial introduction to the index. In other words, about a quarter of items in the CPI are newly introduced within any given year. After the BLS identifies a new outlet and product, the new product is not included in the CPI until a price is recorded for two consecutive periods, thereby creating a record for the item's initial price change, $p_{i,t}/p_{i,t-1}$, for inclusion in the CPI.

Forced Substitutions and Imputations. When a product is unable to be priced in a given month, the BLS implements one of two types of procedures. If the product is only temporarily unavailable, then the BLS imputes a value to the missing price observation. This value tends to be the average price change of all available products, which is therefore equivalent to dropping that product's price change from the index for the period. If the product is no longer available at an outlet, then there are two types of substitutions. The first is a "comparable substitution", which replaces the previous item with one that is similar in sufficiently many dimensions to consider it the same fundamental item. In this case there is no quality adjustment applied to the prices of the new or old product versions. The second is a "non-comparable substitution", which occurs when there is no available item that is a sufficiently close substitute to the old. In this case, the BLS implements a quality adjustment to net out the difference in price between old and new version that can be attributed to differences in underlying product characteristics. We refer the reader to Moulton and Moses (1997) for a discussion of each type of non-comparable substitution in practice. Bils et al. (2012) document that from 1990-2009 the monthly rate of forced item substitutions is approximately 3 percent and the monthly rate of temporary unavailability is 12 percent.

From the perspective of price index construction, the quality adjustment can be understood as follows. Suppose product i is currently in its v-th vintage or version. Let φ_i^v be consumers' perceived quality from version v of product i. Likewise $p_{i,t}^v$ is the price of the v-th vintage or version of product

[10] This result holds under standard regularity conditions (see Borusyak and Jaravel (2021), Proposition 1).

[11] Indeed, using $\log(\frac{p_{i,t}}{p_{i,t-1}})=\frac{p_{i,t}}{p_{i,t-1}}-1$ and $\exp(\frac{p_{i,t}}{p_{i,t-1}}-1)=\frac{p_{i,t}}{p_{i,t-1}}$ for small price changes, we have:

$$I_t \equiv \exp\left(\sum_{i=1}^{N}\omega_{i,t-1}\log\left(\frac{p_{i,t}}{p_{i,t-1}}\right)\right)=\exp\left(\sum_{i=1}^{N}\omega_{i,t-1}\left(\frac{p_{i,t}}{p_{i,t-1}}\right)-1\right)=\sum_{i=1}^{N}\omega_{i,t-1}\left(\frac{p_{i,t}}{p_{i,t-1}}\right)=1-d\log U_t\ .$$

[12] Assuming there is no change in product variety, the Geometric Laspeyres Index is the exact price index for a CES price aggregator when the elasticity of substitution is 1. Indeed, the general CES price index is,

$$I_t^{CES}=\frac{P_t}{P_{t-1}}=\sum_{i=1}^{N}\omega_{i,t-1}\left(\frac{p_{i,t}}{p_{i,t-1}}\right)^{1-\sigma}\ ,$$

where σ is the elasticity of substitution and the CES price aggregator is $P_t=\left(\sum_{i=1}^{N}p_{i,t}^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$.

i. The quality-adjusted price is $p^v_{i,t}/\varphi^v_i$ and therefore the associated quality-adjusted price change in the absence of a substitution is,

$$\frac{p^v_{i,t}/\varphi^v_i}{p^v_{i,t-1}/\varphi^v_i} = \frac{p^v_{i,t}}{p^v_{i,t-1}} \ .$$

When the BLS initiates a substitution, it compares two versions (denote them v and $v+1$) that have different underlying product characteristics and therefore different perceived quality from consumers. The quality-adjusted price change during a product substitution from version v to version $v+1$ of product i is,

$$\frac{p^{v+1}_{i,t}/\varphi^{v+1}_i}{p^v_{i,t-1}/\varphi^v_i} = \frac{1}{\varphi^{v+1}_i/\varphi^v_i} \times \frac{p^{v+1}_{i,t}}{p^v_{i,t-1}} \ ,$$

where $\varphi^{v+1}_i/\varphi^v_i$ is the ideal quality adjustment that the BLS approximates and nets out from price change at substitution.

Product Turnover. Sample attrition can stem from both planned rotations and forced substitutions. In both cases, an exiting product is replaced by a new product in order to maintain sample size by ELI. For the typical ELI, since about 25% of products are subject to a planned rotation each year and about 5% of products face a forced substitution, a given cohort of products is characterized by the following hazard rates: about 70% are still observed after 1 years, 34% after 3 years, 17% after 5 years, and a product is rotated out after this point. Most of the attrition is due to planned rotations.

Alternative Indices. We leverage the price micro data to build alternative category-level price indices, which we use for various robustness tests and extensions. Alternative category-level price indices help us address potential measurement issues. For example, the baseline CPI uses quality adjustments when the BLS data collector is unable to find the exact same product in the exact same store from one period to the next (e.g., the 500 ml bottle of Coca-Cola might no longer be on the shelf at Whole Foods and might have been replaced with a 500 ml bottle of Pepsi). Given that BLS quality adjustments may not perfectly account for potential changes in underlying product characteristics in such cases, we build an alternative price index based solely on price changes for "continued products" (i.e., those instances when the same item in the same store is observed from one period to the next). We also leverage the micro data to build alternative price indices that help decompose the sources of the price effects we document. For example, we can isolate the role of the price response of products made in the United States.[13]

Specification Checklists. When a BLS price collector prices an item for the first time, they

[13]See Section IV for a complete discussion. Note that such robustness tests and statistical decompositions would not be possible by using the publicly-available inflation series from the BLS. Another downside of the public data from the BLS, relative to the CPI-RDB data set, is that the publicly available product categories are coarser than ELIs and their definitions change over time; as a result it is difficult to build a balanced panel of detailed product categories over a long time horizon in this data set.

create a detailed description of its characteristics. This description is partially contained in a pre-written checklist that ensures the price collector records information that is necessary to identify that item upon returning to the outlet, or to identify an appropriate substitute for that item if it is no longer available. A specification checklist can be also be used to prevent inconsistencies in price collection from month to month.

This paper utilizes specification checklists to identify imported goods in the CPI. For some product categories, there is an explicit field for denoting the product's country of origin (either as a pre-designed checkbox or as a write-in field). For all categories, there are open fields that price collectors use to write information that has not been explicitly coded into check-boxes and will include country of origin if the United States is not the explicit product manufacturer. Product categories that tend to contain an explicit check box for country of origin are apparel, non-perishable food items, furniture and household furnishings, electronics, and motor vehicles.[14]

The procedure by which we identify country of origin is as follows. If a checkbox exists then we can find out whether the U.S. produced the product. If not, then we must rely on a write-in text field. We use a fuzzy text match to identify country of origin in such cases. Though no special denotation is required for domestically produced products, we searched for text strings that denote domestic production such as "United States", "USA", "US", "U.S", "U.S.A", "U.SA", "US.A", "USA.", "U. S.", "domestic", "Alabama", "Alaska", "Arizona", etc. as well as state abbreviations and the names of major U.S. cities. We also searched for text strings denoting non-domestic production such as "Import", "Impt", "Imprt", "Foreign" and the names and abbreviations of possible importing countries (including countries that existed earlier in the sample but do not exist in 2019) and foreign cities. Remaining cases for which a country of origin was not explicitly identified was assigned to the United States. To validate text matches, a random 10% sample of text fields was manually inspected and multiple text matching algorithms were implemented to ensure robustness.

Data Processing. Price collectors flag substitutions and abnormally large price changes, which analysts use to implement quality adjustments and imputations. However, in order to reduce the sensitivity of price indexes to exceptionally large price changes as well as avoid respondent disclosure (as per BLS disclosure avoidance policy) we exclude positive or negative price changes greater than 500%. These outliers occur rarely in the sample.

ELI definitions changed in 1998 and new ELIs have been introduced since 2007. There are nearly 360 ELIs between 1988-1998 and 270 ELIs after the 1998 revision of definitions. We collapse the number of ELIs to 222 in order to maintain a consistent definition before and after a 1998 revision to the ELI structure. To do this, we matched ELI categories based on the category descriptions

[14] In fact, the apparel industry lobbied Congress in the 1970s to require that country of origin be placed on tags by law.

available in the BLS' documentation for the CPI Research Database. The full list of ELIs and their average consumption weights over the sample period is available upon request.

Finally, we define a set of ELIs as *Durable Goods*, motivated by the set of products Bils (2009) studied. This set of products tend to be durable goods that require more use of quality adjustments than the rest of the product categories in the CPI. The ELIs in this list include "personal computers and peripherals", "telephones", "watches", "electric appliances", "refrigerator", "washers and dryers", "microwave ovens", "small kitchen appliances", "clocks", "televisions", "audio equipment", and "other video equipment"; the full list is available from Bils (2009). We use fixed effects for these durable goods to control for inflation trends that may be introduced through methodological issues in the construction of inflation measures for these products.

Data Limitations. Certain data limitations motivate the specifications we estimate. First, the CPI sampling methods limit what can be said of product variety growth over time. The CPI chooses a set of products and follows those products over time. Planned item rotations introduce new products to the sample, but the number of new items is pre-selected. While forced item substitutions could entail product turnover from old to new varieties, the number of price quotes in the CPI is not changed through forced substitution. Therefore, the CPI introduces new product varieties in a way that does not explicitly track the number of varieties in the economy's consumption set.

Second, we follow the aggregation procedures used by the BLS in constructing the CPI. This means that price change within an ELI assumes a particular elasticity of substitution, specifically a unit elasticity. This elasticity is motivated by a desire to allow for consumer substitution in utility, but without the proper data for identifying the structural elasticity of substitution in each ELI. However, this specification constitutes a general first-order approximation to the exact price indices of the quantitative models we study, which is sufficient for our main purposes. Moreover, a unit elasticity of substitution obviates us from identifying the elasticity's value in each ELI. In order to test the robustness of this assumption, we have built a CES index from the micro data, with an elasticity of substitution $\sigma = 5$ or $\sigma = 3$. The alternative elasticities of substitution yield IV point estimates that are very similar in magnitude to our headline estimates reported in Table in the main text.

Aggregation. Whenever we need to aggregate measured inflation from ELIs to a higher level of aggregation, we use weights based on Consumer Expenditure Surveys for each year from 1988–1995, 1999–2004 and 2008-2012. For all other years, we set weights equal to the most recently available year's weights (e.g., assign 1995 weights to 1997). We follow Bils and Klenow (2004), Bils et al. (2012) and Gagnon et al. (2013) in using weights based on the Consumer Expenditure Survey. These weights are also used as regression weights in the baseline specifications.

B.B Historical CPI Data

To check for pre-trends in the CPI data, it is useful to have a long time series. Accordingly, for pre-trend exercises in the CPI data we incorporate pre-1988 inflation by Entry Level Item into our analysis.

The CPI-RDB was extended to cover 1977-1987 by Nakamura et al. (2018). This data was scanned from microfilm cartridges and converted to digital format using Optical Character Recognition software. The final data set contains prices for 80,000 to 100,000 products per month. For each product, the data set contains the product's price (in level and percent change from the preceding period), a product identifier, the Entry Level Item (ELI) classifier for the product, an outlet identifier, the location of the outlet, a flag indicating whether the product was on sale, and a flag indicating whether the product underwent an item replacement procedure (and, if so, the flag indicates what type of quality adjustment or imputation was made).

Because ELI definitions changed after 1987, Nakamura et al. (2018) created a concordance that maps pre-1988 ELIs into post-1988 ELI definitions. We use their concordance to create a consistent set of ELIs across time.

B.C Crosswalks

Our data building process uses a total of eight crosswalks, including five new crosswalks we build by hand.

We build three many-to-one crosswalks to the ELI product categories that define our main analysis sample: from SIC industries to ELIs, from NAICS industries to ELIs, and from UCC consumption categories to ELIs. Because SIC, NAICS and CEX categories are significantly more detailed than ELI categories, a many-to-one match is convenient. Furthermore, we build a many-to-one match of ELIs to 6-digit IO industries from the BEA's 2007 input-output table (as the BEA's industries are a bit less detailed than ELIs). Finally, for the falsification test using French CPI data, we build a many-to-one crosswalk from the (less detailed) COICOP categories to ELIs. The match is made by hand according to a comparison of the description of the product descriptions (as well as individual item names contained in in the CPI-RDB and discussions with BLS analysts).

Finally, we rely on three crosswalks from prior work: HS to NAICS codes from Pierce and Schott (2012), NAICS to IO codes from Borusyak and Jaravel (2021), and SIC to NAICS codes from the U.S. Census Bureau.

B.D Variables based on the 2007 Input-Output Table

We use the BEA's 2007 input-output table to measure indirect exposure to trade with China. We follow the same data construction steps as in Acemoglu et al. (2016)'s study of input-output linkages, except that they use the 1992 IO table. The 2007 IO table is much more disaggregated and hence potentially more accurate.

Indirect exposure via supplier effects. For indirect exposure via supplier effects, we compute the change in the import penetration rate from China in industry j's total output. By definition this quantity will be small if value-added is a high share of industry j's output. In robustness check, we implement a similar procedure to compute the import penetration rate from China in industry j's intermediate inputs (instead of j's output). We carry out this analysis using the BEA's "Use table." In a pre-processing step, we must obtain a square industry Use Table, denoted U^I. Following the methodology of the BEA, we do so by pre-multiplying the original (non-symmetric) Use Table U by the commodity-normalized Make Table M^C (which is close to an identity matrix). The Make Table gives the share of total production of each commodity across all industries, and each of its columns sums up to one.[15]

Indirect exposure via buyer effects. For indirect exposure via buyer effects, we compute the change in the import penetration rate from China in "buyer industries", scaled by (1 - share of industry sales to final consumers). This quantity is by definition low if an industry is primarily selling to final consumers. We carry out this analysis using the BEA's "Use table."

B.E Producer Price Index Data

This section contains information about the Producer Price Index (PPI). Further information is available from the chapter 14 of the BLS Handbook of Methods (U.S. Bureau of Labor Statistics (2018)).

Overview. We use data from the PPI's Research Database (PPI-RDB) from January 1987 to August 2008.[16] The BLS defines PPI prices as "net revenue accruing to a specified producing establishment from a specified kind of buyer for a specified product shipped under specified transaction terms on a specified day of the month." Accordingly, BLS requests (via fax or email) each establishment in the PPI sample to report the price of actual shipments transacted, as of the Tuesday of the week containing the 13th of the month. If an establishment fails to respond in a given month, the BLS price collector follows up with a phone call.

[15] U is commodity (row) by industries (columns), while U^I is industry by industry. M^C is (row) by commodities (columns). Each column of M^C sums to one because it reflects the share of production of each commodity produced by each industry.

[16] See Nakamura and Steinsson (2008) and Goldberg and Hellerstein (2009) for additional details about the PPI-RDB.

Sample Frame. The BLS collects prices from approximately 25,000 to 30,000 establishments for approximately 100,000 individual items on a monthly basis. The sample is constructed from the universe of establishments in the U.S., derived from the Quarterly Census of Employment and Wages business register that is collected in the enforcement of unemployment insurance programs in each U.S. state. Individual establishments within an industry are chosen probabilistically based on the total value of shipments, or total number of employees. Individual items are then selected by a BLS price collector during a field visit to the establishment according to value of shipment.

Industries are defined as a 6-digit NAICS category and span goods producing sectors (e.g., mining, manufacturing, agriculture, fishing, forestry, energy and construction industries) for the whole sample. Service sector industries were introduced to the PPI in 2005, which we exclude from our analysis.

Index Construction. The PPI constructs a *matched-model* price index, much like the CPI does. Once prices have been recorded for an item i at times $t-1$ and t, we can compute price change as $p_{i,t}/p_{i,t-1}$. These price changes are aggregated to the 6-digit NAICS classifications (or any high aggregations thereof) according to a *Laspeyres* price index formula. The Laspeyres is constructed as,

$$I_t^L \equiv \frac{\sum_{i=1}^N p_{i,t} q_{i,t-1}}{\sum_{j=1}^N p_{j,t-1} q_{j,t-1}} = \sum_{i=1}^N \omega_{i,t-1} \frac{p_{i,t}}{p_{i,t-1}}$$

where

$$\omega_{i,t-1} \equiv \frac{p_{i,t-1} q_{i,t-1}}{\sum_{j=1}^N p_{j,t-1} q_{j,t-1}}$$

is item i's share of total sales in the sample from the NAICS category. To the first order, by Roy's identity the arithmetic Laspeyres index captures the change in consumer welfare caused by small price shocks.[17]

Item Rotation. Establishments continue to report prices for a given item for five to seven years on average. After these five to seven years, the BLS selects a new sample for the 6-digit NAICS industry. Like item rotation in the CPI, the new sample attempts to better reflect the structure of a particular industry in terms of establishments and products over time.

Forced Substitutions. When an item is no longer produced, or future production has incorporated a change in the product's characteristics, the BLS must initiate a substitution. If the updated product is a sufficiently close substitute for the one it replaces, then the two product versions' prices are compared directly. However, when a close substitute is not available, the BLS and the establishment choose a substitute product that possesses as similar product characteristics as possible. The BLS then implements a quality adjustment to eliminate differences in prices across products

[17] Moreover, the Laspeyres index can be re-expressed as an exact CES price index (see Section B.A) with an assumed elasticity of substitution of zero, which is derived from a Leontief aggregate production technology

that are due to changes in underlying product characteristics. When a comparison between the new and old products is not feasible, or the respondent does not provide a price record in a given month, the BLS imputes the change in price, usually as the average price change across all products for which reliable information is available. The quality adjustment can be represented in the index by introducing notation for product quality, as in Section B.A.

Data Processing. In order to reduce the sensitivity of price indexes to exceptionally large price changes as well as avoid respondent disclosure (as per BLS disclosure avoidance policy) we exclude positive or negative price changes greater than 500%. These outliers occur rarely in the sample. Summary Statistics are reported in Appendix Table A12.

B.F Estimated Markups

In this appendix, we describe how we use Compustat to estimate markups and compute profitability ratios.

Estimated Markup. To estimate markup in Compustat, we follow De Loecker et al. (2020). They derive expressions for markups based on observables by exploiting cost minimization of a variable input of production. A key assumption underlying this approach is that the producer is a "price taker" for variable inputs. The main advantages of this approach are twofold: (a) theoretically, it does not require a specific model of how firms compete or specific assumptions about the demand system; (b) empirically, firms' financial statements are sufficient to implement this approach, there is no need for separate information on prices and quantities.

In this framework, the gross markup is defined as the difference between the consumer price P_{it}^{Q} and the shadow value of one more unit of production to the firm, λ_{it}, which is itself defined by the firm's cost-minimization problem. For firm i at time t, output comes from a production function using variable inputs (labor, intermediates, materials, etc.), capital and a fixed cost. Intuitively, the gross markup is the wedge between the willingness to pay of consumers for one more unit of output (P_{it}^{Q}) and the "reservation price" of the supplier to produce one more unit (λ_{it}).

The reservation price, λ_{it}, can be solved for in terms of observables by solving the cost-minimization problem of the producer,

$$\Lambda = \min_{V_{it},K_{it}} P_{it}^{V} V_{it} + r_{it} K_{it} + F_{it} - \lambda_{it} \left(Q(V_{it}, K_{it}) - \overline{Q_{it}} \right),$$

where the first-order condition with respect to variable costs yields an expression for λ_{it},

$$\lambda_{it} = \frac{P_{it}^{V} V_{it}}{Q_{it}} \theta_{it}^{v}, \tag{A1}$$

where the output elasticity of production with respect to variable inputs is,

$$\theta_{it}^{v} \equiv \frac{\partial Q(V_{it}, K_{it})}{\partial V_{it}} \frac{V_{it}}{Q_{it}}.$$

De Loecker et al. (2020) estimate the output elasticity of production with respect to variable inputs, θ_{it}^{v}, using various production functions. Defining the gross markups as $\mu_{it} \equiv \frac{P_{it}^{Q}}{\lambda_{it}}$ and substituting into equation (A1) yields

$$\mu_{it} = \theta_{it}^{v} \underbrace{\frac{P_{it}^{Q} Q_{it}}{P_{it}^{V} V_{it}}}_{\equiv \frac{SALES_{it}}{COGS_{it}}}. \tag{A2}$$

For intuition, it is instructive to consider the following heuristic derivation of the markup. Conceptually, the markup can be viewed as the wedge between the reservation price the consumer would be willing to pay to have the producer use one more unit of variable inputs ($P_{it}^{Q} \times \frac{\partial Q(.)}{\partial V_{it}}$), and the reservation price of the producer for doing so (P_{it}^{V}). The wedge is therefore given by $\mu_{it} = \frac{P_{it}^{Q}}{P_{it}^{V}} \frac{\partial Q(.)}{\partial V_{it}} = \frac{P_{it}^{Q} Q_{it}}{P_{it}^{V} V_{it}} \cdot \theta_{it}^{v}$, as in (A2). A positive wedge between reservation prices constitutes an inefficiency. In recent work, De Ridder et al. (2022) provide empirical support for the accuracy of the methodology of De Loecker et al. (2020).

We follow the data construction steps of De Loecker et al. (2020) in the Compustat North America – Fundamentals Annual data set (obtained through WRDS). The steps we take are identical to De Loecker et al. (2020), except that we assign a firm to a unique 6-digit NAICS industry (instead of the focus on more aggregated 2-digit industries in their paper). We eliminate firms with reported cost-of-goods to sales, and SG&A to sales ratio's in the top and bottom 1 percent, where the percentiles are computed for each year separately.

Following equation (A2), two key variables are used two compute markups: total sales (variable "SALE" in Compustat) and the total cost of goods sold (variable "COGS" in Compustat). Furthermore, for our baseline specification we use a time-invariant and sector-invariant elasticity $\theta_{it}^{v} = 0.85$, because De Loecker et al. (2020) show that their results are driven by changes in the ratio of sales to cost of goods sold, and remain very similar whether or not the elasticity is allowed to vary over time and across sectors (see their Appendix B1). Panel A of Online Appendix Figure A1 reports trends in markups over time: the average markup is increasing, but this increase is not uniform across the markup distribution and is driven by the top of the distribution.

Profitability Ratios. For robustness analysis, we compute two profitability ratios, profits over sales and profits over assets. To calculate profits, we use the markup measure in equation (A2) and account for all costs, including fixed costs:

$$\Pi_i = S_{it} - P_t^{V} V_t - r_t K_{it} - P_t^{X} X_t.$$

With Π_i denoting profits, the profit rate as a share of sales is $\pi_{it} = \frac{\Pi_i}{SALE_{it}}$. This measure scales the profits by firm size as measured by its revenue. From an investment viewpoint, we may want to measure the return on assets. Therefore we also compute an analogous measure, now dividing by total capital instead of total sales: $\pi_{it} = \frac{\Pi_i}{K_{it}}$.

Panel B of Online Appendix Figure A1 reports the trends in profitability over time. The trends paint a picture similar to the markup trends: profitability has been increasing over time, and particularly so for the most profitable firms.

C Additional Empirical Results

C.A Robustness of Main Estimates

In this section, we describe several robustness tests.

Falsification tests. Appendix Figure A2 reports additional pre-trend tests, suggesting that it is important to at least include fixed effects for apparel and durable to avoid pre-trends. Panel A repeats the estimation of equation (2), but without fixed effects for apparel or durable goods. The figures exhibits pre-trends: ELI categories with a higher NTR gap had lower inflation even prior to 2000. Panel B of Figure A2 shows that once fixed effects for apparel and durable goods are included, the estimated coefficients do not exhibit pre-trends. These result indicate that including fixed effects for apparel and durable goods is important to ensure that a causal interpretation of the estimates is plausible. The validity of the research design would be doubtful if changing the set of controls in (2) had a large impact on pre-trends, but we find that the the results are stable with alternative sets of fixed effects, as long as there are controls for apparel and durables. Table A13 reports the same pre-trend tests, with fixed effects for apparel and durable goods, using using the extended CPI sample going back to 1977. Year-by-year estimates are reported in Figure A3 and show no signs of pre-trends. Figure A4 document similar patterns using publicly-available data from the NBER-CES database, measuring inflation as the change in the NBER-CES price index for the value of shipments.

As an additional falsification test, Table A14 shows that the NTR gap does not predict import penetration from China in France.

Sensitivity and heterogeneity analysis with the NTR gap instrument. We conduct several sensitivity test. Table A3 shows the IV estimates from a specification with more periods, rather than two long periods as in the main text. The period-specific IV estimates are similar to the specification with long differences reported in the main text. Table A15 shows that the NTR gap does not predict changes in import penetration in the U.S. prior to 2000. Table A4 shows that the IV estimates are similar whether or not the specification uses consumption weights, and whether or not the sample is restricted to goods only. This table also shows that the IV estimates remain similar with additional specifications with apparel and durable by period FE, and 2-digit ELI FEs interacted with period dummies. Table A16 is similar to Table 3 in the main text, but follows Pierce and Schott (2016) by including additional controls in Column (2): exposure to MFA quota reductions, Chinese import tariffs from Brandt and Morrow (2017), data on export licensing requirements from Bai et al. (2017), and data on production subsidies from China's National

Bureau of Statistics. The IV estimate remains similar. Furthermore, Online Appendix Table A17 shows that the results are similar when controlling for exports. This finding addresses the potential concern that import penetration from China in the domestic market may mis-measure changes in import competition, which also occurs in foreign markets for exporting firms.

Next, we examine heterogeneity across industries. We investigate the mechanism behind the relevance of the NTR gap instrument. With non-convex adjustment costs, a fall in uncertainty should boost capital investment (e.g., Dixit and Pindyck (1994)), which should especially matter for capital-intensive industries. Consistent with this idea, Online Appendix Table A18 shows that the first-stage relationship is stronger for capital-intensive industries with the NTR gap instrument, but not with the change in import penetration from China in other countries. Finally, using the estimated elasticities from Broda and Weinstein (2006), Online Appendix Table A19 reports that the IV coefficient is stable across subsamples with different trade elasticities.

Sensitivity analysis with the ADH instrument. Online Appendix Table A20 shows that the results are similar when the instrument is the change in import penetration in other developed economies.

Alternative measures of import penetration. We adjust our measure of import penetration from China to account for distribution margins. Intuitively, consider a product category like apparel. China substantially increases its market share in the production of apparel. But a substantial share of the retail price of apparel results from retail and transportation costs, implying that China's "market share" increases much less in the consumer basket. China's market share only increases at the production stage, while, by definition, retail and transportation costs continue to be incurred domestically.

We therefore use the IO sample and adjust the denominator in equation (1) for distribution margins, which are given by the ratio of total output in purchaser prices to total output in producer prices. The results are reported in Online Appendix Table A21. As expected, the IV coefficient becomes more negative when purchaser prices are used (-4.37, s.e. 0.852) rather than producer prices (-2.44, s.e. 0.431), because the effective change in import penetration from China is smaller with purchaser prices. Although it is instructive to note that the price effects become even stronger with the adjustment for distribution margins, for comparability with prior work we focus on the IV estimates with the baseline import penetration measure.

Additional results for specifications with continued products. Online Appendix Table A22 focuses on the subset of goods that existed prior to the "China shock" (specifically, they were available as of 2000). We still find a large response of continued products inflation, which shows that pre-existing varieties are affected by Chinese import competition. This result shows that the patterns of lower continued products inflation documented in the main text are not due to goods that were introduced after the China shock, implying that "reallocation effects" do not drive the observed price response. Furthermore, Online Appendix Table A23 documents that trade with China led to increased product turnover, consistent with the notion that Chinese products displace domestic varieties.

C.B The Role of Imported Intermediate Inputs

This appendix presents additional results about the role of imported intermediate inputs.

To examine whether buyer-supplier linkages affect our results, we first compute the correlations between our baseline measure of import competition and indirect exposure via domestic suppliers or domestic buyers. We conduct this analysis with our Input-Output sample, using the BEA's IO table with the standard proportionality assumptions. To measure industry j's exposure via domestic suppliers, denoted "ΔChina IP Supplier", we compute the change in the share of spending on intermediate inputs from China in industry j's total sales. For exposure via domestic buyers, denoted "ΔChina IP Buyer", we compute the change in the import penetration rate from China in industry j's domestic buyer industries, multiplied by the share of domestic buyer industries in industry j's total sales. By definition, both "ΔChina IP Supplier" and "ΔChina IP Buyer" are low if an industry has a high share of value added or sells primarily to final consumers.

Table A8 reports the correlations between direct and indirect exposure to trade with China. Column (1) shows the raw relationship without any controls: the coefficient is positive and significant, but small in magnitude. When the import penetration rate from China increases by 1 percentage point in industry j, the share of intermediate inputs from China in industry j's total output increases by only ten basis points. The relationship decreases further when we introduce the same set of controls as in our baseline IV specification in Column (2), and when we exclude intra-industry buyer-seller relationships (the diagonal component of the IO table) in Column (3).

For buyer effects, Columns (4) shows that the raw relationship is also positive but even smaller: a 1 percentage point increase in import competition from China in industry j is associated with a further 2 basis point increase in import competition via domestic buyer industries. Columns (5) and (6) report that the relationship becomes a precisely estimated zero with the other specifications. These results indicate that direct import competition is not correlated with indirect effects, therefore the price effects we document are unlikely to be explained by these channels.

Table A24 shows that the results are similar to the estimates reported in the main text when accounting for higher-order I-O linkages. Table A25 reports similar results in an augmented IV framework, where we instrument for direct and indirect exposure measures simultaneously.

C.C The Role of Markups

In this section, we discuss additional results on the role of markups.

First, we study changes in the distribution of estimated markups. Table A26 document changes in the distribution of markups across industries with heterogeneous exposure to the NTR gap. In Column (1), the reduced form coefficient for the 90th percentile of markups is -17.42 (s.e. 7.28); in Column (2) the effect gets attenuated by a factor of over 50%, with a coefficient of -7.97 (s.e. 4.83) for the 50th percentile; in Column (3) the coefficient for the 10th percentile becomes insignificant and is close to zero (-0.84, s.e. 4.023). Consistent with the predictions of the model, the response of markups is much stronger at the top of the markup distribution.

Second, we analyze the response of firm profitability to increased import penetration from China. We compute the ratios of total profits to total sales and to total assets, where profits are computed inclusive of fixed costs incurred by the firm (in contrast, the markup measure from equation (A2) does not use information about fixed costs). Columns (4) and (5) of Table A26 report the corresponding point estimates, which are statistically significant at the 1% level.

Finally, we measure heterogeneity in our IV coefficient using CPI data (in complement to the PPI data analyzed in the main text). To measure domestic market concentration, we obtain data on Herfindahl indices by 6-digit NAICS industries from the Census for 1997 (as in Grullon et al. (2018)), which we link to our CPI sample.

Table A27 presents the results. Consistent with the predictions, Column (1) shows that the price response to a one percentage point increase in the import penetration rate from China is much larger when the domestic market is more concentrated, and is much smaller when the initial China share is small. The price decline is 1.29 percentage points larger in the set of more concentrated industries, and it is attenuated by 1.50 percentage points in the set of industries initially more exposed to trade with China.

Columns (2) to (5) show the robustness of these results by repeating the IV specifications in subsamples. The IV coefficient is large in the sub-sample of product categories above median concentration (Column (2)), while it becomes insignificant for those below (Column (3)). The point estimate in Column (3) is close to the prediction from the class of models without strategic interactions characterized by Arkolakis et al. (2012). Columns (4) and (5) show that the estimated effect is over twice as large for industries that were initially less exposed to trade with China.

D Theory Appendix

D.A Connecting the IV Specification to the Melitz-Chaney Model

In this section, we first set up the model in section D.A.1, deriving the key sector-level relationships that hold in the model. We then use these relationships in the following sections, expressing the reduced-form regression coefficient from specification (3) in terms of structural parameters of the model.[18]

D.A.1 Setting

We use the setting of Melitz (2003) and Chaney (2008).

Consumer's Problem. There are n countries. A representative consumer in country j has Cobb-Douglas preferences over K product categories,

$$U_j = \prod_{k=1}^{K} \left(Y_j^k\right)^{\mu_k} \quad \text{such that} \quad \sum_{k=1}^{K} \mu_k = 1$$

Each product category consists of differentiated items over which the consumer has CES preferences with an elasticity of substitution between items of $\sigma > 1$,

$$Y_j^k = \left(\sum_{i=1}^{n} \int_{\Omega_{ij}^k} y_{ij}^k(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}}$$

whereΩ_{ij}^k is the set of varieties from product category k available to the consumer in country j that are produced in country i, and $y_{ij}^k(\omega)$ is the quantity of each such variety $\omega \in \Omega_{ij}^k$. The associated aggregate price for product category k is,

$$P_j^k = \left(\sum_{i=1}^{n} \left(P_{ij}^k\right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} ,$$

where P_{ij}^k is the accompanying aggregate price of varieties from country i and is given by,

$$P_{ij}^k = \left(J_i^k \int_{z_{ij}^k}^{\infty} p_{ij}^k(z)^{1-\sigma} g_i^k(z) dz \right)^{\frac{1}{1-\sigma}} , \tag{D1}$$

and where J_i^k is the mass of firms from country i and, as discussed below, $g_i^k(z)$ is the density function over country i firms' idiosyncratic productivity. The consumer maximizes utility subject

[18] The purpose of this section is to describe which structural parameters of the model are recovered by our regression coefficient across product categories. We thus demonstrate that the reduced-form empirical specification can be precisely linked to standard trade model despite the "missing intercept" inherent in well-identified empirical work. To focus on this goal, in the derivations below we intentionally abstract from domestic shocks and potential endogeneity and expenditure shares, which is addressed by our IV strategy.

to their budget constraint. Given CES preferences over the differentiated goods, the demand function for variety ω is,

$$p_{ij}^k(\omega)y_{ij}^k(\omega) = \left(\frac{p_{ij}^k(\omega)}{P_j^k}\right)^{1-\sigma} X_j^k \tag{D2}$$

where $p_{ij}^k(\omega)$ is the price of each variety $\omega \in \Omega_{ij}^k$, and $X_j^k \equiv P_j^k Y_j^k$ is total expenditure on varieties in product category k. Total expenditures on varieties in product category k are the sum of domestically produced varieties and imported varieties from each country, $X_j^k = \sum_{i=1}^n X_{ij}^k$. Finally, S_{ij}^k is country j's expenditure share of imported goods in product category k from country i,

$$S_{ij}^k \equiv \frac{X_{ij}^k}{X_j^k} = \left(\frac{P_{ij}^k}{P_j^k}\right)^{1-\sigma} \tag{D3}$$

where the second equality holds due to the consumer demand function in equation (D2). Moreover, by Shephard's Lemma, the total change in country j's aggregate price in category k following the foreign supply shocks is,

$$d\log(P_j^k) = \sum_{i=1}^n S_{ij}^k \; d\log(P_{ij}^k) \ . \tag{D4}$$

Firms. Firms within a product category k are monopolistic competitors. Firms in country i have aggregate productivity A_i^k and idiosyncratic productivity $z \sim G_i^k(z)$. These firms produce a variety according to a linear production technology that takes labor as an input, $y_{ij}^k(z) = A_i^k z \ell_{ij}^k$, take the economy-wide wage w_i as given, face an iceberg cost, τ_{ij}^k, and pay a labor-denominated fixed cost $w_i f_{ij}^k$ to operate in the market, which they pay as long as their profits are positive. Firm profits are,

$$\pi_{ij}^k(z) = p_{ij}^k(z)y_{ij}^k(z) - \frac{\tau_{ij}^k w_i}{A_i^k z} y_{ij}^k(z) - w_i f_{ij}^k$$

so that, taking consumer demand in equation (D2) as given, the optimal price is,

$$p_{ij}^k(z) = \frac{\sigma}{\sigma - 1}\frac{\tau_{ij}^k w_i}{A_i^k z} \ . \tag{D5}$$

This implies an entry cutoff for productivity defined by $\pi_{ij}^k(z_{ij}^k) = 0$ such that,

$$z_{ij}^k = \frac{\sigma}{\sigma - 1}\frac{\tau_{ij}^k w_i}{A_i^k P_j^k}\left(\frac{\sigma w_i f_{ij}^k}{X_j^k}\right)^{\frac{1}{\sigma-1}} \ . \tag{D6}$$

Trade Elasticities. Expenditures on items from product category k produced in country i and sold in country j are obtained by aggregating individual firms' sales,

$$X_{ij}^k = J_i^k \int_{z_{ij}^k}^{\infty} p_{ij}^k(z)y_{ij}^k(z)g_i^k(z)dz$$

where J_i^k is the mass of firms from country i and $g_i^k(z)$ is the density function over country i firms' idiosyncratic productivity. Using equation (D5), the partial elasticity of relative trade flows from country i to j with respect to an increase in A_i^k, the aggregate productivity of country i firms,[19] which we refer to as *the trade elasticity*, is

$$\frac{\partial \log(X_{ij}^k/X_{jj}^k)}{\partial \log(A_i^k)} = (\sigma-1) + \frac{g_i^k(z_{ij}^k)z_{ij}^k}{1-G_i^k(z_{ij}^k)} \cdot \frac{(z_{ij}^k)^{\sigma-1}}{\int_{z_{ij}^k}^{\infty} z^{\sigma-1}\frac{g_i^k(z)}{1-G_i^k(z_{ij}^k)}dz}$$

where the first term is the *intensive margin* change in sales and the second term is the *extensive margin* change in sales. The intensive margin indicates that firms that were already producing now make more revenue, while the extensive margin measures the additional revenue from new firms that are now producing.

If firm heterogeneity is characterized by a Pareto distribution, $G_i^k(z) = 1-(a_i^k)^\theta z^{-\theta}$ such that $\theta > \sigma - 1$, then the trade elasticity is determined solely by the Pareto tail parameter,

$$\frac{\partial \log(X_{ij}^k/X_{jj}^k)}{\partial \log(A_i^k)} = (\sigma-1) + \frac{\theta(a_i^k)^\theta(z_{ij}^k)^{-\theta-1}}{(a_i^k)^\theta(z_{ij}^k)^{-\theta}} \cdot \frac{(z_{ij}^k)^{\sigma-1}}{\int_{z_{ij}^k}^{\infty} z^{\sigma-1}\frac{\theta(a_i^k)^\theta z^{-\theta-1}}{(a_i^k)^\theta(z_{ij}^k)^{-\theta}}dz} z_{ij}^k \;=\; \theta\ ,$$

as in Chaney (2008).

D.A.2 Proof of Proposition 1 with homogeneous trade elasticities across sectors

We now allow price indices to change in all countries because of wage changes and and entry-exit of firms. The endogenous change in the domestic wage is denoted $d\log(w_j)$.

To characterize the response of the theoretical price index to foreign supply shocks, $\{d\log(A_i^k)\}_{i\neq j}$, we first consider a setting where firm heterogeneity is characterized for all sectors by the same Pareto distribution, $G_i^k(z) = 1-(a_i^k)^\theta z^{-\theta}$, with $\theta > \sigma - 1$. From equation (D1) we have

$$d\log(P_{ij}^k) = \frac{1}{1-\sigma} d\log\left(\int_{z_{ij}^k}^{\infty} p_{ij}^k(z)^{1-\sigma} g_i^k(z)dz\right).$$

Evaluating the integral with the Pareto productivity distribution $g_i^k(z)$, and substituting the expression for each variety's price from equation (D5) yields:

$$d\log(P_{ij}^k) = d\log(w_i) - d\log(A_i^k) + \frac{\theta-(\sigma-1)}{\sigma-1} d\log(z_{ij}^k). \tag{D7}$$

Using the expression for the entry cutoff in equation (D6), we have:

$$d\log(z_{ij}^k) = \frac{\sigma}{\sigma-1} d\log(w_i) - d\log(A_i^k) - d\log(P_j^k) - \frac{1}{\sigma-1} d\log(X_j^k).$$

[19] We can analogously posit a decrease in iceberg trade costs for country i, τ_{ij}, and obtained the same trade elasticity below. The necessary condition is that all country i firms face an equal change in marginal costs under the candidate perturbation.

For the domestic entry cutoff, this expression simplifies to:

$$d\log(z_{jj}^k) = d\log(w_j) - d\log(P_j^k), \tag{D8}$$

where we use the fact that there is no domestic supply shock by assumption, i.e. $d\log(A_j^k) = 0$, and that $d\log(X_j^k) = d\log(w_j)$, since by Cobb-Douglas preferences across sectors we have $X_j^k = \mu_k (w_j L_j)$, with inelastic labor supply. Substituting (D8) into (D7), using equation (D3) and rearranging yields:

$$d\log(S_{jj}^k) = -\theta \left(d\log(w_j) - d\log(P_j^k) \right). \tag{D9}$$

Using this expression, we can characterize the cross-sector relationship between changes in the domestic expenditure share and changes in the price index for domestic consumers, which is emerges from the underlying unobserved foreign supply shocks $\{d\log(A_i^k)\}_{i\neq j}$, as $d\log(P_j^k) = d\log(w_j) + (1/\theta) d\log(S_{jj}^k)$. With the reduced-form specification $d\log(P_j^k) = \alpha + \beta d\log(S_{jj}^k) + \varepsilon_j^k$, the OLS estimate of β is given by:

$$\begin{aligned}
\hat{\beta} &\equiv \frac{\text{Cov}\left(d\log(P_j^k), d\log(S_{jj}^k)\right)}{\text{Var}\left(d\log(S_{jj}^k)\right)} \\
&= \frac{-\theta\left(\text{Cov}\left(d\log(P_j^k), d\log(w_j)\right) - \text{Var}\left(d\log(P_j^k)\right)\right)}{\theta^2\left(\text{Var}\left(d\log(w_j)\right) + \text{Var}\left(d\log(P_j^k)\right)\right)} \\
&= \frac{1}{\theta},
\end{aligned}$$

where we have used the fact that the log change in the domestic wage, $d\log(w_j)$, is identical across sectors in this model with a single labor type. Intuitively, since the the change in the wage is common across sectors, it is differenced out in the regression and does not enter the estimated regression coefficient.

D.A.3 Proof of Proposition 1 with heterogeneous trade elasticities across sectors

To complete the proof of Proposition 1 under heterogeneous trade elasticities, we first present a general derivation for the OLS estimand under heterogeneous treatment effects, which we then apply to our setting.

OLS estimand under heterogeneous treatment effects. To derive the OLS estimand under heterogeneous treatment effect, we follow Yitzhaki (1996) and Angrist and Krueger (1999), expressing the OLS regression coefficient as a weighted average of state-specific causal responses.

We use the following notation: an outcome, y, depends on an initial state, s, and its exposure to a treatment, u, such that the new state after treatment is $s' = s + u$. The initial state is distributed according to $s \sim F_s(s)$ on the domain $\mathcal{S}$ and the treatment is distributed according to the conditional CDF $u \sim F(u|s)$ which we assume satisfies the mean independence property that $\mathbb{E}[u|s] = \int u \; dF(u|s) = \mu \in \mathbb{R}$ for all $s \in \mathcal{S}$. Heterogeneous treatment effects are characterized by the structural relationship

$$y = g(s,u) \equiv \alpha + \delta(s)u \ ,$$

where the marginal effect of the treatment, $\delta(s)$, is heterogeneous and is a function of the initial state.

While this relationship holds in each state s, with state-specific coefficients $\delta(s)$, we run the reduced-form specification $y = \alpha + \beta u + \epsilon$ with a constant coefficient β across states. The OLS regression coefficient takes the form $\hat{\beta} = \frac{\text{Cov}(u,y)}{\text{Var}(u)}$, so by the law of total covariance we have:

$$\hat{\beta} = \frac{\mathbb{E}[\text{Cov}(u, g(s,u)|s)] + \text{Cov}(\mathbb{E}[u|s], \mathbb{E}[g(s,u)|s])]}{\mathbb{E}[\text{Var}(u|s)] + \text{Var}(\mathbb{E}[u|s])}.$$

Using the hypothesized mean independence property, $\mathbb{E}[u|s] = \mu$ for all s, we get $\text{Cov}(\mathbb{E}[u|s], \mathbb{E}[g(s,u)|s])] = 0$ and $\text{Var}(\mathbb{E}[u|s]) = 0$.

Moreover, through integration by parts, the conditional covariance term can be expressed as,

$$\begin{aligned}
\text{Cov}(u, g(s,u)|s) &= \int_{-\infty}^{\infty} g(s,u)\,(u - \mathbb{E}[u|s])\, dF(u|s) \\
&= \int_{-\infty}^{\infty} \frac{\partial g(s,u)}{\partial u} \left(\int_{u}^{\infty} (\tilde{u} - \mu) dF(\tilde{u}|s) \right) du \\
&= \delta(s) \int_{-\infty}^{\infty} \int_{u}^{\infty} (\tilde{u} - \mu) dF(\tilde{u}|s) du.
\end{aligned}$$

By a similar derivation, the conditional variance can be expressed as,

$$\text{Var}(u|s) = \int_{-\infty}^{\infty} \int_{u}^{\infty} (\tilde{u} - \mu) dF(\tilde{u}|s) du.$$

Thus, the estimated coefficient is the weighted average of state-specific treatment effects,

$$\hat{\beta} = \int_{\mathcal{S}} \omega(s)\delta(s) ds \tag{D10}$$

where $\omega(s)$ are weights over initial states such that,

$$\omega(s) \equiv \frac{\text{Var}(u|s) f_s(s)}{\int_{\mathcal{S}} \text{Var}(u|\tilde{s}) f_s(\tilde{s}) d\tilde{s}}.$$

The weights can be understood by expressing the conditional variance of u as follows,

$$\begin{aligned}
\text{Var}(u|s) &= \int_{-\infty}^{\infty} \Big(\mathbb{E}[\tilde{u}|s, \tilde{u} \geq u] - \Big(\mathbb{E}[\tilde{u}|s, \tilde{u} < u] F(u|s) + \mathbb{E}[\tilde{u}|s, \tilde{u} \geq u](1 - F(u|s)) \Big) \Big) (1 - F(u|s)) du \\
&= \int_{-\infty}^{\infty} \Big(\mathbb{E}[\tilde{u}|s, \tilde{u} \geq u] - \mathbb{E}[\tilde{u}|s, \tilde{u} < u] \Big) \text{Prob }(\tilde{u} < u|s) \text{Prob }(\tilde{u} \geq u|s) \; du
\end{aligned}$$

The expression shows that more weight is given to the sectors that are close to the median of the treatment u conditional on an initial state s, since $Prob(\tilde{u} < u|s)Prob(\tilde{u} \geq u|s)$ is maximized at $Prob(\tilde{u} \geq u|s) = 1/2$. The weight is also larger when there is a larger asymmetry in the distribution of the treatment observed above and below that of the candidate treatment u, leading to a larger spread between conditional expectations $\mathbb{E}[\tilde{u}|s \geq u] - \mathbb{E}[\tilde{u}|s < u]$.

In our setting, the derivations above apply to our cross-sector regressions with outcome $y = d\log(P_j^k)$, treatment $u = d\log(S_{jj}^k)$, and initial state $s = S_{jj}^k$. The mean independence property $\mathbb{E}[u|s] = \mu$ is supported empirically by the fact that the log change in the domestic expenditure share is not correlated with the initial domestic expenditure share, as shown in Appendix Figure A5.

Application to heterogeneous trade elasticities across sectors. With heterogeneous Pareto shape parameters θ^k across sectors, following the same steps as in Subsection D.A.2 we find that in each sector the following relationship holds:

$$d\log(P_j^k) = d\log(w_j) + \frac{1}{\theta^k} d\log(S_{jj}^k).$$

While this relationship holds in each sector s, we run the reduced-form specification $d\log(P_j^k) = \alpha + \beta d\log(S_{jj}^k) + \varepsilon_j^k$, with a constant coefficient β for all sectors. To apply the above results about the OLS estimand under heterogeneous treatment effects, we denote by $S \equiv d\log(S)$ the random variable following the distribution of $d\log(S_{jj}^k)$ across sectors. We then apply the estimator in equation (D10) of Section D.A.3 to obtain the estimated regression coefficient:

$$\hat{\beta} = \sum_k \omega_k \frac{1}{\theta^k},$$

where the weights satisfy $\sum_k \omega_k = 1$, and are given by $\omega_k \equiv \nu_k / \sum_\ell \nu_\ell$ such that:[20]

$$\nu_k \equiv \left(\mathbb{E}[\tilde{S}|S \geq d\log(S_{jj}^k)] - \mathbb{E}[\tilde{S}|S < d\log(S_{jj}^k)]\right) \left(1 - P(\tilde{S} \geq d\log(S_{jj}^k))\right) \left(P(\tilde{S} \geq d\log(S_{jj}^k))\right).$$

Thus, the estimated regression coefficient is a weighted average of the sector-specific inverse trade elasticities.[21]

[20] Given continuous valued treatments that are sampled in a finite number of sectors, each ν_k is associated with one initial state S_{jj}^k and one treatment $d\log(S_{jj}^k)$. Therefore we have suppressed the conditioning on the initial state in the expression. Furthermore, the density function over initial states cancels out in this discrete sampling case, since the initial states are uniformly distributed with $f_s(S_{jj}^k)$= 1/N.

[21] More weight is given to the sectors that are close to the median of $d\log(S_{jj}^k)$, since $\left(1 - P(S \geq d\log(S_{jj}^k))\right)\left(P(S \geq d\log(S_{jj}^k))\right)$ is maximized at $P(S \geq d\log(S_{jj}^k)) = 1/2$. The weight is also larger when there is a larger asymmetry in the change in trade flows observed above and below that of sector k, $d\log(S_{jj}^k)$.

D.A.4 Connecting the IV Specification to Arkolakis et al. (2012)

Arkolakis et al. (2012) prove how models that satisfy a set of common assumptions yield the same welfare implications, despite possessing different microfoundations. Here we show that the cross-sector relationship between changes in prices and changes in expenditure shares are also common across such environments. In particular, we show that,

$$d\log(P_j^k) = \frac{1}{\epsilon} d\log(S_{jj}^k) + d\log(w_j),$$

where ϵ is the partial trade elasticity with respect to variable costs, whose characterization may vary across economic environments. Thus, applying our estimator would yield an estimate of the price effect of $1/\epsilon$.[22]

We characterize below the cross-sector relationships between price change and changes in domestic share for some of the main models within the Arkolakis et al. (2012) class. As before, consider a shock from country i, which we denote more generally as an iceberg cost relative to productivity $d\log(\tau_{ij}^k/A_i^k)$. Since there is no such shock for the domestic economy, we have $d\log(\tau_{jj}^k/A_j^k) = 0$.

(i) In the *Armington model*, country j's expenditure shares on items from country i and associated price are,

$$\begin{aligned} d\log(S_{ij}^k) &= (1-\sigma)\Big(d\log(P_{ij}^k) - d\log(P_j^k)\Big) \\ d\log(P_{ij}^k) &= d\log(w_i) + d\log(\tau_{ij}^k/A_i^k) \end{aligned}$$

and combining the domestic share and price responses yields a version of our empirical specification,

$$d\log(P_j^k) = \frac{1}{\sigma-1} d\log(S_{jj}^k) + d\log(w_j)$$

where $\epsilon \equiv \sigma - 1$ is the trade elasticity.

(ii) In the *Melitz-Chaney model* with Pareto distributed firm heterogeneity governed by tail parameter $\theta > \sigma - 1$, import shares and prices from country i are,

$$\begin{aligned} d\log(S_{ij}^k) &= (1-\sigma)\Big(d\log(P_{ij}^k) - d\log(P_j^k)\Big) \\ d\log(P_{ij}^k) &= \frac{\theta}{\sigma-1}\Big(d\log(w_i) + d\log(\tau_{ij}^k/A_i^k)\Big) + \left(1 - \frac{\theta}{\sigma-1}\right)\left(d\log(P_j^k) - \frac{1}{\sigma-1} d\log(w_i/w_j)\right) \end{aligned}$$

Again, combining the domestic share and price responses yields,

$$d\log(P_j^k) = \frac{1}{\theta} d\log(S_{jj}^k) + d\log(w_j)$$

[22] If the microfoundations for the the trade elasticity include heterogeneity across sectors, denoted ϵ_k, then our estimator would identify the weighted cross-sector average of $1/\epsilon_k$.

with trade elasticity $\epsilon \equiv \theta$.

(iii) In the Eaton and Kortum (2002) model, perfect competition combined with Fréchet distributed firm heterogeneity (with tail parameter $\theta > \sigma - 1$) induces a distribution over least-cost competitors that results in a CES import demand system.[23] In this model, a firm in country i exporting to country j in sector k has a draw of idiosyncratic productivity z and then charges a price $P_{ij}^k(z) = w_i\tau_{ij}^k/(A_i^k z)$ if it is the least-cost provider of the item, and does not produce otherwise. Integrating over country i's least-cost producers exporting to sector k variety in country j and then taking the log differences yields,

$$d\log(P_{ij}^k) = d\log(w_i) + d\log(\tau_{ij}^k/A_i^k)$$

with aggregated price index of $P_j^k = \gamma(\sum_{i=1}^N (P_{ij}^k)^{-\theta})^{-1/\theta}$ and where γ is a parameter resulting from integration over least-cost producers. Thus, the change in expenditure shares on items from country i in country j and sector k take the form,

$$d\log(S_{ij}^k) = -\theta\Big(d\log(P_{ij}^k) - d\log(P_j^k)\Big)$$

and combining the domestic shares and prices responses yields,

$$d\log(P_j^k) = \frac{1}{\theta} d\log(S_{jj}^k) + d\log(w_j)$$

with trade elasticity $\epsilon \equiv \theta$.

It is instructive to note that the relationship between prices and expenditure shares is tied to the gravity equation. In Head and Mayer (2014), a model satisfies a gravity equation if its bilateral trade flows from country i to j are decomposable into an importer-specific component, an exporter-specific component, and a component specific to the bilateral trade flow between countries i and j,

$$\log(X_{ij}^k) = \log(\alpha_i^k) + \log(\delta_j^k) + \log(\phi_{ij}^k)$$

where α_i is the contribution of exporter i to trade flows across all destinations, δ_j are importer characteristics that induce inflows from all exporters, and ϕ_{ij} captures the origin-destination specific component of trade flows between countries i and j.

As shown in Arkolakis et al. (2012) and in Table 1 of Head and Mayer (2014), the set of models in the class defined by ACR satisfy the gravity equation.[24] Specifically, the trade flows X_{ij}^k for the

[23] Bernard et al. (2003) has the same estimating equation despite introducing imperfect competition and non-constant markups. The model's assumptions on productivity give it the same structure as Eaton and Kortum (2002) in our application, the only practical difference concerns price aggregation that yields a different form for the parameter γ.

[24] Arkolakis et al. (2012) points to an exception when the Melitz model has an additional assumption on the allocation of the domestic labor supply, but notes that the partial trade elasticity is still recoverable from such a setup.

set of models illustrated in Arkolakis et al. (2012) can be arranged to satisfy the gravity equation as follows:

(i) In the Armington model we have,

$$\log(X_{ij}^k) = \log\left(w_i^{1-\sigma}\right) + \log\left((P_j^k)^{\sigma-1}X_j^k\right) + \log\left((\tau_{ij}^k/A_i^k)^{1-\sigma}\right)$$

(ii) In the Melitz-Chaney model, given the productivity threshold in equation (D6) the bilateral trade flow is,

$$\begin{aligned}\log(X_{ij}^k) &= constant + \log\left((w_i)^{1-\sigma}a_i^\theta\right) + \log\left((P_j^k)^{\sigma-1}X_j^k\right) + \log\left((\tau_{ij}^k/A_i^k)^{1-\sigma}(z_{ij}^k)^{-\theta+(\sigma-1)}\right)\\ &= constant + \log\left(w_i^{-\theta}a_i^\theta w_i^{-\frac{\theta}{\sigma-1}+1}\right) + \log\left((P_j^k)^\theta(X_j^k)^{\frac{\theta}{\sigma-1}}\right) + \log\left((f_{ij}^k)^{-\frac{\theta}{\sigma-1}+1}(\tau_{ij}^k/A_i^k)^{-\theta}\right)\end{aligned}$$

(iii) In the Eaton and Kortum (2002) model, the bilateral trade flow is,

$$\log(X_{ij}^k) = \log(w_i^{-\theta}) + \log\left((P_j^k)^\theta X_j^k\right) + \log\left((\tau_{ij}^k/A_i^k)^{-\theta}\right)$$

For relating the gravity equation to our empirical specification, a crucial feature of each model considered above is multiplicative separability of the form,

$$\delta_j^k = (P_j^k)^\epsilon \cdot \tilde{\delta}_j^k,$$

where ϵ is the partial trade elasticity. This form of separability implies that trade flows can be rewritten as,

$$\log(X_{ij}^k) = \log(\alpha_i^k) + \left(\epsilon\log(P_j^k) + \log(\tilde{\delta}_j^k)\right) + \log(\phi_{ij}^k)$$

and rearranging yields (noting that $S_{ij}^k = X_{ij}^k/X_j^k$),

$$\log(P_j^k) = \frac{1}{\epsilon}\log(S_{ij}^k) - \frac{1}{\epsilon}\Big(\log(\alpha_i^k) + \log(\tilde{\delta}_j^k/X_j^k)\Big) - \frac{1}{\epsilon}\log(\phi_{ij}^k)$$

Considering the domestic trade flow and differencing yields (where $\phi_{jj}^k = 1$),

$$d\log(P_j^k) = \frac{1}{\epsilon}d\log(S_{jj}^k) - \frac{1}{\epsilon}\Big(d\log(\alpha_j^k) + d\log(\tilde{\delta}_j^k/X_j^k)\Big)$$

and given that $d\log(X_j^k) = d\log(w_j)$, it is straightforward to show that in each of these models,

$$d\log(\alpha_j^k) + d\log(\tilde{\delta}_j^k/X_j^k) = -\epsilon\ d\log(w_j).$$

Thus, the connection between gravity and our empirical specification follows,

$$d\log(P_j^k) = \frac{1}{\epsilon}d\log(S_{jj}^k) + d\log(w_j).$$

D.A.5 Proof of Corollary 1.1

While Proposition 1 characterized the response of the exact CES price index to foreign supply shocks, Corollary 1.1 focuses on the measured Consumer Price Index.

Connecting the measured CPI to the Melitz-Chaney exact price index. The measured CPI can formally be linked to the Melitz-Chaney model using decompositions of price responses into intensive margin and extensive margin price changes. The change in the aggregate price of country i's imports by j is

$$d\log(P_{ij}^k) = \underbrace{\int_{z_{ij}^k}^{\infty} \left(s_{ij}^k(z) g_i^k(z)\right) d\log\left(p_{ij}^k(z)\right) dz}_{\equiv d\log(I_{ij}^k)} + \underbrace{s_{ij}^k(z_{ij}^k) g_i^k(z_{ij}^k) \frac{dz_{ij}^k}{\sigma - 1}}_{\equiv d\log(E_{ij}^k)} , \tag{D11}$$

where $s_{ij}^k(z)$ is country j's expenditure share of the given variety from country i and using $d\log(J_i^k) = 0$ for all i by free entry and labor market clearing. The first term is the *intensive margin* price change, $d\log(I_{ij}^k)$, which gives price changes due to incumbents firms from country i in product category k. The second term is the *extensive margin* price change, $d\log(E_{ij}^k)$, which measures the price change in product category k due to entry-exit from country i firms.[25]

The intensive margin price change, $d\log(I_{ij}^k)$, corresponds to the BLS' measured price index for a given country i in a product category k. Indeed, under the BLS methodology the Cost of Living Index (COLI) is constructed as a general first-order approximation to a matched-model index.[26] Thus, the overall change in P_j^k can be decomposed into an intensive margin price change corresponding to the measured CPI, $d\log(\widetilde{P}_j^k)$,[27] and an unobserved extensive margin price change, $d\log(E_j^k)$, as,

$$d\log(P_j^k) = \underbrace{\sum_{i=1}^{n} S_{ij}^k \ d\log(I_{ij}^k)}_{\equiv d\log(\widetilde{P}_j^k)} + \underbrace{\sum_{i=1}^{n} S_{ij}^k \ d\log(E_{ij}^k)}_{\equiv d\log(E_j^k)} .$$

Connecting the IV specification to the measured CPI. Using the equation above and equation (D9), we obtain that each sector satisfies:

$$d\log(\widetilde{P}_j^k) = \frac{1}{\theta} d\log(S_{jj}^k) + d\log(w_j) - d\log(E_j^k) . \tag{D12}$$

Using this equation, we can obtain a bound on the potential bias from the extensive margin, i.e. we quantify whether the fact that the measured CPI misses the extensive margin response could help explain the magnitude of the IV coefficient. Conceptually, an unobserved fall in product

[25]In Corollary 1.1 in the main text, to streamline notation we use $d\log(E_{si})$ instead of $d\log(E_{ij}^k)$.

[26]The BLS measures a first order approximation to the COLI using a geometric Laspeyres index, which by Roy's identity is the theory-consistent price index for small shocks, as discussed in Appendix B.A.

[27]In Corollary 1.1 in the main text, to streamline notation we use $d\log(\widetilde{P}_s)$ instead of $d\log(\widetilde{P}_j^k)$.

variety when the foreign expenditure shares decreases would mute the change in the exact price index, potentially addressing the fact that the measured price index falls about ten times more than predicted by the Melitz-Chaney model.

To obtain a bound, we assume that there is no fall in the number of foreign varieties imported from the country experiencing the positive supply shock, e.g. China, which is consistent with empirical evidence (e.g., Broda and Weinstein (2006)). This assumption means that $d\log(E_{jc}^k) \leq 0$, since a weakly positive change in product variety lowers the exact price index. Using this assumption and equation (D12), and denoting by m the set of countries excluding the country where the supply shock occurs (China) as well as the domestic economy (the US), we obtain:

$$d\log(\widetilde{P}_j^k) < \frac{1}{\theta} d\log(S_{jj}^k) + d\log(w_j) - S_{jj}^k d\log(E_{jj}^k) - S_{jm}^k d\log(E_{jm}^k).$$

From the inequality above, and noting that

$$\begin{aligned} d\log(E_{ij}^k) &\equiv s_{ij}^k(z_{ij}^k) g_i^k(z_{ij}^k) \frac{dz_{ij}^k}{\sigma - 1} \\ &= \frac{\theta - (\sigma - 1)}{\sigma - 1} \left(d\log(w_i) - d\log(A_i^k) - d\log(P_j^k) + \frac{1}{\sigma - 1} d\log(w_i/w_j) \right), \end{aligned}$$

we obtain that the estimated regression coefficient, $\hat{\beta}$, in the reduced-form specification $d\log(\widetilde{P}_j^k) = \alpha + \beta d\log(S_{jj}^k) + \varepsilon_{sd}$ satisfies the following bound:

$$\hat{\beta} < \sum_k \omega_k \left(\frac{1}{\theta} - S_{jj}^k \frac{\partial[d\log(E_{jj}^k)]}{\partial[d\log(S_{jj}^k)]} - S_{jm}^k \frac{\partial[d\log(E_{jm}^k)]}{\partial[d\log(S_{jj}^k)]} \right) = \left(1 + \frac{\theta - (\sigma - 1)}{\sigma - 1} \sum_k \omega_k (1 - S_{jm}^k) \right) \frac{1}{\theta}$$

Setting $\theta = 4.25$ (Simonovska and Waugh (2014)), $S_{jm}^k = 0.0452$ (Acemoglu et al. (2016), for 1999), and $\sigma = 5$, we obtain $\hat{\beta} < 0.24$, while our IV estimates are one order of magnitude larger. This result shows that it is implausible that extensive margin adjustments can explain the magnitude of our IV estimate.

Connecting the IV specification to the measured domestic CPI. We denote by $\widetilde{P}_{jj}^k$ the measured price index over domestically-produced goods.[28] We have:

$$d\log(\widetilde{P}_{jj}^k) = d\log(I_{ij}^k) = d\log(w_j).$$

The estimated regression coefficient in the reduced-form specification $d\log(\widetilde{P}_{jj}^k) = \alpha + \beta d\log(S_{jj}^k) + \varepsilon_{kj}$ is:

$$\hat{\beta} \equiv \frac{\mathrm{Cov}\left(d\log(\widetilde{P}_{jj}^k), d\log(S_{jj}^k) \right)}{\mathrm{Var}\left(d\log(S_{jj}^k) \right)} = \frac{\mathrm{Cov}\left(d\log(w_j), d\log(S_{jj}^k) \right)}{\mathrm{Var}\left(d\log(S_{jj}^k) \right)} = 0.$$

Intuitively, the theory-consistent measured domestic CPI from this model only responds through

[28] In Corollary 1.1 in the main text, to streamline notation we use $d\log(\widetilde{P}_{sd})$ instead of $d\log(\widetilde{P}_{jj}^k)$.

changes in domestic wages, which cannot contribute to price differences across sectors in the baseline Melitz-Chaney model with a single labor type; this feature generates a cross-sectoral price response of zero, which is inconsistent with our IV estimates.

D.B Connecting the IV Specification to Trade Models with Intermediate Inputs

D.B.1 Setting

To connect our IV specification to models with intermediate inputs, we use the same setting as in Section D.A.1, except that the production function includes intermediate inputs across sectors, as in Caliendo and Parro (2015) and Ossa (2015). We now have $y_{ij}^k(z) = A_i^k z(\ell_{ij}^k)^{1-\alpha_k} \left[\Pi_{k'=1}^K \left(m_{ij}^{k,k'}(z)\right)^{\zeta_{k,k'}}\right]^{\alpha_k}$, where $m_{ij}^{k,k'}$ denotes the composite intermediate good from sector k' used in production by firms from country i selling in j in sector k. The share of intermediate inputs in value added is denoted by α_k, and intermediate inputs have shares that sum to one, i.e. $\sum_{k'} \zeta_{k,k'} = 1$.

Firms price at a constant markup $\frac{\sigma}{\sigma-1}$ over the unit cost such that, by cost minimization, the price for the good from a firm from country i in sector k selling to country j is:

$$p_{ij}^k(z) = \frac{\sigma}{\sigma - 1}\frac{\tau_{ij}^k \phi_k c_i^k}{z},$$

with

$$\begin{aligned}
\phi_k &\equiv (1-\alpha_k)^{-(1-\alpha_k)} \left(\Pi_{k'} \left(\zeta_{k,k'}\alpha_k\right)^{\zeta_{k,k'}}\right)^{-\alpha_k} \\
c_i^k &\equiv \frac{1}{A_i^k} w_i^{1-\alpha_k} \left(\Pi_{k'} \left(P_i^{k'}\right)^{\zeta_{k,k'}}\right)^{\alpha_k}.
\end{aligned}$$

and the entry threshold for productivity is:

$$z_{ij}^k = \left(\frac{\sigma w_i f_{ij}}{X_j^k}\right)^{\frac{1}{\sigma-1}} \frac{\frac{\sigma}{\sigma-1}\tau_{ij}^k \phi_k c_i^k}{P_j^k}.$$

D.B.2 Proof of Corollary 1.2

The proof of Corollary 1.2 follows the same steps as the Proof of Proposition 1. First, we derive the relationship between changes in the price index and changes in the domestic expenditure share that each sector satisfies. Substituting the log change in the entry threshold for domestic firms,

$$d\log(z_{jj}^k) = -d\log(P_j^k) + (1-\alpha_k)d\log(w_j) + \alpha_k \sum_{k'} \zeta_{k,k'} d\log(P_j^{k'}),$$

into the price change for domestic firms,

$$d\log(P_{jj}^k) = (1-\alpha_k)d\log(w_j) + \alpha_k \sum_{k'} \zeta_{k,k'} d\log(P_j^{k'}) + \frac{\theta - (\sigma - 1)}{\sigma - 1} d\log(z_{jj}^k),$$

yields a change in the domestic sales share of,

$$d\log(S^k_{jj}) = -\theta\left((1-\alpha_k)d\log(w_j) - (1-\alpha_k\zeta_{k,k'})d\log(P^k_j) + \alpha_k\sum_{k'\neq k}\zeta_{k,k'}d\log(P^{k'}_j)\right).$$

Rearranging, we obtain:

$$d\log(P^k_j) = \frac{1}{\theta(1-\alpha_k\zeta_{k,k})}d\log(S^k_{jj}) + \frac{1-\alpha_k}{1-\alpha_k\zeta_{k,k}}d\log(w_j) + \frac{\alpha_k}{1-\alpha_k\zeta_{k,k}}\Sigma_{k'\neq k}\zeta_{k,k'}d\log(P^{k'}_j).$$

This equation implies that regression coefficient in the reduced-form specification $d\log(\widetilde{P}^k_{jj}) = \alpha + \beta d\log(S^k_{jj}) + \varepsilon_{kj}$ will be biased if there is a non-zero covariance between $d\log(S^k_{jj})$ and price changes for the bundle of intermediates, $\Sigma_{k'\neq k}\zeta_{k,k'}d\log(P^{k'}_j)$.

Empirically, when the import penetration rate from China increases by 1 percentage point in industry j, the share of intermediate inputs from China in industry$j$'s total output increases by only ten basis points. Thus, empirically industries that are more exposed to import competition are only slightly more exposed to the imported intermediate inputs channel. This fact motivates the assumption that the price change in intermediates experienced by a given industry is *smaller* than the price change of the industry itself, which we formalize by assuming $\Sigma_{k'\neq k}\zeta_{k,k'}d\log(P^{k'}_j) < d\log(P^k_j)$, using $\sum_{k'}\zeta_{k,k'} = 1$. Thus, we obtain that each sector satisfies the bound:

$$d\log(P^k_j) < \frac{1}{\theta(1-\alpha_k)}d\log(S^k_{jj}) + d\log(w_j).$$

Therefore, given that the above inequality holds for each sector k, the estimated regression coefficient in the reduced-form specification, $d\log(P^k_j) = \alpha + \beta d\log(S^k_{jj}) + \varepsilon_{kj}$, satisfies:

$$\hat{\beta} < \sum_k \omega_k \frac{1}{\theta(1-\alpha_k)}$$

where ω_k are sector weights derived from the heterogeneous treatment estimator detailed in Section D.A.3 above.

Thus, the price response for domestic consumers is magnified in a way that is proportional to the use of intermediate inputs. To parameterize this upper bound, we use the national accounts table from the Bureau of Economic Analysis (BEA) to obtain the share of intermediate goods, which yields $\alpha_k = 56.4\%$ at the mean across industries. With $\theta = 4.25$, we obtain $\hat{\beta} < 0.54$, which is much smaller than our IV estimates. This bound is conservative as it consider a case where price indices in $k' \neq k$ fall to the same extent as the price index for k, even though, empirically, the change in import penetration is about ten times smaller (see Figure 4).

D.C Connecting the IV Specification to Trade Models with Oligopolistic Competition

D.C.1 Setting

We use the setting of Edmond et al. (2015).[29] The economy consists of two countries, a domestic country and a foreign country, denoted by d and f respectively. In what follows, we focus on the domestic economy.

Consumer's Problem. A representative consumer in the domestic country has CES preferences over a continuum of differentiated product categories, with an elasticity of substitution between categories of $\epsilon > 1$,

$$Y = \left(\int_{\Omega} (Y_s)^{\frac{\epsilon-1}{\epsilon}} ds \right)^{\frac{\epsilon}{\epsilon-1}},$$

where Y_s is the consumption bundle of product category s, and Ω is the set of product categories. The associated price index is $P = \left(\int_{\Omega} (P_s)^{1-\epsilon} ds \right)^{\frac{1}{1-\epsilon}}$. The consumer maximizes utility subject to their budget constraint, which yields the following demand function for differentiated product category s,

$$S_s \equiv \frac{P_s Y_s}{X} = \left(\frac{P_s}{P} \right)^{1-\epsilon} \tag{D13}$$

where $X \equiv PY$ is total expenditure, and S_s is the domestic country's share of expenditures on product category s.

Within each of product category s there are n_{sd} domestic firms and n_{sf} foreign firms that produce closely related product varieties. The elasticity of substitution between the products within a product category is denoted γ, with $\gamma > \epsilon$. Consumers have CES demand for products within product category s that are produced in country $i \in \{d, f\}$

$$Y_{si} = \left(\sum_{j=1}^{n_{si}} (y_{sij})^{1-\gamma} \right)^{\frac{1}{1-\gamma}}$$

and face associated prices,

$$P_{si} = \left(\sum_{j=1}^{n_{si}} (p_{sij})^{1-\gamma} \right)^{\frac{1}{1-\gamma}}$$

where y_{sij} is the output of firm j of country i within category s, and p_{sij} is the associated price. The domestic consumer's demand for firm j's variety from country i within product category s yields the consumer's expenditure share,

$$S_{sij} \equiv \frac{p_{sij} y_{sij}}{X_s} = \left(\frac{p_{sij}}{P_s} \right)^{1-\gamma} \tag{D14}$$

[29] This setting is also very close to Atkeson and Burstein (2008).

where $X_s \equiv P_s Y_s$ is the total expenditure on products from product category s. Likewise, the domestic consumer's expenditure share of country $i \in \{d, f\}$'s products within category s is,

$$S_{si} = \left(\frac{P_{si}}{P_s}\right)^{1-\gamma} \tag{D15}$$

and we denote the share of the domestic consumer's expenditures on firm j's variety relative to all products from its origin country i within product category s as $\tilde{S}_{sij} \equiv S_{sij}/S_{si} = p_{sij}y_{sij}/P_{si}Y_{si}$.

Product Competition. Within a product category, firms produce according to a linear production technology that takes labor as an input, $y_{sij}(z) = zA_{si}\ell_{sij}$, where A_{si} is aggregate productivity of firms from country i in product category s and z is idiosyncratic productivity that is distributed according to $z \sim G_i^s(z)$. Firms take their country's wage, w_i, as given, face an iceberg cost, $\tau_{si} \geq 1$ with $\tau_{sd} = 1$, and pay a labor-denominated fixed cost $w_i f_{si}$ to operate in the market.

Firms compete oligopolistically within a sector. In Section D.C.2, we derive a proposition based on a non-parametric representation of oligopolistic competition using markup elasticities, as in Amiti et al. (2019a). Let p_{sij} and μ_{sij} denote respectively the price and markup of firm j from country i in sector s. A change in a firm's marginal cost passes-through into its own price at a rate given by $\frac{1}{1+\Gamma_{sij}}$, where $\Gamma_{sij} \equiv -\frac{\partial \log(\mu_{sij}(.))}{\partial \log(p_{sij})}$ is the firm's "own-price markup elasticity". Under perfect and monopolistic competition with CES preferences, $\Gamma_{sij} = 0$. In Section D.C.4, we derive a corollary under more specific assumptions, using Cournot competition.

D.C.2 Proof of Proposition 2

Price and Expenditure Share Responses. The price response of product category s is obtained by applying Shephard's lemma,

$$d\log(P_s) = S_{sd}d\log(P_{sd}) + S_{sf}d\log(P_{sf}) \ . \tag{D16}$$

where the change in the price index for firms from countries $i \in \{d, f\}$ operating within the product category s is given by

$$d\log(P_{si}) = \sum_{j=1}^{n_{si}} \tilde{S}_{sij} d\log(p_{sij}) \ .$$

The change in country i's expenditure share within product category s is,

$$d\log(S_{si}) = (1-\gamma)\Big(d\log(P_{si}) - d\log(P_s)\Big) \tag{D17}$$

and since $dS_{sd} + dS_{sf} = 0$, for the domestic country this can be expressed as,

$$\frac{d\log(P_s)}{dS_{sf}} = \frac{-1}{(\gamma-1)S_{sd}} + \frac{d\log(P_{sd})}{dS_{sf}} \ . \tag{D18}$$

Thus, the price effect consists of two components: the first term is the direct effect of the foreign productivity shock that would be recovered in an economy with constant markups, $-1/(\gamma-1)S_{sd}$, and the second term is an indirect effect from price response of domestic firms, $d\log(P_{sd})/dS_{sf}$.

Characterizing the domestic price response with markup elasticities. To characterize the price response from domestic firms, we follow Amiti et al. (2019a) by log-linearizing the consumer's demand functions and firms' optimality conditions. Each firm's price response is a weighted average of the change in the firm's own marginal cost $c_{sij} \equiv \tau_{si} w_i/(z_{sij} A_{si})$ and the change in the product category's price index,

$$d\log(p_{sij}) = \frac{1 - S_{sij}}{1 - S_{sij} + \Gamma_{sij}} d\log(c_{sij}) + \frac{\Gamma_{sij}}{1 - S_{sij} + \Gamma_{sij}} d\log(P_s) \tag{D19}$$

where Γ^k_{sij} is the firm's own-price markup elasticity, and where $d\log(c_{sij}) \equiv d\log(c_{si}) = d\log(w_i/A_{si})$ since firms' idiosyncratic productivity shocks and the iceberg costs do not respond to the foreign supply shock.[30]

Within sector s, country i's aggregated price response to the foreign supply shock is given by,

$$d\log(P_{si}) = \Omega_{si} d\log(w_i/A_{si}) + (1 - \Omega_{si}) d\log(P_s) \tag{D20}$$

and the overall price response for the sector is,

$$d\log(P_s) = \sum_{i \in d,f} \left(\frac{S_{si}\Omega_{si}}{\sum_k S_{sk}\Omega_{sk}} \right) d\log(w_i/A_{si}) \tag{D21}$$

where, for notational compactness, we define the aggregated cost pass-through as,

$$\Omega_{si} \equiv \sum_{j=1}^{n_{si}} \frac{S_{sij}}{S_{si}} \frac{1 - S_{sij}}{1 - S_{sij} + \Gamma_{sij}}.$$

Since $\Gamma_{sij} \geq 0$, we know that $\Omega_{si} \leq 1$. Under monopolistic competition, $\Gamma_{sij} = 0$ for all firms and thus $\Omega_{si} = 1$. To illustrate the role of strategic interactions, consider the limit case $\Gamma_{sij} \to \infty$: we then have $\Omega_{si} \to 0$ and, from equation (D20), the country price index responds one-for-one to the competitor price index.

From equations (D15) and (D20), the change in the expenditure share on country i's products is,

$$d\log(S_{si}) = (1 - \gamma) \left(d\log(w_i/A_{si}) - d\log(P_s) \right).$$

We further assume that the foreign country receives a positive productivity shock such that $d\log(A_{sf}) > 0$ for foreign firms, while domestic firms do not experience a productivity shock so that $d\log(A_{sd}) = 0$. Accordingly, rearranging the change in the domestic expenditure share in

[30] In equation (D19) we have applied the accounting decomposition of Amiti et al. (2019a) in which the own-price markup elasticity equals the cumulative competitor markup elasticity, e.g., $-\partial \log(\mu^k_{sij})/\partial \log(p^k_{sij}) = \sum_{l \neq j} \partial \log(\mu^k_{sil})/\partial \log(p^k_{sil})$.

response to the foreign shock yields the following price response expression,

$$d\log(P_s) = d\log(w_d) + \frac{1}{(\gamma-1)\Omega_{sd}} d\log(S_{sd}), \tag{D22}$$

and thus the domestic price response to the shock is given by substituting equation (D22) into equation (D20),

$$d\log(P_{sd}) = d\log(w_d) + \frac{1-\Omega_{sd}}{(\gamma-1)\Omega_{sd}} d\log(S_{sd}). \tag{D23}$$

Substituting the expression for Ω_{sd} into the above equation, and noting that $dS_{sd} + dS_{sf} = 0$, we obtain the expression for the domestic price response in Proposition 2,

$$\begin{aligned} \frac{d\log(P_{sd})}{dS_{sf}} &= \frac{d\log(w_d)}{dS_{sf}} - \frac{1-\Omega_{sd}}{(\gamma-1)\Omega_{sd}S_{sd}} \\ &= \frac{d\log(w_d)}{dS_{sf}} + \frac{-1}{(\gamma-1)S_{sd}} \sum_{j=1}^{n_{sd}} \left(\frac{\frac{1-S_{sdj}}{1-S_{sdj}+\Gamma_{sdj}}}{\sum_{k=1}^{n_{sd}} \frac{1-S_{sdk}}{1-S_{sdk}+\Gamma_{sdk}}} \right) \frac{\Gamma_{sdj}}{1-S_{sdj}}, \end{aligned} \tag{D24}$$

where the second term is the domestic price effect with respect to the change in foreign expenditure share expressed in terms of each domestic firm's market share and markup elasticity within the product category. In particular, the domestic price effect is a weighted average of each domestic firm's competitor-share normalized markup elasticity. Thus, in the absence of strategic interactions (if $\Gamma_{sdj} = 0$ for all $j = 1, \ldots, n_{sd}$), the domestic price response would only equal the wage response and we would recover an expression identical to the price effect from the constant markup case in Section D.A.

Trade elasticity. The trade elasticity depends on firms' market power, which generates imperfect passthrough of changes in trade costs into prices, as well as strategic pricing. From (D15), we have

$$d\log(S_{sf}/S_{sd}) = (1-\gamma)\left(d\log(P_{sf}) - d\log(P_{sd})\right).$$

Substituting equations (D20) and (D21) into this expression, we obtain the trade elasticity:

$$\eta_s \equiv \frac{\partial \log(S_{sf}/S_{sd})}{\partial \log(A_{sf})} = (\gamma-1)\left(S_{sf}\Omega_{sd}^{-1} + S_{sd}\Omega_{sf}^{-1}\right)^{-1} \tag{D25}$$

Bounds for the difference between the overall price index and the domestic price index. Rearranging equation (D17) and substituting $dS_{sd} + dS_{sf} = 0$ yields,

$$d\log(P_s) = \frac{-1}{(\gamma-1)S_{sd}} dS_{sf} + d\log(P_{sd})$$

which implies the lower bound,

$$\frac{-d\log(P_s)}{dS_{sf}} - \frac{-d\log(P_{sd})}{dS_{sf}} = \frac{1}{(\gamma-1)S_{sd}} \geq 0$$

since $S_{sd} > 0$ and $\gamma > 1$.

To obtain the upper bound, we first show that $\eta_s \leq (\gamma-1)$. Since $\Omega_{si} \leq 1$ for each country

$i \in \{d, f\}$, we have $\Omega_{si}^{-1} \geq 1$ and we can create the following weighted average,

$$S_{sf}\Omega_{sd}^{-1} + S_{sd}\Omega_{sf}^{-1} \quad \geq \quad S_{sf} + S_{sd} = 1 \ .$$

By equation (D25), the left-hand side of the above equation equals $(\gamma - 1)/\eta_s$ and thus $1/\eta_s \geq 1/(\gamma - 1)$. Therefore, for each sector s we obtain

$$0 \leq \frac{-d\log(P_s)}{dS_{sf}} - \frac{-d\log(P_{sd})}{dS_{sf}} = \frac{1}{(\gamma-1)S_{sd}} \leq \frac{1}{\eta_s S_{sd}} \ . \tag{D26}$$

Collecting equations (D22), (D23), (D24) and (D26) completes the proof of Proposition 2.

D.C.3 Proof of Corollary 2.1

In this section, we derive the structural interpretation for the estimated regression coefficient $\hat{\beta}$ in a general oligopolistic competition model. By assumption the productivity shock only affects the foreign country, therefore the change in the domestic cost of production is given by $d\log(c_{sdj}^k) = d\log(w_j)$ for all domestic firms $j = 1, \ldots, n_{sd}$. We can express equation (D22) such that each sector k satisfies:

$$d\log(P_{sd}) = d\log(w_j) + \frac{1}{(\gamma-1)\Omega_{sd}S_{sd}}dS_{sd}.$$

This equation shows that, despite incomplete passthrough of changes in costs into prices, the endogenous domestic wage change enters our regression specification additively, without sector-specific coefficient, such that it is differenced out in estimation.

While this relationship holds in each sector s, we run the reduced-form specification $d\log(P_{sd}) = \alpha + \beta S_{sd} + \varepsilon_{sd}$, with a constant coefficient β for all sectors. Applying the estimator in equation (D10) for regressions with heterogeneous treatment effects to this setting, we obtain the desired expression:

$$\hat{\beta} = \int_0^1 \omega_s \frac{1}{(\gamma-1)\Omega_{sd}S_{sd}} ds,$$

where the sectoral weights ω_s are the same as in Section D.A.3. This expression that the regression coefficient is larger when markups are more responsive. In the limit, the regression coefficient can be unboundedly large: as $\Gamma_{sdj} \to \infty$, we have $\Omega_{sd} \to 0$ and $\hat{\beta} \to \infty$.

Likewise, applying the estimator to the domestic price response in equation (D23) yields a coefficient of,

$$\hat{\beta}_{dom} = \int_0^1 \omega_s \frac{1-\Omega_{sd}}{(\gamma-1)\Omega_{sd}S_{sd}} ds.$$

Finally, since equation (D26) tells us that $0 \leq -(d\log(P_s) - d\log(P_{sd}))/dS_{sf} = \frac{1}{(\gamma-1)S_{sd}} \leq \frac{1}{\eta_s S_{sd}}$ for each sector s, the estimated regression coefficients in reduced-form specifications using the overall price index and the domestic price index as dependent variables satisfy the bound $0 \leq \hat{\beta} - \hat{\beta}_{dom} \leq \frac{1}{\eta \cdot S_d}$, where η denotes the average trade elasticity, defined such that $\frac{1}{\eta S_d} = \int_0^1 \omega_s \frac{1}{\eta_s S_{sd}} ds$.

D.C.4 Proof of Corollary 2.2

Price effects with head-to-head Cournot competition. We now consider the special case of Cournot competition, which yields:

$$\Gamma_{sij} \equiv -\frac{d\log(\mu_{sij})}{d\log(p_{sij})} = \frac{(\gamma-\epsilon)(1-S_{sij})}{\gamma(\epsilon-1)+(\gamma-\epsilon)(1-S_{sij})}(\gamma-1)S_{sij} \ . \tag{D27}$$

With head-to-head competition between one domestic firm and one foreign firm, substituting (D27) into equations (D18) and (D24) completes the proof of Corollary 2.2.

Calibration. We choose two moments to discipline the parameters ϵ and γ. First, we match the elasticity of trade flows in equation (D25) to the empirical estimate of 4.25 from Simonovska and Waugh (2014). Second, we match the estimated price effect of 1.91 in the model from equation (D18). Finally, we parameterize the share of expenditures on items produced in China by $S_{fs} \equiv 0.0452$, which is the average value for 1999 from Acemoglu et al. (2016). Our calibration matches both moments[31] and recovers parameter values of $\epsilon = 1.43$ and $\gamma = 8.73$. These estimates are similar those in Edmond et al. (2015), who estimate their parameters from highly disaggregated 7-digit Taiwanese manufacturing data in a quantitatively model with rich firm heterogeneity. Furthermore, these elasticity parameters imply a markup elasticity $\Gamma_{sd} = 0.59$ that is close to that estimated by Amiti et al. (2019a) on disaggregated 10-digit Belgian data. The empirical estimate of Amiti et al. (2019a), which was not targeted in our calibration, is $\widehat{\Gamma_{sd}} = 0.62$.[32]

Price effects with head-to-head Bertrand competition. The preceding results remain similar with Bertrand competition instead of Cournot competition. We have

$$\Gamma_{sij} \equiv \frac{(\gamma-\epsilon)(1-S_{sij})}{(S_{sij}\epsilon+(1-S_{sij})\gamma)(S_{sij}\epsilon+(1-S_{sij})\gamma-1)}(\gamma-1)S_{sij} \ ,$$

and we repeat our calibration of parameters ϵ and γ to match the trade elasticity of 4.25 and price effect of 1.91. Our calibration again tightly matches these empirical moments and recovers parameters values $\epsilon = 2.08$ and $\gamma = 7.47$, implying a markup elasticity $\Gamma_{sd} = 0.48$. As in the Cournot case, these parameters are in line with the literature, highlighting the plausibility of the oligopolistic competition channel to explain our large price effects.

Interpretation of head-to-head competition effects. Referring to the domestic and foreign variables as belonging to the "US" and "China" respectively, the relationship between the industry price index and the change in the import penetration rate from China is:[33]

$$d\log(P_s^{US}) = -\frac{1+\Gamma_s^{US}/S_s^{China}}{(\sigma_s-1)(1-S_s^{China})}\cdot dS_s^{China}.$$

Intuitively, Chinese producers reduce prices when they experience a positive productivity shock,

[31] The model fits the data moments very tightly, within 6 decimal points of the target.

[32] We obtain this estimate from Column 5 of their Table 1: $\hat{\Gamma} = 1/\hat{\alpha} - 1$, with $\hat{\alpha} = 0.616$.

[33] Here we abstract from the change in the domestic wage for the sake of explication.

which leads U.S. producers to also reduce prices due to strategic interactions. Because of the U.S. price response, the equilibrium change in the spending share on the product from China is lower than it would be absent this price response. As a result, the relationship between changes in import penetration from China and price changes (our IV coefficient) can be large.

To illustrate the logic, consider a limiting case with an extremely high markup elasticity, $\Gamma_s^{US} \to \infty$, in which the two producers supply highly substitutable products and become Bertrand competitors. In such a case, the U.S. producer matches the fall in price from the Chinese producer almost entirely, i.e. $d\log\left(P_s^{US}\right) \approx d\log\left(P_s^{China}\right)$. This can be readily seen by substituting the components of the overall price index, $d\log(P_s) = S_s^{US} d\log(P_s^{US}) + S_s^{China} d\log(P_s^{China})$, into equation (D19) under head-to-head competition,

$$d\log(P_s^{US}) = \frac{1}{1+\Gamma_s^{US}} d\log(c_s^{US}) + \frac{\Gamma_s^{US}}{1+\Gamma_s^{US}} d\log(P_s^{China})$$

where $\Gamma_s^{US} \to \infty$ implies that $\frac{1}{1+\Gamma_s^{US}} \to 0$ and $\frac{\Gamma_s^{US}}{1+\Gamma_s^{US}} \to 1$. Because the relative price of the two producers remains almost unchanged, the import penetration rate from China barely changes. Since both P_s^{US} and P_s^{China} fall, so does the industry price index P_s. Therefore we get price effects despite no changes in expenditure shares: $\frac{d\log(P_s)}{dS_s^{China}} \to \infty$ as $\Gamma_s^{US} \to \infty$. In this limiting case, the reduced-form relationship between price changes and changes in trade with China across industries can be unboundedly large.

Contributions of Domestic vs. Foreign Price Changes. This setting is also useful to illustrate the relative contributions of domestic and foreign price changes. The model predicts a larger decrease in the price of a given Chinese firm than from a given US firm, because a reduction in the Chinese firm's marginal cost creates a larger downward price adjustment relative to U.S. firms'. However, US firms account for approximately 95% of initial market share, which puts much greater quantitative weight on the domestic price response.

To see this, suppose that the U.S. firm's marginal cost does not change while the foreign firm's marginal cost does, e.g. $d\log(c_s^{US}) = 0$ and $d\log(c_s^{China}) \neq 0$. It is immediate that the magnitude of the U.S. price change is indeed lower than that of the Chinese price change because $d\log(p_s^{US}) = \frac{\Gamma_s^{US}}{1+\Gamma_s^{US}} d\log(p_s^{China})$, with $\Gamma_s > 0$ but markups are not perfectly elastic.

We can write the aggregate price response as,

$$\begin{aligned} d\log(p_s) &= S_s^{China} d\log(p_s^{China}) + (1 - S_s^{China}) d\log(p_s^{US}) \\ &= (1 - S_s^{China}) \frac{\Gamma_s^{US}}{1+\Gamma_s^{China}} d\log(p_s^{China}) + S_s^{China} d\log(p_s^{China}). \end{aligned}$$

If $(1 - S_s^{China}) d\log(p_s^{US}) > S_s^{China} d\log(p_s^{China})$, then the domestic price response dominates the foreign price response in the aggregate. This condition is met empirically because we have

$$\Gamma_s^{US} > \frac{\frac{S_s^{China}}{1-S_s^{China}}}{1-\frac{S_s^{China}}{1-S_s^{China}}} = 0.05$$

given that market shares satisfy $S_s^{China} = 0.0452$, and that the evidence in Amiti et al. (2019a) provide a range of estimates for the markup elasticity that satisfies $\Gamma_d^{US} > 1/3$ (with a preferred estimate of $\Gamma_d^{US} = 0.6$). Thus, the domestic price change dominates the foreign price change in accounting for the aggregate price response.

D.D Connecting the IV Specification to Trade Models with Monopolistic Competition and Variable Elasticities of Substitution

For completeness, we now draw a connection between our results and the predictions of Arkolakis et al. (2019). Arkolakis et al. (2019) study a model in which the univariate distribution of markups is independent of the level of trade costs. Instead of obtaining variable markups by using CES utility with a departure from monopolistic competition such as assuming Bertrand or Cournot competition (as in Bernard et al. (2003), Atkeson and Burstein (2008) and Edmond et al. (2015)), Arkolakis et al. (2019) endogenize markups through consumer preferences featuring varying elasticities of substitution at the firm-level with monopolistic competition.

Monopolistic competition yields markups $m_{si} = \varepsilon_{D_{si}}/(\varepsilon_{D_{si}}-1)$, where, in the notation of Arkolakis et al. (2019), ε_D denotes the elasticity of demand $D(p_{si}/P_s)$ and $D(.)$ is a strictly decreasing function. Arkolakis et al. (2019) focus on additively separable preferences in the Pollak family, specifying:[34]

$$q_{si} \equiv D(p_{si}/P_s) = (p_{si}/P_s)^{1/\gamma} - \alpha,$$

for firm i operating in category s, which nests the CES case with $\alpha = 0$. Arkolakis et al. (2019) specify $\alpha = 1$ and $\gamma = -0.35$. Under monopolistic competition, the firm's price elasticity is $\frac{\partial \log(p_{si})}{\partial \log(q_{si})} = \gamma \frac{q}{q+\alpha} < 0$ and $\frac{\partial^2 \log(p_{si})}{\partial \log(q_{si})^2} = \gamma \frac{\alpha}{(q+\alpha)^2} < 0$, i.e. demand becomes less elastic as quantities increase, since $\alpha > 0$. In that case, larger firms face a more inelastic demand and have larger markups. In this setting, an increase in foreign productivity leads to a fall in domestic production, leading to a more elastic demand for domestic firms and a fall in domestic markups.

Next, we show that this framework makes two counterfactual predictions. First, we derive the predicted domestic price response to a foreign supply shock in this framework, so that we can compare the prediction to our empirical estimates.[35] We assume that domestic production costs remain unchanged, i.e. we only need to compute the change in markups for a domestic firm, dm_{si}, induced by a foreign supply shock. With $\varepsilon_{D_{sd}} \equiv -\frac{q+\alpha}{q}\frac{1}{\gamma}$ as in Arkolakis et al. (2019), we get

[34] See Section 5.1 in Arkolakis et al. (2019).

[35] We focus on the response of the measured domestic price index, rather than the exact price index.

$d\varepsilon_{D_{sd}} = \frac{dq_{sd}\alpha}{(q_{sd})^2}\frac{1}{\gamma} > 0$, with $dq_{sd} < 0$ due to the foreign supply shock. Thus, demand becomes more elastic and domestic markups decrease, with $dm_{si} = \frac{-1}{(\varepsilon_{D_{si}}-1)^2}d\varepsilon_{D_{sd}} < 0$. Thus,

$$\begin{aligned}
\frac{d\log(p_{sd})}{dS_{sf}} &= \frac{d\log(m_{sd})}{dS_{sf}} \\
&= -\frac{1}{(\varepsilon_{D_{si}}-1)^2}\frac{d\varepsilon_{D_{sd}}}{dS_{sf}}\frac{1}{m_{sd}} \\
&= -\frac{\alpha/\gamma}{\varepsilon_D \cdot (\varepsilon_{D_{si}}-1)\cdot q_{sd}}\frac{d\log(q_{sd})}{dS_{sf}} \\
&= \frac{\alpha/\gamma}{(\varepsilon_D \cdot (\varepsilon_{D_{si}}-1)\cdot q_{sd} - \alpha/\gamma)\cdot S_{sd}}
\end{aligned}$$

where the fourth line uses $d\log(q_{sd}) = d\log(S_{sd}) - d\log(p_{sd})$[36] and $dS_{sd} = -dS_{sf}$.

To evaluate this expression, we use Arkolakis et al. (2019)'s preferred parameters of $\alpha = 1$ and $\gamma = -0.35$ and $S_{sf} = 0.0452$ from Acemoglu et al. (2016). Furthermore, we consider parameter values for $\varepsilon_{D_{si}}$ and q_{sd} that are consistent with evidence on U.S. markups. According to De Loecker et al. (2020), markups were about 45% above marginal cost in the 2000s, implying $\varepsilon_{D_{si}} = 3.2$ and $q_{sd} = 8$.[37] With these parameters, we obtain the prediction:

$$\frac{\widehat{dlog(p_{sd})}}{dS_{sf}} = -0.0506.$$

Thus, the model predicts that domestic prices across sectors should fall by 5 basis points when the foreign expenditure share increases by 1 percentage point, which is more than one order of magnitude smaller than our IV estimate in Table 6.

In a robustness check, we obtain a similarly small price effect with an alternative calibration strategy, choosing parameter values to match the markup elasticity of $\Gamma_{sd} = 0.6$ in Amiti et al. (2019a) instead of the markup from De Loecker et al. (2020).

Intuitively, models of endogenous markups using oligopolistic competition and strategic interactions are better able to match our large IV estimate (as Bernard et al. (2003), Atkeson and Burstein (2008) and Edmond et al. (2015)), compared to models using variable elasticities of substitution (as Arkolakis et al. (2019)), because strategic interactions can deliver a large fall in prices with small changes in trade flows. As discussed in Section D.C.4, with oligopolistic competition the reduced-form relationship between price changes and changes in trade with China across industries can be unboundedly large in limiting cases. In contrast, with variable elasticities of substitution, conventional parameter values require a large change in trade flows to induce a sizable change in markups.

Second, the framework counterfactually predicts a larger response of domestic markups for

[36] Here we use the fact that, with Cobb-Douglas preferences across categories indexed by s, in general equilibrium total category-level expenditures remain fixed, i.e. $d\log(X_s) = 0$ (see Section D.A.4).

[37] Indeed, with $q_{si} = 8$ we have $\varepsilon_{D_{sd}} \equiv -\frac{8+1}{8}\frac{1}{-0.35} = 3.2$ and $m_{si} = 3.2/(3.2-1) = 1.45$, as desired.

smaller firms. For example, there is a choke price such that demand is zero when $p_{si} \geq P_s$: at this level of null demand ($q_{sd} = 0$), the demand elasticity is infinite ($\varepsilon_{D_{sd}} \equiv -\frac{q_{sd}+\alpha}{q_{sd}}\frac{1}{\gamma}$), and the rate of increase in the demand elasticity is locally infinite ($d\varepsilon_{D_{sd}} = \frac{dq_{sd}\alpha}{(q_{sd})^2}\frac{1}{\gamma}$), so the response of markups is largest for infinitesimal firms. This framework's prediction is counterfactual. Instead of finding that the increase in markups is larger for smaller firms in response to an increase to the foreign expenditure share, in Section IV.D we find the opposite: the price response is higher for larger firms, as well as when the market is more concentrated, which is consistent with the predictions of oligopolistic competition models.

E Online Appendix Tables and Figures

Figure A1: Summary Statistics for Estimated Markups and Profitability Measures

Panel A: Estimated Markups

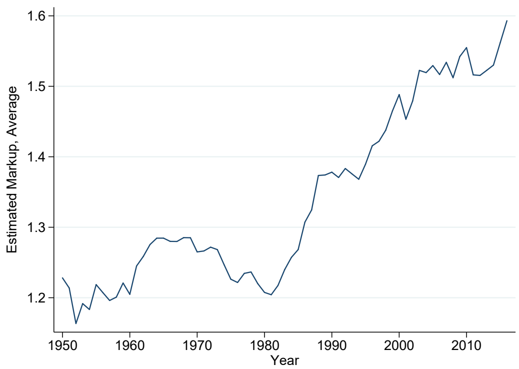

(a) Average

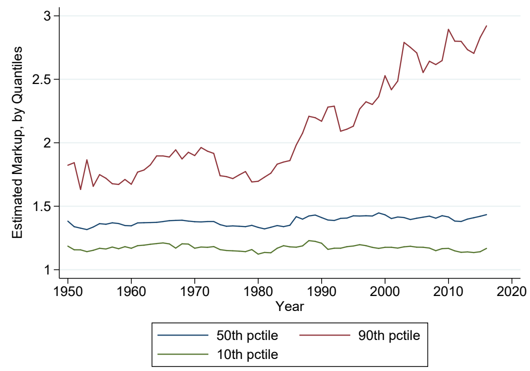

(b) Quantiles

Panel B: Profitability Measures

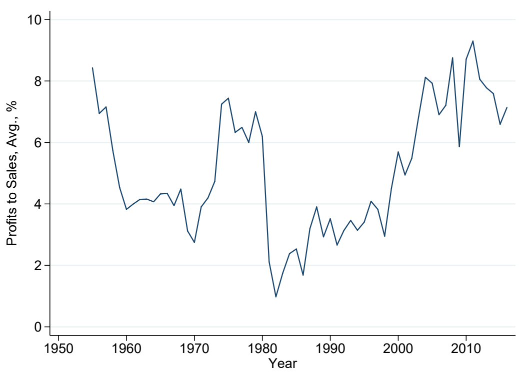

(a) Profits/Sales, Average

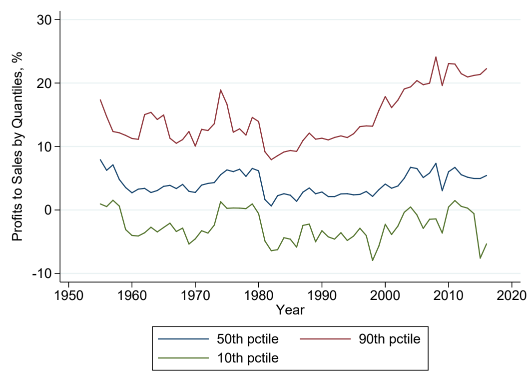

(b) Profits/Sales, Quantiles

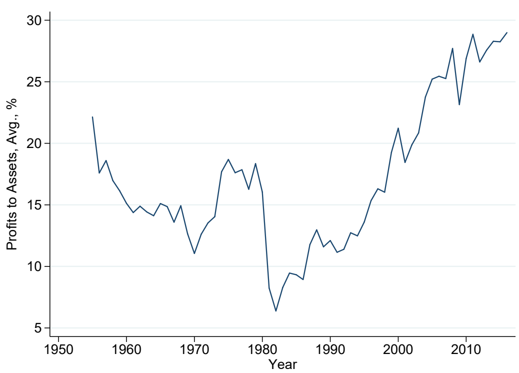

(c) Profits/Assets, Average

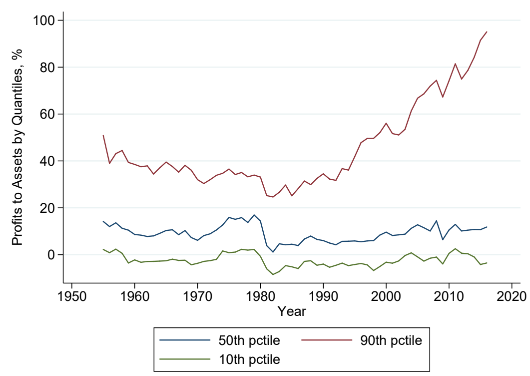

(d) Profits/Assets, Quantiles

Notes: This figure reports trends in estimated markups. We follow the methodology of De Loecker et al. (2020) (described in Appendix B.F) and obtain similar results. Statistics are computed using sales as weights.

Figure A2: Additional Pre-trends Tests

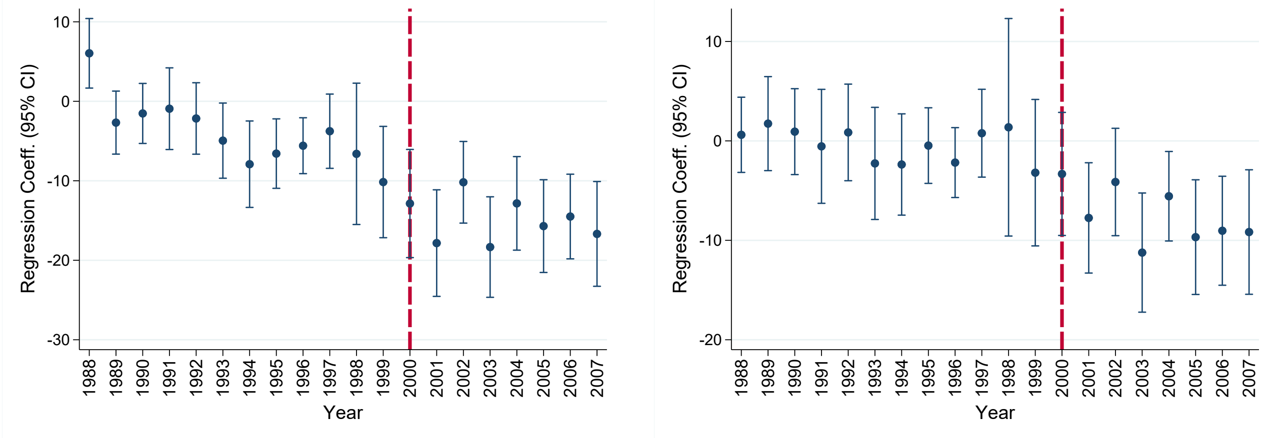

(A) Without Controls

(B) With Apparel & Durables F.E.

Notes: In Panel A, the specification is the same as described in Section III.B, but without any fixed effects. In Panel B, the specification is the same as in Section III.B, but with fixed effects for apparel and durables instead of ELI fixed effects. Panel A exhibits pre-trends, in contrast with Panel A and Figure 1 in the main text. Panel B shows no evidence for pre-trends. These results indicate that including fixed effects for apparel and durables is important to ensure that a causal interpretation of the estimates is plausible, and suggest that ELI fixed effects may not be necessary.

Figure A3: Testing for Pre-trends in the Extended CPI Sample

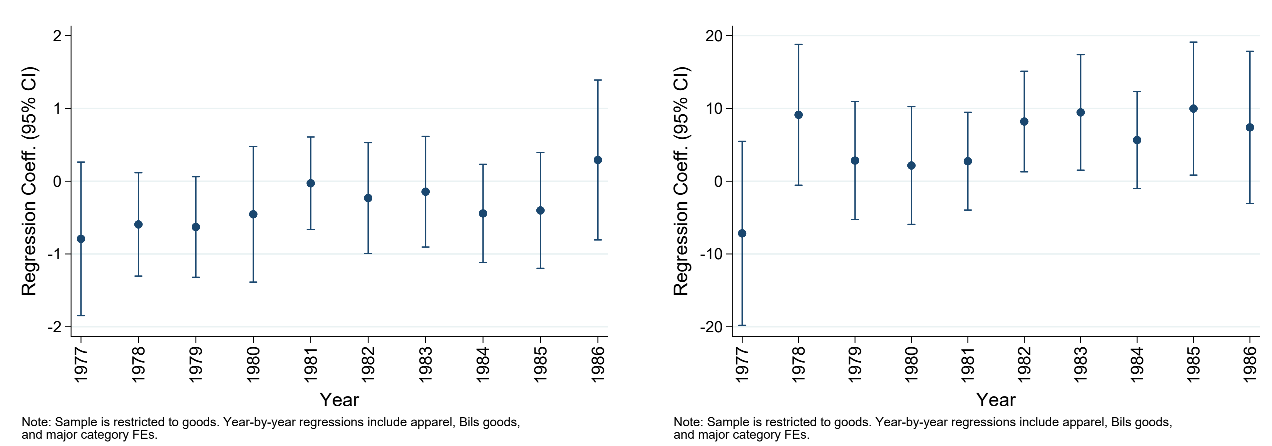

(A) China IP in other developed economies

(B) NTR gap

Notes: This figure uses the extended CPI sample with the specification described in Online Appendix B.C. F-tests indicate that we cannot reject that the estimated coefficients are jointly insignificant.

Figure A4: Long-Run Event Studies in NBER-CES Manufacturing Database

Panel A: Excluding NAICS 334, Computers and Electronics

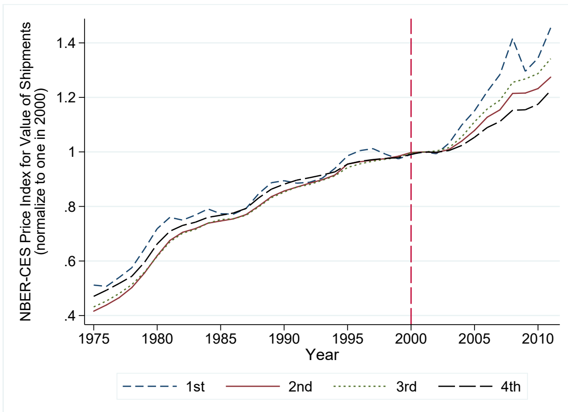

(a) NTR gap, by quartiles

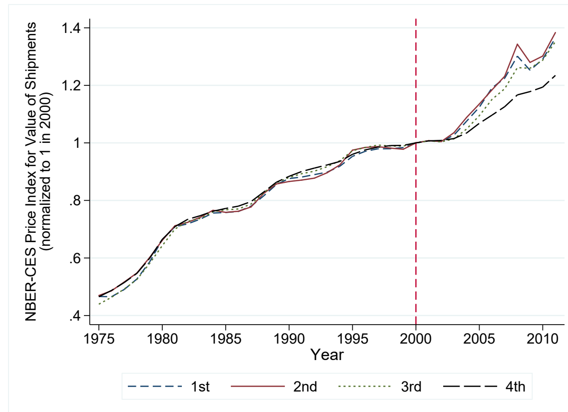

(b) China IP in other developed economies, by quartiles

Panel B: Including NAICS 334, Computers and Electronics

(a) NTR gap, by quartiles

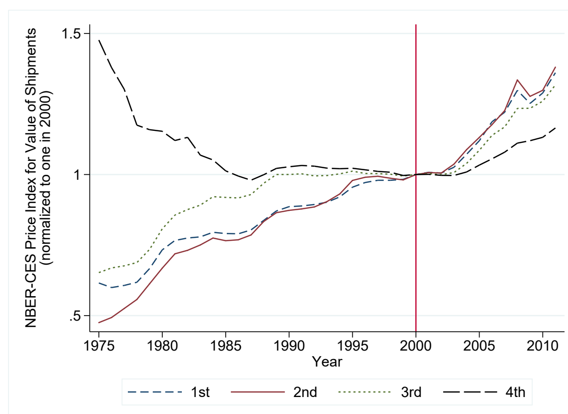

(b) China IP in other developed economies, by quartiles

Notes: This figure reports a long-run analysis of price trends by quartiles of the instruments for trade with China. A higher quartile indicates higher exposure. We thank Teresa Fort for recommending us to conduct this analysis. The data source for prices is the NBER-CES Manufacturing Industry database, which provides a price index for the value of shipments for each 6-digit NAICS industry in each year from 1975 to 2011. All industries within manufacturing are covered, including those providing intermediate inputs. We match this data set to the instruments for trade with China: the NTR gap is available across 6-digit NAICS codes; using the SIC-NAICS described in Online Appendix B.C, we link the data set to the 2000-2007 change in the import penetration rate from China in other developed economies. In all panels, the price index for the value of shipments is normalized to one in 2000 and the price trends are reported by quartiles of exposure to the instruments. Panel A excludes industries belonging to the 3-digit NAICS category "Computers and Electronics" (NAICS 334). In this panel, industries across quartiles of exposure are on similar price trends up to the treatment period (starting in 2000) and start diverging afterwards. With both instruments, more exposed industries have a lower inflation rate after 2000. These results are consistent with the estimates presented in Section III. The specification reported in this figure helps reduce noise by showing cumulative price differences over time. When we run an analysis with fixed effects analogous to specification (2) in the main text, the year-specific estimates are too noisy to discern a statistically significant pattern. Panel B includes industries within "Computers and Electronics": when doing so, large pre-trends appear because these industries are more exposed to the instruments and have been on lower inflation trends for decades. These results indicate the importance of excluding these categories or including suitable controls, as we do in Section III. The results are similar with a median split or by deciles of exposure of the instruments, instead of quartiles (not reported).

Figure A5: Mean independence between the log change in import penetration and the initial import penetration across product categories

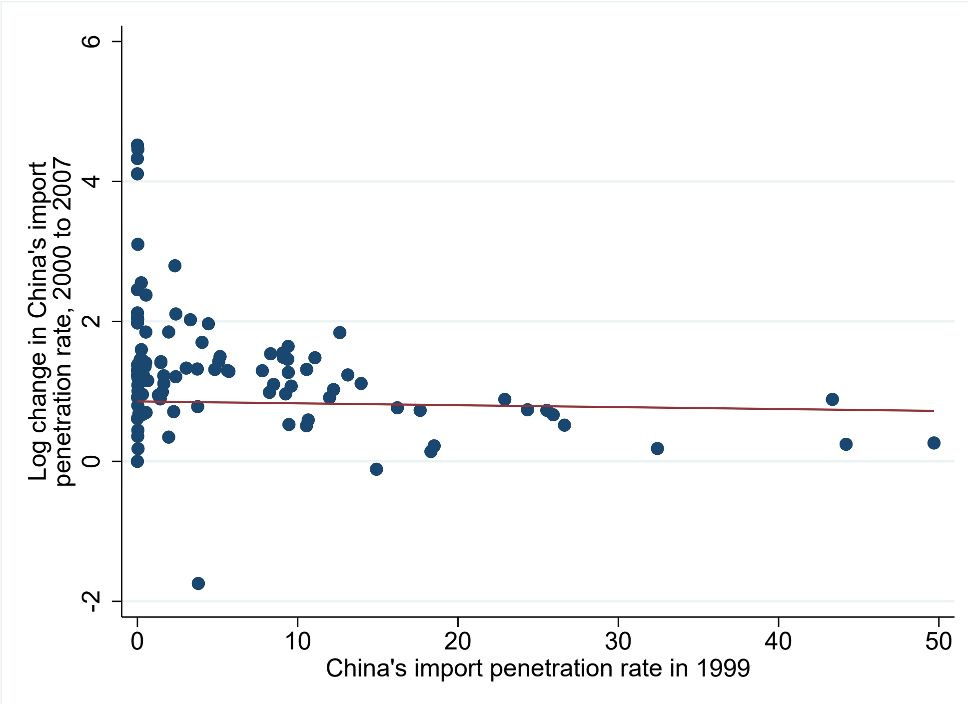

Notes: This figure is a binned scatter plot presenting the relationship between the log change in the import penetration rate from China between 2000 and 2007 and the initial import penetration rate from China, in 1999. The lack of relationship motivates our mean independence assumption in Section D.A.3, which is necessary to derive Proposition 1. The OLS relationship in this figure is depcited in red and is not statistically significant, with a point estimate of -0.00271 (s.e. 0.00427).

Table A1: IV Estimates in the Subsample of Consumer Packaged Goods, as in Bai and Stumpner (2019)

	U.S. CPI Inflation (pp)
	IV
Δ China IP (pp)	-0.685^{*}
	(0.39)
First-stage F	15.29
2000-2007 only	✓
Goods, Durables & Apparel F.E.	✓
Instruments:	
Δ China IP Other	✓
N	82

Notes: To assess whether our results are in line with those of Bai and Stumpner (2019), this table reports the IV estimates in the sample of product categories corresponding to consumer packaged goods. To carry out this analysis, we keep ELIs belonging to the following product categories: food at home (starting with the following two letters: FA, FB, FC, FD, FE, FF, FG, FH, FI, FJ, FK, FL, FM, FN, FO, FP, FQ, FR, FS, FT), tobacco products (starting with letters GA), personal care products (GB), appliances (HK), household equipment (HL), housekeeping suppliers (HN), pet food (RB011), toys, games, and playground equipment (RE011), and magazines (RG011). We then implement our IV specification in this subsample. Specifically, we use the specification closest to Bai and Stumpner (2019), using only cross-sectional variation after 2000 with the instrument of Autor et al. (2013). The point estimate reported in this table is -0.685 (s.e. 0.39), i.e. about 50% smaller than the corresponding estimate in the full sample (see Panel B of Table 2, Col. 2, where the point estimate is 1.27 (s.e. 0.28)). The 95% confidence interval spans the values for the regression coefficient implied by the Melitz-Chaney model under homogeneous or heterogeneous trade elasticities (equal to 0.23 and 0.38 respectively, as discussed in Section IV.A). These results highlight the importance of considering the full consumption basket, in particular given that most of trade with China occurs in product categories that do not belong to consumer packaged goods.

Table A2: List of Durable Goods, reproduced from Bils (2009)

Good	Spending share
Watches	0.069
Jewelry	0.416
Personal computers & equipment	0.370
Telephone & equipment	0.080
Calculators, typewriters, etc.	0.014
Electric personal care products	0.021
Luggage	0.027
Infant's equipment	0.018
Curtains & drapes	0.064
Window coverings	0.053
Mattresses & springs	0.146
Bedroom furniture	0.193
Sofas & slipcovers	0.276
Living room chairs	0.122
Living room tables	0.057
Kitchen & dining room furniture	0.162
Infant's furniture	0.025
Occasional furniture	0.148
Refrigerators & home freezers	0.083
Washers & dryers	0.103
Stoves	0.030
Microwaves	0.029
Vacuums	0.064
Small kitchen appliances	0.034
Other electric appliances	0.079
Lamps & lighting	0.040
Clocks & decorative items	0.325
Dishes	0.083
Flatware	0.014
Nonelectric cookware	0.039
Tableware & nonelectric kitchenware	0.057
Power tools	0.058
Misc. hardware	0.096
Nonpowered hand tools	0.026
Medical equipment for general use	0.011
Supportive & convalescent equipment	0.031
Televisions	0.246
Other video equipment	0.104
Audio equipment	0.164
Bicycles	0.044
General sports equipment	0.229
Hunting, fishing, & camping equipment	0.086
Photography equipment	0.057
Sewing machines	0.044
Musical instruments & accessories	0.069

Notes: This table presents the lists of durable goods reported in Table 1 of Bils (2009)

Table A3: IV Estimates from Panel Specification with Additional Periods

Panel A: With NTR gap Instrument

	U.S. CPI Inflation (pp)
	IV
	(1)
Δ China IP (pp) — 2000-2002	-2.03^{**}
	(0.94)
Δ China IP (pp) — 2003-2005	-2.62^{**}
	(1.29)
Δ China IP (pp) — 2006-2007	-2.28^{**}
	(1.01)
First-stage F	17.43
ELI F.E.	✓
Period-specific Goods F.E.	✓
N	1332

Panel B: With Change in Import Penetration in Other Countries

	U.S. CPI Inflation (pp)
	IV
	(1)
Δ China IP (pp) — 2004-2007	-1.64^{**}
	(0.75)
Δ China IP (pp) — 2000-2003	-1.12^{**}
	(0.53)
Δ China IP (pp) — 1996-1999	-1.42^{**}
	(0.72)
Δ China IP (pp) — 1991-1995	-1.24^{**}
	(0.52)
First-stage F	15.21
ELI F.E.	✓
Period-specific Goods F.E.	✓
N	888

Notes: The level of observation is an ELI-by-period cell. Panel A uses the NTR gap as the instrument. The periods are indexed by t, with $t \in \{1991-1993\}, \{1994-1996\}, \{1997-1999\}, \{2000-2002\}, \{2003-2005\}, \{2006-2007\}$. The IV specification is:

$$\pi_{it} = \sum_{k=1991-1993}^{2006-2007} \beta_k \Delta ChinaIP_{it} \cdot 1_{\{k=t\}} + \delta_i + \delta_t + \varepsilon_{it},$$

$$\Delta ChinaIP_{it} = \sum_{k=1991-1993}^{2006-2007} \gamma_k \mathbf{Z_{it}} \cdot 1_{\{k=t\}} + \tilde{\delta}_i + \tilde{\delta}_t + \eta_{it},$$

Panel B uses the same specification with the change in import penetration in other economies as the instrument. The periods are $t \in \{1991-1995\}, \{1996-1999\}, \{2000-2003\}, \{2004-2007\}$. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A4: Sensitivity Analysis

	U.S. CPI Inflation (pp)				
	IV	IV	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)
Δ China IP (pp)	-2.23^{***}	-2.20^{***}	-2.31^{***}	-1.78^{***}	-1.86^{***}
	(0.47)	(0.34)	(0.49)	(0.59)	(0.54)
First-stage F	38.14	52.88	32.69	24.0	26.5
ELI F.E.	✓	✓	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓	✓	✓
Consumption weights	✓		✓	✓	✓
No weights		✓			
Full sample	✓	✓		✓	✓
Goods only			✓		
Periode-specific Durables & Apparel F.E.				✓	
Period-specific Major Category F.E.					✓
N	444	444	344	444	444

Notes: The level of observation is an ELI-by-period cell. Standard errors are clustered by ELIs. The instrument is the NTR gap. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A5: OLS Estimand with Homogeneous or Heterogeneous Trade Elasticities in the Melitz-Chaney Model

Panel A: Distribution of Broda and Weinstein (2006)'s trade elasticities, θ^s, across ELIs

	Broda and Weinstein (2006)'s trade elasticities
Mean	4.25
Standard deviation	3.36
p25	2.01
p50	3.99
p75	4.82

Panel B: Implied OLS regression coefficients, $\hat{\beta}$, in the Melitz-Chaney Model

	$\hat{\beta}$
Homogeneous trade elasticity, $\hat{\beta} = 1/\mathbb{E}[\theta^s]$	0.2352
Heterogeneous trade elasticities, $\hat{\beta} = \sum_s \omega_s \frac{1}{\theta^s}$	0.3877

Notes: Panel A reports the distribution of trade elasticities from Broda and Weinstein (2006), averaged at the ELI level using our crosswalk from HS6 categories to ELIs. The panel shows substantial heterogeneity across ELIs, considering only the sample of ELIs with positive trade flows. Using the results in Proposition 1, panel B reports the implied OLS regression coefficients in the cross-ELI regression, $\Delta \log(P_s) = \alpha + \beta \Delta \log(S_{sd}) + \varepsilon_{sd}$, according to the Melitz-Chaney model. The weights ω_s are provided in Appendix D.A.3. The results reported in this panel show that the implied OLS regression coefficient is 64% larger when considering heterogeneous trade elasticities, compared with a counterfactual case with homogeneous trade elasticities for all product categories equal to the average trade elasticity. While the coefficient is magnified, it remains about five times smaller than our empirical estimates.

Table A6: The Roles of Continued and Domestic Goods, using the Change in Import Pentration in Other Developed Economics as the Instrument

Panel A: IV Estimates for Continued Goods in Main Sample (CPI)

	U.S. CPI Inflation, Continued Products (pp)		Contribution to U.S. CPI Inflation (pp) *[%]*	
	(1)	(2)	(3)	(4)
Δ China IP (pp)	-1.58^{*}	-2.15^{***}	-0.65 *[45%]*	-1.25^{***} *[98.4%]*
	(0.91)	(0.71)	(0.49)	(0.33)
ELI F.E.	✓		✓	
Period-specific Goods F.E.	✓		✓	
2000-2007 only		✓		✓
Goods, Durables & Apparel F.E.		✓		✓
N	444	222	444	222

Panel B: IV Estimates for Domestic Goods in Main Sample (CPI)

	U.S. CPI Inflation, Domestic Products (pp)		Contribution to U.S. CPI Inflation (pp) *[%]*	
	(1)	(2)	(4)	(5)
Δ China IP (pp)	-1.26^{***}	-1.32^{***}	-0.92^{***} *[64%]*	-1.08^{***} *[85%]*
	(0.48)	(0.31)	(0.38)	(0.25)
ELI F.E.	✓		✓	
Period-specific Goods F.E.	✓		✓	
2000-2007 only		✓		✓
Goods, Durables & Apparel F.E.		✓		✓
N	444	222	444	222

Panel C: IV Estimates for Continued and Domestic Goods in PPI Sample

	U.S. PPI Inflation (pp)		U.S. PPI Infl., Continued Products (pp)	
	(1)	(2)	(4)	(5)
Δ China IP (pp)	-1.45^{**}	-2.13^{*}	-1.02^{**}	-0.96^{**}
	(0.67)	(1.22)	(0.48)	(0.45)
First-stage F	652.20	521.31	652.20	521.31
NAICS F.E.	✓		✓	
Period-specific Goods F.E.	✓		✓	
2000-2007 only		✓		✓
Goods, Durables & Apparel F.E.		✓		✓
N	550	275	550	275

Notes: The specifications are the same as for Table 6 in the main text, except that the instrument is the change in import penetration in other developed economie. Standard errors are clustered by industries. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A7: Summary Statistics on Country of Origin Flags

	Number of ELIs with flags	Share of Expenditures with flags	
	All	All	Tradables
Year	(1)	(2)	(3)
2000	51	0.1830	0.3703
2001	59	0.1760	0.3555
2002	59	0.1828	0.3630
2003	63	0.1929	0.3959
2004	62	0.1860	0.3865
2005	65	0.2016	0.4300
2006	60	0.1832	0.3877
2007	61	0.1743	0.3668

Notes: This table presents summary statistics on the number of ELIs with a country of origin flag. This ELIs explicitly gather country of origin information (e.g., "Was the product made in the United States; Yes or No?" or "Write in the country in which the product was made."). Country of origin flags are obtained from specification checklists, as explained in Online Appendix B.A.

Table A8: Correlations between Direct and Indirect Exposure to Trade with China

	Δ China IP Supplier, First-order IO (pp)			Δ China IP Buyer, First-order IO (pp)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ China IP (pp)	0.099***	0.073***	0.046***	0.018***	0.01241	−0.00081
	(0.014)	(0.0154)	(0.010)	(0.0062)	(0.0098)	(0.00191)
6-digit IO F.E.		✓	✓		✓	✓
Period-specific Goods F.E.		✓	✓		✓	✓
Excl. diagonal of IO Table			✓			✓
N	170	170	170	170	170	170

Notes: The specifications are described in Section IV.C. The level of observation is a 6-digit IO industry-by-period cell. The instrument is the NTR gap. Standard errors are clustered by IO industries. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A9: Summary Statistics on Related-Party Trade

	Mean	S.D.	p10	p50	p90	Sample
Share of related trade, All Countries, 2005, %	48.08	25.37	11.85	47.99	82.61	NAICS
Share of related trade, China, 2005, %	26.07	23.49	2.148	17.70	65.55	
Share of related trade, All Countries, 2015, %	50.61	24.86	17.43	50.24	86.20	
Share of related trade, China, 2015, %	27.61	19.84	4.90	25.69	56.49	
Share of related trade, China, 2005, %	11.45	17.24	1.227	4.098	38.38	ELI matched sample

Notes: This table reports summary statistics on the share of U.S. imports occuring between related parties in trade, with all trading partners and with China specifically. The data source is the related-party trade database of the U.S. Census. The original data is provided across NAICS codes, but the patterns are similar once we match the data to our ELI sample (as shown in the fifth row). The average share of related-party trade is smaller in our ELI sample (11% in 2005) than in the full NAICS sample (26% in 2005) because our sample covers final goods and there tends to be more trade between related parties for intermediate products. Although the average share of related-trade is is small, there is susbtantial variation across ELIs. For example, the share of related-party trade from China is particularly high for computer storage devices (72%) and other computer equipment (65%), while is it low for "men's suits and coats" (1.9%) and "women's suits and coats" (2.1%).

Table A10: Employment Effects of Trade

Panel A: With the NTR Gap

	Δ Non-Prod Emp. (pp)	Δ Prod Emp. (pp)	Δ Total Emp. (pp)
	(1)	(2)	(3)
Δ China IP (pp)	-2.591	-1.893	-1.834
	(0.789)	(0.648)	(0.537)
First-stage F	25.464	25.464	25.464
ELI F.E.	✓	✓	✓
Durables & Apparel Time-Varying F.E.	✓	✓	✓
N	306	306	306

Panel B: With the Change in Import Penetration from China in Other Developed Economies

	Δ Non-Prod Emp. (pp)	Δ Prod Emp. (pp)	Δ Total Emp. (pp)
	(1)	(2)	(3)
Δ China IP (pp)	-2.319	-2.137	-1.774
	(1.174)	(1.002)	(0.930)
First-stage F	13.860	13.860	13.860
ELI F.E.	✓	✓	✓
Period-specific F.E. for Durables/Apparel	✓	✓	✓
N	306	306	306

Panel C: With Both Instruments

	Δ Non-Prod Emp. (pp)	Δ Prod Emp. (pp)	Δ Total Emp. (pp)
	(1)	(2)	(3)
Δ China IP (pp)	-2.506	-1.970	-1.815
	(0.635)	(0.574)	(0.498)
First-stage F	16.234	16.234	16.234
ELI F.E.	✓	✓	✓
Period-specific F.E. for Durables/Apparel	✓	✓	✓
N	306	306	306

Notes: This table reports the results from our baseline IV specification (3), except that the outcome is the change in industry employment (expressed in %). Panel A uses the NTR gap instrument, Panel B uses changes in import penetration in other developed economies, and Panel C uses both instruments jointly. The employment outcomes are measured in the NBER CES database, which distinguishes between "production" and "non-production" workers. We consider in turn employment for production workers, non-production workers, and total employment as outcomes. The results indicate that employment falls by 1.77% to 2.59%, depending on the specification, for each one percentage point increase in the import penetration rate from China. We obtain similar results when we use total employment from the County Business Patterns Database instead (not reported). Standard errors are clustered at the level of ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A11: Estimates of the Increase in Consumer Surplus from Increased Trade with China from 2000-2007, in 2007 U.S. Dollars

	(1)	(2)	(3)
Annual Increase in Consumer Surplus, \$/Household	1,711	1,105	1,466
Calibration Parameters:			
- Observed Cumulative Change in China IP, 2000-2007, within Goods (=6.15pp)	✓	✓	✓
- Average Household Spending on Goods in 2007, CEX (=\$12,479)	✓	✓	✓
IV Estimates:			
- NTR gap: $\beta_{price} = -2.23$	✓		
- ΔChina IP Other: $\beta_{price} = -1.44$		✓	
- Both: $\beta_{price} = -1.91$			✓

Notes: This table estimates the gains to U.S. consumers from the fall in prices induced by the increase in trade with China from 2000 to 2007. The results are expressed in 2007 dollars of consumer surplus per U.S. household. Assuming that there are no GE effects affecting prices in all product categories, our cross-industry IV estimates accurately reflect the price changes induced by increasing trade with China at the level of the whole economy. If prices do not revert back in the future, the estimated annual gains reported in the table should persist going forward. Under these assumptions, the estimate in Column (1) should be interpreted as follows: from 2007 onward, the annual purchasing power of the average U.S. consumer is \$1,711 higher thanks to the increase in trade with China between 2000 and 2007 (which is about 2% of total consumption expenditures). The increase in consumer surplus is computed based on three components: (a) the increase in import penetration from China between 2000 and 2007 in the set of all tradable product categories (denoted $\Delta_{2000-2007}ChinaIP$); (b) the IV estimates for the price response (β); and (c) average household spending on tradable product categories in 2007, which we measure in the 2007 Consumer Expenditure Survey (denoted C_{2007}). A product of these three ingredients gives a first-order approximation to the annual consumer surplus created by falling prices from increase trade with China: $CS = \frac{-\beta}{100} \cdot \Delta_{2000-2007}ChinaIP \cdot C_{2007}$. The three columns of the table consider different estimates for the price response, which we apply to the observed cumulative change in import penetration from China between 2000 and 2007 (equal to 6.15 percentage points within the set of tradable product categories). In Column (1), using the IV estimate from the NTR gap instrument, we obtain an increase in consumer surplus per household of \$1,711 ($= \frac{-2.23}{100} \cdot 6.15 \cdot 12479$). Columns (2) and (3) report the results using alternative instruments for the price effects. The increase in consumer surprlus is \$1,105 per U.S. household using the change in trade with China in other developed economies (Column (2)), and \$1,466 using both instruments jointly (Column (3)). These estimates are much larger than predicted by the class of trade models nested by Arkolakis et al. (2012). Using a standard trade elasticty of $\theta = -4$, the predicted price effect in these models is $\beta = \frac{1}{\theta} = -0.25$, implying an increase in consumer surplus of \$192 per U.S. household ($= \frac{-0.25}{100} \cdot 6.15 \cdot 12479$). In robustness checks, we find that these results are similar (i) when we use the BEA's measure of average personal consumption expenditures on tradable goods in 2007 (equal to $\$11,153$) instead of the estimates from the CEX, as well as (ii) when we use an adjusted measure for the change in trade with China as in Acemoglu et al. (2016) (they attempt to isolate the share of the observed increase in trade with China that was caused by increased productivity in China, rather than by other factors such as a fall in productivity in the U.S.). The calculations underlying this table rest on two simplifying assumption that may understate the magnitude of the gains to U.S. consumers: (1) the first-order approximation does not allow U.S. consumers to reallocate their expenditures toward product categories that become relatively cheaper – including these second-order gains would further increase consumer surplus; (2) we assumed away GE effects affecting all product categories, but if increasing import penetration induces an overall fall in domestic prices to restore trade balance, then the increase in purchasing power for U.S. consumers would increase further.

Table A12: Summary Statistics for CPI and PPI samples

Panel A: CPI Sample

	Mean	S.D.	p10	p50	p90	Observations	
						N	Aggreg. Level
Inflation, continued products (%)	-3.77	10.75	-19.04	1.18	7.36	3,774	ELI-by-year

Panel B: PPI Sample

	Mean	S.D.	p10	p50	p90	Observations	
						N	Aggregation Level
Inflation, all (%)	0.04	10.19	-7.19	0.84	7.28	1,044	NAICS6-by-period
ΔChina IP in U.S., direct	0.39	0.69	0.00	0.11	1.04		
NTR Gap	0.24	0.19	0.00	0.29	0.45	522	NAICS6

Notes: Panel A presents summary statistics for inflation for continued products in the CPI sample. Panel B presents summary statistics for the PPI sample, which is described in Section II.A and Online Appendix B.E. The sample extends between 1991 and 2007 and is divided into two periods, 1991-1999 and 2000-2007.

Table A13: Testing for Pre-trends in the Extended CPI Sample (1977-1986)

	Annual U.S. CPI Inflation	
	(1)	(2)
NTR Gap	1.798	
	(2.285)	
Δ China IP Other		−0.2863
		(0.5016)
N	156	156

Notes: This table reports the reduced-form specifications in the extended CPI sample. The level of observation is an ELI and heteroeksedasticty-robust standard errors are reported in parentheses. The corresponding binned scatter plots are shown in Panel B of Figure 1 in the main text. The extended CPI sample is described in Online Appendix B.B.

Table A14: Placebo First-stage with French Trade Data

	Δ China IP (in France, pp)
	OLS
	(1)
NTR Gap	0.63
	(0.82)
COICOP F.E.	✓
Period-specific Goods F.E.	✓
N	264

Notes: The level of observation is a COICOP-by-period cell. Standard errors are clustered by COICOPs.*** denotes statistical significance at the 1% level, ** at the 5% level.

Table A15: First Stage for Panel Specification with Additional Periods

	Δ China IP (pp)					
	$\gamma_{1991-1993}$	$\gamma_{1994-1996}$	$\gamma_{1997-1999}$	$\gamma_{2000-2002}$	$\gamma_{2003-2005}$	$\gamma_{2006-2007}$
NTR Gap	0.61	-0.13	0.43	2.65^{***}	3.23^{***}	3.73^{***}
	(0.51)	(0.67)	(0.54)	(0.56)	(0.64)	(0.71)
ELI F.E.			✓			
Period-specific Goods F.E.			✓			
N			1332			

Notes: The level of observation is an ELI-by-period cell. The periods are indexed by t, with $t \in \{1991-1993\}, \{1994-1996\}, \{1997-1999\}, \{2000-2002\}, \{2003-2005\}, \{2006-2007\}$. The IV specification is:

$$\pi_{it} = \sum_{k=1991-1993}^{2006-2007} \beta_k \Delta ChinaIP_{it} \cdot 1_{\{k=t\}} + \delta_i + \delta_t + \varepsilon_{it},$$

$$\Delta ChinaIP_{it} = \sum_{k=1991-1993}^{2006-2007} \gamma_k \mathbf{Z_{it}} \cdot 1_{\{k=t\}} + \tilde{\delta}_i + \tilde{\delta}_t + \eta_{it},$$

The table reports the first-stage estimates γ_k. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A16: IV Estimates with NTR Gap Instrument and Additional Controls

	U.S. CPI Inflation (pp)	
	IV	IV
	(1)	(2)
Δ China IP (pp)	-2.10^{***}	-2.03^{***}
	(0.62)	(0.58)
First-stage F	23.19	19.42
ELI F.E.	✓	✓
Period-specific Goods F.E.	✓	✓
Time-varying controls for High-tech, Contract intensity and Union membership	✓	✓
Controls for MFA quota reductions, Chinese import tariffs, export licensing requirements, and Chinese production subsidies		✓
Instrument: NTR Gap	✓	✓
N	444	444

Notes: This tables is similar to Table 3 in the main text, but follows Pierce and Schott (2016) by including additional controls in Column (2): exposure to MFA quota reductions, Chinese import taris from Brandt et al. (2012), data on export licensing requirements from Krishna, Bai, and Ma (2015), and data on production subsidies from China's National Bureau of Statistics. The level of observation is an ELI-by-period cell. Standard errors are clustered by ELIs.*** denotes statistical significance at the 1% level, ** at the 5% level.

Table A17: IV Results with Controls for Exports

	U.S. CPI Inflation	
	IV	IV
	(1)	(2)
Δ China IP	-1.805^{***}	-1.447^{***}
	(0.474)	(0.3358)
First-stage F	25.628	205.028
Controls:		
Change in exports to China, 1991-1999	✓	✓
Exports to China in 1992	✓	✓
Instruments:		
NTR Gap	✓	
Δ China IP Other		✓
N	306	306

Notes: This table reports the IV estimates with specifications similar to Section III.C but including controls for exports from the U.S. to China across product categories. Exports to China are measured in trade data recorded under HS codes (which we link to NAICS industries and to ELIs using the crosswalks from Online Appendix B.C). The controls include the log change in exports to China from 1991 to 1999, as well as the level of exports to China in 1992. The results are similar when repeating the anlaysis in subsamples (above vs. below median of exports), when including controls in level and changes for exports to China for other years, and when including all exports instead of exports to China specifically (not reported). These results indicate that the baseline IV estimates are not confounded by differences in export dynamics across product categories. The level of observation is an ELI-by-period cell. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A18: A Test of the Uncertainty Channel — First-stage Heterogeneity by Capital Intensity

	ΔChina IP	
	OLS	OLS
	(1)	(2)
Δ NTR Gap	3.861***	
	(1.361)	
Δ NTR Gap $\times$ Capital Intensity	4.786**	
	(2.308)	
Δ China IP Other		1.353***
		(0.331)
Δ China IP Other $\times$ Capital Intensity		-0.249
		(0.357)
N	306	306

Notes: This table reports the results from first-stage regressions with interaction terms for capital intensity. The interaction term is the only difference with the specifications described in Section III.C. Capital intensity is measured in the NBER-CES database. The sample is restricted to ELIs that can be matched to this data set; the crosswalk is described in Online Appendix B.C. The NTR gap instrument corresponds to a fall in uncertainty over import tariffs applied by the U.S. to China. In a situation of uncertainty, standard models generate a region of inaction in investment space due to nonconvex adjustment costs (e.g., Dixit and Pindyck (1994)). If the relevance of the NTR gap instrument is driven by the uncertainty channel, we expect the first stage to be stronger in product categories that are more capital intensive. Column (1) confirm that this is the case in the data. Column (2) reports a placebo test and shows that the first stage features no heterogeneity by capital intensity when the instrument is the change in import penetration from China in the set of developed economies other than the United States. The level of observation is an ELI-by-period cell. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A19: Testing for Heterogeneity by Trade Elasticities

	U.S. CPI Inflation	
	IV	IV
	(1)	(2)
Δ China IP	-2.363^{***}	-1.911^{**}
	(0.399)	(0.816)
Subsample:	Trade Elasticity $\geq$ p50	Trade Elasticity $<$ p50
Instrument: NTR Gap	✓	✓
N	140	140

Notes: This table reports the IV estimates from the baseline specification from Section III.C in two subsamples, above and below the median trade elasticity as estimated by Broda and Weinstein (2006). The trade elasticities were estimated by Broda and Weinstein (2006) for the period 1990 - 2001 across HS codes, which we match and aggrgeate to the level of ELIs. The IV estimates are a bit larger in the subsample with a higher trade elasticitity. The estimates in both subsamples are similar to the baseline IV results from Table 2. In theory, the relationship between changes in import penetration from China and U.S. consumer prices could have widely varied depending on the trade elasticity. This table indicates that in practice the magnitudes are relatively stable, implying that our baseline IV estimate provides a meaningful summary measure. In other (unreported) IV specifications, we find that when interacting the estimated trade elasticity with the change in import penetration from China, the interaction term is not statistically significant. The level of observation is an ELI-by-period cell. Standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A20: Alternative IV Specifications with the Change in Import Penetration in Other Developed Economies

	U.S. CPI Inflation			
	(1)	(2)	(3)	(4)
Δ China IP	-1.43^{***}	-1.43^{**}	-1.58^{***}	-1.78^{***}
	(0.34)	(0.61)	(0.48)	(0.65)
First-stage F	357.70	14.84	21.50	48.07
Major Category F.E.	✓			
ELI F.E.		✓	✓	
Period-specific Goods F.E.	✓	✓	✓	✓
Durables & Apparel Time-Varying F.E.		✓		
Excluding Deflationary ELIs			✓	
6-digit IO industry F.E.				✓
Instrument: Δ China IP Other	✓	✓	✓	✓
N	444	444	444	170

Notes: The specifications reported in this table are described in Section III.C. They are identical to Panel A of Figure 3 in the main text, except that we use the change in import penetration from China in other developed economies as the instrument, instead of the NTR gap. The level of observation is an ELI-by-period cell and the sample includes all ELIs from 1991 to 2007, with variables averaged over two periods, 1991-1999 and 2000-2007. Column (4) of Panel A is an exception: the data is aggregated from ELIs to 6-digit industries defined in the BEA's IO table. Consumption weights are used. Standard errors are clustered by ELIs or 6-digit IO industries. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A21: IV Results with Purchaser vs. Producer Pricers

	U.S. CPI Inflation	
	IV	IV
	(1)	(2)
Δ China IP, Producer Prices	-2.44^{***}	
	(0.431)	
Δ China IP, Purchaser Prices		-4.37^{***}
		(0.852)
First-stage F	111.71	31.56
Hansen J	0.881	0.459
Instruments: NTR Gap & Δ China IP Other	✓	✓
N	170	170

Notes: The specifications reported in this table are described in Section III.C, except that the data is aggregated from ELIs to 6-digit industries defined in the BEA's IO table. Column (1) uses the baseline definition for the change in import penetration rate from China (defined in Section II.A). Column (2) adjusts this definition by accounting for distribution margins. Distribution margins correspond to the costs associated with transportation and retail, which inflate the denominator in the definition of China IP in equation (1) in the main text. For each 6-digit IO industry, we estimate distribution margins as the ratio of purchaser prices to producer prices observed in the BEA's 2007 IO table. When accounting for distribution margins, the change in the import penetration rate from China decreases, and accordingly the IV coefficient is larger in Column (2) than in Column (1). These IV specifications use both instruments jointly (the NTR gap and the change in import penetration from China in other developed economies). The Hansen J statistics indicate that we cannot reject the over-identification restrictions. The results are similar when using the 1992 IO table, where the available industries are more aggregated (not reported). The level of observation is a 6-digit IO industry-by-period cell. Standard errors are clustered by 6-digit IO industries. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A22: IV Estimates for Continued Goods in Balanced Sample

	2000-2003 CPI Infl. for Contined Goods		2000-2005 CPI Infl. for Contined Goods	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Δ China IP	-1.98^{***}	-2.63^{***}	-2.00^{***}	-2.66^{***}
	(0.71)	(0.98)	(0.72)	(1.01)
First-stage F		215.40		214.06
2000-2007 only	✓	✓	✓	✓
Goods, Durables & Apparel F.E.	✓	✓	✓	✓
N	222	222	222	222

Notes: This table reports OLS and IV estimates using inflation for continued products as the outcome variable. We consider a balanced sample of products that are continuously available from 2000 to 2003 (Columns (1) and (2)) or from 2000 to 2005 (Columns (3) and (4)). The NTR gap is used as an instrument. The price effects are not sensitive to the period we choose for the balanced sample. The magnitudes are similar to the estimates in Panel A of Table 6, which indicates that changes in composition do not drive our results for continued products. This result can help discipline quantitative trade models, because it shows that "reallocation effects" (entry or exit of more/less productive products of firms in response to trade shocks) are not the leading force in the data. Instead, there is a large response of pre-existing varieties (continued products inflation). The level of observation is an ELI-by-period and the standard errors are clustered by ELIs. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A23: The Effect of Trade with China on Product Turnover

	Product Turnover (pp)	
	(1)	(2)
Δ China IP (pp)	1.72	1.45^{***}
	(1.533)	(0.403)
Instrument:		
NTR gap	✓	
Δ China IP Other		✓
N	444	444

Notes: This table investigates the impact of trade with China on product turnover. Product turnover is measured as the rate of "product substitutions" in the BLS data. Product substitutions occur when price collectors can no longer find the product they were pricing in a given store (for instance, this could happen because this product was displaced by foreign competition). The table shows that product turnover increases substantially in response to trade with China, consistent with the notion that Chinese products displace domestic varieties. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A24: The Role of Higher-Order Input-Output Linkages

Panel A: Direct and Indirect Higher-Order Exposure to Trade with China

	Δ China IP Supp, Higher-order IO			Δ China IP Buyer, Higher-order IO		
	(1)	(2)	(3)	(3)	(4)	(6)
Δ China IP (pp)	0.0567^{***}	0.0394^{***}	0.03607^{***}	0.0163^{***}	0.01127	-0.0013
	(0.0077)	(0.0097)	(0.0072)	(0.0061)	(0.0098)	(0.00149)
6-digit IO F.E.		✓	✓		✓	✓
Period-specific Goods F.E.		✓	✓		✓	✓
Excl. diagonal of IO Table			✓			✓
N	170	170	170	170	170	170

Panel B: IV Results Controlling for Indirect Higher-Order Exposure to Trade with China

	U.S. CPI Inflation		
	(1)	(2)	(3)
Δ China IP	-3.143^{**}	-2.831^{**}	-3.196^{**}
	(1.451)	(1.383)	(1.515)
First-stage F	7.110	8.497	6.321
Controls:			
Δ China IP Supplier. Full IO	✓		✓
Δ China IP Buyer. Full IO		✓	✓
ELI F.E.	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓
N	170	170	170

Notes: The sample and specification are the same as in Panel A of Table 7 in the main text, except that the IO-adjusted measures including higher-order IO linkages instead of first-order linkages only. The level of observation is a 6-digit IO industry-by-period cell. The instrument is the NTR gap. Standard errors are clustered by IO industries. *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A25: IV Estimates for Input-Output Effects

	U.S. CPI Inflation			
	First-order I-O Linkages		Higher-order I-O Linkages	
	(1)	(2)	(3)	(4)
Δ China IP	-1.454^{***}	-1.458^{***}	-1.441^{***}	-1.478^{***}
	(0.402)	(0.408)	(0.389)	(0.399)
Δ China IP Supplier	-4.602	-5.552	-8.688	-9.208^{*}
	(3.821)	(3.633)	(5.348)	(5.310)
Δ China IP Buyer	-4.868	-8.102	-0.383	-1.489
	(9.348)	(10.614)	(9.942)	(11.373)
First-stage F	45.773	40.567	45.773	40.567
Period-specific Goods F.E.	✓	✓	✓	✓
Tech-by-period F.E.	✓	✓	✓	✓
IO2 F.E.		✓		✓
N	170	170	170	170

Notes: This table reports IV estimates with three endogenous variable: direct trade exposure, indirect exposure via intermediate inputs ("supplier effect") and indirect exposure via domestic buyer industries ("buyer effect"). Columns (1) and (2) use first-order IO linkages only, while columns (3) and (4) use higher-order IO linkages. The supplier and buyer effects across specifications are computed as explained in Online Appendix B.D. We include alternative sets of fixed effects across specifications. With 6-digit IO fixed effects, the IV becomes weak (not reported). *** denotes statistical significance at the 1% level, ** at the 5% level.

Table A26: Markup Quantiles and Profitability Ratios

	U.S. Markups by Quantiles (pp)			Profitability	
	p90	p50	p10	Profits/Sales	Profits/Assets
	(1)	(2)	(3)	(4)	(5)
NTR Gap	-17.42^{**}	-7.97^{*}	-0.84	-7.47^{***}	-4.91^{***}
	(7.28)	(4.83)	(4.023)	(2.66)	(2.23)
NAICS F.E.	✓	✓	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓	✓	✓
N	796	796	796	796	796

Notes: The level of observation is a 6-digit NAICS-by-period cell. Standard errors are clustered by 6-digit NAICS industries.*** denotes statistical significance at the 1% level, ** at the 5% level, * at the 10% level. Source: Compustat North America Fundamentals Annual Data, (Wharton Research Data Services) and authors' calculations.

Table A27: Heterogeneity by Market Structure in CPI Sample

	U.S. CPI Inflation (pp)				
	Interacted Specs.	Subsample Specs.			
	(1)	(2)	(3)	(4)	(5)
Δ China IP (pp)	-0.70	-1.58^{***}	-0.34	-0.77^{**}	-1.61^{***}
	(0.53)	(0.40)	(0.48)	(0.35)	(0.33)
Δ China IP$\times$High Concentration	-1.29^{**}				
	(0.53)				
Δ China IP$\times$High China IP	1.50^{**}				
	(0.60)				
First-stage F	5.76	28.77	8.26	2.64	34.23
ELI F.E.	✓	✓	✓	✓	✓
Period-specific Goods F.E.	✓	✓	✓	✓	✓
Subsample	All	High Conc.	Low Conc.	High China IP	Low China IP

Notes: "High Concentration" product categories have a level of domestic market concentration above median in 1997 (resp. below for "Low Concentration"). "High China IP" product categories have an import penetration rate from China above the 75th percentile in 1999 (resp. below for "Low China IP"). The level of observation is an ELI-by-period cell. Standard errors are clustered by ELIs.*** denotes statistical significance at the 1% level, ** at the 5% level, * at the 10% level.