

The Distributional Effects of Trade: Theory and Evidence from the United States*

Kirill Borusyak
Harvard University

Xavier Jaravel
London School of Economics

January 14, 2018

[Please Click Here for the Latest Version](#)

Abstract

Are the gains from trade unequally distributed in society? This paper presents new evidence on the distributional effects of trade on education groups in the U.S. through both consumer prices (*expenditure channel*) and wages (*earnings channel*). Our analysis, guided by a simple quantitative trade model, leverages linked datasets that cover the entire U.S. economy and include detailed spending data on consumer packaged goods and automobiles. First, we show that the expenditure channel is distributionally neutral due to offsetting forces. College graduates spend more on services, which are largely non-traded; however, their spending on goods is skewed towards industries, firms, and brands with higher import content. Second, on the earnings side, we find that college graduates work in industries that (1) are less exposed to import competition, (2) export more, (3) are more income-elastic, and (4) use fewer imported inputs. The first three forces cause trade liberalizations to favor college graduates; the fourth has the opposite effect. Finally, we combine and quantify the expenditure and earnings channels using the model. A 10% reduction of all import and export barriers generates a modest increase in inequality between education groups, primarily due to the earnings channel. Welfare gains are 16% higher for college graduates, whose real income increases by 2.02% compared to 1.74% for individuals without a college degree. Reductions of import barriers with China have qualitatively similar implications.

Keywords: Trade liberalization, Inequality, Non-homothetic preferences, Skill premium, Scanner data
JEL Codes: F14, F16, F60, D63

*Borusyak: borusyak@fas.harvard.edu; Jaravel: x.jaravel@lse.ac.uk.

We are particularly indebted to Pol Antràs, Ed Glaeser, Gita Gopinath, Elhanan Helpman, Oleg Itskhoki, Larry Katz, Danial Lashkari, Marc Melitz, and Oren Ziv. We are also grateful to seminar participants at Harvard, SIEPR, and the EIIT conference for insightful comments and conversations. We thank Jim Davis for his help at the Boston Census Research Data Center and James O'Brien for sharing a concordance between CEX categories and NAICS industries with us. The research in this paper was conducted while Borusyak was a Special Sworn Status researcher of the U.S. Census Bureau at the Boston Census Research Data Center. Any opinions and conclusions expressed herein are those of the authors and do not necessarily reflect the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Part of the results in the paper are calculated based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

Are the gains from trade, and the losses from protectionism, unequally distributed in society? Despite extensive research, the answer to this question remains debated. In the United States, policymakers on both sides of the aisle have recently proposed increasing import tariffs with major trading partners amid growing concerns over the impact of trade on inequality.¹ Both canonical and more recent trade theories predict that trade should negatively impact the earnings of low-skilled U.S. workers.² However, an emerging line of research suggests that the benefits of trade from falling consumer prices may disproportionately accrue to low-income populations (Fajgelbaum and Khandelwal, 2016). The net effect on inequality is therefore ambiguous.

This paper provides new evidence on the distributional effects of trade through both consumer prices (expenditure channel) and wages (earnings channel), and thus on the *net* distributional effects. Our analysis is based on linked datasets that cover the consumption and production sides of the entire U.S. economy and include expenditure microdata on consumer packaged goods and automobiles merged with restricted access customs data. To preserve tractability of the labor market analysis, we focus on the effects on two groups of Americans: those with and without a college degree. We show that the expenditure channel is distributionally neutral, while the earnings channel moderately favors college graduates. A set of intuitive reduced-form statistics—moments of the data that capture the differential exposure of education groups to trade—plays a central role in our analysis. The key statistic that governs the expenditure channel is the differential share of spending on imports, both directly via purchases of imported products and indirectly via imported intermediate inputs. The group that spends relatively more on imports enjoys greater purchasing-power benefits of trade liberalization. Similarly, the earnings channel depends on labor market differences in the exposure to trade; specifically, differences in import penetration, export shares, usage of imported intermediate inputs, and income elasticities between the industries that employ workers of different education.

These statistics emerge from a simple quantitative trade model. We build this model with three goals: to guide the reduced-form measurement, to perform the counterfactual analysis of trade policies, and to isolate the contributions of different mechanisms to the distributional effects. The counterfactuals necessarily rely on the model assumptions, which we discuss in detail below.

The first part of the paper compares spending on imports (both direct and indirect) between education groups using three new datasets that integrate expenditure microdata with importing statistics. First, we match spending categories in the Consumer Expenditure Survey (CEX) to detailed goods and service

¹Donald Trump proposed a 45% tariff on imports from China during his campaign and a 20% tariff on imports from Mexico in the first days of his presidency (Haberman, 2016; Porter, 2017). Concurrently, Senate Democrats proposed protectionist measures against China, in an attempt to “outdo Trump on trade” (Appelbaum, 2017). This pushback against free trade stems from the perception that a host of social ills, including growing inequality, are due to the globalized marketplace (Bohlen, 2016).

²In the Heckscher-Ohlin model with two sectors and two factors of production, the Stolper and Samuelson (1941) theorem implies that opening up to trade reduces the wage of the relatively scarce factor. Since the U.S. is skill-abundant, wages of low-skilled workers would fall. See Burstein and Vogel (2017), Caron et al. (2017), and Cravino and Sotelo (2017), among others, for different mechanisms.

industries in the national accounts and trade statistics from the Bureau of Economic Analysis (BEA). The merged dataset covers the universe of spending, accounting for trade in intermediate goods and in services. Second, for consumer packaged goods, we match products (barcodes) from the Nielsen Homescan Consumer Panel to their manufacturers or distributors in the confidential U.S. Economic Census and Customs microdata.³ We proxy for product import content by the share of imports in sales of the corresponding firm, and for the country of origin by the source of these imports.⁴ Finally, we link automobile brands from the CEX questionnaire on car purchases to Ward’s Automotive statistics on U.S. imports of assembled cars and to the Census of Manufactures and Customs data to account for imported car parts. Consumer packaged goods and automobiles, for which our data are particularly detailed, cover around 40% of total expenditure on goods.

We show that the expenditure channel is distributionally neutral as a result of offsetting forces. On the one hand, college graduates consume a larger share of services, which are largely non-traded. On the other hand, spending of college graduates on goods is skewed toward industries with higher import penetration rates, such as electronics relative to food. Moreover, within consumer packaged goods and automobiles, purchases of college graduates are skewed toward imported brands, particularly those that come from developed countries and tend to be more expensive. These forces largely compensate each other, resulting in similar overall spending shares on imports across education groups.

The patterns are similar when we compare import spending across income groups. Our results stand in contrast with [Fajgelbaum and Khandelwal \(2016\)](#) who infer consumption baskets of different groups based on aggregate international trade flows and a structural model. While their model has attractive aggregation properties and can be estimated without detailed data, it predicts substantially larger import spending by low-income groups, which we do not find in the U.S. data.⁵

We also document differences in spending on imports from specific countries of origin, in particular from China. Consistent with our prior, Chinese brands of consumer packaged goods are less expensive and tend to be purchased more often by individuals without a college degree. However, imports from China are concentrated in industries with a higher expenditure share by college graduates (e.g. electronics), providing an offsetting force.

The second part of the paper presents the reduced-form patterns governing the earnings channel. We find that college graduates work in industries that (a) are less exposed to import competition (overall and from China), (b) export more, (c) sell income-elastic products, and (d) have a lower share of imported intermediate inputs. Viewed through the lens of the model, trade liberalization increases the college wage premium through the first three effects and decreases it through the fourth. Indeed, import-competing

³Our data cover food, beverages, household supplies, health and beauty products, and other supermarket items.

⁴With this strategy, we capture imports of both final products and intermediate inputs. The bridge is constructed at the firm level (by name and address), so barcodes that belong to the same firm are assigned the same import intensity. We address the attenuation bias that may result from such aggregation.

⁵Their estimates imply that the spending share on imports is around 8 *percentage points* higher for consumers at the 75th percentile of income distribution relative to the 25th percentile, on average across countries (Figure XII). While comparable figures are not available for the U.S., they are likely to be as strong, based on the large pro-poor gains relative to autarky reported in Table V of their paper.

industries shrink due to trade, which reduces labor demand for the skill group they are intensive in. Furthermore, trade induces an expansion of industries which export, are income-elastic or use imported inputs; these industry expansions affect relative labor demand accordingly.⁶

The finding that college graduates are less exposed to import competition may appear consistent with the familiar Heckscher-Ohlin prediction: as a skill-abundant country, the U.S. should import in low skill-intensive industries. However, this interpretation is largely misguided: we document that the differential exposure is caused by service industries, which hire more college graduates and are subject to less import competition. Offsetting this effect, college-educated workers are employed in goods-producing industries that are subject to *more* import competition.

The third part of the paper performs counterfactual analyses, based on a model of a small open economy similar to Caron et al. (2017), Cravino and Sotelo (2017), and Morrow and Trefler (2017). The economy is populated by skilled and unskilled agents, whose preferences over composite goods of different industries belong to a flexible class that combines non-homotheticities as in Comin et al. (2016) with a nested CES structure. Agents inelastically supply a unit of labor and are freely mobile across industries. Industries, which include goods and services, supply varieties that are differentiated by country of origin, in a way that generates a standard gravity equation. Markets for each variety are perfectly competitive, and production technologies exhibit constant returns to scale. Trade flows in the model are shaped by product differentiation and arbitrary differences in technology and skill endowments. We perturb the equilibrium observed in the data with a counterfactual trade shock and characterize the impact on prices and wages in comparative statics using a first-order log-linear approximation.⁷

The virtue of the model is its tractability: general equilibrium counterfactuals can be computed based on the simple reduced-form statistics documented in the previous parts of the analysis.⁸ Specifically, the small open economy assumption implies that trade policy shocks do not affect foreign factor prices; hence U.S. data are sufficient for the analysis. The assumption of perfect competition in turn generates complete pass-through of trade costs into prices, which implies that differential spending on imports governs the expenditure channel in the log-linear approximation. On the earnings side, the class of preferences we employ simultaneously allows for income effects of trade and rich but tractable import competition effects. Finally, free mobility of workers conveniently reduces inequality to a single dimension—education. Mobility costs are known to be important for the transitory effects of trade (e.g. Artuç et al., 2010, Autor et al., 2014, Traiberman, 2016, and Galle et al., 2017), but should play a smaller role in the long-run.⁹

⁶Income-elastic industries expand because trade generates additional real income for domestic consumers, which is disproportionately spent on income-elastic products.

⁷Our analysis differs from the sufficient statistic approach of Arkolakis et al. (2012): they evaluate unobserved historical shocks, while we focus on pre-specified counterfactual trade policies that perturb the current equilibrium. Accordingly, *changes* in equilibrium objects over time constitute the data for them, while we rely on the *snapshot* of the equilibrium.

⁸In addition, we need standard structural elasticities, namely trade elasticities and elasticities of substitution across the product space on the demand side, as well as the aggregate elasticity of substitution between workers with and without a college degree on the supply side.

⁹Even in the short run, our estimates of the between-group inequality need not be biased: while low-skilled manufacturing workers suffer more in presence of mobility costs, low-skilled non-manufacturing workers are better insulated from trade shocks, creating a countervailing force.

Using the model, we show that the distributional effects of trade moderately favor college graduates, mostly through the earnings channel. A 10% reduction in trade barriers with all U.S. trading partners generates welfare gains that are positive for both groups but 16% higher for college graduates (2.02% vs. 1.74%, measured a fraction of their consumption).¹⁰ Differences in import penetration, export shares, and income elasticity contribute similar amounts to the earnings channel. Imported intermediate inputs mildly reduce it and general equilibrium forces slightly strengthen it. The expenditure channel does not offset the earnings channel; instead, it is also biased in favor of college graduates but is small in magnitude. We also consider a 10% reduction in prices of imports from China specifically and find qualitatively similar effects.¹¹

In sum, this paper makes three contributions to the literature on the distributional effects from trade. Our main contribution is to show that the expenditure channel is distributionally neutral in the United States. Additionally, we quantify the relative importance of various mechanisms contributing to the earnings channel, which favors college graduates. Finally, we develop a simple framework, in which a set of reduced-form statistics governs the distributional effects from trade policies and which can be applied to other countries and time periods.

Related Literature. This paper contributes to the growing literature on the distributional effects of trade through the expenditure channel. A few papers use international trade data to predict spending on imports for different income groups through the structure of the demand system: doing so, [Fajgelbaum and Khandelwal \(2016\)](#) and [He and Zhang \(2017\)](#) found strong pro-poor effects of the expenditure channel for all countries, while the estimates of [Nigai \(2016\)](#) are pro-rich. In contrast, the estimates reported in this paper are based on direct observation of consumption baskets for both domestic and imported products and therefore require minimal structural assumptions to characterize the magnitude of the expenditure channel. Furthermore, a small number of papers directly measure spending on imports and compare them across consumer groups: [Porto \(2006\)](#) for Argentina, [Faber \(2014\)](#) for Mexico, and [Levell et al. \(2017\)](#), [Dhingra et al. \(2017\)](#), and [Breinlich et al. \(2017\)](#) for the U.K. Data limitations force these papers to focus only on particular types of differential spending.¹² In contrast, our paper considers the entire economy, taking into account imports of both final and intermediate goods, and at the same time uses very detailed data on consumer packaged goods and automobiles to address potential aggregation bias.¹³

¹⁰The shock we consider can be interpreted as a 10 percentage point reduction in ad valorem tariffs, ignoring the extra tariff revenue, or as a reduction in transportation costs which makes imported and exported goods 10% cheaper.

¹¹For this shock, the net gains are 38% higher for college graduates, but small in absolute value for both groups (0.197% vs. 0.142% of consumption).

¹²[Porto \(2006\)](#) captures differences in spending across seven large categories of final goods and services, [Faber \(2014\)](#) looks at imported intermediate inputs, [Levell et al. \(2017\)](#) limit their analysis to 9 categories of food, and [Dhingra et al. \(2017\)](#) and [Breinlich et al. \(2017\)](#) consider 12–13 broad groups of goods and services consumed by households. The paper by [Furman et al. \(2017\)](#) is also related: they merge the CEX consumption data by group with import shares for 14 categories of spending but, focusing on the incidence of tariffs, do not report differential import spending.

¹³Our work is also related to papers that measure the impact of trade and exchange rates on prices indices. [Bai and Stumpner \(2017\)](#) and [Broda and Romalis \(2008\)](#) examine distributional effects at the level of product categories, while we extend the analysis to the level of firms within categories. [Cravino and Levchenko \(2017\)](#) and [Hottman and Monarch \(2017\)](#) capture within-category patterns too, but they rely on structural assumptions due to data limitations. [Amiti et al. \(2017\)](#) investigate the effect of China’s WTO entry on the U.S. manufacturing price index, but they do not document distributional

The relationship between our work and the extensive literature on the distributional effects of trade through the earnings channel is twofold. First, the modeling framework allows us to assess the relative importance of the key mechanisms studied in the structural literature on the earnings channel separately. Beyond the role of skill endowment emphasized by the Stolper-Samuelson theorem, more recent papers focus on the contributions of non-homothetic preferences (Caron et al., 2017), complementarity between goods and services (Cravino and Sotelo, 2017), and the skill bias of exporters (Burstein and Vogel, 2017). Second, our findings are consistent with the existing empirical literature. Tests of the Heckscher-Ohlin model have documented that in the United States the skill content of net imports is small, implying that trade does not generate large changes in inequality through the earnings channel, in line with our findings based on a more flexible model. Moreover, Autor et al. (2013) show that trade with China induced a significant fall in U.S. manufacturing employment in the 1990s and 2000s, which is consistent with our calibration.¹⁴

Finally, this paper contributes to the emerging literature on the joint analysis of the expenditure and earnings channels. To the best of our knowledge, there are only two papers in this space: Porto (2006) uses time-series regressions to estimate the impact of trade-induced price changes on wages and domestic prices, while He and Zhang (2017) generalize the structural model of Fajgelbaum and Khandelwal (2016). As previously discussed, we take a different approach by focusing on a set of reduced-form statistics measured in detailed data.

The remainder of the paper is organized as follows. Section 2 presents the model. Sections 3–5 report estimates of spending on imports: Section 3 focuses on patterns across industries covering the entire U.S. economy, Section 4 provides estimates using scanner data for consumer packaged goods, and Section 5 presents the patterns using detailed expenditure data on automobiles. Section 6 reports the reduced-form patterns on import competition, exports, income elasticities and the use of imported intermediate inputs, which together govern the earnings channel. Finally, Section 7 presents the estimates of the distributional effects from trade policies, feeding the reduced-form patterns from the previous sections through the model. Section 8 concludes.

2 Theory

We develop a model to characterize the welfare consequences of counterfactual changes in trade costs across skill groups.¹⁵ This section first presents the model in a special case without input-output linkages and then describes the full model used for counterfactual analyses.

effects.

¹⁴Several approaches in the literature reached the same general conclusion that globalization was not an important cause of the rising skill premium in the U.S. in the 1980s (e.g. Borjas et al., 1997, Krugman, 2000, Lawrence and Slaughter, 1993, and Berman et al., 1994). Autor et al. (2013) estimate the negative impact of trade on U.S. employment at the level of commuting zones but remain largely silent about the effect of trade on inequality, because they do not document the distribution of trade shocks across commuting zones.

¹⁵Appendix A provides the proofs. To facilitate reading, Table A1 provides a catalog of variables used.

2.1 Setup

Trade, Preferences, and Technologies. We study a static global economy with $\mathcal{C} + 1$ countries in which international trade is shaped by product differentiation, cross-country differences in technologies and endowments, and trade costs. The United States is denoted $c = H$ (Home) and the set of all other countries F (Foreign). The home economy is assumed to be sufficiently small, such that shocks to trade costs between Home and any other countries do not affect foreign prices for goods and factors.¹⁶

The home economy is populated by two types of agents: skilled ($i = S$) and unskilled ($i = U$), with measures L_i . They derive utility $\mathcal{U}(Q_1^i, \dots, Q_{\mathcal{J}}^i)$ from consuming composite products of $\mathcal{J}$ industries, which include both goods and services. They spend $X_j^i = p_j Q_j^i$ on the industry j products, which constitutes a share $s_j^i = X_j^i / X_i$ of their total spending $X_i = \sum_j X_j^i$. Agents all inelastically supply one unit of labor. We assume that labor is freely mobile across industries (but not across countries), implying that domestic wages w_i differ only across skill levels, not within. Prices and wages are measured relative to a foreign numeraire. We allow the budget to be imbalanced to account for the large trade deficit in the U.S., assuming that the total expenditure of each agent is a multiple of her wage, $X_i = \zeta w_i$, with an exogenous constant $\zeta > 1$.¹⁷

To model domestic preferences, we use a flexible demand system that embeds non-homothetic CES preferences from [Hanoch \(1975\)](#) in a two-tier nesting structure. This demand system captures complementarity between goods and services (the upper tier) and substitution between industries (the lower tier) within each of the two sectors. It inherits the desirable property of non-homothetic CES that income and price elasticities are shaped by independent parameters (see [Comin et al., 2016](#); [Matsuyama, 2017](#)). These features of the demand system allow us to jointly accommodate and compare several mechanisms that have previous been examined in isolation.¹⁸ Non-homothetic nested CES utility is defined recursively by

$$\begin{aligned} \mathcal{U}_i &= \left(\sum_r (Q_r^i)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)}, \quad r = \text{Goods, Services} \\ Q_r^i &= \left(\sum_{j \in r} \left(a_j \mathcal{U}_i^{\varphi_j - 1} \right)^{1/\varepsilon_r} (Q_j^i)^{(\varepsilon_r - 1)/\varepsilon_r} \right)^{\varepsilon_r/(\varepsilon_r - 1)}, \end{aligned} \quad (1)$$

where ρ and ε_r are elasticities of substitution between and within sectors, respectively, and primitive

¹⁶According to the World Development Indicators database, exports *from* the U.S. constitute only 3.9% of absorption in other countries while exports *to* the U.S. account for only 5.5% of foreign production. In our model with product differentiation, these patterns suggest that the impact of U.S. trade cost shocks on prices outside the U.S. is likely to be limited.

¹⁷U.S. imports were 47% higher than exports in 2007, therefore assuming balanced trade would be counterfactual. Our assumption that the ratio of expenditures to income is constant is still imperfect: net imports as a fraction of GDP ($\zeta - 1$ in the model) have fluctuated historically from 1.1% in 1997 to 5.6% in 2005 and back to 2.9% in 2015 (see [Figure A1](#)). Our approach differs from a more common assumption stating that the absolute value of net imports is fixed ([Dekle et al., 2008](#)), but it is more tractable in a model with multiple factors because we do not need to keep track of income and expenditure changes separately.

¹⁸See [Cravino and Sotelo \(2017\)](#) on the role of complementarity between goods and services and [Caron et al. \(2017\)](#) on non-homotheticities. We also accommodate Stolper-Samuelson type forces, as well as the role of intermediate inputs, to which we return in [Section 2.3](#).

parameters $\{\varphi_j\}$ determine the income elasticity of each industry, which we denote by $\{\psi_j\}$.¹⁹

Each industry j consists of a set of differentiated varieties sold in perfectly competitive markets. The structure of product differentiation we impose gives rise to a standard industry gravity equation. For brevity of notation, we use the [Armington \(1969\)](#) formulation of the product space: there is one variety per country and industry, and varieties are combined with constant elasticity of substitution ξ_j , so that $Q_j = \left(\sum_c b_{jc}^{1/\xi_j} Q_{jc}^{(\xi_j-1)/\xi_j} \right)^{\xi_j/(\xi_j-1)}$, where b_{jc} are taste shifters. Accordingly, the industry price index is $p_j = \left(\sum_c b_{jc} p_{jc}^{1-\xi_j} \right)^{1/(1-\xi_j)}$, where prices of imported products p_{jc} are inclusive of iceberg trade costs τ_{jc} . The microfoundation from [Eaton and Kortum \(2002\)](#) is isomorphic to this setup.²⁰ The results are also unchanged if the number of varieties varies by country and industry, as long as there is no entry or exit.

Domestic production combines labor inputs with materials from various industries, and there are no other factors of production. Output is given by $Q_{jH} = F_{jH} \left(L_S^j, L_U^j, M^j \right)$, where F_{jH} is some constant returns to scale function, L_i^j is type- i labor employed in j , and M^j are materials purchased from other industries.²¹ Mirroring spending shares, $e_i^j = w_i L_i^j / w_i L_i$ denotes the equilibrium share of group i 's earnings that industry j accounts for.

We remain agnostic about foreign endowments, preferences, and technologies, requiring only that foreign buyers aggregate varieties across countries of origin with the same elasticity ξ_j as domestic buyers. Since Home is small and does not affect industry price indices abroad, there are no cross-price effects, and export demand elasticity for domestic products is also ξ_j . Therefore, the quantity of exports satisfies $Q_{jH}^{\text{Export}} = a_j^{\text{Export}} \left(p_{jH} \tau_j^* \right)^{-\xi_j}$, where τ_j^* is the exporting iceberg trade cost and a_j^{Export} are constants.²²

Counterfactuals. The equilibrium is defined by (i) the demand relationships for domestic consumers of each skill type and for foreign consumers, (ii) zero profit conditions for domestic and foreign producers, and (iii) labor and product market clearing conditions (see [Appendix A.1](#)). To characterize how the equilibrium responds to counterfactual trade shocks, we rely on a first-order log-linear approximation around the equilibrium, which is precise when the trade shocks are small.²³ Formally, we use the “hat algebra” of [Jones \(1965\)](#): we consider a set of small price and wage shocks $\hat{p}_j$ and $\hat{w}_i$ induced by trade liberalization, where hats denote relative changes from the original equilibrium.

The distributional effects of a change in trade policy depend on the set of industries and of trading partners that are affected by the policy change. Considering trade policies skewed towards specific indus-

¹⁹Preferences reduce to homothetic nested CES when $\varphi_j \equiv 1$, in which case demand shifters do not depend on the utility level, and to non-homothetic CES from [Hanoch \(1975\)](#) when $\varepsilon_r \equiv \sigma$. See equation (A22) in the Appendix for the general expression for ψ_j . While income elasticities theoretically must depend on income, we will ignore the differences in ψ_j across types and equilibria, viewing them as fixed industry characteristics, as in [Aguiar and Bils \(2015\)](#).

²⁰In the Eaton and Kortum version, the Fréchet parameter plays the role of the trade elasticity $\xi_j - 1$.

²¹We use superscripts (subscripts) to indicate buyers (sellers). Agents are buyers in the product markets and sellers in the labor market.

²²This is the quantity delivered to foreign consumers, at price $p_{jH} \tau_j^*$ per unit.

²³It is instructive to draw a distinction between our counterfactuals and the sufficient statistic approach of [Arkolakis et al. \(2012\)](#). We study the effects of a pre-specified counterfactual shock to trade barriers, while they evaluate the impact of *historical* changes in trade barriers. Accordingly, our results are based on detailed information about the current equilibrium, at which the counterfactual policy is implemented, instead of the *changes* between two equilibria in different years, as in their paper.

tries or trading partners, one could plausibly generate any result. To discipline the analysis, we consider two policies that are uniform across industries: (i) a bilateral reduction of tariffs with all foreign countries, and (ii) a unilateral reduction of tariffs on Chinese imports. The first policy reduces consumer prices for all imported goods at home and exported goods abroad, so we interpret its effects as the overall distributional effects of international trade. The second policy is motivated by the large increase in imports from China in recent years and by the evidence for its sizable employment effects (Autor et al., 2013, 2014; Caliendo et al., 2017). These policies can be thought of as changes in tariffs, but reductions of iceberg transportation costs or of other value-based barriers are isomorphic.²⁴ We parameterize each policy as a combination of a U.S. import tariff change $\hat{\tau}$ that affects some set of foreign countries $\mathbf{c}$ and a tariff change $\hat{\tau}^*$ in all foreign countries for U.S. goods. The bilateral trade shock corresponds to $\mathbf{c} = F$ and $\hat{\tau}^* = \hat{\tau}$, whereas the China shock is parameterized by $\mathbf{c} = \text{China}$ and $\hat{\tau}^* = 0$. Negative values of $\hat{\tau}$ and $\hat{\tau}^*$ correspond to liberalizations.

Defining Welfare. To evaluate welfare changes in a way that is comparable across agents, we follow the standard approach in the literature by using a money metric for utility: $\hat{\mathcal{U}}_i$ is defined as the equivalent variation EV_i divided by original expenditures X_i .²⁵ For example, $\hat{\mathcal{U}}_i$ is equal to 0.01 if the trade liberalization is equivalent, in utility terms, to increasing total spending by 1% at the original prices.

For small shocks, the envelope theorem (Roy’s identity) implies that price reductions benefit each type of consumer in proportion to their spending share for this industry, regardless of the demand system:²⁶

$$\hat{\mathcal{U}}^i \equiv \frac{EV^i}{X^i} = \hat{w}_i - \sum_j s_j^i \hat{p}_j \equiv \hat{w}_i - \hat{\pi}_i, \quad (2)$$

where $\hat{\pi}_i$ is the Laspeyres price index.

From (2), differential gains between two skill groups can be decomposed into components related to changes in prices and wages, which we label the expenditure and earnings channels:

$$\hat{\mathcal{U}}^S - \hat{\mathcal{U}}^U = \underbrace{-(\hat{\pi}_S - \hat{\pi}_U)}_{\text{Expenditure channel}} + \underbrace{\hat{w}_S - \hat{w}_U}_{\text{Earnings channel}}. \quad (3)$$

Positive values indicate the effects that favor the skilled group (“pro-skilled”), and both channels are invariant to the choice of the numeraire. Average welfare gains admit a similar representation as the

²⁴Following most of the trade literature, we ignore revenues generated by tariffs. For a recent treatment on the distinction between tariffs and generic iceberg trade costs, see Felbermayr et al. (2015).

²⁵See Theil (1975); Deaton (1989); Fajgelbaum and Khandelwal (2016); Nigai (2016). Because we allow for trade imbalances, expenditures rather than income is the relevant denominator. With the money metric, we bypass the question of whether the unskilled value an additional dollar more than the skilled.

²⁶Intuitively, because consumers are optimizing, re-optimization of expenditure shares in response to small changes in relative prices has only a second-order effect on welfare. The first-order effect is the money the agent saves by paying lower prices, holding spending shares constant.

difference between average wage growth and inflation.²⁷ Using bars for economy-wide averages,

$$\hat{\mathcal{U}} = \hat{w} - \hat{\pi} \equiv [\bar{v}\hat{w}_S + (1 - \bar{v})\hat{w}_U] - \sum_j s_j^{\text{Final}} \hat{p}_j, \quad (4)$$

where $\bar{v}$ is the income share of the skilled group in the original equilibrium and $s_j^{\text{Final}} = \frac{X_j^S + X_j^U}{\sum_j (X_j^S + X_j^U)}$ is the share of j in total final spending.

2.2 Welfare Effects of Trade

We characterize the average and differential welfare effects of counterfactual trade shocks in two steps. First, we treat wage changes as given and solve for the price changes induced by the shock. Second, we use the labor market clearing condition to solve for wage changes. We focus on the special case without input-output linkages, assuming that production uses only labor but not materials, and return to the general case in the next section.

Step 1: Price Changes Conditional on Wage Changes. Deriving price changes is easy in our model because of perfect competition. First, markups are constant (zero), so prices are determined by marginal costs with a complete pass-through rate. Second, supply is perfectly elastic, so marginal costs are determined by factor prices and not directly affected by demand conditions.²⁸ As a consequence, import prices change one for one with the import tariffs. Similarly, by the envelope theorem (Shephard's lemma), domestic price changes $\hat{p}_{jH}$ are shaped by changes in domestic wages:

$$\begin{aligned} \hat{p}_{jH} &= v_j \hat{w}_S + (1 - v_j) \hat{w}_U \\ &= \hat{w} + (v_j - \bar{v}) (\hat{w}_S - \hat{w}_U), \end{aligned} \quad (5)$$

where $v_j = w_S L_j^S / VA_j$ is the share of value added paid to skilled workers (a model-consistent measure of skill intensity) and value added is defined as $VA_j = w_S L_j^S + w_U L_j^U$ (which coincides with payroll in our model without capital).

The consumer price index in each industry combines domestic and import price changes:

$$\hat{p}_j = IP_{jc} \hat{\tau} + (1 - IP_j) (\hat{w} + (v_j - \bar{v}) (\hat{w}_S - \hat{w}_U)), \quad (6)$$

where weight $IP_{jc} = \sum_{c \in \mathbf{c}} X_{jc} / X_j$ is the industry import penetration from the set of countries affected by the trade liberalization and $1 - IP_j = 1 - \sum_{c \in F} IP_{jc}$ is the domestic share. Plugging the price formulas

²⁷ $\hat{\mathcal{U}}$ measures the total equivalent variation relative to the total expenditure and does not require averaging cardinal utility levels.

²⁸See [Epifani and Gancia \(2008\)](#) on how economies of scale matter for the distributional effects of trade. Monopolistic competition with constant markups and free entry would manifest itself as scale effects (cf. [Costinot and Rodríguez-Clare, 2015](#)).

above into our expressions for the average gains from trade and the expenditure channel yields:

$$\hat{\mathcal{U}} = \mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}}] \cdot (-\hat{\tau}) + \mathbb{E}_{\text{Final}} [IP_j] \cdot \hat{w} - \mathbb{E}_{\text{Final}} [(1 - IP_j)(v_j - \bar{v})] \cdot (\hat{w}_S - \hat{w}_U), \quad (7)$$

$$-(\hat{\pi}_S - \hat{\pi}_U) = \Delta_{\text{Final}} [IP_{j\mathbf{c}}] \cdot (-\hat{\tau}) + \Delta_{\text{Final}} [IP_j] \cdot \hat{w} - \Delta_{\text{Final}} [(1 - IP_j)(v_j - \bar{v})] \cdot (\hat{w}_S - \hat{w}_U), \quad (8)$$

where $\mathbb{E}_{\text{Final}} [z_j] = \sum_j s_j^{\text{Final}} z_j$ is the cross-industry average of variable z_j with total final consumption weights, $IP_{j\mathbf{c}} = \sum_{c \in \mathbf{c}} IP_{jc}$ is the import penetration from countries affected by the import shock, and $\Delta_{\text{Final}} [z_j] = \sum_j s_j^S z_j - \sum_j s_j^U z_j$ denotes the difference between averages in the consumption baskets of the two groups.

Equations (7) and (8) show that the average welfare effect and the expenditure channel are governed by the same three forces. First, an import tariff reduction ($-\hat{\tau}$) directly makes the affected imported varieties cheaper, which benefits a consumer in proportion to her spending share on these imports. Second, an increase in domestic wages ($\hat{w}$) relative to foreign wages makes *all* imports more affordable;²⁹ this channel benefits a consumer in proportion to her overall spending on imports. Finally, an increase in the skill premium ($\hat{w}_S - \hat{w}_U$) hurts consumers of skill-intensive domestic varieties.³⁰ We highlight that through both the first and second forces, the differential spending share on imports is the key statistic for the expenditure channel.³¹

Step 2: Solving for Wage Changes. To solve for wage changes, we proceed in three steps. We first use the labor market equilibrium to relate wage changes to growth in value added across industries. Then we use the demand system to characterize growth of each industry as a function of tariff and wage changes. These steps produce a system of equations for wage changes, which we solve in the final step.

Step 2a: Connecting Wages to Industry Sizes. We start with a simple observation: total value added in all industries is by definition equal to the average wage times the labor supply (total number of workers). Labor supply is fixed, so the change in the average wage equals the change in total value added, which is the weighted average of industry growth rates with pre-shock value added as weights:

$$\hat{w} = \frac{\sum_j VA_j \cdot \widehat{VA}_j}{\sum_j VA_j} \equiv \mathbb{E}_{VA} [\widehat{VA}_j]. \quad (9)$$

²⁹This channel operates even if the shock only affects export tariffs: domestic wages will grow to reduce the trade surplus, which generates welfare gains.

³⁰This force generates distributional effects if the economy is “segregated”: each group of agents tends to consume from industries where they are predominantly employed. Then if the skill premium grows in response to trade shocks, the goods skilled individuals consume will become relatively more expensive, dissipating some of the labor market gains through the expenditure channel. The same force also generates aggregate gains or losses if total domestic demand is skewed towards high or low skill-intensive industries, compared with export demand.

³¹Note that even if pass-through is incomplete, as empirical evidence suggests (e.g. De Loecker and Warzynski, 2012; De Loecker et al., 2016; Arkolakis et al., 2017), but it is not systematically related to the industry consumer mix, then the differential spending on imports would remain the key determinant of the expenditure channel. $\Delta_{VA} [IP_j]$ shapes both the direct effects of falling import prices and the pro-competitive effects on markups, which are absent in our model (see Arkolakis et al., 2017), and the pass-through rate would only rescale those effects.

We obtain a similar representation for the change in the skill premium (see Appendix A.1 for the proof). Intuitively, if high skill-intensive industries expand faster than low skill-intensive ones, the relative demand for skilled labor grows. In equilibrium, labor supply is fixed, so this change in the relative demand has to be offset by a growing skill premium, which makes all industries switch from high-skilled to low-skilled workers. Formally,

$$\hat{w}_S - \hat{w}_U = \frac{\Delta_{VA} [\widehat{VA}_j]}{\sigma_{\text{within}}}, \quad (10)$$

where $\Delta_{VA} [\widehat{VA}_j] = \sum_j e_S^j \widehat{VA}_j - \sum_j e_U^j \widehat{VA}_j$ is the difference in the growth rates of industries where the skilled and unskilled work (using payroll weights), and σ_{within} is the elasticity that captures labor substitution within all industries. Appendix A.1 shows that the elasticity is given by $\sigma_{\text{within}} = 1 + \mathbb{E}_{VA} \left[\frac{v_j(1-v_j)}{v(1-v)} \cdot (\sigma_j - 1) \right]$, where σ_j is the local elasticity of substitution between labor types in domestic industry j .³²

Step 2b: Solving for Changes in Industry Sizes Conditional on Wage Changes. In Step 1, we solved for the price changes induced by trade shocks, conditionally on wage changes. We now map these price changes into changes in industry output using the demand system (of each skill group as well as foreign buyers); without input-output linkages, output equals value added. We express industry growth as a sum of four terms:

$$\widehat{VA}_j = \eta_{jc}^{\text{import}} \cdot \hat{\tau} + \eta_j^{\text{export}} \cdot (-\hat{\tau}^*) + \eta_j^{\text{avg wage}} \hat{w} - \eta_j^{\text{skill prem}} (\hat{w}_S - \hat{w}_U). \quad (11)$$

The first two terms capture the expansion of domestic production as import tariffs grow and exporting barriers fall, holding wages fixed. The other two terms represent the response of demand to wage changes, which affect both purchasing power and domestic prices. Log-linearizing demand, Appendix A.1 characterizes the corresponding elasticities:

$$\eta_{jc}^{\text{import}} = \text{Dom share}_j \cdot \left[(\xi_j - 1) \cdot IP_{jc} + (\varepsilon_r - 1) \cdot (\mathbb{E}_{\text{Final}} [IP_{jc} | r] - IP_{jc}) + (\rho - 1) \cdot (\mathbb{E}_{\text{Final}} [IP_{jc}] - \mathbb{E}_{\text{Final}} [IP_{jc} | r]) - (\psi_j - 1) \cdot \mathbb{E}_{\text{Final}} [IP_{jc}] \right], \quad (12a)$$

$$\eta_j^{\text{export}} = \text{Export share}_j \cdot (\xi_j - 1), \quad (12b)$$

$$\eta_j^{\text{avg wage}} = \text{Dom share}_j - \eta_{jF}^{\text{import}} - \eta_j^{\text{export}}, \quad (12c)$$

where $\text{Dom share}_j = 1 - \text{Export share}_j$ is the share of domestic consumers in the domestic industry output.³³ We now discuss these three expressions in turn.

³² σ_{within} is higher when skills are more substitutable within each industry (higher σ_j) and also when skill intensity is more homogenous across industries. Indeed, when technologies in each industry are such that one type of labor (sometimes skilled and sometimes unskilled) is much more productive, there is little room for adjusting industry labor mixes in response to wage shocks, and $\sigma_{\text{within}} \approx 1$. If instead all industries employ skilled and unskilled labor in the same proportions, σ_{within} is just the size-weighted average of σ_j , which is above one as long as skills are substitutes.

³³ To preserve space, the expression for $\eta_j^{\text{skill prem}}$ is given by (A8) in the Appendix. To facilitate exposition, two quantitatively negligible adjustment terms are dropped from η_j^{import} , as described in Appendix A.1 (see (A7)).

Falling domestic import tariffs ($\hat{\tau} < 0$) lower import prices, which drives the consumer price index down in proportion to import penetration. In our nested demand system, this price change generates cross-price effects for domestic varieties at each tier. First, demand is reallocated away from domestic varieties within each industry, for $\xi_j > 1$. Second, if $\varepsilon_r > 1$, spending is reallocated towards industries with more imports. Finally, since goods rely more on imports and become relatively cheaper, complementarity between goods and services ($\rho < 1$) makes consumers spend more on services. These *import competition effects* are captured by the first three terms in equation (12a). The last term in (12a) is the *income effect*: gains from trade, which depend on the average spending on imports $\mathbb{E}_{\text{Final}}[IP_j]$, lead to higher spending on income-elastic industries. All of these effects only influence domestic consumption, so they are scaled by the domestic share of industry sales.

Reductions in foreign import tariffs ($\hat{\tau}^*$) also matter. As foreign import tariffs fall, export demand grows according to the trade elasticity $\xi_j - 1$. This contributes to the output growth in proportion to the export share (equation (12b)).³⁴

Finally, changes in industry size depend on domestic wages. Growing average wage raises purchasing power and thus demand, captured by the first term in (12c). However, growing wages also raise domestic prices, making domestic varieties less attractive both at home and abroad. This reduces the demand for domestic varieties in the same way as falling import tariffs and growing exporting barriers would. This expenditure switching is captured by the last two terms in (12c).³⁵

Step 2c: Wage Changes in General Equilibrium. The last step of our analysis brings together the preceding formulas to account for feedback effects between price changes and wage changes in general equilibrium. Formally, equations (9), (10), (11) form a linear system. Solving it, we obtain:³⁶

$$\hat{w} = \left(\mathbb{E}_{VA} \left[\eta_j^{\text{import}} \right] \hat{\tau} + \mathbb{E}_{VA} \left[\eta_j^{\text{export}} \right] \cdot (-\hat{\tau}^*) \right) \cdot \text{Multiplier}, \quad (13a)$$

$$\hat{w}_S - \hat{w}_U = \left(\Delta_{VA} \left[\eta_j^{\text{import}} \right] \hat{\tau} + \Delta_{VA} \left[\eta_j^{\text{export}} \right] \cdot (-\hat{\tau}^*) + \Delta_{VA} \left[\eta_j^{\text{avg wage}} \right] \hat{w} \right) / \sigma_{\text{macro}}, \quad (13b)$$

where

$$\text{Multiplier} = 1 / \left(1 - \mathbb{E}_{VA} \left[\eta_j^{\text{avg wage}} \right] \right) \quad (14)$$

and $\sigma_{\text{macro}} = \sigma_{\text{within}} + \Delta_{VA} \left[\eta_j^{\text{skill prem}} \right]$.

Equation (13a) shows that a reduction of domestic (foreign) import tariffs decreases (increases) demand for domestic goods, which reduces (increases) domestic income. The direct effects of changes in domestic and foreign import tariffs on the average wage are magnified by a multiplier effect. This effect is conceptually similar to a local fiscal multiplier (e.g. Chodorow-Reich, 2017): if domestic income goes up

³⁴The small open economy assumption implies that the elasticity of demand equals the trade elasticity.

³⁵Changing skill premium also affects relative prices, as discussed in Step 1, and therefore demand. These effects are captured by $\eta_j^{\text{skill prem}}$.

³⁶To derive these expressions, we make the approximation $\mathbb{E}_{VA} \left[\eta_j^{\text{skill prem}} \right] \cdot (\hat{w}_S - \hat{w}_U) \approx 0$, ignoring the impact of the skill premium change on the average wage. We verify in the calibration that this impact is quantitatively negligible. Without this approximation, the solution to system becomes less transparent.

because of additional spending from foreign, domestic spending goes up and induces a feedback loop.³⁷

Equation (13b) shows that three effects are at play in the earnings channel: the skill premium goes up if skilled workers tend to work in industries that shrink less after the fall in domestic import tariffs ($\Delta_{VA} [\eta_j^{\text{import}}] < 0$), or that expand as a result of export opportunities ($\Delta_{VA} [\eta_j^{\text{export}}] > 0$), or that expand more as a result of a growing average wage ($\Delta_{VA} [\eta_j^{\text{avg wage}}] > 0$).³⁸

The effects entering the earnings channel are all scaled by the endogenous macro elasticity of labor substitution σ_{macro} . Intuitively, if skilled and unskilled workers are more substitutable, a smaller increase in the skill premium is sufficient to compensate the change in labor demand induced by trade and to restore the labor market equilibrium. σ_{macro} generalizes the macro elasticity of substitution between factors derived by Oberfield and Raval (2014) to the open economy and to a more flexible preference structure. Appendix A.1 provides further discussion of σ_{macro} .

Our theoretical results on the earnings channel can be connected to the data as follows. Plugging (12a)–(12c) into (13b), one can see that differential exposure of the two skill groups to import competition ($\Delta_{VA} [IP_{jc}]$), exports ($\Delta_{VA} [\text{Export share}_j]$) and income-elastic industries ($\Delta_{VA} [\psi_j]$) plays the key role in how the skill premium responds to trade shocks. While those are not sufficient statistics, e.g. if trade elasticities vary across industries, our reduced-form analysis of the earnings channel will focus on characterizing the extent to which college graduates work in industries that (1) are less exposed to import competition, (2) export more, (3) are more income elastic. We discuss the role of imported inputs in the next subsection.

2.3 Model with Input-Output Linkages

When measuring the differential import spending and labor market exposure, it is important to account for intermediate inputs, both because trade in intermediates is increasingly important (Feenstra and Hanson, 1996) and because value added, final consumption, and gross output are substantially different from each other in the data. We follow the literature by assuming that production combines value added with intermediate inputs, which are composite goods from various industries (e.g. Caliendo and Parro, 2015).³⁹

³⁷To understand how the multiplier works formally, consider the expression for $\eta_j^{\text{avg wage}}$. It implies that when nominal income grows, for example due to the growth of exports, spending by domestic consumers raises GDP according to Dom share_j—the fraction that domestic sales constitute in GDP—if the expenditure structure is held constant. When the export share of the economy is small, this creates a very strong feedback loop. However, wage growth also raises the prices of domestic goods, which makes them less attractive (relative to imports and also in export markets). In our calibration, we find that these effects substantially weaken the feedback loop. Chodorow-Reich (2017) follow a similar logic when they describe how expenditure switching effects reduce the local fiscal multiplier. If our home country is viewed as a small region of a large country, the export demand shock we consider becomes isomorphic to government purchases financed from the outside. An important difference is that in our case the shock affects welfare through terms-of-trade, whereas in Chodorow-Reich (2017) there is a real output response, due to price rigidities.

³⁸The average wage may fall in response to some trade shocks. In such a case, the average wage effect would contribute to a fall in the skill premium if $\Delta_{VA} [\eta_j^{\text{avg wage}}] > 0$.

³⁹This ensures the proportionality assumption that underlies the U.S. input-output tables: all final and intermediate users of an industry’s output purchase varieties from different countries in the same proportions. World Input-Output Tables depart from the proportionality assumption slightly by allowing for differences in import shares between final and intermediate consumers (but not across different types of each). However, they are more aggregated and have other limitations (Timmer

Appendix A.2 provides the formal exposition of our full model, which differs from the model without input-output linkages in several ways. In the full model, the expenditure channel is governed by spending patterns on imported products (“direct imports”) but also on imported inputs embedded in domestic products (“indirect imports”). For the earnings channel, input-output linkages enrich the result in three ways. First, a fall in domestic import tariffs has a new effect: the prices of domestic products produced with imported inputs fall, which raises demand for these products. Consequently, the skill premium falls if college graduates work in industries that use fewer imported inputs. Second, domestic goods compete with foreign varieties to sell not only to final consumers but also to downstream industries. Finally, all shocks propagate upstream. For example, increased export demand for a final good raises demand for its inputs, whose production grows accordingly, which causes further expansions upstream. As a result, the export share driving the industry expansion includes exports of the industry itself but also of its downstream buyers. The same logic applies to import penetration. Similarly, income effects increase industry size according to the weighted *average* of income elasticities of the industry itself as well as the domestic final industries which buy its output. In the remainder of the paper, we refer to these measures as input-output adjusted measures.

3 Differential Spending on Imports Across Industries

In this section, we use data on 380 detailed industries to compare the spending shares on imports of individuals with and without a college degree. We leave the investigation of differences in spending on imports *within* industries to Sections 4 and 5. We find that spending on imports is similar between education groups, equal to 13.3% for consumers with a college degree and 14.0% for those without, which implies that the expenditure channel is distributionally neutral at the level of industries.

3.1 Data

We conduct the analysis based on a detailed merge of the Consumer Expenditure Survey (CEX) to the BEA Input-Output (I-O) table. The CEX is a survey which measures the universe of personal spending for around 650 detailed categories, covering all categories of goods and services, and records consumers’ characteristics such as education and income.⁴⁰ We merge the CEX to the BEA I-O table to obtain information on domestic production and trade. Specifically, we build a manual concordance from 636 CEX consumption categories into 172 I-O industries.⁴¹ BEA data are the most detailed available accounts of the entire U.S. economy, including non-manufacturing; the most recent detailed I-O table is from 2007, et al., 2015).

⁴⁰The CEX consists of two separate parts, the interview and diary surveys, which we use in combination.

⁴¹This concordance is much more detailed than the one between CEX and NIPA, provided by BLS and used by Buera et al. (2015) and Jaimovich et al. (2015), among others. We thank James O’Brien for providing us with the concordance between CEX interview categories and the 2012 version of NAICS from Levinson and Brien (2016). We use this concordance, converted into 2007 I-O codes, as a starting point. We manually extend it to diary categories as well as some missing interview ones. The concordance is many-to-one, with a few exceptions where we allocate CEX consumption by each group equally across the corresponding I-O codes.

so we center the analysis around that year. To increase sample size, we combine data from the CEX for 2006-2008. Appendix C.1 provides more details on the CEX and the specific sample we use.

We use the I-O table in three ways: as a source of import shares, input-output linkages, and consumption structure. First, for each industry we compute the import penetration as a percentage of absorption (defined as output plus imports minus exports). There are two advantages of using the BEA data to measure import penetration: trade in services is accounted for and trade flows are measured from the same data as domestic output, which improves consistency. Second, we build the input requirement matrix, which measures the composition of suppliers for each buying industry and from which we construct the share of indirect imports (imports of intermediate inputs used in domestic production).⁴² Finally, we use personal final consumption from the I-O table as a measure of total spending in the industry, which is then decomposed into consumption by education group using the CEX. This approach parallels Lebow and Rudd (2003), who show that reweighting the CEX using BEA spending shares yields more accurate inflation estimates, correcting the non-classical measurement error in the CEX (e.g. Garner et al., 2009).⁴³ Appendix C.2 provides additional details on input-output tables and on the data construction.

To characterize which groups of industries drive the effects we will document, we classify industries into standard sectors and sub-sectors. We start by computing all I-O-related objects using all 389 industries and drop 9 special industries at the end.⁴⁴ Manufacturing, agriculture, and mining are classified into goods, while all other industries into services.⁴⁵ Goods and services are further classified into 24 and 15 subsectors based on three-digit I-O codes and two-digit NAICS codes, respectively, which are listed in Table A2.⁴⁶

To conduct the analysis for specific trading partners, we need to merge additional trade data. The I-O table provides total imports for each industry but does not decompose them into countries of origin. To measure import penetration from China, NAFTA countries, and developed economies, we use the 2007 U.S. international trade flows statistics from the Census Bureau by product and source country, which were made available by Schott (2008) and converted into NAICS industry codes by Pierce and Schott (2012). We follow the same strategy as with the consumption data: we keep total imports from the I-O table, which are consistent with the rest of the I-O table, but distribute them across countries of origin using the trade statistics. For each I-O industry, the import penetration from a specific trading partner is computed as the product of total I-O-based import penetration and the fraction of this trading partner

⁴²See Section 2.3 for an informal discussion and Appendix A.2 for formulas.

⁴³Measurement error in the CEX does not create biases for our results as long as it has the multiplicative structure proposed and justified by Aguiar and Bils (2015): there may be industry- and consumer group-specific biases but no interactions between them. Industry-specific biases are corrected by the BEA weights, while consumer-group-specific biases only result in a re-scaling of consumption across groups without systematic effects on the expenditure composition of each group.

⁴⁴These are five government industries, Scrap, Used and secondhand goods, Noncomparable imports, and Rest of the world adjustment.

⁴⁵Construction is sometimes viewed as a good-producing industry (Comin et al., 2016) and sometimes as a service industry (Cravino and Sotelo, 2017). We treat construction as an industry ultimately providing shelter for households and businesses, therefore we classify it into services.

⁴⁶We merge Management and Administrative services (NAICS industries 55 and 56) into Professional, Scientific, and Technical Services (code 54).

in total imports in the NAICS codes that belong to this I-O code.⁴⁷ Armed with this merged dataset, we turn to the computation of import spending shares by education group across industries.⁴⁸

3.2 Results and Mechanisms

We find that spending shares on imports are very similar across consumers with and without a college degree. Table 1 provides the main estimates on the average and differential spending on imports, decomposing the latter into within- and between-components. The first row shows that the total expenditure of U.S. consumers in 2007 includes 13.7% of imports (column (1)), of which 7.3 p.p. are direct imports (column (2)) and 6.4 p.p. are indirect imports via imported intermediate inputs embedded in domestic products.⁴⁹ The overall spending share on imports is slightly lower than the imports-to-GDP ratio, which is 15.6%, because some imports are used in the production of exports. The following rows compute average spending on imports for consumers with and without a college degree. College graduates devote 13.3% of their spending to imports, as opposed to 14.0% for consumers without a college degree. College graduate benefit therefore less when imports become cheaper, but the difference is small, only 0.66 p.p., or 4.8% of the average.⁵⁰ The difference mostly comes both from direct imports (0.28 p.p.) and indirect import spending (the remaining 0.38 p.p.).

The small observed difference in spending shares on imports across education groups is primarily the consequence of two offsetting patterns. On the one hand, college graduates consume more services as a fraction of total expenditure than consumers without a college degree,⁵¹ and services are much less imported than goods. If consumption baskets of the two groups were identical within goods and within services, the import spending share for college graduates would have been 1.47 p.p. lower than for individuals without a college degree.⁵² On the other hand, within goods and services college graduates

⁴⁷Trade flow statistics is available only for trade in goods, we therefore assign zero imports from the specific trading partners of interest in all industries which are not available in the trade flows statistics. This does not constitute an important limitation for China and Mexico. For instance, China constitutes less than 3% of total U.S. imports of services according to the BEA International Services tables for 2007. This limitation is likely to be more important when considering trade with developed economies.

⁴⁸It should be noted that the merged dataset has several imperfections. There is a notable discrepancy between the I-O table and our model, which we ignore in this draft of the paper. In the data total final expenditures differ from personal expenditures because of investment (capital expenditures) and government purchases. From the point of view of our static model, there is no conceptual difference between a firm purchasing materials or building a new factory, but only the former type of spending will appear in the I-O table. A factory building is instead treated as final use of the construction industry and value added for the buyer. There are also discrepancies between the ways certain industries are treated in the CEX and I-O table, ignored in the current version as well. For instance, imagine a t-shirt is sold at a department store. We attribute all of its value to the apparel industry in the CEX, whereas the I-O table moves its wholesale and retail margins into the Wholesale and Retail Trade industries, respectively.

⁴⁹Note that to get this decomposition we did not have to classify goods into final and intermediate ones, which is problematic because many goods are used in both ways. Instead we rely on the input-output table to split the *use* of products as final or intermediate.

⁵⁰This difference is statistically significant. Inference here is based on the randomness of the CEX sample. Error in the industry import penetration is non-statistical (e.g. due to imputations required to prepare input-output tables) and is not taken into account.

⁵¹Boppart (2014) has established the same pattern with respect to consumer income, also using the CEX data.

⁵²Services constitute 81.9% (78.2%) in consumption baskets of individuals with (without) a college degree. The average share of imports is 46.0% for goods and 5.5% for services. See Table A3 for the estimates and Appendix A.4 for the within-between decomposition formula.

spend relatively more on industries with a higher import share, which reduces the difference by more than half (0.81 p.p.). Table 1 also shows that most of the differences within goods and within services (0.65 out of 0.81 p.p.) result from differences in spending patterns across 29 subsectors, while within-subsector heterogeneity is relatively unimportant.⁵³

The differences in spending on imports across education groups within goods and within services are robust patterns, which can be assessed graphically. Panel (a) of Figure 1 shows the relationship between the share of sales to college graduates and the average share of imports, both direct and indirect, in domestic final expenditures. Each dot represents 5% of detailed industries within goods or services (with final consumption weights) and the figure shows that the patterns are not driven by a small number of industries: there is a strong positive slope for goods and a smaller one for services. The interpretation of this graph relies on the insight that the share of spending on imports is higher for college-educated *consumers* if and only if *industries* that sell relatively more to college graduates have higher import shares (see Appendix A.4 for a formal treatment).

We characterize graphically the parts of the product space contributing to the within-sectoral differences in spending on imports in favor of college graduates. Panel (b) of Figure 1 groups industries by subsectors, where the size of each circle indicates their importance in final expenditures. “Food” and “Computers and electronics” are characteristic subsectors within the goods sector: while food is purchased relatively more by non-college consumers and does not have much imports, electronics are represented disproportionately in the consumption basket of college graduates and have a very high share of imports.⁵⁴

We obtain qualitatively similar patterns, with small differences in spending on imports across education groups, when considering trade with specific partners. Columns (3)–(8) of Table 1 show statistics for spending on imports from three sets of countries: China, NAFTA (Canada and Mexico) and a group of 34 developed economies (OECD members, excluding NAFTA, plus Taiwan and Singapore). College graduates spend 0.09 p.p. *more* on imports from China: this finding is explained by imports of electronics, for which China is the key source (see Figure A2). Consumers without a college degree spend more on imports from NAFTA and Developed Economies (0.27 p.p. and 0.33 p.p., respectively), but this is only due to their higher spending on goods, rather than to differences within goods.

Similar patterns hold when considering differences in spending on imports across income groups. In Figure A3 we compare spending shares on imports, overall and from China, NAFTA, and developed countries specifically, across bins of household income. In all cases, we find sufficiently flat patterns, although spending on imports from China is slightly increasing with income, while it is slightly decreasing for imports from developed economies.

In sum, considering spending patterns across 169 categories of final consumption, we have shown that college and non-college educated consumers have similar spending shares on imports, whether overall or from China. Our analysis so far suffers from potential “aggregation bias”: for instance, it could be the

⁵³The subsectors with non-zero final consumption include 17 three-digit I-O codes for goods and 12 two-digit NAICS codes for services, listed in Table A2 without stars.

⁵⁴If these two subsectors are dropped, the regression slope estimate halves.

case that the low-skill group consumes a larger fraction of imported varieties *within* categories. We now turn to this question to provide evidence that there is no such pattern for consumer packaged goods (Section 4) and automobiles (Section 5).

4 Differential Spending on Imports of Consumer Packaged Goods

In this section, we examine within-industry spending on imports for consumer packaged goods—goods that are typically purchased in supermarkets—creating a firm-level dataset in which we observe both consumer type and imports. We find patterns that are similar to those across industries: college graduates spend more imports, but the difference is small. Expressed as percentage of average spending on imports, the difference between consumers with and without a college degree is less than 7%.

4.1 Data

We start from the Nielsen Homescan Consumer Panel (henceforth Nielsen), which measures spending at the level of barcode and provides consumer characteristics such as education and income.⁵⁵ This dataset does not classify products into domestically produced and imported and, as a consumption survey, it is not informative on the share of imported inputs embedded in domestic products. We address both problems at once by linking products (barcodes) to their producers or distributors in the Economic Census by name and address. We measure the sales share of imports of those firms using the U.S. Customs data, capturing both direct and (first-order) indirect imports. We find Census matches for the majority of Nielsen firms, excluding the smaller ones, and cover over 80% of total Nielsen sales.⁵⁶ This novel linked dataset may be a useful platform for future research at the intersection of the consumption and production sides of the economy.

The Nielsen data cover over \$400 billion purchases per year in three classes of products: (i) food, alcohol, and tobacco (henceforth “food”), (ii) health, beauty, and household products (henceforth “health and household”), and (iii) general merchandize, namely other products found in supermarkets such as tableware, stationery, or electronics. Within product classes there are 10 departments (e.g. Frozen Foods), 117 product groups (e.g. Frozen Prepared Foods), and 1,165 product modules (e.g. Frozen Soup). To study within-industry differences in import spending in the way compatible within industry-level analysis of Section 3, we manually convert modules into 71 detailed I-O industry codes.

We attribute barcodes to firms using addition data from GS1, a U.S. non-profit organization than maintains the barcode system. To sell products in supermarkets, a manufacturer or a distributor has

⁵⁵The data are from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business. Information on availability and access to the data is available at <http://research.chicagobooth.edu/nielsen>.

⁵⁶We are aware of two alternative approaches to identify imported products in barcode-level data. First, the first three digits of the barcode identify the country in which the barcode is registered. Bems and Giovanni (2016) exploit this strategy for Latvia, where foreign firms do not register their products locally. But in the U.S. most foreign products have domestic barcodes. Second, Antoniadou and Zaniboni (2016) manually collect the country-of-origin information listed on 3,000 product labels in the United Arab Emirates. The massive number of products sold in the U.S. makes this strategy infeasible for us.

to purchase a block of barcodes from GS1; each barcode can only be registered by one firm. With a small fraction of exceptions, these firms have U.S. addresses—foreign firms do not tend to register barcodes without an affiliate or an intermediary in the U.S. In Appendix D we track several products photographed in a Walmart store to verify that domestically produced goods are normally registered by the manufacturer, while imported ones by the distributor, often a wholesaler.

We then link the Nielsen data to three confidential data products on American businesses collected by the U.S. Census Bureau. First, Business Register, or SSEL, is the comprehensive list of firms and establishments, i.e. locations of economic activity within firms; we use it as a source of names and addresses to merge firms in the Census with Nielsen. Second, the source of production data is the quinquennial Economic Census from 2007 and 2012.⁵⁷ Because we are interested in imports of final products which are often done by dedicated wholesalers or by retailer, it is useful for us to observe establishments in the entire economy rather than the most commonly used Census of Manufacturers only. Finally, LFTTD is the transaction-level data on imports and exports of goods from the U.S. Customs, linked to the other databases by the firm identifier.

Merging firms between Nielsen and the Economic Census is a complex multi-stage procedure. For a brief summary, we use both exact and fuzzy matching on firm names and different components of the address: state, city, street, street number, and zipcode. We develop a set of consecutive merging rules and verify their quality by manual inspection of a sample of merged firms. Out of the total number of 23,300 Nielsen firm-years in 2007 and 2012, we successfully match 12,700, covering 83% of sales.

Appendix C.3 provides details on the data sources, describes the matching process, and presents the match statistics.⁵⁸

Key Variables and Summary Statistics. We use Nielsen to measure total sales for barcodes and firms overall and by consumer education and income. To compute firm import shares, we divide the value of imports from LFTTD by total sales of all firm establishments in the Economic Census. This includes imports of both final products (e.g. by wholesalers and firms with multinational production) and intermediate inputs, except those acquired through domestic intermediaries. We split the total import share into a sum of four components, based on the imports from China, NAFTA (Canada and Mexico), 34 developed economies, and rest of the world.⁵⁹

We adopt a square-root weighting scheme to reduce measurement error. As we cannot attribute firm’s imports to a particular class of products it sells or even to Nielsen products overall, our proxy for the import share is likely to be noisier for large firms, which typically operate in multiple industries. Those same firms play a large role when measuring the average share of import spending for each consumer

⁵⁷To reduce noise we merge three years of the Nielsen data to each Economic Census: 2006–2008 for the 2007 Census and 2011–2013 for 2012.

⁵⁸The appendix also shows that Nielsen firms, which we were able to merge, are larger but have similar consumer characteristics as the unmatched ones. Similarly, Census firms within the best covered Food, Beverage, and Tobacco industry, for which we found a Nielsen match, are larger than other firms in that industry but have similar composition of workers by skill.

⁵⁹As before, developed countries are OECD members (excluding NAFTA), Taiwan, and Singapore.

group, a consequence of granularity that characterizes firm-level datasets (Gabaix, 2011). We make the results more stable by rescaling each firm’s Nielsen sales to its square root, reducing the influence of poorly measured large firms.⁶⁰

The main unit of the analysis will be a firm selling in a given product module. Although import shares only vary at the firm level, decomposing firms into modules allows us to perform a more nuanced analysis: measure differential import spending for each product class separately and decompose it into the within- and between-industry components.

Table 2 shows the summary statistics for the main variables used in the analysis at the level of firms by product modules, for all products together and split by the product class. The average spending on imports is 11.1%, with large differences across product classes, ranging from 6.9% for food to 14.6% for health and household and to 28.0% for general merchandize (with substantial variation across product modules within each class).⁶¹

4.2 Results and Mechanisms

Table 3 measures the average share of spending on imports and the difference between Nielsen panelists with and without degree college. Column (1) estimates that imports constitute 11.1% of the total expenditures, but this number is higher for college graduates than consumers without a college degree: 11.5% and 10.9%, respectively. The pro-skilled difference of 0.59 p.p. is statistically significant and equals 5.4% of the average.⁶²

To make our Nielsen results complementary to the industry-level ones, we need to avoid double-counting. Since the Nielsen sample covers multiple I-O industries, some of the 0.59 p.p. difference may stem from the composition of consumption across I-O industries, which has already been accounted for in Section 3. However, the next row of column (1) shows that double-counting is not a big issue: most of the difference (0.48 p.p.) is a consequence of differential spending on imports within the same industries.⁶³ Moreover, around half of the total difference (0.28 p.p.) is within product modules—the most detailed classification available in Nielsen.

Panel (a) of Figure 2 visualizes the differential import spending. This binned scatterplot groups firm-module cells by their consumer base—the share of college graduates in sales—and correlates it with the

⁶⁰Granularity is a substantial issue: in the full Nielsen sample, top 50 (200) firms capture 46% (70%) of sales in an average year; with square-root weights, they take up only 9% (21%). When decomposing firm sales, we allocate the total firm weight (square-root of sales) to different products and consumers proportionally to their fraction in the total sales, which ensures consistency of results across levels of disaggregation. In Appendix E.1 verify robustness of our results to other weighting schemes: full sales and sales to the powers 1/4 and 3/4. We pick power functions as weights because they are the only ones which satisfy scale invariance: weighted averages are the same whether sales are measured in dollars or thousands of dollars.

⁶¹The composition of imports is also substantially different: while China is unimportant for food, it plays a bigger role for health and household and takes up the majority of imports in general merchandize. On average, 31.2% of firm sales go to college graduates, with small differences across product classes. While all of these statistics are computed using the square-root weighting scheme described above, Appendix Table A11 shows analogous averages with full weights.

⁶²Unlike Section 3, the sample of households in Nielsen is much larger than the set of firms. Therefore, we develop and implement a different approach to inference that is based on the randomness in the sample of firms; see Appendix B.1.

⁶³One caveat is that the within-industry component may be overestimated because large firms span many industries. Our firm-level measure of imports does not allow for differences across industries that a firm is selling in, which potentially attenuates the across-industry component.

average import share for corresponding firms on the vertical axis. Appendix A.4 proves that the differential spending on imports between the two groups can be expressed as the slope of this relationship, multiplied by the consumption segmentation index—a measure of dissimilarity between their consumption baskets. In this graph we residualize the relationship on the I-O industry dummies to capture only the within-industry component of differential spending.

The figure shows that firms selling more to the college graduates, import more. The slope is not only statistically but also economically significant: while firm-modules with less than 20% sales to college graduates have the average import share under 9%, the corresponding import share is around 12% for firm-modules with more than 40% of college-graduate sales. However, the slope of this relationship is not sufficient to generate big differences in import spending between groups, given the consumption segmentation index of 6.5%.⁶⁴

The remaining columns of Table 3 report average and differential spending on imports from China, NAFTA, and developed economies separately. The results are intuitive: almost all of the difference from column (1) is explained by imports from developed countries, with the difference reaching 17.6% of the average. College graduates also spend slightly more on imports from Canada and Mexico. At the same time, they spend around 4.4% less on Chinese products. Panels (b)–(d) of Figure 2 show corresponding graphs.

We highlight that these patterns differ from those across industries (Table 1), complementing those. College-educated consumers relatively more from industries where China is strong, in particular computers and electronics, and cross-industry patterns are weaker for developed economies. However, within consumer packaged goods college graduates buy slightly fewer varieties from China, and more from developed countries. Table A9 verifies that the findings hold qualitatively in each of the three product classes: college graduates spend less on Chinese products (in particular for health and household products, but still only 5.4% difference) and more on other imports (most strongly for food).

A natural mechanism for the patterns we observe is related to product quality, which richer college graduates may value relatively more. We test whether imports from developed countries are of high quality, while Chinese imports have lower quality relative to domestic products. Although we do not model quality explicitly,⁶⁵ we approach this question empirically by proxying for quality with detailed barcode-level prices.⁶⁶ We convert prices into comparable units within product modules, e.g. per ounce of soda rather than per bottle, and split the distribution of prices within the module into deciles.

⁶⁴This value of the segmentation index means that college graduates buy from firms which, on average, sell 6.5 p.p. more to them than the firms which non-college consumers buys from. Panel (a) of Figure 2 can be compared with Figure 2A from Jaravel (2017). Using similar Nielsen data, he finds that the relationship between product module-level inflation and the average consumer income is sufficiently strong to generate big differences in average inflation across groups.

⁶⁵See Fajgelbaum et al. (2011) for a model with this effect.

⁶⁶There is a long literature on trade and product quality. Early contributions, such as Schott (2004) and Hummels and Klenow (2005), similarly used unit values from the trade data as a proxy for quality. More recent papers, such as Khandelwal (2010) and Hallak and Schott (2011), adjust prices by productivity estimate to obtain measures of quality. Lacking proxies for productivity for non-manufacturing firms in the Census data or any measures for detailed products, we do not do this useful adjustment. At the same time, we benefit from directly observing consumer prices, while using trade data creates aggregation concerns.

Panel (a) of Figure 3 confirms that college graduates buy higher-priced products: moving from the lowest to the highest decile shifts the share of college consumers from 25% to almost 40%. The following panels of this figure locate imports in the distribution of prices. Products in the top deciles in their modules tend to have more imports from countries other than China, with most of the effect coming from developed countries (Panels (b) and (c), respectively). Conversely, imports from China in the Health and Household product class are substantially more prevalent at the bottom of the price distribution (Panel (d)). That pattern is not present for imports from China within General Merchandise (Panel (e)), which is consistent with weaker differences in spending between college and non-college consumers in that class (Table A9).

Appendix E.1 presents additional evidence. It shows that the differences we observe are likely to be generated by imports of final, rather than intermediate, products. The patterns are similar when comparing import spending across the income distribution instead of education groups.

This appendix also presents a series of robustness checks. We show that the results are unchanged if we weight firms by their Nielsen sales to the power $1/4$ or $3/4$, and are generally similar with full sales weights. Mismeasurement of import shares for retailers is also not likely to be a problem. Most importantly, we develop a methodology to bound the attenuation bias that may arise because we only observe imports at the firm level, not for individual barcodes. Applying this methodology, we find that attenuation should not reduce the effect by more than 1.5 times.

Overall, we find that college graduates buy more imports, in particular from developed countries, but less of Chinese imports, consistent with differences in product quality. However, these differences are relatively small. Even for China, there is no product class where the anti-skilled bias of import spending exceeds 6% of the mean.

5 Differential Spending on Imports of Automobiles

This section shows that college graduates devote a larger fraction of their spending on cars to imported models. The difference is very large when considering cars assembled outside of U.S. and its NAFTA partners (Canada and Mexico). When imports of assembled cars from Canada and Mexico are included, the difference in import spending shares between education groups is reduced by approximately a factor of two, but it remains significant.

5.1 Data

We conduct the analysis at the level of brands, combining data from the CEX and Ward’s Automotive Yearbooks. The CEX interview survey asks households to report the brands of the cars they own.⁶⁷ Chevrolet and Buick are examples of such brands, which are more detailed than firms (Chevrolet and

⁶⁷The respondents are asked to list all vehicles they own, but we focus on cars, excluding trucks (including SUVs), motorcycles, boats, etc. As previously, we classify households into the two skill groups based on the college education of the respondent. Household income before tax is used as an alternative measure of skill for robustness.

Buick are both produced by GM) but not as detailed as models (e.g. Chevrolet Camaro). We use data from Ward’s Automotive Yearbooks, a leading publication for statistics on the automobile industry, to estimate for each brand the fraction of cars assembled outside of the U.S. or NAFTA in the total number of cars sold in the U.S. To reduce noise in both datasets, we combine years 2009–2015.

We use Ward’s table that reports the number of cars of each model sold in the U.S. with a breakdown into those assembled within and outside NAFTA. Unfortunately these sales statistics are not decomposed by country within NAFTA. However, Ward’s statistics on assembly are more detailed, reported by model and individual country. Most models are only assembled in one of the NAFTA countries, so we assign all of their American sales to that country. For models produced both in the U.S. and in Canada or Mexico, we allocate sales proportionately to assembly in those countries. Finally, we aggregate models belonging to the same brand, yielding brand-level import shares.

Our final sample includes 39 brands and 51,498 vehicles, 38.3% of which were purchased new. Table A15 reports consumer characteristics and import shares across brands. We dropped a small fraction of CEX purchases for brands which we do not observe in Ward’s because their production was discontinued before 2009.⁶⁸ Appendix C.4 provides more detail on the data construction.

5.2 Results and Mechanisms

We find that college graduates are more likely to purchase imported cars. Table 4 measures the average and differential spending on imports of automobiles. Column (1) shows that on average 49.7% of purchased cars are assembled abroad, while the corresponding number is 53.8% for college graduates and 47.2% for those who did not attend college. The spending share of college graduates on imported cars is thus 6.6 p.p. higher, or 13.3% of the average.

The spending differences on imported cars between education groups vary substantially across trading partners. Columns (2) and (3) of Table 4 decompose imports into those from NAFTA and from other countries. Spending on imports from outside NAFTA exhibits striking pro-skilled differences: 32.7% of cars purchased by college graduates are imported from those countries, relative to only 21.3% for those who did not go to college. Imports from Canada and Mexico offset about half of this effect: the spending share from college graduates is 4.8 p.p. lower in this sample.

The patterns are similar for new and used cars. Columns (4) and (5) of Table 4 split the sample into new and used cars and document similar patterns for both types of purchases, with a slightly stronger pro-skilled difference for used cars.

The robustness of the findings can be assessed graphically. Figure 4 plots import share against the fraction of sales to college graduates across brands, reporting the results separately for imports from all countries outside a NAFTA and from all countries. Panel (a) shows a very strong positive relationship between imports and consumer education when excluding NAFTA countries. Two clusters of brands

⁶⁸Oldsmobile is the most frequent brand we have to drop. We also drop CEX observations corresponding to brands which produce only trucks but not cars, such as GMC and Jeep. All dropped brands combined constitute less than 2.5% of the sample.

become apparent: those selling to college graduates are mostly high-end foreign brands (e.g. BMW, Lexus, and Mercedes-Benz), whereas brands selling to the consumers without a college degree are almost all domestic (e.g. Chevrolet, Buick, and Dodge). Panel (b) shows that the total difference in import spending shares across education groups is weaker because of the NAFTA: many domestic brands have 20 to 50% of their cars assembled in Mexico and Canada, creating a partially offsetting effect. Consistent with Table 4, the slope for total imports remains positive and significant in Panel (b).

Appendix E.2 shows that similar patterns hold across income groups. The pro-rich bias is particularly strong for non-NAFTA brands above the 80th percentile the income distribution.

5.3 Imports of Car Parts

Domestically assembled cars may use a substantial amount of imported parts, which the data we have used so far cannot capture: Ward’s only gives information on the country of assembly. To address this potential issue, we use the confidential Census of Manufactures and the Customs import transactions data, where the fraction of both imported cars and car parts in the value of sales can be measured. The direct share of imports is computed as the ratio of imports of assembled cars from the Customs data to the value of car shipments from the Census. The total share of imports additionally includes imports of car parts in the numerator, using the harmonized classification of traded products.⁶⁹ The downside of this approach is that we have to aggregate the sample from the level of brands to firms, overlooking the patterns of consumption and imports across brands of the same firm. For this reason, our main analysis was based on the CEX and Ward’s Automotive Yearbooks.

We find that college graduates spend less on indirect imports of automobiles, but the difference is very small and does not outweigh the fact that they spend more on direct imports. Table 5 reports the results.⁷⁰ In columns (1) and (2), the dependent variables are direct and total imports (as a percentage of total new auto sales) and the regressor is the share of new cars sold to college graduates in the CEX. Columns (3) and (4) repeat the same exercise for used car sales. In both cases, the regression coefficient is reduced when indirect imports are taken into account, but by around 10% only. While these regressions may suffer from aggregation bias, because we are considering firms instead of brands, the comparison between them strongly suggests that indirect imports do not have a substantial offsetting effect on differences in imports spending between education groups, thereby validating our estimates. The results from Sections 3–5 thus consistently find that more educated individuals spend a slightly higher share on imports, both across and within industries.

6 Differences in Exposure to the Labor Market Effects of Trade

Moving from the expenditure side to the earnings side, this section reports reduced-form evidence characterizing the extent to which different education groups are exposed to the labor market effects of trade.

⁶⁹Appendix C.5 describes the data construction in more detail.

⁷⁰Data confidentiality does not allow us to show individual observations, like Figure 4 did.

As discussed in Section 2, the earnings channel of trade liberalization can affect labor demand through intensified import competition, increased opportunities to export, cheaper intermediate inputs, as well as through income effects from the gains from trade. This section documents how these channels vary across industries and shows that they tend to favor college graduates. Our main analysis is based on the industry-level data, in parallel with Section 3, but we return to possible aggregation biases at the end of the section.

6.1 Data

Payroll data by industry and education group is the first ingredient we need to compute the labor market exposure to trade across education groups. We rely on the 2007 American Community Survey (ACS) and the 2007 Quarterly Survey of Employment and Wages (QCEW) to obtain this information.⁷¹ We select all respondents employed in the private sector and aggregate their annual labor income by industry, doing so separately for workers with and without a college degree. Industry skill intensity is computed as the payroll share of college-educated workers. Since the industry classification in the ACS is not very detailed (covering only 253 industry groups), we infer skill intensity at a finer level of disaggregation based on the fact that skill intensity is strongly correlated with average wages across industries. Average wages are available for each detailed six-digit NAICS industry from the 2007 QCEW. For each two-digit NAICS sector, we regress the ACS skill intensity on the QCEW average wage, using ACS industry groups as the unit of observation, and we then use this relationship to predict skill intensity for each six-digit NAICS industry. Appendix C.6 provides more detail about this imputation procedure.

Moreover, we need to compute four industry-level outcomes: import penetration rates, export shares, the shares of imported inputs, and income elasticities. We do so using the same industry-level data as in Section 3. The first three outcomes are measured using the I-O table. To disaggregate import penetration statistics by country of origin, the I-O table is merged with the trade statistics from Schott (2008) and Pierce and Schott (2012). To estimate income elasticities, we use the CEX.⁷²

6.2 Results and Mechanisms

The reduced-form patterns suggest that the earnings channel favors college graduates. Table 6 reports the average and differential exposure of education groups to the various labor market effects of trade. Payroll weights are used for all statistics.

First, we find that college graduates are less exposed to import competition. Column (1) of Table 6

⁷¹The ACS is the long form of the population census that is answered by a random 1% sample of the U.S. population every year. We obtain the 2007 ACS via IPUMS (Ruggles et al., 2015). The QCEW is published by the Bureau of Labor Statistics based on unemployment insurance statistics and has almost universal coverage. We match ACS industries to NAICS using a crosswalk provided by IPUMS (available at <https://usa.ipums.org/usa/volii/indcross03.shtml>).

⁷²We split consumers into income bins and measure bin-specific expenditure shares on 671 CEX spending categories. By definition, higher-income consumers have larger expenditure shares on income-elastic products. Using this logic, we first compute the income *semi*-elasticity for each spending category by regressing spending shares on the logged total expenditure, corresponding to the bin. We convert the estimates to elasticities and aggregate them into the I-O industries. Appendix C.6 describes details of the procedure.

reports import penetration ratios. The average import penetration in domestic industries, weighted by payroll, is only 4.21%. This share is low, for instance in comparison to the imports-to-GDP ratio, for two reasons. First, as prescribed by the model, this share includes only direct imports: import competition stems from imports *of* the industry’s product, rather than from imported inputs *by* the industry. Second, international specialization limits the negative impact of import competition on the industry size because employment and payroll are low in industries in which the U.S. has largely stopped producing, such as toys.⁷³ While consumption-side gains coming from an industry are scaled by *expenditures* in that industry, the reduction in labor demand due to import competition is scaled by the industry’s *payroll*. The following rows of Column (1) establish that industries that employ college graduates are less exposed to import competition: average import penetration is 4.13% for workers with a college degree, compared with 4.28% for those without. Albeit small, the 0.15 p.p. difference is statistically significant.⁷⁴

The lower exposure of college graduates to import competition is interesting to investigate in the context of standard trade theories, as it may appear consistent with a standard Heckscher-Ohlin mechanism. Since the U.S. is a relatively skill-abundant country, its imports may be expected to be higher in low skill-intensive industries. However, the Heckscher-Ohlin interpretation is largely misguided, as can be seen by decomposing the differential exposure to import competition into “between” and “within” components across goods and services.⁷⁵ The lower exposure of college graduates to import competition is driven by the fact that services constitute a larger share of their payroll and are much less imported than goods.⁷⁶ If the two groups had identical compositions of payroll within goods and within services, import competition exposure for college graduates would have been 1.54 p.p. smaller than for workers without a college degree, which is a much bigger difference than what we observe in the data. In other words, within sectors college graduates are in fact more exposed to import competition, in contrast with the prediction of the standard Heckscher-Ohlin mechanism.

The finding that more skill-intensive industries have *higher* import penetration within goods and within services is easy to see graphically. Panel (a) of Figure 5 plots import penetration against college payroll across industries (each dot represents 5% of the data) and shows that the positive slopes are robust and not driven by outliers. The slope is quite large for goods and smaller for services.

To identify the parts of the product space contributing to within-sector differences in exposure to import competition, Panel (b) of Figure 5 groups industries by subsectors. The size of each circle reflects the importance of each industry according to their share in total payroll. Like on the expenditure side, “Food” and “Computers and electronics” are characteristic subsectors within the goods sector: while food manufacturing is low skill-intensive and does not have much import competition, electronics are

⁷³Wood (1995) explains that specialization makes imports beneficial for both types. In our model, it is captured by the fact import penetration is weighted by payroll for the earnings channel but by consumption for the average gains, and those weights diverge as countries specialize more.

⁷⁴Similar to Section (3), inference here is based on the randomness of the ACS sample. As it is a very large sample (1% of U.S. population), standard errors are small. Errors in the industry import penetration and other outcomes are non-statistical (e.g. due to the imperfect measurement of imports by NAICS in the I-O table) and are not taken into account.

⁷⁵Appendix A.4 derives the decomposition formula.

⁷⁶Table A18 reports that services constitute 88.7% (82.0%) of payroll for workers with (without) a college degree. The average import penetration is 23.9% for goods and 6.6% for services.

high skill-intensive and have a high share of imports. The last rows of Column (1) of Table 6 show that the within-sector difference in import competition between education groups mainly occurs across subsectors, which implies that Panel (b) of Figure 5 is an accurate depiction of the main forces at play in the data (heterogeneity across finer industries contributes to the within-sector difference in the same direction).

Appendix E.3 shows that the import competition patterns within the goods-producing sector are a relatively recent phenomenon. Using the NBER CES panel data on manufacturing industries, we establish that the relationship between import penetration and skill intensity was flat in 1992, weakly increasing in 1999, and steep in 2007. Growing imports of machinery and electronics explain why the slope increases over time.

In addition to direct import competition (in the industry of a worker’s employment), indirect import competition in downstream industries also affects employment and wages, potentially differently across education groups. Column (2) of Table 6 takes this into account by using the I-O adjusted import penetration as the outcome variable. This adjustment is performed using the input-output table, adding up import penetration rates in downstream industries.⁷⁷ The differential exposure to import competition quadruples from -0.15 to -0.60 p.p., due to the increase in the between-sector component. Intermediate goods-producing industry mostly sell to other goods-producing industries, therefore they suffer more from import competition. This pattern strengthens the impact of import competition in the goods sector, which accounts for a lower share of payroll for college graduates, and thus reinforces the bias of the import competition channel in favor of college graduates.⁷⁸

The remaining columns of Table 6 conduct a similar analysis for export shares, imported inputs and income elasticities. In all cases, results are reported both with and without the adjustment for input-output linkages. I-O linkages turn out not to make any qualitative difference, therefore we only discuss the I-O adjusted patterns.

We find that the export channel favors college graduates. Column (4) of Table 6 shows that an average industry (weighted by payroll) exports 9.2% of its output, with equal proportion of direct and indirect exports. The exposure of college graduates to exporting opportunities is slightly higher, generating a pro-skilled effect of trade liberalizations. As previously, there are offsetting effects. The between-sector force favors workers without a college degree because goods are more exported, but within sectors workers with a college degree are more likely to work for industries that export more, which happens to more than offset the between-sector force. As illustrated by Panel (a) of Figure 6, “Computer and electronics” are more skill intensive and export more compared with “Food”. Similarly for services, “Professional and business

⁷⁷Note that the I-O adjusted import penetration that we analyze here is different from the “total” share of imports in the industry from Section 2. We are now adding up import penetration in downstream (buying) industries, as demand shocks propagate upward from those industries. Conversely, on the expenditure side we were adding up imports from upstream (supplying) industries, as price shocks propagate downward from those industries (see, for example, Acemoglu et al., 2015 for a detailed discussion of these effects).

⁷⁸In contrast, the magnitude of the within-sector component is left essentially unaffected by IO linkages: within each sector, workers with and without a college degree have a similar propensity to work for industries that supply other industries who are competing with the rest of the world.

services” and “Finance” are more skill intensive are export more in comparison with “Construction” or “Accommodation and food”.

Next, we show that college graduates benefit less from imported inputs. According to Column (6), this pattern is due to both between- and within-sector differences in the use of intermediate inputs. For instance, within services, “Construction” and “Transportation” are the subsectors with the highest shares of imported inputs, and both of them have low skill intensity. The contributions of different subsectors to differential exposure to imported inputs is shown in Panel (b) of Figure 6.

Finally, we find that college graduates work in industries with higher income elasticities. The average income elasticity weighted by the college payroll is 1.14, relative to 1.03 when payroll of non-college workers is used. This finding is consistent with prior research (Caron et al., 2014; Leonardi, 2015) and implies that the income effects of trade favor college graduates (Caron et al., 2017; He and Zhang, 2017). As indicated by Panel (c) of Figure 6, the difference in income elasticities across education groups in our data is largely driven by the “Education” subsector, which is highly income-elastic and skill-intensive.

We obtain qualitatively similar patterns when considering trade with specific trading partners. In Table A17 we report the differential exposure to import competition with China, NAFTA, and Developed Economies separately. The general finding that college graduates are less exposed to the adverse labor market effects of trade holds for each of these specific trading partners. The between-sector force is generally the most important channel and, in the cases of China and Developed Economies, it is partially offset by the positive relationship between import penetration and skill intensity within goods.⁷⁹

The various patterns documented in this section are based on industry-level data and may therefore suffer from aggregation bias. In Appendix E.3, we conduct two additional analyses to address this potential issue.

First, we use the plant-level microdata from the Census of Manufactures and the Management and Organizational Practices Survey (MOPS) to compare the differences in exposure to exports between education groups, as well as between non-production and production workers. We find that more skill-intensive plants within the same industry tend to export more; however, the main difference is across manufacturing industries and the degree of aggregation bias is small.

Second, on the import side, we do not know which workers within industries are particularly vulnerable to import competition and offshoring. Borjas et al. (1997) pointed out that a reasonable proxy for the skill level of the marginal worker displaced by import competition may be the skill intensity of the U.S. industry in the distance past. We embed this idea in our theoretical framework by assuming that each industry has two segments, one competing with imports and the other one being insulated from imports. Skill intensity in the import-competing segment is proxied for using the data from the 2000 and 1990 population censuses. We find that the differential exposure to imports is larger under this set of assumptions, but the difference relative to our benchmark case is primarily due to the overall increase in the U.S. skill

⁷⁹Figure A7 depicts the contributions of subsectors to these patterns. One caveat is that data limitations do not allow us to decompose imports of services by country, so we assume that they are zero. This may bias the findings, but the bias is unlikely to be large for China, which accounts for a small share of U.S. services imports (see footnote 47).

endowment, which may not be related to trade.

7 Estimates of The Distributional Effects of Trade Policies

The previous sections have documented that the share of spending on imports are similar between college and non-college educated consumers, while college graduates are relatively less exposed to the negative labor market effects of trade through most channels. These patterns qualitatively predict that trade policies should favor college graduates. However, reduced-form statistics are not sufficient to quantify the counterfactual effects of trade liberalizations, in particular for the earnings channel, for two main reasons. First, the magnitudes of different earnings channel mechanisms depend on structural parameters of the model—elasticities of substitution in demand and production. Second, the reduced-form patterns do not capture general equilibrium effects that in the model result from changing domestic average wages.

To quantify the average and distributional welfare gains of trade policies, this section uses the model from Section 2, calibrated with the reduced-form estimates from previous sections as well as structural elasticities borrowed from the literature. We decompose the distributional effects into the expenditure and earnings channels, and further into different mechanisms to provide evidence on their relative importance. We focus on two specific trade policies: a 10% bilateral fall in all import and export barriers and a 10% fall in barriers on Chinese imports. Our main calibration is at the industry level, structured around the detailed input-output table, although we incorporate the results from Sections 4–5 as well. We find that college graduates benefit from trade liberalization 16% more than non-college graduates, this difference being primarily driven by the earnings channel and slightly strengthened by the expenditure channel. The pro-skilled bias is stronger in the case of China.

7.1 Elasticities

Demand Elasticities. Besides income elasticities estimated in Section 6, the non-homothetic nested CES demand system (equation (1)) is characterized by elasticities of substitution at each nesting tier: ξ_j between domestic and foreign varieties in industry j , ε_r between I-O industries within goods and services sectors, and ρ between goods and services. In this version of the draft, we report results taking typical values from the literature and check robustness of the results to a range of their values.⁸⁰

For substitution between domestic and foreign varieties, our baseline calibration assumes that this elasticity is 3.5 in all industries (which is equivalent to the trade elasticity of 2.5 for trade flows measured by value rather than quantity). The value we use is near the median elasticity of 3.7 reported in Broda and Weinstein (2006) for ten-digit industries, and of 3.4–3.7 in Soderbery (2015) using the same Broda-Weinstein method but for eight-digit industries and for different years of data, as well as near the mean of 3.6 in Ossa (2015).⁸¹ In robustness checks in Section 7.3, we allow for heterogeneous trade elasticities

⁸⁰In ongoing work, similar results are obtained from direct estimation of elasticities. See Costinot and Rodríguez-Clare (2015, section 5.3) for a recent discussion of the state of the literature on trade elasticities.

⁸¹At the same time, the range of other estimates is wide: it can be as low as 1.9 using the Soderbery (2015) LIML

across the product space.

Estimates of the elasticities of substitution between sectors (goods and services) and industries (IO6) are also required. There is substantial debate on the value of the elasticity of substitution between goods and services, but it is generally recognized that the two sectors are complements, i.e. $\rho < 1$. We follow the recent paper by [Cravino and Sotelo \(2017\)](#) and set $\rho = 0.2$ in our baseline calibration, but also consider a range of other values between zero and one as robustness.

Regarding the elasticities of substitution between industries within each sector, ε_r , the review by [Costinot and Rodríguez-Clare \(2015\)](#) uncovers that the prevalent approach in the literature is the “idiot’s law of elasticities” ([Dawkins et al., 2001](#))—setting the elasticities to one absent more clear evidence. We follow the same approach and, in a robustness exercise, consider a range of $\varepsilon_r \in [0.8, 3.5]$ since this elasticity is likely to be above ρ and below trade elasticities.⁸²

Macro Elasticity of Substitution between Labor Types. On the production side, the elasticities of substitution between skilled and unskilled labor in each industry enter only through the macro elasticity σ_{macro} . We follow [Burstein and Vogel \(2017\)](#), [Cravino and Sotelo \(2017\)](#), and [Caron et al. \(2015\)](#) by calibrating the macro elasticity directly rather than aggregating it from micro estimates. We use an estimate of 1.41 for the baseline calibration, following [Katz and Murphy \(1992\)](#) based on the time series of employment and wages. We also check robustness to the range of 1.41–1.8, considering the upper bound of the range estimated by [Acemoglu \(2002\)](#) and [Acemoglu and Autor \(2011\)](#).

7.2 Results and Mechanisms

Overview. We find that trade liberalization generate a modest increase in inequality, primarily due to the earnings channel. Table 7 presents the key results of our baseline model calibration for a 10% bilateral reduction trade barriers and for a 10% reduction in barriers on Chinese imports in columns (1) and (2), respectively.

A bilateral liberalization creates 1.88% welfare gains on average, as indicated by the first line of Table 7. However, these gains are unequal: 2.02% for college graduates and 1.74% for the others. The pro-skilled difference of 0.285 p.p., which constitutes 15.1% of the average gains,⁸³ is mostly driven by the earnings channel (0.268 p.p.); the expenditure channel strengthens the pro-skilled bias by 0.017 p.p. only. The expenditure channel here is a combination of across-industry differences in import spending from Section 3 and within-industry differences from Section 4–5, as described in Appendix E.4.

A reduction of import barriers with China has much smaller average gains of 0.169%, both because

estimator under the [Feenstra \(1994\)](#) and [Broda and Weinstein \(2006\)](#) assumptions. Or it can also be as high as 9.28 in [Eaton and Kortum \(2002\)](#) using a different estimation approach (one based on price differences between countries), although this estimator has been debated by [Simonovska and Waugh \(2014\)](#) on econometric grounds.

⁸²A recent paper by [Redding and Weinstein \(2017\)](#) estimated the elasticities of substitution between 6- and 4-digit NAICS industries to be 1.47 and 1.34, respectively. The estimate by [Hottman and Monarch \(2017\)](#) using 4-digit HS industries 2.78. The range we use covers all of these values, although they are not directly comparable because of differences between our and their demand systems.

⁸³The same difference equals 16.4% of the gains for the non-college group.

this shock affects only one country and because barriers are reduced only on the import side. However, these gains are also biased towards the skilled: 0.196% for the skilled compared with 0.142% for the unskilled. The difference is even larger as a percentage of the average (32.2%); it is primarily driven by the earnings channel and strengthened by the expenditure channel.

The remainder of this section decomposes the average gains and the distributional effects into different mechanisms. Following the logic of the model, we first explain how the average gains are shaped by the average share of import spending and general equilibrium effects related to endogenous wages. We proceed to a similar analysis for the expenditure channel. Finally, we analyze the earnings channel, decomposing it into the mechanisms related to exporting, import competition, imported intermediate inputs, and non-homothetic preferences, as well as general equilibrium forces.

Average Gains. In general equilibrium, the gains from trade depend on how wages respond to changes in trade costs. Equation (7) in Section 2.2 shows that the average welfare gains of trade liberalizations depend on the total share on spending on imports, which governs the benefits from reduced import prices.⁸⁴ Trade-induced growth of the nominal wages relative to the foreign numeraire creates additional gains from trade, because imported final and intermediate products priced in the foreign numeraire become even more affordable (this is not the case for domestic products, the price of which goes up). The growth of the average wage drives these general equilibrium adjustments, although unequal growth of wages across skill groups has an additional welfare impact if the goods consumed domestically differ in skill intensity from the overall economy.

We find that the average wage response in general equilibrium increases the gains from trade by 37.5%, relative to the average gains. Panel (a) of Figure 7 shows that the total of 1.88% welfare gains from a uniform trade liberalization is a combination of 1.37 p.p. due to the reduction of import prices (a direct consequence of 13.7% average spending on imports from Table 1) and 0.51 p.p. due to the average wage increase, while the adjustment through the skill premium is minimal.

Why does the nominal average wage grow after a uniform trade liberalization? As explained in Section 2.2, the average wage is proportional to GDP, which is the sum of value added in all industries. After a reduction of export barriers, exporting industries grow, increasing the average wage, whereas the effect of import barriers is twofold: import competition reduces domestic income, but cheaper intermediate inputs increase it. Figure A10 provides this decomposition, governed by equation (13a). After a 10% trade liberalization, the average wage grows by 1.51% at impact. This effect is magnified by a multiplier: a growing average wage raises domestic demand, although prices also go up, which creates an offsetting force. The multiplier is estimated to be 2.49, hence the average wage grows by 3.75% (37.5% of the trade shock).

The same logic applies when considering a fall of imports barriers for products from China. After this shock, the domestic economy shrinks because of import competition, so average wages fall by 0.27%.

⁸⁴All estimates in this section use the general model with input-output linkages (Appendix A.2). We provide references to the analogous equations in the simplified model in the main text.

This reduces the total welfare gain from 0.206% (corresponding to the 2.06% average import spending on imports from China in Table 1) to 0.169%. In this case, the average wage response lowers the gains from trade by about 18%.

Expenditure Channel. The computation of the distributional effects from the expenditure channel follows the same logic as the average wage (see Step 1 in Section 2.2). It is primarily driven by differential import spending across skill groups. But it also depends on the growth of domestic wages, which makes imports more affordable and disproportionately benefits skills groups that spend more on imports. Using equation (8), we decompose that college graduates get 1.3 basis points (0.013 p.p.) higher benefit due to falling import prices and 0.4 b.p. (37.5% of 1.3 b.p.) through growing average wages. The segregation force operates if the skill premium changes, and consumption baskets of the two groups differ in terms of skill intensity. We find that it is quantitatively negligible.

For the China shock, falling import prices generate a 0.4 b.p. pro-skilled benefit but general equilibrium forces weaken it by a small 0.04 b.p. because the average wage is falling, making *all* (not just Chinese) imports less affordable. In both cases, the expenditure channel favor college graduates but is very small in magnitude.

Earnings Channel. We find that the earnings channel benefits college graduates relatively more. Considering a bilateral trade liberalization, Table 7 shows that additional gains for college workers through the earnings channel are equal to 14.3% of the average gain (or 26.8 basis points of real income after a 10% liberalization). This section quantifies different mechanisms that contribute to this overall distributional effect.

The distributional effects from the earnings channel stem from five forces. The decomposition in Panel (b) of Figure 7 is guided by the results of Section 2.2 (Step 2) and shows the four forces that we investigated in the reduced form in Section 6, in addition to general equilibrium effects. As predicted, intensified import competition, growing exporting opportunities, and income effects are all pro-skilled. Quantification shows that their contributions are similar: 7.0%, 3.9%, and 5.1% of the average welfare gain, respectively. Since college graduates work in industries that use fewer imported inputs, cheaper imported inputs generate a mild offsetting force, equal to -3.1% of the average welfare gain. Finally, general equilibrium forces caused by the rising average wage generate an additional pro-skilled effect, but it is small, equal to 1.3% of the average gain.

Among these mechanisms, import competition deserves additional discussion because it combines several types of demand reallocation, which we have not emphasized in the reduced-form analysis. Panel (c) decomposes import competition effects into various forces, guided by the generalization of (12a) that accounts for input-output linkages (equation (A14)). First, within industries both final and intermediate buyers substitute domestic varieties with imported varieties, which become cheaper. Such reallocation affects industries with higher I-O adjusted import penetration ratios more and, since they have lower skill

intensity, generates pro-skilled effects of 0.51% and 4.43% of the average gain, respectively.⁸⁵ Second, final demand shifts across industries in response to changes in industry-level price indices induced by the trade liberalization. Such reallocation may occur between industries (IO6) within sectors (goods and services), as well as between goods and services as a whole. In our baseline calibration the elasticity of substitution between industries within sectors is set to one, therefore reallocation between industries is mechanically shut down. We find substantial effects from reallocation across sectors: because services are less imported than goods and given that the two sectors are complements ($\rho < 1$), demand for services grows at the expense of goods following trade liberalization, as pointed out by [Cravino and Sotelo \(2017\)](#). This channel generates an additional pro-skilled effect equal to 2.1% of the average gains.

The finding that all import competition forces are relatively small is interesting in the context of the trade literature. In the traditional two-sector, two-factor formulation of the Heckscher-Ohlin model, the “gains” from trade for non-college graduates should be negative. In fact, we find that the gains from trade for this skill groups are positive and that import competition makes their gains lower by only 7.0% relative to the average gains. The reason mirrors our discussion of the differential exposure to import competition in [Section 6.2](#)—offsetting forces between and within goods and services. After trade liberalization, demand is reallocated from domestic goods to domestic services. This reallocation is primarily a consequence of smaller import penetration in services relative to goods, although an additional force is the complementarity between goods and services. At the same time, within sectors demand is also reallocated toward industries such as food—those with low import penetration and, on average, low skill intensity. This evidence is provided by [Table A21](#), which applies the within-between decomposition to the import competition effects and shows that, absent the within-sectoral offsetting pattern, import competition effects would have been substantially larger, around 17.2% of the average gains.

Considering a fall of import barriers with China, the earnings channel is also pro-skilled and moderately larger in magnitude. The pro-skilled earnings channel is 32.2% of the average welfare gain, which amounts to 5.4 basis points of real income after a 10% barrier reduction, and is driven by import competition (25.0% of the average gain) and income effect (8.3%), while the effects of cheaper imported inputs and general equilibrium wage decline are small anti-skilled (-2.3% and -1.1%, respectively).⁸⁶

7.3 Sensitivity to Choice of Elasticities

The earnings channel depends on a number of elasticities: the elasticity of substitution between goods and services (ρ), the elasticity of substitution between industries within sectors (ε), the trade elasticity (ξ), and the aggregate elasticity of substitution between skilled and unskilled labor (σ_{macro}), and there is substantial disagreement about their values. To investigate whether our results are sensitive to the choice of these parameters, we repeat the calibration under the range of elasticity values mentioned in [Section 7.1](#) on the basis of the literature. We examine the sensitivity of the results varying one elasticity at a

⁸⁵The contribution of intermediate demand is larger because the goods sector has more intermediate sales *and* higher import penetration, strengthening the pro-skilled between-sectoral component, while the reverse happens for final demand.

⁸⁶There is no shock to export barriers in this counterfactual.

time and keeping the other elasticities to their baseline values.

Across the relevant range of elasticities, the upper and lower bounds for the distributional effects from the earnings channel are reported in Panel (b) of Table 7. Due to the log-linearization approach, the distributional effects vary monotonically with the elasticities, the upper and lower bounds are reached for the extreme values of the range of elasticities we consider.⁸⁷ Across all parameter values, the earnings channel is always pro-skilled. In the case of a uniform trade liberalization, the lower bound for the distributional effects of the earnings channel is always much larger than our estimate of the distributional effect from the expenditure channel, reported in Table 7. These results lend support to the robustness of the results of Table 7 that the earnings channel is always pro-skilled and quantitatively more important than the expenditure channel.

Besides the *level* of the trade elasticity, its *variation* across industries may also be important.⁸⁸ In ongoing work, we are investigating whether the results are affected by taking the elasticities using the Broda and Weinstein (2006) method from the original paper as well as from Soderbery (2015).

8 Conclusion

This paper has characterized the distributional effects of trade in the United States, taking into account both changes in consumer prices (*expenditure channel*) and in wages (*earnings channel*). Combining theory and empirics in a simple way, we established three results.

First, on the expenditure side, we documented that spending shares on imports, either directly or via imported inputs embedded in domestic goods, are very similar across education and income groups, slightly higher for more educated and richer consumers. This pattern does not result from the fact that households in different education or income groups tend to purchase similar consumption bundles. Rather, a number of forces across and within sectors offset each other: college graduates spend relatively more on (largely non-traded) services but within tradables they spend more on imports (e.g. they spend more on electronics than on food, and they tend to purchase more imported automobiles and food than less-educated consumers).

Second, we documented a series of reduced-form patterns governing the distributional effects of trade via wages. Three forces contribute to an increase in the college wage premium: college educated-workers work in industries that are less exposed to import competition, that export more, and that are more income-elastic. However, we also found a force operating in the other direction: college-educated workers work in industries that rely less on imported inputs.

Third, we combined and assessed the quantitative importance of the reduced-form findings using a simple quantitative trade model, taking into account general equilibrium effects. We found that the

⁸⁷Technically speaking, the denominator—average welfare gains—also varies with elasticities due to general equilibrium effects (equation (7)). However, we verify that variation in average gains is very small across all robustness checks (between 1.80 and 1.90% of consumption), so this does not affect the results.

⁸⁸Ossa (2015) shows that variation *by itself* is important for evaluating gains from trade based on observed trade shares. In our case only the variation systematically related to import penetration, skill intensity, or other industry characteristics plays a role for evaluating the impact of a trade shock.

expenditure channel is distributionally neutral, while the earnings channel implies that the gains from a uniform trade liberalization (applying to all trading partners) are 16% larger for individuals with a college degree, compared to those without. These findings are qualitatively similar when considering more specific changes in trade policy which have recently been debated in the United States, namely an increase in tariffs on imports from China.

Although these results are specific to the United States and to the period under consideration, the method and tools we employed can be readily applied in other contexts, for instance to investigate the likely distributional impacts of other major changes in trade policy such as Brexit.

In future research, we plan on extending both the model and empirics to allow for richer heterogeneity. On the expenditure side, we can exploit the data further to study differences in import spending by gender, age, family structure, and geography. On the earnings side, a model with imperfect mobility of the labor force could capture important differences in the effect of trade on workers in different industries or occupations, e.g. within and outside manufacturing, potentially generating net losses concentrated in some subpopulations of the U.S.

References

- Acemoglu, Daron**, “Technical Change, Inequality, and the Labor Market,” *Journal of Economic Literature*, 2002, 40 (1), 7–72.
- and **David H. Autor**, “Skills, tasks and technologies: Implications for employment and earnings,” in “Handbook of Labor Economics,” Vol. 4, Elsevier B.V., 2011, pp. 1043–1171.
- , **Ufuk Akcigit**, and **William R. Kerr**, “Networks and the Macroeconomy: An Empirical Exploration,” 2015.
- Aguiar, Mark and Mark Bils**, “Has Consumption Inequality Mirrored Income Inequality?,” *American Economic Review*, 2015, 105 (9), 2725–2756.
- Amiti, Mary, Mi Dai, Robert C. Feenstra, and John Romalis**, “How Did China’s WTO Entry Benefit U.S. Consumers?,” 2017.
- Angrist, Joshua**, “Estimating the labor market impact of voluntary military service using social security data on military applicants,” *Econometrica*, 1998, 66 (2), 249–288.
- Antoniades, Alexis and Nicola Zaniboni**, “Exchange Rate Pass-Through into Retail Prices,” *International Economic Review*, 2016, 57 (4), 1425–1447.
- Antràs, Pol, Davin Chor, Thibault Fally, and Russell Hillberry**, “Measuring the Upstreamness of Production and Trade Flows Three Measures of Upstreamness,” *American Economic Review: Papers & Proceedings*, 2012, 102 (3), 412–416.
- Appelbaum, Binyamin**, “Senate Democrats Seek to Outdo Trump on Trade,” *New York Times*, 2017, (<https://nyti.ms/2ukuXIH>).

- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare**, “New Trade Models , Same Old Gains?,” *American Economic Review*, 2012, *102* (1), 94–130.
- , — , **Dave Donaldson, and Andrés Rodríguez-Clare**, “The elusive pro-competitive effects of trade,” 2017.
- Armington, Paul S.**, “A Theory of Demand for Products Distinguished by Place of Production,” *IMF Staff Papers*, 1969, *16* (1), 159–178.
- Artuç, Erhan, Shubham Chaudhuri, and John McLaren**, “Trade shocks and labor adjustment: A structural empirical approach,” *American Economic Review*, 2010, *100* (3), 1008–1045.
- Autor, David H., David Dorn, and Gordon H. Hanson**, “The China Syndrome: Local Labor Market Impacts of Import Competition in the United States,” *American Economic Review*, 2013, *103* (6), 2121–2168.
- , — , — , and **Jae Song**, “Trade Adjustment: Worker-Level Evidence,” *The Quarterly Journal of Economics*, 2014, *129* (4), 1799–1860.
- , **Lawrence F Katz, and AB Krueger**, “Computing inequality: have computers changed the labor market?,” *The Quarterly Journal of Economics*, 1998, (November), 1169–1213.
- Bai, Liang and Sebastian Stumpner**, “Estimating US Consumer Gains from Chinese Imports,” 2017.
- Bems, Rudolfs and Julian Giovanni**, “Income-Induced Expenditure Switching,” *American Economic Review*, 2016, *106* (12), 3898–3931.
- Berman, Eli, John Bound, and Stephen Machin**, “Implications of skill-biased technological change: international evidence,” *The Quarterly Journal of Economics*, 1998, *113* (4), 1245–1279.
- , — , and **Zvi Griliches**, “Changes in the Demand for Skilled Labor Within U. S. Manufacturing: Evidence From the Annual Survey of Manufactures,” *The Quarterly Journal of Economics*, 1994, *109* (2), 367–397.
- Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott**, “Survival of the best fit: Exposure to low-wage countries and the (uneven) growth of U.S. manufacturing plants,” *Journal of International Economics*, 2006, *68* (1), 219–237.
- , — , and — , “Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods,” in “Producer Dynamics: New Evidence from Micro Data” 2009, pp. 513–552.
- , — , **Stephen J. Redding, and Peter K. Schott**, “Firms in international trade,” *Journal of Economic Perspectives*, 2007, *21* (3), 105–130.
- , — , — , and — , “Global Firms,” *Journal of Economic Literature*, 2018.
- , **Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum**, “Plants and Productivity in International Trade,” *American Economic Review*, 2003, *93* (4), 1268–1290.
- Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, and Ron Jarmin**, “What Drives Differences in Management?,” 2016.
- Bohlen, Celestine**, “A Global Push to Address Trade and Inequality,” *New York Times*, 2016, (<https://nyti.ms/2l95tar>).

- Boppart, Timo**, “Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences,” *Econometrica*, 2014, 82 (6), 2167–2196.
- Borjas, George J, Richard B Freeman, and Lawrence F Katz**, “How Much Do Immigration and Trade Affect Labor Market Outcomes?,” *Brookings Papers on Economic Activity*, 1997.
- Breinlich, Holger, Elsa Leromain, Dennis Novy, and Thomas Sampson**, “The Brexit Vote and Inflation,” 2017.
- Broda, Christian and David E. Weinstein**, “Globalization and the Gains From Variety,” *The Quarterly Journal of Economics*, 2006, 121 (2), 541–585.
- and **John Romalis**, “Inequality and Prices: Does China Benefit the Poor in America?,” *Mimeo*, 2008.
- Buera, Francisco J, Joseph P Kaboski, and Richard Rogerson**, “Skill Biased Structural Change,” *NBER Working Paper*, 2015.
- Burstein, Ariel and Jonathan Vogel**, “International Trade, Technology, and the Skill Premium,” *Journal of Political Economy*, 2017, 125 (5), 1356–1412.
- Caliendo, Lorenzo and Fernando Parro**, “Estimates of the trade and welfare effects of NAFTA,” *Review of Economic Studies*, 2015, 82 (1), 1–44.
- , **Maximiliano Dvorkin, and Fernando Parro**, “Trade and Labor Market Dynamics,” 2017, (December).
- Caron, Justin, Thibault Fally, and Ana Cecília Fieler**, “Home-Market Effects on Innovation,” 2015.
- , — , and **James R Markusen**, “International trade puzzles: A solution linking production and preferences,” *Quarterly Journal of Economics*, 2014, 129 (3), 1501–1552.
- , — , and **James R. Markusen**, “Per Capita Income and the Demand for Skills,” 2017.
- Chetty, Raj and Nathaniel Hendren**, “The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates,” 2017.
- Chodorow-Reich, Gabriel**, “Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?,” 2017.
- Comin, Diego A, Danial Lashkari, and Martí Mestieri**, “Structural Change with Long-run Income and Price Effects,” 2016.
- Costinot, Arnaud and Andrés Rodríguez-Clare**, *Trade Theory with Numbers: Quantifying the Consequences of Globalization*, Vol. 4 2015.
- Cravino, Javier and Andrei A Levchenko**, “The distributional consequences of large devaluations,” 2017.
- and **Sebastian Sotelo**, “Trade Induced Structural Change and the Skill Premium,” 2017.
- Dawkins, Christina, T. N. Srinivasan, and John Whalley**, “Chapter 58 Calibration,” *Handbook of Econometrics*, 2001, 5, 3653–3703.

- De Loecker, Jan and Frederic Warzynski**, “Markups and Firm-Level Export Status,” *American Economic Review*, 2012, 102 (6), 2437–2471.
- , **Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik**, “Prices, Markups, and Trade Reform,” *Econometrica*, 2016, 84 (2), 445–510.
- Deaton, Angus**, “Rice Prices and Income Distribution in Thailand: A Non-Parametric Analysis,” *The Economic Journal*, 1989, 99 (395), 1–37.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum**, “Global Rebalancing with Gravity: Measuring the Burden of Adjustment,” *IMF Staff Papers*, 2008, 55 (3), 511–540.
- DeNavas-Walt, Carmen, Bernadette D. Proctor, and Jessica C. Smith**, “Income, Poverty, and Health Insurance Coverage in the United States: 2010,” *U.S. Census Bureau Current Population Reports*, 2011.
- DeSalvo, Bethany, Frank F. Limehouse, and Shawn D. Klimek**, “Documenting the Business Register and Related Economic Business Data,” 2016.
- Dhingra, Swati, Hanwei Huang, Gianmarco I. P. Ottaviano, Joao Paulo Pessoa, Thomas Sampson, and John Van Reenen**, “The Costs and Benefits of Leaving the EU: Trade Effects,” 2017.
- Duncan, Otis Dudley and Beverly Duncan**, “Residential Distribution and Occupational Stratification,” *American Journal of Sociology*, 1955, 60 (5), 493–503.
- Eaton, Jonathan and Samuel Kortum**, “Technology, geography, and trade,” *Econometrica*, 2002, 70 (5), 1741–1779.
- Epifani, Paolo and Gino Gancia**, “The skill bias of world trade,” *Economic Journal*, 2008, 118 (530), 927–960.
- Faber, Benjamin**, “Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices,” 2014.
- Fajgelbaum, Pablo and Amit K. Khandelwal**, “Measuring the Unequal Gains from Trade,” *Quarterly Journal of Economics*, 2016, 131 (3), 1113–1180.
- , **Gene M. Grossman, and Elhanan Helpman**, “Income Distribution, Product Quality, and International Trade,” *The Journal of Political Economy*, 2011, 119 (4).
- Feenstra, Robert C.**, “New Product Varieties and the Measurement of International Prices,” *American Economic Review*, 1994, 84 (1), 157–176.
- and **Gordon H. Hanson**, “Globalization, Outsourcing, and Wage Inequality,” *American Economic Review: Papers & Proceedings*, 1996, 86 (2), 240–245.
- Felbermayr, Gabriel, Benjamin Jung, and Mario Larch**, “The welfare consequences of import tariffs: A quantitative perspective,” *Journal of International Economics*, 2015, 97 (2), 295–309.
- Furman, Jason, Katheryn Russ, Jay Shambaugh, and Jason Furman**, “US tariffs are an arbitrary and regressive tax,” *VOX CEPR’s Policy Portal*, 2017.

- Gabaix, Xavier**, “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 2011, 79 (3), 733–772.
- Galle, Simon, Andrés Rodríguez-clare, and Moises Yi**, “Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade,” 2017.
- Garner, Thesia I, Robert McClelland, and William Passero**, “Strengths and Weaknesses of the Consumer Expenditure Survey from a BLS Perspective,” 2009.
- Haberman, Maggie**, “Donald Trump Says He Favors Big Tariffs on Chinese Exports,” *New York Times*, 2016.
- Hallak, Juan Carlos and Peter K. Schott**, “Estimating cross-country differences in product quality,” *Quarterly Journal of Economics*, 2011, 126 (1), 417–474.
- Hanoch, Giora**, “Production and Demand Models with Direct or Indirect Implicit Additivity,” *Econometrica*, 1975, 43 (3), 395–419.
- He, Zheli and Feiran Zhang**, “Trade and Real Wages of the Rich and Poor : Cross-Country Evidence,” 2017.
- Horowitz, Karen J. and Mark A. Planting**, *Concepts and Methods of the U.S. Input-Output Accounts*, Washington, D.C.: Bureau of Economic Analysis, Department of Commerce, 2009.
- Hottman, Colin J and Ryan Monarch**, “Estimating Unequal Gains across U . S . Consumers with Supplier Trade Data,” 2017.
- Hummels, David and Peter J Klenow**, “The Variety and Quality of a Nation’s Trade,” *American Economic Review*, 2005, 8712.
- Jaimovich, Nir, Sergio Rebelo, and Arlene Wong**, “Trading Down and the Business Cycle,” *NBER Working Papers*, 2015, September (21539).
- Jaravel, Xavier**, “The Unequal Gains from Product Innovations: Evidences from the US Retail Sector,” 2017.
- Jones, Ronald W.**, “The Structure of Simple General Equilibrium Models,” *Journal of Political Economy*, 1965, 73 (6), 557.
- Kamal, Fariha and Ryan Monarch**, “Identifying Foreign Suppliers in U.S. Merchandise Import Transactions,” *CES working paper*, 2016.
- Katz, Lawrence F and Kevin M Murphy**, “Changes in Relative Wages, 1963-1987 - Supply and Demand Factors,” *Quarterly Journal of Economics*, 1992, 107 (1), 35–78.
- Khandelwal, Amit K.**, “The long and short (of) quality ladders,” *Review of Economic Studies*, 2010, 77 (4), 1450–1476.
- Krugman, Paul**, “Technology, trade and factor prices,” *Journal of International Economics*, 2000, 50 (1), 51–71.
- Lawrence, Robert Z and Matthew J Slaughter**, “International Trade and American Wages in the 1980s: Giant Sucking Sound or Small Hiccup?,” *Brookings Papers: Microeconomics*, 1993, (2), 161–226.
- Lebow, David E and Jeremy B Rudd**, “Measurement Error in the Consumer Price Index: Where Do We Stand?,” *Journal of Economic Literature*, 2003, 41 (1), 159.

- Leonardi, Marco**, “The Effect of Product Demand on Inequality: Evidence from the US and the UK,” 2015, 7 (3), 221–247.
- Levell, Peter, Martin O Connell, and Kate Smith**, “The exposure of households’ food spending to tariff changes and exchange rate movements and exchange rate movements,” *IFS Briefing Note BN213*, 2017.
- Levinson, Arik and James O Brien**, “Environmental Engel Curves,” 2016.
- Matsuyama, Kiminori**, “Engel’s Law in the Global Economy: Demand-Induced Patterns of Structural Change and Trade across Countries,” 2017, pp. 1–46.
- Morrow, Peter M and Daniel Treffer**, “Endowments, Skill-Biased Technology, and Factor Prices: A Unified Approach to Trade,” 2017.
- Nigai, Sergey**, “On Measuring the Welfare Gains from Trade Under Consumer Heterogeneity,” *Economic Journal*, 2016, 126 (593), 1193–1237.
- Oberfield, Ezra and Devesh Raval**, “Micro Data and Macro Technology,” 2014.
- Ossa, Ralph**, “Why trade matters after all,” *The Journal of International Economics*, 2015, p. 17.
- Pierce, Justin R and Peter K. Schott**, “A Concordance Between Ten-Digit U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries,” *Journal of Economic and Social Measurement*, 2012, 37 (1,2), 61–96.
- Porter, Eduardo**, “Making Mexico Pay for Border Wall: Ideas Abound, and So Do Obstacles,” *New York Times*, 2017, (<https://nyti.ms/2jD20Ng>).
- Porto, Guido G.**, “Using survey data to assess the distributional effects of trade policy,” *Journal of International Economics*, 2006, 70 (1), 140–160.
- Redding, Stephen J. and David E Weinstein**, “Aggregating from Micro to Macro Patterns of Trade,” *NBER Working Paper*, 2017.
- Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek**, “Integrated Public Use Microdata Series: Version 6.0,” 2015.
- Schott, Peter K.**, “Across-Product versus Within-Product Specialization in International Trade,” *The Quarterly Journal of Economics*, 2004, 119 (2), 647–678.
- , “The relative sophistication of Chinese exports,” *Economic Policy*, 2008, 23 (53), 6–49.
- Simonovska, Ina and Michael E Waugh**, “The elasticity of trade : Estimates and evidence,” *Journal of International Economics*, 2014, 92 (1), 34–50.
- Soderbery, Anson**, “Estimating import supply and demand elasticities : Analysis and implications,” *Journal of International Economics*, 2015, 96 (1), 1–17.
- Stolper, Wolfgang F. and Paul A. Samuelson**, “Protection and Real Wages,” *The Review of Economic Studies*, 1941, 9 (1), 58–73.

Theil, Henri, *Theory and measurement of consumer demand*, Elsevier Science Publishing Co Inc., 1975.

Timmer, Marcel P., Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries, “An Illustrated User Guide to the World Input-Output Database: the Case of Global Automotive Production,” *Review of International Economics*, 2015, *23* (3), 575–605.

Tintelnot, Felix, Ayumu Ken Kikkawa, Magne Mogstad, and Emmanuel Dhyne, “Trade and Domestic Production Networks,” 2017.

Traiberman, Sharon, “Occupations and Import Competition Evidence from Danish Matched Employee-Employer Data,” 2016.

Wasi, Nada and Aaron Flaaen, “The Stata Journal,” *Stata Journal*, 2015, *15* (3), 672–697.

Wood, Adrian, “How Trade Hurt Unskilled Workers,” *Journal of Economic Perspectives*, 1995, *9* (3), 57–80.

Table 1: Spending on Imports by Education Group, Industry Data

	Imports by Trading Partners							
	All Countries		China		NAFTA		Developed Economies	
	Total (1)	Direct (2)	Total (3)	Direct (4)	Total (5)	Direct (6)	Total (7)	Direct (8)
All, %	13.69	7.28	2.06	1.53	2.79	0.98	3.46	2.15
College, %	13.30	7.12	2.11	1.58	2.63	0.92	3.26	1.99
Non-college, %	13.95	7.39	2.02	1.50	2.90	1.02	3.59	2.27
College minus non-college, p.p.	-0.66 (0.18)	-0.28 (0.15)	+0.09 (0.05)	+0.08 (0.05)	-0.27 (0.04)	-0.09 (0.02)	-0.33 (0.05)	-0.28 (0.05)
<i>as % of avg. import spending</i>	<i>-4.79</i>	<i>-3.83</i>	<i>4.48</i>	<i>5.48</i>	<i>-9.60</i>	<i>-9.63</i>	<i>-9.61</i>	<i>-12.90</i>
→ Between goods and services	-1.47 (0.16)	-1.17 (0.13)	-0.29 (0.03)	-0.27 (0.03)	-0.30 (0.03)	-0.17 (0.02)	-0.43 (0.05)	-0.38 (0.04)
→ Within goods and services	+0.81 (0.05)	+0.89 (0.05)	+0.38 (0.03)	+0.36 (0.03)	+0.03 (0.01)	+0.08 (0.01)	+0.10 (0.03)	+0.11 (0.03)
→ Between subsectors	+0.65 (0.05)	+0.74 (0.04)	+0.43 (0.03)	+0.40 (0.03)	+0.03 (0.01)	+0.07 (0.01)	+0.04 (0.02)	+0.03 (0.02)
→ Within subsectors	+0.16 (0.03)	+0.15 (0.03)	-0.04 (0.02)	-0.05 (0.02)	+0.01 (0.01)	+0.01 (0.01)	+0.07 (0.02)	+0.08 (0.02)

Notes: This table estimates the fraction of imports in expenditure across education groups using the industry-level data from Section 3. Total import spending includes consumption of imported products (direct import spending in 169 industries with positive final consumption) and imported intermediate inputs embedded in domestic products (indirect import spending, measured using the input-output linkages across 380 industries in the input-output table). The table reports the average and differential spending in the entire economy and decomposes the difference into the within- and between- components according to equation (A24). NAFTA stands for Canada and Mexico, whereas Developed Economies are OECD members (excluding NAFTA), Taiwan, and Singapore. Subsectors are shown in Table A2. Standard errors are shown in parentheses.

Table 2: Imports in Merged Nielsen-Census Sample, % of Firms' Sales

	All Products	By Product Class		
		Food	Health & Household	General Merchandize
Total Imports	11.10	6.92	14.58	27.96
Imports from China	4.15	0.88	6.51	17.91
Imports from NAFTA	1.91	1.67	2.19	2.74
Imports from Developed Economies	3.10	2.42	4.24	4.90
% of Firm-Module Sales to College Graduates	31.18	31.34	30.71	31.12
(st.dev.)	(12.62)	(13.54)	(11.05)	(9.44)
% of Product Class in Total Sales	100.00	67.29	20.24	12.48
<i>N</i> firms	8,200	5,700	2,400	2,000
<i>N</i> firm-years	12,700	9,000	3,700	2,800
<i>N</i> firm-module-years	131,000	88,600	29,800	12,500

Notes: This table reports statistics on imports based on the merged Nielsen-Census sample from Section 4, for all products and for three product classes: Food, Alcohol, and Tobacco (“Food”), Health and Beauty Products and Household Supplies (“Health and household”), and General Merchandize. Imports are measured at the firm level and the summary statistics are computed using the square-root of firms’ Nielsen sales as weights. The reported percentage of each product class uses the same weighting scheme. When computing the percentage of firm-module sales to college graduates, weights are decomposed across barcodes of the same firm proportionally to sales. See Appendix Table A11 for statistics with an alternative weighting scheme. Observations are firm-module-year cells and the numbers of observations are rounded to the nearest 100 to preserve confidentiality.

Table 3: Spending on Imports by Education Group, Merged Nielsen-Census Sample

	All Imports (1)	Imports by Trading Partner		
		China (2)	NAFTA (3)	Developed Economies (4)
All, %	11.10	4.16	1.91	3.10
College, %	11.50	4.02	1.95	3.37
Non-college, %	10.91	4.20	1.86	2.82
College minus non-college, p.p.	+0.59 (0.12)	-0.18 (0.07)	+0.09 (0.04)	+0.55 (0.06)
<i>as % of avg. import spending</i>	<i>5.35</i>	<i>-4.37</i>	<i>4.61</i>	<i>17.63</i>
→ Within industries	+0.48 (0.11)	-0.10 (0.05)	+0.06 (0.03)	+0.38 (0.06)
<i>as % of avg. import spending</i>	<i>4.34</i>	<i>-2.44</i>	<i>3.20</i>	<i>12.27</i>
→ Within product modules	+0.28 (0.06)	-0.14 (0.03)	+0.03 (0.02)	+0.28 (0.04)
<i>as % of avg. import spending</i>	<i>2.56</i>	<i>-3.03</i>	<i>1.47</i>	<i>9.17</i>
<i>N</i> firm-years	12,700	12,700	12,700	12,700

Notes: This table reports the fraction of imports in expenditure for different education groups using the merged Nielsen-Census sample from Section 4. Importing is proxied by the share of total imports in firm sales. Differential spending on imports is decomposed into “within” and “between” components for 6-digit I-O codes (“industries”) and for Nielsen product modules (“product modules”) according to equation (A24). Firms are weighted by the square-root of Nielsen sales. Standard errors are shown in parentheses.

Table 4: Spending on Imports by Education Groups, Automobile Sample

	Imports by Trading Partner			Imports by Purchase Type	
	All Countries (1)	Outside NAFTA (2)	NAFTA (3)	New (4)	Used (5)
All, %	49.66	25.47	24.19	50.23	49.30
College, %	53.82	32.68	21.14	53.30	54.27
Non-college, %	47.24	21.26	25.98	47.59	47.06
College minus non-college, p.p.	+6.58	+11.41	-4.84	+5.71	+7.21
	(0.27)	(0.36)	(0.22)	(0.38)	(0.36)
<i>as % of avg. import spending</i>	<i>+13.25</i>	<i>+44.82</i>	<i>-19.99</i>	<i>+11.37</i>	<i>+14.62</i>
<i>N</i> auto purchases	51,498	51,498	51,498	19,617	31,610
<i>N</i> brands	39	39	39	38	39

Notes: This table reports the shares of purchases of imported cars as a fraction of total car purchases, by education group. The merged dataset measures imports at the level of brands, linking the CEX to Ward’s data, as described in Section 5. Columns (1) to (3) use the full sample of purchases and distinguish between imports from NAFTA and other countries. Columns (4) and (5) decompose the sample into purchases of new and used cars (excluding CEX purchases with missing information on whether the vehicle is new or used). All brands are listed in Table A15 (for Austin-Healey, only used purchases are observed). Standard errors are shown in parentheses.

Table 5: Relationship between Consumer Education and Imports of Assembled Cars and Car Parts across Manufacturers

	Imports as % of Car Sales		Imports as % of Car Sales	
	Assembled Cars Only (1)	Assembled Cars & Cars Parts (2)	Assembled Cars Only (3)	Assembled Cars & Cars Parts (4)
% of New Cars Sold to College Graduates	1.031 (0.326)	0.918 (0.317)		
% of Used Cars Sold to College Graduates			1.565 (0.305)	1.425 (0.290)
<i>N</i> firms	20	20	20	20

Notes: This table shows that imports of car parts do not create large biases for the differential spending of college and non-college consumers on imported cars. The dependent variables in OLS regressions are the shares of imports of assembled cars (“Assembled Cars Only”) or of both assembled cars and car parts (“Assembled Cars & Cars Parts”) in the value of car sales. They are computed using the Customs microdata and the Census of Manufactures at the firm level, as described in Section 5.3. The independent variable is the fraction of sales of each firm to college graduates in the CEX sample of car purchases, separately for new cars in columns (1) and (2) and used cars in columns (3) and (4). Each regression is weighted by the number of purchases. The coefficient magnitudes are comparable to the slopes in Figure 4. The sample size is rounded to the nearest 10 to protect confidentiality. Robust standard errors are shown in parentheses.

Table 6: Exposure to Labor Market Effects of Trade by Education Group, Industry Data

	Payroll-weighted Averages							
	Import Penetration		Export Share		Imported Inputs Share		Income Elasticity $\times 100$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All workers, %	4.21	8.15	4.66	9.21	7.08	13.73	109.47	108.28
College-educated workers, %	4.13	7.85	4.90	9.42	6.35	12.69	115.04	113.62
Non-college educated workers, %	4.28	8.45	4.43	9.01	7.79	14.75	103.96	103.01
College minus non-college, p.p.	-0.15	-0.60	+0.47	+0.41	-1.44	-2.07	+11.08	+10.61
	(0.04)	(0.05)	(0.03)	(0.04)	(0.02)	(0.04)	(0.11)	(0.10)
<i>as % of avg.</i>	<i>-3.54</i>	<i>-7.33</i>	<i>10.08</i>	<i>4.48</i>	<i>-20.39</i>	<i>-15.05</i>	<i>10.12</i>	<i>9.80</i>
→ Between goods and services	-1.54	-2.04	-0.92	-1.21	-0.80	-1.28	+1.02	+0.86
	(0.03)	(0.03)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.01)
→ Within goods and services	+1.39	+1.45	+1.39	+1.62	-0.65	-0.79	+10.06	+9.75
	(0.02)	(0.04)	(0.02)	(0.03)	(0.01)	(0.04)	(0.11)	(0.10)
→ Between subsectors	+0.99	+1.22	+0.82	+1.17	-0.61	-0.44	+7.57	+7.48
	(0.02)	(0.03)	(0.02)	(0.03)	(0.01)	(0.02)	(0.09)	(0.09)
→ Within subsectors	+0.40	+0.22	+0.57	+0.45	-0.04	-0.35	+2.49	+2.27
	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)	(0.07)	(0.06)
Adjusted for I-O linkages	No	Yes	No	Yes	No	Yes	No	Yes

Notes: This table reports the payroll-weighted averages of several industry characteristics, overall and for college- and non-college educated workers, using the industry-level data from Section 6, which covers 380 industries. It also decomposes the difference between education groups into the within and between components for sectors (goods and services) and subsectors (listed in Table A2), according to equation (A24). The outcomes are imports as % of absorption, exports as % of industry output, imports of intermediate inputs as % of output, and income elasticities. Even columns account for imports, exports, imported inputs, and income elasticities in downstream industries (see Section 2.3 for details).



Table 7: Calibration of the Welfare Effects of Trade Liberalizations

(a) Main Results		
	10% Reduction in Trade Barriers	
	All Import and Export Barriers (1)	Import Barriers with China (2)
<i>Average welfare effects, equivalent variation, % of spending</i>		
All	1.879	0.169
College	2.024	0.196
Non-College	1.740	0.142
<i>Distributional effects, college minus non-college, p.p. [as % of avg. welfare effect]</i>		
Overall	+0.285 [15.1%]	+0.054 [32.2%]
→ Expenditure channel, pro-skilled	+0.017 [0.9%]	+0.004 [2.2%]
→ Earnings channel, pro-skilled	+0.268 [14.3%]	+0.050 [29.9%]

(b) Sensitivity of the Earnings Channel to Elasticities		
	10% Reduction in Trade Barriers	
	All Import and Export Barriers (1)	Import Barriers with China (2)
<i>Distributional effects from the earnings channel, college minus non-college, % of avg. welfare effect</i>		
Baseline: $\xi = 3.5$, $\varepsilon = 1$, $\rho = 0.2$, $\sigma_{\text{macro}} = 1.41$	+14.3	+29.9
→ Varying ξ from 1.9 to 5.1	+11.6 to +16.9	+17.0 to +43.4
→ Varying ε from 0.8 to 3.5	+14.0 to +17.1	+29.0 to +42.3
→ Varying ρ from 0 to 1	+15.1 to +11.0	+30.9 to +26.0
→ Varying σ_{macro} from 1.41 to 1.8	+14.3 to +11.2	+29.9 to +23.4


Notes: This table calibrates the welfare effects of trade liberalizations across education groups using the model from Section 2 and the reduced-form patterns in Sections 3–6. Panel (a) reports the welfare effects in terms of equivalent variation, expressed as a percentage of initial consumption spending for each education group. The distributional effects are decomposed into the expenditure channel and the earnings channel according to equation (3). Panel (b) shows the sensitivity of the distributional effects from the earnings channel to the choice of structural elasticities. Using the notation from the model in Section 2, ξ denotes the trade elasticity, ε the elasticity of substitution between industries (6-digit I-O codes), ρ the elasticity of substitution between sectors (goods and services), and σ_{macro} the elasticity of substitution between skilled and unskilled labor (workers with and without a college degree). For each of this elasticities, we consider a plausible range of values in light of the literature (see the main text for a complete discussion).

Figure 1: Industry Import Shares and Consumer Base



Notes: The binned scatterplot in Panel (a) groups six-digit I-O industries within each sector into bins by consumer base (% of industry sales to college graduates) and reports the share of total (direct plus indirect) imports from all countries in final expenditures in these industries. In Panel (b), each circle corresponds to a subsector from Table A2, and the circle size indicates final spending. Subsectors that account for less than 3% of the sectoral expenditure are not shown. Industry-level data from Section 3 are employed in both panels.


Figure 2: Import Shares and Consumer Base, Merged Nielsen-Census Sample




Notes: These binned scatterplots group firm-module-year cells in 20 bins by the fraction of Nielsen sales to college graduates. The vertical axis shows the average share of imports in sales measured at the firm level in the Census. Firms are weighted by the square-root of their Nielsen sales, and weights are decomposed across barcodes of the same firm proportionally to sales. Fixed effects of I-O industries by year are absorbed.

Figure 3: The Role of Product Quality


(a) Prices and Consumer Base




(b) Prices and Imports Excluding China




(c) Prices and Imports from Developed Economies



(d) Prices and Imports from China, Health & Household only




(e) Prices and Imports from China, General Merchandise only




Notes: These binned scatterplots show average % of sales to college graduates and import shares by decile of barcode prices within their respective product modules. Import shares are computed at the firm-level. The analysis is performed on the sample of firm-year-module-decile cells. Product modules which include barcodes with quantity measured in different units (e.g. ounces vs. counts) are decomposed by measurement unit. Firms are weighted by the square-root of their Nielsen sales, and weights are decomposed across barcodes of the same firm proportionally to sales. Fixed effects of modules by year are absorbed, which explains that there are more than ten bins.

Figure 4: Imports Shares and Consumer Base across Auto Brands

(a) Imports excluding NAFTA




(b) Total Imports



Notes: Each circle corresponds to a brand of automobile from Table A15. The import shares on the vertical axis are based on the Ward's data, aggregated from models into brands, and the shares of cars of each brand sold to college graduates on the horizontal axis is from CEX. The size of each circle indicates the number of purchases in the CEX data. Brands that account for less than 35 purchases are not shown.


Figure 5: Imports and Skill Intensity across Industries




Notes: The binned scatterplot in Panel (a) groups six-digit I-O industries within each sector into bins by skill-intensity (payroll share of college graduates) and reports the import penetration (the share of direct imports from all countries in absorption) in these industries. In Panel (b), each circle corresponds to a subsector from Table A2, and the circle size indicates payroll (subsectors that account for less than 3% of the sectoral payroll are not shown). Industry-level data from Section 6 are used in both panels.

Figure 6: Additional Industry-Level Outcomes and Skill Intensity


(a) I-O Adjusted Export Share as % of Output



(b) Imported Intermediate Inputs as % of Output




(c) I-O Adjusted Income Elasticity




Notes: This figure uses industry-level data from Section 6 to show the relationship between skill intensity (payroll share of college graduates) and three outcomes: the share of exports (including exports that happen through domestic customers) in industry output in Panel (a), the share of imported intermediate inputs in industry output in Panel (b), and the weighted average income elasticity corresponding to the final demand in the industry and its domestic customers in Panel (c). The outcomes are measured according to the model in Appendix A.2. Each circle corresponds to a subsector from Table A2, and the circle size indicates total payroll. Subsectors that account for less than 3% of the sectoral payroll are not shown. For clarity, Panel (c) only labels the Education subsector that is most important to understand the patterns.

Figure 7: Decomposition of the Welfare Effects of Trade Liberalizations


(a) Decomposition of Welfare Effects from a 10% Fall in Import and Export Barriers



(b) Decomposition of Differential Wage Effects



(c) Decomposition of Import Competition Effects



Notes: This table decomposes the welfare effects of trade liberalizations using the model from Section 2 and the reduced-form patterns in Sections 3–6. Panel (a) reports average welfare gains in terms of equivalent variation, expressed as a percentage of total consumption, and decomposes them into effects coming from cheaper imports, higher wages and higher college-wage premium. Panel (b) decomposes the distributional effects from the earnings channel, and Panel (c) decomposes the distributional effects from intensified import competition.

Appendix

A Theory Appendix

A.1 More on Model without Input-Output Linkages

Equilibrium Conditions. The equilibrium is defined by a set of quantities and prices that satisfy five standard conditions:

1. Profit maximization by domestic and foreign producers;
2. Utility maximization for each type of domestic agents: $\max_{\{Q_{jc}^i\}} \mathcal{U}$ s.t. $\sum_{j,c} p_{jc} Q_{jc}^i = \zeta w_i$;
3. Export demand: $Q_{jH}^{\text{Export}} = a_j^{\text{Export}} \left(p_{jH} \tau_j^* \right)^{-\xi_j}$ for each j ;
4. Product market clearing for domestic varieties: $Q_{jH} = L_S Q_{jH}^S + L_U Q_{jH}^U + Q_{jH}^{\text{Export}} = F_j \left(L_S^j, L_U^j \right)$ for each j ;
5. Labor market clearing for each type of agents: $\sum_j L_i^j = L_i$ for $i = S, U$.

Proof of (10). We start from the labor market clearing condition, expressed in value terms. For the skilled group, $w_S L_S = \sum_j w_S L_S^j = \sum_j V A_j \cdot v_j$, which in log-changes becomes

$$\hat{w}_S = \sum_j e_S^j \cdot \left(\widehat{V A_j} + \hat{v}_j \right). \quad (\text{A1a})$$

Similarly for the unskilled, $w_U L_U = \sum_j w_U L_U^j = \sum_j V A_j \cdot (1 - v_j)$, thus

$$\hat{w}_U = \sum_j e_U^j \cdot \left(\widehat{V A_j} + \widehat{1 - v_j} \right). \quad (\text{A1b})$$

To solve for the change in the payroll shares of the skilled and unskilled groups ($\hat{v}_j$ and $\widehat{1 - v_j}$, respectively), we note that $v_j / (1 - v_j) = w_S L_S^j / w_U L_U^j$. By definition of the local elasticity of substitution σ_j , this implies:

$$\left(\widehat{\frac{v_j}{1 - v_j}} \right) = (1 - \sigma_j) (\hat{w}_S - \hat{w}_U).$$

Expanding the left-hand side using the standard log-differentiation rules, we obtain:⁸⁹

$$\begin{aligned} \hat{v}_j &= (1 - \sigma_j) (1 - v_j) (\hat{w}_S - \hat{w}_U), \\ \widehat{1 - v_j} &= - (1 - \sigma_j) v_j (\hat{w}_S - \hat{w}_U). \end{aligned}$$

⁸⁹We use $\partial \log(z / (1 - z)) / \partial \log z = 1 / (1 - z)$ and $\partial \log(1 / (1 - z)) / \partial \log z = -z / (1 - z)$.

Plugging these into (A1a) and (A1b) and taking the difference between skilled and unskilled groups yields

$$\hat{w}_S - \hat{w}_U = \sum_j \left(e_S^j - e_U^j \right) \widehat{VA}_j + \sum_j (1 - \sigma_j) \left(e_S^j (1 - v_j) + e_U^j v_j \right) \cdot (\hat{w}_S - \hat{w}_U). \quad (\text{A2})$$

Denoting the value added share of industry j by $e_j = VA_j / GDP$ (where $GDP = \sum_j VA_j$), we can represent e_j^S as

$$e_j^S = \frac{w_S L_S^j}{w_S L_S} = \frac{VA_j}{GDP} \cdot \frac{w_S L_S^j / VA_j}{w_S L_S / GDP} = e_j \cdot \frac{v_j}{\bar{v}},$$

and similarly $e_j^U = e_j (1 - v_j) / (1 - \bar{v})$. This implies

$$e_S^j (1 - v_j) + e_U^j v_j = e_j \left(\frac{v_j (1 - v_j)}{\bar{v}} + \frac{v_j (1 - v_j)}{1 - \bar{v}} \right) = e_j \frac{v_j (1 - v_j)}{\bar{v} (1 - \bar{v})}.$$

Plugging this into (A2), we finally obtain (10):

$$\begin{aligned} \hat{w}_S - \hat{w}_U &= \Delta_{VA} \left[\widehat{VA}_j \right] - \mathbb{E}_{VA} \left[\frac{v_j (1 - v_j)}{\bar{v} (1 - \bar{v})} (\sigma_j - 1) \right] \cdot (\hat{w}_S - \hat{w}_U) \\ &= \frac{\Delta_{VA} \left[\widehat{VA}_j \right]}{1 + \mathbb{E}_{VA} \left[\frac{v_j (1 - v_j)}{\bar{v} (1 - \bar{v})} \cdot (\sigma_j - 1) \right]} \equiv \frac{\Delta_{VA} \left[\widehat{VA}_j \right]}{\sigma_{\text{within}}}. \end{aligned}$$

We note that the weights applied to $\sigma_j - 1$ aggregate to

$$\mathbb{E}_{VA} \left[\frac{v_j (1 - v_j)}{\bar{v} (1 - \bar{v})} \right] = \frac{\bar{v} - \bar{v}^2 - \text{Var}[v_j]}{\bar{v} (1 - \bar{v})} = 1 - \text{Segm}_{\text{prod}},$$

where $\text{Segm}_{\text{prod}} = \text{Var}[v] / \bar{v} (1 - \bar{v})$ is the production segmentation index, which measures the heterogeneity of industries by skill intensity. This proves that σ_{within} is higher when skill intensity is more homogenous across industries (production segmentation is lower), provided that labor types are substitutes ($\sigma_j > 1$).

Proof of (11)–(12). From the market clearing condition for domestic products, the value of domestic output (or equivalently value added) combines domestic and export sales: $VA_j = X_{jH}^{\text{Final}} + X_j^{\text{Export}}$, where $X_{jH}^{\text{Final}} = X_{jH}^S + X_{jH}^U$. In log-differences,

$$\begin{aligned} \widehat{VA}_j &= \text{Dom share}_j \cdot \hat{X}_j^{\text{Final}} + \text{Export share}_j \cdot \hat{X}_j^{\text{Export}} \quad \text{and} \\ \hat{X}_j^{\text{Final}} &= \mu_j \hat{X}_{jH}^S + (1 - \mu_j) \hat{X}_{jH}^U. \end{aligned} \quad (\text{A3})$$

We combine the trade-induced price changes from Step 1 with the demand system to solve for the domestic and foreign consumers' demand changes. First, from the assumption on foreign demand and

(5),

$$\begin{aligned}\hat{X}_j^{\text{Export}} &= (1 - \xi_j) (\hat{p}_{jH} + \hat{\tau}^*) \\ &= (1 - \xi_j) (\hat{w} + \hat{\tau}^* + (v_j - \bar{v}) (\hat{w}_S - \hat{w}_U)).\end{aligned}\quad (\text{A4})$$

Second, Appendix A.3 derives the change in the spending of a group- i consumer, who has non-homothetic nested CES preferences, on the industry j 's composite good (see equation (A21)). Combining it with CES preferences over varieties from different countries and using $\hat{X}_i = \hat{w}_i$, we obtain

$$\hat{X}_{jH}^i = \hat{w}_i + (1 - \xi_j) (\hat{p}_{jH} - \hat{p}_j) + (1 - \varepsilon_r) (\hat{p}_j - \hat{p}_{ir}) + (1 - \rho) (\hat{p}_{ir} - \hat{\pi}_i) + (\psi_j - 1) (\hat{w}_i - \hat{\pi}_i) \quad (\text{A5})$$

where $\hat{p}_{ir} = \sum_{j \in r} s_{j|r}^i \hat{p}_j$ is the group-specific sectoral price index and $s_{j|r}^i = s_j^i / s_r^i$ is the spending share within the sector. (A5) represents the change in domestic spending on domestic varieties as a sum of five terms, representing growth of domestic income, reduction of domestic prices relative to the industry composite, cross-price effects across industries and sectors, and income effects. Averaging $\hat{X}_{jH}^i$ between the two consumer groups and expressing everything in terms of averages and differences between the groups, we obtain:

$$\begin{aligned}\hat{X}_{jH}^{\text{Final}} &= \hat{w} + (\mu_j - \bar{\mu}) (\hat{w}_S - \hat{w}_U) \\ &\quad + (1 - \xi_j) (\hat{p}_{jH} - \hat{p}_j) \\ &\quad + (1 - \varepsilon_r) (\hat{p}_j - \hat{p}_r - (\mu_j - \mu_r) (\hat{p}_{Sr} - \hat{p}_{Ur})) \\ &\quad + (1 - \rho) (\hat{p}_r + (\mu_j - \mu_r) (\hat{p}_{Sr} - \hat{p}_{Ur}) - \hat{\pi} - (\mu_j - \bar{\mu}) (\hat{\pi}_S - \hat{\pi}_U)) \\ &\quad + (\psi_j - 1) (\hat{w} - \hat{\pi} + (\mu_j - \bar{\mu}) (\hat{w}_S - \hat{w}_U) - (\mu_j - \bar{\mu}) (\hat{\pi}_S - \hat{\pi}_U)),\end{aligned}\quad (\text{A6})$$

where $\hat{p}_r \equiv \mu_r \hat{p}_{Sr} + (1 - \mu_r) \hat{p}_{Ur} = \mathbb{E}_{\text{Final}} [\hat{p}_j \mid r]$ is the sectoral price index for all final consumption.

It remains to characterize various price indices. Equations (5) and (6) characterize producer and consumer price changes by industry, $\hat{p}_{jH}$ and $\hat{p}_j$, respectively, together implying:

$$\hat{p}_{jH} - \hat{p}_j = -IP_{j\mathbf{c}} \hat{\tau} + IP_j \hat{w} + IP_j (v_j - \bar{v}) (\hat{w}_S - \hat{w}_U).$$

Similarly,

$$\begin{aligned}\hat{p}_j - \hat{p}_r &= (IP_{j\mathbf{c}} - \mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}} \mid r]) \hat{\tau} - (IP_j - \mathbb{E}_{\text{Final}} [IP_j \mid r]) \hat{w} \\ &\quad + ((v_j - \bar{v}) (1 - IP_j) - \mathbb{E}_{\text{Final}} [(v_j - \bar{v}) (1 - IP_j) \mid r]) (\hat{w}_S - \hat{w}_U), \\ \hat{p}_r - \hat{\pi} &= \mathbb{E}_{\text{Final}} [\hat{p}_j \mid r] - \mathbb{E}_{\text{Final}} [\hat{p}_j], \\ \hat{p}_{Sr} - \hat{p}_{Ur} &= \Delta_{\text{Final}} [\hat{p}_j \mid r] \equiv \sum_{j \in r} \left(s_{j|r}^S - s_{j|r}^U \right) \hat{p}_j,\end{aligned}$$

and

$$\hat{\pi}_S - \hat{\pi}_U = \Delta_{\text{Final}} [\hat{p}_j],$$

where in the last three lines industry price indices can again be plugged in from (6). We substitute these expressions in (A6) and then to (A3), and rearrange terms, combining those with $\hat{\tau}$, $\hat{\tau}^*$, $\hat{w}$, and $\hat{w}_S - \hat{w}_U$. The coefficients at these terms are η_j^{import} , $-\eta_j^{\text{export}}$, $\eta_j^{\text{avg wage}}$, and $-\eta_j^{\text{skill prem}}$, respectively. With some algebra, we arrive at (12b), (12c), and

$$\begin{aligned} \eta_j^{\text{import}} &= \text{Dom share}_j \cdot \{(\xi_j - 1) \cdot IP_{j\mathbf{c}} + (\varepsilon_r - 1) \cdot (\mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}} | r] - IP_{j\mathbf{c}}) \\ &\quad + (\rho - 1) \cdot (\mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}}] - \mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}} | r]) - (\psi_j - 1) \cdot \mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}}] \\ &\quad + (\varepsilon_r - \rho) (\mu_j - \mu_r) \Delta_{\text{Final}} [IP_{j\mathbf{c}}] - (\psi_j - \rho) (\mu_j - \bar{\mu}) \Delta_{\text{Final}} [IP_{j\mathbf{c}}]\}. \end{aligned} \quad (\text{A7})$$

The two terms in the last line were ignored in (12a) as negligible: they correspond to the effects of tariff changes that operate through differential price indices; they only affect the earnings channel if consumer base is correlated with skill intensity and if spending on imports is differential across groups, neither of which happens much in our data. Finally, ignoring analogous negligible effects, the elasticity of industry size with respect to the lower skill premium, which is important for σ_{macro} , is given by

$$\begin{aligned} \eta_j^{\text{skill prem}} &= \text{Export share}_j \cdot (\xi_j - 1) (v_j - \bar{v}) + \text{Dom share}_j \cdot \{(\xi_j - 1) IP_j (v_j - \bar{v}) \\ &\quad + (\varepsilon_r - 1) ((1 - IP_j) (v_j - \bar{v}) - \mathbb{E}_{\text{Final}} [(1 - IP_j) (v_j - \bar{v}) | r]) \\ &\quad + (\rho - 1) (\mathbb{E}_{\text{Final}} [(1 - IP_j) (v_j - \bar{v}) | r] - \mathbb{E}_{\text{Final}} [(1 - IP_j) (v_j - \bar{v})]) \\ &\quad + (\psi_j - 1) \mathbb{E}_{\text{Final}} [(1 - IP_j) (v_j - \bar{v})] - \psi_j (\mu_j - \bar{\mu})\}. \end{aligned} \quad (\text{A8})$$

Discussion of σ_{macro} . Recall that $\sigma_{\text{macro}} = \sigma_{\text{within}} + \Delta_{VA} [\eta_j^{\text{skill prem}}]$. The second term can be expressed as a cross-industry average, $\Delta_{VA} [\eta_j^{\text{skill prem}}] = \mathbb{E}_{VA} \left[\eta_j^{\text{skill prem}} \cdot \frac{v_j - \bar{v}}{\bar{v}(1 - \bar{v})} \right]$ (see analogous equation (A26) in Appendix A.4). Plugging in (A8), we obtain

$$\begin{aligned} \sigma_{\text{macro}} &= (1 - \text{Segm}_{\text{prod}}) \bar{\sigma} + \omega_{\xi} \bar{\xi} + \omega_{\varepsilon} \bar{\varepsilon} + \omega_{\rho} \bar{\rho} + \omega_{\psi} \bar{\psi} \\ &\quad - \Delta_{VA} [\text{Dom share}_j \cdot \psi_j (\mu_j - \bar{\mu})]. \end{aligned}$$

Here $\omega_{\xi} + \omega_{\varepsilon} + \omega_{\rho} + \omega_{\psi} = \text{Segm}_{\text{prod}}$, so the first line is a weighted average of elasticities in the model: σ_j , ξ_j , ε_r , ρ , and ψ_j , similar to Oberfield and Raval (2014).⁹⁰ The weights correspond to the importance of different types of labor reallocation: $1 - \eta_{\text{prod}}$ is large when skill-intensities are relatively homogeneous

⁹⁰For instance, $\omega_{\xi} = \mathbb{E}_{VA} \left[(1 - \text{Dom share}_j \cdot (1 - IP_j)) \frac{(v_j - \bar{v})^2}{\bar{v}(1 - \bar{v})} \right] = \Delta_{VA} [(1 - \text{Dom share}_j \cdot (1 - IP_j)) (v_j - \bar{v})]$, and other weights are similarly obtained from $\Delta_{VA} [\eta_j^{\text{skill prem}}]$.

across industries, which creates room for within-industry reallocation. The other terms capture different types of between-industry reallocation. For example, the weight on the trade elasticities (ω_ξ) is larger when the economy is more open on both import and export sides. The second line provides a segregation adjustment.

A.2 General Model

In this section we derive the average gains from trade policy, as well as the distributional effects through the expenditure and earnings channels in the general model, which allows for input-output linkages.

To characterize intermediate demand, we make a functional form assumption on the production function: it combines intermediate inputs from various industries with the composite output of labor (value added) in the Cobb-Douglas way,

$$Q_{jH} = \left(F_j^{\text{VA}} \left(L_S^j, L_U^j \right) \right)^{1-\beta_j} \cdot \prod_l \left(Q_l^j \right)^{\beta_l^j},$$

where F_j^{VA} is a homogeneous of degree one function with local elasticity of substitution σ_j , $\beta_j = \sum_l \beta_l^j$ is the total cost share of intermediates, and matrix $\mathbf{B} = \left(\beta_l^j \right)$ is the input requirement (I-O) matrix. The rows of $\mathbf{B}$ correspond to selling (upstream) industries l , columns to buying (downstream) industries j , and elements measure the fraction of domestic industry j 's costs spent on inputs from industry l . Q_l^j is the quantity of the industry- l composite commodity purchased by j , and it combines varieties from all countries in the same CES way as final consumers aggregate them. That is the standard proportionality condition that the BEA input-output table relies on.⁹¹ The Cobb-Douglas assumption is only required for the earnings channel—as before, the expenditure channel results remain non-parametric.

As in Section 2, we proceed in two steps. We first characterize price changes conditionally on wage changes and solve for the average gains and the expenditure channel. Then we solve for the industry size change and link it to skill premium growth.

Step 1. Price Changes. We show that industry-specific consumer (producer) price indices are determined by the *total* share of spending on imports (on imported intermediates), both direct and indirect. We use tildes to denote objects that account for upstream suppliers, and define the total share of imports from country c in consumption ($\widetilde{IP}_{jc}$) and in domestic production ($\widetilde{IP}_{jc}^{\text{Interm}}$) in an intuitive recursive way:

$$\widetilde{IP}_{jc} = IP_{jc} + (1 - IP_j) \cdot \widetilde{IP}_{jc}^{\text{Interm}}, \quad (\text{A9a})$$

$$\widetilde{IP}_{jc}^{\text{Interm}} = \sum_l \beta_l^j \widetilde{IP}_{jc}, \quad (\text{A9b})$$

⁹¹The proportionality condition is also standard in the literature (e.g. [Caliendo and Parro, 2015](#)). The World Input-Output Database departs from it slightly by allowing for differences in import shares between final and intermediate consumers (but not across different types of each). However, it is more aggregated and has other limitations ([Timmer et al., 2015](#)).

with analogous notation for imports from a set of countries $\mathbf{c}$ or from all foreign countries.⁹² In matrix form, (A9a) solves as $\widetilde{IP}_c = \widetilde{\mathbf{B}} \cdot IP_c$ where $\widetilde{\mathbf{B}} = (\mathbf{I} - \text{diag}(1 - IP_j) \mathbf{B}')^{-1}$ is a Leontief-type inverse matrix.⁹³

By Shephard's lemma, producer price index satisfies

$$\hat{p}_{jH} = (1 - \beta_j) (\hat{w} + (v_j - \bar{v}) (\hat{w}_S - \hat{w}_U)) + \sum_l \beta_l^j \hat{p}_j.$$

And consumer price index combines domestic and foreign price changes:

$$\hat{p}_j = IP_{jc} \hat{\tau} + (1 - IP_j) \hat{p}_{jH}.$$

Combining these expression, we obtain a recursive characterization for the consumer price index that parallels (A9a)–(A9b):

$$\begin{aligned} \hat{p}_j &= IP_{jc} \hat{\tau} + (1 - IP_j) \left((1 - \beta_j) (\hat{w} + (v_j - \bar{v}) (\hat{w}_S - \hat{w}_U)) + \sum_l \beta_l^j \hat{p}_j \right) \\ &= \widetilde{IP}_{jc} \hat{\tau} + (1 - \widetilde{IP}_j) (\hat{w} + (\tilde{v} - \bar{v}) (\hat{w}_S - \hat{w}_U)), \end{aligned} \quad (\text{A10})$$

where $\tilde{v}_j$ is the average skill intensity of the domestic part of the supply chain leading to good j , which is defined in matrix form by $\left\{ (1 - \widetilde{IP}_j) \tilde{v}_j \right\} = \widetilde{\mathbf{B}} \cdot \left\{ (1 - \beta^j) (1 - IP_j) v_j \right\}$. Similarly for producer price index,

$$\hat{p}_{jH} = \widetilde{IP}_{jc}^{\text{Interm}} \hat{\tau} + (1 - \widetilde{IP}_j^{\text{Interm}}) (\hat{w} + (\tilde{v} - \bar{v}) (\hat{w}_S - \hat{w}_U)). \quad (\text{A11})$$

Expressions (A10) and (A11) generalize (6) and (5), respectively. They imply the expressions for the average gains and the expenditure channel that are analogous to (7) and (8), with tildes added.

Step 2. Wage Changes. The connection between the wage changes and differential growth of industries in Step 2a holds in the general model.⁹⁴ To solve for the average wage and skill premium changes, it remains to generalize Step 2b, i.e. equations (12a)–(12c) by characterizing the changes in industry value added. We start from the product market clearing condition; domestic output can be sold to domestic final and intermediate consumers, as well as to exports: $X_{jH} = X_{jH}^{\text{Final}} + X_{jH}^{\text{Interm}} + X_{jH}^{\text{Export}}$, where $X_{jH}^{\text{Interm}} = \sum_k X_{jH}^k$ measures total intermediate sales and k indexes domestic downstream industries, each buying X_{jH}^k from j .

To measure the contribution of different modes of selling to the total, we need to know the weights of

⁹²Tintelnot et al. (2017) use analogous definitions looking at the firm-to-firm input-output network. We have that in mind when computing the import share from the microdata in Sections 4 and 5, accounting for both direct and first-order indirect imports.

⁹³Multiplication of $\mathbf{B}'$ by $\text{diag}(1 - IP_j)$ is the open-economy adjustment to the I-O table described by Antràs et al. (2012).

⁹⁴This requires that the production function can be written as $Q_{jH} = F_j(F_j^{\text{VA}}(L_S^j, L_U^j), Q_1^j, \dots, Q_J^j)$ for some F_j homogenous of degree one, so that substitution patterns between labor types are independent of the choice of intermediate inputs and σ_j is well-defined. The Cobb-Douglas assumption is sufficient but not necessary.

each term in the total before the shock and to predict their changes after the shock. Regarding the former, in the I-O table we observe the share of exports *in output* and the share of final sales and sales to each downstream industry k *in absorption*; purchases of domestic varieties by different buyers are not reported directly. However, they can be computed using the proportionality condition. Specifically, we introduce the intermediate absorption matrix $\mathbf{D} = (\delta_j^k)$. Its rows (columns) correspond to the selling (buying) industries j (k), and typical element $\delta_j^k = X_{jH}^k / \text{Absorption}_j$ measures the share of industry j 's absorption (spending on both domestic and foreign varieties) that is used as intermediate inputs to downstream industry k . While $\mathbf{B}$ looks at industry's suppliers, $\mathbf{D}$ characterizes its buyers. By proportionality, shares δ_j^k equally apply to the domestic sales of *domestic* varieties, $X_{jH}^k / (X_{jH}^{\text{Final}} + X_{jH}^{\text{Interm}}) = \delta_j^k$, hence the share of domestic output that goes to k equals $X_{jH}^k / X_{jH} = \text{Dom share}_j \cdot \delta_j^k$. Similarly, the share of domestic output that is sold to domestic final consumers is $\text{Dom share}_j \cdot (1 - \delta_j) \equiv \text{Dom final share}_j$ with $\delta_j = \sum_k \delta_j^k$ measuring the share of final sales in absorption. As a result,

$$\hat{X}_{jH} = \text{Export share}_j \cdot \hat{X}_{jH}^{\text{Export}} + \text{Dom final share}_j \cdot \hat{X}_{jH}^{\text{Final}} + \text{Dom share}_j \cdot \sum_k \delta_j^k \hat{X}_{jH}^k. \quad (\text{A12})$$

The change of output on the left-hand side equals $\widehat{VA}_j$ due to the Cobb-Douglas assumption. To characterize the change of intermediate sales $\hat{X}_{jH}^k$, we apply Cobb-Douglas again. The share of spending by industry k on *all* varieties of j is fixed, so $\hat{X}_{jH}^k = \hat{X}_k$, and CES aggregation across varieties implies that $\hat{X}_{jH}^k = \hat{X}_k + (1 - \xi_j)(\hat{p}_{jH} - \hat{p}_j)$, where the last term captures import competition in intermediate sales. Plugging these into (A11) yields a recursive characterization for $\widehat{VA}_j$:

$$\widehat{VA}_j = \text{Export share}_j \cdot \hat{X}_{jH}^{\text{Export}} + \text{Dom final share}_j \cdot \hat{X}_{jH}^{\text{Final}} + \text{Dom share}_j \cdot \left(\delta_j (1 - \xi_j) (\hat{p}_{jH} - \hat{p}_j) + \sum_k \delta_j^k \widehat{VA}_k \right).$$

Denoting $\text{Interm share}_j = \text{Dom share}_j \cdot \delta_j$ and solving it in matrix form, we obtain

$$\begin{aligned} \widehat{VA} = \tilde{\mathbf{D}} \cdot \left\{ \text{Export share}_j \cdot \hat{X}_{jH}^{\text{Export}} + \text{Dom final share}_j \cdot \hat{X}_{jH}^{\text{Final}} \right. \\ \left. + \text{Interm share}_j \cdot (1 - \xi_j) (\hat{p}_{jH} - \hat{p}_j) \right\}, \end{aligned} \quad (\text{A13})$$

where $\tilde{\mathbf{D}} = (\mathbf{I} - \text{diag}(\text{Dom share}_j) \mathbf{D})^{-1}$ is the Leontief inverse corresponding to $\mathbf{D}$. Three terms in (A13) correspond to direct changes in export demand, domestic final demand, and competition with foreign varieties in intermediate markets. Pre-multiplication by $\tilde{\mathbf{D}}$ makes the *I-O adjustment* to account for the propagation of shocks from downstream industries up through changes in intermediate demand; algebraically, (A13) is the sum of direct value added changes in the industry itself, its intermediate customers, their customers, etc. For example, elements of $(\tilde{\mathbf{D}} \cdot \text{Export share})$ are shares of domestic output that is exported either directly or indirectly by selling to domestic downstream industries that export. Similarly, elements of $(\tilde{\mathbf{D}} \cdot \text{Dom final share})$ are complementary shares of output ultimately sold to domestic consumers.

Expressions (A4) and (A5) for export and domestic final demand follow from the corresponding demand systems and extend to the general case with no change. We plug in consumer and producer price indices from (A10) and (A11) and combine terms at $\hat{\tau}$, $\hat{\tau}^*$, $\hat{w}$, and $\hat{w}_S - \hat{w}_U$ to arrive at (11), where industry size responses are characterized by the expressions we present in order.⁹⁵ The elasticity with respect to import tariffs (generalization of (12a)) combines import competition effects at the three tiers (varieties, industries, and sectors), income effects, and cheaper intermediate input effects.⁹⁶

$$\eta_j^{\text{import}} = \text{Import comp effect}_j + \text{Income effect}_j - \text{Int.Input effect}_j. \quad (\text{A14a})$$

Here import competition effects are like (12a) before, but I-O adjusted:

$$\begin{aligned} \text{Import comp effect} = & \tilde{\mathbf{D}} \cdot \{[(\xi_j - 1) IP_{j\mathbf{c}} + (\varepsilon_r - 1) (\mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}} | r] - IP_{j\mathbf{c}}) \\ & + (\rho - 1) (\mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}}] - \mathbb{E}_{\text{Final}} [IP_{j\mathbf{c}} | r])] \cdot \text{Dom final share}_j \\ & + (\xi_j - 1) IP_{j\mathbf{c}} \cdot \text{Interm sales share}_j\}. \end{aligned} \quad (\text{A14b})$$

Income effects, which did not exist in the simplified model, are determined by the weighted average income elasticity $\bar{\psi}_j$ that combines income elasticity of the varieties produced by the industry and its downstream clients (“I-O-adjusted income elasticity”):

$$\begin{aligned} \text{Income effect}_j = & \tilde{\mathbf{D}} \cdot \{-(\psi_j - 1) \cdot \text{Dom final share}_j\} \cdot \mathbb{E}_{\text{Final}} [\tilde{IP}_{j\mathbf{c}}] \\ \equiv & -(\bar{\psi}_j - 1) \cdot (\tilde{\mathbf{D}} \cdot \text{Dom final share}_j) \cdot \mathbb{E}_{\text{Final}} [\tilde{IP}_{j\mathbf{c}}]. \end{aligned} \quad (\text{A14c})$$

Finally, due to cheaper intermediate inputs, domestic varieties outcompete foreign varieties as they become cheaper by $\tilde{IP}_{j\mathbf{c}}^{\text{Interm}}$, which is only partially offset by the lower industry price index ($\tilde{IP}_{j\mathbf{c}}^{\text{Interm},H} \equiv \tilde{IP}_{j\mathbf{c}}^{\text{Interm}} \cdot (1 - IP_j)$). Moreover, lower price index attracts demand from other industries to j . These effects are formalized by:

$$\begin{aligned} \text{Int.Input effect}_j = & \tilde{\mathbf{D}} \cdot \left\{ \left[(\xi_j - 1) \left(\tilde{IP}_{j\mathbf{c}}^{\text{Interm}} - \tilde{IP}_{j\mathbf{c}}^{\text{Interm},H} \right) + (\varepsilon_r - 1) \left(\tilde{IP}_{j\mathbf{c}}^{\text{Interm},H} - \mathbb{E}_{\text{Final}} [\tilde{IP}_{j\mathbf{c}}^{\text{Interm},H} | r] \right) \right] \right. \\ & + (\rho - 1) \left(\mathbb{E}_{\text{Final}} [\tilde{IP}_{j\mathbf{c}}^{\text{Interm},H} | r] - \mathbb{E}_{\text{Final}} [\tilde{IP}_{j\mathbf{c}}^{\text{Interm},H}] \right) \cdot \text{Dom final share}_j \\ & + (\xi_j - 1) \left(\tilde{IP}_{j\mathbf{c}}^{\text{Interm}} - \tilde{IP}_{j\mathbf{c}}^{\text{Interm},H} \right) \cdot \text{Interm sales share}_j \\ & \left. + (\xi_j - 1) \tilde{IP}_{j\mathbf{c}}^{\text{Interm}} \cdot \text{Export share}_j \right\}. \end{aligned} \quad (\text{A14d})$$

The responses of industry size to the lower export tariff and to higher domestic average wage are similar to (12b)–(12c) but I-O adjusted. In particular, the export effect includes exports in downstream

⁹⁵Details of the algebra are available from the authors upon request.

⁹⁶This ignores negligible terms analogous to those discussed in Appendix A.1.

industries:

$$\eta_j^{\text{export}} = \tilde{\mathbf{D}} \cdot \{(\xi_j - 1) \text{ Export share}_j\}, \quad (\text{A15})$$

$$\eta^{\text{avg wage}} = \tilde{\mathbf{D}} \cdot \text{Dom final share} - \eta^{\text{import}} - \eta^{\text{export}}. \quad (\text{A16})$$

The generalization of (A8) for $\eta_j^{\text{skill prem}}$ is obtained analogously. This concludes Step 2b. Step 2c is unchanged, and equations (13a)–(14) characterize the changes in the average wage and the skill premium.

A.3 Non-Homothetic Nested CES

In this section we show that non-homothetic nested CES preferences defined implicitly by (1) yield a very intuitive log-linear approximation for the *change* in demand after a set of wage and price shocks, even though demand functions do not have a closed-form solution.

We will fix one type of agents and therefore suppress the i index, and to understand the demand behavior in the general case, we treat the set of shocks to consumer's expenditure and prices, $\hat{X}$ and $\hat{p}_j$, as exogenous. We also assume the parameter restrictions which guarantee that preferences are well-defined. In particular this requires $\varepsilon_r \neq 1$, since non-homothetic Cobb-Douglas preferences cannot be globally defined (Hanoch, 1975; Comin et al., 2016).⁹⁷

Conditional on the utility level $\mathcal{U}$ at the optimal bundle of goods, preferences reduce to nested CES, which has a well-known demand structure. Define the sectoral and overall price indices as

$$\begin{aligned} p_r^* &= \left(\sum_{j \in r} a_j \mathcal{U}^{\varphi_j - 1} p_j^{1 - \varepsilon_r} \right)^{1/(1 - \varepsilon_r)}, \\ \pi^* &= \left(\sum_r p_r^{*1 - \rho} \right)^{1/(1 - \rho)}. \end{aligned} \quad (\text{A17})$$

Then

$$\mathcal{U} = X/\pi^*, \quad (\text{A18})$$

and spending on good j satisfies

$$X_j = X \cdot s_j \equiv X \cdot \underbrace{\frac{a_j \mathcal{U}^{\varphi_j - 1} p_j^{1 - \varepsilon_r}}{p_r^{*1 - \varepsilon_r}}}_{s_{j|r}} \cdot \underbrace{\frac{p_r^{*1 - \rho}}{\pi^{*1 - \rho}}}_{s_r}. \quad (\text{A19})$$

Define $\lambda_j = (\varphi_j - 1) / (1 - \varepsilon_r)$, $\lambda_r = \sum_{j \in r} s_{j|r} \lambda_j$, and $\lambda = \sum_r s_r \lambda_r$, which are observable at the

⁹⁷In the application we will allow for $\varepsilon_r = 1$, interpreted as $\varepsilon_r \rightarrow 1$, which is sufficient for us since we are only interested in the local behavior of demand.

original equilibrium, given preference parameters. Then log-differentiating (A17) yields:

$$\begin{aligned}\hat{p}_r^* &= \sum_j s_{j|r} \left(\hat{p}_j + \lambda_j \hat{\mathcal{U}}^* \right) \quad \text{and} \\ \hat{\pi}^* &= \sum_r s_r \hat{p}_r^* = \hat{\pi} + \lambda \hat{\mathcal{U}}^*,\end{aligned}$$

where $\hat{\mathcal{U}}^* = d \log \mathcal{U}$ is the relative change in the *cardinal utility*. Together with (A18) this implies

$$\hat{\mathcal{U}}^* = \hat{X} - \hat{\pi}^* = \frac{\hat{X} - \hat{\pi}}{1 + \lambda}. \quad (\text{A20})$$

This equation relates changes in the cardinal utility to observable objects only: money metric (change in the total expenditure minus the Laspeyres price index) and the spending shares at the original equilibrium (which enter λ). We can now solve for $\hat{p}_r^*$, $\hat{\pi}^*$, and ultimately for the change in demand, also in terms of observables:

$$\begin{aligned}\hat{p}_r^* &= \hat{p}_r + \frac{\lambda_r}{1 + \lambda} \left(\hat{X} - \hat{\pi} \right), \\ \hat{\pi}^* &= \hat{\pi} + \frac{\lambda}{1 + \lambda} \left(\hat{X} - \hat{\pi} \right), \quad \text{and} \\ \hat{X}_j &= \hat{X} + (\varphi_j - 1) \hat{\mathcal{U}}^* + (1 - \varepsilon_r) (\hat{p}_j - \hat{p}_r^*) + (1 - \rho) (\hat{p}_r^* - \hat{\pi}^*) \\ &= \hat{X} + (1 - \varepsilon_r) (\hat{p}_j - \hat{p}_r) + (1 - \rho) (\hat{p}_r - \hat{\pi}) + (\psi_j - 1) \left(\hat{X} - \hat{\pi} \right),\end{aligned} \quad (\text{A21})$$

where

$$\psi_j = 1 + \frac{(1 - \varepsilon_r) (\lambda_j - \lambda_r) - (1 - \rho) (\lambda_r - \lambda)}{1 + \lambda}. \quad (\text{A22})$$

According to (A21), the change in spending on industry j has four components. The first three are identical to homothetic nested CES, capturing the change in total expenditure (i.e., in income), reallocation of demand within the sector and across sectors. The fourth is the income effect, shaped by the income elasticity ψ_j . When the money metric of utility, which is an observable measure of real income, goes up, spending on income-elastic products with $\psi_j > 1$ increases. Within each sector, income elasticities are higher in industries with higher φ_j , but the comparison across sectors is less straightforward.

A.4 Two Decompositions

To empirically investigate differences in spending on imports ($\Delta_{\text{Final}} [IP_{j\mathbf{c}}]$) and differences in exposure to the labor market effects of trade through various channels ($\Delta_{\text{VA}} [\cdot]$), we use two decompositions.

First, we note that there is a convenient way of visualizing the main patterns in the data. It is intuitive that the share of spending on imports is higher for the skilled than unskilled *consumers* (in the agent space) if and only if *industries* that sell relatively more to the skilled group have higher import shares (in

the product space). We formalize this idea by defining μ_j as the fraction of industry j 's domestic final sales that goes to the skilled group. We will refer to μ_j as the industry's consumer base, and its parallel in the labor market is the skill intensity v_j . Then the difference in spending on imports can be represented as the slope of the regression of import penetration on the consumer base, rescaled by the *consumption segmentation index*:

$$\Delta_{\text{Final}}[IP_{\mathbf{c}}] = \frac{\text{Cov}[\mu, IP_{\mathbf{c}}]}{\text{Var}[\mu]} \cdot \text{Segm}_{\text{cons}}, \quad (\text{A23})$$

where $\text{Segm}_{\text{cons}} = \frac{\text{Var}[\mu]}{\bar{\mu}(1-\bar{\mu})} = \Delta_{\text{Final}}[\mu]$ measures the difference between consumption baskets of the two types in a model-consistent way.⁹⁸ Segmentation equals zero when all industries have the same mix of final consumers and attains the maximum value of one when each industry sells only to one group. Decomposition (A23) shows that skilled consumers spend more on imports when segmentation is sufficiently high and import shares are higher in industries with skilled consumers. The regression slope can be visualized using scatterplots and other standard tools.

Second, differential effects can be represented as a sum of the components arising “between” and “within” more aggregated groups of products (for instance, sectors). Consider the expenditure side:

$$\Delta_{\text{Final}}[IP_{j\mathbf{c}}] = \sum_j (s_j^S - s_j^U) IP_{j\mathbf{c}} = \Delta_{\text{Final}}^{\text{between}}[IP_{\mathbf{c}}] + \Delta_{\text{Final}}^{\text{within}}[IP_{\mathbf{c}}], \quad (\text{A24})$$

where $\Delta_{\text{Final}}^{\text{between}}[IP_{\mathbf{c}}] = \sum_g (s_g^S - s_g^U) IP_{g\mathbf{c}}$ and $\Delta_{\text{Final}}^{\text{within}}[IP_{\mathbf{c}}] = \sum_j (s_j^S - s_j^U) (IP_{j\mathbf{c}} - IP_{g\mathbf{c}})$. In these expressions g indexes product groups, s_g^i is the share of spending of type i on all products within group g , and $IP_{g\mathbf{c}}$ is the average import share for all products in group g , with total final expenditures weights. The “between” component ignores compositional differences within product groups, while the “within” component only captures those.

Analogous decompositions hold on the earnings side, with three differences: industries are weighted by value added instead of final consumption, spending shares are replaced with payroll shares e_i^j , and consumer base μ_j is replaced by the skill intensity v_j . The counterpart to $\text{Segm}_{\text{cons}}$ is the production segmentation index $\text{Segm}_{\text{prod}} = \text{Var}[v] / \bar{v}(1 - \bar{v})$, which measures the heterogeneity of industries by skill intensity. Oberfield and Raval (2014) consider an equivalent index based on capital intensity, which they call the heterogeneity index.

⁹⁸The final consumption-weighted average $\bar{\mu}$ represents the fraction of the skilled population in total expenditures. By assumption, expenditures are proportionate to income for each group, which implies that the overall expenditure and income shares of the skilled group are equal, i.e. $\bar{\mu} = \bar{v}$. Consumption segmentation index is related to, but conceptually distinct from, a commonly used dissimilarity index Duncan and Duncan (1955). Appendix A.5 explains the relationship between the two measures.

Proofs. To establish decomposition (A23), note that the share of an industry in the college graduates' spending can be represented in the following way:

$$\begin{aligned}
s_j^S &= \frac{X_j^S}{X_S} \\
&= \frac{X_j^S + X_j^U}{X_S + X_U} \cdot \frac{X_j^S / (X_j^S + X_j^U)}{X_S / (X_S + X_U)} \\
&= s_j^{\text{Final}} \cdot \frac{\mu_j}{\bar{\mu}}.
\end{aligned} \tag{A25a}$$

This expression implies, for instance, that industry j is overrepresented in consumption of the skilled group ($s_j^S > s_j^{\text{Final}}$) if and only if $\mu_j > \bar{\mu}$. Similarly for the unskilled group,

$$s_j^U = s_j^{\text{Final}} \cdot \frac{1 - \mu_j}{1 - \bar{\mu}}. \tag{A25b}$$

Plugging (A25a) and (A25b) into the differential spending formula yields a representation of differential spending as a rescaled covariance between import penetration and the consumer base across industries:

$$\begin{aligned}
\Delta_{\text{Final}}[IP_{\mathbf{c}}] &= \sum_j (s_j^S - s_j^U) IP_{j\mathbf{c}} \\
&= \sum_j s_j^{\text{Final}} \left(\frac{\mu_j}{\bar{\mu}} - \frac{1 - \mu_j}{1 - \bar{\mu}} \right) IP_{j\mathbf{c}} \\
&= \mathbb{E}_{\text{Final}} \left[\frac{\mu_j - \bar{\mu}}{\bar{\mu}(1 - \bar{\mu})} IP_{j\mathbf{c}} \right] \\
&= \frac{\text{Cov}[\mu_j, IP_{j\mathbf{c}}]}{\bar{\mu}(1 - \bar{\mu})},
\end{aligned} \tag{A26}$$

where covariance is weighted by total final consumption. The last step to (A23) follows from the canonical formula for the least square regression coefficient, $\text{Cov}[\mu_j, IP_{j\mathbf{c}}] / \text{Var}[\mu_j]$.

The within-between decomposition is straightforward. By definition of the spending share of the group of products, $s_g^i = \sum_{j \in g} s_j^i$. This implies

$$\Delta_{\text{Final}}^{\text{between}}[IP_{\mathbf{c}}] = \sum_j (s_j^S - s_j^U) IP_{g\mathbf{c}}.$$

Adding it up with $\Delta_{\text{Final}}^{\text{within}}[IP_{\mathbf{c}}]$, one immediately gets $\Delta_{\text{Final}}[IP_{\mathbf{c}}]$, as required.

A.5 Consumption Segmentation and Dissimilarity Index

The dissimilarity index (Duncan and Duncan, 1955) is defined as the minimal share of consumption of one group that needs to be reallocated to different products to reach the consumption basket of the other

group, which can be formalized as

$$\text{Dissim} = \frac{1}{2} \sum_j |s_j^S - s_j^U|. \quad (\text{A27})$$

Here we show that it admits a representation similar to the consumption segmentation index, but with the variance replaced by mean absolute deviation, which is the first moment rather than the second. Plugging in (A25a) and (A25b) into (A27) yields:

$$\begin{aligned} \text{Dissim} &= \frac{1}{2} \sum_j \left| s_j^{\text{Final}} \cdot \left(\frac{\mu_j}{\bar{\mu}} - \frac{1 - \mu_j}{1 - \bar{\mu}} \right) \right| \\ &= \frac{1}{2} \frac{\mathbb{E}_{\text{Final}} [|\mu_j - \bar{\mu}|]}{\bar{\mu} (1 - \bar{\mu})}, \end{aligned}$$

where the numerator is indeed the (weighted) mean absolute deviation.

Because $\mu_j \in [0, 1]$, consumption segmentation and dissimilarity coincide in extreme cases: they both are zero if and only if consumer base is the same in all industries, and they are both one if any only if consumption baskets do not overlap, i.e. $\mu_j \in \{0, 1\}$. More interestingly, they also coincide in the case when industries are of two types: some sell to only one type of consumers ($\mu_j \in \{0, 1\}$), while the others sell to everyone in proportion to their income ($\mu_j = \bar{\mu}$). Then it is straightforward to verify that both segmentation and dissimilarity equal to the share of total spending on the non-overlapping industries.

B Econometric Appendix

B.1 Inference for the Expenditure Channel Decomposition

In this section we describe how standard errors can be constructed for the differential shares of import spending from the firm sample in Section 4.2. Suppose there is a set of firms indexed by f . Each of them is characterized by the outcome y_f (the import share) and non-negative spending levels in dollars by the skilled and unskilled people, denoted S_f and U_f , which are positive in expectation.⁹⁹ Our population object of interest is the differential expectation of the outcome in consumption baskets of the skilled and unskilled:¹⁰⁰

$$\theta = \mathbb{E} \left[\frac{S_f}{\mathbb{E}[S_f]} y_f - \frac{U_f}{\mathbb{E}[U_f]} y_f \right]. \quad (\text{A28})$$

We observe an i.i.d. sample of N firms characterized by (y_f, S_f, U_f) . This captures three types of randomness in the data: in the outcome variable, consumer base of the firm, and firm size, which we denote by $X_f = S_f + U_f$. The plug-in estimator for θ is

$$\hat{\theta} = \frac{\sum_f S_f y_f}{\sum_f S_f} - \frac{\sum_f U_f y_f}{\sum_f U_f}. \quad (\text{A29})$$

⁹⁹We will assume that all regularity conditions, such as finiteness of second moments, are satisfied.

¹⁰⁰To define the estimand with our square-root weighting scheme, we divide S_f and U_f by $\sqrt{S_f + U_f}$, so that they add up to the square root of the actual firm sales. This weighting scheme helps make regularity conditions hold in the data despite the skewed distributions of S_f and U_f .

We explain how to do inference for $\hat{\theta}$ as well as for the between- and within-group components of it.

Inference on the Weighted Mean. The estimator in (A29) is a difference of two objects. We first show how to do inference on each of them separately, for instance the first one:

$$\hat{\theta}_S = \frac{\sum_f S_f y_f}{\sum_f S_f},$$

which estimates $\theta_S = \mathbb{E}[S_f y_f] / \mathbb{E}[S_f]$. The asymptotic behavior of $\hat{\theta}_S$ follows from Central Limit Theorem and the Slutsky's theorem:

$$\begin{aligned} \sqrt{N} (\hat{\theta}_S - \theta_S) &= \sqrt{N} \left(\frac{\frac{1}{N} \sum_f S_f y_f}{\frac{1}{N} \sum_f S_f} - \theta_S \right) \\ &= \frac{\frac{1}{\sqrt{N}} \sum_f S_f (y_f - \theta_S)}{\frac{1}{N} \sum_f S_f} \\ &\rightarrow^p \mathcal{N} \left(0, \frac{\text{Var}[S_f (y_f - \theta_S)]}{(\mathbb{E}[S_f])^2} \right). \end{aligned}$$

The asymptotic variance of $\hat{\theta}_S$ can then be consistently estimated by

$$\begin{aligned} \widehat{\text{Var}}(\hat{\theta}_S) &= \frac{1}{N} \frac{\sum_f \left(S_f (y_f - \hat{\theta}_S) \right)^2}{\left(\frac{1}{N} \sum_f S_f \right)^2} \\ &= \frac{\sum_f \left(S_f (y_f - \hat{\theta}_S) \right)^2}{\left(\sum_f S_f \right)^2}. \end{aligned}$$

An analogous expression holds for the estimator $\hat{\theta}_U$ of the weighted mean for the unskilled θ_U .

Inference on Differential Means. Now come back to the estimator (A29), which satisfies

$$\hat{\theta} - \theta = \frac{\frac{1}{N} \sum_f S_f (y_f - \theta_S)}{\frac{1}{N} \sum_f S_f} - \frac{\frac{1}{N} \sum_f U_f (y_f - \theta_U)}{\frac{1}{N} \sum_f U_f}.$$

Because both denominators converge in probability to a non-zero number and numerators have zero expectations, noise in the estimation of the denominator does not increase the variance. Hence we can write:

$$\begin{aligned} \sqrt{N} (\hat{\theta} - \theta) &= \frac{1}{\sqrt{N}} \sum_f \left\{ \frac{S_f}{\mathbb{E}[S_f]} (y_f - \theta_S) - \frac{U_f}{\mathbb{E}[U_f]} (y_f - \theta_U) \right\} \\ &\rightarrow^p \mathcal{N} \left(0, \text{Var} \left[\frac{S_f}{\mathbb{E}[S_f]} (y_f - \theta_S) - \frac{U_f}{\mathbb{E}[U_f]} (y_f - \theta_U) \right] \right). \end{aligned}$$

The variance of $\hat{\theta}$ can be consistently estimated as

$$\begin{aligned}\widehat{Var}(\hat{\theta}) &= \frac{1}{N^2} \sum_f \left\{ \frac{S_f}{\frac{1}{N} \sum_f S_f} (y_f - \hat{\theta}_S) - \frac{U_f}{\frac{1}{N} \sum_f U_f} (y_f - \hat{\theta}_U) \right\}^2 \\ &= \sum_f \left\{ \frac{S_f}{\sum_f S_f} (y_f - \hat{\theta}_S) - \frac{U_f}{\sum_f U_f} (y_f - \hat{\theta}_U) \right\}^2.\end{aligned}\tag{A30}$$

This formula can be extended to the case of clustering in the standard way.

Regression Representation. While formula (A30) provides the analytical expression for $\widehat{Var}(\hat{\theta})$, a slightly different (in finite samples) but also consistent variance estimator can be obtained using the standard regression toolkit. Note that $\hat{\theta}_S$ is the slope of a simple regression of $y_f \sqrt{S_f}$ on $\sqrt{S_f}$ without a constant, and $\hat{\theta}_U$ can be obtained in an analogous way. Estimating these two regressions simultaneously, e.g. using the `suest` command in Stata, we can get robust or clustered standard errors for $\hat{\theta}_S - \hat{\theta}_U = \hat{\theta}$ without implementing (A30) manually.

Within and Between. Now suppose that firms are classified into groups g , such as industries. We will think of the sample as a sample of groups within which we observe all firms. We also assume that there is a large sample of groups. The objects of interest are the between- and within-group components of $\hat{\theta}$:

$$\begin{aligned}\hat{\theta}^{\text{between}} &= \frac{\sum_g S_g \bar{y}_g}{\sum_g S_g} - \frac{\sum_g U_g \bar{y}_g}{\sum_g U_g}, \\ \hat{\theta}^{\text{within}} &= \frac{\sum_f S_f (y_f - \bar{y}_g)}{\sum_f S_f} - \frac{\sum_f U_f (y_f - \bar{y}_g)}{\sum_f U_f},\end{aligned}$$

where $S_g = \sum_{f \in g} S_f$ and $U_g = \sum_{f \in g} U_f$ are the group sizes and $\bar{y}_g = \sum_{f \in g} (S_f + U_f) y_f / (S_g + U_g)$ is the group-level average outcome. One can easily verify that $\hat{\theta} = \hat{\theta}^{\text{between}} + \hat{\theta}^{\text{within}}$.

The between part is just a version of our estimator (A29) defined at the group level. With a large sample of groups, we can directly apply the variance estimator (A30). The within-part is the same as $\hat{\theta}$ with the outcome variable $y_f - \bar{y}_g$. This outcome variable is correlated within the group. Still, a clustered version of (A29) is consistent for the variance of $\hat{\theta}^{\text{within}}$.

In practice we use the regression representation while clustering at the firm-level for the overall effect and clustering at the group level for the within and between components.

B.2 Attenuation Bias from Product Aggregation

In this section we show that if the relationship between import shares and consumer base is weaker within firms than between, and goes in the same direction, then the differential import spending measured at the more aggregated firm level is attenuated by a factor which is below the ratio of the consumption segmentation indices across barcodes and across firms.

Consider the covariance representation of the differential import spending at the level of barcodes j (equation (A26)):

$$\Delta_{Final} [IP] = \frac{\text{Cov} [IP_j, \mu_j]}{\bar{\mu} (1 - \bar{\mu})}.$$

The denominator does not suffer from any aggregation bias. This covariance can be decomposed using the law of total covariances into the components within and between firms, denoted f , implying:

$$\frac{\text{Cov} [IP_j, \mu_j]}{\bar{\mu} (1 - \bar{\mu})} = \frac{\text{Cov} [IP_f, \mu_f]}{\bar{\mu} (1 - \bar{\mu})} + \frac{\mathbb{E} [\text{Cov} [IP_j, \mu_j \mid f]]}{\bar{\mu} (1 - \bar{\mu})}. \quad (\text{A31})$$

The across-firms covariance is related to the slope of the regression of IP_f on μ_f , as in equation (A23):

$$\frac{\text{Cov} [IP_f, \mu_f]}{\bar{\mu} (1 - \bar{\mu})} = \beta_{\text{between}} \cdot \eta_{\text{cons}}^{\text{between}},$$

where “between” indicates objects from the firm-level data. The second term in (A31) admits a similar decomposition:

$$\frac{\mathbb{E} [\text{Cov} [IP_j, \mu_j \mid f]]}{\bar{\mu} (1 - \bar{\mu})} = \bar{\beta}_{\text{within}} \cdot \eta_{\text{cons}}^{\text{within}},$$

where

$$\bar{\beta}_{\text{within}} = \frac{\mathbb{E} [\beta_f \cdot \text{Var} [\mu_j \mid f]]}{\mathbb{E} [\text{Var} [\mu_j \mid f]]}$$

and

$$\eta_{\text{cons}}^{\text{within}} = \frac{\mathbb{E} [\text{Var} [\mu_j \mid f]]}{\bar{\mu} (1 - \bar{\mu})}.$$

The former is the weighted average of firm-specific relationships between import shares and consumer base across barcodes within firms. Weights are proportional to the firm sales but also to the variance of consumer base, which is analogous to the well-known result by Angrist (1998) that OLS with fixed effects weights group-specific treatments by the group-specific variance of the treatment status. Correspondingly, $\eta_{\text{cons}}^{\text{within}}$ is the component of consumption segmentation that related to the within-firm differences in consumption baskets between the two groups. By the law of total variance, $\eta_{\text{cons}}^{\text{between}} + \eta_{\text{cons}}^{\text{within}} = \eta_{\text{cons}}$ is the full consumption segmentation at the barcode level.

Combining these characterizations, we conclude that the differential import spending at the detailed level can be represented as

$$\Delta_{Final} [IP] = \beta_{\text{between}} \eta_{\text{cons}}^{\text{between}} + \bar{\beta}_{\text{within}} \eta_{\text{cons}}^{\text{within}}.$$

Now assume that the regression slope within firms, $\bar{\beta}_{\text{within}}$, is between zero and β_{between} . Without loss of generality, suppose that $\beta_{\text{between}} > 0$, so $0 < \bar{\beta}_{\text{within}} < \beta_{\text{between}}$. Then $\Delta_{Final} [IP]$ lies between $\beta_{\text{between}} \eta_{\text{cons}}^{\text{between}}$, which is the estimate at the firm level, and $\beta_{\text{between}} \eta_{\text{cons}}$, which is larger by a factor of $\eta_{\text{cons}} / \eta_{\text{cons}}^{\text{between}}$ —the ratio of consumption segmentation indices across barcodes and across firms, which is

always above or equal to 1.

B.3 Bias-Corrected Estimation of Consumption Segmentation

For the population of goods j , which may be individual barcodes or firm-level composites, the consumption segmentation index is defined as

$$\eta_{\text{cons}} = \frac{\text{Var} [\mu_j^*]}{\bar{\mu}^* (1 - \bar{\mu}^*)},$$

where μ_j^* is the population share of skilled consumers in purchases of good j and the variance is weighted by some measure ω_j of the importance of the good (the square-root of sales with our main weighting scheme). We only observe a random sample of consumers h and their expenditures, from which we estimate the consumer base:

$$\mu_j = \sum_h s_{jh} \text{College}_h,$$

where College_h is an indicator for whether the consumer is a college graduate and s_{jh} is the share of consumer h in the observed sales of good j .

The section argues that the variance of μ_j across goods is an inconsistent, upwardly biased, estimate of $\text{Var} [\mu_j^*]$ and develops a bias-corrected estimator. We make two simplifying assumptions. First, we assume that μ_j is independent across goods, which would be true if the sets of consumers were not overlapping. Second, we treat consumption shares s_{jh} as non-random, only studying sample variation coming from the college dummy. An approach similar to the one in Appendix B.1 can be potentially developed to take random shares into account.

Since the sample of consumers is i.i.d. for each good, μ_j is unbiased for μ_j^* , and we can write $\mu_j = \mu_j^* + \epsilon_j$ with a mean-zero noise $\epsilon_j = \sum_h s_{jh} (\text{College}_h - \mathbb{E} [\text{College}_h | j])$ that is uncorrelated with μ_j^* . As a result, the weighted variance of observed μ_j across goods includes the fundamental variance of μ_j^* and the average variance of the noise:

$$\text{Var} [\mu_j] = \text{Var} [\mu_j^*] + \mathbb{E} [\sigma_j^2], \quad (\text{A32})$$

where $\sigma_j^2 = \mathbb{E} [\epsilon_j^2]$. If unbiased estimates $\hat{\sigma}_j^2$ are available for σ_j^2 , averaging them across goods yields an unbiased estimate of $\mathbb{E} [\sigma_j^2]$, which is also consistent if the set of goods is growing. Subtracting it from $\text{Var} [\mu_j]$ then yields a bias-correct estimate of $\text{Var} [\mu_j^*]$.

An unbiased estimator for σ_j^2 can be obtained by noticing that ϵ_j is just a weighted average of a random sample, hence

$$\sigma_j^2 = HHI_j \cdot \text{Var} [\text{College}_h | j], \quad (\text{A33})$$

where $HHI_j = \sum_h s_{jh}^2$ is the Herfindahl index that measures the (inverse) effective number of consumers of this good. Here $\text{Var} [\text{College}_h | j]$ depends on the good—for example, a good which fundamentally sells 99% to college graduates will have very little sample variation in μ_j . This fundamental variance can be

estimated as a transformation of the sample variance of College_h among observed consumers:

$$\begin{aligned}\mathbb{E} \left[\sum_h s_{jh} \left(\text{College}_h - \sum_{h'} s_{jh'} \text{College}_{h'} \right)^2 \right] &= \left(\sum_h s_{jh} (1 - s_{jh})^2 + \sum_h s_{jh} \sum_{h' \neq h} s_{jh'}^2 \right) \cdot \text{Var} [\text{College}_h | j] \\ &= (1 - HHI_j) \cdot \text{Var} [\text{College}_h | j].\end{aligned}$$

Since we treat weights, and therefore HHI_j , as non-random, $\sum_h s_{jh} (\text{College}_h - \mu_j)^2 / (1 - HHI_j)$ is unbiased for $\text{Var} [\text{College}_h | j]$. Plugging this into (A33), we get

$$\hat{\sigma}_j^2 = HHI_j \cdot \frac{\sum_h s_{jh} (\text{College}_h - \mu_j)^2}{1 - HHI_j}.$$

With equal weights, $\hat{\sigma}_j^2$ becomes a familiar unbiased variance estimator with $N - 1$ in the denominator. Plugging it into (A32), we obtain a consistent estimator of $\text{Var} [\mu_j^*]$:

$$\widehat{\text{Var} [\mu_j^*]} = \sum_j \omega_j (\mu_j - \bar{\mu})^2 - \sum_j \omega_j s_j^2.$$

Dividing through by $\bar{\mu}(1 - \bar{\mu})$ transforms it into the consumption segmentation index.

B.4 Estimation of Income Elasticities

Here we describe the procedure used to estimate income elasticities for each I-O industry in Section 6, based on the CEX data from Section 3.1. Our approach is inspired by equation (A22), which shows that ψ_j can be estimated directly from the relationship between spending and consumer expenditure, as long as the log-linear approximation works well and different consumers face the same prices. We take this approach in Section 6 instead of estimating the primitive parameters φ_j structurally.

We split households in the CEX sample into 11 bins by the reported pre-tax household income and compute consumption shares across 671 spending categories j for each of the bins i separately (s_j^i) and overall (s_j). Then for each spending category we estimate the income *semi*-elasticity by regressing, across income bins, spending shares on the log of total expenditure in this income group, averaged across households:

$$s_j^i = \text{constant}_j + \beta_j \log \text{Expenditures}_i + \text{error term}_{ij}.$$

Observations are weighted by number of households in each income bin. For an income-elastic spending category, the share is increasing in the total expenditures, so $\beta_j > 0$, and the reverse holds for income-inelastic products. We then convert the semi-elasticity into the elasticity ψ_j for an average consumer of product j :

$$\psi_j = 1 + \frac{\hat{\beta}_j}{s_j}.$$

The intermediate step with semi-elasticities guarantees that the spending-weighted average of income

elasticities across all spending categories is equal to one, as it should be theoretically:

$$\sum_j \psi_j s_j = \sum_j s_j + \sum_j \hat{\beta}_j = 1 + 0 = 1,$$

where $\sum_j \hat{\beta}_j = 0$ because spending shares sum up to a constant (one) for each income group, and the regression of a constant on $\log \text{Expenditures}_i$ yields a zero slope.

Expenditures are used on the right-hand side instead of income because in the CEX, total expenditures do not vary one-to-one with reported income. The relationship is increasing but much less than proportionate, which may be a consequence of imperfect measurement of income—either because current income is a bad proxy for permanent income, or for pure measurement error reasons. In either case, income elasticity estimates would be biased towards one if income was used on the right-hand side.

Table A16 presents the averages of income elasticities by classes of CEX spending categories, which are broadly consistent with the estimates from Aguiar and Bils (2015).

C Data Appendix

C.1 Consumer Expenditure Survey

CEX is a stratified household survey conducted by the U.S. Bureau of Labor Statistics that measures the universe of personal spending by around 650 detailed categories. CEX consists of two separate parts: interview and diary surveys, which we use in combination. Quarterly interviews cover the complete range of expenditures, whereas diaries focus on some categories, such as food and clothing, in much greater detail. The interview panel include around 6,900 households per quarter, each surveyed for four consecutive quarters. Diaries are collected for roughly the same number of distinct households per year but capture only two weeks of consumption. We select categories of spending (UCC) from both surveys according to the Integrated Stub file provided by the CEX, so that they cover all categories without double-counting.

The key advantage of CEX is that consumption structure can be measured separately for different groups of households. We split panelists by education of the household’s reference person answering the interview, defining college education as bachelor’s degree or higher, and by bins of household income before tax.¹⁰¹

To increase the sample size, we combine data from 2006–2008. We drop all households with income below \$5,000. Our final interview sample includes 32,668 unique households with average annualized spending of \$35,351, while the diary sample has 16,901 households spending \$13,384 per household per year. The distribution of spending by large groups is given in Table A16.

Expenditure on housing services requires special treatment. The range of CEX spending categories

¹⁰¹We use education of the reference person (variable EDUC_REF), although the results are nearly identical if we use education of both male and female heads of the household. For income, we use FINCBTXM in the interview survey and FINCBEX in the diary survey. Eleven income bins are defined by the following cutoffs (in \$000): 10, 20, 30, 40, 50, 60, 75, 90, 110, and 150.

includes rents and mortgage interest, but not the mortgage principal payments. However, an addendum section of the interview survey provides information on the self-reported rental value of owned property. In our static setup that is the closest analog to annual expenditures on housing for home-owners, so we add imputed rents to the set of UCC we consider.¹⁰²

C.2 Input-Output Table

We use the most recent detailed I-O table, which dates from 2007 . While BEA publishes annual tables with 71 relatively coarse three-digit industries, the 2007 one is disaggregated into 389 six-digit industries. These industries are groups of six-digit NAICS codes: while NAICS includes 581 goods and 565 service industries, the I-O classification includes 258 and 122, respectively, plus 9 special industries such as government and non-comparable imports. Some I-O industries are as detailed as NAICS (e.g. Electronic computer manufacturing), but in other cases aggregation is quite strong (e.g. 24 NAICS codes within Apparel manufacturing become a single category).

The use of the I-O table is complicated by the fact that the same product (“commodity”) can be produced by different industries: for example, trucks are manufactured by both truck and car manufacturing establishments. We follow the standard procedure to address this issue by using the Supplementary Tables after Redefinitions (Horowitz and Planting, 2009) and combining the Make and Use tables to produce a square commodity-by-commodity matrix.

C.3 Merged Nielsen-Census Sample

Data Sources. The Nielsen company asks around 55,000 U.S. households per year to record all purchases within certain classes of products. Consumers scan purchased goods using handheld barcode scanners provided by Nielsen. They also manually enter products that do not have barcodes, such as fresh produce. Nielsen obtains price information from a combination of store data and manual entry by households. The stratified sample of households is representative of the U.S. population in terms of income, education, age, race, household size, and other characteristics when using the Nielsen-provided projection weights.

GS1 maintains the concordance between barcodes and firm names and addresses; the version we obtained is complete as of February 2016. We drop 5.2% Nielsen barcodes which we could not link to GS1 (they constitute 1.8% of total sales in Nielsen). In most cases GS1 firms are located within the U.S., although there are some exceptions, mostly with Canadian addresses. We drop firms with addresses outside 50 U.S. states and Washington, D.C. or with missing state information, which constitute 4.3% of all Nielsen firms but only 0.75% of total sales.

Business Register, or SSEL, is the comprehensive list of establishments, with names and addresses, assembled using Census surveys, Internal Revenue Service tax data, and other data sources at the annual

¹⁰²Aguiar and Bils (2015) follow a similar approach. The information on the mortgage principal is also collected, but is less suitable for our purposes.

frequency (DeSalvo et al., 2016). Because firms change names and addresses over time, while GS1 provides only one observation per firm, we use addresses in the SSEL for all years from 1991–2014, which improves the quality of the merge.

The Economic Census is the survey of all business establishments in the U.S. It is conducted by the Census Bureau in years that end with 2 or 7, and participation is required by law. The content of the questionnaire varies across sectors and industries but all of them include questions on the total revenue. We primarily use Censuses of Manufacturing, Wholesale, and Retail. Establishments in Services, Finance, and Utilities are also part of our Economic Census sample, but they are rarely matched to Nielsen.

Finally, LFTTD (Linked/Longitudinal Firm Trade Transaction Database) is the microdata on all international trade transactions, based on the import declarations and shippers export declarations. It has been matched to the Census by firm identifier (see Bernard et al., 2009).

Sample Construction. We predict total sales of each Nielsen barcode by applying projection weights provided by Nielsen to the purchases by each household and, using the GS1 crosswalk, aggregate them to firms and firm-module cells. We classify households into college- and non-college by using education of both male and female heads. If they are both present but only one has college degree, we attribute half of the purchases to each education group. Income is reported in 16 discrete bins, and we use their midpoints.¹⁰³ Income is reported with a two-year lag, so we use the value from two years after, whenever available.

We apply several filters to Nielsen. First, we drop households with reported income below \$5,000. Second, we drop “magnet data”—products that do not use standard barcodes, such as fresh fruits and vegetables. Finally, we also drop firm-years with less than five unique barcode-household pairs and those with total spending by Nielsen panelists under \$100—we label those as “tiny” Nielsen firms. From now on, we will suppress mentioning years.

We then compute import shares for each Census firm. The numerator is total imports from LFTTD. To measure the total firm output in the denominator, we aggregate revenue of all establishments belonging to the firm. However, this creates double-counting if a manufacturing company ships its products to its own wholesalers or retailers and then sells them. Therefore, we only count the total revenue in the largest NAICS2 sector in which the firm operates, although the results are not substantially different without this correction. We drop firms for which imports exceed 200% of annual sales, indicating an imperfect match between LFTTD and the Census.

Finally, we merge name and addresses in GS1 with the Census firms—a procedure we describe next. Once done, we implement a consistency filter. Some firms, particularly large ones, span many industries, so their scope may not be covered well by the set of products covered by Nielsen. As a result, the overall importing behavior may be a very bad proxy for the set of products covered by Nielsen. We therefore

¹⁰³The cutoffs in \$000 are: 5, 8, 10, 12, 15, 20, 25, 30, 35, 40, 45, 50, 60, 70, and 100. In some years, the top-income group is decomposed further, but we use a consistent classification. We assign the top-income group the value of \$140,000, based on the average income in the years when we have more detailed data.

require that Nielsen sales of a firm are within the range of 1% and 300% of the Census sales. Although still wide, this range excludes strongest violations of consistency in both directions and makes our results robust to using the square-root of Nielsen or Census sales as weights.

Merging Process. We match names and addresses between GS1 and each year of SSEL from 1991–2014 separately. The process consists of three steps. First, we pre-process names and addresses in both datasets to maximize the probability of exact matches. Second, we develop a series of matching rules and apply them starting from the strictest, giving priority to multi-establishment Census firms. Third, because names and addresses change over time, some matches will only be found in some years. We extrapolate them to other years wherever possible. We now describe each step in detail.

Pre-processing. We use the algorithms from the `reclink2` package from [Wasi and Flaaen \(2015\)](#), with minor modifications. For company names, the `stnd_compname` command removes special symbols, makes standard substitutions (e.g., INTL to International), and isolates the entity type (e.g., INC) into a separate variable. Pre-processing of addresses is particularly important. The `stnd_address` command parses them into several parts: the main address variable (where special symbols are removed, street types are converted to their abbreviations, e.g., Street into ST, etc.), as well as the post office box, unit (e.g. SUITE 1400), and building numbers, if present. We implement an important addition to this parsing procedure by also extracting the house number from the address. We define it as the number at the beginning of the address or, if the address starts with a letter, the largest number in the address.¹⁰⁴

Matching Algorithm. The SSEL consists of records of three types: multi-unit (one per establishment for firms with multiple establishments), “submaster” (one per tax identifier of a multi-unit firm, created for consistency with the IRS), and single-unit. We give priority to multi-unit and submaster records by first attempting to match GS1 firms to them. For GS1 firms that are still not merged, we try matching to single-unit firms that are part of the LBD (the Longitudinal Business Database, which links SSEL records across years). The lowest priority is given to single-unit firms outside of the LBD.¹⁰⁵

Within each priority level, we apply consecutive matching rules, starting from the strictest one. Once a GS1 firm finds an SSEL match, it is removed from the process. This guarantees that each GS1 firm is matched to only one Census firm, except for rare cases when we find several matches using the same matching rule. At the same time, we allow several GS1 firms to be matched to the same Census firm, as should be the case for subsidiaries of the same firm that appear in GS1 separately.

We developed seven matching rules by trial and error and manually checked samples of matched firms

¹⁰⁴Extracting the largest number is inspired by the the addresses of foreign firms are treated in the LFTTD (see [Kamal and Monarch, 2016](#)). With fuzzy matching, matching on the house number ensures that buildings like 47 Main St. and 49 Main St. are distinguished. It is also very useful for parts of Wisconsin and Illinois which use alphanumeric addresses, e.g. “W190 N10768 Commerce Cir, Germantown, WI.”

¹⁰⁵One SSEL record may list up to two addresses per establishment (physical and mailing) and sometimes specifies two zipcodes (one reported and one inferred automatically based on the rest of the address). We use all available versions of the address to increase the probability of the match.

to verify that each of them mostly produces correct matches. Each rule requires an exact match and non-missing values for some key variables, an exact match on additional variables where missing values are allowed, and a bigram probabilistic (“fuzzy”) match on other variables with a specified match score threshold. The implementation is again based on the `reclink2` package from [Wasi and Flaaen \(2015\)](#). While we kept its logic, we substantially improved computational efficiency; the modified code is available upon request.

Table A4 lists the rules. The two strictest rule require a non-missing match is required for the 9-digit zipcode (ZIP+4). Although available only for some firms, it generally identifies the building or a post box precisely. The first rule additionally requires an exact (possibly missing) match for the firm name, house number, address, PO Box, unit, and building, standardized as previously described, while the second rule only requires an exact match on the house number, while the other variables are matched in a fuzzy way. The least restrictive seventh rule requires exact matches on the firm name, its entity type, and state, still delivering high quality of matches for the records that have not been matched using stricter rules.

Extrapolation of Matches. Matching with GS1 is done separately for each year of the SSEL. If a GS1 firm does not find any SSEL match in a given year t , we turn to the matches that were found for this firm in other years, with preference to the closest years.¹⁰⁶ If some match is found in year t' , we check in the LBD whether the matched firm existed in t and, if so, use this match for year t .

Match Statistics. Panel (a) of Table A5 shows that the majority of Nielsen firms, excluding tiny ones, is matched, covering over 83% of total Nielsen sales.¹⁰⁷ In 2007, there were 26,900 Nielsen firms, and elimination of the tiny ones leaves us with 11,000 without any significant loss in total projected sales. Out of them we are able to find a Census match in the same year of the Census Business Register for 7,600, while using names and addresses from other years adds another 600 firms, making it 8,200 total. Although all firms are supposed to fill out Census forms, not all of them do, so we find 7,200 Nielsen firms in at least one of the Censuses, and of them 6,100 pass the consistency filter. Although there are a few cases where we find two Census matches for the same Nielsen firms, the number of Nielsen firms with single matches is the same 6,100 after rounding. Statistics are similar for 2012, increasing the sample size to 12,700 firm-years.

Panel (b) of Table A5 shows merging statistics starting from Census firms. Since Nielsen only covers consumer packaged goods, we do not expect a high match rate in most industries. However, Nielsen coverage is strongest for food, alcohol, and tobacco. This panel starts from all 51,500 firms in the Census of Manufactures in the corresponding NAICS code 312. Out of them, 8,900 (or 17.3%) are merged to any Nielsen firm, including the tiny ones, and the merged ones account for 79% of the total sales. After dropping small Nielsen firms and implementing the consistency filter, we match only 9.3% of the firm count but still 58.7% of sales by all manufacturers in the industry. Note that we also merge many wholesalers

¹⁰⁶For year t , we start with $t + 1$, then use $t - 1$, $t + 2$, $t - 2$, $t + 3$, etc.

¹⁰⁷The match rate is above 83% of sales for food and health and household products, but a bit worse for general merchandize, at 76%.

and retailers selling food, not accounted for in this table. Table A6 shows that multi-establishment firms are a minority in the matched sample (29%), but they cover 93% of sales. Within both multi- and single-establishment matched firms, the strictest matching rule 1 captures the largest share of firms, but all rules contribute to the sample.

Table A7 shows the fractions of firms operating in different sectors, defined by their 2-digit NAICS codes, in the sample.¹⁰⁸ The manufacturing sectors constitutes the largest fraction of the sample (57.2% with square-root weighting), followed by wholesaling (29.0%) and retailing (8.7%). The smaller share of retailers is in part determined by their large average sales, which imply that the square-root weighting scheme reduces their importance. At the same time, it is important to understand that most products sold by retailers are registered by other firms. Appendix D shows that this is true even for products manufactured for and distributed exclusively by Walmart. Among the 3-digit NAICS codes, Food Manufacturing and Nondurable Goods Wholesalers are the most prevalent ones, followed by Chemical Manufacturing (which includes soap, shampoos, etc.) and Beverage and Tobacco Manufacturing.

The last column of Table A7 presents a nice test on the quality of the match. Nielsen data allow us to identify products that are branded by the retail chain that sells them (“private label brands”). We find that over 99% sales of barcodes registered by food and beverage stores, according to their main NAICS code in the Economic Census, are private label brands. For comparison, this share is only 7.9% for wholesalers and mere 1.2% for manufacturers.

Table A8 examines how representative the matched sample is. Panel (a) compares firms in Nielsen, excluding tiny ones, that found a match to those that did not. Median firms in the merged sample have about twice as large Nielsen sales relative to the firms that did not find a match. Matched firms also sell to slightly, but statistically significantly, poorer and less educated consumers. For example, 29.1% of sales of matched firms is to college graduates, as opposed to 30.7% for firms that we did not match. However, these differences can largely be explained by the size difference; they are reduced when controlling for a quadratic term in log Nielsen sales. Panel (b) provides evidence on sample selection for the firms in the Census of Manufactures producing food, alcohol, and tobacco. Again, merged firms are much larger, with median sales of \$13.3 million, payroll of \$1.9 million and 54 employees, as opposed to \$606,000 sales, \$113,000 payroll and 4 employees for a median Census firm that we did not merge. Comparing these sets of firms by skill intensity (the payroll share of non-production workers) does not reveal statistically or economically significant differences.

C.4 CEX and Ward’s Data on Automobiles

To measure consumer characteristics by auto brand, we use the OVB file (“Owned Vehicles Detailed Questions”) from the CEX Interview Survey, which asks respondents to provide information about all cars they own, including the brand, whether the car was purchased new or used, and in some cases the price.

¹⁰⁸Because Census data provides NAICS codes for establishments not firms, we classify firms by the 2- and 3-digit NAICS in which they have the largest payroll, excluding NAICS code 55 “Management of Companies and Enterprises”.

The data are available since 2006 but we use it for 2009–2015 for consistency with the Ward’s sample. Each household is expected to participate in the survey for four consecutive quarters, so to avoid duplication we only use the most recent survey in which the OVB survey is filled. Still, the same household may own several cars. Like in other datasets we build, we drop cars owned by households with income (before tax) below \$5,000. We only include automobiles and exclude all other types of owned vehicles.

Importing data come from Ward’s Automotive Yearbooks. We use the electronic versions of the 2011, 2013, 2014, and 2016 yearbooks. Each of them shows the statistics for the previous two years, thus covering the 2009–2015 period. In each year we use four Ward’s tables. One is on sales (U.S. Car Sales by Line by Month) and for each model (“line”, e.g. Chevrolet Camaro) it decomposes the number of sold cars into those built within and outside NAFTA. We use the other three tables (U.S. Vehicle Production by Line by Month and same for Mexico and Canada) to decompose cars produced within NAFTA into those built domestically versus imported from Canada or Mexico.

We first aggregate all years of Ward’s data to measure, for each model, the number of cars sold, the share of them from outside NAFTA, and the shares of production within NAFTA that comes from the U.S., Canada, and Mexico separately. We then compute the domestic share of each model sales as the product of those from within NAFTA (from the sales table) and the share of U.S. within NAFTA production (from production tables). For two models only, the sales table reports some NAFTA production, but production data are missing, in which case we assign all NAFTA production to the U.S.

At the end we aggregate all models by brand using sales weights from Ward’s. We find 35 brands in both CES and Ward’s data. We also keep four brands (Daewoo, MG, Land Rover, and Austin-Healey) which are in CEX but not in Ward’s, and are fully imported. This results in the sample of 39 brands listed in Table [A15](#).

C.5 Microdata on Automobiles

We use the 2012 version the Census of Manufacturers and the Customs data for the same year. At the same time, to increase sample size we measuring consumer base using all years of the CEX when the brand variable is available, 2006–2015 (see Section [C.4](#) on the data description).

To match domestic car producers in the CEX, we first link each car brand to the firm that owned it in 2012, using the Ward’s Automotive Yearbook and Internet search. Then we manually search for firm names in the 2012 Business Register (SSEL)—the list of all establishments in the U.S., and obtain the firm identifier or identifiers for all firms that participated in the Census.

Our sample includes two types of observations. If a firm has no production in the U.S., we keep its brands separately and assign 100% imports, both direct and total. And if a firm has some U.S. production (and participated in the 2012 Census of Manufacturers),¹⁰⁹ we aggregate its brands together and measure import shares.

¹⁰⁹Participation in the quinquennial Census is required by law, so the vast majority of firms reply. However, not all of them do, and the information on participation is confidential.

The value of imports of assembled cars is defined as total imports in the Customs data in the Harmonized Trade Classification (HS) code 8703 “Motor cars and other motor vehicles principally designed for the transport of persons”.¹¹⁰ Imports of car parts are defined as those in HS codes 8706 (chassis fitted with engines), 8707 (bodies for motor vehicles), 8708 (parts and accessories of motor vehicles), 84 (machinery), 85 (electrical machinery and equipment), 90 (measuring and other instruments), 39 (plastics), 40 (rubber), 73 (articles of iron and steel), 83 (miscellaneous articles of base metal), and 94 (furniture).

We measure car sales by the sum of total shipments of domestically assembled cars and the imports of assembled cars. The former is defined as the total value of shipments from all of the firm’s establishments which belong to NAICS code 33611 (Automobile and light duty motor vehicle manufacturing) in the Census of Manufactures. Then direct (total) import share is the ratio of imports of cars (cars plus parts) in car sales.

C.6 Imputation of Skill Intensity for Detailed Industries

Our goal is to decompose the total payroll in each detailed I-O industry by education group. To do so, we first do it for more detailed six-digit NAICS (N6) industry codes and then aggregate up by I-O. We use two pieces of data. First, from QCEW we observe total payroll and the average wage $\bar{w}_{N6}$ for each N6 industry. Second, from IPUMS ACS we know payroll of college- and non-college workers separately, from which we compute skill intensity—the college share of payroll v_{IND} , but it is only available for more aggregated industries, based on the ACS variable IND.¹¹¹

In the model, skill intensity of any industry or group of industries j is linearly related to the average wage:

$$\begin{aligned}\bar{w}_j &= w_U (1 - v_j) + w_S v_j, & \text{hence} \\ v_j &= \alpha_0 + \alpha_1 \bar{w}_j,\end{aligned}\tag{A34}$$

where $\alpha_0 = -w_U / (w_S - w_U)$ and $\alpha_1 = 1 / (w_S - w_U)$. Empirically, we recognize that this relationship differs across sectors, so we allow α_0 and α_1 to vary across two-digit NAICS sectors, denoted N2.

Equation (A34) holds in theory both across more aggregate IND industries, where we observe both v_j and $\bar{w}_j$, and across detailed N6, where we only observe the right-hand side variable. Therefore, we estimate this equation at the IND level (by least squares with payroll weights) and then impute v_{N6} from it. The prediction equals:

$$v_{N6}^{\text{imputed}} = v_{IND} + \alpha_{1,N2} \cdot (\bar{w}_{IND} - \bar{w}_{N2}).$$

We verify that estimates of $\alpha_{1,N2}$ are positive for all sectors and that wages have substantial predictive power: the adjusted R^2 of regression of v_j on N2 fixed effects goes up from 65.1% to 82.8% when wages (with N2-specific coefficient) are included. We constrain the imputed skill intensity to lie between 0 and

¹¹⁰This HS code includes some vehicles besides cars (e.g. SUVs and ambulances), which may create some upward bias.

¹¹¹Only in one case (NAICS industry 519130), the same N6 code corresponds to two IND codes. We split this N6 code into two proportionately to the IND payroll.

1 in rare cases where the prediction is outside this interval.

This imputation preserves the average skill intensity from ACS. This allows us to build a weighted crosswalk to assign for each ACS respondent a set of probabilities that this person works in each N6 and thus I-O industry, based on their IND industry and college dummy. These probabilities differ between college and non-college workers: for example, a N6 industry with low average wages will have lower weight for college graduates. We use this crosswalk to do inference for measures of differential labor market exposure in Section 6.2.

C.7 Manufacturing Microdata

To measure the relationship between skill intensity and exporting at the plant level in Appendix E.3, we use Census microdata. We focus on the manufacturing sector because it is the only one where information of the worker types is available,¹¹² and it is the most tradable sector, too.

Until recently, Census surveys did not ask establishments about education of their workers, which led to a long tradition to proxy for skill intensity by the payroll or employment share of non-production workers (e.g. Berman et al., 1994; Autor et al., 1998), who are considered to be more skilled than production workers (Berman et al., 1998). The situation has changed with the arrival of the 2010 MOPS survey, which is a supplement to the Annual Survey of Manufactures (ASM), which covers all largest firms as well as a sample of smaller ones.

We use MOPS questions 32–35, which ask for number of managers and employees, as well as the share of managers and non-managers with a college (bachelor) degree.¹¹³ The shares are listed in terms of discrete bins, so we use the midpoints of those bins.¹¹⁴ This yields an estimate of the share college graduates in total employment, v_j^{Emp} . Unfortunately we do not observe wages of college- and non-college workers. Therefore, to impute the payroll share we use the economy-wide average wages of these groups from the U.S. Census Bureau (DeNavas-Walt et al., 2011). They show that the median wage of college graduates is about 80% higher than that of non-college workers (considering individuals in the labor force and 25 years or older), so we measure the payrolls share of college graduates in each establishment j as

$$v_j = \frac{1.8 \cdot v_j^{\text{Emp}}}{1.8 \cdot v_j^{\text{Emp}} + (1 - v_j^{\text{Emp}})}.$$

It is very strongly correlated with v_j^{Emp} , so the details of imputation are not important.

Besides the MOPS sample, we use the 2010 ASM and the full 2007 CMF. We match them to the Customs microdata (LFTTD) to measure exports. Like Bernard et al. (2018), we do not use the CMF and

¹¹²Analysis on other sectors could be possible by the Longitudinal Employer-Household Dynamics (LEHD) data, but we do not have access to it.

¹¹³The questionnaire is available at <https://www2.census.gov/programs-surveys/mops/technical-documentation/questionnaires/mop-2010.pdf>; also see Bloom et al. (2016). We drop observations where any of these questions is missing.

¹¹⁴The bins are under 20%, 21–40%, 41–60%, 61–80%, and over 80% for managers and 0%, 1–10%, 11–20%, and over 20% for non-managers (we assign 25% to the last category).

ASM question about plant exports, which is less reliable than direct observation of trade transactions. For firms with multiple establishments, we attribute firm exports proportionately to the value of establishment sales (shipments). We drop firms where exports exceed twice the total value of manufacturing sales, as those are likely to result from measurement error or other firm establishments which are not part of the sample (e.g. the non-manufacturing ones). We compute the export share of an establishment relative to the value of shipments.

D Examples of Products

To verify our understanding of which firms register barcodes, we visited a Walmart store and photographed a sample of products, both domestic and imported according to their labels. Then we identified them in the GS1 database by the barcode and searched for the information about the firms that registered them in the Internet. Figure A11 shows pictures of the five of them that illustrate well different situations we observed.¹¹⁵

Panels (a) and (b) show two plates labeled as Made in the USA, one from an independent brand and the other one distributed by Walmart. According to the GS1 data, the blocks of barcodes they belong to were registered by World Kitchen, LLC and Merrick Engineering Inc., respectively. Internet search verified that both of them are manufacturing firms in the U.S. , so we will recognize these products as domestic, unless their firms import a lot of materials.

Three remaining products are imported. Bed sheets in Panel (c) belong to the same brand as (b), distributed by Walmart, yet they are made in China. Correspondingly, the barcode is registered by Jiangsu Royal Home USA, Inc which according to Internet sources belongs to the NAICS code 423220 “Home furnishing merchant wholesalers” and imports from China.

Plates in picture (d), also made in China, are registered by First Design Global, Inc, which is a manufacturing firm but it imports tableware and kitchenware from China. We will therefore attribute these plates partially to imports, in proportion to the fraction of imports of this firm to its sales. That does not introduce a bias if this firm manufactures other products domestically and they have similar buyer characteristics to the imported ones.

Finally, the Canadian hair conditioner from picture (e) is distributed by Walmart and, unlike previously considered products, was registered by Walmart itself. Therefore, in the Nielsen-Census merge we will view the probability that it is imported as the fraction of Walmart’s direct imports relative to its total Census sales. This may be an underestimate if Walmart’s direct imports in the Customs data mostly cover its own-registered products, whereas its sales include all products, e.g. those from all previous pictures.

¹¹⁵We have not used any Nielsen data in this section. These products may or may not be in our final sample.

E Additional Evidence and Robustness Checks

E.1 Differential Spending on Imports of Consumer Packaged Goods

In this section we check robustness of the results to different skill measures and weights and discuss some concerns with the estimation, stemming from underestimating of retailer imports and measuring imports at the firm rather than barcode level. Finally we exploit differences in the magnitude of the differential import spending across product modules to make a humble prediction for the product categories outside Nielsen.

Product Classes. Table A9 verifies that the patters of differential import spending hold across product classes, both for imports from China and for other imports grouped together. Columns (1)–(3) show that college graduates buy more imports from countries other than China in all product classes, although the effect is particularly strong for food (17.1% of the average). Similarly, in all of columns (4)–(6) college graduates spend less on Chinese products. The strongest effect is in health and household products, but its magnitude is only 5.4% of the average share of Chinese imports in that class.

Final and Intermediate Products. Table A10 attempts to isolate direct and indirect imports, i.e. imports of final goods and intermediate inputs. We do not classify *products* into final and intermediate. Instead we consider the main activity of the firm that registered the barcode. We view imports by wholesalers and retailers as direct imports. We would like to interpret imports by U.S. manufacturing firms as imports of intermediates, but if a manufacturer is engaged in multinational production, it may import final products as well. This indeterminacy turns out to be relatively unimportant, as most of the differential import spending comes from direct wholesaler imports. They capture the entire pro-skilled effect for total imports (columns (1)–(3)) and over 60% of the anti-skilled effect for Chinese imports in health & household products (columns (4)–(6)).

Patterns Across Income Groups. While the main results split consumers into two groups by college education, we check robustness of the results to using income as a measure of skill. Figure A4 splits consumers into 15 household income bins (which is how income is reported in Nielsen) and measures the average spending on imports for them. The patterns are monotonic, with the fraction of non-China imports varying between 6.3% for the very poor to 7.6% for the very rich in all products covered by Nielsen (Panel (a)). Similarly, for China the fraction of imports falls monotonically with income from 6.8% to 6.4% for health and household products (Panel (b)) and from around 19% to 17.5% for general merchandize (Panel (c)).¹¹⁶

¹¹⁶Because of the compositional differences across product classes (rich people buy relatively more general merchandize than food), the fraction of spending on Chinese products in the full sample is not decreasing in income (see Panel (d)). This does not affect our results since we are interested in within-industry patterns.

Weighting Schemes. All main results are based on the square-root weighting scheme, which reduces the influence of a small number of giant firms for which our proxy for the import share is noisier. In unreported results we verify that all findings are very similar, both qualitatively and quantitatively, when firms are weighted by their Nielsen sales to the power of $1/4$ or $3/4$ instead of $1/2$, as well as by square-root of the firm’s sales in the Census.

Table A12 presents the estimates of the differential import spending using full Nielsen sales weights. Within-industry differences follow the same pattern as with square-root weights in Tables 3 and A9: imports from China are anti-skilled while imports from other countries are pro-skilled, and the latter dominates in the total. However, the magnitudes of within-industry differences are weaker than before: they are under 6% of the average pro-skilled for all cases excluding China and under 2.5% anti-skilled for China. Also, across-industry differences play a much bigger role.

Import Shares for Retailers. We measure the import share for a barcode as the fraction of imports in sales of the firm that registered it. This proxy underestimates the total share of imports in the firm’s sales if the domestic firm is buying imported products through domestic intermediaries. In the industry-level analysis of Section 3, such higher-order indirect spending contributes only 2.9 p.p. to the total spending on imports of 13.7% (see Table 1).¹¹⁷ However, it is likely to be much more important in the retail sector, where firms obtain products they sell from wholesalers, including foreign products. As a consequence, retailers in our Nielsen sample (with square-root weights) have a low average import share of 2.2%, compared with 7.4% for manufacturers and 16.5% for wholesalers (Table A13).

It is important to highlight that this underestimated import share applies only for barcodes that are *registered* by retailers—think of a subset of Walmart own brands rather than everything Walmart sells (see Appendix D). Although most of Nielsen sales happen *through* retailers, they involve products registered by other firms, for which we have independent measures of import shares. Only 18.6% of total Nielsen sales is in products registered by retailers (8.7% with square-root weights).

Still, retailer brands target less skilled clients, with the average share of college graduates of 28.6% for retail relative to 28.9% for manufacturers and 29.7% for wholesalers (Table A13). Therefore, adding missing retailers imports could potentially create an anti-skilled pattern of trade. However, columns (3) and (6) of Table A10 suggest that such bias is unlikely to be important. The differential spending that is generated by retailers is under -0.02 p.p., both for total imports and for imports from China in health & household products. Even if imports by retailers were underestimated by a factor of ten, this would not make any significant difference.

Attenuation Bias from Aggregation within Firms. One limitation of our data is that imports are measured at the firm level. Although the analysis is much more detailed than any industry-level study (recall that we have 12,700 firm-year observations for only 71 I-O industries), the import proxy is aggregated across barcodes within each firm. We use our regression representation (A23) to develop

¹¹⁷7.2% is direct imports of final goods and 6.5% is captured by direct imports of intermediates.

a theoretical characterization for the attenuation bias such aggregation creates. We find that the true differential import spending may be at most twice as large as the previous results found.

To build intuition, imagine the most extreme case where import spending of one type, for instance the unskilled, is actually zero, implying a very strong anti-skilled effect of trade even if cross-firm analysis does not find any. This scenario requires that each firm sells some products only to the skilled and the others only to the unskilled, and imports are used only for the latter type. However, segmentation of consumption between barcodes within firms is observable in the Nielsen data, and it is quite moderate—there is strong overlap between the consumption baskets of the two types.

If segmentation is not very large, trade may still be anti-skilled if the true relationship between import shares and the consumer base at the barcode level is negative and steep. We find this implausible given that Figure 2 (Panel (a)) revealed a moderate positive slope across firms. Firms tend to have a general sourcing strategy which spreads to its products, even those selling to different types of consumers. Therefore, we expect differences between import shares to be smaller within firms than between, conditional on the difference in the consumer base, and have the same sign.

Appendix B.2 formalizes this assumption and proves that in that case, aggregation to the level of firms attenuates the differential import spending by a factor that is bounded between one and the ratio between barcode- and firm-level consumption segmentation indices. If there are much stronger differences between consumption baskets of the two groups across barcodes than across firms, attenuation may be substantial.

If we could observe consumer base of each barcode perfectly, computing this ratio would be a trivial exercise. Unfortunately, for many barcodes the fraction of college-graduated consumers has to be estimated from a small sample of consumers, which leads to an excess variance and upward biased consumption segmentation. To address this issue, Appendix B.3 develops a methodology to provide an unbiased estimate of segmentation. The key insight is that consumer base is just a sample of average of college dummies for all observed consumers, so an unbiased variance estimate can be computed for each barcode. The excess sample variance of consumer base, relative to the true variance, equals the average of noise variances for each barcode. In a large sample of barcodes, unbiasedness of each variance estimate is sufficient for consistency of the average. Then excess consumption segmentation is just the excess variance rescaled in the same way as in (A26).¹¹⁸

Table A14 estimates segmentation at the barcode and firm-module levels for all products and for each of the three product classes separately. Without bias correction, it appears that segmentation is more than twice as high at the barcode level, which would allow for strong attenuation of the differential import spending. However, almost half of the barcode-level segmentation turns out to be noise, and the ratio of bias-corrected segmentation indices is only 1.42. It is slightly larger for health & household (1.61) and even more so in general merchandize (2.17), which may partially explain why we were finding smaller patterns there.

Multiplying estimates from Table 3 by these ratios, we get the upper bound on differential import

¹¹⁸A similar approach to bias correction is found in Chetty and Hendren (2017, sec V.A).

spending at the barcode level under our assumptions. The maximum anti-skilled effect is for China in health & household products, and it is bounded by 8.7% of the average. At the same time, the pro-skilled differential spending on imports excluding China in food may potentially be as large as 22.8% of the average.

Extrapolation Outside Nielsen. Industries covered, at least partially, by the Nielsen data account for over 40% the total expenditure in goods. This includes food, alcohol, and tobacco industries, but also a wide range of other manufacturing of final goods: chemical manufacturing (e.g. shampoos and cleaning supplies), fabricated metals (e.g. cutlery and cookware), electrical equipment (e.g. small appliances and bulbs), some electronics, etc. These industries are quite diverse in terms of the imports shares, overall and from China. We exploit this diversity to show that the differential import spending as a fraction of the average is weaker in product modules which have higher import penetration—which is where the effect is most important.

Panel (a) of Figure A5 bins product modules by their overall import penetration and shows on the vertical axis the differential import spending share between the skilled and unskilled groups relative to the average. There is a clear pattern: trade is up to 10% pro-skilled but only in the *least* exposed product modules. Product modules with import penetration above 10% exhibit differential effects under 5% of the mean. Panel (b) of Figure A5 repeats the same exercise for Chinese imports within health & household products—the product class where we found strongest anti-skilled result. Here the anti-skilled bias may reach 10–15% of the average, but again only in the modules with the lowest import penetration.

These patterns, particularly if they hold outside Nielsen products too, make a substantial expenditure channel is even less likely.

E.2 Differential Spending on Imports of Automobiles

Patterns Across Income Groups. Figure A6 shows shares of imports by bins of the household income of the owner. As in Table 4, rich consumers buy more imported cars overall, driven by their spending on imports from outside NAFTA. An interesting nuance revealed by this graph is that the bias towards non-NAFTA imports is particularly strong at the higher end of the income distribution. The fraction of spending on these imports gradually grows with income from around 20% in the lower tercile to around 30% at the 80th percentile. Then inflection happens, and the import share exceeds 45% at the very top.

E.3 Differences in Labor Market Exposure to Trade

Import Penetration and Skill Intensity over Time. The positive relationship between skill intensity and import penetration within the goods-producing sector appears to be a relatively recent phenomenon. To show this, Figure A9 groups manufacturing industries into bins by their skill intensity in 1992, 1999, and 2007 (Panels (a)–(c), respectively) and reports the average import penetration for each bin in the corresponding year. Because we do not have the I-O tables for years other than 2007, we use a

combination of the NBER CES dataset to measure domestic output and the Schott (2008) imports data, both for 6-digit manufacturing NAICS industries.¹¹⁹ Skill intensity in the NBER CES is measured as the payroll share of non-production workers.

The figure shows that the relationship between import penetration and skill intensity was flat in 1992, weakly increasing in 1999, and became quite steep in 2007: moving from the 25th percentile of skill intensity to the 75th percentile, the import penetration grows by 6.6 p.p., or 38% of its mean.

Skill Bias of Exporters within Industries. It is well-known that exporting firms, which are more productive than non-exporters, tend to also have higher skill intensity (for U.S. evidence, see for example Bernard et al., 2007). Column (1) or Table 6 captures the across-industry component of this skill bias of exports. Burstein and Vogel (2017) find in a multi-country Bernard, Eaton, Jensen and Kortum (2003) type setting that the within-industry skill bias of productivity and exports is sufficient to generate a sizable pro-rich earnings channel. However, reduced-form evidence in the U.S. is mixed. Bernard et al. (2007) find in the 2002 Census of Manufactures that exporting establishments have 19% higher skill intensity relative to the non-exporting ones, and 11% of it survives when industry fixed effects are included. In contrast, the estimates for 2007 from Bernard et al. (2018) suggest a 6% difference, out of which only a statistically insignificant 1% is within industries. Skill intensity is measured in both cases as the employment share of non-production workers.¹²⁰

We address the same question in a way consistent with Table 6, measuring the average export share for establishments employing skilled and unskilled people and comparing the two. We estimate the payroll of college-educated and other workers for over 33,000 manufacturing establishments in the 2010 Management and Organizational Practices Survey, which is a supplement to the Annual Survey of Manufactures (MOPS, see Bloom et al., 2016). We also check robustness to using the payroll of non-production and production workers in the full 2007 Census of Manufactures (same data as in Bernard et al., 2018), 2010 Annual Survey of Manufactures which covers larger firms, and its MOPS subsample. Appendix C.7 describes data construction.

Table A19 presents the results. In all cases exporters are more skill-intensive than non-exporters but the difference is mostly across industries, so it was already captured in Table 6. The fraction of the within-industry component of the difference varies across specifications from 0.2% to 17.1%. The largest number corresponds to the MOPS case where skill is defined by college education, but is still relatively small.¹²¹

Moreover, theoretically speaking these differences are likely to overestimate the importance of the

¹¹⁹Import penetration is measured slightly differently than before as the ratio of imports to imports plus domestic output. We do not subtract exports because, due to imperfect concordance from the HS codes, absorption becomes negative in some industries. We verify that the same pattern holds in the SIC-level data from Autor et al. (2013).

¹²⁰There are other differences in the calculations, e.g. in how exports are measured in the Census data (see Bernard et al., 2018).

¹²¹Figure A8 visualizes this pattern. For each establishment we compute the export share and the six-digit I-O industry mean export share. Then averages of both measures are plotted for each bin of establishment skill intensity. If most differences were within industries, the industry-level export share would be unrelated to skill intensity, while on the picture the slope is almost as strong as when the establishment export share is used.

within-industry skill bias of exporters for the earnings channel. In models with the extensive margin of exporting, such as [Eaton and Kortum \(2002\)](#), the elasticity of exports with respect to export barriers reduction is smaller for more productive firms. In that model, most productive firms already export, and therefore the liberalization only has an intensive margin effect on them. At the same time, less productive firms are more likely to enter exporting, which creates an additional source of export growth lower in the distribution (while the intensive margin is the same, determined by the elasticity of demand). As a result, productive skill-intensive firms do not expand as much as they would have with a constant trade elasticity, and the demand for skilled labor does not grow as much.

We conclude that the earnings channel may be stronger than we find in [Table 6a](#) because of the [Burststein and Vogel \(2017\)](#) mechanism, but the bias is unlikely to be strong.

Skill Bias of Importers. It is impossible to observe which domestic firms face more import competition within an industry. Yet, if these firms on average have lower skill intensity (e.g. [Bernard et al., 2006](#)), our measure of differential exposure to import competition, based on the industry import penetration, would be biased. Here we use the insight from [Borjas et al. \(1997\)](#) to address this issue in a somewhat heuristic way. Specifically, we assume that the marginal domestic worker displaced by import competition is representative of the skill mix of her industry in the past—in 2000 or in 1990 in our estimation.

We embed this idea in our theoretical framework in the following way. Assume that each industry has two segments: traditional segment A and hi-tech segment B. skill intensity (the college payroll share) of segment A is v'_j , and we assume it is observable; the fraction of this segment in the industry output is denoted λ_j , which is not observed. We do not make restrictions on the skill composition of segment B. We examine the case where all imports in the industry are concentrated in segment A, while B is completely insulated from competition with foreign varieties. This implies that the marginal mix of workers affected by trade has skill intensity v' . Besides differences in import penetration and skill intensity, these two segments are identical: they have the same share and composition of intermediate imports, the same export share, etc. Both final and intermediate consumers view composite products of A and B as two generic industries within their sector with the same preference parameters.

This setup implies that all the theoretical results of the paper go through, but there is aggregation bias: we need to measure the differential exposure to trade at the level of segments, while the data do not have that level of detail. However, the exposure of each group to exports, imported inputs, and income effects is unchanged because segments A and B are identical in those respects. We will show now the knowledge of v' is sufficient to measure the differential exposure to imports.

We first establish that the import penetration from some foreign country c within segment A equals $IP_{jA,c} = IP_{jc} / (IP_j + \lambda_j (1 - IP_j))$. Indeed, if the industry absorption is normalized to one, IP_j is the total value of imports in the industry, and therefore in its segment A, while the denominator is the segment's absorption, which consists of imports as well as share λ_j of the industry's output purchased domestically, $1 - IP_j$. Import penetration in B is zero by construction.

Before measuring the exposure to import competition by worker type, we note that the average

exposure in the industry equals

$$IP_{jc} \cdot \frac{\lambda_j}{IP_j + \lambda_j(1 - IP_j)} = IP_{jc} \cdot \frac{\lambda_j}{\lambda_j + (1 - \lambda_j)IP_j} < IP_{jc}.$$

By allowing for heterogeneity of import penetration within the industry, and therefore for specialization between countries, we reduced the effect of import competition even if $v'_j = v_j$. For the same reason the payroll-weighted average import penetration was found to be small in Section 6.2. To focus on the skill bias of import competition, we assume that either segment A is sufficiently bigger than segment B or import penetration is sufficiently low, so $\frac{1-\lambda_j}{\lambda_j} \cdot IP_j \ll 1$, and the average exposure to imports is unaffected by having two segments per se.

Given that, the average exposure to imports of skilled workers in the industry equals the product of the payroll share of segment A for these workers and the import penetration in A:

$$\frac{\lambda_j v'_j}{v_j} \cdot IP_{jA,c} = \frac{v'_j}{v_j} \cdot IP_{jc} \cdot \frac{\lambda_j}{IP_j + \lambda_j(1 - IP_j)} \approx \frac{v'_j}{v_j} \cdot IP_{jc}. \quad (\text{A35a})$$

Similarly, exposure of the unskilled workers equals

$$\frac{\lambda_j(1 - v'_j)}{1 - v_j} \cdot IP_{jA,c} \approx \frac{1 - v'_j}{1 - v_j} \cdot IP_{jc}. \quad (\text{A35b})$$

Equations (A35a)–(A35b) constitute a very simple result: to measure the differential exposure in presence of the skill-bias of import competition, one needs to adjust exposure of the skilled group up by a factor v'_j/v_j , the exposure of the unskilled group down by $(1 - v'_j)/(1 - v_j)$, and the knowledge of the shares of segments is not required as long as segmentation of imports within the industry does not change the average exposure. Since we assume that the two segments are identical in other respects, including the types of downstream domestic firms that sell from them, when adjusting for input-output linkages we use the overall import penetration in downstream industries.

We implement these adjustments using the 2000 and 1990 IPUMS ACS data from the population censuses, constructing the samples in the same way as in Section (6). Table A20 present the results. As a benchmark, Columns (1) and (2) use the 2007 skill intensity, so the results are unchanged relative to the case with just one segment per industry, as in Table 6. College workers have lower exposure to imports but the difference is only 3.5% of the average without the I-O adjustment, or 7.3% with it. Columns (3) and (4) use the 2000 data, where college graduates accounted for 46.1% of the total labor income, compared to 49.7% in 2007. Correspondingly, the difference in exposure to import competition grows substantially, to 25.3% (18.5%) of the average without (with) the I-O adjustment. The pattern becomes even more striking when using the data from 1990 in columns (5)–(6), with the average skill intensity of only 41.4%. The differential exposure equals 59.4% (36.1%) of the average without (with) the I-O adjustment.

These patterns suggest that a sufficiently strong skill bias of importing can generate a strong mech-

anism for the pro-skilled effects of trade liberalizations.¹²² However, this analysis attributes the entire growth of skill intensity in the economy to trade, which leads to overestimation of this mechanism. Indeed, the higher differential exposure in columns (3)–(6) relative to (1)–(2) is mostly due to the shift in the economy-wide mean, which may be due to better education or other factors. The differential change in skill intensity across industries, where trade may play a bigger role, plays little role in the patterns from Table A20. In ongoing work we are exploring more realistic assumptions that would enable us to predict the marginal worker displaced by import competition.

E.4 Combining Estimates on Differential Spending on Imports

Sections 3 and 4–5 estimated the differences in spending on imports across and within industries, respectively. In this section, we combine those estimates using the within-between decomposition (A24). Because our microdata cover only a subset of industries, this requires extrapolation of the patterns within consumer packaged goods and automobiles into other traded industries. We assume that covered industries are representative in terms of the relative difference in import spending.

Formally, we rewrite (A24) in terms of differences as fractions of the average spending on imports:

$$\frac{\Delta_{\text{Final}}[IP_{\mathbf{c}}]}{\mathbb{E}_{\text{Final}}[IP_{\mathbf{c}}]} = \frac{\Delta_{\text{Final}}^{\text{between}}[IP_{\mathbf{c}}]}{\mathbb{E}_{\text{Final}}[IP_{\mathbf{c}}]} + \mathbb{E}_{\text{Imports}}[\omega_g \cdot \text{Rel diff}_{g\mathbf{c}}],$$

where $\mathbb{E}_{\text{Imports}}[\cdot]$ is the average across sectors weighted by spending on imports and $\text{Rel diff}_{g\mathbf{c}} = \frac{\Delta_{\text{Final}}[IP_{\mathbf{c}}|g]}{\mathbb{E}_{\text{Final}}[IP_{\mathbf{c}}|g]}$ is the difference in import spending between the two consumer groups within group of industries g , as a fraction of the average. The adjustment term is $\omega_g = \frac{\mu_g(1-\mu_g)}{\bar{\mu}(1-\bar{\mu})} \approx 1$.¹²³

The between-term equals to -4.79%, according to Table 1. Tables 3 and 4 estimate that Rel diff_g is +4.34% in the Nielsen data and +13.25% for automobiles, respectively.¹²⁴ Averaging those weighted by the total spending on imports in those categories, we get +5.71%.¹²⁵ Therefore, our final estimate of the differential spending on imports is +0.92% of the average.

When looking at imports from China, automobiles do not play a significant role. Combining the across component +4.48% from Table 1 with the within component of -2.44% from Table 3, we get the final difference of +2.05% of the average.

¹²²Indeed, we estimated the counterfactual effects of a bilateral trade liberalization with all trading partners, as in Section 7, accounting for the skill-bias of importing and found that the pro-skilled earnings channel increases from 14.3% of the average gains to 26.5% using the 2000 skill intensity and to 47.2% using the 1990 data.

¹²³To prove this representation, we use (A26). Applying the law of total covariance to its numerator, we isolate the within component: $\Delta_{\text{Final}}^{\text{within}}[IP_{\mathbf{c}}] = \mathbb{E}_{\text{Final}}\left[\frac{\text{Cov}[\mu_j, IP_{j\mathbf{c}}|g]}{\bar{\mu}(1-\bar{\mu})}\right] = \mathbb{E}_{\text{Final}}\left[\omega_g \cdot \frac{\text{Cov}[\mu_j, IP_{j\mathbf{c}}|g]}{\mu_g(1-\mu_g)}\right] = \mathbb{E}_{\text{Final}}[\omega_g \cdot \Delta_{\text{Final}}[IP_{\mathbf{c}}|g]] = \mathbb{E}_{\text{Final}}[IP_{\mathbf{c}}] \cdot \mathbb{E}_{\text{Imports}}[\omega_g \cdot \text{Rel diff}_{g\mathbf{c}}]$.

¹²⁴For Nielsen, we only use the within-industry component of the difference to avoid double-counting (see Section 4.2).

¹²⁵Nielsen data should cover all food and beverages, but its coverage in other industries, e.g. chemicals, is only partial. Therefore, we reduce the weight of all non-food Nielsen categories by a factor of two.

Additional Tables and Figures

Table A1: Notation for Model without Input-Output Linkages

Indices	$i \in \{S, U\}$ j, r c	Group of agents: S (skilled) or U (unskilled) Industry; Sector (goods or services) Country: H (Home), F (Foreign), $\mathbf{c}$ (affected by import shock)
Prices, Quantities, Transaction Values	p_{jc} p_j Q_{jc}^i, X_{jc}^i Q_j^i, X_j^i $Q_{jH}, Q_{jH}^{\text{Export}}$ X_j X_i X_{jH}, VA_j	Consumer price of country- c variety in industry j Consumer price index in industry j Quantity and expenditure of group i on variety jc Composite quantity and expenditure of group i in industry j Domestic output and exports in j (quantity) Total domestic expenditure on industry j Total domestic expenditure by group i Total value of output and value added by domestic industry*
Wages, Employment	$w_i, \bar{w}$ L_i, L_i^j	Wage for group i , average wage Measure of group i consumers; industry employment
Equilibrium Shares	$s_j^i, s_j^{\text{Final}}$ e_i^j $\mu_j, \bar{\mu}$ $v_j, \bar{v}$ $IP_{jc}, IP_{j\mathbf{c}},$ IP_j	Share on industry j in expenditure of group i , both groups Share of industry j in group i earnings and employment Share of skilled consumers in final sales, by industry and overall Share of skilled workers in payroll, by industry and overall Import penetration: share of country c , set of countries $\mathbf{c}$, or all foreign countries in domestic expenditure in industry j
Counterfactual Changes	Hats $\hat{\tau}$ $\hat{\tau}^*$ $\hat{U}_i$ $\hat{\pi}_i, \hat{\pi}$	Relative change from original to counterfactual equilibrium Counterfactual growth of import barriers (with countries $\mathbf{c}$) Counterfactual growth of export barriers (with all partners) Money metric of welfare growth Laspeyres price index for group i and both groups together
Elasticities	$\xi_j, \varepsilon_r, \rho$ ψ_j $\sigma_j, \sigma_{\text{macro}}$ η_j^{import} η_j^{export} $\eta_j^{\text{avg wage}}$	Substitution between country varieties in j ; between industries within sector r ; between goods and services Income elasticity of domestic varieties in j Elasticity of substitution between labor types in domestic production in j and at macro level Negative elasticity of industry VA w.r.t. import tariff Elasticity of industry VA w.r.t. export tariff Elasticity of industry VA w.r.t. domestic average wage
Averages and Differences Across Industries	$\mathbb{E}_{\text{Final}} [\cdot]$ $\Delta_{\text{Final}} [\cdot]$ $\mathbb{E}_{\text{VA}} [\cdot]$ $\Delta_{\text{VA}} [\cdot]$	Average weighted by s_j^{Final} (domestic final expenditure) Difference between averages weighted by s_j^S and s_j^U Average weighted by domestic value added Difference between averages weighted by e_S^j and e_U^j

*Equal to each other in the model without input-output linkages.

Notes: This table lists the notation for the model in Section 2.1. “With respect to” is abbreviated to “w.r.t.”

Table A2: Classification of Subsectors

Goods	Services
Apparel and leather and allied products	Accommodation and food services
Chemical products	Arts, entertainment, and recreation
Computer and electronic products	Construction*
Electrical equipment, appliances, and components	Educational services
Fabricated metal products	Finance and insurance
Farms	Government
Food and beverage and tobacco products	Health care and social assistance
Forestry, fishing, and related activities*	Information
Furniture and related products	Other services, except government
Machinery	Professional, scientific, and technical services
Mining, except oil and gas*	Real Estate, rental and leasing
Miscellaneous manufacturing	Retail trade*
Motor vehicles, bodies and trailers, and parts	Transportation and warehousing
Nonmetallic mineral products	Utilities
Oil and gas extraction*	Wholesale trade*
Other transportation equipment	
Paper products	
Petroleum and coal products	
Plastics and rubber products	
Primary metals*	
Printing and related support activities*	
Support activities for mining*	
Textile mills and textile product mills	
Wood products*	

* Subsectors with zero final consumption (either in the input-output table or in the CEX, or both).

Notes: This table lists subsectors within the goods-producing and service sectors according to the detailed 2007 BEA input-output table. Goods-producing services include agriculture, manufacturing, and mining. Subsectors are defined by the 3-digit input-output codes for goods and 2-digit NAICS codes for services (except Management and Administrative Services, which are included in the Professional, Scientific, and Technical Services).

Table A3: Differential Import Spending by Consumer Education: The Role of Goods and Services

	Share of Services in Total Spending	Share of Imports in Spending by Sector			
		From All Countries		From China	
		Goods	Services	Goods	Services
	(1)	(2)	(3)	(4)	(5)
All consumers, %	79.70	45.96	5.47	8.48	0.42
College consumers, %	81.87	48.15	5.58	9.60	0.44
Non-college consumers, %	78.24	44.72	5.40	7.80	0.41
College minus non-college, p.p.	+3.63	+3.43	+0.18	+1.90	+0.02

Notes: This table provides evidence on the role of goods and services in explaining the differential spending share on imports between consumer education groups. Using industry-level data from Section 6, it shows that services constitute a large share of spending for college graduates (column (1)) and have lower share of total imports (including imported intermediate inputs and measured as % of absorption). Yet, within goods and to a smaller extent within services, the share of spending on imports is larger for college graduates (columns (2)–(5)).

Table A4: Nielsen-Census Matching Rules

	Non-Missing Exact Match	Exact and [Fuzzy] Match
Rule 1	Zip-9	House, Name, Address, PO Box, Unit, Bldg
Rule 2		House; [Name, Address, PO Box, Unit, Bldg]
Rule 3	Zip-5, House	Name, Address, PO Box, Unit, Bldg
Rule 4		[Name, Address, PO Box, Unit, Bldg]
Rule 5	Zip-5	Name
Rule 6	City	Name, State
Rule 7	State	Name, Entity

Notes: This table lists the rules used to match names and addresses in the Nielsen and Census samples. Each rule requires an exact match and non-missing values of the variables listed in the first column, as well as an exact or probabilistic (fuzzy) match on the variables from the second columns (missing values are allowed). Variables where fuzzy match is allowed are listed in brackets. For fuzzy matching, a 75% threshold is chosen for the match quality score assigned by the `reclink2` package from [Wasi and Flaaen \(2015\)](#).

Table A5: Nielsen-Census Match Statistics

(a) Nielsen Firms				
	2007		2012	
	Firms	% of Sales	Firms	% of Sales
All Nielsen	26,900	100.00	28,600	100.00
Nielsen with size filter	11,000	99.77	12,100	99.82
Matched to SSEL, same year	7,600	83.19	8,900	87.29
Matched to SSEL, any year	8,200	90.76	9,300	91.86
Matched to Economic Census	7,200	88.68	7,800	88.57
Passed consistency filter	6,100	83.02	6,600	83.61

(b) Census Firms in Food, Alcohol, and Tobacco			
	All years		
	Firms	% of Sales	
All Census	51,500	100.00	
Matched to Nielsen	8,900	78.96	
Matched to Nielsen with size filter	5,200	75.57	
Passed Consistency Filter	4,800	58.73	

Notes: This table reports the number of firms and percentage of total sales remaining after each step of the merging process between the Nielsen and Census samples, explained in detail in Appendix C.3. Panel (a) measures these statistics relative to the full Nielsen sample (for 2007 and 2012 Economic Censuses separately), while Panel (b) measures them relatively to the set of Census firms active in the Food, Alcohol, and Tobacco Manufacturing industries (NAICS codes 311 and 312). The last line of panel corresponds to the final merged sample, for all firms in Panel (a) and for those in food, alcohol, and tobacco in Panel (b). The numbers of firms are rounded to the nearest 100 to preserve confidentiality.

Table A6: Distribution of Match Types, Merged Nielsen-Census Sample

	% of Matched Firms (1)	% of Sales (2)	% of $\sqrt{\cdot}$ Sales (3)
<i><u>Multi-establishment firms</u></i>			
Rule 1	10.30	19.72	17.88
Rule 2	4.12	18.99	10.76
Rule 3	5.21	19.86	12.77
Rule 4	3.87	18.54	9.82
Rule 5	2.54	4.12	4.46
Rule 6	1.72	6.80	4.75
Rule 7	1.65	5.11	4.34
Total multi-establishment	29.42	93.14	64.79
<i><u>Single-Establishment Firms</u></i>			
Rule 1	33.87	3.42	17.23
Rule 2	10.27	0.87	4.87
Rules 3–7	26.44	2.57	13.12
Total single-establishment	70.58	6.86	35.21

Notes: This table shows the fractions of the Nielsen-Census merged sample corresponding to each of the merging rules, described in Appendix C.3. Column (1) shows the raw fraction of Nielsen firms in each category, while column (2) shows the share of total Nielsen sales, and column (3) weights firms by the square-root of Nielsen sales.

Table A7: Distribution of NAICS Industries, Merged Nielsen-Census Sample

NAICS Industry		% of Firms	% of Sales	% of $\sqrt{\text{Sales}}$	% of Private Label Brands
Code	Description	(1)	(2)	(3)	(4)
<i><u>2-digit NAICS codes</u></i>					
31-33	Manufacturing	49.78	61.63	57.17	1.21
42	Wholesale	39.37	16.02	29.00	7.90
44-45	Retail	4.80	18.55	8.66	93.74
—	Other	6.04	3.80	5.18	5.19
<i><u>3-digit NAICS codes</u></i>					
311	Food Manufacturing	31.16	36.74	34.78	0.73
312	Beverage and Tobacco Manufacturing	5.73	6.68	6.26	0.30
322	Paper Manufacturing	0.75	4.76	1.96	1.86
325	Chemical Manufacturing	5.36	8.18	6.97	2.79
423	Durable Goods Wholesalers	8.34	2.20	5.86	5.91
424	Nondurable Goods Wholesalers	29.96	15.24	23.05	6.49
445	Food and Beverage Stores	2.24	9.82	4.97	99.10
—	Other	16.44	16.38	16.16	49.69

Notes: Columns (1)–(3) of this table report the fractions of the Nielsen-Census merged sample corresponding to selected 2- and 3-digit NAICS sectors. Each firm in the Economic Census is classified into the sector where its establishments have the highest total payroll. Column (1) shows the raw fraction of firms in each sector, while column (2) shows the share of total Nielsen sales, and column (3) weights firms by the square-root of Nielsen sales. Column (4) measures, for firms in each sector, the sales share of Nielsen barcodes that are classified as private label brands—brands that belong to the retail store. We identify them in the Nielsen data as those which contain “CTL BR” in the barcode description.

Table A8: Nielsen-Census Sample Selection

(a) Nielsen Firms						
	<i>N</i>	% of Total Sales	Median Sales, \$k	% of Sales to College Grads	Mean HH Income, \$k	
Matched	12,700	83.50	1,904	29.14	67.63	
Didn't Match	10,400	16.50	981	30.71	69.75	
P-value of t-test				[0.009]	[0.008]	
P-value controlling for size				[0.425]	[0.028]	

(b) Census Firms in Food, Alcohol, and Tobacco						
	<i>N</i>	% of Sales	Median Sales, \$k	Median Payroll, \$k	Median Employment	Mean Skill Intensity
Matched	4,800	58.73	13,303	1,889	54	0.336
Didn't Match	46,600	41.27	606	113	4	0.341
P-value of t-test						[0.744]

Notes: This table compares firms in the matched Nielsen-Census sample to other firms in Nielsen (Panel (a)) and in the Economic Census (Panel (b)) which did not find a match, in terms of size, consumer, and producer characteristics. The universe of firms in Panel (a) is all Nielsen firms that passed the size filter, while in Panel (b) it is all firms in the Economic Census active in Food, Alcohol, and Tobacco Manufacturing. P-values for t-tests for equality of means between the matched and unmatched samples are shown in brackets. The last row of Panel (a) performs such t-test controlling for a quadratic polynomial in log firm sales. The numbers of firms are rounded to the nearest 100 and medians are computed as geometric means of the 45 and 55 percentiles to protect confidentiality.

Table A9: Spending on Imports by Education Group and Product Class,
Merged Nielsen-Census Sample

	Imports Excluding China by Product Class			Imports from China by Product Class		
	Food	Health & Household	General Merchandize	Food	Health & Household	General Merchandize
	(1)	(2)	(3)	(4)	(5)	(6)
All, %	6.04	8.07	10.04	0.88	6.51	17.91
College, %	6.75	8.16	10.41	0.85	6.27	17.65
Non-college, %	5.72	8.03	9.88	0.90	6.62	18.03
College minus non-college, p.p.	+1.03	+0.13	+0.54	-0.04	-0.35	-0.38
	(0.12)	(0.17)	(0.22)	(0.03)	(0.14)	(0.28)
<i>as % of avg. import spending</i>	<i>17.06</i>	<i>1.61</i>	<i>5.37</i>	<i>-4.98</i>	<i>-5.38</i>	<i>-2.13</i>
→ Within industries	+0.73	+0.21	+0.36	+0.01	-0.33	-0.30
	(0.113)	(0.11)	(0.17)	(0.023)	(0.18)	(0.13)
<i>as % of avg. import spending</i>	<i>12.14</i>	<i>2.55</i>	<i>3.62</i>	<i>0.57</i>	<i>-5.12</i>	<i>-1.67</i>
→ Within product modules	+0.498	+0.13	+0.39	-0.00	-0.43	-0.32
	(0.080)	(0.09)	(0.15)	(0.015)	(0.08)	(0.18)
<i>as % of avg. import spending</i>	<i>8.25</i>	<i>1.56</i>	<i>3.88</i>	<i>-0.11</i>	<i>-6.59</i>	<i>-1.76</i>
<i>N</i> firm-years	9,000	3,700	2,800	9,000	3,700	2,800

Notes: This table reports the fraction of imports in expenditure for different education groups using the merged Nielsen-Census sample from Section 4. Importing is proxied by the share of total imports in firm sales. Differential spending on imports is decomposed into “within” and “between” components for 6-digit I-O codes (“industries”) and for Nielsen product modules (“product modules”) according to equation (A24). The same firm may operate in more than one product class. Firms are weighted by the square-root of Nielsen sales. Standard errors are shown in parentheses.

Table A10: Consumer Spending on Imports by Firm Activity: Manufacturing, Wholesale, and Retail
(Merged Nielsen-Census Sample)

	Total imports, All products			Imports from China, Health & Household		
	MFG (1)	WH (2)	RT (3)	MFG (4)	WH (5)	RT (6)
All, %	4.37	5.82	0.30	1.98	3.99	0.28
College minus non-college, p.p.	-0.09 (0.07)	0.62 (0.10)	-0.01 (0.02)	-0.11 (0.08)	-0.21 (0.11)	-0.01 (0.02)
→ Within industries	-0.05 (0.05)	0.47 (0.11)	-0.00 (0.01)	-0.10 (0.04)	-0.19 (0.16)	-0.01 (0.02)
<i>N</i> firm-years	12,700	12,700	12,700	3,700	3,700	3,700

Notes: This table estimates the average and differential fraction of imports in spending, decomposed by the main activity of the firm that registered the product: manufacturing (MFG), wholesale (WH), or retail (RT). Other activities are not shown. Each firm is assigned the main activity based on the total payroll of establishments in the corresponding NAICS sectors. Each block of three columns is based on the same data: we decompose import spending into components, without amending the sample. Firms are weighted by the square-root of Nielsen sales. Standard errors are shown in parentheses.

Table A11: Summary Statistics, % of Firms' Sales
(Merged Nielsen-Census Sample, Full Weights)

	All products	Food	Health & Household	General Merchandise
Total Imports	8.16	5.15	11.23	29.53
Imports from China	2.25	0.65	2.69	17.17
Imports from NAFTA	2.05	1.62	2.92	3.81
Imports from Developed Economies	2.41	1.54	4.13	5.98
% of Firm-Module Sales to College Graduates (st.dev.)	29.14 (8.29)	28.59 (8.62)	30.17 (7.37)	31.64 (6.58)
% of Product Class in Total Sales	100.00	71.73	21.21	7.06
<i>N</i> firms	8,200	5,700	2,400	2,000
<i>N</i> firm-years	12,700	9,000	3,700	2,800
<i>N</i> firm-module-years	131,000	88,600	29,800	12,500

Notes: This table is analogous to Table 2, except using Nielsen sales instead of the square-root of sales as weights. It reports statistics on imports based on the merged Nielsen-Census sample from Section 4, for all products and for three product classes: Food, Alcohol, and Tobacco (“Food”), Health and Beauty Products and Household Supplies (“Health and household”), and General Merchandise. Imports are measured at the firm level. Observations are firm-module-year cells and the numbers of observations are rounded to the nearest 100 to preserve confidentiality.

Table A12: Spending on Imports by Education Group,
Merged Nielsen-Census Sample, Full Weights

	All products			Food	Health & Household		General Merchandize	
	All	China	Excl. China		China	Excl. China	China	Excl. China
	(1)	(2)	(3)		(5)	(6)	(7)	(8)
All, %	8.16	2.25	5.91	4.50	2.69	8.55	17.17	12.36
College, %	8.58	2.34	6.24	4.83	2.65	8.56	16.81	12.57
Non-college, %	7.99	2.21	5.78	4.37	2.70	8.54	17.34	12.26
College minus non-college, p.p.	+0.59	+0.13	+0.46	+0.46	-0.05	+0.02	-0.53	+0.31
	(0.15)	(0.06)	(0.13)	(0.13)	(0.05)	(0.25)	(0.43)	(0.28)
<i>as % of avg. import spending</i>	<i>+7.20</i>	<i>+5.60</i>	<i>+7.81</i>	<i>+10.16</i>	<i>-1.97</i>	<i>+0.25</i>	<i>-3.10</i>	<i>+2.53</i>
→ Within industries	+0.18	-0.02	+0.20	+0.26	-0.07	+0.04	-0.28	+0.07
	(0.07)	(0.03)	(0.06)	(0.05)	(0.05)	(0.18)	(0.27)	(0.21)
<i>as % of avg. import spending</i>	<i>+2.14</i>	<i>-1.02</i>	<i>+3.35</i>	<i>+5.82</i>	<i>-2.46</i>	<i>+0.41</i>	<i>-1.64</i>	<i>+0.57</i>
N firm-years	12,700	12,700	12,700	9,000	3,700	3,700	2,800	2,800

Notes: This table is analogous to Tables 3 and A9, except using Nielsen sales instead of the square-root of sales as weights. It reports the fraction of imports in expenditure for different education groups using the merged Nielsen-Census sample from Section 4. Column (1) measures all imports; columns (2)–(3) decompose imports into those from China and other countries (“Excluding China”), and the following columns report import shares separately by product class. Imports of food from China are small and not shown. Standard errors are shown in parentheses.

Table A13: Summary Statistics by Firm Activity, Merged Nielsen-Census Sample

	Firm Activity			
	MFG	WH	RT	Other
Total Imports, % of Firms’ Sales	7.37	16.46	2.23	14.98
Imports from China, % of Firms’ Sales	1.38	5.67	1.34	6.33
% of Firm-Module Sales to College Graduates	28.98	29.67	28.55	32.35
% of Firm Group in Total Sales	61.63	16.02	18.55	3.80
% of Firm Group in Total $\sqrt{\text{Sales}}$	57.17	29.00	8.66	5.18
N firm-years	6,300	5,000	600	800

Notes: This table reports statistics on imports and consumer education by the main activity of the firm in the merged Nielsen-Census sample: manufacturing (MFG), wholesale (WH), retail (RT), and Other. Each firm is assigned the main activity based on the total payroll of establishments in the corresponding NAICS sectors. Summary statistics are computed using the square-root of firms’ Nielsen sales as weights. Numbers of observations are rounded to the nearest 100 to preserve confidentiality.

Table A14: Bias Correction for the Consumption Segmentation Index, Merged Nielsen-Census Sample

	All products (1)	Food (2)	Health & Household (3)	General Merchandise (4)
<i>Naïve estimates of consumption segmentation index, %</i>				
Across barcodes	15.03	15.09	14.10	16.13
Across firm-modules	7.42	8.52	5.74	4.16
<i>Bias-corrected estimates, %</i>				
Across barcodes	8.22	9.21	6.58	5.46
Across firm-modules	5.77	6.87	4.09	2.52
Attenuation ratio	1.424	1.339	1.610	2.165

Notes: This table estimates the consumption segmentation index defined by (A23), first using a naïve plug-in estimator and then with the bias-correction procedure described in Appendix B.3. The ratio of bias-corrected estimates of segmentation across barcodes and firm-modules, presented in the last row, bounds the attenuation bias in the differential import spending, as explained in Appendix B.2. Square-root weighting scheme is used throughout. The uncorrected consumption segmentation index for firm-modules is different from the one reported in Section 4.2 because this table does not subtract the within-industry component.

Table A15: Automobile Brands

Brand Code	Brand	N	Consumer Characteristics		Imports as % of Sales	
			% of Sales to College Graduates	Average Percentile of Consumer Income	All Imports	Outside NAFTA
TOY	Toyota	6,964	44.7	56.5	44.2	25.4
HON	Honda	5,961	44.6	58.2	35.3	9.9
CHE	Chevrolet	5,489	22.4	48.3	36.0	4.8
FOR	Ford	5,022	27.5	49.8	45.0	0.0
NIS	Nissan	3,253	33.4	53.9	43.2	3.8
PON	Pontiac	2,070	23.1	48.6	17.5	17.5
BUI	Buick	1,958	23.0	41.7	19.0	5.9
HYU	Hyundai	1,747	38.5	55.3	46.0	46.0
DOD	Dodge	1,695	21.6	47.3	52.5	0.0
CHR	Chrysler	1,689	26.6	50.7	33.2	0.0
SUB	Subaru	1,635	57.9	63.1	61.4	61.4
VOK	Volkswagen	1,597	47.0	63.3	77.8	17.4
MAZ	Mazda	1,219	45.3	57.5	87.3	80.6
BMW	BMW	1,086	58.7	70.9	100.0	100.0
MEC	Mercury	1,044	24.1	43.4	94.6	0.0
MRB	Mercedes-Benz	1,002	57.4	65.4	91.3	91.3
SAT	Saturn	917	30.9	49.6	19.8	19.8
LEX	Lexus	905	58.6	67.2	99.8	99.8
CAD	Cadillac	903	27.6	49.3	17.9	0.0
ACU	Acura	840	52.0	64.5	34.7	34.7
KIA	KIA	793	27.6	49.8	73.3	73.3
MIT	Mitsubishi	648	25.5	52.3	70.4	70.4
LIN	Lincoln	613	28.7	44.9	73.7	0.0
VOV	Volvo	568	54.6	64.1	100.0	100.0
INF	Infiniti	390	56.9	66.4	100.0	100.0
AUD	Audi	390	62.6	69.8	100.0	100.0
SAA	Saab	197	61.4	68.8	100.0	100.0
MIN	Mini	174	63.2	73.6	100.0	100.0
SCI	Scion	170	37.1	58.9	100.0	98.1
SUZ	Suzuki	159	26.4	49.2	100.0	100.0
POR	Porsche	147	67.3	73.4	100.0	100.0
JAG	Jaguar	140	43.6	61.9	100.0	100.0
DAW	Daewoo	32	15.6	48.3	100.0	100.0
FIA	Fiat	25	52.0	55.0	92.5	14.9
SMA	Smart	20	20.0	54.5	100.0	100.0
MGA	MG	16	68.8	62.9	100.0	100.0
TES	Tesla	10	70.0	71.0	0.0	0.0
LAN	Land Rover	6	33.3	47.4	100.0	100.0
AUS	Austin-Healey	4	50.0	33.3	100.0	100.0
Total		51,498	36.9	54.4	49.7	25.5

Notes: This table lists 39 brands of cars in the sample on auto purchases, described in Section 5.1. For each brand, it reports the total number of purchases and average consumer characteristics in the CEX sample, as well as the fraction of imported assembled cars from Ward's Automotive reports. Light trucks (including SUVs) are not included.

Table A16: Spending Shares and Average Income Elasticities in the CEX

Category of spending	Share of total spending, %	Average income elasticity
Housing*	25.17	1.060
Transportation	19.29	1.093
Utilities	8.62	0.721
Health and insurance	7.13	0.775
Food at home	6.83	0.427
Entertainment, reading	6.17	1.262
Food away from home	5.37	1.138
Cash contributions	4.10	1.374
Furnishings and equipment	3.63	1.201
Apparel	2.43	1.019
Domestic services and childcare	2.21	1.416
Education	2.09	1.382
Miscellaneous	1.54	1.099
Housekeeping supplies	1.28	0.719
Personal care	1.26	1.004
Shoes and other apparel	1.22	1.003
Alcohol	0.94	1.143
Smoking	0.72	0.078

* Unlike the I-O table, the Housing category in the CEX include maintenance and repairs.

Notes: This table shows spending shares and average income elasticities by group of CEX spending categories, weighted by total spending. Grouping is based on the Integrated Stub file provided by the CEX (see <https://www.bls.gov/ce/pumd.doc.htm>). Income elasticities are estimated using the methodology described in Section B.4.

Table A17: Exposure to Import Competition by Education Group and Trading Partner, Industry Data

	Payroll-weighted Import Penetration by Trading Partner (even columns adjust for input-output linkages)					
	China		NAFTA		Developed Economies	
	(1)	(2)	(3)	(4)	(5)	(6)
All workers, %	0.65	1.19	0.85	1.87	1.41	2.87
College-educated workers, %	0.59	1.08	0.69	1.61	1.37	2.75
Non-college educated workers, %	0.72	1.29	1.00	2.13	1.45	2.98
College minus non-college, p.p.	-0.13	-0.21	-0.31	-0.52	-0.08	-0.23
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
<i>as % of avg.</i>	-19.48	-17.50	-37.13	-27.81	-5.85	-8.15
→ Between goods and services	-0.30	-0.36	-0.39	-0.52	-0.64	-0.81
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
→ Within goods and services	+0.17	+0.16	+0.07	-0.00	+0.56	+0.58
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
→ Between subsectors	+0.17	+0.19	+0.08	+0.04	+0.41	+0.47
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
→ Within subsectors	-0.00	-0.03	-0.00	-0.05	+0.15	+0.11
	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)

Notes: This table reports average import penetration statistics weighted by total payroll and payroll of college- and non-college educated workers separately, using the industry-level data from Section 6, which covers 380 industries. It also decomposes the difference between education groups into the within and between components for sectors (goods and services) and subsectors (listed in Table A2), according to equation (A24). Import penetration is measured as imports from a given set of countries as % of industry absorption. Even columns account for import penetration in downstream industries (see Section 2.3 for details).

Table A18: Differential Import Penetration by Worker Education: the Role of Goods and Services

	Share of Services	Import Penetration by Sector	
	in Total Payroll	Goods	Services
	(1)	(2)	(3)
All workers, %	85.32	23.88	6.56
College workers, %	88.68	28.09	7.17
Non-college workers, %	81.99	21.27	5.92
College minus non-college, p.p.	+6.69	+6.81	+1.25

Notes: This table provides evidence on the role of goods and services in explaining the differential exposure to import competition between worker education groups. Using industry-level data from Section 6, it shows that services constitute a large share of payroll for college graduates (column (1)) and have lower import penetration (direct imports as % of absorption). Yet, within goods and within services, the exposure to import competition, measured as the average import penetration weighted by group-specific payroll shares, is larger for college graduates (columns (2)–(3)).

Table A19: Skill-Bias of Exporters in Census Microdata

	Measure of skill intensity: payroll share of			
	College graduates	Non-production workers		
	MOPS 2010 (1)	CMF 2007 (2)	ASM 2010 (3)	MOPS 2010 (4)
Average export share, %	22.84	14.70	19.47	22.84
Differential export share, skilled minus unskilled, p.p.:				
Overall	+5.26	+4.50	+4.52	+5.36
→Between industries	+4.49	+4.09	+4.51	+5.20
→Within industries	+0.77	+0.41	+0.01	+0.16
<i>N</i> establishments	33,400	294,200	50,500	33,400

Notes: This table shows the average export shares (exports as % of sales) for three samples of manufacturing establishments: the 2010 MOPS (columns (1) and (4)), the 2007 Census of Manufactures (column (2)) and the 2010 Annual Survey of Manufactures (column (3)). It also shows the differential exposure for skilled and unskilled workers and decomposes it into “between” and “within” components across six-digit industries, according to equation (A24). Skilled workers are defined as college graduates in column (1) and non-production workers in the other columns. Establishments are weighted by their total payroll. See Section E.3 for details of data construction.

Table A20: Differential Exposure to Skill-Biased Import Competition

	Marginal Displaced Worker is Representative of Skill Intensity in					
	2007		2000		1990	
	(1)	(2)	(3)	(4)	(5)	(6)
All workers, %	4.21	8.15	4.21	8.15	4.21	8.15
College-educated workers, %	4.13	7.85	3.67	7.39	2.95	6.67
Non-college educated workers, %	4.28	8.45	4.74	8.90	5.45	9.62
College minus non-college, p.p.	-0.15	-0.60	-1.06	-1.51	-2.50	-2.95
<i>as % of avg.</i>	<i>-3.54</i>	<i>-7.33</i>	<i>-25.25</i>	<i>-18.53</i>	<i>-59.37</i>	<i>-36.14</i>
Average College Payroll Share	49.70		46.09		41.43	


Notes: This table reports the payroll-weighted exposure to import competition by education group assuming that import competition happens in segments of each industry that have the skill mix as in the U.S. in the past. Columns (1) and (2) repeat the results from Table 6 without the skill bias of importing. Columns (3)–(4) use the skill intensity from 2000 as a proxy for the mix of marginally affected workers, and columns (5)–(6) use year 1990. Columns (2), (4), and (6) adjust for input-output linkages. The methodology is described in Appendix E.3.

Table A21: Decomposition of Import Competition Effects of a Bilateral Trade Liberalization
(as % of average welfare gains)

	Total import competition effects (1)	Final demand reallocation within IO6 (2)	Final demand reallocation across goods/services (3)	Intermediate demand reallocation within IO6 (4)
College minus non-college p.p.	7.04	0.51	2.10	4.43
→ Between goods and services	17.18	6.69	2.85	7.63
→ Within goods and services	-10.13	-6.18	-0.75	-3.20
→ Between subsectors	-7.79	-4.21	-0.54	-3.04
→ Within subsectors	-2.34	-1.97	-0.21	-0.16


Notes: This table presents additional evidence on the import competition effects of a bilateral trade liberalization from Panel (c) of Figure 7. Each of the effects is decomposed into the within- and between- components according to equation (A24).

Figure A1: U.S. Net Imports as % of GDP over Time




Notes: This figure shows the evolution of the ratio of net imports to GDP in the U.S., using BEA Summary I-O Tables for 1997–2015. Net imports are computed as total imports minus total exports. The vertical line corresponds to the year of our main analysis, 2007.

Figure A2: Imports from China and Consumer Base across Subsectors




Notes: This figure shows the relationship between consumer base (% of industry sales to college graduates) and the share of total (direct plus indirect) imports from China in final expenditure using industry-level data from Section 3. Each circle corresponds to a subsector within the goods-producing sector (listed in Table A2), and the circle size indicates final spending. Subsectors that take up less than 3% of the sectoral spending are not shown. Only the composite of the service sector is shown because direct imports of services from China are zero in our data by construction and indirect import penetration is very small.

Figure A3: Import Spending by Consumer Income Bin, Industry-Level Data




Notes: These binned scatterplots group CEX panelists into 11 bins by household income before tax. They report the average share of total (direct and indirect) imports in the spending of each bin, computed using the industry-level data from Section 3. Panel (a) accounts for all imports (including services), whereas Panel (b) measures only imports of goods from developed economies, NAFTA countries, and China.

Figure A4: Import Spending by Consumer Income Bin, Merged Nielsen-Census Sample




Notes: These binned scatterplots group Nielsen panelists into 15 bins by household income. They report the average share of imports in the spending of each bin, computed using the merged Nielsen-Census sample from Section 4. Panel (a) accounts for imports from countries other than China. The other panels measure imports from China: for Health and Household products (Panel (b)), General Merchandize (Panel (c)), and overall (Panel (d)). The upward slope in Panel (d) is a consequence of compositional differences only: rich people spend more on general merchandize and less on food (see Appendix E.1).

Figure A5: Variation in Differential Import Spending Across Product Modules, Merged Nielsen-Census Sample



Notes: These binned scatterplots group Nielsen product modules into bins by the average import share, computed using the merged Nielsen-Census sample. For each bin, it reports the differential import spending between college and non-college consumers as percentage of the average import spending. Panel (a) uses all modules and looks at imports from all countries, while Panel (b) focuses on imports from China within the Health and Household product class, where differences are most significant (see Table A9). The table shows that difference in import spending are weaker in modules with more imports.

Figure A6: Fraction of Imported Cars by Household Income Bins



Notes: These binned scatterplots split car purchases in the CEX into equally-sized bins by the percentile of the owner's household income (among all surveyed households) in the year of the survey. Each car in the data is assigned a probability of being imported (overall or from Canada and Mexico specifically) based on the average import share of the car brand in the Ward's data.

Figure A7: Import Penetration by Country and Skill Intensity






Figure A8: Skill Bias of Exporters in the 2010 MOPS Survey



Notes: This figure shows that the within-industry association between exporting and skill intensity does not create large biases for the differential exposure of college- and non-college-graduated workers to exporting opportunities. Based on column (1) of Table A19, it groups manufacturing establishments in the 2010 MOPS survey into bins by skill intensity (fraction of payroll to college graduates). Circles show the average export share (exports as % of sales) of this establishment, whereas crosses use the six-digit I-O industry average export share instead. Each establishment is weighted by its payroll. The lines are similar when within-industry patterns are minor relative to the across-industry ones. See Section E.3 for details of data construction.


Figure A9: Import Penetration and Skill Intensity over Time




Notes: These binned scatterplots group 6-digit manufacturing NAICS industries into bins by the measure of skill intensity available in the NBER CES dataset (payroll share of *non-production* workers). It reports import penetration measured as total imports divided by the sum of imports and domestic output. Industries are weighted by payroll ($N = 462$ in 1992 and 465 in the 1999 and 2007). Payroll and skill intensity are from the NBER CES database, and imports at the NAICS level are from [Schott \(2008\)](#) and [Pierce and Schott \(2012\)](#). Several industries in the NBER CES have been aggregated to match those in the Pierce-Schott data (see Appendix [E.3](#) for details of data construction).

Figure A10: Decomposition of Average Wage Effects of Trade Liberalizations

(a) Direct Effects of a bilateral 10% liberalization on the average wage
(excluding the Feedback Effect)



(b) Decomposition of the Feedback Effect
(Multiplier = $1 / (1 - 0.598) = 2.49$)



Notes: This figure uses equation (13a) to measure the contribution of different mechanisms to the average domestic wage change (relative to a foreign numeraire) caused by a 10% bilateral reduction of tariffs. Panel (a) measures the terms in parentheses in (13a) and Panel (b) estimates the multiplier. See Section 2.2 for the intuition behind each mechanism.

Figure A11: Examples of Products

Domestic Products

(a) Plates “Corelle”



UPC 071160 015449
World Kitchen, LLC

(b) Plates “MainStays”



UPC 018643 157371
Merrick Engineering, Inc.

Imported Products

(c) Bed Sheets “MainStays”,
Made in China



UPC 844178 030335
Jiangsu Royal Home USA, Inc.

(d) Plates “Better Homes”,
Made in China



UPC 855602006 567
First Design Global, Inc.

(e) Conditioner “Equate Beauty”,
Made in Canada



UPC 681131 124836
Wal-Mart Stores, Inc.

Notes: These products were photographed in a Wal-Mart store on September 16, 2017. Each barcode (UPC) is split by a space into the firm prefix in the GS1 database and the part which identifies the product within a firm. The country of origin (U.S., China, Canada) is from the product label, whereas the firm information is from GS1 record corresponding to the barcode prefix.