



Are trade wars class wars? The importance of trade-induced horizontal inequality[☆]

Kirill Borusyak^{a,b,1}, Xavier Jaravel^{c,b,*}

^a UC Berkeley, United States of America

^b CEPR, United Kingdom

^c LSE, United Kingdom

ARTICLE INFO

Dataset link: <https://data.mendeley.com/datasets/vrzvbmvgngw/1>

JEL classification:

F14

F16

F60

D63

Keywords:

Trade liberalization

Distributional effects

Inequality

ABSTRACT

What is the nature of the distributional effects of trade? This paper demonstrates conceptually and empirically the importance of “trade-induced horizontal inequality,” i.e. inequality that occurs among workers with the same level of earnings before the trade shock. This type of inequality does not affect the income distribution but generates winners and losers at all income levels. To quantify the horizontal inequality and changes in the income distribution induced by trade in a data-driven way, we develop a characterization of the welfare impacts, governed by simple and intuitive statistics of labor market and consumption exposure to trade. In the U.S., we find substantial heterogeneity in exposure and thus in the welfare effects of trade shocks across workers. Over 99% of the variance of welfare changes from trade shocks arises *within* income deciles. These findings run against a popular narrative that “trade wars are class wars.”

1. Introduction

Characterizing the effects of international trade shocks on the economy, in particular on welfare and inequality, is one of the key tasks of economists. Indeed, a large body of work studies changes in the income distribution induced by trade (e.g., [Helpman, 2018](#)). In this paper, we show that changes in the income distribution constitute only one of the dimensions of inequality brought about by trade, which we call “vertical inequality”. We highlight conceptually and empirically the importance of “horizontal inequality”, i.e. inequality occurring among workers with the same level of earnings prior to the shock.² This type of inequality generates winners

[☆] We are grateful to Pol Antras, Ariel Burstein, Grigory Franguridi, Oleg Itskhoki, Kalina Manova, Marc Melitz, and Jonathan Vogel for insightful comments, and to Josh Feng for helping with the analysis using U.S. Census Bureau data. This paper supersedes part of our earlier work, disseminated as “The Distributional Effects of Trade: Theory and Evidence from the United States.” Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2229 (CBDRB-FY24-0037). This research uses data from the Census Bureau’s Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. Jaravel gratefully acknowledges financial support from UKRI under grant agreement EP/X016056/1.

* Correspondence to: Houghton St, London WC2A 2AE, United Kingdom.

E-mail addresses: k.borusyak@berkeley.edu (K. Borusyak), x.jaravel@lse.ac.uk (X. Jaravel).

¹ Address: 222 Giannini Hall, Berkeley, CA, United States.

² Horizontal inequality is distinct from “residual” inequality, i.e. rising wage dispersion within occupations and industries (e.g., [Helpman et al., 2017](#)). While a change in residual inequality contributes to changes in the overall income distribution, shock-induced horizontal inequality does not, at the first order. Similarly, horizontal inequality differs from the “horizontal inequalities”, understood as inequalities across ethnic or other social groups (e.g., [Stewart, 2005](#)).

<https://doi.org/10.1016/j.jinteco.2024.103935>

Received 1 September 2022; Received in revised form 10 April 2024; Accepted 13 April 2024

Available online 25 April 2024

0022-1996/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

and losers at all income levels, without affecting the income distribution. Trade-induced horizontal inequality may be of particular importance to understand political support for trade policy, to the extent that people may express support based on their economic self-interest. Using data from the United States, we find that the horizontal, rather than vertical, component is the dominant force. In our baseline specification, over 99% of the variance of welfare changes from trade shocks arises within income deciles.

To quantify horizontal and vertical inequality, we develop a data-driven characterization of the welfare impacts of trade shocks governed by simple and intuitive statistics of labor market and consumption exposure to trade and by microeconomic primitives (i.e., income and substitution elasticities). This characterization holds in a class of quantitative trade models allowing for a broad set of preferences, including non-homothetic, and production functions. Our approach thus advances an emerging, influential literature (e.g., [Adão et al., 2022b](#), [Baqae and Farhi, 2024](#)) that bridges the gap between large-scale computational general equilibrium (CGE) models, which are traditionally used for quantitative trade policy analysis but lack transparency, and simple stylized models, which are linked to data in a transparent way but may lead to less reliable quantitative predictions.³

Our theoretical analysis characterizes the unequal welfare effects of trade shocks in terms of changes in wages (the “earnings channel”) and consumer prices (the “expenditure channel”).⁴ We show that changes in factor demand in response to small changes in iceberg trade costs can be decomposed into several terms corresponding to different economic forces: exports, import competition, imported intermediate inputs, non-homotheticities, and cross-industry substitution effects. Each of these terms is governed by an intuitive statistic measuring exposure to international trade. For instance, a factor whose employment is concentrated in industries that have high export ratios, directly or indirectly, will see factor demand grow after a trade liberalization, *ceteris paribus*. How factor prices respond to factor demand in turn depends on the elasticities of aggregate factor demand, which we characterize in terms of microeconomic primitives. We show that factor demand tends to be less elastic in non-traded industries, for which the wage impacts of labor demand changes are therefore stronger, providing another mechanism for the unequal effect of trade shocks.⁵ Using this theoretical framework, we formally define horizontal and vertical inequality and link it to heterogeneous exposures to trade shocks, within and across income groups.

Taking our characterization to the data, we evaluate a counterfactual where trade costs fall by 10% with all trading partners of the United States.⁶ We implement this analysis with a linked dataset measuring trade flows and input–output linkages across about four hundred industries covering the entire U.S. economy, matched to the heterogeneous employment structure and, following [Borusyak and Jaravel \(2023\)](#), consumption baskets of different groups of consumers and workers. While our theoretical results allow for any factor types, empirically we consider different groups of workers in the main analysis and study capital in a robustness check.

We first analyze a setting in which there is no mobility of workers across industries. In that case, each worker’s labor market exposure is simply her industry’s exposure.⁷ The key lesson emerging from this empirical analysis is that exposure differences and the corresponding distributional effects are primarily concentrated within income groups, rather than across. Over 99% of the variance of welfare changes arises within income deciles, generating horizontal inequality.

Despite the substantial distributional effects, which generate sizeable changes in relative earnings as well as winners and losers at all income levels, we find little impact of a fall in trade costs on the shape of the income distribution. Indeed, the Gini index falls by 0.0001 only, and the standard deviation of (real) log-wages falls by 0.0004. The spread between the 10th and 90th percentiles of welfare effects is over 2 percentage points within each decile, while variation across deciles is much smaller: all groups benefit on average and the gains are slightly *higher* for poorer households, ranging from 2.0% in the bottom decile to 1.8% for the top decile.⁸

To confirm that there are no strong distributional effects *across* groups of ex-ante similar workers, we conduct a robustness analysis across education groups. In this analysis, we consider two groups of workers — those with and without a college degree — and assume perfect mobility across industries. We again find that the effects are very similar across groups. The welfare gain from the 10% fall in trade costs is 1.7% for college-educated workers, compared with 1.6% for those without a college degree. In other extensions, we continue to find that horizontal distributional effects are the driving force and that vertical inequality is small. We study non-uniform changes in trade costs, as well as within-industry heterogeneity across firms,⁹ and we demonstrate the robustness of our results to choices of elasticities.

³ More precisely, relative to CGE modeling, our formulas explicitly highlight the economic mechanisms at play. Relatively to simple stylized models, such as [Arkolakis et al. \(2012\)](#), we incorporate more mechanisms important for policy counterfactuals and link them directly to moments of the detailed data. We note that our characterization obtains in a first-order approximation, which is guaranteed to hold for small shocks. Large shocks can be analyzed by integrating over a sequence of small shocks, as is common in CGE models (e.g., [Dixon et al., 2013](#)) and recently done by [Adão et al. \(2022b\)](#) and [Baqae and Farhi \(2024\)](#).

⁴ We thus extend the analysis of [Borusyak and Jaravel \(2023\)](#), who focus exclusively on the expenditure channel and on the partial equilibrium.

⁵ This mechanism has been informally proposed by [Rodrik \(1997\)](#) and shown to be relevant when analyzing the effects of labor supply shocks from immigration by [Burstein et al. \(2020\)](#). We demonstrate its importance for labor demand shocks, specifically those due to trade.

⁶ This is a 10% in the gross iceberg trade costs, similar in magnitude to a 10 percentage points reduction in tariffs. Following [Borusyak and Jaravel \(2023\)](#), we also assess the impacts of other shocks, including a trade liberalization with China specifically, historical reductions in trade costs, and the introduction of the “Trump tariffs” in 2018. We treat all of these shocks as changes in iceberg trade costs which do not generate tariff revenue, and use a first-order approximation.

⁷ While the assumption of no mobility may be most appropriate in the short-run, this analysis can be generalized to the case where labor mobility follows a Roy model with a finite but non-zero elasticity of industry labor supply, as in [Galle et al. \(2022\)](#). We have no reason to expect that the key lesson presented below would change qualitatively in that medium-run model, provided the industry labor supply elasticity is not too large, which is indeed the case according to the estimates of [Galle et al. \(2022\)](#).

⁸ Higher gains for poor households may look surprising, in particular in light of the canonical Heckscher–Ohlin model. Consistent with the Stolper–Samuelson theorem, relative labor demand for low-income workers falls after the trade shock in this analysis. Yet, an offsetting force dominates: low-income workers are employed relatively more in service industries, which have lower labor demand elasticities; as a result, a given labor demand shock induces, on average, a stronger wage response for them.

⁹ Using confidential plant-level microdata from the Census of Manufactures and the Management and Organizational Practices Survey, we find that more skill-intensive plants within the same industry tend to export more, consistent with [Burstein and Vogel \(2017\)](#). Nonetheless, this force is small relative to between-industry differences, on which we focus in the main analysis.

Thus, all our findings run against a popular narrative that “trade wars are class wars” (Klein and Pettis, 2020). Rather, trade causes winners and losers at all income levels. To the extent that losers are more likely to oppose trade liberalizations than winners are likely to express support, the large magnitude of horizontal distributional effects may lead to waning support for free trade.

Related literature. This paper contributes to an emerging literature that develops new theoretical tools to connect the effects of economic shocks in quantitative macroeconomic or trade models to detailed micro-level data, characterizing and decomposing welfare impacts in general equilibrium in terms of microeconomic sufficient statistics (Baqae and Farhi, 2019, Baqae and Farhi, 2024, Adão et al., 2022b, Borusyak and Jaravel, 2023, Oberfield and Raval, 2021, Kleinman et al., 2022). We incorporate additional channels, in particular via non-homothetic preferences, and propose a characterization of trade-induced horizontal inequality. Our decomposition allows us to assess the relative importance of the role of skill endowment emphasized by the Stolper–Samuelson theorem, the contributions of non-homothetic preferences (Caron et al., 2020), and the complementarity between goods and services (Cravino and Sotelo, 2019).

In independent work subsequent to ours, Adão et al. (2022b) develop a different decomposition for factor price changes due to trade into import and export channels and apply it to detailed firm-level data in Ecuador, focusing on vertical inequality. Relative to them, our characterization of trade-induced vertical inequality also captures the effects of import competition in intermediate demand (rather than in final demand only), making our model consistent with the standard industry-level gravity equation.¹⁰ Furthermore, we isolate the negative effects of import competition from the positive productivity effects of imported intermediate inputs.

Our results complement the analysis of Borusyak and Jaravel (2023), who characterize the effects of trade on inequality in a sufficient statistics framework that focuses entirely on heterogeneous consumption baskets.¹¹ Their analysis accounts for the effect of trade shocks on imported goods prices in a partial equilibrium framework holding factor prices fixed. In this paper, while mainly focusing on factor price responses to trade shocks, we account for the fact that endogenous changes in factor prices may also change households’ price indices in heterogeneous ways. In this sense, we provide an exposure-based framework accounting for both the earnings and expenditure channels of trade on inequality, as well as their interaction. We thus contribute to a small literature that analyzes the expenditure and earnings channels jointly, in a unified framework.¹²

Furthermore, our analysis relates to a broad literature characterizing the effects of trade on income inequality. First, a longstanding literature uses statistics on the net factor content of trade, guided by the Heckscher–Ohlin model (e.g., Katz and Murphy, 1992; Deardorff and Staiger, 1988, Krugman, 2000). More recently, Burstein and Vogel (2017) and Cravino and Sotelo (2019) show that the net factor content of trade is not an appropriate statistic for welfare in richer models. Our characterization provides a set of sufficient statistics in a modern, multi-sector gravity model. Second, many studies analyze the impact of trade shocks on relative factor prices, such as the skill premium (e.g., Burstein and Vogel, 2017; Costinot and Vogel, 2009) or wage differences across space (Autor et al., 2013; Caliendo et al., 2019; Adão et al., 2022a). Relative to them, we link the heterogeneous effects to the initial earnings distribution and provide a methodology to map changes in factor prices into changes in inequality.¹³ In this respect, our analysis is similar to Helpman et al. (2017), although they do not consider trade-induced horizontal inequality. Third, recent work by Galle et al. (2022) shows, by using exact hat algebra in a multi-sector gravity model, that the “China shock” generates strong distributional effects. Relative to them, our contribution is to quantify the extent to which the distributional effects of the shocks we study are horizontal rather than vertical.

Finally, our focus on the winners and losers from trade shocks speaks to the growing empirical literature on attitudes toward trade (e.g., Mayda and Rodrik, 2005; Jäkel and Smolka, 2017) and voting on trade (e.g., Van Patten and Méndez-Chacón, 2022). This literature assesses whether individual-level exposure to trade predicts support for trade agreements. Our analysis provides a new justification for these analyses by demonstrating, with a combination of theory and quantitative analysis, that individual-level exposure statistics govern the unequal welfare effects of trade across workers in a general equilibrium trade model. In addition, the rich set of statistics in our framework could be leveraged in future work on the political economy of trade.

The remainder of the paper is organized as follows. Section 2 presents the theoretical framework, Section 3 describes the data and elasticities used to take the theory to the data, and Section 4 presents the estimates of the distributional effects from counterfactual trade shocks in general equilibrium. Section 5 concludes.

¹⁰ Their assumption of Cobb–Douglas preferences across all intermediate inputs implies that domestic firms using foreign inputs in their supply chains are forced to exit in autarky. Our framework makes a more realistic assumption that domestic firms can substitute away from foreign inputs.

¹¹ Regarding the distributional effects of trade via heterogeneous consumption baskets, see also Porto (2006), Faber (2014), Fajgelbaum and Khandelwal (2016), Nigai (2016), Hottman and Monarch (2020), and Jaravel and Sager (2023). Atkin et al. (2018) document how foreign direct investment in retail also have distributional effects via consumption baskets.

¹² 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 (2020) generalizes the structural model of Fajgelbaum and Khandelwal (2016). We take a different approach by focusing on a set of exposure statistics measured in detailed data.

¹³ The regional dimension of the effects of trade could be studied using our exposure-based approach given appropriate data; we do not pursue this path due to data limitations. So far, this line of work has largely been silent about the effect of trade on wage inequality: Autor et al. (2013) do not find significantly different cross-sectional effects on skilled and unskilled wages (see Tables 6 and 7) and do not document the distribution of trade shocks across commuting zones. They instead find negative effects of trade with China on manufacturing employment in the U.S. at the level of commuting zones (and industries in Acemoglu et al., 2016), which is consistent with our model.

2. Theory

In this section, we theoretically characterize the distributional effects of counterfactual trade shocks in general equilibrium. We build on [Borusyak and Jaravel \(2023\)](#) who show that a simple statistic — the fraction of imports in spending — governs the effects of trade shocks on the purchasing power of consumers (the “expenditure channel”), holding domestic wages fixed. We extend their analysis to a general equilibrium setting in which both relative wages (the “earnings channel”) and relative purchasing power adjust. We show that intuitive observable statistics of labor market exposure to trade govern the distributional effects of trade through the earnings channel. Moreover, the earnings and expenditure channels can interact: domestic wage changes affect prices, which can lead to unequal purchasing power effects across consumers. We then link them to the two types of inequality induced by trade shock, horizontal and vertical. Details and proofs are relegated to Appendix A.

2.1. Setting

In this section, we describe the assumptions underlying our model, regarding preferences, the labor market, and technology. We then specify the trade shocks and welfare effects we study in the remainder of the paper.

Preferences. We study a neoclassical static global economy with $C + 1$ countries in which international trade is shaped by product differentiation, cross-country differences in technologies and endowments, and trade costs. We denote the Home economy $c = H$, the United States, with the set of foreign countries denoted by F . Product varieties (j, c) are defined as pairs of industry $j = 1, \dots, J$ and country of origin c , as in multi-sector versions of the [Armington \(1969\)](#) model.¹⁴ Preferences of domestic consumers across industries are characterized by a utility function $\mathcal{U}(q_1, \dots, q_J)$, which is left unrestricted in this section, and in particular can be non-homothetic. Preferences over varieties within each industry have constant elasticity of substitution (CES):

$$q_j = \left(\sum_{c=1}^{C+1} a_{jc}^{1/\xi_j} q_{jc}^{(\xi_j-1)/\xi_j} \right)^{\xi_j/(\xi_j-1)}, \quad (1)$$

where a_{jc} are taste shifters and $\xi_j - 1$ is the industry’s trade elasticity.¹⁵ The corresponding spending shares are denoted s_{jc}^x for a consumer with earnings x .¹⁶

Labor markets. In the labor market, workers are immobile across countries and exogenously grouped into types $i = 1, \dots, I$ with wages w_i per efficiency unit.¹⁷ Workers of the same type are perfectly substitutable in the labor market, supply labor inelastically, and are endowed with heterogeneous efficiency units, capturing within-type income differences. We denote worker’s earnings at the initial equilibrium by x_i ; workers are thus uniquely characterized by a pair of type and earnings, (i, x) . Type- i workers are freely mobile within a set of industries \mathcal{J}_i , but are not employed outside it. This formulation allows for a scenario with no mobility across industries (i.e., i are industry groups and $\mathcal{J}_i = \{i\}$), as well as a scenario in which i corresponds to education groups freely mobile across all domestic industries, as in our calibrations below.¹⁸

Technology. Domestic production in industry j combines primary factors L_ℓ^j with composite inputs Q_ℓ^j from all industries ℓ . We assume a Cobb–Douglas production function in terms of value added and intermediate inputs, $Q_{jH} = F_j^{VA} \left(L_1^j, \dots, L_I^j \right)^{1-\beta_j}$.

$\Pi_{\ell=1}^J \left(Q_\ell^j \right)^{\beta_\ell^j}$, with $\sum_\ell \beta_\ell^j = \beta_j$, and allow for any constant returns to scale and differentiable value-added aggregator across industries, F_j^{VA} . The Cobb–Douglas assumption for intermediate inputs is standard (e.g., [Acemoglu et al., 2012](#); [Caliendo and Parro, 2015](#)) and consistent with the stability of input shares in the U.S. input–output (IO) table over time. Within each industry, firms aggregate varieties from different countries in the same way as final consumers, according to (1). This assumption ensures that our model implies a standard industry-level gravity equation and satisfies a proportionality assumption embedded in country-specific IO tables.

¹⁴ As usual, the microfoundation from [Eaton and Kortum \(2002\)](#) is isomorphic, and [Krugman \(1980\)](#) or Melitz-Pareto ([Chaney, 2008](#)) versions would only differ by introducing home market effects, which tend to have small welfare impacts ([Costinot and Rodríguez-Clare, 2015](#)).

¹⁵ For consistency with industry-level data, we follow, e.g., [Caron et al. \(2014\)](#) in allowing for non-homothetic utility across industries but not within. This specification is in line with the finding of [Borusyak and Jaravel \(2023\)](#) that import spending shares within industries do not vary systematically across income groups, and it is necessary for our model to be consistent with the proportionality assumptions embedded in U.S. input–output tables. Non-homothetic demand within industries can be theoretically accommodated and would yield a characterization similar to [Propositions 1 and 2](#) below.

¹⁶ Throughout the paper, we indicate buyers in the superscripts and sellers in the subscripts. Agents are buyers in the product markets and sellers in the labor market.

¹⁷ Although we focus on workers and wages for the main analysis, our theoretical framework can accommodate arbitrary factors, such as capital.

¹⁸ This approach can be generalized to a finite elasticity of labor supply in each industry via a Roy model; see Appendix A.4 in our early draft ([Borusyak and Jaravel, 2018](#)) and [Galle et al. \(2022\)](#).

Trade and counterfactual trade shocks. We consider how domestic goods and factor prices adjust in general equilibrium (GE) following a set of infinitesimal reductions in trade costs between Home and some country (or set of countries) c . Specifically, products imported from $c \neq H$ to Home are subject to an iceberg trade cost τ_{jc} , while domestic firms face iceberg costs τ_{jc}^* when exporting to c , and both types of trade costs are allowed to change.

Since our detailed data only cover the U.S., we rule out changes in relative factor prices abroad by assuming that for every industry and foreign country, exports to Home are a small fraction of sales, and imports from Home are a small fraction of industry absorption. This assumption implies that relative product demand and price indices abroad do not significantly move after the trade shock with Home. We thus take all foreign prices as the single numeraire, while domestic factor and product prices are still fully endogenous, and trade shocks lead to terms-of-trade adjustments, as in the analysis of a small economy by Demidova and Rodríguez-Clare (2009). With our focus on a uniform change in trade costs with all countries, the assumption of fixed relative factor prices outside Home appears plausible, as the U.S. economy accounts for less than 6% of sales and absorption in the rest of the world.¹⁹

With changes in prices abroad ruled out, we remain agnostic about foreign endowments, preferences, technologies, and labor markets. We only require that foreign buyers aggregate varieties across countries of origin with the same elasticity ξ_j as domestic buyers, resulting in constant-elasticity export demand for domestic producers.

Finally, we allow for a trade imbalance in the domestic economy assuming, as in Costinot and Rodríguez-Clare (2015), that it is fixed in proportion to Home's GDP. Specifically, we assume that every consumer spends the same exogenous multiple of their earnings, ζx .²⁰

Equilibrium and welfare. The equilibrium is defined by (i) utility-maximizing allocation of spending across varieties by domestic consumers of each income level, (ii) profit maximization by domestic producers, (iii) labor and product market clearing conditions, and (iv) constant-elasticity export demand with exogenous shifters (see Appendix A.1 for details).

We adopt a cardinalization of welfare by applying a monotonic transformation to utility such that at the initial prices p , welfare equals nominal expenditures for all consumers. Specifically, for any combination of earnings x' and prices p' , welfare is defined as

$$\mathcal{W}(x', p') \equiv C(\mathcal{V}(\zeta x', p'), p), \tag{2}$$

where C and \mathcal{V} are the cost function and the indirect utility function corresponding to U , respectively.

We follow the standard approach defining the *change* in welfare for a given consumer as the equivalent variation divided by initial expenditures, $(\mathcal{W}(x', p') - \zeta x)/\zeta x$ (e.g. Fajgelbaum and Khandelwal, 2016; Borusyak and Jaravel, 2023). For example, the welfare change is equal to 0.01 if the trade liberalization is equivalent, in utility terms, to increasing total spending by 1% at the original prices. Cardinalization (2) ensures that locally around the initial equilibrium the welfare change simply equals $d \log \mathcal{W}$. By the envelope theorem (Roy's identity), a set of consumer price changes $d \log p_{jc}$ affects each consumer h in proportion to the spending shares, such that for a consumer of type i with initial income x ,

$$d \log \mathcal{W}_h = \underbrace{d \log w_i}_{\text{earnings channel}} - \underbrace{\sum_{j,c} s_{jc}^x d \log p_{jc}}_{\text{expenditure channel}}, \tag{3}$$

regardless of the demand system.²¹ Here the first term captures the change in the per-efficiency-unit wage (and thus in total earnings and expenditures) which only depends on the worker type and the second term is the change in the cost of living (the Laspeyres price index) determined solely by the initial earnings level.²² From (3), differential welfare gains between labor types can be decomposed into components related to changes in wages and prices, which we respectively call the earnings and expenditure channels of the distributional effects of the trade shock. As we discuss below, these channels can interact: domestic wage changes affect consumer prices and thus purchasing power, and potentially differentially so across consumer groups.

¹⁹ While the U.S. accounts for a substantial fraction of world GDP, exports from the U.S. constitute only 3.9% of absorption in other countries according to the World Development Indicators database for 2007. Exports to the U.S. similarly account for only 5.5% of foreign production. Similar to Demidova and Rodríguez-Clare (2013), our economy can be viewed as a limit with the shares of trade with the U.S. going to zero in foreign countries but U.S. trade shares held fixed.

²⁰ The U.S. has run a trade deficit every year since 1976, with imports exceeding exports by 48% in 2007. Although the ratio of net imports to GDP fluctuates over time, our assumption provides a better fit to the data than balanced trade. It is more common to assume that the absolute value of net imports is fixed (Dekle et al., 2008), but our approach is more tractable in a model with heterogeneous agents, as we do not need to keep track of income and expenditure changes separately.

²¹ Borusyak and Jaravel (2023) show the conditions under which (3) applies beyond neoclassical models, e.g. with endogenous product entry and exit, as in a generalization of Eaton and Kortum (2002), or the generalized Melitz-Pareto model of Kucheryavyi et al. (2023).

²² Note that in our model expenditure shares are entirely determined by initial income x , since all agents have identical non-homothetic preferences. In reality, different households with the same incomes are likely to have different expenditure shares on foreign goods due to taste heterogeneity. Our theoretical approach can be readily generalized to allow for unrestricted preference heterogeneity, using household-specific expenditure share s_{jc}^h . However, it is challenging to accurately measure preference heterogeneity because our data on expenditures do not allow us to follow the same households for more than a year. Annual expenditure patterns may overstate the true, long-run heterogeneity in spending shares on imports across households due to lumpy spending on tradables (e.g., purchasing a car from Germany) and measurement error. We conjecture that preference heterogeneity unexplained by income differences would strengthen the importance of horizontal inequality from trade shocks, relative to vertical inequality; in this sense, our approach can be viewed as conservative.

Notation. We finally introduce some notation. On the import side, we define IP_{jc} as the share of imports from c in domestic absorption of j at the initial equilibrium, with IP_j for the total import penetration. Let \widetilde{IP}_{jc} be the share of imports from c in industry absorption, both directly and indirectly via IO linkages, i.e. including imports of intermediate inputs. Similarly, the share of inputs imported from c , both directly and indirectly, in the cost structure of domestic production in industry j is denoted $\widetilde{IP}_{jc}^{Int}$. These shares are defined recursively as $\widetilde{IP}_{jc} = IP_{jc} + (1 - IP_j) \cdot \widetilde{IP}_{jc}^{Int}$ and $\widetilde{IP}_{jc}^{Int} = \sum_{\ell=1}^J \beta_{\ell}^j \widetilde{IP}_{\ell c}$. Finally, we define the share of imports from c in expenditures of consumers with income x by

$$ImpSh_c^x = \underbrace{\sum_j s_j^x \widetilde{IP}_{jc}}_{\text{Direct}} + \underbrace{\sum_j s_j^x (1 - IP_j) \widetilde{IP}_{jc}^{Int}}_{\text{Indirect}}, \tag{4}$$

where s_j^x is the spending share on industry j for those consumers. On the export side, $ExSh_{jc}$ denotes the share of exports to country c in j 's domestic output. $DomSalesSh_j$ denotes the share of domestic sales (both final and intermediate) in j 's total sales. The share of final domestic customers in total sales is DFS_j , and $\mu_{x|j}$ are the shares of sales to consumers with income x in j 's final sales.²³ We characterize domestic final demand of consumers with income x at the industry level by the income elasticity of expenditure shares $\psi_{xj} = \partial \log s_j^x / \partial \log x$ and the own- and cross-price elasticities $\epsilon_{xjk} = \partial \log s_j^x / \partial \log p_k$ measuring the response of j 's expenditure share to the change in industry k 's CES price index.

2.2. Exposure-based decomposition of labor demand and wage changes

We now state our main theoretical result on the wage responses to the trade shock. For notational brevity we present it for a shock that is uniform across industries and applies to bilateral trade with country (or set of countries) c : $d \log \tau_{jc} = d \log \tau_{jc}^* = d \log \tau < 0$ for all j ; an immediate generalization to non-uniform shocks across industries is found in Appendix A.4.

Proposition 1. *After a uniform reduction in bilateral trade costs with country c , changes in domestic wages $w = (w_1, \dots, w_I)$ satisfy*

$$\frac{d \log w}{-d \log \tau} = \underbrace{\widetilde{G}}_{\substack{\text{inverse labor demand} \\ \text{elasticity matrix}}} \cdot \underbrace{E \widetilde{D} \eta}_{\substack{\text{labor demand} \\ \text{response}}}. \tag{5}$$

Here η is a $J \times 1$ vector of direct industry exposure to the shock via several mechanisms:

$$\begin{aligned} \eta_j = (\xi_j - 1) & \left[\underbrace{ExSh_{jc}}_{\text{export effect}} - \underbrace{IP_{jc} \cdot DomSalesSh_j}_{\text{import competition effect}} + \underbrace{\widetilde{IP}_{jc}^{Int} \cdot (ExSh_j + IP_j \cdot DomSalesSh_j)}_{\text{intermediate input effect}} \right] \\ & + DFS_j \cdot \sum_x \mu_{x|j} \left[\underbrace{\psi_{xj} ImpSh_c^x}_{\text{income effect}} - \underbrace{\sum_{k=1}^J \epsilon_{xjk} (\widetilde{IP}_{kc} - ImpSh_c^x)}_{\text{substitution effects}} \right]. \end{aligned} \tag{6}$$

The “IO adjustment” $J \times J$ matrix \widetilde{D} is such that $(\widetilde{D} \eta)_j$ is the sum of direct industry j exposure η_j and indirect exposure in industries downstream from j . The “payroll composition” $I \times J$ matrix E captures the shares of industries j in type i payroll, such that $E \widetilde{D} \eta$ measures the payroll-weighted average shock exposure by labor type. Finally, \widetilde{G} is the (negative of the) $I \times I$ inverse matrix of macro labor demand elasticities with respect to w , given by (A14) in the Appendix.

The intuition behind Eq. (5) is that, with fixed labor supply, trade shocks affect wages via shifts in labor demand. Shifts in labor demand arise from product demand in industries which employ each type of labor. The novel characterization in Eq. (6) shows that the product demand response to a small shock can be decomposed into several mechanisms, each driven by observable exposure measures scaled by corresponding elasticities.

In the rest of this section we (i) provide the intuition for each of the margins of exposure, (ii) explain how industry exposure to trade translates into labor demand changes across worker types and, in turn, to the wage incidence of the shock, and (iii) characterize the overall distributional effects via purchasing power in addition to wages.

Margins of industry exposure to trade. The first two terms in (6) show the *export* and *import competition effects*. As export trade costs fall, export demand grows according to the trade elasticity $\xi_j - 1$, contributing to industry labor demand growth in proportion to the export share $ExSh_{jc}$. Similarly, falling import trade costs lower import prices, which drives the industry price index down in proportion to import penetration IP_{jc} . This leads to reallocation of spending between domestic and foreign varieties within each industry. Because this effect only influences domestic consumption, it is scaled by the domestic share of industry sales, $DomSalesSh_j$.²⁴

²³ We view x as discrete only for notational clarity and because we discretize income into bins in the empirical analysis.

²⁴ Unlike traditional factor content statistics, the measure of exposure to import competition in Proposition 2 is valid in the presence of international specialization. Consider an industry, such as toys, in which the U.S. has largely stopped producing. Then its factor intensity is largely irrelevant for domestic demand for labor types that span many industries (e.g. skill groups), and thus for their wages. Accordingly, such an industry does not have a sizeable effect on our exposure measure, while it can have large effects on the factor content of trade (e.g. Wood, 1995).

The third term relates to *imported intermediate inputs*. Access to cheaper intermediate inputs makes domestic varieties more competitive, helping them gain market shares both abroad and at home. Industries are more exposed to this mechanism when they have a higher share of imported inputs $\widetilde{TP}_{jc}^{\text{Int}}$ in production costs.

The final terms are the *income and substitution effects*. Partial equilibrium welfare gains, driven by the import share $ImpSh_c^x$ (Borusyak and Jaravel, 2023, Proposition 1), lead to higher spending on income-elastic industries (those with $\psi_{xj} > 0$). Moreover, demand for a domestic industry falls if substitute industries k (those with $\varepsilon_{xjk} > 0$) become relatively cheaper, due to their above-average import share, and if complement industries have below-average import shares. Both effects only influence domestic final sales, as combining consumers of different income, hence the scaling by their shares in total sales.

From industry exposure to labor demand and wage changes. We now explain how payroll-weighted average exposure determines the labor demand response to trade shocks in our model and which forces determine how it translates into incidence. To do so, we describe two key equations (proved in the Appendix), respectively characterizing product and labor market equilibria in changes in response to the shocks. With L denoting the vector of labor endowments in efficiency units, we first have:

$$d \log w + d \log L = E \cdot d \log VA + V \cdot d \log w. \tag{7}$$

This shows that the change in the total earnings of each labor type, which equals the wage change in our counterfactual since $d \log L = 0$, can arise from two sources. First, changing industry value added, $d \log VA$, and holding the labor mix of each industry fixed, mechanically changes the earnings of each labor type according to the payroll composition matrix

$$E = (e_{j|i})_{i,j},$$

where $e_{j|i}$ is the share of industry j in type i 's earnings. Second, changes in relative wages lead to reallocation of payroll shares across labor types, which is governed by the labor substitution matrix

$$V = \left(\sum_j e_{j|i'} \frac{\partial \log v_{ij}}{\partial \log w_{i'}} \right)_{i,i'}.$$

The (i, i') element of this matrix characterizes the cross-wage elasticity of the fraction of j 's value added that accrues to value type i , v_{ij} , with respect to the wages $w_{i'}$, averaged across industries j with i 's payroll composition weights. Intuitively, the earnings of type i grow if workers substitutable to them become more expensive. While E is observed, V is not, but its measurement and parameterization simplify substantially in the two special cases described in Section 3.3.

The second equation behind Proposition 1 characterizes changes in the product market equilibrium:

$$d \log VA = \eta \cdot (-d \log \tau) + D \cdot d \log VA + G \cdot d \log w. \tag{8}$$

The first term in this equation shows that an industry's value added is directly affected by the trade shock through the various mechanisms in (6). There are also two indirect effects. First, changes in value added propagate from downstream industries up through the demand for intermediate inputs, with D summarizing the shares of domestic intermediate buyers in j 's sales. Second, wage changes affect both the purchasing power of domestic consumers and the prices of domestic goods and thus demand for them, as summarized by the G matrix.²⁵

Setting the wage changes on the right-hand of (7) and (8) to zero, we obtain the expression $E \bar{D} \eta \cdot (-d \log \tau)$ for the change in labor demand. Here the exposure of industries to trade, $\eta \cdot (-d \log \tau)$, propagates upstream through the IO matrix and directly affects each labor type according to their payroll composition. Further accounting for the feedback loops via labor market competition, consumer purchasing power, and the production costs (hence prices) of domestic goods, we obtain the incidence of these shocks on wages in (5).

We note that the general equilibrium feedback loops necessarily generate a gap between labor demand shocks and their wage incidence. One manifestation of this gap is that a version of Lerner's symmetry holds in our model. Specifically, in Appendix A.5 we establish an equivalence, in terms of welfare and relative wages across all workers, between reducing all importing costs by a constant $d \log \tau$ and reducing all exporting costs by industry-dependent values $\frac{\xi_j}{\xi_j - 1} d \log \tau$.²⁶ While labor demand changes induced by these two shocks are very different, their incidence is the same.

Welfare responses. Given wage changes characterized by Proposition 1, we can obtain welfare changes for all agents. We have:

Proposition 2. *After a uniform reduction in bilateral trade costs with country c , the welfare change for agent h of type i with initial income x satisfies*

$$d \log \mathcal{W}_h = \underbrace{(d \log w_i - d \log \bar{w})}_{\text{earnings channel}} - \underbrace{ImpSh_c^x d \log \tau}_{\text{price effect of trade costs}}$$

²⁵ Appendix A.2 formally defines D and G and uses them to derive \bar{D} and \bar{G} , which enter (5).

²⁶ There are two differences of this result from the standard Lerner symmetry, even in its modern general treatment by Costinot and Werning (2019): it applies to iceberg trade costs, rather than changes in import tariffs and export subsidies, and the changes in the importing and exporting costs are not equal in magnitude. Unlike the standard theorem, our result relies on the assumption of a small economy in the sense of Section 2.1.

$$+ \underbrace{ImpSh^x d \log \bar{w}}_{\text{price effect of avg.wages}} - \underbrace{\sum_j s_j^x (1 - \bar{IP}_j) (d \log \bar{W}_j - d \log \bar{w})}_{\text{segregation effect}}, \tag{9}$$

where $d \log \bar{w} = \sum_i v_i d \log w_i$ is the change in the average wage in the economy, v_i denotes the initial payroll share of type i , $ImpSh^x$ is the spending share on imports from all countries, $1 - \bar{IP}_j$ is the share of domestic factors in the cost structure of the industry (accounting for its full supply chain) and $d \log \bar{W}_j$ is the average change in the wage of those factors defined by (A7).

The first term in (9) is the earnings channel, capturing the wage growth of type i (relative to the average wage change in the economy). The other three terms jointly determine the expenditure channel, i.e. the purchasing power effects of the shock. Specifically, the second term is the partial equilibrium effect of trade costs on prices, which is governed by the consumer's spending share on imports from c , directly or indirectly (Borusyak and Jaravel, 2023). The third term, governed by the share of spending on imports from all foreign countries, can be viewed as a terms-of-trade adjustment: if domestic wages grow on average, all imports become relatively cheaper, benefiting the consumers who spend relatively more on them. The final term, which we label a "segregation effect", captures the idea that, if some group of consumers tends to buy goods from industries where wages grow relatively more after the shock (directly or in their supply chains), this group will benefit less. For instance, if low-income households predominantly purchase products with low skill-intensity, then a trade shock leading to a fall in the relative wage of low-skill workers would also have a positive indirect effect on the purchasing power of low-income households through the price changes induced by these wage changes.²⁷

2.3. Horizontal and vertical inequality induced by trade shocks

Given welfare effects for all agents h , $d \log \mathcal{W}_h$, which is determined by their labor market type and initial earnings, we analyze two types of estimands of independent importance. First, we characterize the heterogeneity of welfare effects across workers, decomposing it into "vertical" and "horizontal" components. Second, we characterize the impact of the shock on the income distribution, showing that only the vertical component of heterogeneous welfare effects contributes to changes in the income distribution.

The variation in welfare effects across workers arises both across and within groups of initial earnings, x . Formally, we use a variance decomposition for welfare effects into the components across and within groups of initial earnings, x , using a variance decomposition:

$$\text{Var} [d \log \mathcal{W}] = \underbrace{\text{Var} [\mathbb{E} [d \log \mathcal{W} | x]]}_{\text{"vertical" inequality}} + \underbrace{\mathbb{E} [\text{Var} [d \log \mathcal{W} | x]]}_{\text{"horizontal" inequality}}. \tag{10}$$

Here the variances and expectations are taken across the distribution of workers h .

In contrast, the effects of the trade shock on the distribution of (real) incomes, at the first order, arise from the vertical component only. For instance, the following result holds for the standard deviation of log-earnings:²⁸

Proposition 3. *The change in the standard deviation of log-earnings due to a small shock can only be non-zero, at the first order, if the welfare changes are correlated with the initial earnings:*

$$\text{SD}(\log x + d \log \mathcal{W}) - \text{SD}(\log x) = \text{Corr} [d \log \mathcal{W}, \log x] \cdot \text{SD}(d \log \mathcal{W}) + o(d \log \tau). \tag{11}$$

This result follows because the heterogeneity of $d \log \mathcal{W}$ has only a second-order effect on $\text{Var} [\log x + d \log \mathcal{W}]$, unless the welfare change $d \log \mathcal{W}$ and the initial log-earnings $\log x$ are correlated; see Appendix A.6 for the proof.

The finding that horizontal inequality does not matter for the income distribution holds more generally. Proposition 4 in the Appendix shows that if there is no vertical inequality, the cumulative distribution function of real incomes is unchanged, at the first order. Thus, other conventional statistics of inequality, such as the Gini index, are also unchanged.

We note that the distinction between horizontal and vertical inequality is important for how optimal policy responses to trade shocks may differ from those demanded by voters. In the presence of a distortionary income tax-transfer system, in standard models the social welfare weights across agents for monetary transfers are determined by disposable income (e.g., Mirrlees, 1971). Thus, the social planner would adjust taxes and transfers in response to trade-induced vertical inequality, as in Antràs et al. (2017). Horizontal distributional effects do not create such a motive for additional redistribution, given that the social marginal utility of income is the same for workers with the same initial earnings. However, horizontal distributional effects may be a key determinant of political support for trade, to the extent that people express support for trade liberalizations based on economic self-interest (e.g., Mayda and Rodrik, 2005).²⁹

²⁷ For the analyses of the segregation effect, see Clemens et al. (2019) and Wilmers (2017), as well as our early draft (Borusyak and Jaravel, 2018).
²⁸ With our cardinalization, real expenditures (i.e., welfare) equal ζx rather than x , but the constant multiplier ζ is immaterial for inequality. For notation brevity, in Proposition 3 we add the welfare change to workers' earnings, rather than expenditures.
²⁹ Similarly, while vertical distributional effects may dissipate the gains from a trade liberalization when the social welfare function takes inequality into account (as in Atkinson, 1970 and Galle et al., 2022), horizontal distributional effects do not have such implications.

3. From theory to data

Our theoretical results provide a transparent way of connecting theory to data and guide our empirical analysis, which proceeds in six steps. First, we measure each statistic of direct industry exposure to trade in (6). Second, we adjust these statistics for input–output linkages, via the \tilde{D} matrix, for example measuring the share of industry output that is exported to c not only directly but also in downstream industries. Third, we obtain labor demand shifts for each group by averaging industry exposure using the fractions of different industries in the group’s payroll, captured by the E matrix, as weights. Fourth, we translate these labor demand shifts into the general equilibrium wage changes by applying the \tilde{G} matrix, given particular assumptions on the labor market substitution introduced below. Fifth, we measure the welfare effects for all workers, $d \log \mathcal{W}_h$, accounting for changes in both wages and cost-of-living in general equilibrium, via Proposition 2. Finally, we decompose the distributional effects of the shock into the horizontal and vertical components, as well as the impacts of the shock on measures of earnings equality.

To implement these steps, Propositions 1 and 2 require three types of inputs, which we discuss in turn in this section. Section 3.1 describes the data sources used to measure statistics characterizing the initial equilibrium: the payroll composition of each type of workers, trade shares of each industry and IO linkages between them, and consumption baskets by consumer group. We supplement these statistics by specifying the relevant elasticities: on the demand side in Section 3.2 and in the labor market in Section 3.3.

3.1. Measuring exposure to trade

Our main dataset which characterizes the initial equilibrium at the level of detailed industries in year 2007 builds on the data from Borusyak and Jaravel (2023). They focus on heterogeneous exposure to trade shocks across consumer groups; we link their data to the payroll composition by worker type and initial earnings, which is key to characterize worker exposure to trade in the labor market, per Proposition 1. We first summarize the industry-level data construction from Borusyak and Jaravel (2023) and how we link them to worker exposure; Appendix B.1 reports the details. We then describe the confidential firm-level data on within-industry heterogeneity of worker exposure to trade from the U.S. Census Bureau used for our robustness analysis.

Our main dataset is based on the detailed BEA input–output table from 2007. The IO table provides the most detailed available accounts of the entire U.S. economy at the level of 389 six-digit industries. We use it as a source of trade shares, input–output linkages, total consumer expenditures by industry, and total payroll by industry. Measuring trade shares from the BEA data is attractive, as trade in services is accounted for and trade flows are measured from the same data as domestic output, ensuring consistency. For each industry, import penetration is computed as the fraction of total imports in absorption, defined as output plus imports minus exports, and the export share as a fraction of total exports in output. The IO table further yields the composition of suppliers for each buying industry and the composition of buyers for each supplying industry, allowing us to measure indirect import shares in Eq. (4) and perform the IO adjustment to labor demand in (5).

We match several other datasets to the IO table. First, we decompose industry-level trade flows by trading partner using the 2007 U.S. international trade flow tabulations from the Census Bureau made available by Schott (2008). This allows us to measure trade shares for specific countries or groups of countries separately, specifically China, NAFTA countries (Mexico and Canada), and 34 developed economies (OECD members, except NAFTA, plus Taiwan and Singapore), except for services. Second, we decompose industry-level personal consumption expenditures into bins of household income using detailed expenditure categories from the Consumer Expenditure Survey (CEX) by the Bureau of Labor Statistics. Finally, we decompose the payroll of each industry into the contributions of workers by bin of their earnings and by other characteristics we use for robustness: two education levels (workers with and without a college degree), Census region, race, age, gender, and occupation. To do so, we use the American Community Survey (ACS), which is the long form of the population census answered by a random 1% sample of the U.S. population every year (Ruggles et al., 2015). In all of these cases we maintain the internal consistency of our dataset by keeping the industry totals — of trade flows, personal spending, and compensation of employees, respectively — from the IO table and applying group-specific percentages to those totals.³⁰

We finally supplement industry-level data with a plant-level dataset using the confidential Census Bureau data to empirically assess the importance of the within-industry heterogeneity of worker exposure to trade. Specifically, using the Economic Census and Customs microdata (LFTTD), we measure export shares by firm. We complement these data with the Longitudinal Employer–Household Dynamics (LEHD) sample to match firms to workers in different deciles of earnings, and therefore compute the differences between deciles in their exposure to exporting both within and across industries. Appendix B.2 describes the data construction steps for this analysis.

3.2. Demand system

Taking Proposition 1 to the data requires specifying a demand system, to characterize both the substitution patterns across countries of origin within each industry (i.e., the trade elasticities) and the income and substitution effects across domestic industries. We first discuss the choice of the trade elasticities. We then discipline the analysis of cross-industry demand with nested non-homothetic CES preferences. We introduce this demand system and its relevant properties and describe the choice of parameters. Income elasticities are estimated within our sample; for substitution elasticities across and within industries, we take prevailing values in the literature and consider robustness to a range of values.

³⁰ For 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 non-classical measurement error in the CEX (e.g., Garner et al., 2009); see Borusyak and Jaravel (2023) for further discussion.

Trade elasticities. Our baseline value for the elasticity of substitution between domestic and foreign varieties, ξ_j , is 3.5 in all six-digit IO industries, which is equivalent to a trade elasticity of $\xi_j - 1 = 2.5$. This baseline value is close to the median elasticity of 3.7 reported in Broda and Weinstein (2006) for ten-digit commodities, and of 3.4–3.7 in Soderbery (2015) using the same Broda–Weinstein method but for eight-digit commodities and for different years of data. It is also close to the mean elasticity of 3.6 in Ossa (2015).

In robustness checks, we consider values of ξ_j between 1.9 and 5.1. This range corresponds to the estimates from the Soderbery (2015) LIML procedure and from Simonovska and Waugh (2014), respectively, and also covers the typical values in Feenstra et al. (2018). Moreover, we also allow ξ_j to vary across three-digit IO industries according to the estimates from Broda and Weinstein (2006).³¹

Nested non-homothetic CES demand. To characterize cross-industry demand, we consider a nested version of the non-homothetic CES (NHCES) utility function (Hanoch, 1975; Comin et al., 2021), proposed by Borusyak and Jaravel (2023). Under NHCES, income and substitution elasticities are governed by distinct parameters, allowing us to specify and analyze them separately.³² The nested version of NHCES, which features the complementarity of goods and service consumption in the outer nest, further allows us to incorporate the mechanism from Cravino and Sotelo (2019) who show that the substitution of consumers towards services in response to a trade liberalization can affect relative wages.

Indexing the goods and services sectors by r , our nested NHCES utility function is defined recursively by

$$\begin{aligned}
 \mathcal{U} &= \left(\sum_{r=\text{Goods, Services}} q_r^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)}, \\
 q_r &= \left(\sum_{j \in r} \left(a_j \mathcal{U}^{\varphi_j(\varepsilon_r-1)} \right)^{1/\varepsilon_r} q_j^{(\varepsilon_r-1)/\varepsilon_r} \right)^{\varepsilon_r/(\varepsilon_r-1)}.
 \end{aligned}
 \tag{12}$$

The outer nest in the first line is a CES aggregate of consumption between the goods and services sectors with elasticity $\rho \neq 1$, and the inner nest in the second line aggregates consumption across industries within each sector ($j \in r$) with elasticity $\varepsilon_r \neq 1$. The industry consumption q_j is the homothetic CES aggregator across country varieties from (1). Primitive parameters $\varphi_{jc} < 1$ determine how tastes $a_j \mathcal{U}^{\varphi_j(\varepsilon_r-1)}$ vary with consumer utility: a higher $\varphi_{jc}(\varepsilon_r - 1)$ translates into a higher income elasticity of the industry within each sector (see Equation (A20) for the formula). Whenever $\varphi_{jc} \neq 0$, income affects tastes via its impact on utility, i.e. preferences are non-homothetic.³³

Substitution effects. With this demand system, the own- and cross-price elasticities are governed by ε_r and ρ exactly in the same way as they would under the standard (homothetic) nested CES preferences. In particular, in Appendix A.8 we use the results from Borusyak and Jaravel (2023) to show that, for industry j that belongs to sector r , substitution effects in (6) satisfy:

$$\begin{aligned}
 \sum_{k=1}^J \varepsilon_{xjk} \left(\widetilde{IP}_{kc} - ImpSh_c^x \right) &= (1 - \varepsilon_r) \left(\widetilde{IP}_{jc} - ImpSh_{rc}^x \right) \\
 &+ (1 - \rho) \left(ImpSh_{rc}^x - ImpSh_c^x \right),
 \end{aligned}
 \tag{13}$$

where $ImpSh_{rc}^x = \sum_{k \in r} s_k^x \widetilde{IP}_{kc} / \sum_{k \in r} s_k^x$ is the total (direct and indirect) average share of imports in expenditures on sector r for a consumer with income x .

Eq. (13) shows that, provided $\varepsilon_r > 1$, in response to a uniform trade liberalization consumers reallocate consumption towards industries (within sectors) with the share of imports that is higher than the sectoral average, other things equal. Similarly, provided $\rho < 1$ (a condition satisfied under the prevalent estimates, as discussed below), they substitute away from the manufacturing sector, where the import share is larger than in services.

We set the elasticity of substitution between goods and services to $\rho = 0.6$, obtained from Cravino and Sotelo (2019). For robustness we consider the range between 0.2 and 0.85, as in Comin et al. (2021) (see also Cravino and Sotelo, 2019). These values all indicate complementarity between goods and services in consumption.

The elasticities of substitution between industries within goods and services, ε_r , are more difficult to obtain (Dawkins et al., 2001; Costinot and Rodríguez-Clare, 2015). As they are expected to lie between ρ and ξ_j , we set $\varepsilon_r = 2$ in the baseline analysis and consider values between 0.6 and 3.5. A recent paper by Redding and Weinstein (2017) estimates 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 (2020) using 4-digit HS industries is 2.78. The range of elasticities we use covers all of these values.

³¹ Specifically, we take the median elasticity value across all ten-digit commodities corresponding to the three-digit IO code. To further reduce noise, we take the original values from Broda and Weinstein (2006) and the reanalysis of their method by Soderbery (2015) and average the two numbers.

³² NHCES possesses other attractive properties compared with alternative non-homothetic demand systems: income effects are non-negligible in presence of long-run economic growth (Comin et al., 2021) and the expenditure channel of trade is not mechanically biased in favor of low-income consumers (Borusyak and Jaravel, 2023).

³³ This way of writing the utility function, following Borusyak and Jaravel (2023), differs slightly from that of Comin et al. (2021) to better resemble traditional nested CES when $\varphi_j \equiv 0$. A single-tier version of (12) is equivalent to Eq. (1) in Comin et al. (2021) with $1 - \varphi_{jc}$ as relevant income elasticity parameters. As in Comin et al. (2021), a parameter restriction $\varphi_{jc} < 1$ is required for integrability.

Income effects. Proposition 1 requires the knowledge of income elasticities, which characterize industry-level Engel curves. We estimate them directly using CEX data on the heterogeneity of consumption baskets across income groups. Specifically, we first estimate income semi-elasticities by regressing, for each spending category in the CEX, spending shares on the log of average expenditure across income bins. We then convert them to elasticities ψ_j by IO industries; Appendix B.3 provides a complete discussion of the estimation steps. The resulting estimates are in line with those reported by Aguiar and Bils (2015) at a more aggregated level.

As Comin et al. (2021) point out, the empirical evidence suggests that the slopes of the Engel curves (in logs) are stable with respect to income. We therefore ignore the possibility that income elasticities can vary by income levels and set $\psi_{x_j} \equiv \psi_j$ for all x .³⁴

3.3. Labor market mobility and labor substitution

To implement Proposition 1, we also need to specify substitution patterns in the labor market. To assess the nature of the distributional effects of trade, and in particular whether they occur primarily within or across ex-ante similar groups of workers, we consider two polar cases. We first present the setting for the worker-level analysis, assuming no mobility of workers across industries; then we turn to a setting with perfect labor mobility across industries and two labor types defined by education.

3.3.1. Setting for the worker-level analysis

We first consider a setting in which workers cannot move across industries, i.e. $J(i)$ is a singleton for any type i . The lack of mobility simplifies the analysis in several ways. Since E maps each labor type to its industry, the change in labor demand is the same for all worker types in the same industry, and it equals the industry exposure from (6). Moreover, we show in Appendix A.9.1 that wage responses to the trade shock are identical for all workers in the same industry, equal to the (endogenous) growth of industry value added, and unaffected by within-industry labor substitution in V . Thus, we can equivalently view each worker as a distinct type in the labor market (as our “worker-level analysis” label suggests) or group all workers from the same industry into the same type. For computational tractability, we choose the latter option, setting $I = J$, $E = \mathbb{I}_J$, and $V = 0_{J \times J}$. Finally, the inverse labor demand matrix \tilde{G} is fully characterized by observable statistics of the initial equilibrium and the demand parameters from Section 3.2, in the same way as η in (6).

While the inverse labor demand matrix is generally a complex $J \times J$ matrix, this tractable setting allows us to make progress in understanding how exposure to trade maps into the incidence of trade shocks. Specifically, we show that in industries with lower trade shares (i.e., export shares and import penetration rates), wages are more responsive to shifts in labor demand, compared to more traded industries, as Rodrik (1997) informally argued (see also Slaughter, 2001 and Burstein et al., 2020). Intuitively, the own-wage elasticity of labor demand is higher for more traded industries because the strongest substitutability in our model is between domestic and foreign varieties within an industry. When domestic wages grow, prices of domestic varieties increase, inducing shifts in demand. Both domestic and foreign buyers can substitute away to foreign varieties in highly traded industries, but not so much in others.

We prove this result in a restricted model, in which only the export and import competition effects arise, while intermediate inputs, income, and substitution effects are shut down. We find in our analysis of Section 4.1 that this result holds qualitatively even when all mechanisms are operative. Formally, suppose $\tilde{D} = \mathbb{I}_J$ (i.e., there are no intermediate inputs), $\psi_{x_j} \equiv 0$, and $\varepsilon_{xjk} \equiv 0$ (i.e., preferences are Cobb–Douglas across industries), and consider a set of shifts to labor demand $d \log L_j^D$ (or, equivalently, a similar reduction in labor supply). Then we prove in Appendix A.9.1 that the wage response to these shifts, $d \log w = \tilde{G} \cdot d \log L^D$, satisfies

$$d \log w_j = \frac{d \log L_j^D}{1 + (\xi_j - 1) T_j} + \frac{DomSalesSh_j}{\zeta_2 (1 + (\xi_j - 1) T_j)} \cdot \left(\sum_{k=1}^J \frac{e_k d \log L_k^D}{1 + (\xi_k - 1) T_k} \right), \tag{14}$$

where $T_j = ExSh_j + IP_j \cdot DomSalesSh_j$, $\zeta_2 = 1 - \sum_j \frac{e_j DomSalesSh_j}{1 + (\xi_j - 1) T_j} \in (0, 1)$, and e_j is the payroll share of industry j in the economy.

The first term in (14) shows that wages in more traded industries are less responsive to shifts in labor demand in their own industry, via the T_j term which increases in both the export share and the import penetration rate.³⁵ The second term shows that they are also less sensitive to the economy average shift in labor demand, via both higher T_j and lower $DomSalesSh_j$.

3.3.2. Setting for the robustness analysis across two education groups

To assess whether there are distributional effects across groups of ex-ante similar workers, we consider for robustness a setting with full labor mobility across industries and two labor types — education groups in our empirical analysis — which we denote H and L (high and low skilled). The general result of Proposition 1 simplifies again, but in different ways. We first show that this model features a sufficient statistic which summarizes labor substitution in all industries, which we refer to as the macro elasticity

³⁴ An alternative approach would be to estimate the primitive φ_j parameters of NHCES and learn the heterogeneity of income elasticities across income levels from the structure of the demand system. That approach, however, would rely more heavily on parametric assumptions.

³⁵ This term directly parallels Eq. (13) in Burstein et al. (2020): both equations show how the occupation (or industry) labor demand elasticity is a weighted average between the elasticity of substitution across occupations (or industries; in our case here it equals to one) and between domestic and foreign goods, with weights depending on tradability.

of labor substitution, σ_{macro} . The local elasticities of substitution between the two labor types at the initial equilibrium in industries j , σ_j , enter the V matrix, and therefore \bar{G} , through this scalar parameter. Specifically, Appendix A.9.2 proves that

$$V = (\sigma_{\text{macro}} - 1) \begin{pmatrix} -v_L & v_L \\ v_H & -v_H \end{pmatrix}, \tag{15}$$

where

$$\sigma_{\text{macro}} - 1 = \sum_j e_j \frac{v_{H|j} v_{L|j}}{v_H v_L} (\sigma_j - 1), \tag{16}$$

e_j is the payroll share of industry j in the economy, and as before $v_{i|j}$ and v_i are payroll shares of type $i = H, L$ in industry j and the overall economy, respectively.³⁶

We follow Burstein and Vogel (2017), Cravino and Sotelo (2019), and Caron et al. (2020) in calibrating the macro elasticity directly rather than aggregating it from micro estimates. For the baseline analysis, we use the estimate of $\sigma_{\text{macro}} = 1.41$ obtained by Katz and Murphy (1992). We check robustness to the range of [1.41, 1.8], with the upper bound corresponding to the estimates from Acemoglu (2002) and Acemoglu and Autor (2011).

The simple structure of the V matrix allows us to gain insight into the wage responses to the trade shock in this setting. Appendix A.9.2 shows that the skill group that is initially specialized in industries that will grow faster after the shock will experience a higher wage growth, all the more so if σ_{macro} is small:

$$d \log \frac{w_H}{w_L} = \frac{1}{\sigma_{\text{macro}}} \left(\sum_j v_{H|j} d \log V A_j - \sum_j v_{L|j} d \log V A_j \right). \tag{17}$$

According to (8), which industries grow faster depends on their exposure to the shock, although it also depends on the general equilibrium feedback loops encoded in the G matrix.

4. Quantitative results

To assess the importance of both vertical and horizontal distributional effects, we first consider the “worker-level” setting with no mobility of workers across industries. Second, to confirm the lack of distributional effects across groups of ex-ante similar workers, we consider the setting with two education groups, assuming perfect mobility of each group of workers across industries. Third, we present several additional extensions and robustness checks. In all analyses, we find little role for the distributional effects across income and education groups.

4.1. Trade-induced inequality across workers

In this section, we conduct the worker-level analysis of the welfare effects of a counterfactual 10% fall in iceberg trade costs.³⁷ We first quantify the relative importance of horizontal and vertical distributional effects in Section 4.1.1. We then decompose each type of distributional effects into several mechanisms, starting with vertical distributional effects in Section 4.1.2 and then turning to horizontal distributional effects in Section 4.1.3.

4.1.1. The relative importance of trade-induced horizontal and vertical inequality

Fig. 1 depicts the measures of worker exposure to trade by decile of the income distribution, showing that exposure varies primarily within deciles rather than across. This is the key data fact that explains why trade-induced inequality is primarily horizontal. Using Proposition 1, we present the five components of worker exposure, $E \bar{D} \eta$, multiplying these terms by the 10% change in trade costs. The results are directly informative about the drivers of the labor demand response to trade liberalizations for different workers. For each income decile, we report average worker exposure along with the 10th and 90th percentiles of the worker-level exposure distribution. The within-decile variation arises from different industries employing workers from the same income decile.

Panel A shows changes in labor demand resulting from the export effect after a 10% fall in trade costs. The increase in labor demand is larger for higher-income workers, ranging from about +1.2% for the average worker in the first (i.e., bottom) decile to about +2.5% on average within the top decile. The change in labor demand varies substantially more across workers within deciles, with 90–10 gaps between 3.6p.p. and 5.1p.p. Panel B reports the changes in labor demand from the import competition effect: the fall in labor demand is more pronounced for higher-income workers, with a change of about −1.7% in the top decile compared with −0.9% in the bottom decile. Heterogeneity in the labor demand effects of import competition is large within each

³⁶ We note that $\sigma_{\text{macro}} - 1$ is generally not a weighted average of $\sigma_j - 1$: the sum of weights, $\sum_j e_j \frac{v_{H|j} v_{L|j}}{v_H v_L}$, is smaller than one unless all industries have the same skill composition.

³⁷ Since we use the first-order approximation, the choice of the size of the shock is immaterial and is based on presentation clarity only.

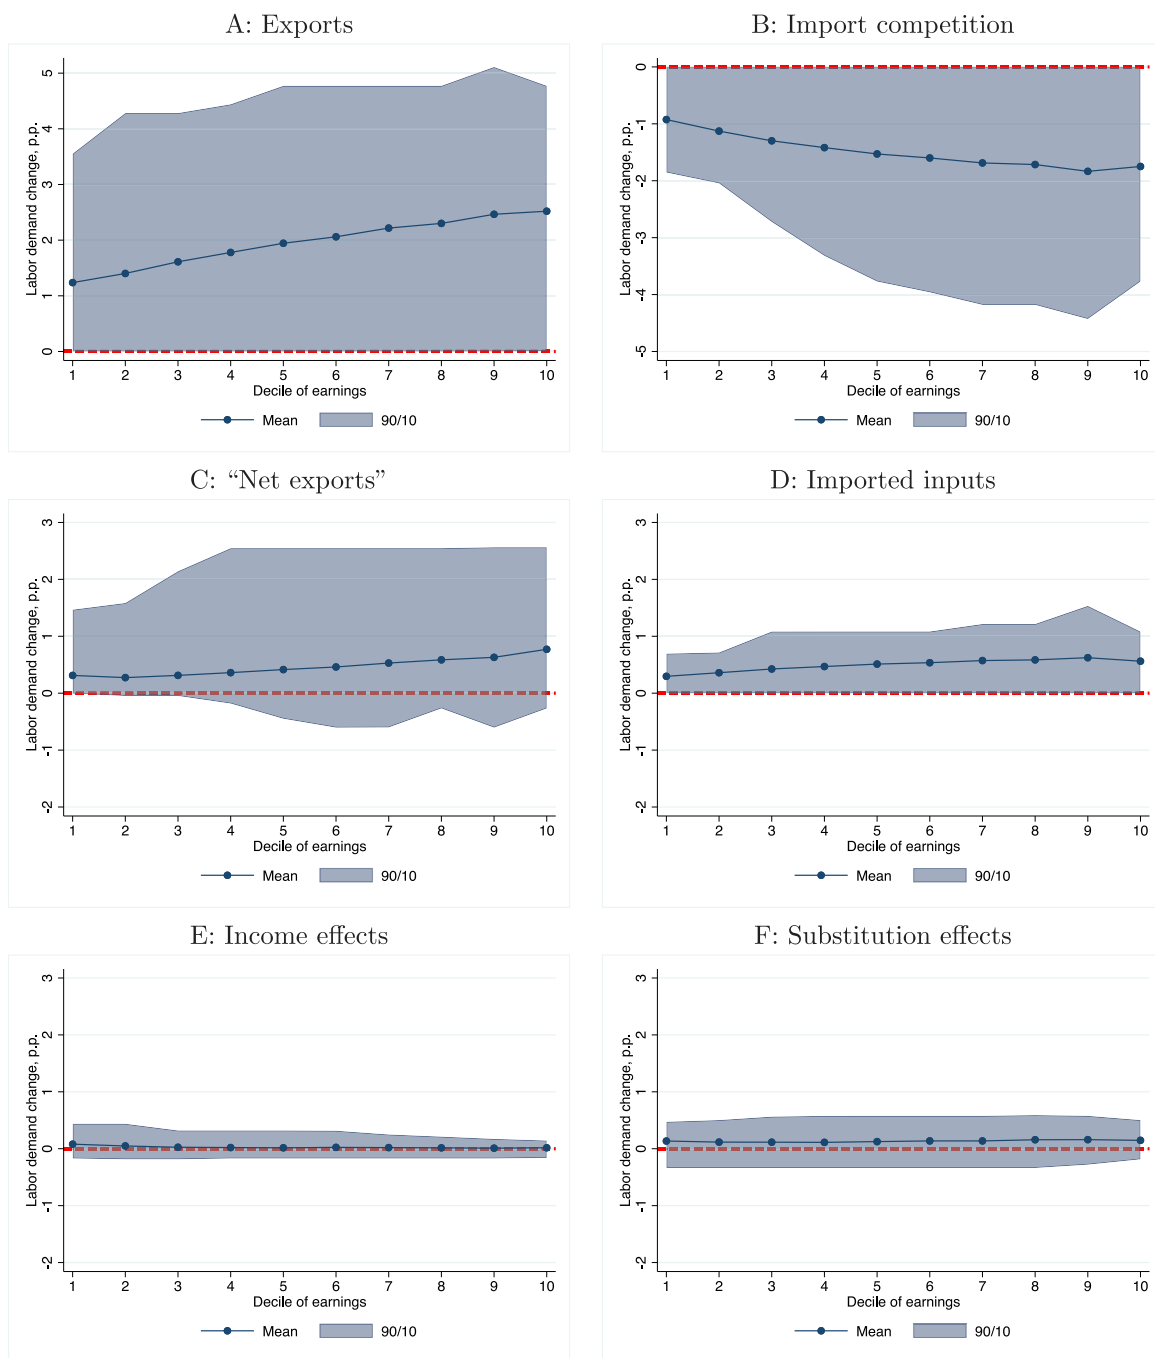


Fig. 1. Margins of worker-level exposure to a fall in trade costs, across and within income deciles. *Notes:* This figure groups workers from the ACS data by decile of earnings and plots the margins of the labor demand response following a uniform 10% fall in trade costs. Panels A–B and D–F correspond to the five components of $E\bar{D}\eta \cdot 10\%$ in Proposition 1, while Panel C shows the sum of exposures to exports and import competition. Each panel reports the average, the 10th percentile, and the 90th percentile across workers in each bin.



Fig. 2. Worker-level effects of a 10% fall in trade costs on labor demand and wages, across and within income deciles. *Notes:* For the worker-level analysis of Section 4.1, Panel A plots the labor demand responses following a uniform 10% fall in trade costs, while Panel B shows general equilibrium welfare changes, defined as the equivalent variation as a fraction of initial expenditures. Each panel reports the average effects by decile of worker initial earnings, along with 10th and 90th percentiles within each decile.

decile, with 90–10 gaps of 1.9 to 4.4p.p.³⁸ On net, increases in labor demand from exposure to both export opportunities and import competition, reported in Panel C, are stronger for richer workers. This “net exports” composite effect ranges from about 0.3% on average in the bottom decile to 0.8% in the top decile, while the variation within each decile is substantial, with 90–10 gaps over 1.5p.p.³⁹

Next, Panel D shows that the increase in labor demand from the imported inputs effect is also largest in the top decile relative to the bottom (at 0.6% vs. 0.3%), again with large heterogeneity within deciles shown by the 90–10 gaps of 0.7–1.5p.p. Panel E reports changes in labor demand from income effects, which are relatively flat across deciles and close to zero on average, but vary substantially within each decile, with 90–10 gaps of 0.3–0.6p.p. Similarly, Panel F shows that changes in labor demand from substitution effects are essentially flat across the distribution, with 90–10 gaps of 0.7–0.9p.p.⁴⁰

Following Propositions 1 and 2, Fig. 2 reports the overall change in labor demand on Panel A, combining the five mechanisms from Fig. 1, and in welfare on Panel B, i.e. accounting for general equilibrium changes in both wages and price indices. Panel A shows that there is higher growth of labor demand at higher income deciles, from +0.8% at the bottom to +1.5% at the top. Again, heterogeneity between workers occurs primarily within rather than across deciles: the spread between the 10th and 90th percentiles is 2.2–4.2p.p. For the welfare change in general equilibrium on Panel B, heterogeneity in the equivalent variation is also much larger within income deciles than across. Within each decile, the 10–90 gap in welfare effects is over 2 percentage points, while variation across income deciles is much smaller, from 2.0% in the first decile to 1.8% at the tenth.⁴¹

Table 1 quantifies the relative importance of trade-induced horizontal and vertical inequality formally. Panel A shows that the shock has very heterogeneous effects across workers: while the median welfare gain is 2.29%, it is below 0.75% for 10% of workers and over four times larger, above 3.15% for another 10% of them. The standard deviation of the welfare changes is 1.44p.p. Despite this large heterogeneity in welfare gains, Panel B shows that the effect of the shock on the shape of the income distribution is very small. To measure this effect, we add the estimated welfare change to the initial (nominal) income of each worker, leveraging the cardinalization of utility from Section 2.1. We then obtain changes in the distribution of welfare, e.g. $SD(\log x + d \log W) - SD(\log x)$. The shock leaves the shape of the income distribution essentially unchanged: the Gini index falls by 0.0001, while the standard deviation of (real) log-wages falls by 0.0004. As shown by Proposition 3 the standard deviation of real log-wages remains unchanged despite large distributional effects of the shocks if the magnitude of welfare gains does not covary with the initial level of income. Accordingly, we find that the standard deviation of welfare effects is 36 times larger than the change in the standard deviation of

³⁸ The finding that high-earning workers are on average more exposed to import competition contrasts with the traditional two-sector, two-factor formulation of the Heckscher–Ohlin model, in which low-paid workers are more exposed to import competition and lose from trade. Instead, our results highlight the importance of heterogeneous trade costs across industries: high-earning workers (except the very top earnings decile) are more likely to be employed in the more tradable manufacturing sector (see Appendix Figure S1), as well as in more tradable industries within both manufacturing and services and are therefore more exposed to import competition.

³⁹ We put net exports in quotes because, as Footnote 24 explains, the relevant industry exposure measure from Proposition 1 which is reported here is not based on the difference between the values of exports and imports.

⁴⁰ We find that the patterns in Fig. 1 are driven primarily by the heterogeneity of worker exposure to export ratios, import penetration, cost shares of intermediate inputs, and income elasticities of their industries, rather than by IO and other adjustments from Proposition 1. We show this result in Appendix Figure S2, which reports “raw” exposure statistics and finds patterns similar to Fig. 1, both within and across income groups.

⁴¹ Given the importance of horizontal inequality, one may ask whether its magnitude is particularly large for certain income groups. We analyze this possibility in Appendix Figure S3 but we find no systematic pattern.

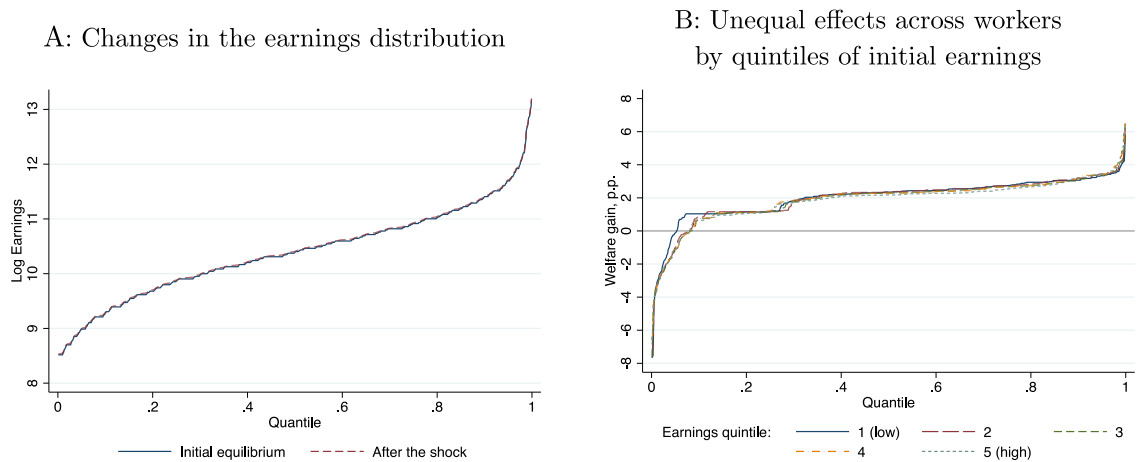


Fig. 3. Distributions of worker-level effects of a 10% fall in trade costs. Notes: For the worker-level analysis of Section 4.1, this figure plots the change in the (real) earnings distribution (Panel A) and the distribution of welfare changes by quintile of worker initial earnings (Panel B), following a uniform 10% fall in trade costs. Welfare changes are defined as the equivalent variation as a fraction of initial expenditures. A worker's log-earnings after the trade shock is defined as the initial (nominal) log-earnings plus the welfare change, leveraging the cardinalization of utility from Section 2.1. In both panels, the distributions are reported as quantile functions, i.e. with quantiles on the horizontal axis.




Fig. 4. Decomposing the effects of a 10% fall in trade costs on labor demand across income deciles. Notes: For the worker-level analysis of Section 4.1, this figure plots the average labor demand responses following a uniform 10% fall in trade costs across deciles of worker initial earnings. The bottom decile is normalized to zero and the overall change relative to the bottom decile is decomposed into different mechanisms according to Proposition 1.

the log-income distribution. In this sense, a fall in trade costs primarily generates horizontal inequality. Similarly, using Eq. (10), we estimate that only 0.20% of the variance of worker-level welfare changes from trade shocks arises across income deciles.⁴²

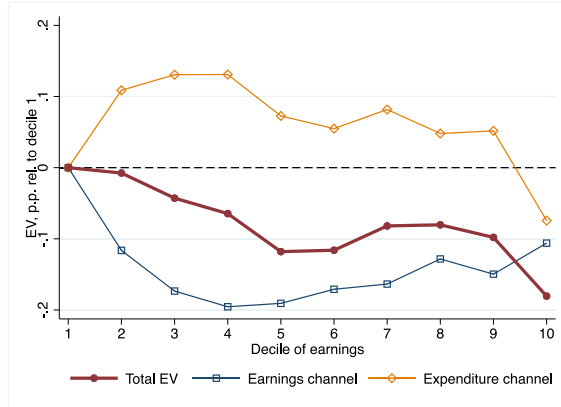
Fig. 3 depicts these patterns graphically. Panel A reports the distribution of real wages before and after the 10% fall in trade costs. It shows that the wage distribution remains essentially unchanged, even in the tail. In contrast, Panel B shows that some workers experience substantial changes in welfare. Plotting the distribution of the worker-level welfare changes by quintile of their initial earnings, this panel shows no differences across quintiles, including in the tail of the distributions. As in Panel B of Fig. 2, about 80% of workers in each income quintile experience welfare gains ranging between 1% and 3%. The welfare gains or losses become large in the tails of the distribution, with welfare losses over 7% at the bottom and welfare gains over 6% at the top, which are observed for all income groups. The full distribution of welfare changes confirms the finding from Table 1: the distributional effects of trade can be large but are entirely driven by heterogeneity across workers within the same income group.

4.1.2. Decomposing trade-induced vertical inequality


Although we found that trade shocks primarily generate horizontal rather than vertical inequality, it is instructive to decompose vertical inequality to understand the mechanisms at play.

⁴² To contextualize our results on counterfactual trade shocks, we consider typical income changes in the Panel Study of Income Dynamics (PSID) between 1999 and 2005. We find that the share of the variance in 1999–2005 log income changes explained by 1999 log income is 19.6%. The share increases to 21.5% when using 1999 income deciles as predictors. Much of the predictive power of initial income is tied to mean reversion in income. These findings suggest that trade shocks may generate relatively more horizontal inequality than other economic forces. Appendix B.4 provides details on our PSID analysis.

A: Decomposition of welfare into earnings and expenditure channels



B: Decomposition of the earnings channel into mechanisms



C: Decomposition of the expenditure channel into mechanisms



Fig. 5. Decomposing the effects of a 10% fall in trade costs on welfare across income deciles. Notes: For the worker-level analysis of Section 4.1, this figure plots the average welfare responses following a uniform 10% fall in trade costs across deciles of worker initial earnings, with the bottom decile normalized to zero. Welfare changes are defined as the equivalent variation as a fraction of initial expenditures. Panel A decomposes welfare effects into the expenditure and earnings channels using Eq. (3). Panel B decomposes the earnings channel into four mechanisms using Eq. (6) in Proposition 1, while Panel C decomposes the expenditure channel into three mechanisms, using Eq. (9) in Proposition 2.




Fig. 6. Unequal effects across workers by sector. Notes: For the worker-level analysis of Section 4.1, this figure plots the distribution of worker-level welfare changes by sector (goods or services) following a uniform 10% fall in trade costs. Welfare changes are defined as the equivalent variation as a fraction of initial expenditures, and are reported by quantiles.

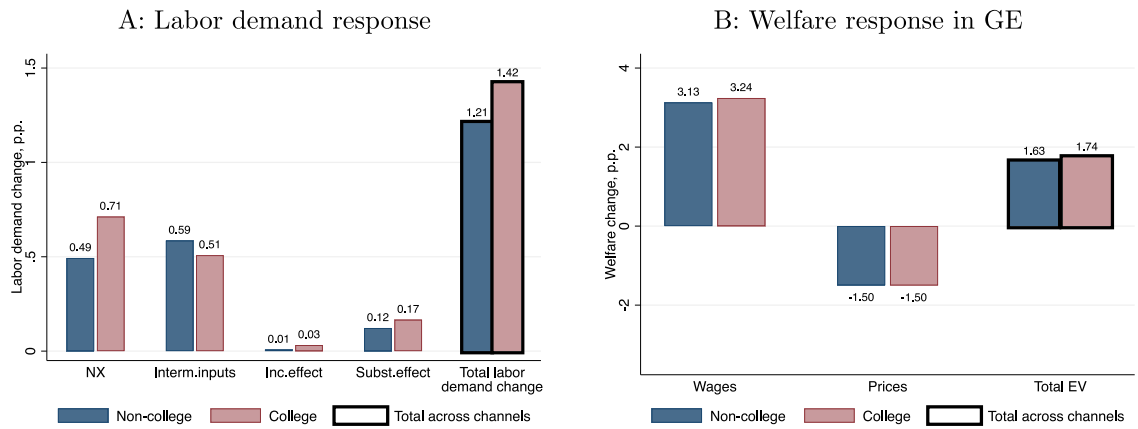


Fig. 7. Welfare effects of a 10% fall in trade costs across education groups. *Notes:* For the analysis across education groups of Section 4.2, this figure plots the partial equilibrium labor demand response (Panel A) and the welfare response in GE (Panel B) to a uniform 10% reduction in trade costs for workers with and without a college degree. Welfare changes are defined as the equivalent variation as a fraction of initial expenditures. Panel A decomposes the effects into several mechanisms according to Proposition 1, and Panel B uses Eq. (3) for the decomposition.

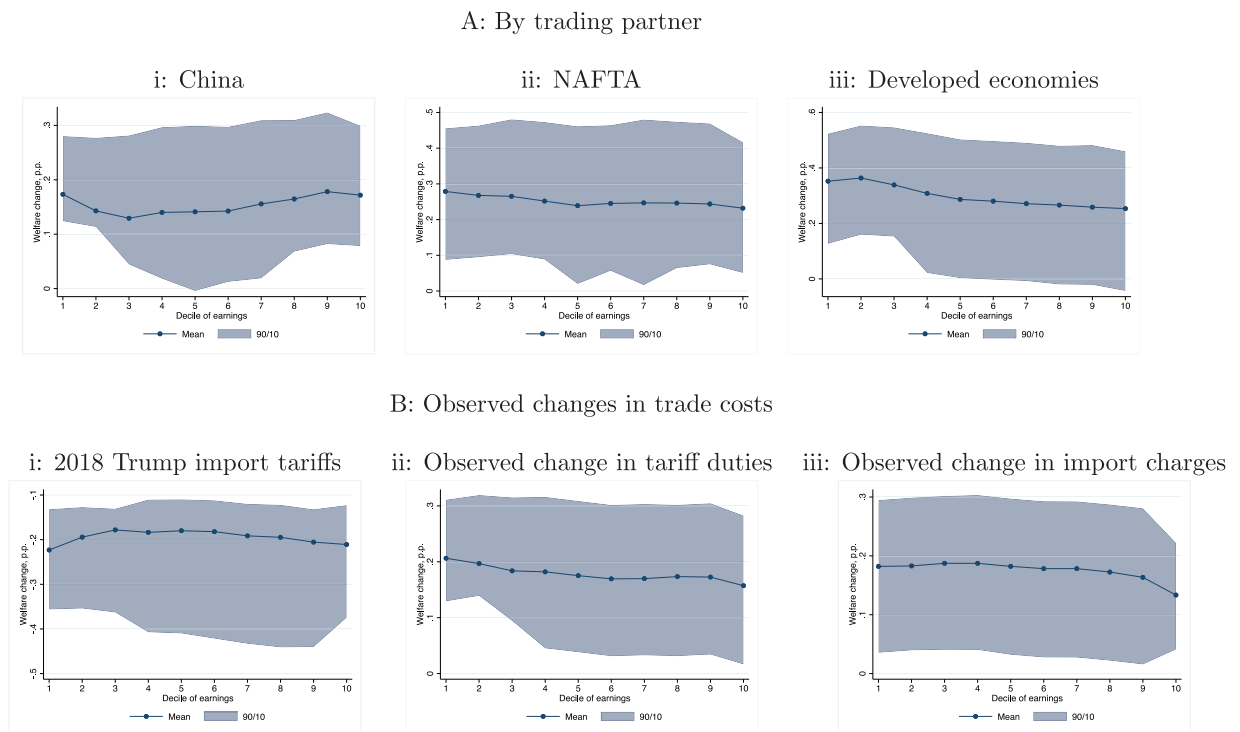


Fig. 8. Worker-level welfare effects of non-uniform trade shocks. *Notes:* For the worker-level analysis of Section 4.1, Panel A plots the welfare effects of a 10% fall in trade costs for goods imported from specific trading partners. NAFTA corresponds to Canada and Mexico, while 34 developed economies are OECD members, excluding NAFTA, plus Taiwan and Singapore. Panel B plots the welfare effects from observed shocks to the costs of importing goods: the introduction of import tariffs in 2018 by the Trump administration, changes in U.S. tariff duties between 1992 and 2007, and changes in import charges (i.e., total transportation and insurance costs) between 1992 and 2007; see Appendix B.5 for details. Each panel reports average effects by decile of initial earnings, along with the 10th and 90th percentile within each decile. Panel A follows Proposition 1, while Appendix B.5 describes the methodology for Panel B.

First, it is notable that, in contrast to the change in labor demand in Panel A of Fig. 2, which is higher at the top of the income distribution, the average welfare gains in Panel B are slightly *higher* at the bottom of the income distribution. The change in slope when accounting for the \bar{G} matrix is explained by the role of the service sector. As Section 3.3.1 showed, when labor demand grows, less traded industries, such as services, experience a larger increase in wages. Services also tend to have relatively more lower-income workers, as Appendix Figure S1 reports. Thus, this larger wage response benefits the low-income group more.

Table 1
Unequal effects across workers vs. Changes in the income distribution.

A: Unequal effects of the shock across workers					
	SD (1)	p10 (2)	p50 (3)	p90 (4)	
Welfare change, p.p.	1.42	0.75	2.29	3.15	
B: Effects of the shock on the income distribution					
	SD(log wage) (1)	p10 (2)	p50 (3)	p90 (4)	Gini index (5)
Initial income level	0.8230	10,700	32,500	90,000	0.4509
Counterfactual	0.8226	10,835	33,101	90,563	0.4508
Change	-0.0004	+1.27%	+1.85%	+0.63%	-0.0001

Notes: Panel A reports statistics of the distribution of welfare changes across workers after a uniform 10% fall in trade costs in the worker-level analysis of Section 4.1. Panel B shows how the same shocks affects the distribution of (real) incomes, by reporting statistics for two income distributions: the one observed in the data and the counterfactual one, in which the estimated welfare effects are added to each worker's initial earnings. Both panels show the standard deviation and 10th, 50th, and 90th percentiles, while Panel B additionally reports Gini indices.

Table 2
Accounting for the distributional effects of trade through exposure.

	Labor demand change		Wage change		Wage change	
	Coef.	R^2 shapley decomp., %	Coef.	Coef.	R^2 shapley. decomp., %	
	(1)	(2)	(3)	(4)	(5)	
Exports	1.000	54.43		0.310 (0.025)	12.19	
Import competition	1.000	21.26		0.800 (0.029)	50.88	
Intermediate inputs	1.000	19.19		0.702 (0.095)	13.99	
Income effects	1.000	1.86		0.957 (0.035)	4.69	
Substitution effects	1.000	3.6		1.171 (0.081)	18.25	
Labor demand change, all five mechanisms			0.372 (0.036)			
R^2	1.000		0.385	0.959		

Notes: For the worker-level analysis of Section 4.1, this table reports ordinary least square regressions of the labor demand response (column (1)) and the wage response (columns (3), and (4)) to a uniform 10% fall in trade costs. The explanatory variables in columns (1) and (4) are the components of the labor demand response corresponding to each term in Eq. (6). In column (3) the explanatory variable is the total labor demand response. Columns (2) and (5) report the Shapley decomposition of the R^2 , in percent of the total regression R^2 . The analysis is implemented using the 1,126,226 workers in the ACS data. Column (1) suppresses standard errors because of the perfect fit; columns (3) and (4) show standard errors clustered by worker's industry.

Next, we investigate the economic forces driving the differences in the labor demand and wage effects of the shock across deciles of initial earnings. Starting with the change in labor demand, i.e., $E\tilde{D}\eta \cdot (-d \log \tau)$, Fig. 4 reports the contribution of the mechanisms corresponding to the components of η from Eq. (6), normalizing the bottom decile to zero. The key differences arise in the exposure to net exports and intermediate inputs, which both favor richer workers. Income and substitution effects do not play a significant role.

The decomposition for welfare changes in GE is richer, following Eq. (3) and Proposition 2. Panel A of Fig. 5 decomposes welfare changes in GE, again relative to the first decile, into the earnings and expenditure channels. The panel shows that the earnings and expenditure channels both have non-monotonic welfare effects across deciles. The expenditure channel generates the largest welfare gains for the third decile of the income distribution, and the smallest gains for the top income decile. In contrast, the earnings channel generates the largest welfare gains for the bottom income decile, and the smallest for the fourth decile. Together, the two channels generate a near-monotonic fall in welfare effects across deciles, driven by the earnings channel at the bottom and the expenditure channel at the top. To understand these patterns, Panels B and C of Fig. 5 further decompose the earnings and expenditure channels, respectively.

Panel B decomposes the wage response in GE into the mechanisms from Eq. (6). Compared with the partial equilibrium labor demand response in Fig. 4, substitution effects now play a much larger role than in the partial equilibrium analysis of labor demand, contributing to a fall in wage inequality. Intuitively, substitution effects compound the larger increase in wages in general equilibrium for less traded industries like services: since the elasticity of substitution between goods and services is below one,




Fig. 9. Export intensity at the firm and industry level by worker earnings decile. *Notes:* This figure displays the average firm- and industry-level export intensity for workers in each decile of the earnings distribution, as measured by the LEHD in 2007. Panel A displays all firms in the LEHD and Panel B focuses only on firms with manufacturing as the main sector. For each decile, we calculated a weighted average of the export statistic, where the weights are the earnings of a given worker. Industry-level export shares are calculated by taking a weighted average across all firms in the NAICS-6 industry, where the weights are total payroll for the firm. Throughout, we only use data from individuals with recorded wages in all four quarters of 2007. Source: authors' calculations based on U.S. Census data (FSRDC Project Number 2229, Disclosure Review Board (DRB) approval number: CBDRB-FY24-0037).

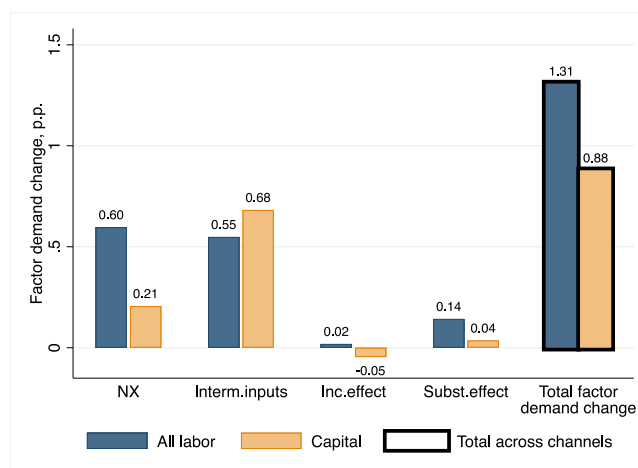


Fig. 10. Changes in factor demand for capital owners and workers for a uniform fall in trade costs. *Notes:* This figure reports the partial equilibrium change in factor demand for labor and capital for a uniform 10% fall in trade costs, decomposing the change into the several mechanisms as in Proposition 1. The composition of industries in payments to capital owners is obtained from the “Gross operating surplus” row in the IO Table, similarly to how “Compensation of employees” is used for labor in the main industry-level dataset.

consumers reallocate spending toward services as wages in this sector increase, leading to a further increase in wages for service workers.

Finally, Panel C of Fig. 5 decomposes the expenditure channel into three mechanisms following Eq. (9) in Proposition 2, isolating the contribution of the price of imports through falling trade costs, the change in average wage, and the segregation effect. This panel shows that the lower welfare gain for higher-income workers via the expenditure channel is explained primarily by the segregation effect: domestic wages increase more in sectors that cater to richer households.⁴³ The other two mechanisms push slightly in the opposite direction. Compared with the bottom income decile, the fall in prices from lower trade costs benefits richer workers slightly more because the import share of their consumption baskets, including both direct and indirect imports, is slightly higher, as shown by Borusyak and Jaravel (2023).⁴⁴ Prices also change because average domestic wages increase in general equilibrium, through the terms-of-trade effect. The response of domestic prices to this change is proportional to and smaller than the effects of lower trade

⁴³ In unreported analyses, we find that the segregation effect stems from two forces: first, wages (and thus domestic prices) grow more in services, whose expenditure share is higher for richer workers; second, within services domestic prices grow more for those services consumed more by richer workers.

⁴⁴ Specifically, Borusyak and Jaravel (2023) show that import shares are similar across the distribution of household incomes, with a slight increase from the bottom to the top income deciles. As explained in Appendix B.1, we match household income deciles to worker earnings deciles.

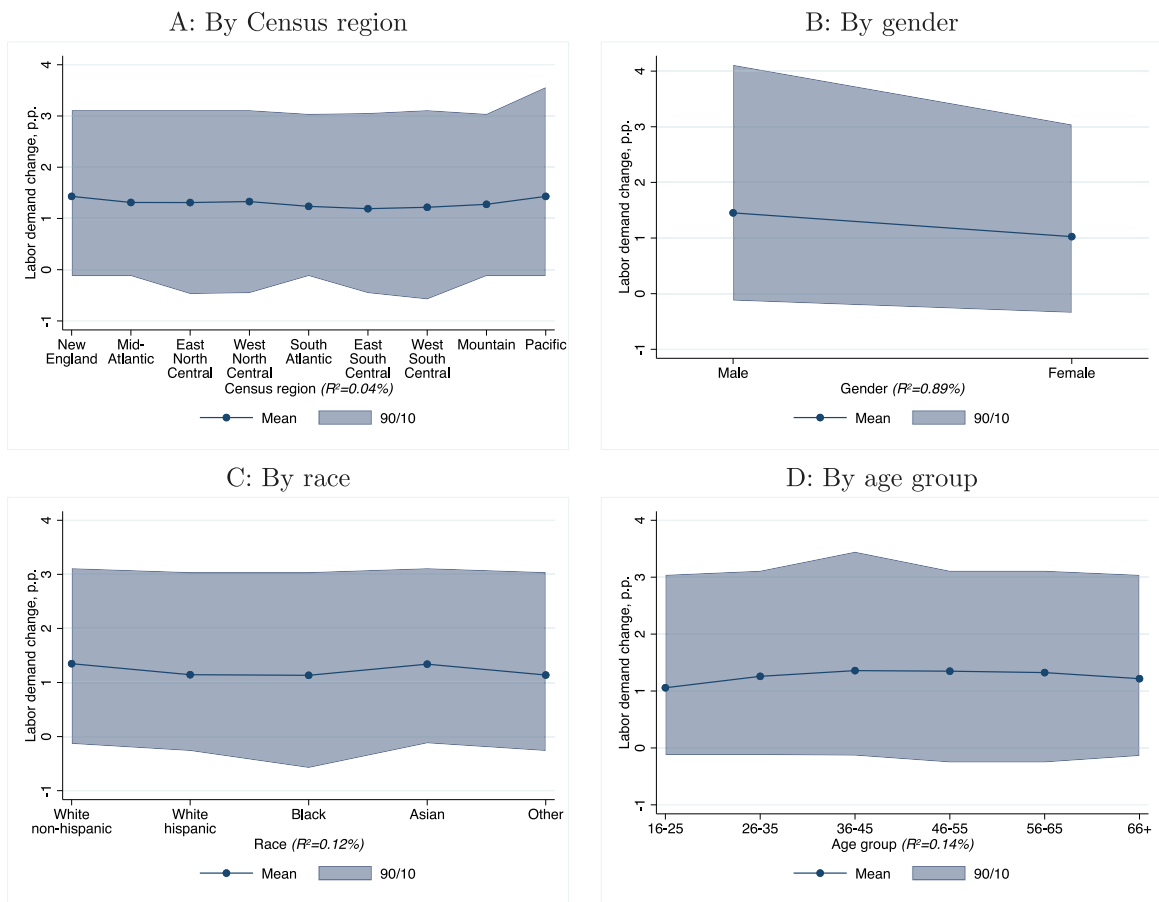


Fig. 11. Exposure to a fall in trade costs by socio-demographic groups. *Notes:* This figure reports the average, the 10th and the 90th percentile of the worker-level labor demand change in response to a 10% fall in trade costs, grouping workers by socio-demographic characteristics in the ACS data. The R^2 reported in parentheses corresponds to the fraction of variance in labor demand responses across workers explained by the socio-demographic group dummies.

costs, with a limited impact on inequality across deciles. Overall, absent the segregation effect, the welfare gains of falling trade costs would be nearly equal across income deciles, instead of being larger at the bottom.

4.1.3. *Decomposing trade-induced horizontal inequality*

In this section, we use several decompositions to understand the horizontal inequality induced by falling trade costs. We document the role of sectoral exposure for the heavy left tail of worker-level welfare changes and use Proposition 1 to decompose the variance of worker-level distributional effects.

We first assess the role played by sectoral differences in the tails of the distribution of welfare changes. Fig. 6 depicts the worker-level distribution of welfare changes within goods-producing industries and within services. The panel shows that the large welfare changes in the tails are driven by goods. While workers in services all experience welfare gains, within goods-producing industries about 60% of workers experience a welfare loss, including 5% with a welfare loss in excess of 4%. Large welfare gains are also more common for goods-producing industries (e.g. due to the growing export demand), with about 5% of workers experiencing a welfare gain above 5%, which is higher than for any worker within services. Sectoral differences thus play a key role for the tails of worker-level welfare changes.

Next, we characterize the relative importance of different economic forces for the heterogeneity of labor demand and wage effects of trade across workers. In Column (1) of Table 2, we regress the worker-level labor demand response on the five margins of exposure to the labor market effects of trade: exports, import competition, intermediate inputs, income effects, and substitution effects. This analysis parallels Fig. 4, but the regression is estimated across workers rather than income deciles to focus on the horizontal inequality induced by trade. In our model, the heterogeneous labor demand effects across workers are fully explained by the five margins of their industry’s exposure. Therefore, the regression coefficients in Column (1) are all equal to 1 and the R^2 is 100%. Column (2) uses the Shapley decomposition of the R^2 to assess the relative importance of each mechanism in explaining the overall variation. Exports explain 54% of the variation, imports 21% and intermediate inputs 19%, while income effects and substitution effects play more minor roles. These results thus suggest that export exposure is the key driving force.

In general equilibrium, the wage effects also depend on the \tilde{G} matrix. We therefore run a regression to assess whether the labor demand response is a good predictor of the overall wage effect, at the worker level. Column (3) of Table 2 reports the result: the R^2 falls to 39%, indicating that the overall exposure to trade shocks is an imperfect predictor of the wage effect. However, Column (4) shows that the R^2 increases to 96% when using the five exposure margins from Column (1) as separate regressors, allowing for different regression coefficients for each margin. The estimated coefficients now range from 0.31 for exports to 1.17 for substitution effects.⁴⁵ This result shows that the earnings channel of trade in general equilibrium can be almost perfectly recovered using as linear predictors the margins of exposure we study, as long as the coefficients on the different margins are allowed to vary.⁴⁶ These findings speak to a growing empirical literature on attitudes toward trade (e.g., Mayda and Rodrik, 2005; Jäkel and Smolka, 2017, Van Patten and Méndez-Chacón, 2022), which assesses whether individual-level exposure to trade predicts support for trade agreements, typically using linear regressions. Our analysis provides a novel foundation for these analyses, since we find that partial equilibrium exposure measures retain very strong predictive power in general equilibrium in a standard, fully-specified quantitative trade model.

4.1.4. Takeaways

The analysis in this section yields three main lessons. First and foremost, the distributional effects of trade shocks are primarily concentrated within income deciles, i.e. trade-induced inequality is primarily horizontal. There is little impact of a fall in trade costs on the shape of the income distribution, but there are substantial distributional effects creating sizeable changes in relative wages, as well as winners and losers at all income levels. This finding is not a mechanical feature of the model but results from the fact that the welfare effects of trade shocks are only weakly correlated with income. If specialization patterns had been sufficiently different across income deciles, we could have found an effect across deciles as large as the effect obtained within deciles.⁴⁷ Second, the average gains from trade liberalizations are positive for all income deciles. Third, direct exposure measures turn out to be sufficient to closely approximate the overall welfare effects of trade across workers, offering new a justification for the growing literature studying attitudes toward trade depending on individual exposure statistics.

4.2. Robustness: Trade-induced inequality across education groups

To verify that there is a robust pattern of weak distributional effects across groups of observably similar workers in general equilibrium, we now conduct the analysis with two groups, considering workers with or without a college degree and assuming they are freely mobile across industries. Our focus is therefore on the college wage premium, which has played an important role in the evolution of U.S. income inequality (e.g. Autor et al., 2008; Goldin and Katz, 2007). Fig. 7 reports the effects of a 10% reduction in trade costs in this setting.

Using Proposition 1, Panel A reports shifts in labor demand and their drivers across education groups. We find that labor demand grows by more for college graduates, mainly because they are employed in industries with higher “net exports”. Favorable income and substitution effects magnify the difference slightly, while exposure to imported inputs is lower for college graduates, which partially offsets the gap. In total, labor demand grows by 1.4% for the group of college graduates and 1.2% for workers without a college degree in response to the shock.

Panel B reports welfare changes across education groups in general equilibrium using Propositions 1 and 2. We find that both groups benefit from reduced trade costs and the college wage premium remains almost unchanged. The equivalent variation is 1.7% for college-educated workers, compared with 1.6% for those without a college degree; the small difference of 0.11p.p. arises from the earnings channel. The expenditure channel, also reported in this figure, is distributionally neutral.

Taken together, the results of our two analyses – across workers and across education groups – show that trade can induce sizeable horizontal inequality when labor mobility is limited; vertical inequality is not found either with or without labor mobility. These findings also illustrate how our theoretical results can be used to assess the importance of different mechanisms and different labor market assumptions in governing the distributional effects of trade shocks.

4.3. Extensions and further robustness

In this section, we consider several extensions, allowing for other counterfactual shocks, within-industry heterogeneity, and capital as a separate factor, and analyzing the heterogeneity in impacts by worker characteristics other than income. We also demonstrate the robustness of the results to different choices of elasticities.

⁴⁵ The heterogeneous coefficients imply that the relative predictive importance of different margins is very different from the labor demand analysis of Column (2). Using the Shapley decomposition of the R^2 , Column (5) shows that exposure to exports explains only 12.2% of the variation, consistent with the lower coefficient on this margin, while import competition now explains 50.9%, followed by substitution effects at 18.3%, and intermediate inputs at 14.0%. Income effects are still relatively unimportant, at 4.7%.

⁴⁶ This result is not a mechanical feature of the model since, for each worker, the welfare effect is determined by the exposure of their own industry with industry-specific coefficients (diagonal elements of \tilde{G}) as well as exposure of other industries (non-diagonal elements). Appendix Table S1 shows that the five margins of labor demand exposure very accurately predict not only the overall wage response to a trade shock but also each of the five margins of the wage response separately, with the heterogeneous coefficients shown in the table. Thus, one can approximately view the mapping \tilde{G} from labor demand to wage changes as a rotation of the five mechanisms through which labor demand shocks happen in response to a trade shock.

⁴⁷ This first lesson from our analysis echoes the empirical findings of Hummels et al. (2014) who estimate the effects of exports and offshoring on wages of different groups of workers in a reduced-form framework. The economic mechanisms they study are different, as our framework does not incorporate offshoring (although it can be introduced by modeling it as skill-biased import competition, as in our early draft, Borusyak and Jaravel, 2018). They find that the distributional effects of globalization arise primarily *within* groups of ex-ante similar workers because of their heterogeneous exposure (Table 6).

Non-uniform changes in trade costs. We follow [Borusyak and Jaravel \(2023\)](#) and consider reductions of trade costs with specific trading partners, as well as counterfactuals inspired by recent changes in trade policy and trade costs. Panel A of [Fig. 8](#) analyzes a 10% fall in iceberg costs for imports from China, NAFTA or 34 advanced economies separately, for the worker-level analysis. Panel B investigates the impact of the import tariffs introduced by the Trump administration in 2018–2019 (on solar panels, washing machines, steel and aluminum products, and a large set of products from China), the observed change in U.S. import tariffs in 1992–2007, and the observed change in transportation and insurance costs in the same period.⁴⁸ Appendix Figure S4 repeats the same analyses across education groups. The results are similar to the baseline analyses: substantial distributional effects of trade are found only within income deciles in the absence of labor mobility.

Within-industry heterogeneity. To assess the potential importance of within-industry heterogeneity, in [Fig. 9](#) we use the plant-level microdata from the Census of Manufactures and the Longitudinal Employer-Household Dynamics (LEHD) dataset. These data allow us to analyze, at a more granular level, one of the mechanisms from [Proposition 1](#): the difference in exposure to exports between workers in different earnings deciles. We find that higher-earnings workers tend to be employed at firms that export more within the same industry, in line with [Burstein and Vogel \(2017\)](#). However, this within component is smaller than the difference arising across industries, which we analyzed previously. This result holds both across all sectors (Panel A) and in manufacturing specifically (Panel B).

Relative factor demand for capital and labor. [Fig. 10](#) documents changes in the partial equilibrium factor demand for capital vs. labor after a fall in trade costs, quantifying all mechanisms from [Proposition 1](#): exposure to net exports, intermediate inputs, income, and substitution effects. We find that relative factor demand remains relatively similar after a uniform fall in trade costs.

The role of other worker characteristics. [Fig. 11](#) compares the partial equilibrium change in labor demand in response to a trade shock across several socio-demographic characteristics of workers other than income and education. Specifically, we consider locations (Census regions), gender, race, and age groups. We find the demand for male workers increasing on average after a trade liberalization, relative to females; yet, gender explains only 0.9% of the variance in labor demand impacts across workers. There are no notable differences across other demographics.

Additionally, we find that the most detailed occupation category codes in ACS explain 15% of the variance in the labor demand effects of the trade shocks, i.e. the majority of the variance remains unexplained because the same occupation exists in industries with very different trade exposure.

Robustness of main results to choice of elasticities. Finally, Appendix Figure S5 shows that the welfare effects of the uniform 10% shock remain similar in both our quantitative analyses when we vary the trade elasticity $\xi - 1$, substitution elasticities in demand (ρ and ε) or the labor substitution elasticity σ_{macro} within the ranges used in the literature. Since exposure to trade is similar across worker groups, elasticities do not play a decisive role.

5. Conclusion

In this paper, we conceptualized the notion of trade-induced horizontal and vertical inequality, and studied their relative importance in the United States. We accounted for changes in both prices and wages in a unified general equilibrium framework and found empirically that horizontal inequality is the dominant force, because exposure to trade is heterogeneous primarily within groups of ex-ante similar workers. Since trade shocks generate winners and losers at all income levels but do not change the shape of the income distribution, our findings run against a popular narrative that “trade wars are class wars” ([Klein and Pettis, 2020](#)). Our exposure-based theoretical framework contributes to an emerging literature characterizing the welfare effects of trade or other macroeconomic shocks in terms of microeconomic sufficient statistics and elasticities (e.g., [Baqaee and Farhi, 2019, 2024](#); [Adão et al., 2022b](#); [Kleinman et al., 2022](#)).

Our framework can be readily applied to study the effects of other counterfactual shocks of interest, beyond trade shocks. For example, the framework can accommodate any productivity shock (e.g., carbon abatement costs due to the net-zero transition) or labor supply shock (e.g., shocks to education or immigration). Expanding the framework to account for the role of local labor markets and within-country trade would be another interesting direction for further research. Finally, a fruitful area of inquiry would be to link our concepts of horizontal and vertical inequality to studies of income dynamics and persistence (e.g., [Auten and Gee, 2009](#); [Arellano et al., 2017](#); [Güvenen et al., 2021](#)), building on the analysis from our footnote 42. For example, one could estimate the extent to which realized trade shocks contribute to overall horizontal inequality, i.e. all income changes unrelated to initial income. One could also assess whether other economic shocks are similar to trade shocks in that they generate much more horizontal than vertical inequality. Our framework thus offers several avenues for novel applications and extensions in future work.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

⁴⁸ Appendix B.5 describes the data sources, and Appendix A.4 explains how to apply [Propositions 1](#) and [2](#) to shocks which are not uniform across industries.

Data availability

<https://data.mendeley.com/datasets/vrzbmvgnw/1>.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jinteco.2024.103935>.

References

- Acemoglu, Daron, 2002. Technical change, inequality, and the labor market. *J. Econ. Lit.* 40 (1), 7–72.
- Acemoglu, Daron, Autor, David H., 2011. Skills, tasks and technologies: Implications for employment and earnings. In: *Handbook of Labor Economics*, vol. 4, pp. 1043–1171.
- Acemoglu, Daron, Autor, David H., Dorn, David, Hanson, Gordon H., Price, Brendard, 2016. Import competition and the great U.S. employment sag of the 2000s. *J. Labor Econ.* 34 (1), S141–S198.
- Acemoglu, Daron, Carvalho, Vasco, Ozdaglar, Asuman, Tahbaz-Salehi, Alireza, 2012. The network origins of aggregate fluctuations. *Econometrica* 80 (5), 1977–2016.
- Adão, Rodrigo, Arkolakis, Costas, Esposito, Federico, 2022a. General Equilibrium Effects in Space: Theory and Measurement. Working Paper.
- Adão, Rodrigo, Carrillo, Paul, Costinot, Arnaud, Donaldson, Dave, Pomeranz, Dina, 2022b. Imports, exports, and earnings inequality: Measures of exposure and estimates of incidence. *Q. J. Econ.*
- Aguiar, Mark, Bils, Mark, 2015. Has consumption inequality mirrored income inequality? *Amer. Econ. Rev.* 105 (9), 2725–2756.
- Antràs, Pol, de Gortari, Alonso, Itskhoki, Oleg, 2017. Globalization, inequality and welfare. *J. Int. Econ.* 108, 387–412.
- Arellano, Manuel, Blundell, Richard, Bonhomme, Stéphane, 2017. Earnings and consumption dynamics: a nonlinear panel data framework. *Econometrica* 85 (3), 693–734.
- Arkolakis, Costas, Costinot, Arnaud, Rodríguez-Clare, Andrés, 2012. New trade models, same old gains? *Amer. Econ. Rev.* 102 (1), 94–130.
- Armington, Paul S., 1969. A theory of demand for products distinguished by place of production. *IMF Staff Pap.* 16 (1), 159–178.
- Atkin, David, Faber, Benjamin, Gonzalez-Navarro, Marco, 2018. Retail globalization and household welfare: Evidence from Mexico. *J. Polit. Econ.* 126 (1), 1–73.
- Atkinson, A.B., 1970. On the measurement of inequality. *J. Econom. Theory* 2, 244–263.
- Auten, Gerald, Gee, Geoffrey, 2009. Income mobility in the United States: New evidence from income tax data. *Natl. Tax J.* 62 (2), 301–328.
- Autor, David H., Dorn, David, Hanson, Gordon H., 2013. The China syndrome: Local labor market impacts of import competition in the United States. *Amer. Econ. Rev.* 103 (6), 2121–2168.
- Autor, David H., Katz, Lawrence F., Kearney, Melissa S., 2008. Trends in u. s. Wage inequality: Revising the revisionists. *Rev. Econ. Stat.* 90 (2), 300–323.
- Baqee, David Rezza, Farhi, Emmanuel, 2019. JEEA-FBBVA lecture 2018: The microeconomic foundations of aggregate production functions. *J. Eur. Econom. Assoc.* 17 (5), 1337–1392.
- Baqee, David Rezza, Farhi, Emmanuel, 2024. Networks, barriers, and trade. *Econometrica* 92 (2), 505–541.
- Borusyak, Kirill, Jaravel, Xavier, 2018. The Distributional Effects of Trade: Theory and Evidence from the United States. Mimeo.
- Borusyak, Kirill, Jaravel, Xavier, 2023. The Distributional Effects of Trade: Theory and Evidence from the United States. SSRN Working Paper 3269579.
- Broda, Christian, Weinstein, David E., 2006. Globalization and the gains from variety. *Q. J. Econ.* 121 (2), 541–585.
- Burstein, Ariel, Hanson, Gordon H., Tian, Lin, Vogel, Jonathan, 2020. Tradability and the labor-market impact of immigration: Theory and evidence from the U.S.. *Econometrica* 88 (3), 1071–1112.
- Burstein, Ariel, Vogel, Jonathan, 2017. International trade, technology, and the skill premium. *J. Polit. Econ.* 125 (5), 1356–1412.
- Caliendo, Lorenzo, Dvorkin, Maximiliano, Parro, Fernando, 2019. Trade and labor market dynamics: General equilibrium analysis of the China trade shock. *Econometrica* 87 (3), 741–835.
- Caliendo, Lorenzo, Parro, Fernando, 2015. Estimates of the trade and welfare effects of NAFTA. *Rev. Econ. Stud.* 82 (1), 1–44.
- Caron, Justin, Fally, Thibault, Markusen, James R., 2014. International trade puzzles: A solution linking production and preferences. *Q. J. Econ.* 129 (3), 1501–1552.
- Caron, Justin, Fally, Thibault, Markusen, James R., 2020. Per capita income and the demand for skills. *J. Int. Econ.* 123 (103306).
- Chaney, Thomas, 2008. Distorted gravity: The intensive and extensive margins of international trade. *Amer. Econ. Rev.* 98 (4), 1707–1721.
- Clemens, Jeffrey, Gottlieb, Joshua D, Olsen, Morten, David, H, 2019. The Spillover Effects of Top Income Inequality. Mimeo.
- Comin, Diego, Lashkari, Danial, Mestieri, Martí, 2021. Structural change with long-run income and price effects. *Econometrica* 89 (1), 311–374.
- Costinot, Arnaud, Rodríguez-Clare, Andrés, 2015. Trade theory with numbers: Quantifying the consequences of globalization. In: *Handbook of International Economics*, vol. 4, pp. 197–261.
- Costinot, Arnaud, Vogel, Jonathan, 2009. Matching and inequality in the world economy. *J. Polit. Econ.* 118 (4), 747–786.
- Costinot, Arnaud, Werning, Ivan, 2019. Lerner symmetry: A modern treatment. *Am. Econ. Rev.: Insights* 1 (1), 13–26.
- Cravino, Javier, Sotelo, Sebastian, 2019. Trade induced structural change and the skill premium. *Am. Econ. J.: Macro* 11 (3), 289–326.
- Dawkins, Christina, Srinivasan, T.N., Whalley, John, 2001. Calibration. *Handb. Econometr.* 5, 3653–3703.
- Deardorff, Alan V., Staiger, Robert W., 1988. An interpretation of the factor content of trade. *J. Int. Econ.* 24 (1–2), 93–107.
- Dekle, Robert, Eaton, Jonathan, Kortum, Samuel, 2008. Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Pap.* 55 (3), 511–540.
- Demidova, Svetlana, Rodríguez-Clare, Andrés, 2009. Trade policy under firm-level heterogeneity in a small economy. *J. Int. Econ.* 78, 100–112.
- Demidova, Svetlana, Rodríguez-Clare, Andrés, 2013. The simple analytics of the Melitz model in a small economy. *J. Int. Econ.* 90 (2), 266–272.
- Dixon, Peter B., Koopman, Robert B., Rimmer, Maureen T., 2013. The MONASH style of computable general equilibrium modeling: a framework for practical policy analysis. In: *Handbook of Computable General Equilibrium Modeling*, vol. 1, Elsevier, pp. 23–103.
- Eaton, Jonathan, Kortum, Samuel, 2002. Technology, geography, and trade. *Econometrica* 70 (5), 1741–1779.
- Faber, Benjamin, 2014. Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices. Mimeo.
- Fajgelbaum, Pablo, Khandelwal, Amit K., 2016. Measuring the unequal gains from trade. *Q. J. Econ.* 131 (3), 1113–1180.
- Feenstra, Robert, Luck, Philip, Obstfeld, Maurice, Russ, Katheryn, 2018. In search of the armington elasticity. *Rev. Econ. Stat.* 100 (1), 135–150.
- Galle, Simon, Rodríguez-Clare, Andrés, Yi, Moises, 2022. Slicing the pie: Quantifying the aggregate and distributional effects of trade. *Rev. Econ. Stud.*
- Garner, Thesia I., McClelland, Robert, Passero, William, 2009. Strengths and Weaknesses of the Consumer Expenditure Survey from a BLS Perspective. Mimeo.
- Goldin, Claudia, Katz, Lawrence, 2007. Long-run changes in the wage structure: Narrowing, widening, polarizing. *Brook. Pap. Econ. Act.* 38 (2), 135–168.
- Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar, Song, Jae, 2021. What do data on millions of US workers reveal about lifecycle earnings dynamics? *Econometrica* 89 (5), 2303–2339.
- Hanoch, Giora, 1975. Production and demand models with direct or indirect implicit additivity. *Econometrica* 43 (3), 395–419.
- He, Zheli, 2020. Trade and Real Wages of Skilled and Unskilled Workers: Cross-Region Evidence. Mimeo.

- Helpman, Elhanan, 2018. *Globalization and Inequality*. Harvard University Press.
- Helpman, Elhanan, Itzhakovi, Oleg, Muendler, Marc-Andreas, Redding, Stephen J., 2017. Trade and inequality: From theory to estimation. *Rev. Econ. Stud.* 84 (1), 357–405.
- Hottman, Colin J., Monarch, Ryan, 2020. A matter of taste: Estimating import price inflation across US income groups. *J. Int. Econ.* 127, 103382.
- Hummels, David, Jorgensen, Rasmus, Munch, Jakob, Xiang, Chong, 2014. The wage effects of offshoring: Evidence from danish matched worker-firm data. *Amer. Econ. Rev.* 104 (6), 1597–1629.
- Jäkel, Ina C., Smolka, Marcel, 2017. Trade policy preferences and factor abundance. *J. Int. Econ.* 106, 1–19.
- Jaravel, Xavier, Sager, Erick, 2023. What are the Price Effects of Trade? Evidence from the U.S. and Implications for Quantitative Trade Models. Mimeo.
- Katz, Lawrence F., Murphy, Kevin M., 1992. Changes in relative wages, 1963–1987 - supply and demand factors. *Q. J. Econ.* 107 (1), 35–78.
- Klein, Matthew C., Pettis, Michael, 2020. *Trade Wars are Class Wars: How Rising Inequality Distorts the Global Economy and Threatens International Peace*. Yale University Press.
- Kleinman, Benny, Liu, Ernest, Redding, Stephen J., 2022. *International Friends and Enemies*. Working Paper.
- Krugman, Paul, 1980. Scale economies, product differentiation, and the pattern of trade. *Am. Econ. Rev.* 70 (5), 950–959.
- Krugman, Paul, 2000. Technology, trade and factor prices. *J. Int. Econ.* 50 (1), 51–71.
- Kucheryavyy, Konstantin, Lyn, Gary, Rodríguez-clare, Andrés, 2023. Grounded by gravity: A well-behaved trade model with industry-level economies of scale. *Am. Econ. J.: Macroecon.* 15 (2), 372–412.
- Lebow, David E., Rudd, Jeremy B., 2003. Measurement error in the consumer price index: Where do we stand? *J. Econ. Lit.* 41 (1), 159–201.
- Mayda, Ana Maria, Rodrik, Dani, 2005. Why are some people (and countries) more protectionist than others? *Eur. Econ. Rev.* 49 (6), 1393–1430.
- Mirrlees, James A., 1971. An exploration in the theory of optimum income taxation. *Rev. Econ. Stud.* 38 (2), 175–208.
- Nigai, Sergey, 2016. On measuring the welfare gains from trade under consumer heterogeneity. *Econ. J.* 126 (593), 1193–1237.
- Oberfield, Ezra, Raval, Devesh, 2021. Micro data and macro technology. *Econometrica* 89 (2), 703–732.
- Ossa, Ralph, 2015. Why trade matters after all. *J. Int. Econ.* 97, 266–277.
- Porto, Guido G., 2006. Using survey data to assess the distributional effects of trade policy. *J. Int. Econ.* 70 (1), 140–160.
- Redding, Stephen J., Weinstein, David E., 2017. *Aggregating from Micro to Macro Patterns of Trade*. NBER Working Paper.
- Rodrik, Dani, 1997. *Has Globalization Gone Too Far?* Institute for International Economics, Washington, D.C., p. 128.
- Ruggles, Steven, Genadek, Katie, Goeken, Ronald, Grover, Josiah, Sobek, Matthew, 2015. *Integrated Public Use Microdata Series: Version 6.0*. University of Minnesota, Minneapolis.
- Schott, Peter K., 2008. The relative sophistication of Chinese exports. *Econ. Policy* 23 (53), 6–49.
- Simonovska, Ina, Waugh, Michael E., 2014. The elasticity of trade: Estimates and evidence. *J. Int. Econ.* 92 (1), 34–50.
- Slaughter, Matthew J., 2001. International trade and labor-demand elasticities. *J. Int. Econ.* 54 (1), 27–56.
- Soderbery, Anson, 2015. Estimating import supply and demand elasticities : Analysis and implications. *J. Int. Econ.* 96 (1), 1–17.
- Stewart, Frances, 2005. Horizontal inequalities: A neglected dimension of development. *Wider Perspect. Glob. Dev.* (81), 101–135.
- Van Patten, Diana, Méndez-Chacón, Esteban, 2022. Voting on a Trade Agreement: Firm Networks and Attitudes Toward Openness. NBER Working Paper 30058.
- Wilmers, Nathan, 2017. Does consumer demand reproduce inequality? High-income consumers, vertical differentiation and the wage structure. *Am. J. Sociol.* 123 (1), 178–231.
- Wood, Adrian, 1995. How trade hurt unskilled workers. *J. Econ. Perspect.* 9 (3), 57–80.