Essays In International Trade and Labor Markets

by

Rodrigo Rodrigues Adão

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

Author ..Signature redacted..................
Department of Economics
May 18, 2016

Certified by.Signature redacted.................
Arnaud Costinot
Professor of Economics
Thesis Supervisor

Certified by...Signature redacted..................
Daron Acemoglu
Elizabeth and James Killian Professor of Economics
Thesis Supervisor

Accepted by...Signature redacted..................
Ricardo Caballero
Ford International Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Essays In International Trade and Labor Markets

by

Rodrigo Rodrigues Adão

Submitted to the Department of Economics
on May 13, 2016, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis develops empirical methodologies to investigate the effect of globalization on welfare and inequality both between- and within-countries.

The first essay proposes a Roy-like model where workers are heterogeneous in terms of their comparative and absolute advantage. We show that the schedules of comparative and absolute advantage (i) determine changes in the average and the variance of the log-wage distribution, and (ii) are nonparametrically identified from the cross-regional variation in the sectoral responses of employment and wages to observable sector-level demand shifters. Applying these results, we find that the rise in world commodity prices accounts for 5–10% of the fall in Brazilian wage inequality between 1991 and 2010.

The second essay develops a methodology to construct nonparametric counterfactual predictions, free of functional-form restrictions on preferences and technology, in neoclassical models of international trade. First, we establish the equivalence between such models and reduced exchange models in which countries directly exchange factor services. This equivalence implies that, for an arbitrary change in trade costs, counterfactual changes in factor prices, and welfare only depend on the shape of a reduced factor demand system. Second, we provide sufficient conditions for the nonparametric identification of this system. Together, these results offer a strict generalization of the parametric approach used in so-called gravity models. Finally, we use China's recent integration into the world economy to illustrate the feasibility of our approach.

The third essay investigates the connection between the recent rise in services trade and changes in labor market outcomes in different countries. We develop a theoretical framework where trade in services arises from the spatial unbundling of workers' task output. Transmission costs endogenously determine the magnitude of between-sector task trade both within a country ("outsourcing") and between countries ("offshoring"). We show that, while differentials in sectoral task prices decrease in response to outsourcing, they increase in response to offshoring. The heterogeneity in the composition of workers' task endowments controls responses in between- and within-sector wage inequality across countries.

Thesis Supervisor: Arnaud Costinot
Title: Professor of Economics

Thesis Supervisor: Daron Acemoglu
Title: Elizabeth and James Killian Professor of Economics

Acknowledgments

First and foremost, I am deeply indebted to my advisor, Arnaud Costinot, for my academic development throughout the past five years. Since my first semester in MIT, he has been extremely generous in sharing his time, and has provided thoughtful comments that helped transform my crude ideas into the essays in this thesis. During my graduate years, I learned from him as a student, as a research assistant, as a teacher assistant, and as a coauthor. I benefited immensely from his advice, guidance, and support.

I am profoundly grateful to my advisors, Daron Acemoglu and Dave Donaldson, for their invaluable advice during my graduate career. Their comments were instrumental in the successful completion of this thesis. I would like to especially thank Daron for motivating me to focus on ambitious research questions, and for pointing directions at various crossroads in the development of my research. I also greatly benefited from coauthoring with Dave the second essay of this thesis.

I would like to thank the entire MIT Economics community for creating a remarkable intellectual environment where research ideas flourish. In particular, I thank David Atkin for the excellent advice on the presentation of the first essay in this thesis. I also thank David Autor and Ivàn Werning for their comments at different stages of this thesis. Finally, I thank my fellow graduate students for the fruitful discussions during my years in MIT. In special, I thank Dejanir Silva for taking the time to listen to my vague ideas and for sharing his unique insights on my research.

I owe special gratitude to my family for their endless love and support. I dedicate this thesis to my parents whose unconditional support was essential in achieving my dream of obtaining a doctoral degree. I thank my wife for her support in every step of the Ph.D. program, and for putting up with my physical and mental absence as I wrote this thesis. I also thank my sister, Camila, and my brother, Filipe, for being on my side throughout my entire academic trajectory. Lastly, I thank all my friends for the many happy moments that made my graduate years much more pleasant.

Contents

List of Figures

List of Tables

Chapter 1

Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil

1.1 Introduction

In a global economy, changes in good prices caused by shocks in one part of the world have the potential to affect factor prices in another. As shown in Panel A of Figure 1.1, between 1981 and 2010, increases in the world prices of basic commodities were accompanied by reductions in Brazilian wage inequality. Given the importance of the commodity sector in employment of low-income workers, this correlation suggests that changes in world demand for basic goods plausibly contributed to changes in wage inequality in Brazil.[1] Panel B reinforces this view by showing that increases in world commodity prices were also associated with increases in both the relative employment and the relative wage in the commodity sector. In this paper, I develop a new empirical strategy to quantify the causal effect of global shocks in commodity

[1]Production of agricultural and mining products constitutes an important share of the Brazilian economy, representing, in 2010, 58.5% of exports and 19.9% of employment. Commodity sector employees earned, on average, 28.1% less than employees of other sectors in 2010. In Appendix 1A.3, I show that the component associated with workers' observable characteristics was the main driver of the movements in log-wage variance between 1981 and 2009.

prices on Brazilian wage inequality.

My starting point is a theoretical framework where Brazil is assumed to be a collection of small open economies with segmented labor markets. Each regional economy is populated by workers of multiple demographic groups that can be employed either in the commodity or in the non-commodity sectors. The central feature of the model is a Roy's (1951) structure of within-group worker heterogeneity in terms of sector-specific productivity. Conditional on sectoral wages per efficiency unit, workers self-select into sectors according to their comparative advantages; defined as the productivity ratio in the commodity and the non-commodity sectors. In the model, workers' labor income depend on their comparative advantages as well as their absolute advantages; defined as the productivity in the non-commodity sector.

In this environment, comparative and absolute advantage have distinct roles in determining sectoral responses of employment and wages following shocks to the world prices of goods. By affecting the marginal value of labor in each sector, world price shocks induce changes in the sectoral relative wage per efficiency unit. This causes between-sector worker reallocation with magnitude regulated by the comparative advantage distribution, which I refer to as the schedule of comparative advantage. The subsequent between-sector response in average wage combines two terms. The first term is the impact of the change in the relative wage per efficiency unit for a given allocation of workers across sectors. The second term is the compositional effect stemming from the difference in the average sector-specific efficiency of sector-switchers relative to that of sector-stayers. The magnitude of this compositional effect depends on the average of the absolute advantage distribution conditional on comparative advantage, which I refer to as the schedule of absolute advantage.

These sectoral shocks trigger changes in wage inequality, both between and within worker groups. To quantify such distributional effects in the model, I focus on the shock's impact over the average and the variance of the log-wage distribution of different demographic groups. Following sectoral demand shocks, I show that responses in these outcomes are exclusively determined by the schedules of comparative and absolute advantage. Thus, knowledge of these two schedules permits a quantitative

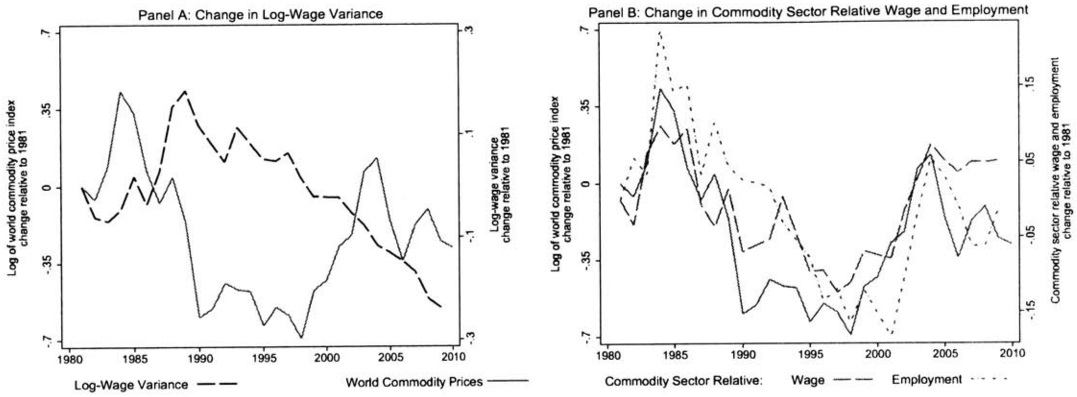

Figure 1.1: World Commodity Prices and the Brazilian Labor Market, 1981-2010

Note: World Commodity Price is the log of the commodity price index computed with the world price of agriculture and mining products converted to Brazilian currency and deflated by the Brazilian consumer price index. Sample of full-time employed males is extracted from the National Household Sample Survey (PNAD). Commodity sector relative wage is the coefficient of the dummy for employment in commodity sector from the regression of log wage on worker attributes. Change in commodity sector relative employment is the average of the change in the log of the employment ratio in the commodity and non-commodity sectors for High School Graduates and Dropouts weighted by the group size in 1981. Details in Appendix 1A.3.

evaluation of the impact of world price shocks on wage inequality.

I then turn to the problem of recovering the schedules of comparative and absolute advantage from observable labor market outcomes. The challenge inherent in identifying these functions is conveyed by Heckman and Honoré's (1990) result that, in the context of the Roy model, the sector-specific productivity distribution is not nonparametrically identified in a single cross-section of individuals. In this paper, I establish the identification of the schedules of comparative and absolute advantage in a set of regional economies. For any number of demographic worker groups, my identification result allows the two schedules to have an arbitrary shape. But it requires two central assumptions. First, I assume that observed covariates and unobserved shocks are additive shifters of the two schedules across regions. Second, given the unobserved productivity shocks, I make the standard assumption that there are excludable shifters of sector labor demand across regions. Under these assumptions, the two schedules are nonparametrically identified from cross-regional variation in sectoral responses of employment and average wages to changes in sectoral wages per efficiency unit induced by the observable sector-level demand shifter.

My nonparametric identification result is critical to inform the source of variation in the data that separately uncovers comparative and absolute advantage. An empirical application based on this result accounts for the conceptually distinct roles of these two schedules in the model. This approach contrasts with recent empirical applications of Roy-like models that build upon a productivity distribution of the Fréchet family; e.g., see Hsieh, Hurst, Jones, and Klenow (2013), Burstein, Morales, and Vogel (2015), and Galle, Rodriguez-Clare, and Yi (2015). This distribution, although highly tractable, mixes the channels of comparative and absolute advantage with strong consequences for the model's predictions: it implies that both sector wage differentials and log-wage variance are invariant to labor demand conditions. To incorporate these potentially important channels while maintaining tractability, my empirical application relies on a parsimonious log-linear system that strictly generalizes the system implied by the Fréchet distribution. The log-linear system contains two structural parameters that specify constant-elasticity schedules of comparative and absolute advantage. In comparison, the Fréchet distribution restricts these two elasticities to have the same absolute value.[2]

Armed with these theoretical results, I apply the framework to investigate the effect of commodity price shocks on wage inequality in Brazil. To this extent, I estimate the schedules of comparative and absolute advantage in a panel of Brazilian regional economies for two demographic groups, High School Graduates and High School Dropouts. In the empirical application, two variables are needed. First, a regional shifter of sectoral demand, which I construct by interacting the change in commodities' world prices and the pre-shock participation of corresponding commodities in the region's labor payroll.[3] Second, a measure of the sector wage per efficiency

[2]In the spirit of the series estimator proposed by Newey and Powell (2003b), the system could be augmented to include higher-order polynomials. In practice, data limitations constitute an important challenge to the implementation of a fully flexible instrumental variable estimator. As Newey (2013b) pointed out, the estimation of nonlinear terms with instrumental variables tends to be accompanied by sharp increases in standard errors. For this reason, my benchmark specification is based on a parsimonious log-linear system with constant-elasticity schedules of comparative and absolute advantage.

[3]My demand shifter is implied by the assumption that production of basic commodities utilizes immobile factors like soil fertility and oil reserves whose endowment varies across regions. As a result, following world price shocks, the regional response of the commodity sector labor demand

unit that is not immediately available in survey datasets. To estimate changes in the sector wage per efficiency unit, I propose a strategy that builds upon the model's predicted relation between wage growth and initial sector employment across quantiles of the wage distribution. For each group and region, I implement this strategy as a first-step regression using repeated cross-section data on wage and employment at the individual level.

I start by investigating the effect of exposure to commodity price shocks on sectoral labor market outcomes across Brazilian regional economies. This reduced-form exercise establishes the basic relations in the data that drive the estimation of the structural parameters of comparative and absolute advantage. For both worker groups, I find that regional economies exposed to stronger price shocks experienced stronger expansions in the commodity sector relative employment. In addition, shock exposure induced increases in the relative wage per efficiency unit of the commodity sector. The combination of these two responses determines the elasticity of the comparative advantage schedule. Lastly, I investigate the effect of shock exposure on the commodity sector wage differential, finding a positive and statistically significant response for High School Graduates and a small and statistically non-significant response for High School Dropouts. Following the commodity price shock, the change in the relative sector average wage was smaller than the change in the relative wage per efficiency unit. This wedge corresponds to the compositional effect that determines the elasticity of the absolute advantage schedule. Results are robust to the inclusion of region fixed-effects, initial region socio economic characteristics interacted with period dummies, and region-specific time trends.

Having established these reduced-form patterns in the sample of Brazilian regions, I turn to the estimation of the structural parameters separately for High School Graduates and Dropouts. Results indicate that the two groups have similar comparative advantage schedules, implying that they exhibit comparable degrees of between-sector mobility. The distinct responses in the sector wage differential for the two groups leads to different estimated coefficients of absolute advantage. Among High School

depends on the initial industry composition within the commodity sector.

Dropouts, estimates are consistent with those of a Fréchet distribution and, for this group, compositional effects completely offset the impact of price shocks on sector wage differentials. The estimated selection pattern for High School Graduates, however, differs from that implied by the Fréchet system. For this group, the log-linear system is able to replicate the estimated effect on log-wage variance associated with exposure to higher commodity prices across Brazilian regions. Such a response is ruled out by the parametric restrictions imposed by the Fréchet distribution.

I conclude the paper by applying the framework to answer one counterfactual question: "In 1991, how would wage inequality change if commodity prices were equal to those of 2010?" To answer this question, I provide two alternative procedures to obtain changes in sectoral wages per efficiency unit stemming from shocks in world commodity price. The first relies on a reduced-form pass-through estimated from the effect of price shock exposure on the wage per efficiency unit in the sample of Brazilian regional economies. While this approach is robust to the specific production structure of the economy, it is not able to capture nationwide effects and it may not hold for shocks on other products and other periods. To address these shortcomings, the second approach relies on a fully specified general equilibrium model where I calibrate the economy's structure of production. This procedure takes inspiration from the exact hat algebra used in recent international trade papers — see, for example, Dekle, Eaton, and Kortum (2007) and, for a review, Costinot and Rodríguez-Clare (2013).

The counterfactual analysis yields similar results with both approaches, delivering two main insights. First, changes in world commodity prices have sizable distributional effects in Brazil. As a result of the 1991–2010 rise in world commodity prices, the relative wage per efficiency unit in the commodity sector increased by 8%–16%. Yet the subsequent worker reallocation created compositional effects that offset most of the shock's impact on between-sector wage differentials. In terms of overall wage inequality, the price shock accounts for 5%–10% of the decline in Brazilian log-wage variance between 1991 and 2010. Second, flexible functional forms that separate the roles of comparative and absolute advantage are quantitatively important. For High School Graduates, the log-linear model captures 10% of the decrease in log-wage variance, but the Fréchet model implies no change in log-wage variance. In contrast, both specifications yield similar counterfactual changes in the average and the variance of the log-wage distribution for High School Dropouts, reflecting the similarity between the estimated structural parameters obtained with the two parametrizations.

This paper is related to an extensive literature on the labor market effects of international trade. Research on the topic has traditionally relied on neoclassical environments that yield stark predictions regarding the changes in relative wages across worker groups (Stolper and Samuelson, 1941; Jones, 1965) and relative factor prices across industries (Jones, 1975). However, empirical studies concluded that the forces highlighted by these models were, at best, secondary drivers of the changes in wage inequality in the 1980s and early 1990s. For instance, a number of authors have documented (i) movements in wage inequality correlated in both developed and developing countries (Goldberg and Pavcnik, 2007); (ii) movements in the skill wage premium uncorrelated with changes in the relative price of skill-intensive products (Lawrence and Slaughter, 1993) while correlated with changes in the skill intensity of production within industries (Berman, Bound, and Machin, 1998); and (iii) limited between-sector responses in employment and wages following trade shocks (Wacziarg and Wallack, 2004 and Goldberg and Pavcnik, 2007).

This evidence motivated departures from the neoclassic environment, giving rise to a body of work analyzing the effect of international trade on workers employed in different firms within industries (see, for example, Verhoogen, 2008; Helpman, Itskhoki, and Redding, 2010; Frias, Kaplan, and Verhoogen, 2012; Helpman, Itskhoki, Muendler, and Redding, 2015; and Burstein and Vogel, 2015) and on the transitional dynamics in the reallocation of workers across sectors and markets (Kambourov, 2009; Artuç, Chaudhuri, and McLaren, 2010; Dix-Carneiro, 2014; Dix-Carneiro and Kovak, 2015a; and Caliendo, Dvorkin, and Parro, 2015). In this paper, I build upon the neoclassical channel that emphasizes the effect of international trade on relative good prices and, consequently, on relative factor prices. Yet my framework augments traditional models with a flexible structure of sector-specific factor productivity. This idea goes back to the work of Mussa (1982) and Grossman (1983), and its implications for international trade is explored in several recent papers — see, for example, Ohnsorge and Trefler (2007), Costinot and Vogel (2010), Acemoglu and Autor (2011), and, for comprehensive reviews, Grossman (2013) and Costinot and Vogel (2014).[4] Relative to these papers, my main contribution is to develop a novel empirical methodology that allows the model to be applied in the quantification of distributional effects of international trade shocks. In the context of the shock in world commodity prices of 1991–2010, my framework indicates sizable distributional effects in the Brazilian labor market.

Two recent papers impose a productivity distribution of the Fréchet family to quantify the portion of changes in between-group wage inequality associated with technological progress in the United States (Burstein, Morales, and Vogel, 2015) and import competition in Germany (Galle, Rodriguez-Clare, and Yi, 2015). My paper differs from these studies in two central aspects of methodology. First, my analysis clearly delineates the distinct roles played by comparative and absolute advantage in determining sectoral employment and sectoral wages, showing how these schedules affect both within and between-group wage inequality. Second, my nonparametric identification result sheds light on the source of variation that uncovers comparative and absolute advantage within each demographic worker group, leading to a new estimation strategy based on cross-market variation in sectoral demand shifters. The results of my counterfatual analysis suggest that the restrictive distributional assumptions imposed by these papers have the potential to significantly affect the quantitative predictions of the model.

This paper is also related to the empirical literature that examines the impact on labor market outcomes of heterogeneous exposure to import competition in terms of sector of employment (Menezes-Filho and Muendler, 2011; and Autor, Dorn, Hanson, and Song, 2014), and region of residence (Topalova, 2010; Kovak, 2013; Autor, Dorn, and Hanson, 2013; Costa, Garred, and Pessoa, 2014; and Dix-Carneiro and Kovak, 2015b). I complement this literature by providing new evidence of sectoral responses

[4]Also, the Roy model has been recently applied in the investigation of the determinants of aggregate productivity — e.g., see Lagakos and Waugh (2013), Hsieh, Hurst, Jones, and Klenow (2013), and Young (2014).

of employment and wages using cross-regional variation in exposure to commodity price shocks in a developing country.[5] My theoretical framework, moreover, connects these responses to structural parameters of comparative and absolute advantage. The structural estimates indicate that, due to compositional effects, the impact of world price shocks on sectoral wage per efficiency unit is larger than one would have inferred from reduced-form regressions based on sector average wages.

Lastly, this paper is related to the large literature investigating the consequences of self-selection based on unobservable characteristics to observable components of labor income — see French and Taber (2011) for a review. In the context of the Roy model, Heckman and Honoré (1990) offer a number of results regarding the nonparametric identification of the sector-specific productivity distribution. By focusing on the schedules of comparative and absolute advantage, my nonparametric identification result relies on weaker assumptions than those imposed by Heckman and Honoré (1990); in particular, I allow for cross-market variation in sectoral efficiency in the form of unobserved additive shifters of comparative and absolute advantage. In this environment, I show that the supply equations relating sector employment and sector average wages to the schedules of comparative and absolute advantage belong to the class of separable models studied by Newey and Powell (2003b), being nonparametrically identified under the same exogeneity and completeness conditions outlined by these authors. To the extent that workers have different levels of sectoral productivity across markets, the flexibility implied by the environment in this paper is important in empirical applications of the Roy model. In fact, very different estimates of the structural parameters are obtained without sector demand shifters.

The rest of the chapter is organized as follows. Section 1.2 presents the model and its implications for the equilibrium structure of employment and wages. Section 1.3 establishes the nonparametric identification of comparative and absolute advantage. Section 1.4 presents the estimation of these schedules in the panel of Brazilian regional

[5]Previous studies analyzing the adjustment of labor markets during trade liberalization episodes in developing countries did not find evidence of responses in employment and wages; e.g., see Wacziarg and Wallack (2004) and Goldberg and Pavcnik (2007). By contrast, evidence of sectoral responses in labor market outcomes has been documented in developed countries; e.g., see Revenga (1992), Gaston and Trefler (1997), and Autor, Dorn, Hanson, and Song (2014).

labor markets differentially exposed to shocks in world commodity prices. Section 1.5 presents the counterfactual analysis of the effect of changes in world commodity prices on changes in Brazilian wage inequality. Section 1.6 offers some concluding remarks.

1.2 Model

My goal is to develop a framework to quantify the effect of shocks in world commodity prices on Brazilian wage inequality. For this purpose, I assume that Brazil is a collection of small open economies with segmented labor markets. Consequently, good prices are exogenously determined internationally, but factor prices are endogenously determined regionally.[6]

1.2.1 Environment

Each regional economy contains workers of multiple demographic groups, $g \in \{1, ..., G\}$, and two aggregate sectors, the commodity sector ($k = C$) and the non-commodity sector ($k = N$). Within each demographic group, there is a continuum of heterogeneous individuals, $i \in \mathcal{I}_g$, endowed with a bivariate skill vector, $(L_g^C(i), L_g^N(i))$, that determines their productivity if employed in each aggregate sector of the economy. This is the core assumption of a large class of Roy-like (1951) models, and it is central in my analysis of the distributional effects of sectoral demand shocks.

In order to incorporate the various commodity categories in the empirical application, I assume that each aggregate sector comprises multiple perfectly competitive industries, $j \in \mathcal{J}^k$, that produce homogeneous goods freely traded in the world market at price p^j. In every industry j of the aggregate sector k, individuals have an identical level of sector-specific productivity. The production technology in industry j utilizes the total number of sector-specific efficiency units supplied by employees, L_g^j, and an industry-specific nonlabor input, X^j. Specifically, the production function

[6] This paper abstracts from migration flows between regional labor markets in Brazil. This simplification is motivated by the empirical analysis below, where I find weak migration responses following regional shocks to the labor demand in the commodity sector. This point is carefully discussed in Section 1.4.

is given by

$$q^j = Q^j\left(L_1^j, ..., L_G^j, X^j\right) \quad \text{where} \quad L_g^j \equiv \begin{cases} \int_{\mathcal{S}_g^j} L_g^C(i)\, di & \text{if } j \in \mathcal{J}^C \\ \\ \int_{\mathcal{S}_g^j} L_g^N(i)\, di & \text{if } j \in \mathcal{J}^N \end{cases} \tag{1.1}$$

and $\mathcal{S}_g^j$ is the set of individuals of group g employed in industry j. Function $Q^j(.)$ is strictly increasing, concave, differentiable, and homogeneous of degree one. The technology allows, but does not require, the effective labor supply of workers of different groups to be imperfect substitutes in production.

This production structure determines the effect of price shocks at the product-level on the demand for labor at the sector-level. In the empirical analysis, I explore this structure to obtain a regional shifter of the commodity sector's labor demand following shocks in world commodity prices. The cross-regional variation in this shifter is implied by the limited supply of the industry-specific nonlabor factor that, in this context, corresponds to the regional endowment of natural resources necessary for production of agricultural and mining goods — e.g., fertile soil, rainfall, metal reserves, or oil reserves.[7]

The analysis is greatly simplified by working with a log-linear transformation of individuals' sector-specific productivities. Define individual i's comparative advantage as $s_g(i) \equiv \ln[L_g^C(i)/L_g^N(i)]$, and absolute advantage as $a_g(i) \equiv \ln[L_g^N(i)]$. In a given group, suppose individuals independently draw their productivity vector from a common bivariate distribution such that, without loss of generality,

$$s_g(i) \sim F_g(s) \quad \text{and} \quad \{a_g(i)|s_g(i) = s\} \sim H_g\left(a|s\right) \tag{1.2}$$

where, for simplicity, $F_g(s)$ is assumed to have full support in $\mathbb{R}$.

[7]Alternatively, one could consider any environment with a generic sector demand for labor efficiency units. For instance, it is straight forward to allow for non-competitive product markets and other mobile factors of production. These extensions do not affect the main insights discussed in this section.

1.2.2 Competitive Equilibrium

In the competitive equilibrium, producers maximize profits conditional on both world product prices and local factor prices. In all industries of the aggregate sector k, producers face an identical labor cost: the sector's wage per efficiency unit, w_g^k. As a result, conditional on world product prices, the labor demand in industry j of sector k is given by, for all $g = 1, ..., G$,

$$w_g^k = p^j \cdot \frac{\partial Q^j}{\partial L_g^j} \quad \text{if } \ j \in \mathcal{J}^k \tag{1.3}$$

where $X^j = \bar{X}^j$, with $\bar{X}^j$ denoting the economy's endowment of the industry-specific nonlabor input.

To determine the supply of effiency units of labor in each sector, consider the employment decision of workers seeking to maximize total labor income. Individual i of group g, if employed in any industry j of sector k, receives w_g^k for each sector-specific efficiency unit supplied. Let $y_g^k(i)$ denote the potential log-wage of individual i in any industry of sector k. Using the log-transformation above, these potential log-wages are given by

$$y_g^N(i) \equiv \omega_g^N + a_g(i) \quad \text{and} \quad y_g^C(i) \equiv \omega_g^C + s_g(i) + a_g(i) \tag{1.4}$$

where $\omega_g^k \equiv \ln w_g^k$.

Because all industries of an aggregate sector yield the same labor income, individuals are indifferent between them. Yet individuals receive different wages in the two sectors and, for this reason, they self-select into the sector where their labor income is higher.[8] Hence, the set of individuals employed in sector k, $\mathcal{S}_g^k$, is given by

$$\mathcal{S}_g^k \equiv \left\{ i \in \mathcal{I}_g : \; k = \text{argmax}\{y_g^C(i), y_g^N(i)\} \right\}. \tag{1.5}$$

In the competitive equilibrium of this economy, sectoral wages per efficiency unit guarantee factor market clearing in the two sectors. Specifically, $\{(w_g^C, w_g^N)\}_g$ are such that, for all g and k,

$$\sum_{j \in \mathcal{J}^k} L_g^j = \int_{\mathcal{S}_g^k} L_g^k(i) \; di \tag{1.6}$$

where, in every industry j, condition (1.3) determines L_g^j; and condition (1.5) determines $\mathcal{S}_g^k$. In order to satisfy labor demand at the industry level, individuals employed in each sector are allocated across industries to satisfy conditions (1.1) and (1.3).

1.2.3 Sectoral Log-Wages and Employment

To determine workers' sectoral employment decisions in the model, I consider a graphical representation of the economy where individuals are ranked according to their level of comparative advantage. For each quantile $q \in [0, 1]$, there is a set of individuals in group g whose level of comparative advantage is $\alpha_g(q) \equiv (F_g)^{-1}(q)$. By construction, $\alpha_g(q)$ is increasing in q so that individuals in higher quantiles are relatively more efficient in the commodity sector than those in low quantiles. Among individuals in quantile q, there is a conditional distribution of absolute advantage, $H_g\left(a|\alpha_g(q)\right)$, with average and variance respectively denoted by $A_g(q)$ and $V_g(q)$. In the rest of this paper, $\alpha_g(.)$ is the schedule of comparative advantage, and $A_g(.)$ is the schedule of absolute advantage.

Figure 1.2 exhibits the average potential log wage in each sector for individuals of group g distributed across quantiles of comparative advantage. Immediately from

[8]This particular formulation closely follows the environment in the extensive literature inspired by the seminal work of Roy (1951). By introducing worker heterogeneity entirely on sector-specific productivity, the distributive impact of a trade shock is completely captured by the behavior of observable labor income. Notice that this model abstracts from between-sector mobility costs. In Appendix 1A.2, I explore an extension that incorporates such a feature into the model in the form of heterogeneity in non-monetary private benefits of employment across individuals. The extended model yields similar conclusions as those outlined in this section.

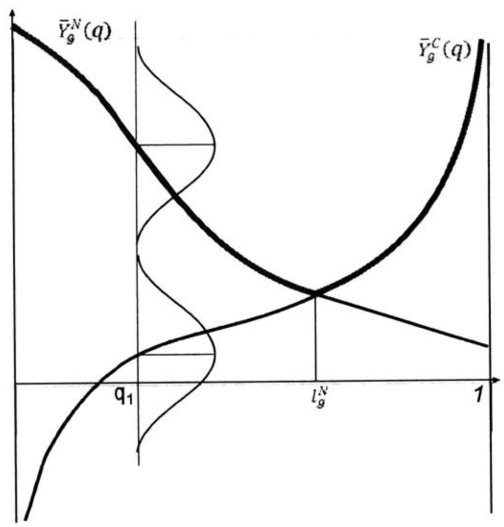

Figure 1.2: Sectoral Log-Wages and Employment in Equilibrium

Red: non-commodity sector employees in group g. Blue: commodity sector employees in group g.

expression (1.4), the average log-wage of workers in quantile q is $\bar{Y}_g^N(q) = \omega_g^N + A_g(q)$ if employed in the non-commodity sector. Alternatively, these workers earn an average log-wage of $\bar{Y}_g^C(q) = \omega_g^C + \alpha_g(q) + A_g(q)$ if employed in the commodity sector. In a particular quantile, the unique source of dispersion in potential sector wages is the dispersion of absolute advantage, $V_g(q)$ — illustrated by the hump-shaped curves in quantile q_1. Lastly, it is important to notice that the two potential log-wage curves exhibit the single-crossing property, because $\alpha_g(q)$ is increasing in q.[9]

The importance of Figure 1.2 lies in the fact that it simultaneously illustrates sectoral employment and sectoral wages for any given level of (ω_g^C, ω_g^N). All individuals in a particular quantile q choose to be employed in the same sector since, for all of them, the potential log-wage premium in the commodity sector is $\omega_g^C + \alpha_g(q) - \omega_g^N$.[10] In high

[9]To simplify the analysis, Figure 1.2 imposes that (ω_g^C, ω_g^N) are such that these curves cross at least once. Inada conditions on the production technology are sufficient for this to occur in equilibrium.

[10]To formalize this claim, consider individual i with comparative advantage $s_g(i) = \alpha_g(q)$. For this individual, potential sector wages in (1.4) correspond to vertical shifts of those of a worker with the same level of comparative advantage but a different level of absolute advantage. Consequently, the sectoral choice of individual i with $s_g(i) = \alpha_g(q)$ is identical to that of a hypothetical individual i' with $s_g(i') = \alpha_g(q)$ and $a_g(i') = A_g(q)$.

quantiles of comparative advantage, the relatively higher efficiency in the commodity sector yields a relatively higher wage in that sector, implying self-selection into the commodity sector — i.e., the blue portion of the potential average wage curve in Figure 1.2. In contrast, individuals in low quantiles of comparative advantage obtain a relatively lower wage in the commodity sector, finding it optimal to self-select into the non-commodity sector — i.e., the red portion of the potential average wage curve in Figure 1.2. Finally, the marginal individuals at the intersection of the two curves have exactly the same potential wage in the two sectors, being indifferent between them. Thus, I establish the following result.

Proposition 1. *Conditional on (ω_g^C, ω_g^N), the allocation of individuals to sectors depends exclusively on their level of comparative advantage. In particular, individual i with $s_g(i) = \alpha_g(q)$:*

i. self-selects into the commodity sector if $\alpha_g(q) > \omega_g^N - \omega_g^C$;

ii. self-selects into the non-commodity sector if $\alpha_g(q) < \omega_g^N - \omega_g^C$; and

iii. is indifferent between the two sectors if $\alpha_g(q) = \omega_g^N - \omega_g^C$.

Proposition 1 indicates the central role played by comparative advantage in determining the sectoral allocation of workers in the model. In equilibrium, the sector employment composition is determined by marginal individuals with comparative advantage equal to the relative wage per efficiency unit, $\omega_g^N - \omega_g^C$. As a result, the share of individuals of group g employed in the non-commodity sector, l_g^N, is determined by the intersection of the two sectoral curves of potential average log-wages:

$$\omega_g^N - \omega_g^C = \alpha_g\left(l_g^N\right). \tag{1.7}$$

Given the sectoral employment decision described in Proposition 1, Figure 1.2 immediately yields the average log-wage of workers in each quantile of comparative advantage. Aggregating across the quantiles allocated to each sector, I obtain the sector average log-wage, $\bar{Y}_g^k$, which is given by

$$\bar{Y}_g^k = \omega_g^k + \bar{A}_g^k(l_g^N) \quad \text{where} \quad \bar{A}_g^k(l) \equiv \begin{cases} \frac{1}{l}\int_0^l A_g(q)\, dq & \text{if } k = N \\ \\ \frac{1}{1-l}\int_l^1 [\alpha_g(q) + A_g(q)]\, dq & \text{if } k = C. \end{cases} \tag{1.8}$$

In expression (1.8), there are two determinants of the average sector log-wage. The first is the sector wage per efficiency unit, ω_g^k, that directly affects the log-wage of all sector employees symmetrically. The second is the sector employment composition, l_g^N, that affects the average efficiency of sector employees through the function $\bar{A}_g^k(.)$. This compositional effect is generated by the variation in the average sector-specific efficiency of workers in different quantiles of comparative advantage. That is, it depends on the shape of the sectoral curves of average efficiency: $A_g(.)$ in the non-commodity sector and $\alpha_g(.) + A_g(.)$ in the commodity sector. In Figure 1.2, $A_g(q)$ is decreasing and $A_g(q) + \alpha_g(q)$ is increasing. This case entails "positive selection into both sectors" because the average sector employee is *more* efficient than marginal workers indifferent between the two sectors (i.e., those in quantile l_g^N). In this case, the average sector-specific efficiency decreases as employment expands in the two sectors: $\bar{A}_g^N(l_g^N)$ is decreasing, and $\bar{A}_g^C(l_g^N)$ is increasing. The model, however, imposes only weak restrictions on the shape of $A_g(q)$ and $\alpha_g(q) + A_g(q)$ since comparative and absolute advantage can be arbitrarily related. As discussed below, the different possible shapes of these functions imply qualitatively different compositional effects in the adjustment of sector average wages to sectoral demand shocks.

Proposition 2. *Conditional on* (ω_g^C, ω_g^N)*, the average sector log-wage,* $\bar{Y}_g^k$*, depends on the sector employment composition,* l_g^N*, through the average efficiency of sector employees,* $\bar{A}_g^k(l_g^N)$*, in equation* (1.8)*. In the non-commodity sector, this compositional effect depends on* $A_g(.)$*; in the commodity sector, it depends on* $\alpha_g(.) + A_g(.)$*.*

1.2.4 Sectoral Demand Shocks and Sectoral Changes in Wages and Employment

In order to illustrate the mechanics of the model, let us analyze the adjustment of sector labor market outcomes following changes in sectoral wages per efficiency unit triggered by a positive shock in world commodity prices. This exercise delineates the distinct roles of comparative and absolute advantage in determining sectoral responses in terms of employment and average wage.

An increase in world commodity prices translates into higher marginal value of labor in the commodity sector. To fix ideas, I consider in this section a partial equilibrium exercise in which, after the shock, ω_g^C increases and ω_g^N remains constant. Figure 1.3 displays the induced movements in the curves of potential sector wages in three cases. Panel (a) illustrates the case analyzed above, where $A_g(q)$ is decreasing and $A_g(q) + \alpha_g(q)$ is increasing. Panels (b) and (c) present other possible shapes for these functions that will be representative of the different qualitative patterns of compositional effects allowed in the model.

The shock causes an increase of $\Delta\omega_g^C$ in the log-wage of all commodity sector employees as represented by the upward shift of the blue curve on Figure 1.3. Since

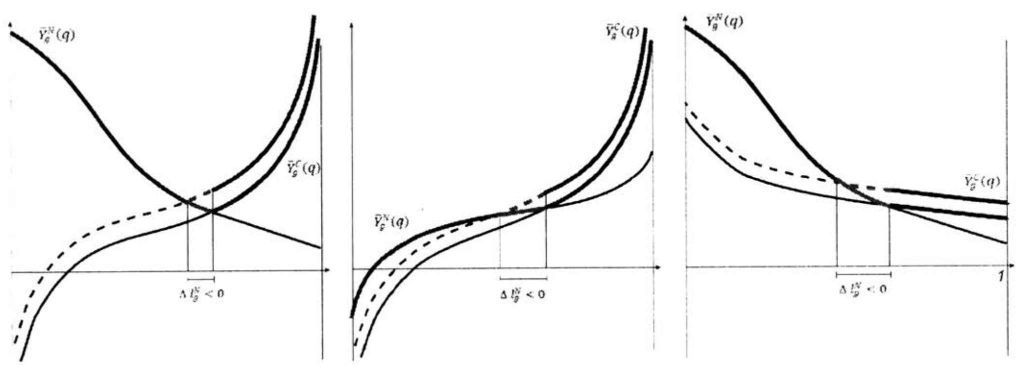

(a)
Positive selection into both sectors

(b)
Negative selection into the non-commodity sector

(c)
Negative selection into the commodity sector

Figure 1.3: Comparative Statics - increase in ω_g^C

Red: sector-stayers in the non-commodity sector. Blue: sector-stayers in the commodity sector. Green: switchers from the non-commodity sector to the commodity sector.

$\Delta\omega_g^N = 0$, the shock does not affect the wage of non-commodity sector employees and the red curve remains unchanged. Only those non-commodity sector employees who decide to switch into the commodity sector benefit from the shock. These sector-switchers are represented in green on Figure 1.3. Their wage gain is bounded from below by $\Delta\omega_g^N$, and from above by $\Delta\omega_g^C$. This is illustrated by the difference between the solid and dashed green curves on Figure 1.3.

In the model, the decision of sectoral allocation is entirely determined by each worker's comparative advantage. Thus, the mass of sector-switchers that benefit from the shock depends on the dispersion of comparative advantage among marginal workers. As implied by equation (1.7), this is captured by the slope of the comparative advantage schedule, $\alpha_g(.)$:

$$\Delta\left[\omega_g^N - \omega_g^C\right] = \int_{l_g^N}^{l_g^N + \Delta l_g^N} \frac{\partial \alpha_g(u)}{\partial q}\, du \approx \frac{\partial \alpha_g(l_g^N)}{\partial q} \cdot \Delta l_g^N \tag{1.9}$$

where $\frac{\partial \alpha_g(q)}{\partial q} \geq 0$ for all q.

Although the wage per efficiency unit remains constant in the non-commodity sector, the implied outflow of employees affects the sector's employment composition and, consequently, the sector's average wage. To the extent that the absolute advantage of sector-switchers differs from that of sector-stayers, the change in sector employment triggers a change in sector average efficiency. Intuitively, this is conveyed by the first-order expansion of equation (1.8):

$$\Delta\bar{Y}_g^N - \Delta\omega_g^N = \int_{l_g^N}^{l_g^N + \Delta l_g^N} \frac{\partial \bar{A}_g^N(u)}{\partial q}\, du \approx \left[A_g(l_g^N) - \bar{A}_g^N(l_g^N)\right] \cdot \Delta \ln(l_g^N). \tag{1.10}$$

where, by definition, $\bar{A}_g^N(l_g^N) = (1/l_g^N) \cdot \int_0^{l_g^N} A_g(q)\, dq$.

The right-hand side of equation (1.10) is the compositional effect implied by the outflow of non-commodity sector employees. This effect is proportional to the average absolute advantage of sector-switchers, $A_g(l_g^N)$, relative to that of sector-stayers, $\bar{A}_g^N(l_g^N)$. To see this, consider the three cases in Figure 1.3. In Panels (a) and (c), the decreasing schedule of absolute advantage $A_g(q)$ implies that non-commodity sector stayers (red) have a *higher level of absolute advantage* than sector-switchers (solid green). In this case, $A_g(l_g^N) < \bar{A}_g^N(l_g^N)$ and the compositional effect is positive. In words, the outflow of workers leaves the non-commodity sector with employees whose average absolute advantage is relatively higher than before. However, this is not the only possibility. In Panel (b), $A_g(q)$ is increasing so that $A_g(l_g^N) > \bar{A}_g^N(l_g^N)$ and the outflow of workers lowers the average wage in the non-commodity sector. In general, $A_g(.)$ determines the magnitude of the compositional effect in the non-commodity sector, which can be either negative, as in Panel (b), or positive, as in Panels (a) and (c).

In the commodity sector, the shock has two effects on the average wage. First, there is an increase in the log-wage of commodity sector employees implied by $\Delta\omega_g^C > 0$ — i.e., the vertical shift of the blue curve in Figure 1.3. Second, there is a compositional effect driven by the inflow of new employees whose average sector-specific efficiency differs from that of original employees in the commodity sector. As in the non-commodity sector, the sign of this effect is ambiguous, and it is determined by the slope of $\bar{A}_g^C(.)$. In Panels (a) and (b), new commodity sector employees (dashed green) are *less efficient* than original commodity sector employees (blue) and, therefore, the employment expansion leads to a negative compositional effect. In Panel (c), alternatively, new workers are *more efficient* than the original commodity sector employees, implying a positive compositional effect.[11]

To summarize, an increase in world commodity prices that causes an increase in the commodity sector's relative wage per efficiency unit, $\omega_g^C - \omega_g^N$, affects both sectoral employment and sectoral wages. The increase in $\omega_g^C - \omega_g^N$ triggers an increase in the relative employment of the commodity sector whose magnitude is regulated by the

[11]Intuitively, the compositional effect in the commodity sector is captured by a first-order expansion of expression (1.8):

$$\Delta\bar{Y}_g^C - \Delta\omega_g^C = \int_{l_g^N}^{l_g^N+\Delta l_g^N} \frac{\partial\bar{A}_g^C(u)}{\partial q}\, du \approx \left[\alpha_g(l_g^N) + A_g(l_g^N) - \bar{A}_g^C(l_g^N)\right] \cdot \Delta\ln(l_g^C).$$

where $\bar{A}_g^C(l_g^N) = (1/(1-l_g^N)) \cdot \int_{l_g^N}^1 \alpha_g(q) + A_g(q)\, dq$. Notice that the average efficiency in the commodity sector is related to the sum of the schedules of comparative and absolute advantage, $\alpha_g(.) + A_g(.)$.

schedule of comparative advantage, $\alpha_g(.)$. The between-sector worker reallocation introduces compositional effects in the response of the commodity sector's relative average wage. Such an effect may reinforce or diminish the positive impact of the increase in $\omega_g^C - \omega_g^N$. The magnitude of the compositional effect in sectoral average wages is determined, in the non-commodity sector, by $A_g(.)$ and, in the commodity sector, by $\alpha_g(.) + A_g(.)$.

1.2.5 Sectoral Demand Shocks and Aggregate Changes in Wage Inequality

I now turn to movements in wage inequality stemming from the sectoral shock analyzed above. In this analysis, there are many ways of quantifying changes in wage inequality. I focus on the responses in the average and the variance of the log-wage distribution of workers in different demographic groups. The main result of this section establishes that these responses are determined by the schedules of comparative advantage, $\alpha_g(.)$, and absolute advantage, $A_g(.)$.

Let us first analyze the average of the log-wage distribution among workers of group g. For these workers, the average log-wage is $\bar{Y}_g \equiv l_g^N \cdot \bar{Y}_g^N + l_g^C \cdot \bar{Y}_g^C$ which, by equation (1.8), is equivalent to

$$\bar{Y}_g = \omega_g^C \cdot l_g^C + \omega_g^N \cdot l_g^N + \int_{l_g^N}^{1} \alpha_g(q)dq + e_g$$

where $e_g \equiv \int_0^1 A_g(q)dq$.

Following a demand-driven shock in (ω_g^C, ω_g^N), this expression implies that

$$\Delta\bar{Y}_g = \left[\Delta\omega_g^C \cdot l_g^C + \Delta\omega_g^N \cdot l_g^N\right] + \left[\alpha_g\left(l_g^N + \Delta l_g^N\right) \cdot \Delta l_g^N - \int_{l_g^N}^{l_g^N + \Delta l_g^N} \alpha_g(q)\ dq\right]. \quad (1.11)$$

In equation (1.11), the first term is the direct effect on the wage of sector employees if they were unable to reallocate between sectors. This direct effect depends solely on the pre-shock employment composition and the change in sectoral wages per efficiency unit. Nevertheless, this is not the only effect, because workers respond to the shock by switching sector. This composition effect is captured by the second term which, intuitively, depends on the schedule of comparative advantage, $\alpha_g(.)$.[12]

Finally, I compute the log-wage variance among individuals of group g. There are two sources of wage dispersion: the between-sector average wage differential and within-sector wage dispersion. By the law of total variance, these two components imply that the log-wage variance in group g, V_g, is given by

$$V_g = l_g^N l_g^C \cdot \left(\bar{Y}_g^C - \bar{Y}_g^N\right)^2 + l_g^N \cdot V_g^N + l_g^C \cdot V_g^C$$

where V_g^k corresponds to the log-wage variance among individuals of group g employed in sector k.

As indicated in Figure 1.2, the within-sector wage variance, V_g^k, combines the variation in average wage of individuals distributed across the quantiles allocated to the sector, $Var[\bar{Y}_g^k(q)]$, and the absolute advantage dispersion in any particular comparative advantage quantile, $V_g(q)$. Consequently, the log-wage variance of group g is given by

$$V_g = l_g^N l_g^C \cdot \left(\bar{Y}_g^C - \bar{Y}_g^N\right)^2 + l_g^N \cdot Var\left(A_g(q) \middle| q < l_g^N\right) + l_g^C \cdot Var\left(\alpha_g(q) + A_g(q) \middle| q \geq l_g^N\right) + \nu_g \tag{1.12}$$

where the variance is taken over the conditional uniform distribution of quantiles allocated to each sector, and $\nu_g \equiv \int_0^1 V_g(q)\, dq$ is the average dispersion in absolute advantage. Because absolute advantage affects log-wage dispersion equally in the two sectors, ν_g does not depend on the sector employment composition.

Expression (1.12) immediately implies that the change in the log-wage variance is determined by the schedules of comparative advantage, $\alpha_g(.)$, and absolute advantage,

[12]The compositional effect is second-order: for small shocks, sector-switchers are the marginal individuals with the same potential wage in the two sectors, so their reallocation does not affect the log-wage distribution. As a result, the change in the average log-wage, up to a first-order approximation, only depends on the initial allocation of workers across sectors. This intuition can be extended to the wage growth across quantiles of the log-wage distribution. In a first-order approximation, it depends exclusively on the pre-shock sectoral allocation of workers in each quantile. I return to this discussion in detail in Section 1.4.2 and in Appendix 1A.1.2.

$A_g(.)$. Such a change comprises two terms: the change in the sector average wage differential, and the change in the log-wage variance within each sector. Both terms are affected by the compositional effects generated by the sectoral reallocation of workers. The following proposition summarizes this discussion.

Proposition 3. *Conditional on demand-driven changes in (ω_g^C, ω_g^N), the schedules of comparative advantage, $\alpha_g(.)$, and absolute advantage, $A_g(.)$, determine the changes in the average and the variance of the log-wage distribution of workers in group g.*

In the rest of the paper, I build upon Propositions 1–3 to construct an empirical strategy to quantify the distributional effects of shocks in world commodity prices. First, I show how Propositions 1–2 can be used to establish the nonparametric identification of the schedules of comparative advantage, $\alpha_g(.)$, and absolute advantage, $A_g(.)$. This result relies on the intuition of the comparative statics exercise in Section 1.2.4, where these schedules determine the magnitude of the sectoral responses in employment and average wage implied by sector demand shocks. Second, I use the model's predicted response in the average and the variance of the log-wage distribution in Proposition 3 to quantify the effect on wage inequality of shocks in world commodity prices. This delivers an empirical framework to analyze the effect of world commodity prices on Brazilian wage inequality.

1.3 Identification of Comparative and Absolute Advantage

The goal of this section is to establish the nonparametric identification of the schedules of comparative and absolute advantage. The challenge inherent in identifying these functions is illustrated by Heckman and Honoré's (1990) result that, in the context of the Roy model, the sector-specific productivity distribution cannot be nonparametrically identified in a single cross-section of individuals. Thus, this section represents an important first step in the empirical application of the model. The nonparametric identification result indicates the source of variation in the data that uncovers comparative and absolute advantage. Such a result does not impose additional restrictions beyond those implied by the theory. In contrast, as noted by Matzkin (2007), the credibility of the empirical analysis would be significantly hindered if identification could only be achieved under restrictive parametric assumptions.

To this extent, I explore the distinct roles of comparative and absolute advantage in determining sectoral employment and sectoral average wage, as described in Propositions 1 and 2. Following a sector demand shock, the schedule of comparative advantage determines the between-sector response of employment. Simultaneously, the schedule of absolute advantage determines the compositional effects embedded in the response of sector average wages. Reflecting these conceptually different effects, the main result of this section establishes that the schedules of comparative and absolute advantage are nonparametrically identified from cross-regional variation in the sectoral responses of employment and wages to observable sector-level demand shifters.

1.3.1 Assumptions

In order to establish identification of comparative and absolute advantage, I make additional assumptions regarding observable labor market outcomes, as well as their relation to unobservable variables.

Segmented Labor Markets. Consider the set of regional economies with segmented labor markets generated by the model in Section 1.2. Each regional market is indexed by m. For workers in a demographic group g, I assume that there is observable information on sector employment composition, $l^k_{g,m}$, sector average wages, $\bar{Y}^k_{g,m}$, and sector wage per efficiency unit, $\omega^k_{g,m}$.[13] In addition, I assume that the productivity distribution in every market m satisfies the following conditions.

[13]In this section, I treat $\omega^k_{g,m}$ as observable variables determined in the competitive equilibrium of each region. Section 1.4.2 provides a methodology to estimate changes in the wage per efficiency unit, $\Delta\omega^k_{g,m}$, based on the model's predicted relation between wage growth and initial sector employment across quantiles of the log-wage distribution.

Assumption 1. *Individual i in market m, $i \in \mathcal{I}_{g,m}$, independently draws $(s_g(i), a_g(i))$ as follows.*

i. ***Comparative Advantage:***

$$s_g(i) = \tilde{s}_g(i) + \tilde{u}_{g,m} \quad \text{and} \quad \{\tilde{s}_g(i)\} \sim F_g(s)$$

where $\tilde{u}_{g,m}$ is a group-market shifter of comparative advantage, and $\alpha_g(q) \equiv (F_g)^{-1}(q)$.

ii. ***Absolute Advantage:***

$$\{a_g(i)|\tilde{s}_g(i) = s\} \sim \mu H^a_g(a|s) + (1-\mu)H^e_{g,m}(a)$$

where $H^e_{g,m}(a) \equiv H^e(a|\tilde{u}_{g,m}, \boldsymbol{\theta}_{g,m})$ is a group-market mixing distribution of absolute advantage such that

$$A_g(q) \equiv \mu \int a \; dH^a_g(a|\alpha_g(q)) \quad \text{and} \quad \tilde{v}_{g,m} \equiv (1-\mu)\int a \; dH^e_{g,m}(a).$$

Assumption 1 imposes no restrictions on the shape of the productivity distribution, allowing it to vary arbitrarily between worker groups. However, the productivity distribution is assumed to only vary across markets with respect to market-specific shifters in comparative and absolute advantage. Specifically, $\tilde{u}_{g,m}$ represents a shock to the relative efficiency of workers in the commodity sector. Also, $\tilde{v}_{g,m}$ is a shifter of the average absolute advantage of workers in the market, capturing supply shocks to workers' productivity in the non-commodity sector. Assume that these supply shocks combine observable and unobservable components as follows.

Assumption 2. *The shifters of comparative and absolute advantage, $(\tilde{u}_{g,m}, \tilde{v}_{g,m})$, are given by*

$$\tilde{u}_{g,m} = \mathbf{X}_{g,m}\gamma^u_g + u_{g,m}, \quad \text{and} \quad \tilde{v}_{g,m} = \mathbf{X}_{g,m}\gamma^v_g + v_{g,m}$$

where $\mathbf{X}_{g,m}$ is an observable vector of group-market variables; and $(u_{g,m}, v_{g,m})$ is an *unobservable vector of group-market supply shocks. These shifters are normalized such that* $E[\tilde{u}_{g,m}] = E[\tilde{v}_{g,m}] = 0$.

Under Assumptions 1 and 2, the supply equations determining sector employment composition in (1.7) and sector average wage in (1.8) are equivalent to

$$\left[\omega^N_{g,m} - \omega^C_{g,m}\right] = \alpha_g\left(l^N_{g,m}\right) + \mathbf{X}_{g,m}\gamma^u_g + u_{g,m} \tag{1.13}$$

$$\left[\bar{Y}^N_{g,m} - \omega^N_{g,m}\right] = \bar{A}^N_g\left(l^N_{g,m}\right) + \mathbf{X}_{g,m}\gamma^v_g + v_{g,m} \tag{1.14}$$

$$\left[\bar{Y}^C_{g,m} - \omega^C_{g,m}\right] = \bar{A}^C_g\left(l^N_{g,m}\right) + \mathbf{X}_{g,m}(\gamma^u_g + \gamma^v_g) + (u_{g,m} + v_{g,m}). \tag{1.15}$$

Equations (1.13)–(1.15) highlight the importance of Assumption 1: it implies identical patterns of selection across markets in the form of the common schedules of comparative advantage, $\alpha_g(.)$, and absolute advantage, $A_g(.)$. In this context, the pair of unobservable productivity shifters, $(u_{g,m}, v_{g,m})$, generates variation in the sector-specific productivity distribution across markets. Accordingly, Assumptions 1 and 2 are weaker than Heckman and Honoré's (1990) restriction of an identical sector-specific productivity distribution in every market. To the extent that workers of a particular group are different in terms of sectoral labor efficiency across markets, the flexibility implied by the unobserved productivity shifters is important in the empirical application of the model. In fact, variation in these shifters translates into variation in the effective labor supply in the two sectors, generating simultaneous general equilibrium responses in sector wage per efficiency unit, $\omega^k_{g,m}$, and sector employment composition, $l^N_{g,m}$. As a result, identification of $\alpha_g(.)$ and $A_g(.)$ based on the supply equations (1.13)–(1.15) requires a sector demand shock that is orthogonal to the productivity shifters in the cross-section of markets.[14]

[14]Thorough the lens of Assumption 1, Heckman and Honoré's (1990) restriction is equivalent to imposing $\tilde{u}_{g,m} = 0$ and $\mu = 1$. In this case, markets are not subject to unobserved supply shocks and, therefore, any cross-market variation in sector employment composition is necessarily generated by sectoral demand shocks. Consequently, the cross-market variation in wage per efficiency unit leads to the identification of $\alpha_g(.)$ and $A_g(.)$ from equations (1.13)–(1.15) with $u_{g,m} = v_{g,m} = 0$.

Instrument: Sector Demand Shifter. Consider an observable vector, $\mathbf{Z}^k_{g,m}$, of sector demand shifters across markets. To be a valid instrument in the supply equations (1.13)–(1.15), this sector demand shifter has to be mean independent from unobserved shocks to the productivity distribution, $(u_{g,m}, v_{g,m})$. Thus, I assume that $\mathbf{Z}^k_{g,m}$ satisfies the following exogeneity restriction.

Assumption 3. $E\left[u_{g,m}|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}\right] = E\left[v_{g,m}|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}\right] = 0.$

Additionally, $\mathbf{Z}^k_{g,m}$ has to induce enough exogenous variation in the endogenous sector composition $l^N_{g,m}$ to uniquely discriminate the underlying productivity distribution of the economy. In the environment introduced in Section 1.2, this shifter must affect sectoral labor demand differentially across markets. Formally, the instrument has to satisfy the equivalent of a rank requirement in the context of nonparametric models. As shown by Newey and Powell (2003b), the necessary and sufficient completeness condition that guarantees identification of the class of models covering equations (1.13)–(1.15) is described as follows.

Assumption 4. *For any $f(.)$ with finite expectation, $E\left[f(l^N_{g,m}, \mathbf{X}_{g,m})\big|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}\right] = 0$ implies $f(l^N_{g,m}, \mathbf{X}_{g,m}) = 0$ almost surely.*

1.3.2 Nonparametric Identification of Comparative and Absolute Advantage

With Assumptions 1-4, I now establish the identification of the schedules of comparative and absolute advantage. Under Assumptions 3–4, the observable shifter $\mathbf{Z}^k_{g,m}$ generates exogenous variation in the sector composition $l^N_{g,m}$ that can be used to identify equations (1.13)–(1.15). To formalize this intuition, I demonstrate in Appendix 1A.1.1 the following particular case of the general result in Newey and Powell (2003b) regarding the nonparametric identification of separable models with endogenous variables.

Lemma 1. ***[Newey and Powell (2003)]*** *Consider a model of the form*

$$y_{g,m} = \Phi_g\left(l_{g,m}^N\right) + \mathbf{X}_{g,m}\gamma_g + u_{g,m},$$

and a vector $\mathbf{Z}_{g,m}^k$ *satisfying Assumptions 3–4. Then, the function* $\Phi_g(.)$ *is identified up to a constant. With the normalization* $E[\mathbf{X}_{g,m}\gamma_g] = 0$*, the constant in* $\Phi_g(.)$ *is also identified.*

Notice that, under Assumptions 1–2, the supply equations in (1.13)–(1.15) belong to the class of models covered by Lemma 1. Thus, the instrument $\mathbf{Z}_{g,m}^k$ satisfying Assumptions 3-4 identifies $\alpha_g(.)$ from equation (1.13). Similarly, Lemma 1 establishes the identification of $\bar{A}_g^N(.)$ and $\bar{A}_g^C(.)$ respectively from equations (1.14) and (1.15). This leads to the identification of $A_g(.)$ and $A_g(.) + \alpha_g(.)$ since, by the definition in (1.8),

$$A_g(q) = \frac{\partial}{\partial q}\left[q \cdot \bar{A}_g^N(q)\right] \quad \text{and} \quad \alpha_g(q) + A_g(q) = -\frac{\partial}{\partial q}\left[(1-q)\cdot \bar{A}_g^C(q)\right].$$

Hence, I establish the following theorem.

Theorem 1. *Consider a set of segmented markets, m, subject to a sector demand shifters,* $\mathbf{Z}_{g,m}^k$*, such that Assumptions 1–4 hold. For each worker group g,*

i. $\alpha_g(.)$ *is identified from equation* (1.13)*;*

ii. $A_g(.)$ *is identified from equation* (1.14)*; and*

iii. $\alpha_g(.) + A_g(.)$ *is identified from equation* (1.15).

Theorem 1 is directly related to the comparative statics exercise in Section 1.2.4. Figure 1.4 illustrates a situation where the demand shifter $\mathbf{Z}_{g,m}^k$ induces a change in the commodity sector's wage per efficiency unit. The vertical shift in the commodity sector curve of potential wage triggers between-sector worker reallocation represented by $\Delta l_{g,m}^N$. Conditional on $\Delta\omega_{g,m}^C - \Delta\omega_{g,m}^N$, the magnitude of the change in sector employment composition is determined by the difference between the slopes of the two curves of potential sector wages: the schedule of comparative advantage, $\alpha_g(.)$, in equation (1.13). The subsequent change in the composition of sector employees triggers an observable response in the measured sector average efficiency, $\Delta[\bar{Y}^k_{g,m} - \omega^k_{g,m}]$. In each sector, the magnitude of this compositional effect is determined by the function $\bar{A}^k_g(.)$ in equations (1.14)–(1.15). Because the compositional effect corresponds to the difference between the average sector-specific efficiency of switchers and stayers, $A_g(.)$ is identified from the average efficiency change in the non-commodity sector, and $\alpha_g(.) + A_g(.)$ is identified from the average efficiency change in the commodity sector.

As a corollary of Theorem 1, the schedules of comparative advantage, $\alpha_g(.)$, and absolute advantage, $A_g(.)$, are identified with only two out of the three supply equations in (1.13)–(1.15). In other words, the model is overidentified whenever employment and average wages are available for the two sectors of the economy. The model's overidentification relies on the fact that sector-specific efficiency is the sole determinant of both sectoral wages and sectoral employment. Accordingly, the presence

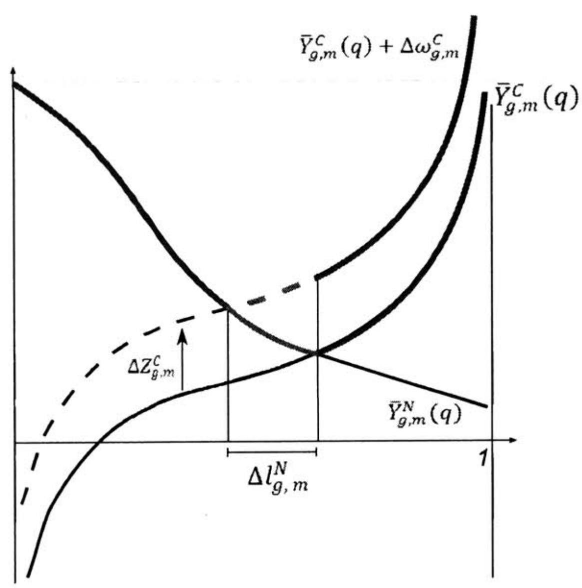

Figure 1.4: Identification of Comparative and Absolute Advantage
caption*Red: sector-stayers in the non-commodity sector. Blue: sector-stayers in the commodity sector. Green: switchers from the non-commodity sector to the commodity sector.

of non-monetary employment benefits breaks overidentification. Appendix 1A.2 establishes the nonparametric identification of the extended model, where workers have heterogeneous private values of employment. In this case, I define a generalized notion of comparative advantage that includes these private benefits. The implied schedule of comparative advantage is identified from equation (1.13). In addition, the schedules of sector-specific average efficiency are identified from equations (1.14)–(1.15).

1.4 Empirical Application

The above result establishes the nonparametric identification of the schedules of comparative and absolute advantage using cross-market variation in observable shifters of sectoral labor demand. Armed with this theoretical result, I now estimate these schedules in a sample of Brazilian regional labor markets differentially exposed to shocks in world commodity prices. I then use these estimates to investigate the effect on Brazilian wage inequality of shocks in world commodity prices.

1.4.1 Sample of Regional Labor Markets and Exposure to World Commodity Price Shocks

The empirical application relies on wage and employment data from the Brazilian Census collected by the Brazilian Institute of Geography and Statistics (IBGE) for 1991, 2000, and 2010. In order to implement the identification strategy outlined above, it is necessary to construct a sample of segmented labor markets. To this end, I use Brazilian regional labor markets as implied by the microregion concept in the Census. The IBGE defines these microregions by aggregating economically integrated municipalities with similar production and geographic characteristics. For each microregion, I select a sample of full-time white employed males aged 16–64. Workers in the sample have strong labor force attachment, diminishing the importance of endogenous responses in total labor supply. I allocate individuals to a group of education (High School Graduates and High School Dropouts) and a sector of employment (commodity and non-commodity).[15] Industries specialized in the production of agricultural and mining products are included in the commodity sector. All manufacturing and service industries are included in the non-commodity sector. In 1991, the commodity sector accounted for 5.2% of employment among High School Graduates (HSG) and 26.2% among High School Dropouts (HSD). In the analysis, I only consider those microregions with positive employment in the commodity sector for all years and groups. As a result, the final sample contains 518 microregions that represented 98.4% of the country's population in 1991. Appendix 1A.3 discusses details on the construction and measurement of labor market outcomes.

As a sectoral demand shifter, I construct a regional measure of exposure to shocks in world commodity prices separately for HSG and HSD. Specifically, the exposure vector of group g in microregion r to commodity price shocks at year t is given by

$$\Delta \mathbf{Z}^{C}_{g,r,t} = \left\{ \phi^{C,j}_{g,r} \cdot \Delta \ln p^{j}_{t} \right\}_{j \in \mathcal{J}^{C}} \tag{1.16}$$

where $\Delta \ln p^{j}_{t}$ is the log-change in the international price of product j between years $t-1$ and t; and $\phi^{C,j}_{g,r}$ is the share of industry j in total labor payments of the commodity sector to individuals of group g in microregion r on the initial year of 1991.

I construct the exposure measure in equation (1.16) using world prices of five major commodity groups: Grains, Soft Agriculture, Livestock, Mining, and Energy. As described in Appendix 1A.3, I compute price indices for each category with data on commodity transactions in the main exchange markets of the United States. To replicate relative prices faced by producers in Brazil, I convert world commodity prices to Brazilian currency and deflate by the Brazilian consumer price index.

The exposure measure in (1.16) is based on the intuition that the response of the commodity sector's labor demand is stronger in regions that specialize in the

[15]These two educational groups are representative of the Brazilian workforce: among male workers, the High School graduation rate was 22.2% in 1991 and 44.5% in 2010. I restrict the benchmark sample to include only white and male individuals because of the strong declines in gender and race wage differential between 1995 and 2010; see Ferreira, Firpo, and Messina (2014). The model in this paper does not speak directly to these components of wage inequality and, therefore, I exclude their behavior from the baseline empirical analysis. In robustness exercises, I extend the benchmark sample to also include female and non-white workers.

production of basic products experiencing stronger international price gains. In the model of Section 1.2, this demand shifter is generated by the limited supply of natural resources specific to commodity production — e.g, fertile soil, or oil and metal reserves. In such an environment, an increase in the world product price triggers an increase in the labor demand of firms producing that product whose effect on the commodity sector's labor demand is proportional to the product's importance in local employment.[16]

Given the shock to world commodity prices, the cross-regional variation in $\Delta \mathbf{Z}^{C}_{g,r,t}$ depends entirely on the cross-regional variation in initial industry composition among workers of a particular group. As shown in Table 1A.3, the great extent of such variation in Brazil implies large variation in shock exposure across regions. This is illustrated in Figure 1.5, which exhibits the total shock exposure across microregions for HSG (left panel) and HSD (right panel). As a consequence of the difference in industry allocation for the two groups, shock exposure differs significantly between HSG and HSD — specifically, the correlation in group exposure is .493.

Exogeneity Assumption. In the empirical application, the regional shifter of sector labor demand must satisfy the central exogeneity restriction imposed in Section 1.3: $\Delta \mathbf{Z}^{C}_{g,r,t}$ has to be uncorrelated with regional shocks to sectoral worker efficiency. This requirement is likely to hold for the following three reasons.

First, Brazilian regions are small relative to the world market of basic commodities, implying that local supply shocks are unlikely to affect international prices. Any national shock correlated across microregions is captured by the time fixed effect included in the specification below. Furthermore, the 1991–2010 period was marked by strong growth in Chinese imports of agriculture and mining products. A growth

[16]Recent empirical papers have built on related measures of local shock exposure in order to investigate the labor market effect of import competition — e.g., see Topalova (2010), Kovak (2013), and Autor, Dorn, and Hanson (2013). In my model, the fixed supply of industry-specific factors guarantees the finiteness of the elasticity of industry labor demand to product price shocks. Thus,

$$\frac{\partial \log \left(\sum_{j \in \mathcal{J}^C} L^j_{g,r}\right)}{\partial \log p^j} = \phi^{C,j}_{g,r} \cdot \frac{\partial \log L^j_{g,r}}{\partial \log p^j}.$$

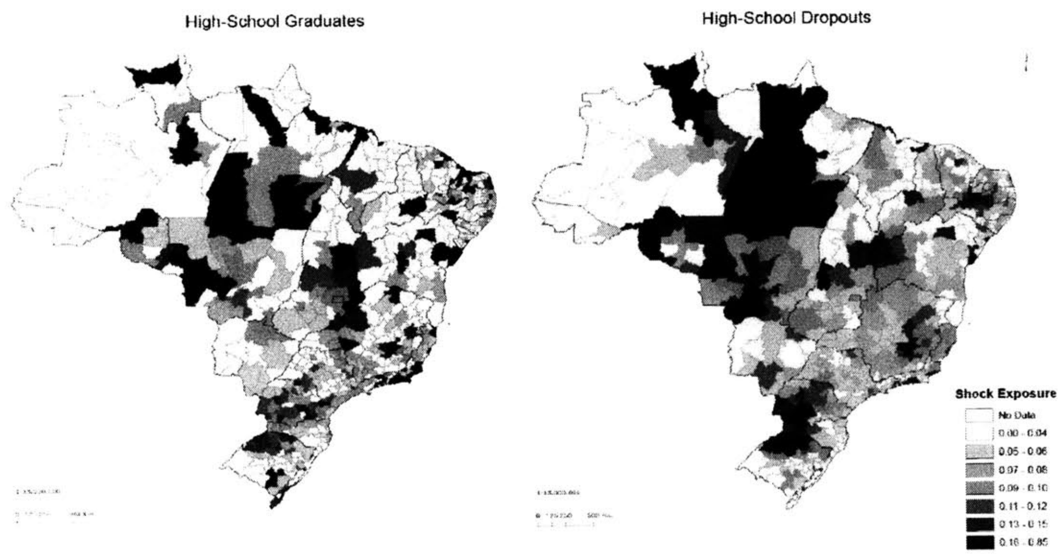

Figure 1.5: Exposure to Commodity Price Shock, 1991-2010

Note. For each microregion, the map presents the total exposure to the commodity price shock between 1991 and 2010: $\sum_{j\in\mathcal{J}^C} \phi_{g,r}^{C,j} \cdot \Delta \ln p_t^j$ where $\Delta \ln p_t^j$ is the log-change in the world price of product j in 1991–2010.

which, arguably, represents an exogenous demand shock to the relative price of raw materials.[17]

Second, the exogeneity restriction requires regional shock exposure to not affect the productivity distribution of workers. This requirement would be violated if the pool of workers in the market varies in response to commodity price shocks because of changes in the labor supply of either native or immigrant workers. In the empirical application, such a concern is unlikely to be important, because the correlation between regional shock exposure and changes in the labor supply of both native and immigrants is small and nonsignificant. This result is partially driven by the inclusion of only full-time prime-aged males in the benchmark sample. See Table 1A.8 in Appendix 1A.4.2.[18]

[17]Between 1992 and 2010, the average annual growth rate of Chinese imports was 17.2% for all products, 16.2% for Agriculture, and 28.3% for Mining. Over the period, Hanson (2012) provides a careful discussion of the transformation in the profile of international trade of emerging economies and, in particular, China. To the extent that this transformation was mainly driven by internal changes in the production structure of China, this large demand shock represented an exogenous impulse to world commodity prices in the period.

[18]Recently, Dix-Carneiro and Kovak (2015a) also find weak responses in migration flows across Brazilian regional labor markets differentially exposed to the tariff reductions during the trade

Third, the empirical application includes a variety of controls intended to capture changes in the productivity distribution potentially correlated with exposure to higher commodity prices. In particular, the control vector includes region-group fixed effects and period dummies interacted with initial regional characteristics (e.g., sector composition and socio economic variables). In this context, identification relies exclusively on the cross-region variation in the exposure to shocks in relative product prices within the commodity sector, allowing for arbitrary shocks to relative prices products in the non-commodity sector.

1.4.2 Estimation of Sector Wage per Efficiency Unit

The identification strategy of Section 1.3 requires information on sector wage per efficiency unit, $\omega^k_{g,r,t}$, for each triple of group-region-period. In this section, I propose a methodology to estimate $\Delta\omega^k_{g,r,t}$ using available information on labor income and employment at the individual level.

In the model, comparative advantage determines worker allocation across sectors, implying that the wage of sector employees is only exposed to changes in the wage per efficiency unit of their own sector of employment. Following changes in $\omega^k_{g,r,t}$, this observation implies that, across different parts of the wage distribution, variation in the pre-shock sector employment composition translates into variation in the growth of wages. Intuitively, if all individuals at the bottom of the distribution are employed in the commodity sector, then the wage gain at the bottom is entirely attributed to the change in the commodity sector's wage per efficiency unit. In such case, an increase in the non-commodity sector's wage per efficiency unit has no impact on the wage of individuals at the bottom of the wage distribution.

To formalize this intuition, let $Y_{g,r,t}(\pi)$ denote the π-quantile of the log-wage distribution of group g in region r at year t. For small shocks, I show in Appendix 1A.1.2 that the wage growth between periods t_0 and t in quantile π of the log-wage

liberalization of 1990–1995.

distribution is given by

$$\Delta Y_{g,r,t}(\pi) = \Delta\omega^{C}_{g,r,t} + \left[\Delta\omega^{N}_{g,r,t} - \Delta\omega^{C}_{g,r,t}\right] \cdot l^{N}_{g,r,t_0}(\pi) + \mu_{g,r,t} \cdot \tilde{\mathbf{X}}_{g,r,t}(\pi) + \Delta v_{g,r,t}(\pi) \quad (1.17)$$

where, at quantile π of the log-wage distribution, $l^{N}_{g,r,t_0}(\pi)$ is the initial employment share of the non-commodity sector and $\tilde{\mathbf{X}}_{g,r,t}(\pi)$ is a set of observable controls. In equation (1.17), $\Delta v_{g,r,t}(\pi)$ is a shock to the absolute advantage of workers in quantile π of the log-wage distribution.[19]

For each group-region-period, equation (1.17) implies that $\Delta\omega^{k}_{g,r,t}$ can be consistently estimated from the relation between the initial sector composition, $l^{N}_{g,r,t_0}(\pi)$, and the wage growth, $\Delta Y_{g,r,t}(\pi)$, across quantiles π of the log-wage distribution. In this context, an estimator of $\Delta\omega^{k}_{g,r,t}$ based on equation (1.17) relies on the assumption that, conditional on the set of controls $\tilde{\mathbf{X}}_{g,r,t}(\pi)$, pre-shock variation in sector employment composition is uncorrelated with variation in labor efficiency shocks among individuals with different levels of labor income in a particular group-region-period. This estimator hinges on a central feature of the Roy model embedded in equation (1.17): the indifference of marginal individuals between the two sectors. For small price shocks, sector-switchers are the marginal individuals with an identical potential wage in the two sectors, implying that their reallocation has no first-order impact on the group's log-wage distribution.[20]

Armed with the model's prediction in equation (1.17), I proceed to estimate $(\Delta\omega^{C}_{g,r}, \Delta\omega^{N}_{g,r,t})$ by regressing wage growth between two consecutive years of the Census, $\Delta Y_{g,r,t}(\pi)$, on the initial year's sector employment composition, $l^{N}_{g,r,t_0}(\pi)$, in a

[19]In Appendix 1A.1.2, I show that equation (1.17) is generated by a first-order expansion of the implicit equation defining $Y_{g,r,t}(\pi)$. In this context, $\Delta v_{g,r,t}(\pi)$ is a shock to the absolute advantage of individuals spread across quantiles of the log-wage distribution. It is introduced by shocks to $(\tilde{u}_{g,r,t}, \boldsymbol{\theta}_{g,r,t})$ that affect the market-specific mixing distribution of absolute advantage, $H^{e}_{g,r,t}(a) \equiv H^{e}(a|\tilde{u}_{g,r,t}, \boldsymbol{\theta}_{g,r,t})$ in Assumption 1. The change in the mixing distribution of absolute advantage has consequences for the labor efficiency of individuals at different income levels. As a result, the process generating innovations in $(\tilde{u}_{g,r,t}, \boldsymbol{\theta}_{g,r,t})$ creates, through $\Delta v_{g,r,t}(\pi)$, idiosyncratic shocks to wage growth across quantiles of the wage distribution.

[20]Expression (1.17) is modified whenever there exists a wedge in sector potential wage of sector-switchers. This is the case in the presence of non-monetary benefits of employment. To the extent that sector-switchers are spread over the wage distribution, the wedge affects wage gains across quantiles. Consistent with this intuition, there is a new term in equation (1.17) that is proportional to the fraction of sector-switchers among individuals at the π-quantile of the log-wage distribution.

set of wage distribution quantiles. For each of the 2,072 group-region-period triples, I implement this regression with 88 bins of 1 p.p. width between the 6^{th} and the 94^{th} percentiles of the wage distribution. The baseline specification contains a set of controls intended to capture potential confounding effects related to differential efficiency growth across workers of various levels of income. These controls include dummies for wage distribution ranges (bottom, middle, top), and dummies for earnings below the minimum wage. Thus, $\Delta\omega_g^k$ is identified from the variation in pre-shock sector employment across small neighborhoods of the log-wage distribution in a group-region-period. Appendix 1A.4.1 provides details regarding the implementation of this methodology along with an investigation of the robustness of estimates to implementation choices.

To evaluate the impact of exposure to commodity price shocks on changes in sectoral wage per efficiency unit, consider the following regression:

$$\Delta\omega_{g,r,t}^k = \beta_g^k \cdot \left[\sum_{j\in\mathcal{J}^C} \phi_{g,r}^{C,j} \cdot \Delta \ln p_t^j\right] + \Delta\mathbf{X}_{g,r,t}\gamma_g^k + \Delta e_{g,r,t} \tag{1.18}$$

where $\Delta\omega_{g,r,t}^k$ is the estimated wage per efficiency unit; and $\mathbf{X}_{g,r,t}$ is a control vector of group-region characteristics potentially correlated with the exposure measure. In the baseline specification, I include period dummies interacted with five macroregion dummies, and I weight microregions by their 1991 share in the national population.[21] Also, I cluster standard errors by microregion to account for serially correlated shocks.

Table 1.1 reports the estimation of equation (1.18) in the sample of Brazilian microregions in 1991–2000 and in 2000–2010. The positive and statistically significant coefficients in column (1) indicate that, for both HSG and HSD, regional exposure to

[21] As discussed in Appendix 1A.4.1, the precision of the estimated sectoral wage per efficiency unit is related to the number of individuals in the microregion. For efficiency purposes, I follow the standard approach of weighting regressions by the population size of the microregion. Alternatively, regions could be weighted by the inverse of the standard error of the estimated wage per efficiency unit. In the baseline specification, I adopted the simple weight by population share for two reasons. First, sectoral regressions would entail different regional weights due to the difference in standard errors in the estimate of each sector's wage per efficiency unit. Second, inference in equation (1.17) is nonstandard, requiring a computationally burdensome bootstrap procedure for each of the 2,072 triples of group-period-region.

Table 1.1: **Exposure to Commodity Price Shocks and Sector Wage per Efficiency Unit**

	Commodity sector			Non-commodity sector		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: High School Graduates						
Commodity price shock	0.960+	1.369**	0.962**	0.410**	0.351**	0.282**
	(0.530)	(0.405)	(0.359)	(0.088)	(0.072)	(0.062)
R^2	0.200	0.550	0.598	0.552	0.560	0.592
Panel B: High School Dropouts						
Commodity price shock	1.977*	1.651*	1.381*	-0.239	-0.028	-0.021
	(0.835)	(0.732)	(0.624)	(0.167)	(0.108)	(0.087)
R^2	0.272	0.646	0.673	0.193	0.484	0.575
Baseline Controls						
Initial commodity sector size	No	Yes	Yes	No	Yes	Yes
Period dummy interaction:						
Initial commodity sector size	No	No	Yes	No	No	Yes
Initial labor market conditions	No	No	Yes	No	No	Yes
Initial manufacturing sector size	No	No	Yes	No	No	Yes

Note. Stacked sample of 518 microregions in 1991–2000 and in 2000–2010. All regressions are weighted by the microregion share in national population in 1991 and include ten macroregion-period dummies. Commodity sector size controls: quadratic polynomial of commodity sector share in group labor income and dummy for commodity sector share in group labor income in the bottom and top deciles of national distribution. Labor market conditions: quadratic polynomial of per-capita income, share of white employees, share of employed individuals, share of formal sector employees, share of individuals earning less than the federal minimum wage. Standard Errors clustered by microregion ** $p<0.01$, * $p<0.05$, + $p<0.10$

higher commodity prices triggers an increase in the commodity sector's wage per efficiency unit. With the aim of eliminating potentially confounding effects, I augment the model with a set of flexible controls for the initial sector composition in the region. In this case, estimation relies on cross-regional variation in exposure to higher relative product prices within the commodity sector. Although these additional controls absorb a large part of the cross-section variation in $\Delta\omega^C_{g,r,t}$, they do not substantially alter evaluated estimates, which actually become more precisely estimated.

Lastly, column (3) includes period dummies interacted with initial labor market conditions and non-commodity sector composition controls. These controls represent period-specific effects projected on initial region characteristics, capturing, for example, effects related to the introduction of cash transfer programs and secular differences in sector productivity growth. In column (3), the response of the commodity sector's wage per efficiency unit to shock exposure is economically large: a 10% increase in commodity prices induces an increase in the commodity sector's wage per efficiency unit of 9.6% for HSG and 14% for HSD.

Columns (4)–(6) of Table 1.1 present the estimation of equation (1.18) for the non-commodity sector's wage per efficiency unit. The estimated coefficients indicate that exposure to higher commodity prices entails a much weaker effect on the wage per efficiency unit in the non-commodity sector. Consequently, there is an increase in the relative wage per efficiency unit of the commodity sector following an increase in world commodity prices. This movement in $\Delta\omega^N_{g,r,t} - \Delta\omega^C_{g,r,t}$ is inconsistent with perfect substitutability of workers in the two sectors of the economy. Therefore, it cannot be generated in traditional trade models in which workers are perfectly exchangeable between sectors. As discussed below, the magnitude of $\Delta\omega^N_{g,r,t} - \Delta\omega^C_{g,r,t}$ is central in the estimation of the schedule of comparative advantage from the subsequent response in sector employment composition.[22]

1.4.3 Parametric Restrictions: Log-Linear System

The nonparametric identification result in Section 1.3 is critical to inform the source of variation in the data that recovers comparative and absolute advantage. In practice, data limitations are severe and they may prevent the implementation of a fully flexible estimator capable of nonparametrically recovering the functions of interest. In such cases, auxiliary functional form assumptions on $\alpha_g(.)$ and $A_g(.)$ are particularly useful to increase estimation precision. It is important, however, that these parametric assumptions do not impose artificial restrictions on the model. In the particular case analyzed in this paper, it is particularly relevant that functional forms allow

[22]In the general equilibrium of the model in Section 1.2, the commodity sector demand shifter affects the wage per efficiency unit in the two sectors of the economy. First, there is a response in the commodity sector's wage per efficiency unit implied by the shift in the sector's labor demand. Second, the wage per efficiency unit in the non-commodity sector responds because, as workers move to the commodity sector, the lower supply of effective labor units in the non-commodity sector can only be an equilibrium if firms in the sector face a higher wage per efficiency unit.

for separate roles for comparative and absolute advantage, since they are related to distinct predictions of the model. Accordingly, the benchmark specification in the empirical application is based on the following parametric assumption.

Assumption 5. *Suppose that the schedules of comparative and absolute advantage are given by*

$$\alpha_g(q) \equiv \alpha_g \cdot [\ln(q) - \ln(1-q)] \quad \textit{and} \quad A(q) \equiv \tilde{A}_g + A_g \cdot \ln(q)$$

where $\alpha_g \geq 0$.

Assumption 5 commands constant-elasticity schedules of comparative and absolute advantage. Following the discussion in Section 1.2, the positive parameter α_g controls the dispersion of comparative advantage; alternatively, the parameter A_g controls the pattern of variation in average absolute advantage of individuals distributed across quantiles of comparative advantage.[23] In the empirical application, the parametric restrictions in Assumption 5 are useful for its dimensionality reduction: there are only two parameters to capture separately comparative and absolute advantage.[24]

The system in Assumption 5 is a strict generalization of the system obtained under a Fréchet distribution of sector-specific productivity. Because of its tractability, this distributional assumption is the basis of numerous recent empirical applications of the Roy model — see, for example, Hsieh, Hurst, Jones, and Klenow (2013); Burstein, Morales, and Vogel (2015); and Galle, Rodriguez-Clare, and Yi (2015). As discussed in Appendix 1A.5, the Fréchet distribution leads to a similar log-linear system, but

[23]In Assumption 5, the distributions of comparative and absolute advantage have finite moments for every $\alpha \geq 0$ and $A_g \in \mathbb{R}$. But this is not necessarily true for its moment generating function. As discussed in Appendix 1A.5.3, finite moment generating functions can be guaranteed with bounds on the support of comparative and absolute advantage. Alternatively, one could impose parameter restrictions, $0 \leq \alpha_g < 1$ and $A_g > -1$.

[24]In the spirit of the series estimator proposed by Newey and Powell (2003b), the log-linear system could be augmented to include higher-order polynomials. In the limit, such an expansion would recover nonparametrically functions $\alpha_g(.)$ and $A_g(.)$. Yet, as pointed out by Newey (2013b), the estimation of nonlinearities tends to be accompanied by sharp increases in standard errors, requiring multiple strong instruments. The application in this paper is no exception and, for this reason, the constant-elasticity specification in Assumption 5 is particularly attractive.

it contains a single parameter to control both comparative and absolute advantage. In terms of the system above, the Fréchet distribution requires that $\alpha_g = -A_g$ where $\alpha_g < 1$.

The Fréchet distribution mixes the distinct roles of comparative and absolute advantage emphasized in this paper, with strong consequences for the model's predictions. Namely, it imposes constraints not only on the magnitude of the between-sector reallocation, but also on the pattern of selection into both sectors. In fact, $\alpha_g = -A_g$ implies that the sector wage differential is constant, being unable to replicate the positive correlation between world commodity prices on the relative average wage of commodity sector employees documented below. Also, the Fréchet distribution implies that the log-wage variance is constant among workers of the same demographic group. These implications of the Fréchet distribution are not generated by robust features of the model, and they may prevent the model from capturing the full extent of wage inequality movements observed in the data.[25]

In contrast, the more general log-linear system in Assumption 5 contains parameters that separately control comparative and absolute advantage. This additional degree of freedom enhances the model's ability to capture movements in wage inequality. In particular, the parameters α_g and A_g allow for much more flexible patterns of selection, generating responses in both sector wage differentials and log-wage variance that would not emerge under the Fréchet distribution.

1.4.4 Estimation Procedure

Now we are ready to propose an estimator for the schedules of comparative and absolute advantage directly related to the identification result in Theorem 1. Towards this goal, I take advantage of the parametric restrictions in Assumption 5 to construct a

[25]Appendix 1A.5 provides a detailed discussion on the pattern of sector selection implied by the Fréchet distribution. While the restriction of $\alpha_g = -A_g$ is a direct implication of assuming a Fréchet distribution, the restriction of $\alpha_g < 1$ is necessary to guarantee a finite effective labor supply in each sector. Appendix 1A.5 also discusses the system implied by normally distributed sector-specific productivities — as in Heckman and Sedlacek (1985) and Ohnsorge and Trefler (2007). Although the normal distribution leads to distinct functional forms, the implied system also entails two parameters that parametrize the slopes of the schedules of comparative and absolute advantage.

consistent GMM procedure with moment conditions that use the differential exposure of Brazilian microregions to the variation in international commodity prices in the two period windows of 1991–2000 and 2000–2010.

To this end, let us combine equations (1.13)–(1.15) with the functional forms in Assumption 5 to write the following first-difference system:

$$\Delta\omega^N_{g,r,t} - \Delta\omega^C_{g,r,t} = \alpha_g \Delta \ln\left(l^N_{g,r,t} \big/ l^C_{g,r,t}\right) + \Delta\mathbf{X}_{g,r,t}\gamma^u_g + \Delta u_{g,r,t} \tag{1.19}$$

$$\Delta\bar{Y}^N_{g,r,t} - \Delta\omega^N_{g,r,t} = \Lambda_g \Delta \ln\left(l^N_{g,r,t}\right) + \Delta\mathbf{X}_{g,r,t}\gamma^v_g + \Delta v_{g,r,t} \tag{1.20}$$

$$\Delta\bar{Y}^C_{g,r,t} - \Delta\omega^C_{g,r,t} = -\left(\alpha_g + \Lambda_g\right)\Delta\left[\frac{l^N_{g,r,t}}{l^C_{g,r,t}} \ln l^N_{g,r,t}\right] - \alpha_g \Delta \ln\left(l^C_{g,r,t}\right) + \Delta\mathbf{X}_{g,r,t}\gamma^e_g + \Delta e_{g,r,t} \tag{1.21}$$

where $\mathbf{X}_{g,r,t}$ is a control vector of group-microregion-period variables that include group-microregion fixed effects; and $\Delta\omega^k_{g,r,t}$ is the change in the wage per efficiency unit of sector k estimated with the procedure described in Section 1.4.2.

Conditional on the parameter vector $\Theta_g \equiv (\alpha_g, \Lambda_g, \gamma^u_g, \gamma^v_g, \gamma^e_g)$, equations (1.19)–(1.21) immediately allow the computation of the vector of structural errors: $\mathbf{e}_g(\Theta_g) \equiv [\Delta u_{g,r,t},\ \Delta v_{g,r,t},\ \Delta e_{g,r,t}]_{r,t}$. I combine this error vector with the matrix of instruments in (1.16), $\mathbf{W}_g \equiv \left[\Delta\mathbf{Z}^C_{g,r,t}, \Delta\mathbf{X}_{g,r,t}\right]_{r,t}$, to obtain moment conditions that allow the consistent estimation of Θ_g. Specifically, I use the following GMM estimator:

$$\hat{\Theta}_g = \arg\min_{\Theta_g} \ \mathbf{e}_g(\Theta_g)'\mathbf{W}_g\mathbf{\Phi}\mathbf{W}_g'\mathbf{e}_g(\Theta_g) \tag{1.22}$$

where $\mathbf{\Phi}$ is a matrix of moment weights.[26] As above, microregions are weighted by their share in the national population of 1991, and standard errors are clustered by microregion.

[26] In the baseline specification, I use the optimal weights implied by the two-stage GMM estimator. Below, I attest that similar results are obtained using other matrices of moment weights.

1.4.5 Results

Reduced-Form Estimates

Before turning to the estimates of the comparative and absolute advantage, I investigate the effect of exposure to commodity price shocks on sectoral employment and wages with the following specification:

$$\Delta Y_{g,r,t} = \beta_g \cdot \left[\sum_{j \in \mathcal{J}^C} \phi_{g,r}^{C,j} \cdot \Delta \ln p_t^j \right] + \Delta \mathbf{X}_{g,r,t} \gamma_g + \Delta e_{g,r,t} \tag{1.23}$$

where $\Delta Y_{g,r,t}$ is the change in a labor market outcome for individuals of group g in microregion r between years $t-1$ and t.

Table 1.2: Exposure to Commodity Price Shocks and Sector Employment and Wages

	Commodity sector			Non-commodity sector		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: High School Graduates						
Commodity price shock	0.039**	0.031**	0.035**	0.369**	0.302**	0.441**
	(0.010)	(0.011)	(0.010)	(0.097)	(0.096)	(0.108)
R^2	0.217	0.291	0.413	0.115	0.145	0.217
Panel B: High School Dropouts						
Commodity price shock	0.187**	0.067*	0.061*	-0.158	-0.127	0.030
	(0.068)	(0.030)	(0.029)	(0.149)	(0.150)	(0.155)
R^2	0.236	0.515	0.561	0.148	0.170	0.215
Baseline Controls						
Initial commodity sector size	No	Yes	Yes	No	Yes	Yes
Period dummy interaction:						
Initial commodity sector size	No	No	Yes	No	No	Yes
Initial labor market conditions	No	No	Yes	No	No	Yes
Initial manufacturing sector size	No	No	Yes	No	No	Yes

Note. Stacked sample of 518 microregions in 1991–2000 and in 2000–2010. All regressions are weighted by the microregion share in national population in 1991 and include ten macroregion-period dummies. Commodity sector size controls and labor market conditions as in Table 1.1. Standard Errors clustered by microregion ** p<0.01, * p<0.05, + p<0.10

Columns (1)–(3) of Table 1.2 present the estimation of equation (1.23), with the dependent variable being the commodity sector employment share. Panel A presents estimates for HSG and Panel B for HSD. In line with the correlation documented in Figure 1.1, the positive and statistically significant coefficients indicate that, for both groups, exposure to higher commodity prices induces workers to reallocate from the non-commodity to the commodity sector.[27] In the structural estimation below, the parameter of comparative advantage is implied by the combination of the response in sectoral employment in Table 1.2 and the response in the relative wage per efficiency unit in Table 1.1.

Turning to the impact of commodity prices on sectoral wages, I estimate equation (1.23), with the dependent variable being the commodity sector's relative average wage. Results in columns (4)–(6) indicate different qualitative responses for the two worker groups. The price shock triggers a significant positive response of the commodity sector wage differential for HSG; in contrast, there is only a small and imprecisely estimated response for HSD. For both groups, however, these estimated responses are much smaller than the estimated response in the relative wage per efficiency unit presented in Table 1.1. In the model, the difference between the commodity sector's response in terms of relative average wage and relative wage per efficiency unit corresponds to the compositional effect induced by worker reallocation between sectors. In fact, the magnitude of this difference determines the magnitude of the structural parameters of comparative and absolute advantage presented below.

Lastly, it is important to notice that the positive response in sectoral wages for HSG is inconsistent with the selection pattern implied by a Fréchet distribution. Below, this leads to the rejection of the parametric restrictions required by the Fréchet model in the HSG's structural estimates. For HSD, the weaker response of sectoral wages yields a selection pattern similar to that implied by the Fréchet distribution.

In Appendix 1A.4.2, I investigate the robustness of these results. In particular,

[27]In Table 1A.8 of Appendix 1A.4.2, I show that regional exposure to higher commodity prices does not induce higher labor supply of both native and migrant workers. Thus, the expansion in the commodity sector employment share is driven by the between-sector reallocation of individuals in the market.

Table 1A.6 shows that results are similar if the baseline specification is extended to include the additional period of 1980–1991 and microregion-specific time trends. In addition, Table 1A.7 reports similar qualitative patterns of sectoral responses in employment and wages for additional demographic groups, including female and non-white workers.

Estimated Parameters of Comparative and Absolute Advantage

I now present the estimates of the comparative and absolute advantage parameters obtained with the procedure described in Section 1.4.4. These estimates are reported in Table 1.3 together with their standard errors clustered by microregion. Column (1) reports the structural parameters implied by the estimation of equations (1.19)–(1.21) under the parametric restriction imposed by the Fréchet distribution, $\alpha_g = -A_g$. Estimated parameters indicate that an increase of 1% in the relative wage per efficiency unit of the commodity sector triggers an increase in the relative employment in the commodity sector of approximately 1.2% for both groups (i.e., the inverse elasticity $1/\alpha_g$).

Column (2) presents the estimates obtained under the unrestricted log-linear system in (1.19)–(1.21). In this case, comparative advantage parameters indicate between-sector employment reallocation whose magnitude is similar to that of the Fréchet model in column (1) for both groups. Nevertheless, the additional degree of freedom is important for HSG, as the estimated parameter of absolute advantage changes substantially. Among HSG, the strong response of sectoral wage differentials documented above yields an absolute advantage parameter that indicates negative selection into the non-commodity sector. This parameter implies curves of potential sector wages similar to those displayed in case (b) of Figure 1.3. In contrast, the weak response of sector wage differentials for HSD drives an absolute advantage parameter that indicates a pattern of selection similar to that implied by the Fréchet model. Indeed, the Fréchet restriction cannot be rejected at usual significance levels. For HSD, there is positive selection into both sectors with curves of potential sector wages similar to those shown in case (a) of Figure 1.3.

Appendix 1A.4.3 investigates the robustness of the results presented in Table 1.3. To increase confidence in the baseline GMM estimator, Table 1A.9 presents the estimated structural parameters obtained with the separate 2SLS estimation of equations (1.19) and (1.20). Such an estimator is less efficiency, because it ignoring the overidentification restriction provided by the response in the commodity sector's average wage in (1.21). Despite this fact, point estimates are not only similar in magnitude, but also qualitative conclusions are similar with inference methods robust to weak instruments. Also, I find similar estimated parameters when, as in the reduced-form regressions, the unique instrument is the aggregate exposure to commodity price shocks. Table 1A.10 reports that similar results are implied by the GMM estimator with restricted vectors of excluded instruments and alternative matrices of moment weights. Finally, Table 1A.11 shows that similar results are obtained with alternative specifications in the estimation of sectoral wages per efficiency unit.

1.4.6 Model Fit

In order to build confidence in the model, I investigate the model's ability to generate responses in the log-wage distribution that are consistent with those observed in the data. Thus, I estimate equation (1.23) using both actual data and the model's predictions regarding changes in the average and the variance of log wages across Brazilian microregions. Since the estimation of the structural parameters relied on sectoral responses in terms of employment and average wages, this exercise constitutes a test of the model's goodness of fit.

To implement this test, I compute the model's predicted changes in the average and the variance of the log-wage distribution using, respectively, equations (1.11) and (1.12) derived in Section 1.2. Expressions (1.11)–(1.12) require the changes in sectoral wages per efficiency unit generated by the shock in world commodity prices. I obtain these responses directly from the predicted changes implied by the estimates in columns (3) and (6) of Table 1.1.

Table 1.4 presents the results of this exercise. Let us first analyze the response in the average log-wages presented on the top row of each panel. In this case, both the

Table 1.3: **Estimated Parameters of Comparative and Absolute Advantage**

	Fréchet model $\alpha_g = A_g$	Log-linear model
	(1)	(2)
Panel A: High School Graduates		
α_{HSG}	0.819**	0.835**
	(0.192)	(0.212)
A_{HSG}	-0.819**	1.966*
	(0.192)	(0.935)
Test of Fréchet restriction (p-value)	-	0.005
Panel B: High School Dropouts		
α_{HSD}	0.856**	0.916*
	(0.140)	(0.399)
A_{HSD}	-0.856**	-0.727**
	(0.140)	(0.142)
Test of Fréchet restriction (p-value)	-	0.644

Note. Stacked sample of 518 microregions in 1991–2000 and 2000–2010. Two-Step GMM estimator with microregions weighted by their share in the 1991 national population. All equations include macroregion-period dummies, initial commodity sector size controls, and initial labor market conditions as in Table 1.1. Excluded instruments: quadratic polynomial of regional exposure to world product prices. Standard Errors clustered by microregion ** p<0.01, * p<0.05

Fréchet and the log-linear models deliver responses whose cross-regional relation to shock exposure is consistent with the cross-regional relation in the data. The similar responses with the two specifications follow from the similar estimated parameters of comparative advantage reported in columns (1) and (2) of Table 1.3. To see this, recall that the change in the average log-wage in equation (1.11) only depends on the schedule of comparative advantage.

When we turn to the variance of log-wages within each group, we see in Table 1.4 that the two specifications yield very different responses. For HSG, the log-linear model implies a negative relation whose magnitude is similar to the negative and statistically significant relation in the data. In contrast, the Fréchet model is unable to generate this relation, since it entails a constant log-wage variance. For HSD in Panel B, the cross-region response of log-wage variance to shock exposure is small and imprecisely estimated. This is consistent with the prediction of the Fréchet model. In this case, the log-linear model yields a small positive response.

1.5 Counterfactual Simulation: Effect of World Commodity Prices on Brazilian Wage Inequality

To conclude, I use the estimated schedules of comparative and absolute advantage to investigate the consequences to the Brazilian wage distribution of shocks in world commodity prices. Precisely, I ask: "In 1991, how would wage inequality change if commodity prices were equal to those of 2010?" In order to answer this question, I proceed in two steps. In the first step, I compute the change in sectoral wages per

Table 1.4: Average and Variance of Log-Wages, Model Predictions and Actual Data

	Predicted change Fréchet model	Predicted change Log-linear model	Actual data
	(1)	(2)	(3)
Panel A: High School Graduates			
Change in Log-Wage Average	0.331	0.331	0.262**
			(0.054)
Change in Log-Wage Variance	0.000	-0.121	-0.117*
			(0.049)
Panel B: High School Dropouts			
Change in Log-Wage Average	0.140	0.138	0.166*
			(0.077)
Change in Log-Wage Variance	0.000	0.075	-0.005
			(0.071)
Baseline Controls			
Controls in Table 1.1	Yes	Yes	Yes

Note. Estimated coefficient of the regression of the dependent variable on shock exposure using the stacked sample of 518 microregions in 1991–2000 and 2000–2010. Regressions are weighted by the microregion share in national population in 1991 and include the baseline controls in Table 1.1 and the initial wage dispersion (log-wage variance regressions). Dependent variables in columns (1)–(2) are counterfactual changes implied by the model. Dependent variables in column (3) are actual data. Standard Errors clustered by microregion ** p<0.01, * p<0.05

efficiency unit implied by the shock to world commodity prices. In the second, I use the model's predictions to compute the counterfactual change in the average and the variance of the log-wage distribution implied by the change in wage per efficiency unit. While the second step is a straight forward application of the sufficiency result in Proposition 3 of Section 1.2, the first step is not. In this section, I present the counterfactual changes in Brazilian wage inequality obtained with two alternative procedures to compute changes in sectoral wages per efficiency unit.

As in Section 1.4.6, the first approach uses the estimated pass-through in Table 1.1 to compute the effect of regional exposure to price shocks on the sector wage per efficiency unit. This methodology has the advantage of being robust to parametric restrictions on the economy's structure of production, but it is subject to two shortcomings. First, it does not capture nationwide effects on sectoral wages per efficiency unit, since these are absorbed by period fixed effects included in the regressions. Second, the estimated pass-through is a reduced-form relation that may not hold for different price shocks and different periods. To address these deficiencies, the second approach relies on a fully specified general equilibrium model to obtain the endogenous change in the wage per efficiency unit following exogenous shocks in world commodity prices. This procedure takes inspiration from the exact hat algebra used in recent quantitative papers in international trade — see, e.g., Dekle, Eaton, and Kortum (2007) and Costinot and Rodríguez-Clare (2013). It illustrates how the structural parameters of comparative and absolute advantage can be combined with specific assumptions regarding the production and market structures to investigate the distributional effects of sectoral shocks.

1.5.1 Counterfactual Simulation: Reduced-Form Pass-Through from World Commodity Prices to Sector Wage per Efficiency Unit

In this section, I present the main results of the paper regarding the counterfactual change in Brazilian wage inequality implied by the shock in world commodity prices.

I compute changes in sectoral wages per efficiency unit from the reduced-form pass-through in columns (3) and (6) of Table 1.1. With these variables, I then compute the predicted changes in the average and the variance of the log-wage distribution using equations (1.11) and (1.12) and the estimated parameters of comparative and absolute advantage in Table 1.3.

Counterfactual Changes in Between-Sector Wage Differentials

Figure 1.6 reports the changes in between-sector wage differentials across Brazilian microregions. The top panel shows that there are large changes in the relative sectoral

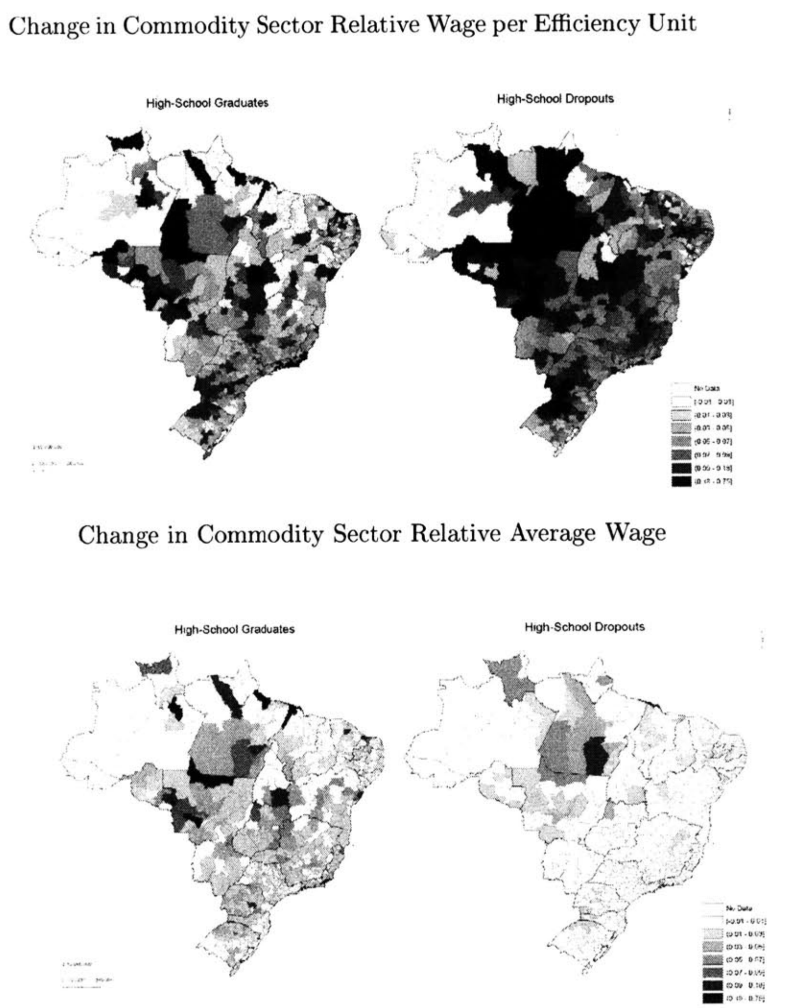

Figure 1.6: Counterfactual Change in Between-Sector Wage Differentials

wages per efficiency unit. Such a change corresponds to the effect of the world price shock on the wage differential of those who do not switch sectors; that is, sector-stayers in the commodity sector versus sector-stayers in the non-commodity sector. At the national level, the change in the relative sectoral wages per efficiency unit is 8% for HSG and 16% for HSD. Reflecting the strong compositional effects implied by the sectoral reallocation of workers, the bottom panel shows that the response in the commodity sector average wage premium is much smaller, with a national average of 1% for both groups. Following the shock, the predicted expansion of the commodity sector in terms of relative employment is 9.2% for HSG and 13.7% for HSD. This implies an average increase in the commodity sector employment share from 5.2% to 5.7% among HSG, and from 26.2% to 28.9% among HSD.

Counterfactual Changes in Brazilian Wage Inequality

Figure 1.7 presents the counterfactual change in average log-wage implied by the rise in commodity prices between 1991 and 2010. The positive price shock triggers average wage gains for the two worker groups. Yet the wage gain is more pronounced among HSD, due to their higher employment share in the commodity sector. Consequently, the shock causes a decrease in the HSG-HSD wage premium of approximately 1.1%.

Table 1.5 reports these counterfactual responses at the national level, where the aggregate log-wage variance is computed with the total variance formula and the microregion's employment share in 1991. Column (2) shows that the price shock triggers a decrease in the log-wage dispersion in the two worker groups. Such a response arises for two reasons. First, regions that specialized in commodity production had lower wages initially, but they experienced stronger wage growth following the positive price shock. Second, the estimated schedules of comparative and absolute advantage imply that, within groups and regions, the shock affects the log-wage variance due to movements in sectoral wage differentials and sectoral employment composition. At the national level, this effect is reinforced by the reduction in the HSG-HSD wage premium triggered by the shock. Panel C shows that 5.6% of the fall in Brazilian log-wage variance is related to the increase in world commodity prices.

Table 1.5: Effect of 1991–2010 Rise in World Commodity Prices on Brazilian Wage Inequality

	Change in Log-Wage Average	Change in Log-Wage Variance (percent of 1991-2010 change)
	(1)	(2)
Panel A: High School Graduates		
	0.039	-0.014
		7.91%
Panel B: High School Dropouts		
	0.050	-0.012
		3.97%
Panel C: All Workers		
	0.047	-0.017
		5.55%

Note. Estimated parameters of comparative and absolute advantage of the log-linear model in column (2) of Table 1.3.

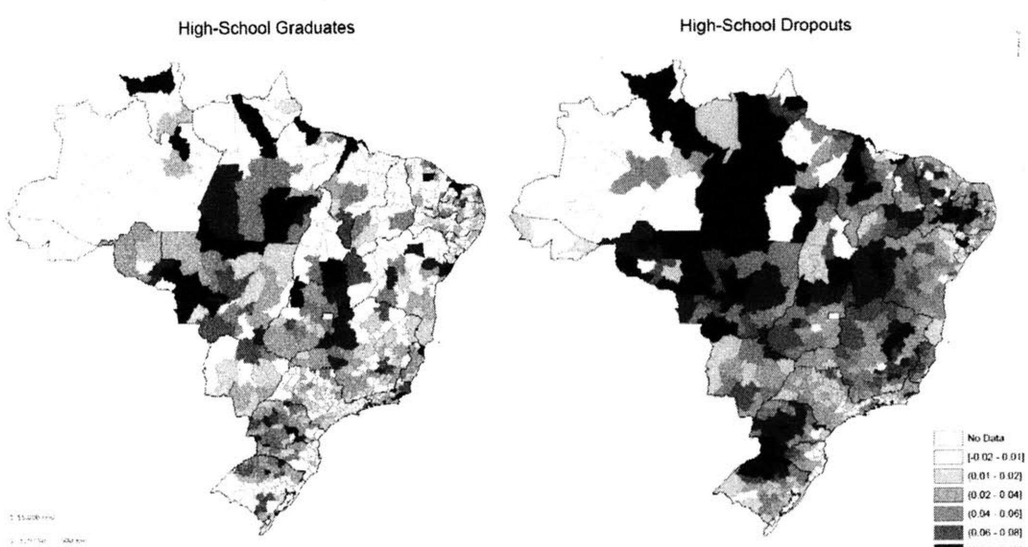

Figure 1.7: Counterfactual Change in Average Log-Wage

1.5.2 Counterfactual Simulation: General Equilibrium Model

In this section, I introduce additional functional form assumptions in the economy presented in Section 1.2. These assumptions allow the computation of changes in sectoral wages per efficiency unit implied by changes in final product prices in a fully specified general equilibrium model. I denote by a "hat" the change in each variable between the initial and the counterfactual equilibrium.

Additional Parametric Assumptions

In the environment of Section 1.2, I assume that there are two demographic groups, High School Graduates and High School Dropouts. In each industry, production utilizes the effective labor supplied by employees of both groups and an industry-specific input. In particular, assume that the production function has the following nested CES structure:

$$q^j = \left(L^j\right)^{\eta^j} \left(X^j\right)^{1-\eta^j} \text{ where } \quad L^j \equiv \left[\beta^j_{HSG} \left(L^j_{HSG}\right)^{\frac{\rho-1}{\rho}} + \beta^j_{HSD} \left(L^j_{HSD}\right)^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}},$$

η^j is the labor share in total revenue of industry j, and X^j is an industry-specific input with fixed supply. In this production function, High School Graduates and High School Dropouts are imperfect substitutes with ρ denoting the constant elasticity of substitution between the effective labor supplied by the two groups.[28]

With this production technology, shocks in product and factor prices cause responses in the labor demand of industry j of sector k that are given by

$$\hat{L}^j_g = \left(\hat{w}^k_g\right)^{-\rho} \left(\hat{W}^j\right)^{\rho-\frac{1}{1-\eta^j}} \left(\hat{p}^j\right)^{\frac{1}{1-\eta^j}}, \tag{1.24}$$

[28]This particular production structure is imposed mainly due to the limited availability of production data for Brazilian microregions. First, I introduce an industry-specific factor as a simplifying device to generate curvature in sector labor demand in every region. Similarly, one could allow, as in Costinot, Donaldson, and Smith (2012b), regions to have a continuum of land units with heterogeneous productivity in various industries. Second, I use a Cobb-Douglas function function without other mobile factors of production because of the lack of data on the cost structure of industries at the regional level. Third, an analysis of inter-regional trade linkages as in Caliendo, Parro, Rossi-Hansberg, and Sarte (2014) cannot be implemented due to insufficient data on cross-regional trade in Brazil.

$$\hat{W}^j \equiv \left[\psi^j_{HSG}\left(\hat{w}^k_{HSG}\right)^{1-\rho} + \psi^j_{HSD}\left(\hat{w}^k_{HSD}\right)^{1-\rho}\right]^{\frac{1}{1-\rho}},$$

and ψ^j_g is the share of group g in the total wage bill of industry j at the initial equilibrium.

In addition, I assume that changes in sectoral supply of effective labor units are given by

$$\hat{\bar{L}}^N_g = \frac{\int_0^{l^N_g \hat{l}^N_g} e^{A_g(q)}\, dq}{\int_0^{l^N_g} e^{A_g(q)}\, dq} \quad \text{and} \quad \hat{\bar{L}}^C_g = \frac{\int_{l^N_g \hat{l}^N_g}^1 e^{\alpha_g(q)+A_g(q)}\, dq}{\int_{l^N_g}^1 e^{\alpha_g(q)+A_g(q)}\, dq} \tag{1.25}$$

where the change in sector employment composition is determined by equation (1.7):

$$\ln \hat{w}^N_g - \ln \hat{w}^C_g = \alpha_g\left(l^N_g \hat{l}^N_g\right) - \alpha_g\left(l^N_g\right). \tag{1.26}$$

The schedules of comparative and absolute advantage govern the sector-specific efficiency of workers employed in each sector, determining the sector supply of effective labor units. In general equilibrium, the sector effective labor supply has to be finite and, therefore, the integrals in equation (1.25) have to be well defined.[29]

In this environment, the counterfactual changes in the wage per efficiency unit, $\{\hat{\omega}^C_g, \hat{\omega}^N_g\}_g$, have to guarantee labor market clearing for all sectors and groups:

$$\sum_{j \in \mathcal{J}^k} \phi^{k,j}_g \cdot \hat{L}^j_g = \hat{\bar{L}}^k_g \tag{1.27}$$

where $\hat{L}^j_g$ is the change in the labor demand of industry j given by equation (1.24), and $\hat{\bar{L}}^k_g$ is the change in sector labor supply given by equation (1.25). At the initial equilibrium, $\phi^{k,j}_g$ denotes the share of industry j in total labor payments of sector k to individuals of group g.

[29] In the model, the effective labor supply in the non-commodity sector is $\bar{L}^N_g = \int_0^{l^N_g} E[e^{a(i)}|\alpha_g(q)]\, dq$. To obtain expression (1.25), it is sufficient that $E[e^{a(i)}|\alpha_g(q)] = \kappa_g \cdot e^{A_g(q)}$ for some constant κ_g. As discussed in Appendix 1A.5.3, such a relation holds under a variety of assumptions regarding the conditional distribution of absolute advantage, including the normal distribution and the Gumbel distribution. A similar argument holds for the effective labor supply in the commodity sector. For all possible parameters, I show in Appendix 1A.5.3 that the finiteness of the integrals in (1.25) is guaranteed by bounds on the productivity support.

Calibration. The change in sectoral labor demand in equation (1.24) requires additional paramters that are calibrated as follows. First, the Cobb-Douglas production function implies that the parameter η^j corresponds to the share of labor in the total revenue of industry j. Thus, I calibrate η^j using information on the cost structure of industries in Brazil computed by the IBGE in the 2009 national accounts. Second, sectoral labor demand requires the elasticity of substitution between skilled and unskilled workers, ρ. I calibrate this parameter using the estimated elasticity in Katz and Murphy (1992) for the U.S.: $\rho = 1.8$. Third, I use labor market data from the Census to compute the initial cost structure in each region: (i) the share of industry j in total labor payments of sector k to workers of group g, $\phi^{k,j}_{g,r}$; and (ii) the share of group g in total wage bill of industry j, $\psi^j_{g,r}$.

Counterfactual Changes in Brazilian Wage Inequality

Starting from the initial equilibrium in each microregion, I use equations (1.24)-(1.27) to compute the counterfactual changes in the wage per efficiency unit, $\{\hat{\omega}^C_{HSD,r}, \hat{\omega}^N_{HSD,r}\}$ and $\{\hat{\omega}^C_{HSG,r}, \hat{\omega}^N_{HSG,r}\}$, implied by shocks to final product prices, $\hat{p}^j$. I then proceed as above and compute the counterfactual change in Brazilian wage inequality.

Table 1.6 presents the counterfactual results computed with the general equilibrium calibrated model in columns (1) and (2) and compares them to the results computed with the reduced-form pass-through in columns (3) and (4). The two approaches yield similar counterfactual changes in log-wage variance both in the aggregate and within the two groups. In the context of the shock in world commodity prices, the general equilibrium model does not deliver additional insights beyond those obtained with the reduced-form pass-through. Nevertheless, the general equilibrium model provides a methodology by which the structural parameters of comparative and absolute advantage can be used to evaluate the distributional effects of other shocks in the economy — for example, shocks to tariffs in manufactured good produced by industries in the non-commodity sector.

However, columns (1) and (2) of Table 1.6 indicate that the distinct roles of comparative and absolute advantage are quantitatively important in determining the

Table 1.6: Effect of 1991–2010 Rise in World Commodity Prices on Brazilian Wage Inequality

	Calibrated General Equilibrium Model		Reduced-Form Pass-Through	
	Fréchet	Log-linear	Fréchet	Log-linear
	(1)	(2)	(3)	(4)
Panel A: High School Graduates				
	-0.002	-0.018	-0.002	-0.014
	1.24%	10.06%	1.02%	7.91%
Panel B: High School Dropouts				
	-0.014	-0.009	-0.023	-0.012
	4.60%	2.81%	7.65%	3.97%
Panel C: All Workers				
	-0.022	-0.019	-0.023	-0.017
	7.22%	5.99%	7.52%	5.55%

Note. Table reports the counterfactual change in Brazilian log-wage variance, along with the percentage of the actual change in log-wage variance, between 1991 and 2010. Computation uses the estimated parameters of comparative and absolute advantage on Table 1.3.

predicted change in log-wage variance. With the log-linear model shown in column (2), the shock in world commodity prices triggers a decrease in wage inequality among HSG driven by movements in sectoral wage differentials and within-sector wage dispersion. Such an effect accounts for 10% of the reduction in wage inequality among HSG. In contrast, these responses are ruled out by the Fréchet model and, therefore, column (1) reports almost no change in log-wage variance for HSG. Among HSD, this pattern is inverted: the fall in log-wage variance predicted by the log-linear model is weaker than that predicted by the Fréchet model. The small magnitude of this difference reflects the similar estimates of the structural parameters obtained with the two parametrizations for HSD.

At the national level, the two specifications yield similar predicted changes in log-wage variance. This similarity reflects the fact that the differences in predicted changes go in opposite directions within the two groups. Also, the more pronounced difference among HSG is attenuated in the aggregate since this group represented only 22.2% of the labor force in 1991. Compared to the initial log-wage variance in 1991, the rise in world commodity prices triggers an inequality fall of 2.0% and 1.7% according, respectively, to the Fréchet model and the log-linear model. These figures correspond, respectively, to 7.2% and 6.0% of the total fall in Brazilian wage inequality between 1991 and 2010.

1.6 Conclusion

This paper starts from one observation: movements in world commodity prices tend to be accompanied by changes in Brazilian wage inequality. Motivated by this aggregate correlation, I developed a unified theoretical and empirical framework to quantify the causal effect of shocks to relative good prices in the world market on the wage distribution within Brazil.

I proposed a model featuring workers' heterogeneity in comparative and absolute advantage with respect to their productivities in the two sectors of the economy. In this environment, I clearly delineated the distinct roles played by the schedules of comparative and absolute advantage in determining sectoral employment and sectoral wages, which allowed me to establish their nonparametric identification in a sample of regional economies. Building on this result, I estimated the schedules of comparative and absolute advantage for High School Graduates and Dropouts using the differential exposure of Brazilian regions to shocks in world commodity prices. Because these schedules are sufficient to compute changes in the average and variance of the log-wage distribution, I was able to use the structural estimates in a quantitative investigation of the effect of shocks in world commodity prices on Brazilian log-wage variance. I concluded that the rise in world commodity prices accounted for 5% to 10% of the decline in Brazilian log-wage variance between 1991 and 2010.

To put my results in perspective, I compare them to the distributional effects of the Brazilian trade liberalization estimated by two recent papers. Dix-Carneiro and Kovak (2015b) investigate the effect of cross-regional variation in exposure to the tariff reduction on the HSG-HSD wage premium, finding that 11% of the 1991–2010 drop in this variable can be attributed to the trade liberalization. Focusing on a different channel, Helpman, Itskhoki, Muendler, and Redding (2015) conclude that heterogeneous worker exposure to firms differentially affected by the trade liberalization caused log-wage dispersion to increase by 2% between 1986 and 1994. In contrast, I study an alternative source of international trade shocks in Brazil: the variation in world commodity prices. When I compare my results to those in these two papers, I find that the shocks in world demand for basic products generate changes in wage inequality with magnitude similar to that created by tariff shocks. My results indicate the distributional importance of shocks in world commodity prices in other developing countries where a large fraction of the workforce is employed in the commodity sector.

In this paper, distributional effects arise from worker heterogeneity in terms of comparative and absolute advantage. My analysis highlighted the different roles of these two economic forces, demonstrating the potential harm of ignoring their distinct implications for changes in wage inequality. This is illustrated by the counterfactual change in wage inequality for HSG, which is much lower under the parametric restrictions imposed by the Fréchet distribution. As a result, I developed a flexible methodology that can be readily applied to investigate changes in between- and within-group wage inequality stemming from a variety of sectoral shocks, including higher competition of foreign manufacturing imports and reductions in the trade cost of services.

In order to dispense with parametric assumptions, my methodology restricted the dimensionality of worker heterogeneity to a bivariate vector of sector-specific productivities. In Roy-like models, it remains an open question how to generalize the insights in this paper to an environment with higher heterogeneity dimensionality while achieving both tractability and flexibility in the analysis of wage inequality movements. Such an extension would enhance the range of questions that could be addressed by the model, being particularly useful to quantify distributional effects of sectoral shocks in countries where workers present high spatial mobility.

1A Appendix

1A.1 Proofs

1A.1.1 Proof of Lemma 1

By Assumption 3, $E[u_{g,m}|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}] = 0$ so that $E\left[\Phi_g(l^N_{g,m}) + \mathbf{X}_{g,m}\gamma_g|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}\right] = E\left[y_{g,m}|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}\right]$. Now let us proceed by contradiction. Suppose there exist $\tilde{\Phi}_g(l^N_{g,m})$ and $\tilde{\gamma}_g$ such that

$$E\left[\tilde{\Phi}_g\left(l^N_{g,m}\right) + \mathbf{X}_{g,m}\tilde{\gamma}_g|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}\right] = E\left[y_{g,m}|\mathbf{Z}^k_{g,m}, \mathbf{X}_{g,m}\right]$$

By Assumption 4,

$$\Phi_g\left(l^N_{g,m}\right) - \tilde{\Phi}_g\left(l^N_{g,m}\right) + \mathbf{X}_{g,m}(\gamma_g - \tilde{\gamma}_g) = 0 \quad \text{almost surely.}$$

Take markets m and n such that $\mathbf{X}_{g,m} = \mathbf{X}_{g,n}$. The condition above implies that

$$\Phi_g\left(l^N_{g,m}\right) - \Phi_g\left(l^N_{g,n}\right) = \tilde{\Phi}_g\left(l^N_{g,m}\right) - \tilde{\Phi}_g\left(l^N_{g,n}\right) \quad \text{for all} \quad l^N_{g,m} \text{ and } l^N_{g,n}.$$

Thus, $\Phi_g(.)$ is identified up to a constant. To determine this constant, we can use the normalizations $E[u_{g,m}] = E[\mathbf{X}_{g,m}\gamma_g] = 0$, which imply that $E[\Phi_g(l^N_{g,m})] = E[y_{g,m}]$.
□

1A.1.2 Derivation of Equation (1.17)

Consider shocks to endogenous and exogenous variables in a particular market m where the sector-specific productivity distribution satisfies Assumption 1. To simplify notation, I drop the index m.

Recall that individual i's log-wage is given by $y_g(i) = \max\{\omega^C_g + s_g(i) + a_g(i);\ \omega^N_g + a_g(i)\}$. Under Assumption 1, this implies that the log-wage distribution of group g is given by

$$\begin{aligned} Pr[y_g(i) \leq y] &= Pr[y_g(i) \leq y;\ s_g(i) \leq \omega_g^N - \omega_g^C] + Pr[y_g(i) \leq y;\ s_g(i) > \omega_g^N - \omega_g^C] \\ &= \int_0^{l_g^N} Pr\left[a_g(i) \leq y - \omega_g^N \middle| \alpha_g(q)\right] dq + \int_{l_g^N}^1 Pr\left[a_g(i) \leq y - \omega_g^C - \alpha_g(q) - \tilde{u}_g \middle| \alpha_g(q)\right] dq \end{aligned} \tag{1.28}$$

where $Pr[a_g(i) \leq a|\tilde{s}_g(i) = s] \equiv \mu H_g^a(a|s) + (1-\mu)H^e(a|\tilde{u}_g, \boldsymbol{\theta}_g)$.

By construction, $Y_g(\pi)$ solves $Pr\left[y_g(i) \leq Y_g(\pi)\right] = \pi$. Taking a first-order expansion,

$$\left[f_g^N\left(Y_g(\pi)\right) + f_g^C\left(Y_g(\pi)\right)\right]\Delta Y_g(\pi) = f_g^N\left(Y_g(\pi)\right) \cdot \Delta\omega_g^N + f_g^C\left(Y_g(\pi)\right) \cdot \Delta\omega_g^C - \Delta v_g'(\pi)$$

where $f_g^N(y) \equiv \frac{\partial Pr[y_g(i) \leq y;\ s_g(i) \leq \omega_g]}{\partial y}$, $f_g^C(y) \equiv \frac{\partial Pr[y_g(i) \leq y;\ s_g(i) > \omega_g]}{\partial y}$, and $\omega_g \equiv \omega_g^N - \omega_g^C$.

In this expression, the first-order impact of the endogenous change in l_g^N is eliminated by the employment condition (1.7), reflecting the fact that marginal workers are indifferent between the two sectors. The term $v_g'(\pi)$ incorporates changes to other exogenous parameters of the productivity distribution:

$$\begin{aligned} \Delta v_g'(\pi) \equiv &(1-\mu)\left[\int_0^{l_g^N} \nabla_\theta H^e\left(Y_g(\pi) - \omega_g^N \middle| \tilde{u}_g, \boldsymbol{\theta}_g\right) dq + \right. \\ &\left. + \int_{l_g^N}^1 \nabla_\theta H^e\left(Y_g(\pi) - \omega_g^C - \alpha_g(q) - \tilde{u}_g \middle| \tilde{u}_g, \boldsymbol{\theta}_g\right) dq\right] \cdot \Delta\boldsymbol{\theta}_g \\ &+ \left[f_g^C\left(Y_g(\pi)\right) + (1-\mu)\left(\int_0^{l_g^N} \nabla_u H^e\left(Y_g(\pi) - \omega_g^N \middle| \tilde{u}_g, \boldsymbol{\theta}_g\right) dq \right.\right. \\ &\left.\left. + \int_{l_g^N}^1 \nabla_u H^e\left(Y_g(\pi) - \omega_g^C - \alpha_g(q) - \tilde{u}_g \middle| \tilde{u}_g, \boldsymbol{\theta}_g\right) dq\right)\right] \Delta\tilde{u}_g. \end{aligned}$$

To obtain equation (1.17), notice that $l_g^N(\pi) \equiv P\left[s_g(i) \leq \omega_g^N - \omega_g^C \middle| y_g(i) = Y_g(\pi)\right]$. This is equivalent to

$$l_g^N(\pi) = \frac{Pr\left[y_g(i) = Y_g(\pi);\ s_g(i) \leq \omega_g^N - \omega_g^C\right]}{Pr\left[y_g(i) = Y_g(\pi)\right]} = \frac{f_g^N\left(Y_g(\pi)\right)}{f_g^N\left(Y_g(\pi)\right) + f_g^C\left(Y_g(\pi)\right)}.$$

Thus,

$$\Delta Y_g(\pi) = \Delta\omega_g^C \cdot l_g^C(\pi) + \Delta\omega_g^N \cdot l_g^N(\pi) + \Delta v_g''(\pi)$$

where $\Delta v_g''(\pi) \equiv -\Delta v_g'(\pi) / \left[f_g^N\left(Y_g(\pi)\right) + f_g^C\left(Y_g(\pi)\right)\right]$.

Finally, equation (1.17) is obtained by projecting $\Delta v''_g(\pi)$ on observable covariates and unobservable variables such that

$$\Delta v''_g(\pi) \equiv \mu_g \cdot \tilde{\mathbf{X}}_g(\pi) + \Delta v_g(\pi).$$

1A.2 Model Extension: Non-monetary Employment Benefits

This section extends the model of Section 1.2 by incorporating non-monetary employment benefits — a reduced-form for work conditions and switching cost. The environment of Section 1.2 remains the same except for workers' preference structure. If employed in sector k, assume that individual $i \in \mathcal{I}_{g,m}$ obtains utility $\tau_g^k(i) \cdot u(c)$ from consuming bundle c where $u(.)$ is homogeneous of degree one. Thus, individual i's payoff of employment in sector k is given by

$$U_g^k(i) = \tau_g^k(i) \cdot \frac{w_{g,m}^k L_g^k(i)}{P_m} \tag{1.29}$$

where $\tau_g^k(i)$ is individual i's private benefit of being employed in sector k; and P_m is the price index in market m.

In the presence of non-monetary employment benefits, I extend the notion of comparative advantage to also include relative sectoral preferences. Accordingly, define individual i's comparative advantage as

$$s_g(i) \equiv \ln[L_g^C(i)/L_g^N(i)] + \ln[\tau_g^C(i)/\tau_g^N(i)],$$

and individual i's efficiency in sector k as

$$a_g^k(i) \equiv \ln[L_g^k(i)] \quad \text{for } k = C, N.$$

In a given group and market, consider the following distribution of preferences and productivities.

Assumption 6. *Suppose individual i in market m, $i \in \mathcal{I}_{g,m}$, independently draws $\left(s_g(i), a_g^C(i), a_g^N(i)\right)$ as follows.*

i. ***Comparative Advantage:***

$$s_g(i) = \tilde{s}_g(i) + \tilde{u}_{g,m} \quad \text{and} \quad \{\tilde{s}_g(i)\} \sim F_g(s)$$

where $\tilde{u}_{g,m}$ is a group-market shifter of comparative advantage, and $\alpha_g(q) \equiv (F_g)^{-1}(q)$.

ii. ***Sector-Specific Efficiency:***

$$a_g^k(i) \equiv \tilde{a}_g^k(i) + \tilde{v}_g^k \quad \text{s.t.} \quad \left\{\left(a_g^C(i), a_g^N(i)\right) \middle| \tilde{s}_g(i) = s\right\} \sim H_g^a\left(a^C, a^N \,|s\right)$$

where $\tilde{v}_{g,m}^k$ is a group-market shifter of sector efficiency and

$$A_g^k(q) \equiv \int\int a^k \, dH_g^a\left(a^C, a^N \,|\alpha_g(q)\right).$$

The preference structure in (1.29) immediately implies that utility maximizing individuals choose to be employed in the non-commodity sector if, and only if, $s_g(i) \leq \omega_{g,m}^N - \omega_{g,m}^C$. Thus,

$$\omega_{g,m}^N - \omega_{g,m}^C = \alpha_g\left(l_{g,m}^N\right) + \tilde{u}_{g,m}. \tag{1.30}$$

Given the allocation of workers to sectors, the average log-wage in the non-commodity sector is $\bar{Y}_g^N \equiv E[\bar{Y}_g^k(q)|q < l_{g,m}^N]$, which is equivalent to

$$\bar{Y}_{g,m}^N = \omega_{g,m}^N + \bar{A}_g^N(l_{g,m}^N) + \tilde{v}_{g,m}^N \quad \text{s.t.} \quad \bar{A}_g^N(l) \equiv \frac{1}{l}\int_0^l A_g^N(q)\, dq. \tag{1.31}$$

Also, the average log-wage in the commodity sector is $\bar{Y}_g^C \equiv E[\bar{Y}_g^C(q)|q \geq l_{g,m}^N]$ and, therefore,

$$\bar{Y}_{g,m}^C = \omega_{g,m}^C + \bar{A}_g^C(l_{g,m}^N) + \tilde{v}_{g,m}^C \quad \text{s.t.} \quad \bar{A}_g^C(l) \equiv \frac{1}{1-l}\int_l^1 A_g^C(q)\, dq. \tag{1.32}$$

With a sector demand shifter satisfying Assumptions 3–4, Lemma 1 establishes the identification of $\alpha_g(.)$ from equations (1.30). Also, Lemma 1 establishes that $\bar{A}_g^N(.)$ and $\bar{A}_g^C(.)$ are respectively identified from equations (1.31)–(1.32). To recover $A_g^C(.)$ and $A_g^N(.)$, notice that

$$A_g^N(q) = \frac{\partial}{\partial q}\left[q \cdot \bar{A}_g^N(q)\right] \quad \text{and} \quad A_g^C(q) = -\frac{\partial}{\partial q}\left[(1-q) \cdot \bar{A}_g^C(q)\right].$$

Theorem 2. *Consider a set of segmented markets, m, subject to sector demand shifters, $\mathbf{Z}_{g,m}^k$, such that Assumptions 2–4 and 6 hold. For each worker group g, $\alpha_g(.)$ is identified from equation* (1.30), *$A_g^N(.)$ is identified from equation* (1.31); *and $A_g^C(.)$ is identified from equation* (1.32).

1A.3 Data Construction and Measurement

1A.3.1 World Price of Agriculture and Mining Commodities

To capture Brazil's exposure to world prices of basic products, I build price indices for each commodity category. The first source of international commodity prices is the Commodity Research Bureau, which publishes price indices by commodity group based on product spot prices in the main exchange markets in the United States. In the paper, I use those groups with sizable employment participation in Brazil: Grains (corn, soybeans, and wheat), Soft Agriculture (cocoa, coffee, sugar, orange juice, and others), Livestock (hides, hogs, lard, steers, tallow, and others), and Metals (copper scrap, lead scrap, steel scrap, tin, zinc, and others). In addition, I build price indices for two commodity groups using future prices in the New York Mercantile Exchange: Precious Metals (gold and silver) and Energy (crude oil). Due to their small employment importance in Brazil, I aggregate Metals and Precious Metals into a single Mining category. These series of international nominal prices were converted into local currency using the nominal exchange rate and deflated by the Brazilian consumer price index (IPCA).[30] To avoid short-term price volatility, I use the average price in the six months preceding the process of data collection of the Census; that is, the average price between March and August of each year of the Census.

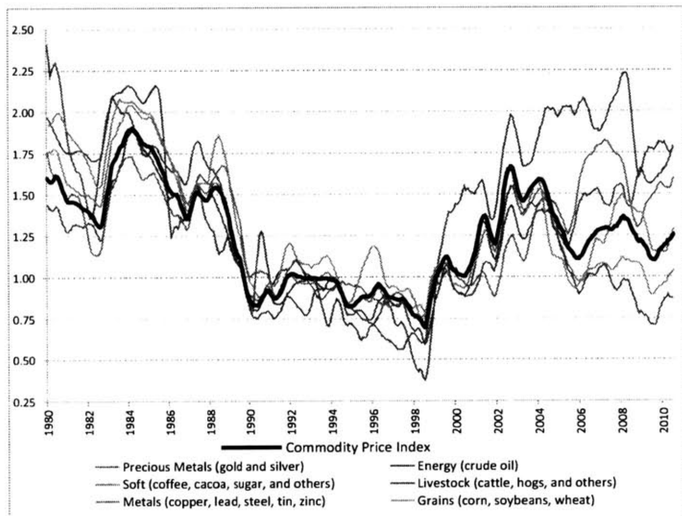

Figure 1A.1: World Price of Agriculture and Mining Commodities

1A.3.2 Industry Composition

Individuals in the sample are allocated to sectors according to their self-reported industry of employment. Table 1A.1 shows the industry classification used in this paper together with corresponding industry codes used by the IBGE in each year of the Census and the PNAD. I use crosswalk tables publicly provided by the IBGE to link the different activity codes across years. The division of industries in the commodity sector accommodates available information on international prices as described above.

[30] All commodity price series were downloaded from the Global Financial Database. In the end of 2008, the soft and grains indices were unified under the foodstuff index. Thus, I build these series for 2009–2010 using each index description. Series of nominal exchange rate and IPCA were downloaded from the IPEADATA.

Table 1A.1: Industry Classification and IBGE Activity Codes

Industry	Atividades	CNEA-Dom	CNEA-Dom 2.0
	PNADs of 1981-2001 1980 Census and 1991 Census	PNADs of 2002-2009 2000 Census	- 2010 Census
Commodity Sector			
Grains (corn, soybeans, and wheat)	20; 21; 22	1102; 1103; 1107	1102; 1103; 1107
Soft (coffee, cocoa, sugar, and others)	11; 12; 14-17; 23; 24	1104; 1105; 1110-1116; 2001; 2002; 15022; 15042	1104; 1105; 1110-1116; 10022; 10093
Livestock (cattle, hogs, and others)	26; 27; 41; 42	1201-1205; 1208; 1209; 1300; 1402; 5001; 5002; 15010; 15030	1201-1205; 1208; 1209; 1402; 1999; 3001; 3002; 10010; 10030
Metals (copper, lead, steel, zinc, and others)	58	13002	7002
Precious Metals (gold and silver)	55	13001	7001
Energy (crude oil)	51	11000	6000
Other agriculture and mining	13; 18-19; 25; 28; 29; 31-37; 50; 52-54; 56; 57; 59; 581	1101; 1106; 1108; 1109; 1117; 1118; 1206; 1207; 1401; 10000; 12000; 14001-14004	1101; 1106; 1108; 1109; 1117-1119; 1206; 1207; 1401; 5000; 8001-9000
Non-Commodity Sector			
Manufacturing	100-300	15021; 15041; 15043; 15050; 16000-37000	10021; 10091; 10092; 10099-32999; 38000
Non-Tradable Goods and Services	340-901	1500; 40010-99000	1500; 33001-37000; 39000-99000

1A.3.3 Trends in Brazilian Wage Dispersion

I obtain annual data on Brazilian labor market outcomes from the National Household Sample Survey (PNAD) collected by the Brazilian Institute of Geography and Statistics (IBGE) between 1981 and 2009.[31] To focus on individuals with strong labor force attachment, I consider a benchmark sample of full-time male employed individuals aged 16–64. I decompose the movement in log-wage variance into observable and residual components by regressing log wages on a full set of dummies for years of experience (0–39 years), years of education (0–16 years), state of residence (27 states), race (white dummy), and sector of employment (commodity sector dummy). Figure 1A.2 presents the trends in Brazilian wage inequality between 1981 and 2009. Throughout the period, observable worker attributes account for a large share of the change in log-wage variance: 73% of the increase in 1981–1990, and 65% of the decline in 1990–2009.

Table 1A.2 presents the full decomposition of log-wage variance. In 1981–1990, the increase in the between-component of wage variance was mainly driven by the covariance term with additional contributions of the terms related to sector and education dummies. In the 1990–2009, the sharp drop in the between-component of wage dispersion is distributed across all terms, with the largest contribution coming

[31] The National Household Sample Survey is not available for the years in which the IBGE published the Brazilian Demographic Census (1980, 1991, 2000, and 2010).

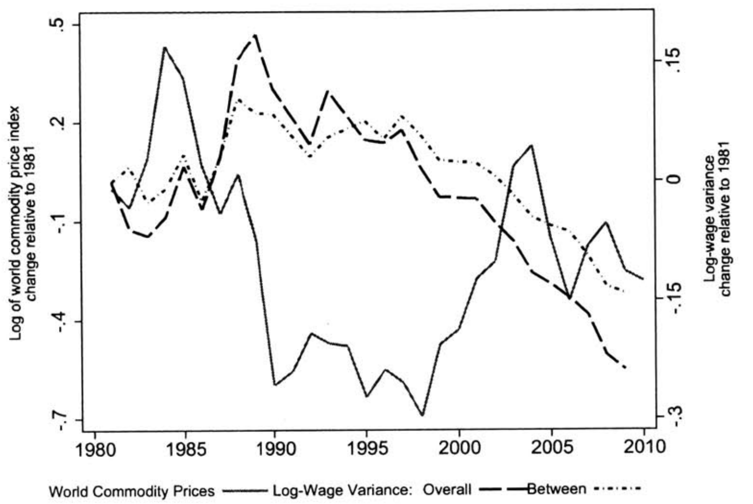

Figure 1A.2: World Commodity Prices and Brazilian Log-Wage Variance, 1981-2010

Note. World Commodity Prices correspond to the log of the commodity price index computed with the world price of agriculture and mining products converted to Brazilian currency and deflated by the Brazilian consumer price index. Sample of male full-time workers extracted from the PNAD. Between component of log-wage variance computed with the predicted values of the regression of log wage on a full set of dummies for years of experience (0–39 years), years of education (0–16 years), state of residence (27 states), race (white dummy), and sector of employment (commodity sector).

from the education dummies. Conclusions in Table 1A.2 are related to results reported elsewhere in the literature. In particular, Ferreira, Firpo, and Messina (2014) highlight the importance of falling educational and state wage gaps for the decrease in Brazilian wage inequality between 1995 and 2012. Using administrative data for formal sector employees, Helpman, Itskhoki, Muendler, and Redding (2015) conclude that observable worker attributes account for roughly half of the increase in log-wage variance between 1986 and 1995. In Table 1A.2, the residual component of log-wage variance presents a lower contribution to log-wage variance movements, due to the inclusion of a more comprehensive set of dummies for state of residence, years of education, and years of experience.

1A.3.4 Empirical Application: Data Construction

Labor Market Data. I obtain data on labor market outcomes from publicly available long versions of the Brazilian Census collected by the Brazilian Institute of

Table 1A.2: **Decomposition of Brazilian Log-Wage Variance, 1981–2009**

	1981	1986	1990	1995	1999	2005	2009
Overall	0.935	0.900	1.053	0.987	0.915	0.805	0.697
Residual	0.459	0.449	0.492	0.444	0.418	0.390	0.366
Between	0.475	0.451	0.561	0.544	0.496	0.415	0.331
Sector	0.024	0.013	0.031	0.033	0.027	0.013	0.012
Education	0.257	0.263	0.278	0.260	0.248	0.232	0.190
State	0.053	0.044	0.052	0.052	0.043	0.045	0.034
Race	-	-	0.005	0.005	0.006	0.003	0.004
Experience	0.109	0.111	0.107	0.078	0.075	0.067	0.053
Covariance	0.034	0.020	0.088	0.115	0.097	0.055	0.037

Note. Sample of male full-time workers extracted from the PNAD. Wage decomposition implied by a regression of log wage on a full set of dummies for years of experience (0–39 years), years of education (0–16 years), state of residence (27 states), race (white dummy), and sector of employment (commodity sector).

Geography and Statistics (IBGE) for the years of 1980, 1991, 2000, and 2010. From the Census, I extract a sample of full-time workers aged between 16 and 64. Full-time workers are defined as those reporting more than 35 weekly worked hours. I restrict the sample to workers with calculated experience between 0 and 39 years. Experience is defined as the individual's age minus a predicted initial working age that equals 23 for college graduates, 18 for High School graduates, and 15 for those with only primary education. The benchmark sample is further restricted to include only white male workers. This restriction allows us to focus on individuals with strong labor force. In addition, it excludes individuals directly affected by the strong declines in gender and race wage differential between 1995 and 2010 (Ferreira et al., 2014). Such movements in wage differentials are not the directly related to the model in this paper. In robustness exercises, I extend the benchmark sample to also include female and non-white workers.

Regional Labor Markets. I use the microregion concept created by the IBGE in the 1991 Census as a regional labor market unit. Each of the 558 microregions corresponds to a set of economically integrated municipalities with interconnected labor markets. This definition was used in a series of recent papers analyzing the response of local labor markets to aggregate trade shocks (e.g., Kovak, 2013 and Dix-Carneiro and Kovak, 2015b,a). The microregion concept in Brazil is similar to the Commuting Zones in the United States used by Autor, Dorn, and Hanson (2013). Despite the sharp increase in the number of municipalities between 1991 and 2010, the IBGE maintained the same microregion definition in the Censuses of 1991, 2000, and 2010.[32] In the 1980 Census, the microregion variable does not exist, so I created it from existing municipalities in 1980. Because of the change in municipality borders between 1980 and 1991, it is only possible to replicate a subset of the microregions using historical administrative borders. To be more precise, I recover 540 microregions in the 1980 Census compared to the 558 microregions in the 1991 Census.[33]

Sample Selection. In the empirical application, I select a baseline sample of 518 microregions with positive employment in the commodity sector for all groups in the 1991–2010 period, covering 98.4% of the country's population in 1991. The 1980–1991 period is excluded from the baseline sample mainly because of the turbulent economic environment in Brazil during the 1980s. The decade was marked by hyper-inflationary episodes, suspension of foreign currency convertibility, and the adoption of restrictive internal controls on prices and wages. In this environment, it is not clear that relative international prices were very informative about relative prices faced by domestic producers when deciding resource allocation. More normal economic conditions returned after the series of structural reforms implemented in 1993–1994 that brought monetary stabilization, eliminated price controls, and restored full currency

[32]There were 4,491 municipalities in 1991 and 5,565 in 2010. Out of the 1074 municipalities created in the period, 998 municipalities had parent municipalities in a single microregion and, therefore, they were allocated to this microregion. The other 76 municipalities had parent municipalities in more than one microregion. These municipalities, which represented .33% of employment in 2000, were allocated to the microregion of the parent municipality that ceded the highest population share to the new municipality. This procedure adopted by the IBGE minimizes any measurement error implied by the border change. In fact, all results in the paper are robust to using a sample of 491 microregions built by aggregating microregions such as to keep borders unchanged in the period.

[33]Out of the 3,991 municipalities in the 1980 Census, I am able to link 3,938 municipalities to at least one of the 4,491 municipalities in the 1991 Census. With these linked municipalities, I construct microregions in 1980 using the microregion assigned to corresponding municipalities in 1991. The main problem of this method is the existence of new municipalities in 1991 that belonged to a different microregion than their parent municipalities in 1980. This is the case for 85 of the 500 municipalities created between 1980 and 1991, accounting for .67% of total employment in 2000.

convertibility. Robustness exercises attest that similar results hold in the extended sample spanning the entire 1980–2010 period.

Sector Demand Shifter. To build the group-region exposure to international commodity prices, I compute total labor income by industry. To this end, I consider the weighted sum of monthly wages of individuals reporting to hold their main job in the industry using Census sampling weights. Denote $Y^j_{g,r,t}$ as the total labor payments of industry j to workers of group g in microregion r at year t. The initial participation of industry j in the labor payments of sector k to group g in microregion r is

$$\phi^{k,j}_{g,r} \equiv \frac{Y^j_{g,r,1991}}{\sum_{j' \in \mathcal{J}^k} Y^{j'}_{g,r,1991}} \quad \text{where} \quad j \in \mathcal{J}^k.$$

Table 1A.3 reports summary statistics of industry composition in the sample of 518 microregions in 1991. Columns (1) and (3) indicate that regions, on average, have a large fraction of their work force allocated to the commodity sector, with agriculture accounting for the bulk of the sector's labor expenditure. Importantly, columns (2) and (4) document great heterogeneity in industry composition across microregions. Comparing columns (1) and (3), it is possible to identify different exposure patterns for the two groups. While HSD are more likely to be employed in the production of grains and soft agricultural items, HSG are more likely to be employed in the production of livestock and crude oil. Due to their small employment share, I aggregate the Metals and the Precious Metals groups into a single Mining category.

Regional Labor Market Outcomes. To calculate wage outcomes, I estimate wage regressions separately for each year using the entire sample of workers in the country. Specifically, I regress the log monthly wage on a full set of experience dummies (0–39 years) interacted with dummies for female and white workers. The residual of this regression corresponds to a wage measure adjusted for variation in these demographic characteristics across groups and microregions. For each triple of group-region-year, the sector average log wage is the weighted average of the adjusted wage among

Table 1A.3: Summary Statistics: Labor Income Share by Industry in Brazil, 1991

Industry	High School Graduates		High School Dropouts	
	Mean	SD	Mean	SD
	(1)	(2)	(3)	(4)
1. Commodity Sector	9.0%	9.6%	21.6%	19.7%
Grains (corn, soybeans, and wheat)	4.7%	11.6%	10.1%	16.3%
Soft (coffee, cocoa, sugar and other)	13.0%	16.1%	19.5%	18.0%
Livestock (cattle, hogs, and others)	35.5%	21.1%	26.8%	15.6%
Metals (copper, lead, steel, zinc, and others)	3.0%	7.2%	1.6%	4.2%
Precious Metals (gold and silver)	1.0%	4.0%	1.8%	4.7%
Energy (crude oil)	8.4%	17.1%	2.3%	6.2%
Other agriculture and mining	34.3%	20.7%	37.9%	19.8%
2. Manufacturing	16.1%	10.7%	18.1%	11.3%
3. Non-Tradable Goods and Services	74.9%	10.8%	60.4%	14.8%

Note. Sample of male white full-time workers extracted from the Brazilian Census of 1991. Statistics weighted by the microregion share in the 1991 national population.

individuals reporting to hold their main job in the sector. Lastly, the average log wage is the weighted average of the adjusted wage among all individuals in the triple of group-region-year. In both cases, the computation uses Census sampling weights.

Following closely Autor, Katz, and Kearney (2008), I use an efficiency-adjusted measure of total hours. With this measure, I compute sectoral employment using the sum of efficiency-adjusted hours supplied by sector employees in a group-region-year. I perform the efficiency adjustment by multiplying individual weekly hours by a time-invariant measure of relative wage for each cell of sex-race-education-experience.[34] I then compute the total sector employment, $H^k_{g,r,t}$, as the weighted sum of efficiency-adjusted hours of individuals reporting to have their main job in the sector. This

[34]I consider 48 cells based on two sex groups, two race groups (white and non-white), three educational groups (high school dropouts, high school graduates, and college graduates), and four experience groups (0–9, 10–19, 20–29, and 30–39 years). For each cell, the relative wage is the average hourly wage divided by the average wage of female non-white High School dropouts with 0–9 years of experience. The cell weight is the average relative wage across regions and years (1991, 2000, 2010). In the 1980 Census, weekly hours are only reported in ranges. To compute efficiency-adjusted hours, I assign 45 and 54 weekly hours to individuals reporting, respectively, 40–48 and 49+ hours.

computation uses the Census sampling weights. Finally, sector employment share is defined as $l^k_{g,r,t} \equiv H^k_{g,r,t}/(H^C_{g,r,t} + H^N_{g,r,t})$.

To obtain total labor supply, I use the aggregate amount of efficiency-adjusted hours in a group-region-year: $\bar{H}_{g,r,t} \equiv H^C_{g,r,t} + H^N_{g,r,t}$. Lastly, the labor supply of immigrants is computed exactly as above in a restricted sample of individuals identified as non-native residents of each microregion.[35]

1A.4 Empirical Application: Sensitivity Analysis

1A.4.1 Estimation of Sector Wage per Efficiency Unit

Data. In the structural exercise, the estimation of sector wage per efficiency unit requires initial sector employment and wage growth across percentiles of the wage distribution for each of the 2,072 group-region-period triples. To create this dataset, I compute percentiles of log hourly wage from the sample of individuals in a group-region-period using Census sampling weights. Individuals are distributed across percentile bins according to their log hourly wage. The sector employment share in each percentile bin corresponds to the fraction of efficiency-adjusted hours reported by sector employees in that percentile bin. Additionally, the wage growth in each percentile corresponds to the difference in the average hourly wage of individuals in that percentile bin between two consecutive years. Since extreme wage values are more likely to be generated by measurement error, I ignore the wage distribution tails by restricting the estimation to bins between the 6^{th} and the 94^{th} percentiles.

Baseline Specification. In principle, equation (1.17) can be implemented with any division of individuals into quantiles; in practice, however, this choice entails a tradeoff. On the one hand, a coarse discretization yields a low number of quantiles with potentially little variation in initial sector employment to precisely estimate the

[35] Given the information in the Census, I can only identify as microregion natives those individuals satisfying one out of two conditions. First, they were born in the same municipality in which they currently live. Second, if they were born in a different municipality, then I also consider microregion natives those that moved into the current municipality from another municipality in the same microregion during the previous ten years.

wage per efficiency unit. On the other hand, a refined discretization exacerbates measurement error of sector employment in each quantile because of the low number of sampled individuals in each sector. With these considerations in mind, I implement the estimation with 88 percentile bins of 1 p.p. width between the 6^{th} and the 94^{th} percentiles. Below, I show that similar results are obtained with bins of 2 p.p. width.

The implementation of expression (1.17) allows for a vector of observable variables that vary with the position in the wage distribution. Accordingly, the baseline specification includes the following dummy variables as nonparametric controls: (i) indicator that wage percentile is at the bottom (P6-P30) or middle (P30-P75) of the log-wage distribution; and (ii) indicator that wage percentile is below the federal minimum wage (pre-year and post-year). These dummies capture, for example, differential efficiency gains for workers in distant parts of the wage distribution, and income gains generated by bunching around the minimum wage. In this specification, sectoral wages per efficiency unit are identified from the variation in pre-shock sector employment in small neighborhoods of the log-wage distribution of workers in the same group-region-period.

Results. Table 1A.4 presents the summary statistics of estimated wages per efficiency unit implied by the baseline specification for each of the 2,072 group-region-period triples. Columns (1)–(2) display statistics of the estimated wage per efficiency unit in the commodity sector, $\Delta\omega^C_{g,r,t}$, and columns (3)–(4) of the estimated relative wage per efficiency unit in the non-commodity sector, $\Delta\omega^N_{g,r,t} - \Delta\omega^C_{g,r,t}$. The commodity sector's wage per efficiency unit presented robust growth in both periods. Between 1991 and 2010, the average increase was 47.0 log-points for HSG and 96.4 log-points for HSD. Simultaneously, the relative wage per efficiency unit in the commodity sector increased sharply. Lastly, column (5) reports the average R^2 of the estimation in the sample of microregions. A large fraction of the variation in wage growth across quantiles of the earnings distribution is captured by equation (1.17); in the two periods, the average R^2 is above 55% for HSG and 71% for HSD.

To address robustness to implementation choices, Table 1A.5 presents the correlation between the estimates of sectoral wages per efficiency unit implied by different specifications of equation (1.17) and those implied by the baseline specification. Columns (1)–(2) and (4)–(5) indicate a high correlation between estimates obtained with different control sets. Notice that, when minimum wage controls are omitted, estimated wages per efficiency unit are very similar to those of the baseline specification. This suggests that quantile range controls absorb much of the variation captured in the minimum wage dummies. Columns (3) and (6) attest that the particular choice of bin width has little impact on estimates: the correlation is above .88 between baseline estimates and those obtained with a coarser discretization of 2 p.p. bins.

Table 1A.4: Summary Statistics: Estimated Change in Wage per Efficiency Unit, 1991–2010

	$\Delta\omega_g^C$		$\Delta\omega_g^C - \Delta\omega_g^N$		R^2
	Mean	SD	Mean	SD	Mean
	(1)	(2)	(3)	(4)	(5)
Panel A: High School Graduates					
1991 - 2000	0.320	0.370	0.151	0.347	0.555
2000 - 2010	0.150	0.645	0.306	0.609	0.758
Panel B: High School Dropouts					
1991 - 2000	0.524	0.579	0.364	0.619	0.715
2000 - 2010	0.440	0.579	0.360	0.634	0.830

Note. Sample of 518 microregions in 1991–2000 and 2000–2010. Statistics are weighted by the microregion share in national population in 1991. Baseline estimates based on the discretization of the wage distribution in 88 bins of 1 p.p. width, including indicator dummies of percentile bins below the federal minimum wage (pre and post years); and percentile bins in bottom, middle, or top of the wage distribution (P6-P30 and P30-P75).

1A.4.2 Reduced-Form Evidence: Sensitivity Analysis

This section investigates the robustness of the reduced-form results reported in Section 1.4.5. To this end, I estimate model (1.23) with additional periods, additional worker groups, and additional labor market outcomes.

Table 1A.5: Estimated Change in Wage per Efficiency Unit, Correlation with Benchmark Specification

	Commodity sector			Non-commodity sector		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: High School Graduates						
Correlation with baseline	0.855	0.973	0.926	0.874	0.969	0.916
Panel B: High School Dropouts						
Correlation with baseline	0.914	0.960	0.886	0.912	0.960	0.893
Baseline Controls						
Percentile below federal minimum wage	Yes	No	Yes	Yes	No	Yes
Percentile in bottom, middle or top	No	Yes	Yes	No	Yes	Yes
Discretization of wage distribution						
Bins of 1 p.p. (N − 88)	Yes	Yes	No	Yes	Yes	No
Bins of 2 p.p. (N − 44)	No	No	Yes	No	No	Yes

Note. Sample of 518 microregions in 1991–2000 and 2000–2010. Statistics are weighted by the microregion share in national population in 1991. Baseline estimates based on the discretization of the wage distribution in 88 bins of 1 p.p. width, including indicator dummies of percentile bins below the federal minimum wage (pre and post years); and percentile bins in bottom, middle, or top of the wage distribution (P6-P30 and P30-P75).

Additional Period. Table 1A.6 estimates the model in the extended sample spanning the entire period of 1980–2010. As argued above, the peculiar economic conditions in Brazil could potentially weaken the connection between domestic and international commodity prices during the 1980s. Yet column (2) indicates very similar responses in terms of commodity sector employment. Differences arise for the response of the commodity sector wage differential in column (8). In this case, the coefficient for HSG falls by 40%, moving towards the lower bound of the baseline confidence interval. For HSD, we obtain a higher and more precise coefficient compared to the nonsignificant coefficient implied by the baseline specification.

Compared to the period of 1991–2010, the 1980s exhibit another important difference: commodity prices experienced strong losses in the decade. Taking advantage of this qualitatively different price behavior, columns (3) and (6) estimate the model with microregion-specific time trends. Such a specification relies exclusively on differential exposure within-microregion across periods. For this reason, it addresses concerns that shock exposure is picking up secular trends in microregions specialized in the commodities with larger price gains in 1991–2010. Although these additional variables absorb much of the cross-section variation in labor market outcomes, they have little effect on estimated coefficients.[36]

Table 1A.6: Exposure to Commodity Price Shocks and Sector Employment and Wages, Additional Period

	Change in commodity sector employment share			Change in commodity sector average log wage premium		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: High School Graduates						
Commodity price shock	0.035**	0.055**	0.063**	0.441**	0.271*	0.256+
	(0.010)	(0.013)	(0.018)	(0.108)	(0.105)	(0.138)
R^2	0.413	0.336	0.430	0.217	0.214	0.314
Panel B: High School Dropouts						
Commodity price shock	0.061*	0.065**	0.061*	0.030	0.322**	0.402**
	(0.029)	(0.021)	(0.028)	(0.155)	(0.095)	(0.128)
R^2	0.561	0.498	0.612	0.215	0.236	0.403
Baseline Controls						
	Yes	Yes	Yes	Yes	Yes	Yes
Additional Controls						
Microregion-specific time trend	No	No	Yes	No	No	Yes
Sample Period						
Baseline: 1991-2010	Yes	Yes	Yes	Yes	Yes	Yes
Extended: 1980-2010	No	Yes	Yes	No	Yes	Yes

Note. Stacked sample of 518 microregions in baseline sample and 503 microregions in extended sample. All regressions are weighted by the microregion share in national population in 1991. Regressions include macroregion-period dummies and the baseline controls in Table 1.1. Industry composition measured in the initial period of 1991 for baseline sample and of 1980 for extended sample. Commodity sector size controls in extended sample: share in group labor income of other agriculture and mining industries in the commodity sector. Standard Errors clustered by microregion ** p<0.01, * p<0.05, + p<0.10

[36] In omitted exercises, I have estimated the model with microeregion fixed effects in the baseline sample of 1991–2010. In this case, standard errors become five to ten times higher. This increase is related to the high correlation in shock exposure between 1991–2000 and 2000–2010 — the autocorrelation of shock exposure is 0.734. Consequently, there is little within-microregion exposure variation to precisely estimate the coefficient of interest. When the 1980–1991 period is included, there is a significant increase in exposure variation within microregions, leading to the more precise results in Table 1A.6.

Additional Worker Groups. In Table 1A.7, I extend the sample to include female and non-white individuals. With this exercise, I evaluate whether these additional worker groups exhibit similar behaviors in the labor market. This possibility is especially relevant given the large changes in gender and race wage gaps in the period (Ferreira, Firpo, and Messina, 2014). In columns (2) and (6), I include female white individuals without significant changes in estimated coefficients. The inclusion of non-white male individuals entails a more intricate change in estimated coefficients, as shown in columns (3) and (7). Responses of sector employment and wages became weaker for HSG in Panel A, but the opposite is true for HSD in Panel B. These different estimated responses are likely related to differences between white and non-white individuals in terms of unobservable characteristics driving their sectoral allocation.

Table 1A.7: Exposure to Commodity Price Shocks and Sector Employment and Wages, Additional Groups

	Change in commodity sector employment share				Change in commodity sector average log wage premium			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: High School Graduates								
Commodity Price shock	0.035**	0.029**	0.022*	0.017*	0.441**	0.499**	0.318**	0.405**
	(0.010)	(0.008)	(0.009)	(0.008)	(0.108)	(0.112)	(0.103)	(0.103)
R^2	0.413	0.431	0.418	0.452	0.217	0.273	0.257	0.340
Panel B: High School Dropouts								
Commodity Price shock	0.061*	0.078**	0.076*	0.085*	0.030	0.059	0.168+	0.150+
	(0.029)	(0.030)	(0.031)	(0.033)	(0.155)	(0.142)	(0.095)	(0.087)
R^2	0.561	0.594	0.627	0.657	0.215	0.259	0.281	0.313
Baseline Controls								
	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Groups								
Baseline: Male / White	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Include Female	No	Yes	No	Yes	No	Yes	No	Yes
Include Non-white	No	No	Yes	Yes	No	No	Yes	Yes

Note. Stacked sample of 518 microregions in 1991–2000 and 2000–2010. All regressions are weighted by the microregion share in national population in 1991. All regressions include macroregion-period dummies and the baseline controls in Table 1.1. Standard Errors clustered by microregion ** p<0.01, * p<0.05, + p<0.10

This is particularly important among HSG because of the extremely low High School graduation rate among non-white individuals in Brazil.

Additional Labor Market Outcomes. Table 1A.8 investigates the effect of shock exposure on the total quantity of hours supplied by workers in a microregion. Such a response is potentially related to changes in the labor supply of native workers and/or changes in the labor supply of immigrant workers. For HSG and HSD, shock exposure presents a small and statistically nonsignificant relation with the total labor supply of both native and non-native workers. This result is consistent with the assumptions required for identification of comparative and absolute advantage: following the commodity price shock, it is unlikely that a market experienced changes in the productivity distribution due to inflow of new workers from either the home sector or other regions.

1A.4.3 Structural Estimation: Sensitivity Analysis

This section investigates the robustness of the estimates of the structural parameters reported in Section 1.4.5. To this end, I present results obtained with alternative estimators, alternative specifications of the GMM estimator, and alternative estimates of sectoral wage per efficiency unit.

Alternative Estimator. Table 1A.9 investigates the robustness of results to the particular choice of estimator. Column (1) replicates the baseline specification obtained from the estimation of equations (1.19)–(1.21), with the Two-Step GMM estimator and the full vector of disaggregated exposure to price shocks. The remaining columns present the estimates of α_g and A_g obtained, respectively, from the separate estimation of equations (1.19) and (1.20). This procedure is less efficient than the baseline specifications, since it does not use the full structure of the model. That is, the estimation equation-by-equation ignores the overidentification restriction provided by the response in the commodity sector's average wage in (1.21). Nevertheless, this

Table 1A.8: Exposure to Commodity Price Shocks and Total Labor Supply

	Change in Log of Total Labor Supply	Change in Log of Immigrants' Labor Supply
	(1)	(2)
Panel A: High School Graduates		
Commodity price shock	0.110	0.082
	(0.141)	(0.161)
R^2	0.798	0.738
Panel B: High School Dropouts		
Commodity price shock	0.180	0.120
	(0.151)	(0.191)
R^2	0.864	0.852
Baseline Controls		
	Yes	Yes

Note. Stacked sample of 518 microregions in 1991–2000 and 2000–2010. All regressions are weighted by the microregion share in national population in 1991. All regressions include macroregion-period dummies, initial relative group size, and the baseline controls in Table 1.1. Standard Errors clustered by microregion ** p<0.01, * p<0.05, + p<0.10

estimator clearly delineates the source of variation driving the structural estimates.

Column (2) shows that the OLS estimation of these equations yields very different results. In this case, OLS is a biased estimator of the structural parameters because supply shocks generate endogenous responses in sectoral wage per efficiency unit and sector employment composition. For the parameter of comparative advantage, the difference in results between columns (1) and (2) has the expected sign. In equation (1.19), a positive shock in workers' comparative advantage in the commodity sector is equivalent to a negative shock to the relative supply of labor in the non-commodity sector, giving rise to a negative bias in the OLS estimator. The sign of the bias in the absolute advantage parameter is less clear, because it depends on the pattern of selection into the two sectors.

Column (3) presents the 2SLS estimation equation-by-equation using the same set of excluded variables of the baseline specification. This estimator yields point estimates that are similar to the baseline but, as expected, estimates have higher standard errors. Because F-stats are low in column (3), weak instruments are a potential concern that I address in two ways. First, I report the 95% confidence intervals computed by conditional likelihood-ratio (CLR), which are similar to those obtained with the usual asymptotic distribution of the 2SLS estimator. Thus, the qualitative selection patterns inferred from the structural parameters are robust to weak instruments. Second, I estimate the same equations where, as in the reduced-form regressions, the unique instrument is the aggregate exposure to commodity price shocks. In this case, the model is just-identified, and the 2SLS is "unbiased." Column (4) shows that this procedure yields similar point estimates, but standard errors are even higher — especially for α_{HSD} that entails a low F-stat. In general, columns (3) and (4) indicate that these restricted estimators yield similar estimated structural parameters as those obtained with the baseline specification in column (1). In this sense, the joint estimation of equations (1.19)–(1.21) by GMM provides efficiency gains that translate into more precise estimates.

Table 1A.9: Parameters of Comparative and Absolute Adavnatage, Alternative Estimator

Estimator:	**Baseline - GMM**	**OLS**	**2SLS**	**2SLS**
Equations:	(1.19)-(1.21)	(1.19)-(1.20)	(1.19)-(1.20)	(1.19)-(1.20)
	(1)	(2)	(3)	(4)
Panel A: High School Graduates				
α_{HSG}	0.835**	0.087	1.088**	1.510**
	(0.212)	(0.073)	(0.234)	(0.468)
CLR (95CI)	-	-	[0.503, 2.007]	[0.036, 3.303]
F excluded	-	-	1.83	6.62
A_{HSG}	1.966*	0.050	2.056*	1.714+
	(0.935)	(0.155)	(0.971)	(0.985)
CLR (95CI)	-	-	[0.480, 5.861]	[-0.355, 5.602]
F excluded	-	-	2.22	10.38
Panel B: High School Dropouts				
α_{HSD}	0.916*	-0.134*	1.655+	2.212+
	(0.399)	(0.059)	(0.919)	(1.242)
CLR (95CI)	-	-	[0.455, 5.255]	[0.097, 7.081]
F excluded	-	-	1.83	2.78
A_{HSD}	-0.727**	-0.442**	-0.814**	-0.955**
	(0.142)	(0.032)	(0.150)	(0.293)
CLR (95CI)	-	-	[-1.401,-0.560]	[-1.778,-0.433]
F excluded	-	-	6.63	14.79
Excluded Instruments				
Disaggregated exposure	Yes	No	Yes	No
Aggregate exposure	No	No	No	Yes

Note. Stacked sample of 518 microregions in 1991–2000 and 2000–2010. All equations are weighted by the microregion share in national population -n 1991 and include the baseline controls in Table 1.3. Disaggregated exposure to price shocks: quadratic polynomial of regional exposure to world product prices. Aggregated exposure to price shocks: sum of regional exposure to world product prices. Standard Errors clustered by microregion ** p<0.01, * p<0.05, + p<0.10

Alternative Specification. In Table 1A.10, I evaluate the robustness of the structural results to specific choices of the estimation procedure regarding the moment weighting matrix and the set of excluded instruments. Again, column (1) replicates the baseline specification obtained with the Two-Step GMM estimator and the full vector of disaggregated exposure to price shocks.

Columns (2)–(3) estimate the model with alternative moment weighting matrices.

Specifically, column (2) imposes that structural errors in the three equations are independent and, in addition, column (3) imposes that structural errors are homoskedastic (i.e., 2SLS weights). Although point estimates are similar, both estimators yield more imprecise estimates. Such a result is expected since these alternative specifications are less efficient under a general structure of error correlation.

In columns (4)–(5), I estimate the model with restricted sets of excluded instruments. The instrument vector in column (3) is restricted to contain only the exposure to commodity price shocks by category. In this case, estimates are similar for HSG, but the comparative advantage parameter for HSD is lower and less precise. Similar conclusions are obtained when the vector of instruments is further restricted to include only the aggregate exposure to price shocks in Agriculture and Mining.

Alternative Estimates of Wage per Efficiency Unit. The estimation of the structural parameters of comparative and absolute advantage relied on estimated dependent variables — the changes in sector wage per efficiency unit. To address concerns regarding the implementation choices adopted in the estimation of these variables, Table 1A.11 presents structural parameters estimated with alternative measures of the sector wage per efficiency unit. Column (2) shows that the particular choice of bin width has little impact on estimates. A coarser discretization of 2 p.p. bins yields similar point estimates with higher standard errors. These more imprecise results reflect the higher measurement error in dependent variables due to fewer data points used in the estimation of $(\Delta\omega^C_{g,r,t}, \Delta\omega^N_{g,r,t})$ for each group-region-period. Columns (3) and (4) display results with estimates of $(\Delta\omega^C_{g,r,t}, \Delta\omega^N_{g,r,t})$ obtained from equation (1.17) using different control sets. These controls capture potential shocks in labor efficiency across workers in various ranges of income. Column (3) indicates estimated coefficients are very similar without the minimum wage controls. This similarity reflects the high correlation shown in Table 1A.5. However, column (4) shows that the percentile range controls are important in the estimation of structural parameters of absolute advantage.

Table 1A.10: Parameters of Comparative and Absolute Adavnatage, Alternative Specification

	Baseline	Matrix of Moment Weights		Vector of Excluded Variables	
	(1)	(2)	(3)	(4)	(5)
Panel A: High School Graduates					
α_{HSG}	0.835**	0.879**	0.997**	0.924*	1.198**
	(0.212)	(0.185)	(0.237)	(0.359)	(0.453)
A_{HSG}	1.966*	1.759*	2.032*	1.501+	1.730+
	(0.935)	(0.834)	(0.978)	(0.873)	(0.980)
Panel B: High School Dropouts					
α_{HSD}	0.916*	1.302*	1.475+	0.394	0.538
	(0.399)	(0.701)	(0.879)	(0.536)	(0.988)
A_{HSD}	-0.727**	-0.640**	-0.795**	-0.811**	-1.072**
	(0.142)	(0.134)	(0.147)	(0.190)	(0.318)
Optimal Matrix of Moment Weights					
Two-Step GMM weights	Yes	No	No	Yes	Yes
Independence	No	Yes	No	No	No
Independece-Homoskedasticity	No	No	Yes	No	No
Excluded Instruments					
Disaggregated exposure (linear)	Yes	Yes	Yes	Yes	No
Disaggregated exposure (quadratic)	Yes	Yes	Yes	No	No
Aggregated exposure	No	No	No	No	Yes

Note. Stacked sample of 518 microregions in 1991–2000 and 2000–2010. All equations are weighted by the microregion share in national population in 1991 and include the baseline controls in Table 1.3. Disaggregated Excluded Instruments: quadratic polynomial of regional exposure to world product prices. Aggregated exposure to price shocks: sum of regional exposure to world product prices (Agriculture, Mining). Standard Errors clustered by microregion ** $p<0.01$, * $p<0.05$, + $p<0.10$

Table 1A.11: Parameters of Comparative and Absolute Adavnatage, Alternative Estimates of Wage per Efficiency Unit

	Baseline	Alternative estimates of wage per efficiency unit		
	(1)	(2)	(3)	(4)
Panel A: High School Graduates				
α_{HSG}	0.835**	0.650	0.658**	1.576**
	(0.212)	(0.438)	(0.359)	(0.489)
A_{HSG}	1.966*	1.552*	2.020*	0.281
	(0.935)	(0.962)	(0.951)	(0.506)
Panel B: High School Dropouts				
α_{HSD}	0.916*	0.950*	1.515**	0.688
	(0.399)	(0.495)	(0.504)	(0.568)
A_{HSD}	-0.727**	-0.623**	-0.716**	-0.321*
	(0.142)	(0.155)	(0.161)	(0.171)
Baseline Controls				
Percentile below federal minimum wage	Yes	Yes	No	Yes
Percentile in bottom, middle or top	Yes	Yes	Yes	No
Discretization of wage distribution				
Bins of 1 p.p. (N – 88)	Yes	No	Yes	Yes
Bins of 2 p.p. (N – 44)	No	Yes	No	No

Note. Stacked sample of 518 microregions in 1991–2000 and 2000–2010. Two-Step GMM estimator with microregions weighted by their share in the 1991 national population. All equations include the baseline controls in Table 1.3. Excluded instruments: quadratic polynomial of regional exposure to world product prices. Standard Errors clustered by microregion ** p<0.01, * p<0.05, + p<0.10

1A.5 Parametric Restrictions on the Distribution of Comparative and Absolute Advantage

This section discusses prominent distributional assumptions that determine the form of $\alpha_g(.)$ and $\Lambda_g^k(.)$. To simplify notation, I omit subscripts for groups, regions and years.

1A.5.1 Normal Distribution

Particularly important in the selection literature is the case of log-normally distributed sector-specific productivity (Roy, 1951; Heckman and Sedlacek, 1985; Borjas, 1987; Ohnsorge and Trefler, 2007; and Mulligan and Rubinstein, 2008). In my model, this is equivalent to assuming that the sector-specific productivity vector is independently drawn from a bivariate log-normal distribution:

$$\left(\ln L^C(i), \ln L^N(i)\right) \sim \mathcal{N}\left(\begin{bmatrix} \mu_C \\ \mu_N \end{bmatrix}; \begin{bmatrix} \sigma_C^2 & \sigma_{CN} \\ \sigma_{CN} & \sigma_N^2 \end{bmatrix}\right).$$

Because the comparative advantage of individual i is defined as $s(i) = \ln L^C(i) - \ln L^N(i)$, it is straight forward to conclude that $s(i) \sim \mathcal{N}(\mu, \sigma^2)$ where $\mu \equiv \mu_C - \mu_N$ and $\sigma^2 \equiv \sigma_C^2 + \sigma_N^2 - 2\sigma_{CN}$. Thus, $(s(i), a(i))$ is jointly normal with covariance of $Cov(s(i), a(i)) = \sigma_{CN} - \sigma_N^2$ and the distribution of $a(i)$ conditional on $s(i) = s$ is normal with conditional mean given by

$$E\left[a(i)|s(i) = s\right] = \tilde{\mu} + \rho \cdot s \quad \text{s.t.} \quad \tilde{\mu} \equiv (1+\rho)\mu_N - \rho\mu_C, \quad \rho \equiv \frac{\sigma_{CN} - \sigma_N^2}{\sigma_C^2 + \sigma_N^2 - 2\sigma_{CN}};$$

and conditional variance given by

$$V[a(i)|s(i) = s] = \sigma_N^2 - \frac{\left(\sigma_{CN} - \sigma_N^2\right)^2}{\sigma_C^2 + \sigma_N^2 - 2\sigma_{CN}}.$$

By definition, $F(s) \equiv \Phi\left(\frac{s-\mu}{\sigma}\right)$ where $\Phi(.)$ is the CDF of the standard normal distribution. Thus,

$$\alpha(q) \equiv F^{-1}(q) = \mu + \sigma \cdot \Phi^{-1}(q). \tag{1.33}$$

Also, notice that

$$\begin{aligned}\bar{A}^N(l^N) \equiv& \frac{1}{l^N}\int_0^l E\left[a(i)|s(i) = \alpha_g(q)\right] dq \\ =& \bar{\mu} + \frac{\rho}{l^N}\int_0^{l^N} \alpha(q)dq = (\tilde{\mu} + \rho\mu) + \frac{\rho\sigma}{l^N}\int_0^{l^N} \Phi^{-1}(q)dq.\end{aligned}$$

Because $\int_0^{l^N} \Phi^{-1}(q)dq = \phi\left(\Phi^{-1}(l^N)\right)$,

$$\bar{A}^N(l) = \bar{\mu} - (\rho\sigma) \cdot \frac{\phi\left(\Phi^{-1}(l^N)\right)}{l^N} \tag{1.34}$$

where $\bar{\mu} \equiv (\tilde{\mu} + \rho\mu) = \mu_N$.

For completeness, consider the average efficiency in the commodity sector:

$$\bar{A}^C(l^N) \equiv \frac{1}{1-l}\int_l^1 \alpha(q) + E\left[a(i)|s(i) = \alpha(q)\right] dq = (\mu + \bar{\mu}) + \sigma(1+\rho) \cdot \frac{\phi\left(\Phi^{-1}(l^N)\right)}{1-l^N}.$$

Equations (1.33)-(1.34) illustrate the connection between the parameters governing the productivity distribution and the schedules of comparative and absolute advantage. First, the dispersion of comparative advantage, σ, controls the magnitude of the between-sector reallocation of individuals in response to changes in the relative wage per efficiency unit. Second, the sensitivity of the mean absolute advantage to the comparative advantage, ρ, controls the compositional effect of employment on sector average wage.

1A.5.2 Extreme Value Distribution

Recent papers have adopted a productivity distribution of the Fréchet family (Hsieh, Hurst, Jones, and Klenow, 2013; Burstein, Morales, and Vogel, 2015; Galle, Rodriguez-Clare, and Yi, 2015). The main advantage of this distribution is its tractability in the multi-dimensional problem of sectoral choice, allowing for an analytical characterization of the equilibrium with an arbitrary number of sectors. As discussed below, this tractability comes at a price: it imposes a restrictive pattern of selection across sectors.

Specifically, assume that sector-specific productivity is independently drawn from a Fréchet distribution:

$$\left(L^C(i), L^N(i)\right) \sim \exp\left[-\sum_{k=C,N}\left(L^k\right)^{-\kappa}\right]$$

where I assume that $\kappa > 1$ to guarantee finiteness of first-order moments.

First, consider the distribution of comparative advantage:

$$F(s) \equiv Pr\left[s(i) < s\right] = \int_{-\infty}^{\infty} e^{-e^{-\kappa(a+s)}} \kappa e^{-\kappa a} e^{-e^{-\kappa a}}\, da = \int_{-\infty}^{\infty} \kappa e^{-\kappa a} e^{-(1+e^{-\kappa s})e^{-\kappa a}}\, da.$$

Define $x \equiv (1+e^{-\kappa s})e^{-\kappa a}$ such that $dx = -\kappa(1+e^{-\kappa s})e^{-\kappa a}da$. Thus, $F(s) = \frac{1}{1+e^{-\kappa s}}$ and, therefore,

$$\alpha(q) \equiv F^{-1}(q) = \frac{1}{\kappa}\ln\left(\frac{q}{1-q}\right). \tag{1.35}$$

Second, consider the joint distribution of absolute and comparative advantage:

$$Pr[a(i) < \bar{a}; s(i) < s] = \int_{-\infty}^{\bar{a}} \kappa e^{-\kappa a} e^{-(1+e^{-\kappa s})e^{-\kappa a}}\, da = \frac{1}{1+e^{-\kappa s}} e^{-(1+e^{-\kappa s})e^{-\kappa \bar{a}}}.$$

To obtain the average efficiency, notice that the productivity distribution in the non-commodity sector is

$$Pr\left[a(i) < a | s(i) < \alpha\left(l^N\right)\right] = e^{-\left(1+e^{-\kappa\alpha\left(l^N\right)}\right)e^{-\kappa a}} = e^{-\frac{1}{l^N}e^{-\kappa a}} = e^{-e^{-\kappa\left(a+\frac{1}{\kappa}\ln l^N\right)}}$$

where the second equality follows from the definition of $\alpha(.)$.

Since this is a Gumbel distribution with parameters $\beta \equiv 1/\kappa$ and $\mu \equiv -\frac{1}{\kappa}\ln l^N$, the average efficiency in the non-commodity sector is

$$\bar{A}^N(l) = \frac{\gamma}{\kappa} - \frac{1}{\kappa}\ln l^N$$

where γ is the Euler-Mascheroni constant. From this expression, we obtain

$$A(q) = \frac{\gamma - 1}{\kappa} - \frac{1}{\kappa} \ln q. \tag{1.36}$$

Analogously, the productivity distribution in the commodity sector is

$$Pr\left[\ln T^C(i) < a^C | s(i) > \alpha\left(l^N\right)\right] = e^{-e^{-\kappa\left(a^C + \frac{1}{\kappa}\ln\left(1 - l^N\right)\right)}}$$

and, therefore,

$$\bar{A}^C\left(l^N\right) = \frac{\gamma}{\kappa} - \frac{1}{\kappa}\ln\left(1 - l^N\right).$$

The schedules of comparative and absolute advantage in equations (1.35)-(1.36) are fully characterized by the dispersion parameter, κ. If productivity dispersion is low (i.e., κ is high), then a small variation in the relative wage per efficiency unit is associated with a large response of sector employment. In addition, a sector employment expansion causes a decrease in the average sector efficiency whose magnitude is also controlled by the productivity dispersion. In other words, the extreme value distribution only allows for positive selection in both sectors. This very particular pattern of selection has strong implications for the log-wage distribution, implying that both sectors exhibit the same distribution of labor earnings. Specifically, the log-wage distribution in sector k is

$$G^N(y) = e^{-e^{-\kappa\left(y - \omega^N + \frac{1}{\kappa}\ln l^N\right)}} = e^{-e^{-\kappa\left(y - \omega^C + \frac{1}{\kappa}\ln(1 - l^N)\right)}} = G^C(y)$$

where the second equality follows from the employment equation in (1.7).

Finally, the log-wage distribution belongs to the Gumber family and, therefore, the log-wage variance is given by $\pi^2/6\kappa$. Thus, this distributional assumption implies that the dispersion of log wages in a demographic group is constant.

1A.5.3 Log-Linear System: An Example

In this section, I describe a distribution that delivers the log-linear functional forms in Assumption 5. To guarantee finite supply of effective labor units for all parameters, assume that the quantile function of comparative advantage is bounded with the following form:

$$\alpha(q) = \begin{cases} \underline{\alpha} & \text{if } 0 \leq q < \varepsilon \\ \alpha \ln\left[q/(1-q)\right] & \text{if } \varepsilon \leq q < 1-\varepsilon \\ \bar{\alpha} & \text{if } 1-\varepsilon \leq q \leq 1 \end{cases}$$

where $\varepsilon \geq 0$, $\bar{\alpha} \equiv \alpha \frac{\varepsilon - 1}{\varepsilon} \ln(1-\varepsilon) - \alpha \ln(\varepsilon)$, and $\underline{\alpha} \equiv \alpha \ln[\varepsilon/(1-\varepsilon)]$.

Although the comparative advantage distribution has finite moments for every $\varepsilon \geq 0$ and $\alpha > 0$, this is not necessarily true for its moment generating function. Accordingly, the upper bound in the support implies a well defined moment generating function for all $\varepsilon > 0$ and, therefore, a finite supply of effective labor units. For ε arbitrarily small, there is posive employment in both sectors and the empirically relevant portion of the quantile function is that presented in Assumption 5.

Also, assume that the conditional distribution of absolute advantage is normal with a linear conditional mean:

$$\{a(i)|\tilde{s}(i) = \alpha(q)\} \sim \mathcal{N}\left(A_g(q), \sigma^2\right) \quad \text{where } A(q) \equiv \begin{cases} \underline{A} & \text{if } 0 \leq q \leq \varepsilon \\ \bar{A} + A \ln q & \text{if } \varepsilon < q \leq 1 \end{cases}$$

with $A \in \mathbb{R}$, and $\underline{A} \equiv (\bar{A} - A) + A \ln \varepsilon$.

Thus,

$$\bar{A}^N(l^N) \equiv \frac{1}{l^N} \int_0^{l^N} A(q)\ dq = A^N + A \cdot \ln l^N$$

where $A^N = \left(\bar{A} - A\right)$.

By assuming that $\varepsilon < l_g^N < 1 - \varepsilon$,

$$\bar{A}^C(l^N) \equiv \frac{1}{1-l^N} \int_{l^N}^1 \alpha(q) + A(q)\ dq = A^C - (\alpha + A) \cdot \frac{l^N}{1-l^N} \ln l^N - \alpha \cdot \ln(1-l^N).$$

where $A^C \equiv \left(\bar{A} - A\right)$.

Chapter 2

Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade

*Joint with Arnaud Costinot (MIT) and Dave Donaldson (Stanford).

2.1 Introduction

Many interesting questions in international economics are counterfactual ones. Consider China's recent export boom. In the last two decades, its share of world exports has increased from 3% in 1995 to 11% in 2011. What if it had not? What would have happened to other countries around the world?

Given the challenges inherent in isolating quasi-experimental variation in general equilibrium settings, the standard approach to answering such questions has been to proceed in three steps. First, fully specify a parametric model of preferences, technology and trade costs around the world. Second, estimate the model's supply- and demand-side parameters. And finally, armed with this complete knowledge of the world economy, predict what would happen if some of the model's parameters were to change. Such Computational General Equilibrium (CGE) models have long been used to answer a stream of essential counterfactual questions; see e.g. Hertel

(2013) for a survey of the influential GTAP model. Over the last ten years or so, this tradition has been enhanced by an explosion of quantitative work based on gravity models, triggered in large part by the seminal work of Eaton and Kortum (2002).

A key difference between old CGE models, like GTAP, and new CGE models, like Eaton and Kortum (2002), is parsimony. The latest version of the GTAP model described in Hertel, McDougall, Narayanan, and Aguiar (2012) has more than 13,000 structural parameters. Counterfactual analysis in the Eaton and Kortum (2002) model can be conducted using knowledge of only one: the trade elasticity. Parsimony is valuable. But it hinges on strong functional form assumptions that may hinder the credibility of counterfactual predictions. The goal of this paper is to explore the extent to which one may maintain parsimony, but dispense with functional-form assumptions. In a nutshell, can we relax Eaton and Kortum's (2002) strong functional-form assumptions without circling back to GTAP's 13,000 parameters?

Our starting point is the equivalence between neoclassical economies and *reduced exchange economies* in which countries simply trade factor services. Formally, we consider a world economy comprising a representative agent in each country, constant returns to scale in production, and perfect competition in all markets. In this general environment we show that for any competitive equilibrium there is an equilibrium in a reduced exchange economy that is equivalent in terms of welfare, factor prices and the factor content of trade—and further, that the converse is also true.

This equivalence is important for its simplifying power: a reduced exchange economy in which countries act as if they trade factor services can be characterized fully by an analysis of the *reduced factor demand system* that summarizes all agents' preferences over factor services. Thus for a number of counterfactual questions, like the effects of uniform changes in trade costs, one does not need the complete knowledge of demand and production functions across countries and industries. For instance, one does not need to know the cross-price elasticity between French compact cars and Italian cotton shirts or between Korean flat screen TVs and Spanish heirloom tomatoes. Similarly, one does not need to know productivity in these various economic activities around the world. All one needs to know is the cross-price elasticity between factors from different countries. This basic observation encapsulates how we propose to reduce the dimensionality of what needs to be estimated for counterfactual analysis—the reduced factor demand system—without imposing strong functional-form assumptions.[1]

Our second theoretical result establishes that, as long as the reduced factor demand system is invertible, knowledge of this demand system as well as measures of the factor content of trade and factor payments in some initial equilibrium are sufficient to construct counterfactual predictions about the effect of changes in trade costs and factor endowments. This result provides a nonparametric generalization of the methodology popularized by Dekle, Eaton, and Kortum (2008). Their analysis focuses on a Ricardian economy in which the reduced labor demand system takes the Constant Elasticity of Substitution (CES) form. This functional form assumption, however, is not a critical condition for the previous approach to succeed; only the invertibility of the reduced factor demand is.

The procedure that we propose to make counterfactual predictions relies on knowledge of the reduced factor demand system. In gravity models, such systems are implicitly assumed to be CES. Hence, a single trade elasticity can be estimated by regressing the log of bilateral flows on an exogenous shifter of the log of bilateral trade costs, like tariffs or freight costs. Our final set of theoretical results demonstrates that this approach can be pushed further than previously recognized. Namely, we provide sufficient conditions under which, given measures of the factor content of trade and observable shifters of trade costs, reduced factor demand systems can be nonparametrically identified using the same exclusion restrictions. As with our counterfactual results, the invertibility of the reduced factor demand remains the critical assumption; strong functional form assumptions can be dispensed with.

We conclude our paper by applying our general results to one particular coun-

[1]It is worth emphasizing that this approach to dimensionality-reduction does not hinge on any assumption about the number of goods and factors in the world. Regardless of whether there are more goods than factors, the point is that one can estimate a single reduced demand system for factors rather than estimate multiple production functions—that determine how factors are demanded by producers of goods—and utility functions—that determine how goods are demanded by consumers. Of course, the fewer factors there are, the easier the estimation of the reduced factor demand system is.

terfactual question: What would have happened to other countries if China had remained closed? In practice, data limitations are severe—Leamer's (2010) elusive land of "Asymptopia" is far away—and estimation of a reduced factor demand system must, ultimately, proceed parametrically. So the final issue that needs to be tackled is how to parametrize and estimate a reduced factor demand system without taking a stance on particular micro-foundations. We offer the following rules of thumb: (*i*) be as flexible as possible given data constraints; (*ii*) allow flexibility along the dimensions that are more likely to be relevant for counterfactual question of interest; and (*iii*) use the source of variation in the data under which demand is nonparametrically identified.[2]

Towards this goal in the present context, we introduce a strict generalization of CES, which we refer to as *mixed CES*, inspired by the work of Berry (1994) and Berry, Levinsohn, and Pakes (1995) in industrial organization. Like in a standard gravity model, we assume the existence of a composite factor in each country so that the factor content of trade between any pair of countries is equal to their bilateral trade flow. Compared to a standard gravity model, however, our demand system features two new structural parameters that measure the extent to which exporters that are closer in terms of either market shares or some observable characteristic, which we take to be GDP per capita, tend to be closer substitutes. Under CES, when China gains market share, Indian and French exports must be affected equally. By contrast, the mixed CES demand system allows data to speak to whether this "independence of irrelevant alternatives" embodied in CES holds empirically or not.

After estimating our mixed CES demand system for 37 large exporters using data on bilateral trade flows and freight costs from 1995 to 2011, we conclude that rich countries tend to gain relatively more than poor countries from China's integration with the rest of the world—that is, rich countries would have been relatively worse off if Chinese trade costs had counterfactually remained at their 1995 value from 1995 to 2011. Under the restriction that demand is CES, no such pattern emerges.

[2]In their original paper on the CES function, Arrow, Chenery, Minhas, and Solow (1961) note that one of its attractive features is that it is "the most general function which can be computed on a suitable slide rule." Computing power has since improved.

Up to this point, we have emphasized the feasibility and potential benefits of our new approach to counterfactual and welfare analysis. It should be clear that our approach also has important limitations. We discuss these further below but four deserve emphasis here. First, the equivalence result on which we build heavily relies on the efficiency of perfectly competitive markets. This does not mean that our approach will necessarily fail if one were to relax the assumption of perfect competition or introduce distortions—indeed, Arkolakis et al. (2012b) and Arkolakis et al. (2012a) offer examples in this vein that cover a number of influential modeling approaches—but it is fair to say that it is much less likely to be useful in such circumstances. Second, the scope of the counterfactual exercises that we consider is limited by the restriction that the shape of the reduced demand system remains stable. Uniform changes in iceberg trade costs satisfy this condition, but many interesting shocks do not, a point we come back to in Section 2.4.3. Third, the restriction that the demand system is invertible implicitly excludes zeros in bilateral factor trade. So our nonparametric approach does not solve the "zeros issue" in standard gravity models.[3] Fourth, the estimation of a reduced factor demand system requires that the factor content of trade be measured accurately. Since the seminal work of Leontief (1953), multiple generations of trade economists have combined input-output matrices with trade data to do so, but the high-level of aggregation of such matrices leaves open the possibility of mis-measurement, a point emphasized more recently by Burstein and Vogel (2010).[4]

[3]This is a version of the "new goods problem" that is common in many demand settings (Bresnahan and Gordon, 2008). Just as in those settings, one can typically place a lower bound on the welfare effects of a counterfactual by requiring that zeros cannot become positive. For our purposes, the more specific question is whether the challenge posed by zeros in the data is alleviated or worsened by the study of reduced factor demand relative to standard gravity approaches. The answer depends on the assumptions that one makes about the number of goods and factors. If one assumes the existence of a composite factor in each country, as we do in our empirical analysis, then focusing on factor demand reduces the prevalence of zeros relative to any analysis that would focus on trade in goods.

[4]In particular, national input-output matrices do not disaggregate factor payments by destination within each producing country-times-industry cell. The implicit assumption used to measure the factor content of trade in the empirical literature therefore is that factor intensity is constant across destinations. Since micro-level evidence, e.g. Bernard and Jensen (1999), suggests systematic variation in factor intensity between firms that serve domestic and foreign markets, one could potentially improve on the measurement of the factor content of trade by combining aggregate data from the national accounts and micro-level data in a consistent way. We do not attempt to do so in

The rest of this chapter is organized as follows. Section 2.2 discusses the related literature. Section 2.3 establishes our main equivalence result. Section 2.4 uses this result to conduct counterfactual and welfare analysis. Section 2.5 provides sufficient conditions for nonparametric identification. Section 2.6 estimates factor demand. Section 2.7 uses these estimates to study the consequences of China's integration with the rest of the world. Section 2.8 offers some concluding remarks.

2.2 Related Literature

This paper combines old ideas from general equilibrium theory with recent methods from industrial organization and international trade to develop a new way of constructing counterfactual predictions in an open economy.

From the general equilibrium literature, we borrow the idea that, for many purposes, production economies may be reduced to exchange economies; see e.g. Taylor (1938), Rader (1972), and Mas-Colell (1991). Early applications of this idea to international trade can be found in Meade (1952), Helpman (1976), Woodland (1980), Wilson (1980), and Neary and Schweinberger (1986). Among those, Helpman (1976), Wilson (1980), and Neary and Schweinberger (1986) are most closely related. Helpman (1976) shows how to reduce computation time necessary to solve for trade equilibria by focusing on the excess demand for factors, whereas Neary and Schweinberger (1986) introduce the concept of direct and indirect factor trade utility functions and use revealed-preference arguments to generalize the Heckscher-Ohlin Theorem. Finally, Wilson (1980) demonstrates that the analysis of the Ricardian model can be reduced to the analysis of an exchange model in which each country trades its own labor for the labor of other countries.

One can think of the starting point of our paper as a generalization of Wilson's (1980) equivalence result to any neoclassical trade model. Compared to the aforementioned papers, our main contribution is to show how the equivalence between

this paper, but we note that, according to our theoretical results, any researcher interested in our counterfactual exercises would also be affected by this issue, albeit perhaps less explicitly.

neoclassical trade models and exchange models can be used as a tool for counterfactual and welfare analysis using commonly available data on trade flows, factor payments, and trade costs. Here, reduced exchange models are a first step towards measurement and estimation, not an analytical device for studying the theoretical properties of competitive equilibria.

We view our paper as a bridge between the recent gravity literature, reviewed in Costinot and Rodríguez-Clare (2013) and Head and Mayer (2013), and the older neoclassical trade literature, synthesized in Dixit and Norman (1980).[5] With the former, we share an interest in combining theory and data to shed light on counterfactual questions. With the latter, we share an interest in robust predictions, free of strong functional form assumptions. Since data is limited, there is a tension between these two goals. To make progress on the first, without giving up on the second, we therefore propose to use factor demand as a sufficient, albeit potentially high dimensional, statistic. This strategy can be thought of as a nonparametric generalization of Arkolakis, Costinot, and Rodríguez-Clare's (2012b) approach to counterfactual and welfare analysis. Ultimately, there is nothing special about gravity models. They are factor demand systems, like any other neoclassical trade model. And like any demand system, factor demand systems can be estimated using data on quantities, prices, and some instrumental variables. Once this basic econometric issue is recognized, it becomes natural to turn to the recent results on the nonparametric identification of demand in differentiated markets; see e.g. Berry, Gandhi, and Haile (2013) and Berry and Haile (2014).

Our analysis is also related to the large empirical literature on the determinants of the factor content of trade. A long and distinguished tradition—e.g. Bowen, Leamer, and Sveikauskas (1987), Trefler (1993), Trefler (1995), and Davis and Weinstein (2001)—aims to test the Heckscher-Ohlin-Vanek model by comparing the factor content of trade measured in the data to the one predicted by the model under various assumptions about technology, preferences, and trade costs (or lack thereof).

[5]Further results about the theoretical properties of gravity models, including sufficient conditions for existence and uniqueness of equilibria, can be found in Allen, Arkolakis, and Takahashi (2014).

Our goal is instead to estimate a factor demand system and use these estimates to conduct counterfactual and welfare analysis. In order to test or assess the fit of the Heckscher-Ohlin-Vanek model in some observed equilibrium, one does not need to know the cross-price elasticities between factors from different countries. Indeed, such tests are often conducted under the assumption that factor price equalization holds, up to some factor-augmenting productivity differences, so that factors from different countries are assumed to be perfect substitutes. For our purposes, knowledge of cross-price elasticities is critical.

Finally, our work has implications for the debate about the extent to which the factor content of trade observed in one equilibrium can be used (or not) for measuring the consequences of international trade on inequality; see e.g. Deardorff and Staiger (1988), Krugman (2000) and Leamer (2000). Such a discussion implicitly boils down to the question of what shape factor demand systems take and whether factors from different countries are perfect substitutes (or not). Our analysis points towards estimating these systems as a way to settle such debates.

2.3 Neoclassic Trade Models as Exchange Models

2.3.1 Neoclassical Trade Model

Consider a world economy comprising $i = 1, ..., I$ countries, $k = 1, ..., K$ goods, and $n = 1, ..., N$ primary factors of production. Factor supply is inelastic. $\boldsymbol{\nu}_i \equiv \{\nu_i^n\}$ denotes the vector of factor endowments in country i.

Preferences. In each country i, there is a representative agent with utility,

$$u_i = u_i(\boldsymbol{q}_i),$$

where $\boldsymbol{q}_i \equiv \{q_{ji}^k\}$ is the vector of quantities consumed in country i and u_i is strictly increasing, quasiconcave, and differentiable. The previous notation allows, but does not require, u_i to depend only on $\{\sum_j q_{ji}^k\}$. Hence, we explicitly allow, but do not require, goods produced in different countries to be imperfect substitutes. Compared to recent quantitative work in the field, we impose no functional-form assumptions on u, though the assumption of a representative agent is by no means trivial.

Technology. Production is subject to constant returns to scale. Output of good k in country i that is available for consumption in country j is given by

$$q_{ij}^k = f_{ij}^k(\boldsymbol{l}_{ij}^k),$$

where $\boldsymbol{l}_{ij}^k \equiv \{l_{ij}^{nk}\}$ is the vector of factors used to produce good k in country i for country j, and f_{ij}^k is strictly increasing, concave, differentiable, and homogeneous of degree one.

Compared to recent quantitative work in the field, we again impose no functional-form assumptions on f_{ij}^k. For instance, it is standard in the existing literature to assume that the difference between production functions across different destinations derive from iceberg trade costs. This special case corresponds to the existence of Hicks-neutral productivity shifters, τ_{ij}^k, such that

$$f_{ij}^k(\boldsymbol{l}_{ij}^k) \equiv f_i^k(\boldsymbol{l}_{ij}^k)/\tau_{ij}^k.$$

In an Arrow-Debreu sense, a good in our economy formally corresponds to a triplet (i, j, k), whereas a factor formally corresponds to a pair (i, n), with the usual wide interpretation. Though we impose constant returns to scale, decreasing returns in production can be accommodated in the usual way by introducing additional primary factors of production. Endogenous labor supply can be dealt with by treating leisure as another nontradable good. Multinational production, as in Ramondo and Rodríguez-Clare (2013), can also be accommodated by expanding the set of goods and using a different index k for goods whose "technologies" originate in different countries. Finally, the assumption of no joint production can be relaxed substantially. The key requirement for our equivalence result is that there is no component of production that is joint across destination markets, as would be the case in the presence of fixed costs of production.[6] Besides the absence of increasing returns in each sector, the only substantial restriction imposed on technology is the absence of intermediate goods. We discuss how to incorporate such goods in Section 2.4.3.

Competitive equilibrium. Goods markets and factor markets are perfectly competitive. We let p_{ij}^k denote the price of good k from country i in country j and w_i^n denote the price of factor n in country i. Letting $\boldsymbol{q} \equiv \{\boldsymbol{q}_i\}$, $\boldsymbol{l} \equiv \{\boldsymbol{l}_{ij}^k\}$, $\boldsymbol{p} \equiv \{p_{ij}^k\}$, and $\boldsymbol{w} \equiv \{w_i^n\}$, we can then define a competitive equilibrium as follows.

Definition 1. *A competitive equilibrium corresponds to* $(\boldsymbol{q}, \boldsymbol{l}, \boldsymbol{p}, \boldsymbol{w})$ *such that:*

i. *consumers maximize their utility:*

$$\boldsymbol{q}_i \in argmax_{\tilde{\boldsymbol{q}}_i} u_i(\tilde{\boldsymbol{q}}_i) \tag{2.1}$$

$$\sum_{j,k} p_{ji}^k \tilde{q}_{ji}^k \leq \sum_n w_i^n \nu_i^n \text{ for all } i; \tag{2.2}$$

ii. *firms maximize their profits:*

$$\boldsymbol{l}_{ij}^k \in argmax_{\tilde{\boldsymbol{l}}_{ij}^k} p_{ij}^k f_{ij}^k(\tilde{\boldsymbol{l}}_{ij}^k) - \sum_n w_i^n \tilde{l}_{ij}^{nk} \text{ for all } i,\ j, \text{ and } k; \tag{2.3}$$

iii. *goods markets clear:*

$$q_{ij}^k = f_{ij}^k(\boldsymbol{l}_{ij}^k) \text{ for all } i,\ j, \text{ and } k; \tag{2.4}$$

iv. *factors markets clear:*

$$\sum_{j,k} l_{ij}^{fk} = \nu_i^n \text{ for all } i \text{ and } n. \tag{2.5}$$

[6]This implies that our theoretical framework can accomodate economies in which there are multiple regions within a country and firms in each region jointly produce goods and amenities.

2.3.2 Reduced Exchange Model

An old idea in general equilibrium theory is that it is often simpler to analyze the competitive equilibrium of a neoclassical model with production by studying instead a fictitious endowment economy in which consumers directly exchange factor services; see e.g. Taylor (1938), Rader (1972), and Mas-Colell (1991). Although this idea is often associated in the trade literature with the Heckscher-Ohlin model, it applies equally well to the Ricardian model of trade; see e.g. Wilson (1980). We now offer a formal proof of the equivalence between a general neoclassical trade model and an exchange economy, in terms of the factor content of trade, factor prices, and welfare. This equivalence result will be the backbone of our approach to counterfactual and welfare analysis in Section 2.4.

Starting from the neoclassical trade model of Section 2.3.1, we can define the reduced utility function over primary factors of production in country i as

$$U_i(\boldsymbol{L}_i) \equiv \max_{\tilde{\boldsymbol{q}}_i, \tilde{\boldsymbol{l}}_i} u_i(\tilde{\boldsymbol{q}}_i) \tag{2.6}$$

$$\tilde{q}_{ji}^k \leq f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) \text{ for all } j \text{ and } k, \tag{2.7}$$

$$\sum_k \tilde{l}_{ji}^{nk} \leq L_{ji}^n \text{ for all } j \text{ and } n, \tag{2.8}$$

where $\boldsymbol{L}_i \equiv \{L_{ji}^n\}$ denotes the vector of total factor demands from country i. It describes the maximum utility that a consumer in country i would be able to achieve if she were endowed with $\boldsymbol{L}_i$ and had access to the technologies of all firms around the world.[7] One can check that $U_i(\cdot)$ is strictly increasing and quasiconcave, though not necessarily strictly quasiconcave, even if $u_i(\cdot)$ is.[8] In particular, $U_i(\cdot)$ is likely to

[7]The above definition is closely related to, but distinct from, the notion of the "direct factor trade utility function" introduced in Neary and Schweinberger (1986). The distinction comes from the fact that Neary and Schweinberger's (1986) factor trade utility function measures the maximum utility attainable if all consumption must be produced using the techniques of the home country. In our definition, each country is assumed to have access to the techniques in all other countries, inclusive of trade costs. This distinction is important. As we will show in a moment, the factor content of trade derived from solving (2.6) coincides with the factor content of trade in the competitive equilibrium. This would no longer be true if one were to maximize Neary and Schweinberger's (1986) factor trade utility function.

[8]The fact U_i is strictly increasing in $\boldsymbol{L}_i$ is trivial. To see that U_i is quasi-concave, take two vectors of factor demand, $\boldsymbol{L}_i$ and $\tilde{\boldsymbol{L}}_i$, and $\alpha \in [0,1]$. Let $(\boldsymbol{q}, \boldsymbol{l})$ and $(\tilde{\boldsymbol{q}}, \tilde{\boldsymbol{l}})$ be the solution of (2.6) associated

be linear whenever production functions are identical around the world. While this situation is obviously knife-edge, this is the special case on which the Heckscher-Ohlin model of trade focuses. We therefore explicitly allow for such situations below.

Letting $\boldsymbol{L} \equiv \{\boldsymbol{L}_i\}$, we can define a competitive equilibrium of the reduced exchange model or, in short, a reduced equilibrium.

Definition 2. *A reduced equilibrium corresponds to* $(\boldsymbol{L}, \boldsymbol{w})$ *such that:*

i. *consumers maximize their reduced utility:*

$$\begin{aligned} \boldsymbol{L}_i \in argmax_{\tilde{L}_i} U_i(\tilde{\boldsymbol{L}}_i) \\ \sum_{j,n} w_j^n \tilde{L}_{ji}^n \leq \sum_n w_i^n v_i^n \textit{ for all } i; \end{aligned} \tag{2.9}$$

ii. *factor markets clear:*

$$\sum_j L_{ij}^n = \nu_i^n \textit{ for all } i \textit{ and } n. \tag{2.10}$$

Our main equivalence result can be stated as follows.

Proposition 1. *For any competitive equilibrium,* $(\boldsymbol{q}, \boldsymbol{l}, \boldsymbol{p}, \boldsymbol{w})$, *there exists a reduced equilibrium,* $(\boldsymbol{L}, \boldsymbol{w})$, *with: (i) the same factor prices,* $\boldsymbol{w}$*; (ii) the same factor content of trade,* $L_{ij}^n = \sum_k l_{ij}^{nk}$ *for all i, j, and n; and (iii) the same welfare levels,* $U_i(\boldsymbol{L}_i) = u_i(\boldsymbol{q}_i)$ *for all i. Conversely, for any reduced equilibrium,* $(\boldsymbol{L}, \boldsymbol{w})$, *there exists a competitive equilibrium,* $(\boldsymbol{q}, \boldsymbol{l}, \boldsymbol{p}, \boldsymbol{w})$, *such that conditions (i)-(iii) hold.*

The formal proof of Proposition 1 can be found in Appendix 2A.1.1. The basic arguments are similar to those used in proofs of the First and Second Welfare Theorems. This should be intuitive. In the reduced equilibrium, each representative agent solves a country-specific planning problem, as described in (2.6). Thus,

with $\boldsymbol{L}_i$ and $\tilde{\boldsymbol{L}}_i$, respectively. Now consider $(\bar{\boldsymbol{q}}, \bar{\boldsymbol{l}}) \equiv \alpha(\boldsymbol{q}, \boldsymbol{l}) + (1-\alpha)(\tilde{\boldsymbol{q}}, \tilde{\boldsymbol{l}})$. By construction, $\bar{l}$ trivially satisfies (2.8). Since f_{ji}^k is concave, we also have $\bar{q}_{ji}^k \leq \alpha f_{ji}^k(\boldsymbol{l}_{ji}^k) + (1-\alpha) f_{ji}^k(\tilde{\boldsymbol{l}}_{jn}^k) \leq f_{ji}^k(\bar{\boldsymbol{l}}_{jn}^k)$ for all j and k. This implies $U_i(\alpha \boldsymbol{L}_i + (1-\alpha)\tilde{\boldsymbol{L}}_i) \geq u_i(\bar{\boldsymbol{q}}) \geq \min\{u_i(\boldsymbol{q}), u_i(\tilde{\boldsymbol{q}})\} = \min\{U_i(\boldsymbol{L}_i), U_i(\tilde{\boldsymbol{L}}_i)\}$ where the second inequality follows from the quasiconcavity of u_i.

showing that any competitive equilibrium is associated with an equivalent reduced equilibrium implicitly relies on the efficiency of the original competitive equilibrium, which the First Welfare Theorem establishes. Similarly, showing that any reduced equilibrium is associated with an equivalent competitive equilibrium implicitly relies on the ability to decentralize efficient allocations, which the Second Welfare Theorem establishes. The key distinction between Proposition 1 and standard Welfare Theorems is that the reduced equilibrium is not a global planner's problem; it remains a decentralized equilibrium in which countries fictitiously trade factor services and budgets are balanced country by country. Broadly speaking, we do not go all the way from the decentralized equilibrium to the global planner's problem, but instead stop at a hybrid reduced equilibrium, which combines country-specific planner's problems with perfect competition in factor markets.[9]

According to Proposition 1, if one is interested in the factor content of trade, factor prices, or welfare, then one can always study a reduced equilibrium—whose primitives are the reduced utility functions, $\{U_i\}$, and the endowments, $\{\boldsymbol{\nu}_i\}$—rather than a competitive equilibrium—whose primitives are the utility functions, $\{u_i\}$, the endowments, $\{\boldsymbol{\nu}_i\}$, and the production functions, $\{f_{ij}^k\}$. In order to do counterfactual and welfare analysis, one does not need to have *direct* knowledge of both the utility functions, $\{u_i\}$, and the production functions, $\{f_{ij}^k\}$. Instead, one merely needs to know how they *indirectly* shape, $\{U_i\}$, and in turn global factor demand—that is, the solution of the reduced utility maximization problem (2.9).[10]

[9]This implies, in particular, that the convexity of preferences and technology, which is central in the proof of the Second Welfare Theorem, plays no role in the proof of Proposition 1. In the Second Welfare Theorem, convexity is invoked for Lagrange multipliers, and in turn, competitive prices to exist. Here, competitive prices for goods can be directly constructed from factor prices in the reduced equilibrium using zero-profit conditions.

[10]This is true regardless of whether the competitive and reduced equilibria are unique. Formally, Proposition 1 establishes that the set of factor prices, factor content of trade, and welfare levels that can be observed in a competitive equilibrium is the same as the set of factor prices, factor content of trade, and welfare levels that can be observed in a reduced equilibrium. Whether the previous sets are singletons is irrelevant for our equivalence result.

2.4 Counterfactual and Welfare Analysis

We start by considering counterfactual shocks to preferences and factor endowments in a reduced exchange model. In this context, we show how to extend the exact algebra popularized by Dekle, Eaton, and Kortum (2008) in the context of a CES demand system to general, non-CES environments. Perhaps surprisingly, the critical assumption required for the previous approach to succeed is not strong functional form assumptions on the structure of factor demand, it is merely its invertibility. Using the equivalence result from Section 2.3, we then show how the previous counterfactual predictions can be used to study the effect of changes in endowments and technology in a general neoclassical model of trade.

2.4.1 Reduced Counterfactuals

Consider a reduced exchange model in which the reduced utility function over primary factors can be expressed as

$$U_i(\boldsymbol{L}_i) \equiv \bar{U}_i(\{L^n_{ji}/\tau^n_{ji}\}), \tag{2.11}$$

where $\bar{U}_i$ is a strictly increasing and quasi-concave utility function and $\tau^n_{ji} > 0$ are exogenous preference shocks. The counterfactual question that we are interested in here is: *What are the effects of a change from $(\boldsymbol{\tau}, \boldsymbol{\nu}) \equiv \{\tau^n_{ji}, v^n_i\}$ to $(\boldsymbol{\tau}', \boldsymbol{\nu}') \equiv \{(\tau^n_{ji})', (v^n_i)'\}$ on trade flows, factor prices, and welfare?*

Trade Flows and Factor Prices

For each country i, let $\boldsymbol{L}_i(\boldsymbol{w}, y_i|\boldsymbol{\tau})$ denote the set of solutions to the utility maximization problem (2.9) as a function of factor prices, $\boldsymbol{w}$, income, $y_i \equiv \sum_n w^n_i \nu^n_i$, and preference parameters, $\boldsymbol{\tau}$. This corresponds to the Marshallian demand for factor services in the reduced equilibrium. The associated vectors of factor expenditure shares are then given by

$$\boldsymbol{\chi}_i(\boldsymbol{w}, y_i|\boldsymbol{\tau}) \equiv \{\{x_{ji}^n\}|x_{ji}^n = w_j^n L_{ji}^n / y_i \text{ for some } \boldsymbol{L}_i \in \boldsymbol{L}_i(\boldsymbol{w}, y_i|\boldsymbol{\tau})\}.$$

Since preference shocks are multiplicative, expenditure shares must depend only on the "effective" factor prices, $\boldsymbol{\omega}_i \equiv \{w_j^n \tau_{ji}^n\}$. We can therefore write, with a slight abuse of notation and without risk of confusion, $\boldsymbol{\chi}_i(\boldsymbol{w}, y_i|\boldsymbol{\tau}) \equiv \boldsymbol{\chi}_i(\boldsymbol{\omega}_i, y_i)$. Using the previous notation, the equilibrium conditions (2.9) and (2.10) can then be expressed compactly as

$$\boldsymbol{x}_i \in \boldsymbol{\chi}_i(\boldsymbol{\omega}_i, y_i) \text{ for all } i, \tag{2.12}$$

$$\sum_j x_{ij}^n y_j = y_i^n, \text{ for all } i \text{ and } n, \tag{2.13}$$

where $\boldsymbol{x}_i \equiv \{x_{ji}^n\}$ denotes the vector of factor expenditure shares in country i and $y_i^n \equiv w_i^n \nu_i^n$ denotes payments to factor n.

A standard gravity model, such as the one developed by Anderson and Van Wincoop (2003) and Eaton and Kortum (2002), corresponds to the special case in which there is only one factor of production in each country and factor demand is CES. Omitting the index for factors, n, such models require

$$\chi_{ji}(\boldsymbol{\omega}_i, y_i) = \frac{(\omega_{ji})^\epsilon}{\sum_l (\omega_{li})^\epsilon}, \text{ for all } j \text{ and } i, \tag{2.14}$$

for some trade elasticity $\epsilon > 1$; see Arkolakis, Costinot, and Rodríguez-Clare (2012b) and Costinot and Rodríguez-Clare (2013) for further discussion.[11] We now proceed under the assumption that χ_i is known, but dispense with any functional-form restriction.

In what follows we refer to χ_i as the factor demand system in country i. The only assumption that we impose on the factor demand system is its invertibility.

A1 [Invertibility]. *In any country i, for any vector of expenditure shares, $\boldsymbol{x} > 0$,*

[11] In this case, total income, y_i, has no effect on factor expenditure shares because of homotheticity.

and any income level, $y > 0$, there exists a unique vector of factor prices (up to a normalization) such that $\boldsymbol{x} \in \chi_i(\boldsymbol{\omega}, y)$, which we denote $\boldsymbol{\chi}_i^{-1}(\boldsymbol{x}, y)$.

In line with A1, we restrict ourselves from now on to equilibria such that $\boldsymbol{x}_i > 0$ for all i.[12] Let $\boldsymbol{x}_i' \equiv \{(x_{ji}^n)'\}$ and w' denote the counterfactual expenditure shares and factor prices in the counterfactual equilibrium with preference parameters and endowments given by $(\boldsymbol{\tau}', \boldsymbol{\nu}')$. The basic idea behind the exact hat algebra of Dekle, Eaton, and Kortum (2008) is twofold: (i) focus on the proportional changes in expenditure shares and factor prices, $\hat{\boldsymbol{x}}_i \equiv \{(x_{ji}^n)'/x_{ji}^n\}$ and $\hat{\boldsymbol{w}} \equiv \{(w_j^n)'/w_j^n\}$, caused by proportional changes in preferences and endowments, $\hat{\boldsymbol{\tau}} \equiv \{(\tau_{ij}^n)'/\tau_{ij}^n\}$ and $\hat{\boldsymbol{\nu}} \equiv \{(\nu_i^n)'/\nu_i^n\}$; and ($ii$) use data on expenditure shares, x_{ij}^n, as well as factor payments, y_i^n, in the initial equilibrium to extract information about the underlying structural parameters of the model. There is nothing in this general strategy that hinges on the demand system being CES. Invertibility of factor demand is the critical assumption.

Let us start by rewriting the equilibrium conditions (2.12) and (2.13) at the counterfactual values of the preference and endowment parameters, $(\boldsymbol{\tau}', \boldsymbol{\nu}')$:

$$\boldsymbol{x}_i' \in \boldsymbol{\chi}_i(\boldsymbol{\omega}_i', y_i') \text{ for all } i,$$

$$\sum_j (x_{ij}^n)' y_j' = (y_i^n)', \text{ for all } i \text{ and } n.$$

These two conditions, in turn, can be expressed in terms of proportional changes,

$$\{\hat{x}_{ji}^n x_{ji}^n\} \in \boldsymbol{\chi}_i(\{\hat{w}_j^n \hat{\tau}_{ji}^n \omega_{ji}^n\}, \sum_n \hat{w}_i^n \hat{\nu}_i^n y_i^n) \text{ for all } i,$$

$$\sum_j \hat{x}_{ij}^n x_{ij}^n (\sum_n \hat{w}_j^n \hat{\nu}_j^n y_j^n) = \hat{w}_i^n \hat{\nu}_i^n y_i^n, \text{ for all } i \text{ and } n,$$

where we have used the fact that total income in the counterfactual equilibrium is equal to the sum of total factor income, $y_i' = \sum_n (y_i^n)'$. Finally, using A1, we can

[12] By $\boldsymbol{x}_i > 0$, we formally mean $x_{ji}^n > 0$ for all j and n. Zero expenditure shares create two issues. First, factor prices can no longer be inferred from expenditure shares. Typically, they can only be bounded from below. Second, proportional changes between the initial and counterfactual equilibrium, on which our analysis focuses, are no longer well-defined. Empirically, zeros are irrelevant for the sample of countries and the level of aggregation at which we will conduct our estimation and counterfactual simulation.

eliminate the effective factor prices in country i, ω_{ji}^n, from the previous expression using initial expenditure shares, $\boldsymbol{x}_i$, and income levels, y_i. This leads to the following proposition.

Proposition 2. *Suppose that A1 holds and the reduced utility function satisfies (2.11). Then the proportional changes in expenditure shares and factor prices, $\hat{\boldsymbol{x}}_i \equiv \{(x_{ij}^n)'/x_{ij}^n\}$ and $\hat{\boldsymbol{w}} \equiv \{(w_j^n)'/w_j^n\}$, caused by the proportional changes in preferences and endowments, $\hat{\boldsymbol{\tau}} \equiv \{(\tau_{ij}^n)'/\tau_{ij}^n\}$ and $\hat{\boldsymbol{\nu}} \equiv \{(\nu_i^n)'/\nu_i^n\}$, solve*

$$\{\hat{x}_{ji}^n x_{ji}^n\} \in \boldsymbol{\chi}_i(\{\hat{w}_j^n \hat{\tau}_{ji}^n (\chi_{ji}^n)^{-1}(\boldsymbol{x}_i, y_i)\}, \sum_n \hat{w}_i^n \hat{\nu}_i^n y_i^n) \textit{ for all } i, \tag{2.15}$$

$$\sum_j \hat{x}_{ij}^n x_{ij}^n (\sum_n \hat{w}_j^n \hat{\nu}_j^n y_j^n) \doteq \hat{w}_i^n \hat{\nu}_i^n y_i^n, \textit{ for all } i \textit{ and } n. \tag{2.16}$$

Once proportional changes in expenditure shares and factor prices have been solved for, the value of imports of factor n from country i in country j in the counterfactual equilibrium, $(X_{ij}^n)'$, can be simply computed as

$$(X_{ij}^n)' = \hat{x}_{ij}^n x_{ij}^n (\sum_n \hat{w}_j^n \hat{\nu}_j^n y_j^n) \text{ for all } i,\, j, \text{ and } n.$$

To sum up, if we know the factor demand system in all countries, $\{\boldsymbol{\chi}_i\}$, and have access to data on expenditure shares and factor payments, $\{x_{ij}^n\}$ and $\{y_j^n\}$, then one can compute counterfactual changes in factor trade and factor prices. Using standard arguments from consumer theory, we establish next that the knowledge of $\boldsymbol{\chi}_i$ is also sufficient for computing welfare changes in country i.

Welfare

Consider an arbitrary country i. We are interested in computing the equivalent variation, ΔW_i, associated with a shock from $(\boldsymbol{\tau}, \boldsymbol{\nu})$ to $(\boldsymbol{\tau}', \boldsymbol{\nu}')$. When expressed as a fraction of country i's initial income, this is given by

$$\Delta W_i = (e_i(\boldsymbol{\omega}_i, U_i') - y_i)/y_i. \tag{2.17}$$

where U_i' denotes the utility level of country i in the counterfactual equilibrium and $e_i(\cdot, U_i')$ denotes the expenditure function,

$$e_i(\boldsymbol{\omega}_i, U_i') \equiv \min_{\tilde{\boldsymbol{L}}_i}\{\sum_{n,j} \omega_{ji}^n \tilde{L}_{ji}^n | \bar{U}_i(\tilde{\boldsymbol{L}}_i) \geq U_i'\}.$$

By construction, ΔW_i measures the percentage change in income that the representative agent in country i would be indifferent about accepting in lieu of the counterfactual change from $(\boldsymbol{\tau}, \boldsymbol{\nu})$ to $(\boldsymbol{\tau}', \boldsymbol{\nu}')$. Note that when preference shocks occur in country i—i.e. when there is a change from τ_{ji}^n for some j and n—the expenditure function implicitly measures the amount of income necessary to reach U_i' given the original preferences, i.e. given utility $\bar{U}_i(\{L_{ji}^n/\tau_{ji}^n\})$, taking into account that after the shock, the consumer maximizes $\bar{U}_i(\{L_{ji}^n/(\tau_{ji}^n)'\})$. Since preference shocks are multiplicative, this is equivalent to a change in effective factor prices from $\boldsymbol{\omega}_i \equiv \{w_j^n \tau_{ji}^n\}$ to $\boldsymbol{\omega}_i' \equiv \{(w_j^n)'(\tau_{ji}^n)'\}$.

To compute ΔW_i, we can solve a system of Ordinary Differential Equations (ODEs), as in Hausman (1981) and Hausman and Newey (1995). Since the expenditure function $e_i(\cdot, U_i')$ is concave in the effective factor prices, it must be differentiable almost everywhere. The Envelope Theorem (e.g. Milgrom and Segal (2002), Theorem 1) therefore implies

$$de_i(\boldsymbol{\omega}, U_i')/d\omega_j^n = L_{ji}^n(\boldsymbol{\omega}, e_i(\boldsymbol{\omega}, U_i')) \text{ for all } j \text{ and } n \text{ and almost all } \boldsymbol{\omega},$$

with $\{L_{ji}^n(\boldsymbol{\omega}, e_i(\boldsymbol{\omega}, U_i'))\}$ that solves (2.9) at the effective factor prices, $\boldsymbol{\omega}$. Given our focus on expenditure shares, it is convenient to rearrange the previous expression in logs. For any selection $\{x_{ji}^n(\boldsymbol{\omega}, y)\} \in \boldsymbol{\chi}_i(\boldsymbol{\omega}, y)$, we must have

$$d\ln e_i(\boldsymbol{\omega}, U_i')/d\ln\omega_j^n = x_{ji}^n(\boldsymbol{\omega}, e_i(\boldsymbol{\omega}, U_i')) \text{ for all } j \text{ and } n \text{ and almost all } \boldsymbol{\omega}. \quad (2.18)$$

By budget balance in the counterfactual equilibrium, we also know that

$$e_i(\boldsymbol{\omega}_i', U_i') = y_i', \quad (2.19)$$

where $\boldsymbol{\omega}'_i$ is the vector of effective factor prices in the counterfactual equilibrium.

The expenditure function $e_i(\cdot, U'_i)$ must be equal to the unique solution of (2.18) satisfying (2.19). This solution can be computed given knowledge of any selection $\{x^n_{ji}(\cdot,\cdot)\} \in \boldsymbol{\chi}_i(\cdot,\cdot)$, country i's income level in the counterfactual equilibrium, $y'_i = \sum_n \hat{w}^n_i \hat{\nu}^n_i y^n_i$, and the effective factor prices in the counterfactual equilibrium, $\boldsymbol{\omega}'_i = \{\hat{w}^n_j \hat{\tau}^n_{ji} (\chi^n_{ji})^{-1}(\boldsymbol{x}_i, y_i)\}$, with $\hat{\boldsymbol{w}}$ given by (2.15) and (2.16). Once $e_i(\cdot, U'_i)$ has been solved for, we can again use the invertibility of demand to substitute for the initial effective factor prices in equation (2.17). This leads to our next proposition.

Proposition 3. *Suppose that A1 holds and the reduced utility function satisfies (2.11). Then the equivalent variation associated with a change from $(\boldsymbol{\tau}, \boldsymbol{\nu})$ to $(\boldsymbol{\tau}', \boldsymbol{\nu}')$, expressed as a fraction of country i's initial income, is*

$$\Delta W_i = (e(\{(\chi^n_{ji})^{-1}(\boldsymbol{x}_i, y_i)\}, U'_i) - y_i)/y_i, \tag{2.20}$$

where $e(\cdot, U'_i)$ is the unique solution of (2.18) and (2.19).

2.4.2 Application to Neoclassical Trade Models

Our goal now is to find structural shocks in a neoclassical trade model that are isomorphic to preference and endowment shocks in a reduced exchange model. By Propositions 1-3, counterfactual predictions about factor content of trade, factor prices, and welfare in neoclassical model can then be computed using equations (2.15)-(2.20).

Consider a neoclassical trade model in which technology can be expressed as

$$f^k_{ij}(\boldsymbol{l}^k_{ij}) \equiv \bar{f}^k_{ij}(\{l^{nk}_{ij}/\tau^n_{ij}\}), \text{ for all } i,\ j, \text{ and } k, \tag{2.21}$$

where τ^n_{ij} denotes factor-augmenting productivity shocks, that are common to all goods for a given exporter-importer pair. Since these productivity shocks are bilateral in nature, we simply refer to them as trade cost shocks from now on.[13]

[13]Formally, a change in iceberg trade costs between countries i and j corresponds to the special case in which productivity shocks are Hicks-neutral for a given exporter-importer pair, i.e., $\tau^n_{ij} = \tau_{ij}$. Note that while the productivity shocks considered in equation (2.21) may not vary across

Given equation (2.21), the reduced utility function over primary factors of production associated with the present neoclassical trade model can be written as

$$
\begin{aligned}
U_i(\boldsymbol{L}_i) \equiv \max_{\tilde{\boldsymbol{q}}_i, \tilde{\boldsymbol{l}}_i} & u_i(\tilde{\boldsymbol{q}}_i) \\
& \tilde{q}_{ji}^k \leq \bar{f}_{ji}^k(\{\tilde{l}_{ji}^{nk}/\tau_{ji}^n\}) \text{ for all } j \text{ and } k, \\
& \sum_k \tilde{l}_{ji}^{nk} \leq L_{ji}^n \text{ for all } j \text{ and } n.
\end{aligned}
$$

A simple change of variable then implies

$$
U_i(\boldsymbol{L}_i) \equiv \bar{U}_i(\{L_{ji}^n/\tau_{ji}^n\}),
$$

with

$$
\begin{aligned}
\bar{U}_i(\boldsymbol{L}_i) \equiv \max_{\tilde{\boldsymbol{q}}_i, \tilde{\boldsymbol{l}}_i} & u_i(\tilde{\boldsymbol{q}}_i) \\
& q_{ji}^k \leq \bar{f}_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) \text{ for all } j \text{ and } k, \\
& \sum_k \tilde{l}_{ji}^{nk} \leq L_{ji}^n \text{ for all } j \text{ and } n.
\end{aligned}
$$

Thus, if technology satisfies (2.21), $U_i(\cdot)$ satisfies (2.11). Not surprisingly, trade cost shocks in a neoclassical trade model are equivalent to preference shocks in the associated reduced exchange model. Since endowment shocks are identical in neoclassical trade models and reduced exchange models, we arrive at the following corollary of Propositions 1-3.

Corollary 1. *Suppose that A1 holds and that technology satisfies (2.21). Then the proportional changes in the factor content of trade, factor prices, and welfare caused by trade cost shocks and endowment shocks in a neoclassical trade models,* $\hat{\boldsymbol{\tau}} \equiv \{(\tau_{ij}^n)'/\tau_{ij}^n\}$ *and* $\hat{\boldsymbol{\nu}} \equiv \{(\nu_i^n)'/\nu_i^n\}$, *are given by (2.15)-(2.20).*

---

goods, equation (2.21) does allow for a very rich set of heterogeneous trading frictions in the initial equilibrium: $\bar{f}_{ij}^k$ may vary with both i and j for all k. Thus, some goods may be more costly to trade than others. Similarly, goods that are exported may have different factor intensity than goods that are sold domestically, as in Matsuyama (2007). Note also that the productivity shocks considered in equation (2.21) may vary across factors n. Hence, our model can accommodate economies in which only a subset of individuals get access to foreign markets.

To sum up, equations (2.15)-(2.20) provide a system of equations that can be used for counterfactual and welfare analysis. It generalizes the exact hat algebra of Dekle, Eaton, and Kortum (2008) developed in the case of Constant Elasticity of Substitution (CES) factor demands to any invertible factor demand system. Namely, given data on expenditure shares and factor payments, $\{x_{ji}^n\}$ and $\{y_i^n\}$, if one knows the factor demand system, χ_i, then one can compute counterfactual changes in factor prices, aggregate trade flows, and welfare.[14] Sections 2.5 and 2.6 discuss identification and estimation, respectively, of the factor demand system, χ_i. Before doing so, we briefly discuss some extensions of the previous results.

2.4.3 Extensions

Sector-specific trade cost shocks

Our approach emphasizes that in any neoclassical trade model, it is as if countries were directly trading factor services. As we have shown in the previous subsection, this approach is well-suited to study factor-augmenting productivity shocks, in general, and uniform changes in iceberg trade costs, in particular. While such shocks are of independent interest, they are restrictive. For instance, one may want to study trade cost shocks that only affect a subset of sectors in the economy. Here, we demonstrate how our analysis can be extended to cover such cases.

Consider the same neoclassical economy as in Section 2.4.2 with technology satisfying (2.21). For expositional purposes, consider a counterfactual shock that only affects productivity, τ_{ij}^{nk}, of one factor, n, for one country pair, i and j, in a single sector, k. To study such a counterfactual scenario, we only need to add one factor and one non-arbitrage condition to our previous analysis. Namely, instead of only having "factor n in country i," we can define "factor n from country i that is used to produce good k for country j" and "factor n from country i that is not." The price of both factors in a competitive equilibrium, of course, should be the same. Given

[14]Like Proposition 1, the above corollary holds whether or not a competitive equilibrium is unique. If there are multiple equilibria, then there is a set of proportional changes in the factor content of trade, factor prices, and welfare caused by $\hat{\tau}$ and $\hat{\nu}$, but this set remains characterized by (2.15)-(2.20).

this new set of factors, all shocks remain uniform across goods. Thus, the results of Section 2.4.2 still apply.

Of course, as trade cost shocks become more and more heterogeneous across sectors, our emphasis on the factor content of trade becomes less and less useful. In the extreme case where all goods are subject to a different shock, it is no simpler to study a reduced exchange model with $K \times N$ factors in each country than the complete neoclassical trade model with K goods and N factors. The flip-side of this observation is that, away from this extreme case, our approach is always useful in the sense that it reduces the dimensionality of what needs to be estimated, i.e., the factor demand system.

Tariffs

Historically, an important application of CGE models has been the analysis of regional trade agreements, such as NAFTA and the European Union, in which the counterfactual shocks of interest were not productivity shocks but rather changes in trade policy; see e.g. Baldwin and Venables (1995) for a survey. We now discuss how our analysis can be extended to analyze the effects of changes in ad-valorem trade taxes. For pedagogical purposes, it is useful to start from a reduced exchange model, as in Section 2.4.1, but one in which a factor n being traded between country i and country j is subject to an ad-valorem import tax or subsidy, t_{ij}^n. Once this case has been dealt with, the empirically relevant case in which tariffs vary across sectors, not factors, can be dealt with by redefining factors appropriately, as in Section 2.4.3.[15]

The key difference between the reduced equilibrium with and without trade taxes is that taxes raise revenue. This needs to be added to factor income in equations (2.15) and (2.16) when computing changes in factor prices and the factor content of trade. Formally, consider a change in trade taxes from $\boldsymbol{t} \equiv \{t_{ij}^n\}$ to $\boldsymbol{t}' \equiv \{(t_{ij}^n)'\}$. The

[15]Wilson (1980) discusses this issue in the context of the Ricardian model.

counterparts of equations (2.15) and (2.16) in this situation become

$$\{\hat{x}_{ji}^n x_{ji}^n\} \in \boldsymbol{\chi}_i(\{\hat{w}_j^n \widehat{(1+t_{ji}^n)}(\chi_{ji}^n)^{-1}(\boldsymbol{x}_i, y_i)\}, \hat{y}_i y_i) \text{ for all } i,$$
$$\sum_j \hat{x}_{ij}^n x_{ij}^n \hat{y}_j y_j = \hat{w}_i^n \hat{\nu}_i^n w_i^n \nu_i^n, \text{ for all } i \text{ and } n,$$

with total income, inclusive of tax revenues, such that

$$y_i = \sum_n y_i^n / (1 - \sum_j \sum_n t_{ji}^n x_{ji}^n / (1 + t_{ji}^n)) \text{ for all } i,$$
$$\hat{y}_i y_i = \sum_n \hat{w}_i^n \hat{\nu}_i^n y_i^n / (1 - \sum_j \sum_n (t_{ji}^n)' \hat{x}_{ji}^n x_{ji}^n / (1 + t_{ji}^n)') \text{ for all } i.$$

Equations (2.18)-(2.20) are unchanged. So, given information on tariffs, $\boldsymbol{t}$ and $\boldsymbol{t}$', changes in the factor content of trade, factor prices, and welfare can still be computed using only: (i) data on initial expenditure shares and factor payments, $\{x_{ij}^n\}$ and $\{y_j^n\}$, and (ii) an estimate of the factor demand system, $\boldsymbol{\chi}_i$, in each country.

Intermediate Goods

The neoclassical trade model of Section 2.3 rules out intermediate goods. We conclude by discussing how our theoretical analysis can be extended to environments with input-output linkages. Consider an economy in which gross output of good k produced in country i that is available in country j—either as a final good for consumers or an intermediate good for firms—is given by

$$q_{ij}^k = f_{ij}^k(\boldsymbol{l}_{ij}^k, \boldsymbol{m}_{ij}^k),$$

where $\boldsymbol{l}_{ij}^k \equiv \{l_{ij}^{nk}\}$ still denotes the vector of factor demands and $\boldsymbol{m}_{ij}^k \equiv \{m_{oij}^{gk}\}$ is the vector of input demands, with m_{oij}^{gk} being the amount of good g from the origin country o that is used as an intermediate good in country i to produce good k and deliver it to country j. In a competitive equilibrium, gross output must then be equal to the total demand by consumers and firms,

$$c_{ij}^k + \sum_{l,d} m_{ijd}^{kl} = q_{ij}^k \text{ for all } i,\ j, \text{ and } k,$$

where $\boldsymbol{c}_i \equiv \{c_{ji}^k\}$ denotes the vector of final demand in country i. All other assumptions are the same as in Section 2.3.1.

In this more general environment, we can still define a reduced utility function over primary factors of production,

$$\begin{aligned} U_i(\boldsymbol{L}_i) \equiv \max&{}_{\tilde{\boldsymbol{q}},\tilde{\boldsymbol{m}},\tilde{\boldsymbol{c}},\tilde{\boldsymbol{l}}} u_i(\tilde{\boldsymbol{c}}_i) \\ &\tilde{q}_{jd}^k \le f_{jd}^k(\tilde{\boldsymbol{l}}_{jd}^k, \tilde{\boldsymbol{m}}_{jd}^k) \text{ for all } d,\ j, \text{ and } k, \\ &\sum_{d,k} \tilde{l}_{jd}^{nk} \le L_{ji}^n \text{ for all } j \text{ and } n, \\ &\tilde{c}_{jd}^k + \sum_{g,r} \tilde{m}_{jdr}^{kg} \le \tilde{q}_{jd}^k \text{ for all } d,\ j, \text{ and } k, \end{aligned}$$

with $\tilde{\boldsymbol{q}} \equiv \{\tilde{q}_{jd}^k\}$, $\tilde{\boldsymbol{m}} \equiv \{\tilde{m}_{jdr}^{kg}\}$, $\tilde{\boldsymbol{c}} \equiv \{\tilde{c}_{jd}^k\}$, and $\tilde{\boldsymbol{l}} \equiv \{\tilde{l}_{jd}^{nk}\}$. Compared to the definition of Section 2.3.2, the control variables now include gross output, intermediate goods, final demands, and primary factors for all destination countries, d, not just country i. This reflects the potential existence of global supply chains in which factors from country j may be used to produce intermediate goods for country d, which are then used to produce final goods for country i.[16]

One can show that Proposition 1 still holds in this economy, with the factor content of trade being computed as in Johnson and Noguera (2012). The only technicality is that the proof now requires the Nonsubstitution Theorem to construct good prices in a competitive equilibrium from factor prices in a reduced equilibrium. Conditional on the new definition of the reduced utility function, Propositions 2 and 3 are unchanged. They can be applied directly to study endowment shocks in a neoclassical model. When intermediate goods are not traded or traded but their factor content is not

[16]Obviously, a solution to the previous maximization problem must always feature $c_{jd}^k = 0$ for all $d \neq i$ since country i cannot benefit from final consumption in other countries.

re-exported, as in Grossman and Rossi-Hansberg (2008a), Propositions 2 and 3 can also be applied directly to the analysis of changes in trade costs. When the factor content of intermediate goods is re-exported, as in Yi (2003), Propositions 2 and 3 can still be used, but they require the space of factors to be augmented, as in Section 2.4.3. Specifically, one needs to treat factors that are imported directly and indirectly differently since they are subject to different (vectors of) iceberg trade costs.

2.5 Identification

2.5.1 Assumptions

In order to go from the economic model of Section 2.3.1 to an econometric model that can be estimated, we need to make additional assumptions on which variables are unobservable and which ones are not as well as the origins of the exogenous shocks generating the observable variables.

Exogenous shocks. Consider a dataset generated by the model of Section 2.3.1 at different dates indexed by t. At each point in time, we assume that preferences and technology in the original neoclassical trade model satisfy

$$u_i(\boldsymbol{q}_{i,t}) = \bar{u}(\{q^k_{ji,t}/\theta_{ji}\}), \text{ for all } i, \tag{2.22}$$

$$f^k_{ij,t}(\boldsymbol{l}^k_{ij,t}) = \bar{f}^k_i(\{l^{nk}_{ij,t}/\tau^n_{ij,t}\}), \text{ for all } i,\ j, \text{ and } k. \tag{2.23}$$

Factor endowments, $\{\nu^n_{i,t}\}$, and trade costs, $\{\tau^n_{ij,t}\}$, are allowed to vary over time, but utility and production functions, $\{u_i\}$ and $\{\bar{f}^k_i\}$, are assumed to be fixed. Differences in preferences across countries take the form of exporter-importer taste shifters, $\{\theta_{ji}\}$, that are common across all goods.[17] In line with the analysis of Section 2.4, equa-

[17]For expositional purposes, we ignore time-varying preference shocks, $\theta_{ji,t}$. They could be dealt with in the exact same way as we dealt with preference shocks in the reduced exchange model of Section 2.4.1. Note also that the absence of sector-specific productivity shocks is sufficient, but not necessary. What is crucial for the analysis below is that sector-specific productivity shocks do not affect the shape of factor demand. For example, if all goods enter symmetrically in the utility function, then a weaker sufficient condition is that the distribution of productivity across sectors is stable over time, though productivity in particular sectors may go up or down at particular points

tions (2.22) and (2.23) lead to the following restriction on the heterogeneity in factor demands across countries.

A2. [Price heterogeneity] *In any country i and at any date t, there exists a vector of effective factor prices, $\boldsymbol{\omega}_{i,t} \equiv \{w_{j,t}^n \theta_{ji} \tau_{ji,t}^n\}$, such that factor demand can be expressed as $\bar{\boldsymbol{\chi}}(\boldsymbol{\omega}_{i,t}, y_{i,t})$.*

Under A2, reduced utility functions over primary factors of production in the reduced exchange model are allowed to vary across countries and over time—either because of primitive differences in preferences or technology—but this heterogeneity can be reduced to differences in effective factor prices, i.e., factor prices adjusted by the relevant preference and trade cost shocks. This implies that in order to identify the shape of factor demand around the world, we only need to identify the shape of $\bar{\boldsymbol{\chi}}$.

An obvious benefit of A2 is that it reduces the dimensionality of the demand system that we want to estimate by a factor I, equal to the number of countries in the world economy. A more subtle, but crucial benefit of A2 is that the global factor demand,$\bar{\boldsymbol{\chi}}$, can be estimated using both time series and cross-sectional variation. This will allow us to control for variations in endogenous factor prices, $w_{j,t}^n$, by including exporter-factor-year dummies when estimating $\bar{\boldsymbol{\chi}}$. Finally, note that A2 holds trivially in a gravity model, as can be seen directly from (2.14).

Observables and unobservables. For any country i and for any date t, we assume that effective factor prices, $\boldsymbol{\omega}_{i,t} \equiv \{\omega_{ji,t}^n\}$, are unobservable and normalized so that:

$$\ln \omega_{1i,t}^1 = 0, \text{ for all } i \text{ and } t, \tag{2.24}$$

$$E[\ln \omega_{ji,t}^n] = 0, \text{ for all } j \text{ and } n. \tag{2.25}$$

The first normalization amounts to expressing effective factor prices relative to factor 1 from country 1 in all markets (i, t). The second normalization is necessary to

in time. Hanson, Lind, and Muendler (2014) offer empirical evidence consistent with that weaker condition.

identify separately effective factor prices from factor-specific taste shifters.[18] The only observables are: (i) factor expenditure shares, $\boldsymbol{x}_{i,t} \equiv \{x^n_{ji,t}\}$; ($ii$) factor payments, $\boldsymbol{y}_{i,t} \equiv \{y^n_{i,t}\}$; and ($iii$) trade cost shifters, $\boldsymbol{z}_{i,t} \equiv \{z^n_{ji,t}\}$.[19] We assume that trade cost shocks in the model, $\tau^n_{ji,t}$, are related to trade cost shifters in the data through

$$\ln \tau^n_{ji,t} = \ln z^n_{ji,t} + \tilde{\varphi}^n_{ji} + \tilde{\xi}^n_{j,t} + \varepsilon^n_{ji,t},$$

where $\tilde{\varphi}^n_{ji}$ and $\tilde{\xi}^n_{j,t}$ are exporter-importer-factor and exporter-factor-year fixed-effects, respectively, and $\boldsymbol{\varepsilon}_{i,t} \equiv \{\varepsilon^n_{ji,t}\}$ are idiosyncratic shocks. In Section 2.6, we will use data on bilateral freight costs as trade cost shifters for (all) factors from a given destination. Combining the previous equation with the definition of effective factor prices, $\omega^n_{ji,t} \equiv w^n_{j,t}\theta_{ji}\tau^n_{ji,t}$, we then obtain

$$\ln\omega^n_{ji,t} = \ln z^n_{ji,t} + \varphi^n_{ji} + \xi^n_{j,t} + \varepsilon^n_{ji,t}, \text{ for all } i,\ j,\ n, \text{ and } t, \tag{2.26}$$

with $\varphi^n_{ji} \equiv \tilde{\varphi}^n_{ji} + \ln\theta_{ji}$ and $\xi^n_{j,t} \equiv \tilde{\xi}^n_{j,t} + \ln w^n_{j,t}$. The first set of fixed-effects, $\{\varphi^n_{ji}\}$, captures—among other things—any source of trading frictions between country i and j that is stable over time. This includes common proxies for trade costs like bilateral distance, whether i and j share a common language, or whether they have colonial ties; see e.g. Anderson and Van Wincoop (2003). Crucially, the second set of fixed effects, $\{\xi^n_{j,t}\}$, captures the variations in factor prices, $\{w^n_{j,t}\}$, which are the key endogenous variables in our model.

Throughout our analysis, we impose the following exogeneity restriction on the vector of idiosyncratic shocks.

[18] One can always start from $\bar{\chi}(\boldsymbol{\omega}_{i,t}, y_{i,t})$ and define $\tilde{\chi}(\tilde{\boldsymbol{\omega}}_{i,t}, y_{i,t}) \equiv \bar{\chi}(\{\tilde{\omega}^n_{ji,t}/\alpha^n_j\}, y_{i,t})$ with $\tilde{\omega}^n_{ji,t} = \alpha^n_j\omega^n_{ji,t}$ for some $\alpha^n_j > 0$. By construction, $\bar{\chi}$ and $\tilde{\chi}$ must generate the exact same observables. However, $\boldsymbol{\omega}_{i,t}$ and $\tilde{\boldsymbol{\omega}}_{i,t}$ cannot both satisfy (2.25).

[19] In principle, data on factor expenditure shares, $\boldsymbol{x}_{i,t} \equiv \{x^n_{ji,t}\}$, and factor payments, $\boldsymbol{y}_{i,t} \equiv \{y^n_{i,t}\}$, can be obtained from sources such as the World Input-Output Database. As already discussed in the Introduction, a practical limitation of such datasets is that they implicitly assume that factor intensity is constant across destinations within the same industry. For the empirical application of Section 2.6, such considerations will be irrelevant since we will assume the existence of a composite factor in each country. Note also that for the purposes of identifying the shape of factor demand, we will only need information on total income, $y_{i,t} = \sum_n y^n_{i,t}$, in each country. Data on factor payments, $\boldsymbol{y}_{i,t} \equiv \{y^n_{i,t}\}$, are only necessary for counterfactual analysis, as shown in Section 2.4.

A3. [Exogeneity] $E[\varepsilon_{i,t}|\ln \boldsymbol{z}_{i,t}, \boldsymbol{d}_{i,t}] = 0$, *where* $\boldsymbol{d}_{i,t}$ *is a full vector of importer-exporter-factor and exporter-factor-year dummies, with* i *as the importer and* t *as the year for all dummies.*

Because of equation (2.26), A3 is stronger than assuming that trade cost shifters, $\boldsymbol{z}_{i,t}$, can be used as instruments for effective factor prices, $\boldsymbol{\omega}_{i,t}$, after controlling for all factors that are either exporter-importer-factor or exporter-factor-year specific. If we think of equation (2.26) as a first-stage, it implies that reduced-form and IV estimates should coincide. Hence, we can infer the impact on factor demand of effective factor prices, which are not observable, by tracing out the impact of trade cost shifters, which are observable. This is the same strategy used for the estimation of (constant) trade elasticities in the gravity literature; see Head and Mayer (2013).[20]

Following Newey and Powell (2003a), we conclude by imposing the following completeness condition.

A4. [Completeness] *For any* $g(\boldsymbol{x}_{i,t}, \boldsymbol{d}_{i,t}, y_{i,t})$ *with finite expectation,*

$$E[g(\boldsymbol{x}_{i,t}, \boldsymbol{d}_{i,t}, y_{i,t})|\ln \boldsymbol{z}_{i,t}, \boldsymbol{d}_{i,t}] = 0 \implies g(\boldsymbol{x}_{i,t}, \boldsymbol{d}_{i,t}, y_{i,t}) = 0.$$

A4 is the equivalent of a rank condition in the estimation of parametric models.[21]

[20] A common finding in the international macro literature is that exporters' costs shocks tend to be incompletely passed through into consumer prices; see e.g. Burstein and Gopinath (2013). This observation does not by itself invalidate the previous strategy. Within the context of a neoclassical model, such findings can be rationalized by assuming that foreign goods need to be distributed, which requires local factors of production, as in Burstein, Neves, and Rebelo (2003). In such a model, there is incomplete pass-through into consumer prices, as observed in the data, yet complete pass-through into effective factor prices, as assumed in equation (2.26).

[21] Going from a finite to an infinite dimensional space of parameters leads to non-trivial issues. Newey (2013a) notes that "In fully nonparametric models (that are infinite dimensional), completeness is not testable, as pointed out by Canay, Santos, and Shaikh (2013). In these models the reduced form is like an infinite dimensional matrix with eigenvalues that have a limit point at zero. Nonidentification occurs when at least one of the eigenvalues equals zero. The problem with testing this hypothesis is that one cannot distinguish empirically a model with a zero eigenvalue from one where the eigenvalues have a limit point of zero. However, completeness is generic, in the sense that it holds for "most" if it holds for one [...]. This is like the discrete, finite support case where most matrices have full column rank if the order condition is satisfied." We have little to add to this discussion.

2.5.2 Identifying Factor Prices and Factor Demand

We are now ready to establish that factor prices and factor demand are identified. The argument follows the same steps as in Berry and Haile (2014). Because our demand system is invertible, we can express each effective factor price as a function of the vector market share plus some error term. Once the estimating equations have been transformed in this way, the completeness condition of Newey and Powell (2003a) provides non-parametric identification.

By A1, we can invert our factor demand system to express effective factor prices, $\boldsymbol{\omega}_{i,t}$ faced by country i at date t as a function of expenditure shares, $\boldsymbol{x}_{i,t}$, and total income, $y_{i,t}$,

$$\omega_{ji,t}^n = (\chi_{ji}^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}), \tag{2.27}$$

with the level of effective factor prices in country i and year t pinned down by (2.24). Taking logs and using equation (2.26), we then have

$$\varepsilon_{ji,t}^n = \ln(\chi_{ji}^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}) - \ln z_{ji,t}^n - \varphi_{ji}^n - \xi_{j,t}^n.$$

By A2, the inverse demand is the same for all importer i and period t,

$$\varepsilon_{ji,t}^n = \ln(\bar{\chi}_j^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}) - \ln z_{ji,t}^n - \varphi_{ji}^n - \xi_{j,t}^n,$$

where $(\bar{\chi}_j^n)^{-1}(\cdot)$ is the inverse demand for factor n from country j. Combining this expression with A3, we obtain the following moment condition

$$E[\ln(\bar{\chi}_j^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}) - \varphi_{ji}^n - \xi_{jt}^n | \ln \boldsymbol{z}_{i,t}, \boldsymbol{d}_{i,t}] = \ln z_{ji,t}^n. \tag{2.28}$$

By A4, there is at most one function g_j^n that satisfies $E[g_j^n(\boldsymbol{x}_{i,t}, y_{i,t}, \boldsymbol{d}_{i,t}) | \ln \boldsymbol{z}_{i,t}, \boldsymbol{d}_{i,t}] = \ln z_{ji,t}^n$.[22] Thus if two inverse demand functions, $(\bar{\chi}_j^n)^{-1}$ and $(\tilde{\chi}_j^n)^{-1}$, satisfy (2.28) for

[22] To see this, suppose that $g_j^n(\boldsymbol{x}_{i,t}, y_{i,t}, \boldsymbol{d}_{i,t})$ and $\tilde{g}_j^n(\boldsymbol{x}_{i,t}, y_{i,t}, \boldsymbol{d}_{i,t})$ satisfy equation (2.28). Then they must also satisfy $E[g_j^n(\boldsymbol{x}_{i,t}, y_{i,t}, \boldsymbol{d}_{i,t}) - \tilde{g}_j^n(\boldsymbol{x}_{i,t}, y_{i,t}, \boldsymbol{d}_{i,t}) | \ln \boldsymbol{z}_{i,t}, \boldsymbol{d}_{i,t}] = 0$, which requires $g_j^n = \tilde{g}_j^n$ by A4.

some $(\varphi_{ji}^n, \xi_{jt}^n)$ and $(\tilde{\varphi}_{ji}^n, \tilde{\xi}_{jt}^n)$, then we must have

$$\ln(\bar{\chi}_j^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}) = g_j^n(\boldsymbol{x}_{i,t}, y_{i,t}, \boldsymbol{d}_{i,t}) + \varphi_{ji}^n + \xi_{jt}^n,$$
$$\ln(\tilde{\chi}_j^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}) = g_j^n(\boldsymbol{x}_{i,t}, y_{i,t}, \boldsymbol{d}_{i,t}) + \tilde{\varphi}_{ji}^n + \tilde{\xi}_{jt}^n.$$

Taking the difference between these two equations, we obtain

$$\ln(\bar{\chi}_j^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}) - \ln(\tilde{\chi}_j^n)^{-1}(\boldsymbol{x}_{i,t}, y_{i,t}) = \varphi_{ji}^n + \xi_{jt}^n - \tilde{\varphi}_{ji}^n - \tilde{\xi}_{jt}^n.$$

Holding $\boldsymbol{x}_{i,t}$ and $y_{i,t}$ fixed, the left-hand side does not vary with i or t. So, $\varphi_{ji}^n + \xi_{jt}^n - \tilde{\varphi}_{ji}^n - \tilde{\xi}_{jt}^n$ cannot vary with i or t either. This establishes that $\ln(\bar{\chi}_j^n)^{-1}$ is identified up to some constant, which equation (2.25) and (2.27) pin down for all j and n. Finally, note that once the inverse factor demand is known, then both factor demand and effective factor prices are known as well, with prices being uniquely pinned down by equation (2.27).

We summarize the previous discussion in the next proposition.

Proposition 4. *Suppose that A1-A4 hold. Then effective factor prices and factor demand are identified, up to the two normalizations (2.24) and (2.25).*

2.5.3 Ricardian Example

The invertibility of demand plays a key role throughout our analysis. We use it to conduct counterfactual and welfare analysis in Section 2.4 and we use it again to establish Proposition 4. We now provide sufficient conditions on the primitives of a neoclassical trade model such that A1 holds. We also show that under the same conditions, a competitive equilibrium is unique. Hence, counterfactual changes in factor prices and welfare are also nonparametrically identified in this environment. We will come back to the same environment for our empirical application in Sections 2.6 and 2.7.

Consider an economy in which utility and productions functions satisfy

$$u_i(\boldsymbol{q}_i) = \bar{u}(\{\sum_j q_{ji}^k\}), \text{ for all } i, \tag{2.29}$$

$$f_{ij}^k(\boldsymbol{l}_{ij}^k) = \alpha_i^k \bar{f}_i(\boldsymbol{l}_{ij}^k)/\tau_{ij}, \text{ for all } i, j, \text{ and } k, \tag{2.30}$$

where $\bar{u}$ is a homothetic utility function that satisfies standard Inada conditions; α_{ij}^k is total factor productivity in country i and sector k when selling to country j; $\bar{f}_i$ is a production function, common to all sectors and destinations; and τ_{ij} is a bilateral iceberg trade cost. Given equation (2.29), the Inada conditions are imposed to rule out zero expenditure shares on all goods.[23] The crucial restriction is imposed in equation (2.30). It states that all goods from country i use factors with the same intensity. Hence, everything is as if there was only one factor per country with price $c_i \equiv \min_{\tilde{l}}\{\sum w_i^n \tilde{l}_i^n | \bar{f}_i(\tilde{\boldsymbol{l}}) = 1\}$ and endowment $\bar{f}_i(\boldsymbol{\nu}_i)$.

In light of the previous discussion, we refer to an economy that satisfies (2.29) and (2.30) as a Ricardian economy. In such an environment, homotheticity and no differences in factor intensity imply that we can write the demand for factors in country i as $\bar{\boldsymbol{\chi}}(\boldsymbol{\omega}_i)$, with $\boldsymbol{\omega}_i \equiv \{\tau_{ji}c_j\}$ the vector of effective prices for the composite factors.

As discussed in Berry, Gandhi, and Haile (2013), a sufficient condition for a demand function to be invertible over its support is that it satisfies the connected substitute property.[24] This property has a long tradition in general equilibrium the-

[23]By itself, the assumption that goods from different exporting countries are perfect substitutes, as described in equation (2.29), is without loss of generality. To see this, note that by assuming that each good k can only be produced in one country, the present model still nests the Armington model. We only impose equation (2.29) to weaken the Inada conditions. Namely, we require all countries to consume all goods, not all goods from all origins.

[24]If we were able to observe the quantities of factor services demanded by each country directly, rather than factor expenditure shares, invertibility would be a straightforward issue in the context of this paper. From Proposition 1, we know that there must be a representative agent whose factor demand solves (2.9). Whenever the reduced utility function is differentiable at the optimum, the first-order conditions of the utility maximization problem (2.9) immediately imply that factor prices are determined (up to a normalization) by the gradient of the reduced utility function, evaluated at the optimal quantities of factor demanded. The case of Cobb-Douglas utility is an extreme example that shows that the previous argument does not carry over to expenditure shares. In that case, there is is uniqueness of prices conditional on quantities demanded, but not conditional on expenditure shares.

ory where it is used to establish the uniqueness of competitive equilibrium prices, through the injectivity of the excess demand function; see Arrow and Hahn (1971), p. 227. For the purposes of this paper, we need a slightly more general version of this property that applies to demand correspondences, not just functions. We focus on the following generalization adapted from Howitt (1980).

Definition 3 [Connected Substitutes]. *A correspondence* $\bar{\boldsymbol{\chi}} : \mathbb{R}^m_{++} \to P(\mathbb{R}^m_{++})$ *satisfies the connected substitute property if for any* $\boldsymbol{\omega}$ *and* $\boldsymbol{\omega}' \in \mathbb{R}^m_{++}$, *any* $\boldsymbol{x} \in \bar{\boldsymbol{\chi}}(\boldsymbol{\omega})$, *any* $\boldsymbol{x}' \in \bar{\boldsymbol{\chi}}(\boldsymbol{\omega}')$, *and any non-trivial partition* $\{M_1, M_2\}$ *of* $M \equiv \{1, ..., m\}$, $\omega'_j > \omega_j$ *for all* $j \in M_1$ *and* $\omega'_j = \omega_j$ *for all* $j \in M_2$ *imply* $\sum_{j \in M_2} x'_j > \sum_{j \in M_2} x_j$.

Our first lemma provides sufficient conditions under which the factor demand system of a Ricardian economy is invertible over its support.

Lemma 1. *Consider a Ricardian economy. If good expenditure shares satisfy the connected substitute property, then for any vector of factor expenditure shares,* $\boldsymbol{x} > 0$, *there is at most one vector (up to a normalization) of effective factor prices,* $\boldsymbol{\omega}$, *such that* $\boldsymbol{x} \in \bar{\chi}(\boldsymbol{\omega})$.

The formal proof can be found in Appendix 2A.1.2. The general strategy is similar to the one used by Scarf and Wilson (2005) to establish the uniqueness of competitive equilibria in a Ricardian model. The key idea is to show that if expenditure shares on goods satisfy the connected property, then the expenditure shares on factors must satisfy the same property. At that point, the invertibility of the factor demand system follows from standard arguments; see e.g. Proposition 17.F.3 in Mas-Colell, Whinston, and Green (1995). The only minor technicality is that the demand function may be a correspondence, which Definition 3 is designed to address.

In light of the above discussion, it should not be surprising that the same sufficient conditions lead to the uniqueness of the competitive equilibrium.

Lemma 2. *Consider a Ricardian economy. If good expenditure shares satisfy the connected substitute property, then the vector of equilibrium factor prices,* $(c_1, ..., c_I)$, *is unique (up to a normalization).*

Let us take stock. Proposition 4 and Lemma 1 imply that factor demand is nonparametrically identified in a Ricardian economy if A2-A4 hold and good expenditure shares satisfy the connected substitute property. Since all the assumptions of Section 2.4 are satisfied, Proposition 2, Proposition 3, and Lemma 2 further imply that proportional changes in factor prices and welfare are uniquely determined given data on initial expenditure shares and factor payments, $\{x_{ij}^n\}$ and $\{y_j^n\}$, and an estimate of factor demand, $\bar{\chi}$.[25] This leads to our final observation.

Corollary 2. *Consider a Ricardian economy. If A2-A4 hold and good expenditure shares satisfy the connected substitute property, then proportional changes in factor prices and welfare caused by trade cost shocks and endowment shocks are nonparametrically identified.*

2.6 Estimation

The above results highlight two important features of neoclassical trade models. First, counterfactual changes in trade costs and factor endowments can be studied with only the knowledge of a reduced factor demand system. Second, this reduced demand system can be nonparametrically identified from standard data sources on international trade in goods and standard exclusion restrictions. Armed with these theoretical results we now turn to a strategy for estimating the reduced demand system, in practice.

2.6.1 From Asymptopia to Mixed CES

Nonparametric identification results, like those presented in Section 2.5, are asymptotic in nature. They answer the question of whether one could point identify each of the potentially infinite-dimensional parameters of a model with a dataset whose sample size tends to infinity—formally, whether there exists a unique mapping from population data to model parameters. As noted by Chiappori and Ekeland (2009),

[25]Proportional changes in the factor content of trade are also unique if factor demand, $\bar{\chi}$, is single-valued at the initial and counterfactual equilibria.

such results are useful because they can help select the most adequate moment conditions; that is, the source of variation in the data directly related to the economic relation of interest.

Of course, datasets in the real world often feature a small number of observations and little exogenous variation. So estimation must inevitably proceed parametrically. Our goal here is to do so in a flexible manner, drawing on recent advances in the area of applied demand estimation; see e.g. Nevo (2011). In the spirit of dimensionality-reduction, we start by making three assumptions. Like in Section 2.5.3, we assume that: (i) preferences are homothetic, so that we can ignore the effect of income on expenditure shares; (ii) all goods have the same factor intensity in each country, so that we can focus on a single composite factor per country; and (iii) cross-country differences in factor demand can be reduced to differences in effective factor prices, so that we can focus on estimating a unique global factor demand. All three assumptions are restrictive, but standard in the existing gravity literature.[26]

Since there is one composite factor in each country, we drop superscripts n from now on. Hence, $\omega_{ji,t}$ stands for the effective price of the composite factor from country j in country i in year t, with $\boldsymbol{\omega}_{i,t} \equiv \{\omega_{ji,t}\}$ being the associated vector of effective prices. Taking inspiration from Berry (1994) and Berry, Levinsohn, and Pakes (1995), we posit that the expenditure share that country i devotes to the factor from country j in year t can be expressed as

$$\bar{\chi}_j(\boldsymbol{\omega}_{i,t}) = \int \frac{(\kappa_j)^{\sigma_\alpha \alpha} (\omega_{ji,t})^{-(\bar{\epsilon} \cdot \epsilon^{\sigma_\epsilon})}}{\sum_{l=1}^{N} (\kappa_l)^{\sigma_\alpha \alpha} (\omega_{li,t})^{-(\bar{\epsilon} \cdot \epsilon^{\sigma_\epsilon})}} dF(\alpha, \epsilon) \tag{2.31}$$

where $\boldsymbol{\kappa} \equiv \{\kappa_j\}$ is a vector of observable country characteristics and $\{\bar{\epsilon}, \sigma_\alpha, \sigma_\epsilon\}$ are structural parameters. The random draws (α, ϵ) can be interpreted as unobserved het-

[26]Fajgelbaum and Khandelwal (2014) is a recent exception that introduces non-homothetic preferences to study how gains from trade vary across income groups. As discussed below, our dataset only includes two importing countries: the United States and Australia. So, there is very little variation that we can use to estimate non-homotheticies. Similarly, introducing differences in factor intensity across sectors would then require estimates of the extent to which multiple factors are substitutable for one another within each country. While in principle this can be achieved with supply-side shifters of relative factor prices, finding such shifters in practice has proven difficult; see, e.g., Oberfield and Raval (2014) for a recent discussion of the capital-labor case.

erogeneity across goods in the elasticities with respect to effective factor prices, $\omega_{ji,t}$, and exporter characteristic, κ_j. We come back to this point below when discussing the relationship between mixed and nested CES.

In our baseline analysis, we assume that κ_j is the per-capita GDP of country j relative to the per-capita GDP of the United States ($j = 1$) in the pre-sample period.[27] We also assume that the joint distribution $F(\alpha, \epsilon)$ is such that α and $\ln \epsilon$ have a joint standard normal distribution with an identity covariance matrix.[28] As a function of effective factor prices, the demand system is completely characterized by three structural parameters: $\bar{\epsilon}$, σ_α and σ_ϵ.

This particular functional form is attractive for two reasons. First, it nests the case of CES demand. That is, in the special case of $\sigma_\alpha = \sigma_\epsilon = 0$, we recover a standard gravity model with trade elasticity $\bar{\epsilon}$, as in Eaton and Kortum (2002) or Anderson and Van Wincoop (2003). When $\sigma_\alpha \neq 0$ or $\sigma_\epsilon \neq 0$, the demand system in equation (2.31) becomes a random coefficients version of CES demand, in the same way that the mixed logit demand system in Berry, Levinsohn, and Pakes (1995) is a random coefficients version of logit demand. For this reason, we refer to our demand system as "mixed CES."

Second, the demand system in equation (2.31) captures flexibly and parsimoniously a number of natural features of demand substitution patterns through the structural parameters σ_α and σ_ϵ. To see this, define the share of the factor from country j in expenditures of country i conditional on (α, ϵ):

$$x_{ji,t}(\alpha,\epsilon) \equiv \frac{(\kappa_j)^{\sigma_\alpha \alpha}(\omega_{ji,t})^{-(\bar{\epsilon}\cdot\epsilon^{\sigma_\epsilon})}}{\sum_{l=1}^{N}(\kappa_l)^{\sigma_\alpha \alpha}(\omega_{li,t})^{-(\bar{\epsilon}\cdot\epsilon^{\sigma_\epsilon})}}. \tag{2.32}$$

[27]More generally, one could incorporate a multivariate set of time-varying characteristics by setting $\kappa_{j,t} \equiv \boldsymbol{\gamma} \cdot \boldsymbol{u}_{j,t}$ where $\boldsymbol{u}_{j,t}$ is a vector of characteristics for exporter j at year t and $\boldsymbol{\gamma}$ is the parameter vector that intermediates the effect of these characteristics on market shares. An alternative modeling strategy would be to organize countries into groups, based on some observed characteristic, and then estimate a nested CES system. This is much like the nested logit approach in Goldberg (1995).

[28]We incorporate the heterogeneity in ϵ with a positive multiplicative shifter to guarantee no sign variation in the trade elasticity. In other words, the sign of the trade elasticity is entirely determined by $\bar{\epsilon}$ but its magnitude is affected the multiplicative shifter, $\epsilon^{\sigma_\epsilon}$, whose distribution is log-normal with mean zero and variance σ_ϵ.

Now take three exporter countries j, l and r competing in the same importing market i in year t. Consider how the demand for the factor from country j relative to the factor from the reference country r depends on the effective price of factor from country l relative to that of country r. This elasticity of relative demand shares to relative prices is given by

$$\frac{\partial \ln\left(\frac{\bar{\chi}_j(\omega_{i,t})}{\bar{\chi}_r(\omega_{i,t})}\right)}{\partial \ln\left(\frac{\omega_{li,t}}{\omega_{rit}}\right)} = \int (\bar{\epsilon}\cdot\epsilon^{\sigma_\epsilon})\left(\frac{x_{ji,t}(\alpha,\epsilon)}{\bar{\chi}_j} - \frac{x_{ri,t}(\alpha,\epsilon)}{\bar{\chi}_r}\right) x_{li,t}(\alpha,\epsilon)dF(\alpha,\epsilon) \tag{2.33}$$

This expression highlights key features of the demand system in equation (2.31). As expected, setting $\sigma_\alpha = \sigma_\epsilon = 0$ recovers the well-know property of independence of irrelevant alternatives (IIA) embedded in the CES demand system: the cross-price elasticity is zero.[29] Departures from this special case yield richer patterns of substitution. The cross-price elasticity is relatively larger when $x_{ji,t}(\alpha,\epsilon)$ and $x_{li,t}(\alpha,\epsilon)$ co-move more than $x_{ri,t}(\alpha,\epsilon)$ and $x_{li,t}(\alpha,\epsilon)$ in the (α,ϵ) space. From equation (2.32), we can see that such a pattern is generated by two channels. Whenever $\sigma_\alpha \neq 0$ and $\sigma_\epsilon = 0$, this is the case if countries j and l are more similar in terms of their characteristics, κ, than countries r and l are (i.e., $|\kappa_j - \kappa_l| < |\kappa_r - \kappa_l|$). Alternatively, whenever $\sigma_\alpha = 0$ and $\sigma_\epsilon \neq 0$, this pattern occurs if countries j and l are more similar in terms of their effective factor price than countries r and l are—this is then intrinsically related to market shares (i.e., $|\bar{\chi}_j - \bar{\chi}_l| < |\bar{\chi}_r - \bar{\chi}_l|$).

One particular set of micro-foundations that would lead to the factor demand system in equation (2.31) is that stemming from: (*i*) a Cobb-Douglas utility with equal weights over a continuum of sectors, with a lower-level CES nest over a continuum of varieties in each sector and (*ii*) country-and-sector-specific Fréchet distributions of productivity across varieties. Under this interpretation, each sector is fully characterized by its corresponding pair (α,ϵ) with $F(\alpha,\epsilon)$ representing the distribution of sector attributes. In this sense, the factor demand system in equation (2.31) is closely

[29] This follows immediately form the observation that $x_{ji,t}(\alpha,\epsilon) = \bar{\chi}_j$ for all j if $\sigma_\alpha = \sigma_\epsilon = 0$. Also, it is straightforward to verify that, in this case, the own-price elasticity is constant and equal to $-\bar{\epsilon}$.

related to the nested CES demand implied by standard multi-sector models in the field; see e.g. Costinot, Donaldson, and Komunjer (2012a) and Caliendo and Parro (2015).

The crucial distinction here concerns the source of variation used for estimation. The results in Section 2.5 demonstrate that the aggregate factor demand system—which, as we have argued, is all that is required to study the counterfactual scenarios we consider here—is nonparametrically identified from aggregate data on factor spending shares. This is the variation that we will use next. Multi-sector level models, in contrast, are estimated using within-sector variation. And while sector-level factor demand relations are identified with sector-level data, the aggregate factor demand function, along with its essential aggregate cross-price elasticities, is not.[30]

Summarizing the above discussion, the "mixed CES" demand in equation (2.31) not only nests commonly used functional forms in the literature but also captures in a parsimonious manner the natural feature that factors similar in the κ-space are closer substitutes.[31] Given the essential role played by these cross-price elasticities of substitution in many counterfactual scenarios of interest, we consider of paramount importance the ability of an estimator to let the data speak directly to these phenomena.

[30]To see why this distinction may matter in practice, suppose that the true factor demand system is CES. In that case, the researcher using sector-level data and positing a nested-CES utility function with an upper-level Cobb-Douglas aggregator would uncover the true lower-level elasticity of substitution, but would wrongly assume that the upper-level elasticity is equal to one. In contrast, the researcher assuming mixed CES would rightly conclude that factor demand is CES. Of course, one could relax the assumption that the aggregator is Cobb-Douglas and attempt to estimate it as well; see e.g. Costinot, Donaldson, and Smith (2015). But at that point, given the dimensionality of the demand system across goods that needs to be estimated, it is not clear what the benefit is compared to estimating the factor demand system directly.

[31]The translog demand system—as used in the Armington context by Novy (2013)—is an important exception not covered by the demand system in (2.31). One way to nest both CES and translog would be to use the CES-Translog demand system introduced by Pollak, Sickles, and Wales (1984). While it is attractive to consider a demand system that nests both CES and translog, the main difficulty with using such a system is designing moment conditions that directly relate to the non-linear parameters of this extended CES-Translog system. One advantage of the "mixed CES" system is the clear connection between parameters and the structure of cross-price elasticities. As discussed below, this provides guidance for the choice of moment conditions.

2.6.2 Estimation Procedure

We now turn to the estimation of the structural parameters $\{\bar{\epsilon}, \sigma_\alpha, \sigma_\epsilon\}$ in equation (2.31). Building on the identification result of Section 2.5, the estimator is based on the existence of an observed and exogenous component of effective factor prices. Later, we take this cost shifter, $z_{ji,t}$, to be the reported freight charges between trading partners.

In order to use the estimation procedure developed by Berry, Levinsohn, and Pakes (1995) in the mixed logit case, it is convenient to focus on the following log-transformation of effective factor prices, $\delta_{ji,t} \equiv -\bar{\epsilon}\ln(\omega_{ji,t}/\omega_{1i,t})$, where $\omega_{1i,t}$ is the effective price of U.S. factor in country i at year t. Expressed in terms of $\boldsymbol{\delta}_{i,t} \equiv \{\delta_{ji,t}\}$, the demand system in equation (2.31) becomes

$$\bar{\chi}_j(\boldsymbol{\delta}_{i,t}|\boldsymbol{\theta}_2) = \int \frac{\exp(\alpha\sigma_\alpha \ln \kappa_j + \epsilon^{\sigma_\epsilon}\delta_{ji,t})}{1 + \sum_{l=2}^{N} \exp(\alpha\sigma_\alpha \ln \kappa_l + \epsilon^{\sigma_\epsilon}\delta_{li,t})} dF(\alpha, \epsilon), \tag{2.34}$$

where $\boldsymbol{\theta}_2 \equiv (\sigma_\alpha, \sigma_\epsilon)$ is the sub-vector of "non-linear" parameters of the model, by which we mean those that will enter non-linearly in the estimation procedure below.

Conditional on the vector $\boldsymbol{\theta}_2$, Berry, Levinsohn, and Pakes (1995) establish the existence of a unique vector $\boldsymbol{\delta}_{i,t}$ that rationalizes expenditure shares in country i and year t, i.e.,

$$\bar{\chi}_j(\boldsymbol{\delta}_{i,t}|\boldsymbol{\theta}_2) = x_{ji,t}, \text{ for all } j.$$

In line with the notation of the previous sections, let $\bar{\boldsymbol{\chi}}^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2) \equiv \{\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2)\}$ denote the solution of this system. By definition, $\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2) = -\bar{\epsilon}\ln(\omega_{ji,t}/\omega_{1i,t})$. Thus, we can use equation (2.26) to write

$$\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2) = -\bar{\epsilon}\ln(z_{ji,t}/z_{1i,t}) + \phi_{ji} + \varsigma_{jt} + e_{ji,t},$$

with $\phi_{ji} \equiv -\bar{\epsilon}(\varphi_{ji} - \varphi_{1i})$, $\varsigma_{jt} \equiv -\bar{\epsilon}(\xi_{jt} - \xi_{1t})$, and $e_{ji,t} \equiv -\bar{\epsilon}(\varepsilon_{ji,t} - \varepsilon_{1i,t})$, or more compactly,

$$\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2) = \boldsymbol{Z}_{ji,t}^1 \cdot \boldsymbol{\theta}_1 + e_{ji,t},$$

where $\boldsymbol{\theta}_1 \equiv (-\bar{\epsilon}, \{\phi_{ji}\}, \{\varsigma_{jt}\})$ denotes the sub-vector of "linear" parameters of the model and $\boldsymbol{Z}^1_{ji,t} \equiv [\ln(z_{ji,t}/z_{1i,t}), \boldsymbol{d}_{i,t}]$ with $\boldsymbol{d}_{i,t}$ denoting a full vector of exporter-importer dummies and exporter-year dummies, with i as the importer and t as the year for all dummies.

Given a vector of instruments $\boldsymbol{Z}_{ji,t}$ that is mean-independent from the structural error term, $E[e_{ji,t}|\boldsymbol{Z}_{ji,t}] = 0$, one can then obtain a consistent GMM estimator of $\boldsymbol{\theta} \equiv [\boldsymbol{\theta}_1 \mid \boldsymbol{\theta}_2]$ by constructing the structural error term $e_{ji,t}(\boldsymbol{\theta}) \equiv \bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2) - \boldsymbol{Z}^1_{ji,t} \cdot \boldsymbol{\theta}_1$ and solving for

$$\hat{\boldsymbol{\theta}} = \arg\min_{\theta} \mathbf{e}(\boldsymbol{\theta})'\mathbf{Z}\Phi\mathbf{Z}'\mathbf{e}(\boldsymbol{\theta}), \tag{2.35}$$

where Φ is a matrix of moment weights. The details of the estimation procedure (as well as our procedure for computing standard errors for $\hat{\boldsymbol{\theta}}_2$) can be found in Appendix 2A.2.

To build instruments for the estimation of $\boldsymbol{\theta}$, we rely on the exogeneity restrictions described in Assumption A3. First, the structural error $e_{ji,t}$ is uncorrelated with the exporter's own freight cost $\ln z_{ji,t}$ and the dummy vector $\boldsymbol{d}_{i,t}$, making $\boldsymbol{Z}^1_{ji,t}$ a natural vector of instruments. This is the usual set of regressors included in the estimation of CES demand. Second, Assumption A3 also entails that $e_{ji,t}$ is uncorrelated with the freight cost of all other competitors in the market, $\{\ln z_{li,t}\}_{l \neq j}$. Following the intuition for the IIA violation implied by $(\sigma_\alpha, \sigma_\epsilon)$, we propose additional instruments for exporter j that are based on the interaction between freight cost of competitors, $\ln z_{li,t}$, and their per-capita GDP difference, $|\kappa_j - \kappa_l|$. Specifically, define the instrument vector $\mathbf{Z}^2_{ji,t} \equiv \{|\kappa_j - \kappa_l| \ln z_{li,t}\}_{l \neq j}$. Intuitively, this choice of instruments is designed to explore the extent to which distance in the characteristic space, $|\kappa_j - \kappa_l|$, affects cross-price elasticities. The final instrument vector combines these two components: $\mathbf{Z}_{ji,t} \equiv [\mathbf{Z}^1_{ji,t} \mid \mathbf{Z}^2_{ji,t}]$.

2.6.3 Data

As described above, our estimation procedure draws on four types of data: (i) data on the total value of bilateral trade in goods, which can then be converted into expenditure shares, denoted by $x_{ji,t}$; (*ii*) data on bilateral freight costs, denoted by $z_{ji,t}$; (*iii*) data on total income by country, denoted by $y_{j,t}$; and (*iv*) data on per-capita GDP, denoted by κ_j.

We obtain data on $x_{ji,t}$ and $y_{j,t}$ from the World Input-Output Database for all years between 1995 and 2011. Following Shapiro (2012), data on $z_{ji,t}$ are available from the publicly available import data for two importers i, Australia and the United States, in all years t from 1990 to 2010.[32] To avoid the possibility of zero trade flows, we focus on the 36 largest exporters to Australia and the United States, and aggregate all other countries up to a single "Rest-of-the-World" unit. In the estimation of $\boldsymbol{\theta}$, we use all years with available information on trade flows and freight costs, 1995-2010. Finally, we obtain the information on per-capita GDP necessary to construct κ_j from the Penn World Table, version 8.0.[33] The list of exporters along with their per-capita GDP values is presented on Table A1 in Appendix 2A.3

2.6.4 Estimation Results

Reduced-Form Evidence

Before turning to our estimates of the structural parameters, we begin with a simpler approach that builds directly on the standard gravity model. Our goal is twofold. First, we illustrate that the deviations from IIA motivated in Section 2.6.1 are a systematic feature of the data. Second, we document that these deviations are directly related to the similarity of competitors in terms of per-capita GDP. To this end, we estimate the following equation:

$$\ln(x_{ji,t}) = \beta \ln z_{ji,t} + \sum_{l \neq j} \gamma_l (|\kappa_j - \kappa_l| \ln z_{li,t}) + \phi_{ji} + \zeta_{jt} + \nu_{it} + \varepsilon_{ji,t}. \tag{2.36}$$

[32]We are grateful to Joe Shapiro for making these data easily accessible to us. For each exporter and year, we compute the freight cost by dividing reported values of total exports CIF by total exports FOB. For domestic sales, we input a freight cost of zero — this is equivalent to assuming a constant (over time) transport cost of domestic sales in the presence of exporter-importer fixed effects.

[33]For each exporter, we compute per-capita GDP by dividing the expenditure-side real GDP at current PPP (USD 2005) by the total population. We then construct κ_j as the average per-capita GDP between 1992 and 1995.

In this specification, $x_{ji,t}$ is the share of country j exports in expenditures of country i at year t and $z_{ji,t}$ is the bilateral freight cost from country j to country i at year t. The terms ϕ_{ji}, ζ_{jt} and ν_{it} represent exporter-importer, exporter-year and importer-year fixed-effects, respectively.

The IIA property implies that competitors' costs affect the spending share of exporter j solely through the importer price index, being fully absorbed by the importer-year fixed effect. In specification (2.36), the IIA property is equivalent to $\gamma_l = 0$ for all l. Alternatively, IIA is violated if the demand for the factor from country j depends also on the price of the factor from country l conditional on the importer-year fixed effect; that is, $\gamma_l \neq 0$ for some exporter l. The interaction between $\ln z_{li,t}$ and $|\kappa_j - \kappa_l|$ relate this third country effect to the proximity of competitors in terms of per-capita GDP.

Table 2.1 reports estimates of various versions of equation (2.36). Column (1) begins by restricting attention to the standard CES case in which $\gamma_l = 0$ for all l. We obtain an estimate of -6.1 for the trade elasticity in line with a vast literature that has estimated such a specification; see e.g. Head and Mayer (2013). Column (2) then includes the interaction terms to estimate the set of coefficients γ_l. Because there are 37 such coefficients and we are only interested in testing whether at least one of them is

Table 2.1: Reduced-Form estimates and violation of IIA in gravity estimation

Dependent variable: log(exports)	(1)	(2)	(3)	(4)
log(freight cost)	-6.103**	-6.347**	-1.301**	-1.277**
	(1.046)	(1.259)	(0.392)	(0.381)
Test for joint significance of interacted competitors' freight costs: $\gamma_l = 0\ \forall l$				
F-stat		42.60**		209.24**
p-value		<0.001		<0.001
Disaggregation level	exp.-imp.		exp.-imp.-ind.	
Observations	1,184		18,486	

Notes. Sample of exports from 37 countries to Australia and USA between 1995 and 2010 (aggregate and 2-digit industry-level). All models include a full set of dummies for exporter-importer(-industry), importer-year(-industry), and exporter-year(-industry). Standard errors clustered by exporter-importer. ** p<0.01.

non-zero, we simply report the value of the F-test for the hypothesis that $\gamma_l = 0$ for all l. This test is comfortably rejected at the one percent level, while clustering standard errors at the exporter-importer level. Columns (3)-(4) estimate the same specification using trade data disaggregated by 2-digit industry. This exercise investigates whether the IIA violation is simply related to industry aggregation. Accordingly, we allow all fixed effects to be industry-specific which implies that parameters are estimated from within-industry variation. For exposition purposes, we impose the same coefficients β and γ_l across sectors. The hypothesis that $\gamma_l = 0$ for all l is again rejected.[34]

To summarize, Table 2.1 supports the relevance of third-country effects as captured by the interaction between competitor's freight costs and distance between per-capita GDPs, $|\kappa_j - \kappa_l| \ln z_{li,t}$. In the structural estimation below, we rely on exactly this variation to obtain estimates of the parameters controlling the cross-price elasticity, σ_α and σ_ϵ.

Structural Estimation

We now turn to our estimates of $\boldsymbol{\theta}$ obtained from the GMM procedure described in Section 2.6.2. These parameters are reported in Table 2.2 along with their accompanying standard errors clustered by exporter-importer pair.[35] In Panel A, we restrict $\sigma_\alpha = \sigma_\epsilon = 0$ in which case we estimate $\bar{\epsilon}$ to be approximately -6. As expected, this value is very similar to the estimate in column (1) of Table 2.1.[36]

Panel B reports our estimates with unobserved heterogeneity only in α, whereas Panel C focuses on our preferred specification with unobserved heterogeneity in both α and ϵ. As can be seen from Panel C, we estimate a value of σ_ϵ close to zero, indicating that deviations from IIA based on market shares are not important. However, the estimate of σ_α is statistically significant which suggests the importance of

[34] We obtain the same conclusion if all coefficients are allowed to vary by industry: the hypothesis that $\gamma_l^k = 0$ for all l and k is rejected at the one percent level.

[35] As noted by Stock and Yogo (2002), research on tests for weak instruments in the non-linear GMM case is still "quite incomplete." In principle, one could calculate the nonlinear Anderson-Rubin statistic proposed by these authors. Given the large number of fixed effects in equation (2.35), computing this statistic has proven too computationally demanding in practice.

[36] The estimates are not identical because they are based on two consistent, but distinct estimators of the same parameter.

Table 2.2: **GMM estimates of mixed CES demand**

	$\bar{\epsilon}$	σ_{α}	σ_{ϵ}
Panel A: CES			
	-5.955**		
	(0.671)		
Panel B: Mixed CES (restricted heterogeneity)			
	-6.115**	2.075**	
	(0.649)	(0.578)	
Panel C: Mixed CES (unrestricted heterogeneity)			
	-6.116**	2.063**	0.003
	(0.671)	(0.647)	(0.175)

Notes. Sample of 1,152 exporter-importer-year triples between 1995 and 2010 (normalizing country is the USA). Importers: Australia and USA. All models include a full set of dummies for importer-exporter and exporter-year. Standard errors (consistent, one-step standard errors, following the procedure in Appendix 2A.2) in parentheses clustered by 72 exporter-importer pairs are reported in parentheses. ** p<0.01.

IIA deviations related to per-capita GDP. To get more intuition about the economic implications of our structural estimates, Figure 2.1 plots the cross price-elasticity in equation (2.33) with respect to a change in Chinese trade costs. While this elasticity is identically equal to zero in the CES system of Panel A, it does not have to be the case for the other specifications. In fact, the parameters estimated in Panel C imply that the elasticity of relative demand to the relative price of Chinese factor is decreasing in per-capita GDP, being statistically different from zero even for low-income countries like China.

2.7 Application: China's Integration in the World Economy

We conclude by applying our methodology to study the consequences of one particular counterfactual: China's integration into the world economy. To shed light on this issue, we proceed in two steps. First, we use the demand system estimated in Section 2.6 to infer the trade costs faced by China, both as an exporter and an importer, at

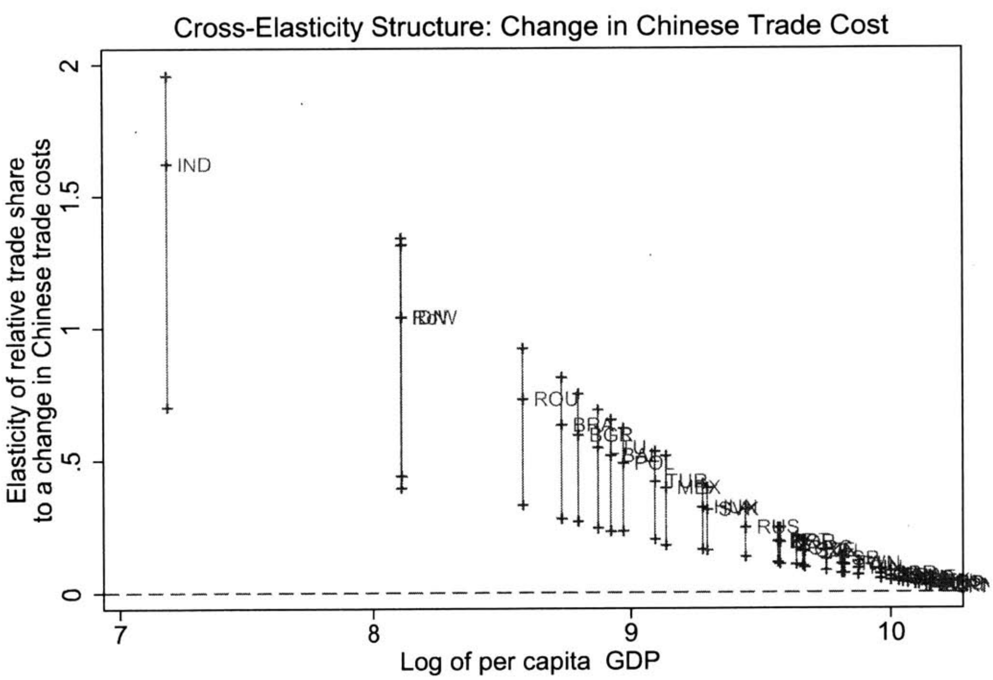

Figure 2.1: Elasticity of demand relative to the U.S. with respect to Chinese factor price.

Notes. Elasticity of U.S. demand for factors from any country relative to U.S. demand for U.S. factors with respect to a change in the Chinese factor price. Elasticities are computed using the estimates of the Mixed CES demand system in Panel C of Table 2.2. 95% confidence intervals shown are computed using the bootstrap procedure described in Appendix 2A.4. Dashed blue line corresponds to the CES case.

different points in time. Given estimates of Chinese trade costs, we then ask: "For any country j, how much higher (or lower) would welfare have been at a given year $t \geq 1995$ if Chinese trade costs were those of 1995 rather than those of year t?" The next subsection focuses on the estimation of trade costs. Counterfactual predictions will be discussed in Section 2.7.2.[37]

2.7.1 Trade Costs

We measure trade costs as follows. For each importer i and each year t in our sample, we start by inverting our demand system, $\bar{\chi}$, to go from the vector of expenditure shares, $\boldsymbol{x}_{i,t}$, to the vector of effective factor prices, $\boldsymbol{\omega}_{i,t} = \bar{\boldsymbol{\chi}}^{-1}(\boldsymbol{x}_{i,t})$, up to a normalization. We then use the time series of effective factor prices, $\{\omega_{ji,t} = \bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t})\}$, and the identity, $\omega_{ji,t} \equiv \tau_{ji,t} c_{i,t}$, to construct the time series of iceberg trade costs, $\{\tau_{ji,t}\}$, such that

$$(\tau_{ji,t}/\tau_{ii,t})/(\tau_{jj,t}/\tau_{ij,t}) = (\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t})/\bar{\chi}_i^{-1}(\boldsymbol{x}_{i,t}))/(\bar{\chi}_j^{-1}(\boldsymbol{x}_{j,t})/\bar{\chi}_i^{-1}(\boldsymbol{x}_{j,t})), \text{ for all } i,\ j, \text{ and } t. \tag{2.37}$$

This (log-)difference-in-differences provides a nonparametric generalization of the Head and Ries's (2001) index used to measure trade costs in gravity models. Compared to the case of a CES demand system, the only distinction is that one cannot directly read the difference-in-differences in effective prices from the difference-in-differences in expenditure shares. Inverting demand now requires a computer.

In order to go from a difference-in-differences to the level of Chinese trade costs,

[37] We follow a two-step procedure because we are interested in quantifying the welfare consequences of China's observed integration—interpreted as changes in iceberg trade costs within our theoretical framework—over the last two decades. Of course, one could dispense with the first step and directly study the effects of arbitrarily chosen changes in trade costs, including those not featuring the normalizations imposed in Section 2.7.1. This is the approach followed in most recent quantitative papers; see e.g. Costinot and Rodríguez-Clare (2013). Note also that our exercise is related to, but distinct from, the simulations in Hanson and Robertson (2010) and Hsieh and Ossa (2011), which evaluate the global consequences of China's sector-wise productivity growth using gravity models.

we follow the same approach as Head and Ries (2001) and further assume that

$$\tau_{ii,t}/\tau_{ii,95} = 1 \text{ for all } i \text{ and } t, \tag{2.38}$$

$$\tau_{ij,t}/\tau_{ij,95} = \tau_{ji,t}/\tau_{ji,95} \text{ for all } t \text{ if } i \text{ or } j \text{ is China.} \tag{2.39}$$

The first condition rules out differential changes in domestic trade costs around the world, whereas the second condition rules out asymmetric changes in Chinese trade costs.[38] Given equations (2.37)-(2.39), we can then measure the proportional changes in Chinese trade costs between 1995 and any period t as

$$\tau_{ji,t}/\tau_{ji,95} = \sqrt{\frac{(\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t})/\bar{\chi}_i^{-1}(\boldsymbol{x}_{i,t}))/(\bar{\chi}_j^{-1}(\boldsymbol{x}_{j,t})/\bar{\chi}_i^{-1}(\boldsymbol{x}_{j,t}))}{(\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,95})/\bar{\chi}_i^{-1}(\boldsymbol{x}_{i,95}))/(\bar{\chi}_j^{-1}(\boldsymbol{x}_{j,95})/\bar{\chi}_i^{-1}(\boldsymbol{x}_{j,95}))}}, \text{ if } i \text{ or } j \text{ is China.}$$

By construction, changes in exporting and importing costs from China are the same, though they may vary across trading partners and over time.

Figure 2.2 reports the arithmetic average of changes in Chinese trade costs across all trading partners. The solid red line corresponds to our baseline estimates, obtained under mixed CES (Table 2.2, Panel C). As can be seen, these are substantial changes in trade costs. Between 1995 and 2007, we estimate that Chinese trade costs decreased by 20.2% on average. If we were to restrict ourselves to a CES demand system (the dashed blue line), the decrease in Chinese trade costs would be equal to 16.7% instead.

[38] Our focus on symmetric changes in Chinese trade costs is partly motivated by the desire stay as close as possible to existing practices in the gravity literature. It should be clear, however, that while some normalization is required to go from differences-in-differences to the levels of trade costs, equations (2.38) and (2.39) provide only one of many possibilities. For example, an alternative would be to allow bilaterally asymmetric changes in Chinese trade costs under the assumption that some reference country's trade costs are constant over time. This is akin to focusing on counterfactuals in which one asks what would have happened if China had integrated with the rest of the world to the same extent as that reference country.

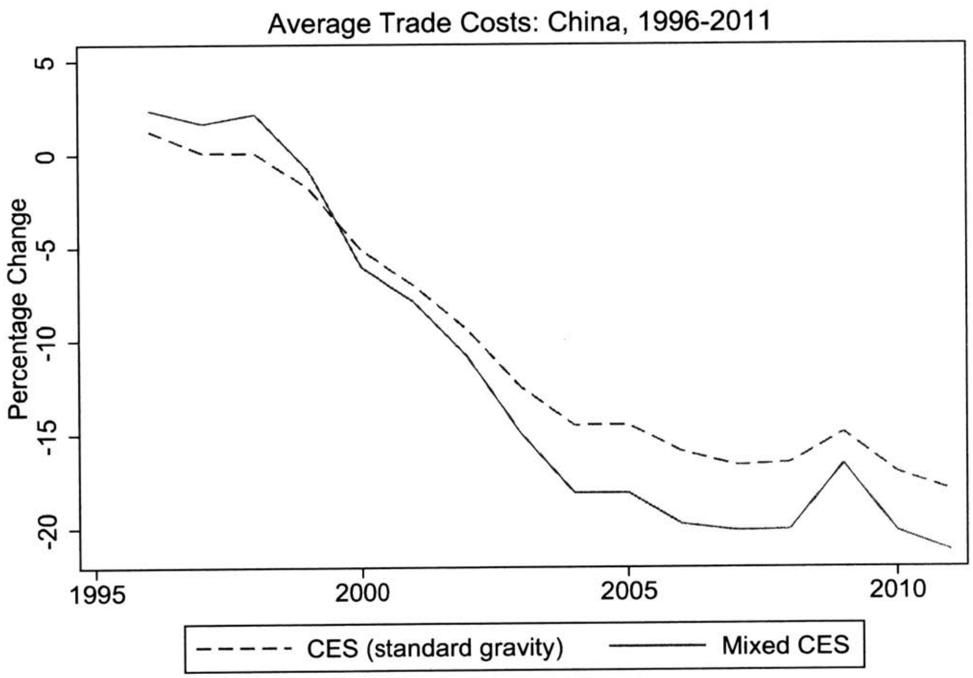

Figure 2.2: Average trade cost changes since 1995: China, 1996-2011.

Notes. Arithmetic average across all trading partners in the percentage reduction in Chinese trade costs between 1995 and each year $t = 1996, \ldots, 2011$. "CES (standard gravity)" and "Mixed CES" plot the estimates of trade costs obtained using the factor demand system in Panels A and C, respectively, of Table 2.2.

2.7.2 Counterfactual predictions

In any year t, we are interested in counterfactual changes in trade costs, $\hat{\tau}_{ji,t}$, such that Chinese trade costs are brought back to their 1995 levels:

$$\hat{\tau}_{ji,t} = \tau_{ji,95}/\tau_{ji,t}, \text{ if } i \text{ or } j \text{ is China,} \tag{2.40}$$

$$\hat{\tau}_{ji,t} = 1, \text{ otherwise.} \tag{2.41}$$

Given estimates of the factor demand system, obtained in Section 2.6, and estimates of trade costs, obtained in Section 2.7.1, we can use Corollary 1 to compute the welfare changes associated with this counterfactual scenario.[39]

Figure 2.3 reports the negative of the welfare changes in China for all years in our sample. A positive number in year t corresponds to the gains from economic integration for China between 1995 and year t. Before the great trade collapse in 2007, we see that the gains from economic integration for China are equal to 1.54%. In line with our estimates of trade costs, we see that imposing CES would instead lead to gains from economic integration equal to 1.04%.

What about China's trading partners? Figure 2.4 reports the welfare change from bringing Chinese trade costs back to their 1995 levels for all other countries in 2007. The bootstrapped 95% confidence intervals corresponding to each of these estimates (as well as those for China) can be found in Table 2A.2 in Appendix 2A.4. Under our preferred estimates (red circles), we see that rich countries tend to gain relatively more from China's integration, with both Indonesia and Romania experiencing statistically significant losses. The previous pattern gets muted if one forces factor demand to be CES instead (blue triangles).

[39]Our counterfactual calculations allow for lump-sum transfers between countries to rationalize trade imbalances in the initial equilibrium. We then hold these lump-sum transfers constant across the initial and counterfactual equilibria. Details on the algorithm for the computation of the counterfactual exercise are described in Appendix 2A.4.

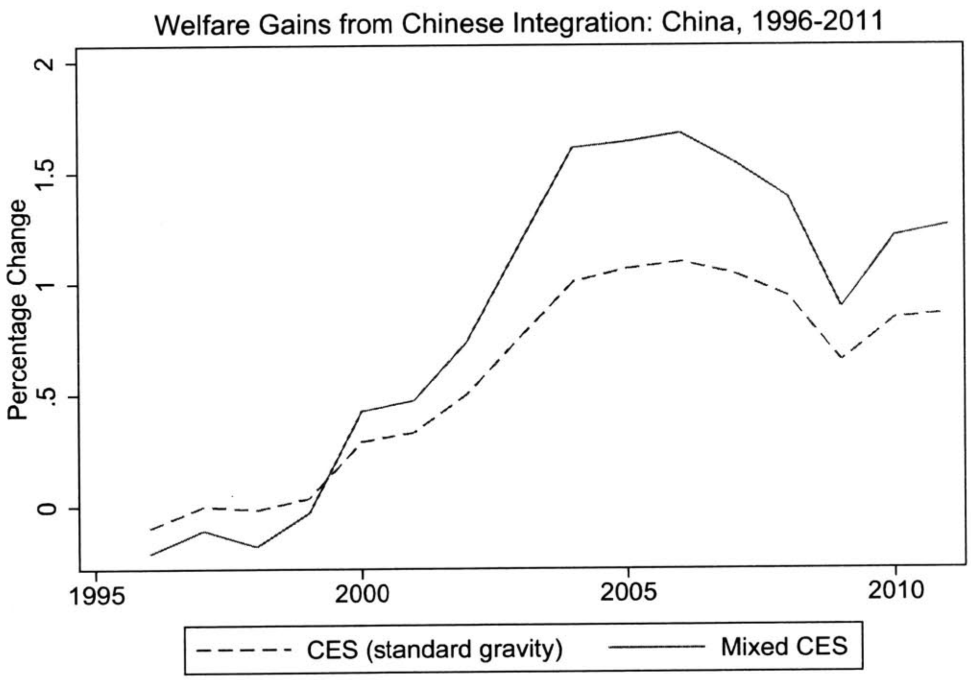

Figure 2.3: Welfare gains from Chinese integration since 1995: China, 1996-2011.

Notes. Welfare gains in China from reduction in Chinese trade costs relative to 1995 in each year $t = 1996, \ldots, 2011$. CES (standard gravity) and Mixed CES plot the estimates of welfare changes obtained using the factor demand system in Panels A and C, respectively, of Table 2.2.

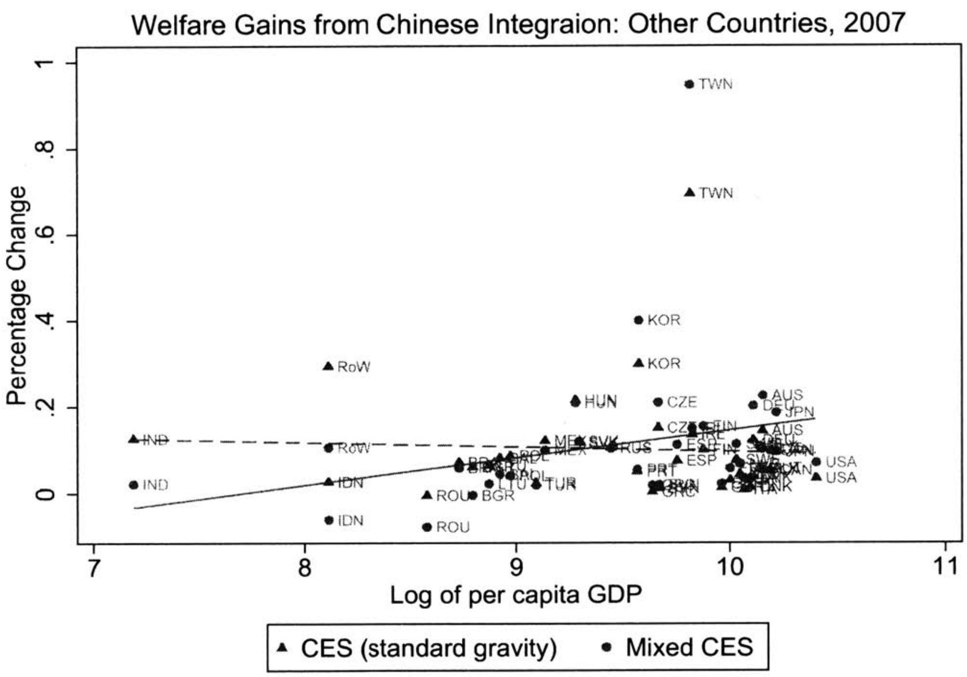

Figure 2.4: Welfare gains from Chinese integration since 1995: other countries, 2007.

Notes. Welfare gains in other countries from reduction in Chinese trade costs relative to 1995 in year $t = 2007$. "CES (standard gravity)" and "Mixed CES" plot the estimates of welfare changes obtained using the factor demand system in Panels A and C, respectively, of Table 2.2. The solid red line shows the line of best fit through the Mixed CES points, and the dashed blue line the equivalent for the CES case. Bootstrapped 95% confidence intervals for these estimates are reported in Table 2A.2.

2.8 Concluding Remarks

This paper starts from a simple observation. If neoclassical trade models are like exchange economies in which countries trade factor services, then the shape of these countries' reduced factor demand must be sufficient for answering many counterfactual questions.

Motivated by this observation, we have developed tools to conduct counterfactual and welfare analysis given knowledge of any factor demand system. Then, we have provided sufficient conditions under which estimates of this system can be recovered nonparametrically. Lastly, we have applied our tools to study a particular counterfactual question: What would have happened to other countries if China had remained closed? Since the answer to this question hinges on how substitutable factors of production from around the world are, we have introduced a parsimonious generalization of the CES demand system that allows for rich patterns of substitution across factors from different countries. The counterfactual results based on estimates of this system illustrate the feasibility and potential benefits of allowing trade data to speak with added flexibility.

Clearly, our emphasis on reduced factor demand also has costs. The demand system in our empirical application remains high-dimensional—we consider a world economy with 37 exporters—but data are limited—freight costs for these 37 exporters are only available for 16 years and 2 importers. So parametric restrictions need to be imposed. The typical approach is to impose such restrictions on deeper primitives of the model, like preferences and technology, and then to use various data sources to estimate or calibrate each of those fundamentals.[40] Here, we propose instead to impose restrictions directly on the factor demand system, while building estimation on precisely the moment conditions under which we have shown this system to be nonparametrically identified. Given data constraints, we do not view our approach as a panacea. But we believe that the tight connection between theory and data that it offers makes it worthy of further investigation.

[40] Bas, Mayer, and Thoenig (2015) provides an interesting example of this approach in the context of monopolistically competitive models of international trade.

An important open question concerns the extent to which one could combine the approach in this paper with additional, more disaggregated data sources. The answer is likely to depend on the additional assumptions that one is willing to impose, with costs and benefits that will need to be weighed. Consider, for instance, the differences in patterns of specialization across sectors and countries. Intuitively, there is a lot of information to be gained from such sector-level data. But if one is interested in aggregate questions, such data never come for free—disaggregated data will need to be aggregated ultimately. One possibility would be to use sector-level data, say in the pre-sample period, to construct additional observed country characteristics in a factor demand system akin to the one introduced in Section 2.6. Another possibility, closer to existing work, would be to maintain strong functional forms on the way that sector-level factor demands are aggregated, but allow for mixed CES demand systems to deal flexibly with the substantial unobserved heterogeneity across goods within narrowly defined sectors; see Schott (2004).

Regardless of the methodology that one chooses, we hope that our theoretical results can make more transparent how CGE models map data into counterfactual predictions. One cannot escape Manski's (2003) "Law of Decreasing Credibility," that "the credibility of inference decreases with the strength of the assumptions maintained" (p. 1). But identifying the critical assumptions upon which counterfactual predictions rely in complex general equilibrium environments can help evaluate their credibility. Once it is established that assumptions about the shape of factor demand—and only the shape of factor demand—determine counterfactual predictions, it becomes easier to ask whether the moments chosen for structural estimation are related to the economic relation of interest and to explore whether functional form assumptions rather than data drive particular results.

In terms of applications, two lines of research seem particularly promising. The first concerns the distributional consequences of international trade. By assuming the same factor intensity in all sectors, our empirical application assumes away distributional issues. None of the theoretical results in Sections 2.3 and 2.4, however, rely on this assumption. Hence the same nonparametric approach could be used to study the impact of globalization on the skill premium or the relative return to capital. The second line of research concerns the consequences of factor mobility, either migration or foreign direct investment. Although factor supply is inelastic in Section 2.3, it would be easy to incorporate such considerations by introducing intermediate goods, as we did in Section 2.4.3. Then either migration or foreign direct investment would be equivalent to trade in intermediate goods, which may be subject to different frictions than trade in final goods.

Finally, while we have emphasized counterfactual and welfare analysis in this paper, the tools that we have developed could be applied more generally. Many questions concerning international trade can be reduced to estimating and inverting a demand system. But this system does not have to be CES. In Section 2.7.1, we have already mentioned the measurement of trade costs, which is an important application of gravity models; see e.g. Anderson and Van Wincoop (2004) and Jacks, Meissner, and Novy (2011). Another natural application is the measurement of comparative advantage; see e.g. Costinot, Donaldson, and Komunjer (2012a) and Levchenko and Zhang (2011). Measures of revealed comparative advantage (RCA) aim to uncover which countries can produce and sell goods relatively more cheaply, and this boils down to a difference-in-differences of (log-)prices. Away from CES, this difference-in-differences will not be proportional to a difference-in-differences of (log-)expenditures. But given estimates of any invertible demand system, RCA remains an easy object to compute.

2A Appendix

2A.1 Proofs

2A.1.1 Proposition 1

Proof of Proposition 1. ($\Rightarrow$) Suppose that $(\boldsymbol{q}, \boldsymbol{l}, \boldsymbol{p}, \boldsymbol{w})$ is a competitive equilibrium. For any country i, let us construct $\boldsymbol{L}_i \equiv \{L_{ji}^n\}$ such that

$$L_{ji}^n = \sum_k l_{ji}^{nk} \text{ for all } i,\ j, \text{ and } n.$$

Together with the factors market clearing condition (2.5), the previous expression immediately implies

$$\sum_j L_{ij}^n = v_i^n \text{ for all } i \text{ and } n.$$

In order to show that $(\boldsymbol{L}, \boldsymbol{w})$ is a reduced equilibrium, we therefore only need to show

$$\begin{aligned} &\boldsymbol{L}_i \in \operatorname{argmax}_{\tilde{\boldsymbol{L}}_i} U_i(\tilde{\boldsymbol{L}}_i) \\ &\sum_{j,n} w_j^n \tilde{L}_{ji}^n \leq \sum_n w_i^n v_i^n \text{ for all } i. \end{aligned} \tag{2.42}$$

We proceed by contradiction. Suppose that there exists a country i such that condition (2.42) does not hold. Since profits are zero in a competitive equilibrium with constant returns to scale, we must have $\sum_{j,k} p_{ji}^k q_{ji}^k = \sum_{j,n} w_j^n L_{ji}^n$. The budget constraint of the representative agent in the competitive equilibrium, in turn, implies $\sum_{j,n} w_j^n L_{ji}^n = \sum_n w_i^n v_i^n$. Accordingly, if condition (2.42) does not hold, there must be $\boldsymbol{L}_i'$ such that $U_i(\boldsymbol{L}_i') > U_i(\boldsymbol{L}_i)$ and $\sum_{j,n} w_j^n (L_{ji}^n)' \leq \sum_n w_i^n v_i^n$. Now consider $(\boldsymbol{q}_i', \boldsymbol{l}_i')$

such that

$$
\begin{aligned}
(\boldsymbol{q}_i', \boldsymbol{l}_i') \in\ & \operatorname{argmax}_{\tilde{\boldsymbol{q}}_i, \tilde{\boldsymbol{l}}_i} u_i(\tilde{\boldsymbol{q}}_i) \\
& \sum_k \tilde{l}_{ji}^{nk} \leq (L_{ji}^n)' \text{ for all } j \text{ and } f, \\
& \tilde{q}_{ji}^k \leq f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) \text{ for all } j \text{ and } k.
\end{aligned}
$$

We must have

$$u_i(\boldsymbol{q}_i') = U_i(\boldsymbol{L}_i') > U_i(\boldsymbol{L}_i) \geq u_i(\boldsymbol{q}_i),$$

where the last inequality derives from the fact that, by construction, $\boldsymbol{L}_i$ is sufficient to produce $\boldsymbol{q}_i$. Utility maximization in the competitive equilibrium therefore implies

$$\sum_{j,k} p_{ji}^k (q_{ji}^k)' > \sum_n w_i^n v_i^n.$$

Combining this inequality with $\sum_{j,n} w_j^n (L_{ji}^n)' \leq \sum_n w_i^n v_i^n$, we obtain

$$\sum_{j,k} p_{ji}^k (q_{ji}^k)' > \sum_{j,n} w_j^n (L_{ji}^n)'.$$

Hence, firms could make strictly positive profits by using $\boldsymbol{L}_i'$, to produce $\boldsymbol{q}_i'$, which cannot be true in a competitive equilibrium. This establishes that $(\boldsymbol{L}, \boldsymbol{w})$ is a reduced equilibrium with the same factor prices and the same factor content of trade as the competitive equilibrium. The fact that $U_i(\boldsymbol{L}_i) = u_i(\boldsymbol{q}_i)$ can be established in a similar manner. If there were $\boldsymbol{q}_i'$ such that $u_i(\boldsymbol{q}_i') = U_i(\boldsymbol{L}_i) > u_i(\boldsymbol{q}_i)$, then utility maximization would imply

$$\sum_{j,k} p_{ji}^k (q_{ji}^k)' > \sum_n w_i^n v_i^n = \sum_{j,n} w_j^n L_{ji}^n,$$

which would in turn violate profit maximization in the competitive equilibrium.

($\Leftarrow$) Suppose that $(\boldsymbol{L}, \boldsymbol{w})$ is a reduced equilibrium. For any positive of vector of output delivered in country i, $\boldsymbol{q}_i \equiv \{q_{ji}^k\}$, let $C_i(\boldsymbol{w}, \boldsymbol{q}_i)$ denote the minimum cost of producing $\boldsymbol{q}_i$,

$$C_i(\boldsymbol{w}, \boldsymbol{q}_i) \equiv \min_{\tilde{l}} \sum_{j,k,n} w_j^n \tilde{l}_{ji}^{nk} \tag{2.43}$$

$$q_{ji}^k \leq f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) \text{ for all } j \text{ and } k. \tag{2.44}$$

The first step of our proof characterizes basic properties of C_i. The last two steps use these properties to construct a competitive equilibrium that replicates the factor content of trade and the utility levels in the reduced equilibrium.

Step 1. *For any country i, there exists $\boldsymbol{p}_i \equiv \{p_{ji}^k\}$ positive such that the two following conditions hold:(i)*

$$C_i(\boldsymbol{w}, \boldsymbol{q}_i) = \sum_{j,k} p_{ji}^k q_{ji}^k, \textit{ for all } \boldsymbol{q}_i > 0, \tag{2.45}$$

and (ii) if $\boldsymbol{l}_i$ solves (2.43), then $\boldsymbol{l}_i$ solves

$$max_{\tilde{\boldsymbol{l}}_{ji}^k} p_{ji}^k f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) - \sum_n w_j^n \tilde{l}_{ji}^{nk} \textit{ for all } j \textit{ and } k. \tag{2.46}$$

For any i, j, and k, let us construct p_{ji}^k such that

$$p_{ji}^k = \min_{\tilde{\boldsymbol{l}}_{ji}^k} \{\sum_n w_j^n \tilde{l}_{ji}^{nk} | f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) \geq 1\}. \tag{2.47}$$

Take $\boldsymbol{l}_{ji}^k(1)$ that solves the previous unit cost minimization problem. Since f_{ji}^k is homogeneous of degree one, we must have $f_{ji}^k(q_{ji}^k \boldsymbol{l}_{ji}^k(1)) \geq q_{ji}^k$. By definition of C_i, we must also have $C_i(\boldsymbol{w}, \boldsymbol{q}_i) \leq \sum_{j,k,n} q_{ji}^k w_j^n l_{ji}^{nk}(1) = \sum_{j,k} p_{ji}^k q_{ji}^k$. To show that equation (2.45) holds, we therefore only need to show that $C_i(w, q_i) \geq \sum_{j,k} p_{ji}^k q_{ji}^k$. We proceed by contradiction. Suppose that $C_i(\boldsymbol{w}, \boldsymbol{q}_i) < \sum_{j,k} p_{ji}^k q_{ji}^k$. Then there must be $q_{ji}^k > 0$ such that

$$\sum_n w_j^n l_{ji}^{nk} < q_{ji}^k \sum_n w_j^n l_{ji}^{nk}(1),$$

where $\boldsymbol{l}_{ji}^k$ is part of the solution of (2.43). Since f_{ji}^k is homogeneous of degree one,

$\boldsymbol{l}^k_{ji}/q^k_{ji}$ would then lead to strictly lower unit cost then $\boldsymbol{l}^k_{ji}(1)$, which cannot be. This establishes condition (i).

To establish condition (ii), we proceed again by contradiction. Suppose that there exists $(\boldsymbol{l}^k_{ji})'$ such that

$$p^k_{ji} f^k_{ji}((\boldsymbol{l}^k_{ji})') - \sum_n w^n_j (l^{nk}_{ji})' > p^k_{ji} f^k_{ji}(\boldsymbol{l}^k_{ji}) - \sum_n w^n_j l^{nk}_{ji}. \tag{2.48}$$

Take the vector of output $\boldsymbol{q}_i$ such that $q^k_{ji} = f^k_{ji}(\boldsymbol{l}^k_{ji})$ and zero otherwise. Condition (i) applied to that vector immediately implies

$$p^k_{ji} f^k_{ji}(\boldsymbol{l}^k_{ji}) = \sum_n w^n_j l^{nk}_{ji}.$$

Combining this observation with inequality (2.48), we get $p^k_{ji} > \sum_n w^n_j (l^{nk}_{ji})' / f^k_{ji}((\boldsymbol{l}^k_{ij})')$, which contradicts the fact that p^k_{ji} is the minimum unit cost.

Step 2. *Suppose that* $(\boldsymbol{q}_i, \boldsymbol{l}_i)$ *solves*

$$\begin{aligned} & max_{\tilde{\boldsymbol{q}}_i, \tilde{\boldsymbol{l}}_i} u_i(\tilde{\boldsymbol{q}}_i) \\ & \tilde{q}^k_{ji} \leq f^k_{ji}(\tilde{\boldsymbol{l}}^k_{ji}) \text{ for all } j \text{ and } k, \\ \sum_k w^n_j \tilde{l}^{nk}_{ji} & \leq \sum_n w^n_i v^n_i. \end{aligned} \tag{2.49}$$

Then $\boldsymbol{q}_i$ *solves*

$$\begin{aligned} & max_{\tilde{\boldsymbol{q}}_i} u_i(\tilde{\boldsymbol{q}}_i) \\ \sum_{j,k} p^k_{ji} \tilde{q}^k_{ji} & \leq \sum_n w^n_i v^n_i, \end{aligned} \tag{2.50}$$

and $\boldsymbol{l}_i$ *solves*

$$max_{\tilde{\boldsymbol{l}}^k_{ji}} p^k_{ji} f^k_{ji}(\tilde{\boldsymbol{l}}^k_{ji}) - \sum_n w^n_j \tilde{l}^{nk}_{ji} \text{ for all } j \text{ and } k. \tag{2.51}$$

If $(\boldsymbol{q}_i, \boldsymbol{l}_i)$ solves (2.49), then

$$\begin{aligned} \boldsymbol{q}_i \in & \ argmax_{\tilde{\boldsymbol{q}}_i} u_i(\tilde{\boldsymbol{q}}_i) \\ & C_i(\boldsymbol{w}, \tilde{\boldsymbol{q}}_i) \leq \sum_n w_i^n v_i^n. \end{aligned}$$

Combining this observation with Step 1 condition (i), we obtain that $\boldsymbol{q}_i$ solves (2.50). Likewise, if $(\boldsymbol{q}_i, \boldsymbol{l}_i)$ solves (2.49), then

$$\begin{aligned} \boldsymbol{l}_i \in & \ \text{argmin}_{\tilde{\boldsymbol{l}}} \sum_{j,k,n} w_j^n \tilde{l}_{ji}^{nk}, \\ & q_{ji}^k \leq f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) \text{ for all } j \text{ and } k. \end{aligned}$$

Combining this observation with Step 1 condition (ii), we obtain that $\boldsymbol{l}_i$ solves (2.51).

Step 3. *For all i, take $(\boldsymbol{q}_i, \boldsymbol{l}_i)$ that solves*

$$\begin{aligned} & max_{\tilde{\boldsymbol{q}}_i, \tilde{\boldsymbol{l}}_i} u_i(\tilde{\boldsymbol{q}}_i) \\ & \quad \tilde{q}_{ji}^k \leq f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k) \textit{ for all } j \textit{ and } k, \\ & \sum_k \tilde{l}_{ji}^{nk} \leq L_{ji}^n \textit{ for all } j \textit{ and } n, \end{aligned} \tag{2.52}$$

and set $\boldsymbol{q} = \sum_i \boldsymbol{q}_i$ and $\boldsymbol{l} = \sum_i \boldsymbol{l}_i$. Then $(\boldsymbol{q}, \boldsymbol{l}, \boldsymbol{p}, \boldsymbol{w})$ is a competitive equilibrium with the same factor prices, $\boldsymbol{w}$; (ii) the same factor content of trade, $L_{ji}^n = \sum_k l_{ji}^{nk}$ for all i, j, and n; and (iii) the same welfare levels, $U_i(\boldsymbol{L}_i) = u_i(\boldsymbol{q}_i)$ for all i.

Since $(\boldsymbol{L}, \boldsymbol{w})$ is a reduced equilibrium, if $(\boldsymbol{q}_i, \boldsymbol{l}_i)$ solves (2.52), then $(\boldsymbol{q}_i, \boldsymbol{l}_i)$ solves (2.49). By Step 2, $\boldsymbol{q}_i$ and $\boldsymbol{l}_i$ must therefore solve (2.50) and (2.51), respectively. Hence, the utility maximization and profit maximization conditions (2.1) and (2.3) are satisfied. Since the constraint $\tilde{q}_{ji}^k \leq f_{ji}^k(\tilde{\boldsymbol{l}}_{ji}^k)$ must be binding for all j and k in any country i, the good market clearing condition (2.4) is satisfied as well. The factor market clearing condition directly derives from the fact that $(\boldsymbol{L}, \boldsymbol{w})$ is a reduced equilibrium and the constraint, $\sum_k \tilde{l}_{ji}^{nk} \leq L_{ji}^n$, must be binding for all j and n in any country i. By construction, conditions (i)-(iii) necessarily hold. $\square$

2A.1.2 Lemma 1

Proof of Lemma 1. We proceed in two steps.

Step 1. *In a Ricardian economy, if good expenditure shares satisfy the connected substitute property, then factor expenditure shares satisfy the connected substitute property.*

Our goal is to establish that factor demand, $\bar{\boldsymbol{\chi}}$, satisfies the connected substitute property—expressed in terms of the effective prices of the composite factors, $\boldsymbol{\omega}_j \equiv \{\tau_{ij}c_i\}$—if good demand, $\bar{\boldsymbol{\sigma}}$, satisfies the connected substitute property, with

$$\bar{\boldsymbol{\sigma}}(\boldsymbol{p}_j) \equiv \{\{s_j^k\} | s_j^k = p_j^k q_j^k / y_j \text{ for some } \boldsymbol{q}_j \in \text{argmax}_{\tilde{\boldsymbol{q}}}\{\bar{u}(\tilde{\boldsymbol{q}}) | \sum_k p_j^k \tilde{q}_j^k \leq y_j\}\}.$$

Note that since $\bar{u}$ is homothetic, $\bar{\boldsymbol{\sigma}}$ does not depend on income in country j. For notational convenience, we omit the importer's index, j, in the rest of this proof.

Consider a change in effective factor prices from $\boldsymbol{\omega}$ to $\boldsymbol{\omega}'$ and a partition of countries $\{M_1, M_2\}$ such that $\omega_i' > \omega_i$ for all $i \in M_1$ and $\omega_i' = \omega_i$ for all $i \in M_2$. Now take $\boldsymbol{x}, \boldsymbol{x}' > 0$ such that $\boldsymbol{x} \in \bar{\boldsymbol{\chi}}(\boldsymbol{\omega})$ and $\boldsymbol{x}' \in \bar{\boldsymbol{\chi}}(\boldsymbol{\omega}')$. For each exporting country i, we can decompose total expenditure shares into the sum of expenditure shares across all sectors k,

$$x_i = \sum_k s^k x_i^k,$$

where s^k denotes the share of expenditure on good k at the initial prices,

$$\{s^k\} \in \bar{\sigma}(\{p^k(\boldsymbol{\omega})\}),$$
$$p^k(\boldsymbol{\omega}) = \min_i \{\omega_i / \alpha_i^k\}.$$

For any good k, there are two possible cases. If no country $i \in M_2$ has the minimum cost for good k at the initial factor prices, $\boldsymbol{\omega}$, then

$$\sum_{i \in M_2} x_i^k = 0, \tag{2.53}$$

$$p^k(\boldsymbol{\omega}) < p^k(\boldsymbol{\omega}'). \tag{2.54}$$

Let us call this set of good K_1. If at least one country $i \in M_2$ has the minimum cost for good k, then

$$\sum_{i \in M_2} (x_i^k)' = 1, \tag{2.55}$$

$$p^k(\boldsymbol{\omega}) = p^k(\boldsymbol{\omega}'). \tag{2.56}$$

Let us call this second set of good K_2. Since $\boldsymbol{x}, \boldsymbol{x}' > 0$, we know that both K_1 and K_2 are non-empty.

Now consider the total expenditure on factors from countries $i \in M_2$ when factor prices are equal to $\boldsymbol{\omega}'$. It must satisfy

$$\sum_{i \in M_2} (x_i)' \geq \sum_{i \in M_2} \sum_{k \in K_2} (s^k)'(x_i^k)' = \sum_{k \in K_2} (s^k)' [\sum_{i \in M_2} (x_i^k)'].$$

Combining the previous inequality with (2.55), we obtain

$$\sum_{i \in M_2} (x_i)' \geq \sum_{k \in K_2} (s^k)'.$$

By the Inada conditions, all goods are consumed. Thus, we can invoke the connected substitute property for goods in K_1 and K_2. Conditions (2.54) and (2.56) imply

$$\sum_{k \in K_2} (s^k)' > \sum_{k \in K_2} s^k.$$

Since $\sum_{i \in M_2} x_i^k \leq 1$, the two previous inequalities further imply

$$\sum_{i \in M_2} (x_i)' > \sum_{k \in K_2} s^k [\sum_{i \in M_2} x_i^k] = \sum_{i \in M_2} \sum_{k \in K_2} s^k x_i^k.$$

Finally, using (2.53) and the fact that $\{K_1, K_2\}$ is a partition, we get

$$\sum_{i \in M_2} (x_i)' > \sum_{i \in M_2} \sum_{k \in K_1} s^k x_i^k + \sum_{i \in M_2} \sum_{k \in K_2} s^k x_i^k = \sum_{i \in M_2} x_i.$$

This establishes that $\bar{\chi}$ satisfies the connected substitute property.

Step 2. *If factor demand $\bar{\chi}$ satisfies the connected substitute property, then for any vector of factor expenditure shares, $\boldsymbol{x} > 0$, there is at most one vector (up to a normalization) of effective factor prices, $\boldsymbol{\omega}$, such that $\boldsymbol{x} \in \bar{\chi}(\boldsymbol{\omega})$.*

We proceed by contradiction. Suppose that there exist $\boldsymbol{\omega}$, $\boldsymbol{\omega}'$, and $\boldsymbol{x}_0 > 0$ such that $\boldsymbol{x}_0 \in \bar{\chi}(\boldsymbol{\omega})$, $\boldsymbol{x}_0 \in \bar{\chi}(\boldsymbol{\omega}')$, and $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$ are not collinear. Since $\bar{\chi}$ is homogeneous of degree zero in all factor prices, we can assume without loss of generality that $\omega_i \geq \omega_i'$ for all i, with at least one strict inequality and one equality. Now let us partition all countries into two groups, M_1 and M_2, such that

$$\omega_i' > \omega_i \text{ if } i \in M_1, \tag{2.57}$$

$$\omega_i' = \omega_i \text{ if } i \in M_2. \tag{2.58}$$

Since $\bar{\chi}$ satisfies the connected substitute property, conditions (2.57) and (2.58) imply that for any $\boldsymbol{x}, \boldsymbol{x}' > 0$ such that $\boldsymbol{x} \in \bar{\chi}(\boldsymbol{\omega})$ and $\boldsymbol{x}' \in \bar{\chi}(\boldsymbol{\omega}')$, we must have

$$\sum_{i \in M_2} x_i' > \sum_{i \in M_2} x_i,$$

which contradicts the existence of $\boldsymbol{x}_0 \in \bar{\chi}(\boldsymbol{\omega}) \cap \bar{\chi}(\boldsymbol{\omega}')$. Lemma 1 follows from Steps 1 and 2. $\square$

2A.1.3 Lemma 2

Proof of Lemma 2. We proceed by contradiction. Suppose that there exist two equilibrium vectors of factor prices, $\boldsymbol{c} \equiv (c_1, ..., c_I)$ and $\boldsymbol{c}' \equiv (c_1', ..., c_I')$, that are not collinear. By Proposition 1, we know that $\boldsymbol{c}$ and $\boldsymbol{c}'$ must be equilibrium vectors of the reduced exchange model. So they must satisfy

$$\sum_j L_{ij} = \bar{f}_i(\boldsymbol{\nu}_i), \text{ for all } i, \tag{2.59}$$

$$\sum_j L_{ij}' = \bar{f}_i(\boldsymbol{\nu}_i), \text{ for all } i, \tag{2.60}$$

where $\{L_{ij}\}$ and $\{L_{ij}'\}$ are the optimal factor demands in the two equilibria,

$$\{L_{ij}\} \in \bar{\boldsymbol{L}}(\boldsymbol{\omega}_j), \text{ for all } j,$$
$$\{L_{ij}'\} \in \bar{\boldsymbol{L}}(\boldsymbol{\omega}_j'), \text{ for all } j,$$

where $\boldsymbol{\omega}_j \equiv \{\tau_{ij} c_i\}$ and $\boldsymbol{\omega}_j' \equiv \{\tau_{ij} c_i'\}$ are the associated vectors of effective factor prices.

We can follow the same strategy as in Step 2 of the proof of Lemma 2A.1.3. Without loss of generality, let us assume that $c_i' \geq c_i$ for all i, with at least one strict inequality and one equality. We can again partition all countries into two groups, M_1 and M_2, such that

$$c_i' > c_i \text{ if } i \in M_1, \tag{2.61}$$

$$c_i' = c_i \text{ if } i \in M_2. \tag{2.62}$$

The same argument then implies that in any country j,

$$\sum_{i \in M_2} x_{ij}' > \sum_{i \in M_2} x_{ij},$$

where $\{x_{ij}\}$ and $\{x_{ij}'\}$ are the expenditure shares associated with $\{L_{ij}\}$ and $\{L_{ij}'\}$, respectively. By definition of the factor expenditure shares, the previous inequality can can be rearranged as

$$\sum_{i\in M_2} c_i' L_{ij}'/(c_j'\bar{f}_j(\boldsymbol{\nu}_j)) > \sum_{i\in M_2} c_i L_{ij}/(c_j\bar{f}_j(\boldsymbol{\nu}_j)).$$

Since $c_i \geq c_i'$ for all i, this implies

$$\sum_{i\in M_2} c_i' L_{ij}' > \sum_{i\in M_2} c_i L_{ij}.$$

Summing across all importers j, we therefore have

$$\sum_{i\in M_2} c_i' \sum_j L_{ij}' > \sum_{i\in M_2} c_i \sum_j L_{ij}.$$

By equations (2.59) and (2.60), this further implies

$$\sum_{i\in M_2} c_i'\bar{f}_i(\boldsymbol{\nu}_i) > \sum_{i\in M_2} c_i\bar{f}_i(\boldsymbol{\nu}_i),$$

which contradicts (2.62). □

2A.2 Estimation

In this section we discuss further details of the estimation procedure outlined in Section 2.6.2.

2A.2.1 GMM Estimator

As in Section 2.6.2, define the stacked matrix of instruments, $\mathbf{Z} \equiv [\mathbf{Z^1} \mid \mathbf{Z^2}]$, and the stacked vector of errors, $\mathbf{e}(\boldsymbol{\theta}) \equiv \bar{\chi}^{-1}(\boldsymbol{x}|\boldsymbol{\theta}_2) - \boldsymbol{Z}^1 \cdot \boldsymbol{\theta}_1$. The GMM estimator is

$$\widehat{\boldsymbol{\theta}} = arg\min_{\theta} \mathbf{e}(\boldsymbol{\theta})'\mathbf{Z}\Phi\mathbf{Z}'\mathbf{e}(\boldsymbol{\theta}).$$

where Φ is the GMM weight. We confine attention to the consistent one-step procedure by setting $\Phi = (\mathbf{Z}'\mathbf{Z})^{-1}$.[41]

2A.2.2 Standard Errors

In our baseline specification, we acknowledge the possibility of autocorrelation in the error term. In particular, we assume that observations are independent across exporter-importer pairs, but there is arbitrary autocorrelation across periods for the same pair. Following Cameron and Miller (2010), we have that

$$\sqrt{M}\left(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}\right) \to N\left[0, \left(B'\Phi B\right)^{-1}\left(B'\Phi\Lambda\Phi B\right)\left(B'\Phi B\right)^{-1}\right]$$

where $B \equiv E\left[Z'_{ji,t}\nabla_\theta e_{ji,t}(\boldsymbol{\theta})\right]$ and $\Lambda \equiv E[(Z'_{ji}e_{ji})(Z'_{ji}e_{ji})']$, with $Z_{ji} = [Z_{ji,t}]_{t=1}^{T}$ and $e_{ji} \equiv [e_{ji,t}]_{t=1}^{T}$ being matrices of stacked periods for exporter-importer pair (j, i).

The covariance matrix can be consistently estimated using

$$\widehat{Avar}(\widehat{\boldsymbol{\theta}}) \equiv \left(\hat{B}'\Phi\hat{B}\right)^{-1}\left(\hat{B}'\Phi\hat{\Lambda}\Phi\hat{B}\right)\left(\hat{B}'\Phi\hat{B}\right)^{-1} \tag{2.63}$$

where $\hat{B} \equiv \left(\mathbf{Z}'\nabla_\theta \mathbf{e}(\widehat{\boldsymbol{\theta}})\right)$, $\nabla_\theta \mathbf{e}(\widehat{\boldsymbol{\theta}}) \equiv [D_{\boldsymbol{\theta}_2}\bar{\boldsymbol{\chi}}^{-1}(\boldsymbol{x}|\boldsymbol{\theta}_2) \mid -\mathbf{Z}^{\mathbf{1}}]$, and $\hat{\Lambda} \equiv \Gamma'\Gamma$ such that $\Gamma \equiv \left[e_{ji}(\widehat{\boldsymbol{\theta}})'Z_{ji}\right]_{ji}$.

This analysis ignored the fact that we take draws of (α_s, ϵ_s) to compute simulated moment conditions in the algorithm described below. Although this simulation step affects standard errors, the asymptotic distribution of the estimator is the same as the number of simulated draws goes to infinite. Thus, we compute the covariance matrix according to expression (2.63) which is assumed to be an appropriate approximation for the large number of simulations (discussed below) used in the empirical implementation.

[41] This is the efficient estimator under homoskedasticity, however this weight matrix leads to an inneficient estimator under a more general covariance structure of errors.

2A.2.3 Estimation Algorithm

The simulated GMM procedure is implemented with the following steps.

Step 0. Draw S simulated pairs $(\alpha_s, \ln \epsilon_s) \sim N(0, I)$. We set $S = 4,000$ and use the same draws for all markets.

Step 1. Conditional on $\boldsymbol{\theta_2}$, compute the vector $\bar{\chi}^{-1}(\boldsymbol{x}_{i,t}|\theta_2) \equiv \{\delta_{ji,t}\}_{j=2}^N$ that solves the following system:

$$\{\bar{\chi}_j(\boldsymbol{\delta}_{i,t}|\boldsymbol{\theta}_2)\}_{j=2}^N = \{x_{ji,t}\}_{j=2}^N$$

where $x_{ji,t}$ is the expenditure share of importer i on exports of j at year t and

$$\bar{\chi}_j(\boldsymbol{\delta}_{i,t}|\boldsymbol{\theta}_2) = \frac{1}{S}\sum_{s=1}^{S} \frac{\exp[\alpha_s \sigma_\alpha \ln \kappa_j + (\epsilon_s)^{\sigma_\epsilon}\delta_{ji,t}]}{1 + \sum_{l=2}^N \exp\left[\alpha_s \sigma_\alpha \ln \kappa_l + (\epsilon_s)^{\sigma_\epsilon}\delta_{li,t}\right]}.$$

Uniqueness and existence of the solution is guaranteed by the fixed point argument in Berry, Levinsohn, and Pakes (1995). To solve the system, consider the fixed point of the following function:

$$G\left(\boldsymbol{\delta}_{i,t}\right) = \left[\delta_{ji,t} + \lambda\left(\ln x_{ji,t} - \ln \bar{\chi}_j(\boldsymbol{\delta}_{i,t}|\theta_2)\right)\right]_{j=2}^N$$

where λ is a parameter controlling the adjustment speed. This fixed point is obtained as the limit of the sequence: $\boldsymbol{\delta}_{i,t}^{n+1} = G\left(\boldsymbol{\delta}_{i,t}^{n}\right)$. Numerically, we compute the sequence until $\max_j \left|\ln x_{ji,t} - \ln \bar{\chi}_j(\boldsymbol{\delta}_{i,t}|\theta_2)\right| < tol$, where tol is some small number that we discuss further below.

This step is implemented as follows. First, the initial guess $\delta_{ij,t}^0$ in the initial iteration is set to be the logit solution $\delta_{ji,t}^0 = \ln x_{ji,t} - \ln x_{1i,t}$. In subsequent iterations, we use the following rule. If $\boldsymbol{\theta_2}$ is close to the parameter vector of the previous iteration, we use the system solution in the last iteration. Otherwise, we use the vector that solved the system for the same importer in the previous year (if it is the first year, we use the logit solution). Second, the speed of adjustment is initially set to $\lambda = 3$. If distance increases in iteration n, then we reduce λ by 5% and compute

$\boldsymbol{\delta}_{i,t}^{n+1}$ again until distance decreases in the step and use the new value of λ until the solution is found. If λ falls below a minimum ($\underline{\lambda} = .001$), then we assume no solution for the system and set the objective function to a high value. Lastly, we set $tol = 10^{-8}$ and, every 20,000 iterations, we increase tolerance by a factor of two. This guarantees that the algorithm does not waste time on convergence for parameter values far away from the real ones.[42]

Step 2. Conditional on $\boldsymbol{\theta_2}$, solve analytically for linear parameters directly from the minimization problem: $\widehat{\boldsymbol{\theta}}_{\mathbf{1}}(\boldsymbol{\theta_2}) = \left(\mathbf{Z^{1\prime}Z\Phi Z^{\prime}Z^{1}}\right)^{-1}\mathbf{Z^{1\prime}Z\Phi Z^{\prime}}\bar{\chi}^{-1}(\boldsymbol{x}|\boldsymbol{\theta}_2)$.

Step 3. Conditional on $\boldsymbol{\theta_2}$, compute the vector of structural errors: $\mathbf{e}(\boldsymbol{\theta}_2) = \bar{\chi}^{-1}(\boldsymbol{x}|\boldsymbol{\theta}_2) - \mathbf{Z}^1 \cdot \widehat{\boldsymbol{\theta}}_{\mathbf{1}}(\boldsymbol{\theta_2})$

Step 4. Numerically minimize the objective function to obtain estimates of $\boldsymbol{\theta}_2$:

$$\hat{\boldsymbol{\theta}}_{\mathbf{2}} \equiv arg \min_{\theta_2} H\left(\boldsymbol{\theta_2}\right) \equiv \mathbf{e}(\boldsymbol{\theta}_2)'\mathbf{Z\Phi Z'}\mathbf{e}(\boldsymbol{\theta}_2).$$

The numerical minimization is implemented using the "trust-region-reflective" algorithm that requires an analytical gradient of the objective function (described below). This algorithm is intended to be more efficient in finding the local minimum within a particular attraction region. First, we solve the minimization problem using a grid of ten initial conditions randomly drawn from a uniform distribution in the parameter space. Second, we solve a final minimization problem using as initial condition the minimum solution obtained from the first-round minimization. Here, we impose a stricter convergence criteria and we reduce the tolerance level of the system solution in Step 1 to $tol = 10^{-12}$.

Objective Function Gradient. The Jacobian is $\nabla H\left(\boldsymbol{\theta}_2\right) = 2 \cdot D\mathbf{e}(\boldsymbol{\theta}_2)'\mathbf{Z\Phi Z'}\mathbf{e}(\boldsymbol{\theta}_2)$ where $D\mathbf{e}(\boldsymbol{\theta_2}) = \left[\frac{\partial e_{ji,t}}{\partial \theta_{21}} \; \cdots \; \frac{\partial e_{ji,t}}{\partial \theta_{2L}}\right]_{ijt}$ is the stacked matrix of Jacobian vectors of the structural error from Step 5. By the envelope theorem, the Jacobian is $D\mathbf{e}(\boldsymbol{\theta_2}) = D\boldsymbol{\delta}(\boldsymbol{\theta_2})$ because $\widehat{\boldsymbol{\theta}}_{\mathbf{1}}(\boldsymbol{\theta_2})$ is obtained from the analytical minimization of the inner

[42] This adjustment procedure follows closely the suggestions in Nevo (2000).

problem restricted to a particular level of $\boldsymbol{\theta}_2$. For each importer-year, the implicit function theorem implies that

$$D\boldsymbol{\delta}_{i,t}(\boldsymbol{\theta}_2) = \begin{bmatrix} \frac{\partial \delta_{2i,t}}{\partial \theta_{21}} & \cdots & \frac{\partial \delta_{2i,t}}{\partial \theta_{2L}} \\ \vdots & \ddots & \vdots \\ \frac{\partial \delta_{Ni,t}}{\partial \theta_{21}} & \cdots & \frac{\partial \delta_{Ni,t}}{\partial \theta_{2L}} \end{bmatrix} = - \begin{bmatrix} \frac{\partial \bar{\chi}_2}{\partial \delta_{2i,t}} & \cdots & \frac{\partial \bar{\chi}_2}{\partial \delta_{Ni,t}} \\ \vdots & \ddots & \vdots \\ \frac{\partial \bar{\chi}_N}{\partial \delta_{2i,t}} & \cdots & \frac{\partial \bar{\chi}_N}{\partial \delta_{Ni,t}} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial \bar{\chi}_2}{\partial \theta_{21}} & \cdots & \frac{\partial \bar{\chi}_2}{\partial \theta_{2L}} \\ \vdots & \ddots & \vdots \\ \frac{\partial \bar{\chi}_N}{\partial \theta_{21}} & \cdots & \frac{\partial \bar{\chi}_N}{\partial \theta_{2L}} \end{bmatrix}$$

where

$$\frac{\partial \bar{\chi}_j}{\partial \delta_{li,t}} = \begin{cases} -\frac{1}{S}\sum_{s=1}^{S}(\epsilon_s)^{\sigma_\epsilon} \cdot x_{ji,t}(\alpha_s, \epsilon_s)x_{li,t}(\alpha_s, \epsilon_s) & \text{if} \quad l \neq j \\ \frac{1}{S}\sum_{s=1}^{S}(\epsilon_s)^{\sigma_\epsilon} \cdot x_{ji,t}(\alpha_s, \epsilon_s)\left(1 - x_{ji,t}(\alpha_s, \epsilon_s)\right) & \text{if} \quad l = j \end{cases}$$

$$\frac{\partial \bar{\chi}_j}{\partial \sigma_\epsilon} = \frac{1}{S}\sum_{s=1}^{S}(\ln \epsilon_s)(\epsilon_s)^{\sigma_\epsilon} \cdot x_{ji,t}(\alpha_s, \epsilon_s) \cdot \left[\delta_{ji,t} - \sum_{l=2}^{N} x_{li,t}(\alpha_s, \epsilon_s) \cdot \delta_{li,t}\right]$$

$$\frac{\partial \bar{\chi}_j}{\partial \sigma_\alpha} = \frac{1}{S}\sum_{s=1}^{S}\alpha_s \cdot x_{ji,t}(\alpha_s, \epsilon_s) \cdot \left[\kappa_i - \sum_{l=2}^{N} x_{li,t}(\alpha_s, \epsilon_s) \cdot \kappa_l\right].$$

2A.3 Sample of Countries

Table 2A.1: List of exporting countries

Abbreviation	Exporter	log(p.c. GDP) [USA=0]
AUS	Australia	-0.246
AUT	Austria	-0.249
BLX	Belgium-Luxembourg	-0.261
BRA	Brazil	-1.666
BGR	Bulgaria	-1.603
CAN	Canada	-0.211
CHN	China	-2.536
CZE	Czech Republic	-0.733
DNK	Denmark	-0.303
BAL	Estonia-Latvia	-1.475
FIN	Finland	-0.522
FRA	France	-0.398
DEU	Germany	-0.290
GRC	Greece	-0.760
HUN	Hungary	-1.121
IND	India	-3.214
IDN	Indonesia	-2.284
IRL	Ireland	-0.574
ITA	Italy	-0.332
JPN	Japan	-0.183
LTU	Lithuania	-1.526
MEX	Mexico	-1.263
NLD	Netherlands	-0.352
POL	Poland	-1.428
PRT	Portugal	-0.830
KOR	Republic of Korea	-0.823
RoW	Rest of the World	-2.286
ROU	Romania	-1.816
RUS	Russia	-0.954
SVK	Slovak Republic	-1.102
SVN	Slovenia	-0.728
ESP	Spain	-0.644
SWE	Sweden	-0.367
TWN	Taiwan	-0.584
TUR	Turkey	-1.305
GBR	United Kingdom	-0.436
USA	United States	0.000

2A.4 Counterfactual Analysis

2A.4.1 Preliminaries

In the counterfactual analysis of Section 2.7, we use the complete trade matrix for the 37 exporters listed in Table 2A.1. In order to reconcile theory and data, we incorporate trade imbalances as follows. For each country, we define $\rho_{j,t}$ as the difference between aggregate gross expenditure and aggregate gross production. We proceed under the assumption that trade imbalances remain constant at their observed level in terms of the factor price of the reference country. Here, the reference country is the United States ($j = 1$) such that its factor price is normalized to one, $\hat{w}_1 = 1$. In particular, the market clearing condition in (2.16) becomes

$$\sum_{i=1}^{N} \hat{x}_{ji,t} x_{ji,t} \left((\hat{w}_i \hat{v}_i) y_{i,t} + \rho_{i,t} \right) = (\hat{w}_j \hat{v}_j) y_{j,t}, \quad \text{for } j = 2, ..., N \tag{2.64}$$

where

$$\hat{x}_{ji,t} x_{ji,t} = \frac{1}{S} \sum_{s=1}^{S} \frac{\exp[\alpha_s \sigma_\alpha \ln \kappa_j + (\epsilon_s)^{\sigma_\epsilon} \left(\bar{\chi}_j^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2) - \bar{\epsilon} \ln(\hat{w}_j \hat{\tau}_{ji}) \right)]}{1 + \sum_{l=2}^{N} \exp\left[\alpha_s \sigma_\alpha \ln \kappa_l + (\epsilon_s)^{\sigma_\epsilon} \left(\bar{\chi}_l^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2) - \bar{\epsilon} \ln(\hat{w}_l \hat{\tau}_{li}) \right)\right]}. \tag{2.65}$$

Notice that, by construction, $\sum_{i=1}^{N} \rho_{i,t} = 0$. Thus, the solution of the system of $N-1$ equations above implies that the market clearing condition for the reference country is automatically satisfied.

2A.4.2 Algorithm

To compute the vector $\hat{\boldsymbol{w}} = \{\hat{w}_j\}_{j=2}^{N}$ that solves system (2.64), we use the same algorithm as in Alvarez and Lucas (2007).

Step 0. Initial guess: $\hat{\boldsymbol{w}}^k = [1, ..., 1]$ if $k = 0$.

Step 1. Conditional on $\hat{\boldsymbol{w}}^k$, compute $\hat{x}_{ji,t} x_{ji,t}$ according to (2.65).

Step 2. Compute the excess labor demand as

$$F_j\left(\hat{\boldsymbol{w}}^k\right) \equiv \frac{1}{y_{j,t}}\left[-(\hat{w}_j\hat{v}_j)y_{j,t} + \sum_{i=1}^{N}\hat{x}_{ji,t}x_{ji,t}\left((\hat{w}_i\hat{v}_i)y_{i,t} + \rho_{i,t}\right)\right]$$

where we divide by $y_{j,t}$ to scale excess demand by country size.

Step 3. If $\max_j \left|F_j\left(\hat{\boldsymbol{w}}^k\right)\right| < tol$, then stop the algorithm. (In practice we set $tol = 10^{-8}$ here.) Otherwise, return to Step 1 with new factor prices computed as

$$\hat{w}_j^{k+1} = \hat{w}_j^k + \mu F_j\left(\hat{\boldsymbol{w}}^k\right)$$

where μ is a positive constant. Intuitively, this updating rule increases the price of those factors with a positive excess demand.

2A.4.3 Welfare

By Proposition 3, we can compute welfare changes in any country i by solving for $e(\cdot, U_i')$. To do so, we guess that for all $\boldsymbol{\omega} \equiv \{\omega_l\}$,

$$e(\boldsymbol{\omega}, U_i') = (y_i')\frac{\exp(\int \frac{1}{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}\ln[\sum_{l=1}^{N}(\kappa_l)^{\sigma_\alpha\alpha}(\omega_l)^{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}]dF(\alpha,\epsilon))}{\exp(\int \frac{1}{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}\ln[\sum_{l=1}^{N}(\kappa_l)^{\sigma_\alpha\alpha}((\omega_{li,t})')^{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}]dF(\alpha,\epsilon))}. \tag{2.66}$$

We then check that our guess satisfies (2.18) and (2.19) if $\bar{\chi}$ satisfies (2.31). By equations (2.20) and (2.66), welfare changes must therefore satisfy

$$\Delta W_i = \frac{(y_i')/\exp(\int \frac{1}{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}\ln[\sum_{l=1}^{N}(\kappa_l)^{\sigma_\alpha\alpha}((\omega_{li,t})')^{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}]dF(\alpha,\epsilon))}{y_i/\exp(\int \frac{1}{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}\ln[\sum_{l=1}^{N}(\kappa_l)^{\sigma_\alpha\alpha}(\bar{\chi}_l^{-1}(\boldsymbol{x}_{i,t}))^{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}]dF(\alpha,\epsilon))} - 1.$$

Using the fact that $(y_i)'/y_i = \hat{w}_i$ and $(\omega_{li,t})' = \hat{w}_l\hat{\tau}_{li}\bar{\chi}_l^{-1}(\boldsymbol{x}_{i,t})$, this finally leads to

$$\Delta W_i = (\hat{w}_i)\frac{\exp(\int \frac{1}{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}\ln[\sum_{l=1}^{N}(\kappa_l)^{\sigma_\alpha\alpha}(\bar{\chi}_l^{-1}(\boldsymbol{x}_{i,t}))^{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}]dF(\alpha,\epsilon))}{\exp(\int \frac{1}{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}\ln[\sum_{l=1}^{N}(\kappa_l)^{\sigma_\alpha\alpha}(\hat{w}_l\hat{\tau}_{li}\bar{\chi}_l^{-1}(\boldsymbol{x}_{i,t}))^{-(\bar{\epsilon}\epsilon^{\sigma_\epsilon})}]dF(\alpha,\epsilon))} - 1,$$

with $\{\hat{w}_l\}$ obtained from the algorithm in Section 2A.4.2.

2A.4.4 Confidence Intervals

The confidence intervals for the counterfactual analysis are computed with the following bootstrap procedure. First, draw parameter values from the asymptotic distribution of the GMM estimator: $\boldsymbol{\theta}(b) \sim N\left(\hat{\boldsymbol{\theta}}, \widehat{AVar}(\hat{\boldsymbol{\theta}})\right)$. Second, compute $\bar{\chi}^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2(b))$ using the algorithm described in Step 1 of Section 2A.2.3. Third, compute the counterfactual exercise with $\boldsymbol{\theta}(b)$ and $\bar{\chi}^{-1}(\boldsymbol{x}_{i,t}|\boldsymbol{\theta}_2(b))$ using the algorithm described in Section 2A.4.2. Lastly, repeat these three steps for $b = 1, ..., 200$. The bootstrap confidence interval corresponds to $[EV^{(.025)}, EV^{(.975)}]$ where $EV^{(\alpha)}$ denotes the α-th quantile value of the equivalent variation obtained across the set of 200 parameter draws.

2A.4.5 Additional Results

Table 2A.2: Welfare gains from Chinese integration since 1995: all countries, 2007

	CES (standard gravity)		Mixed CES	
Exporter	Welfare Gains	95% Confidence Interval	Welfare Gains	95% Confidence Interval
Australia	0.144	(0.120, 0.176)	0.225	(0.163, 0.384)
Austria	0.058	(0.048, 0.071)	0.102	(0.069, 0.187)
Belgium-Luxembourg	0.056	(0.046, 0.069)	0.108	(0.064, 0.195)
Brazil	0.071	(0.059, 0.087)	0.058	(0.051, 0.104)
Bulgaria	0.061	(0.050, 0.075)	-0.005	(-0.050, 0.040)
Canada	0.053	(0.043, 0.065)	0.098	(0.059, 0.176)
China	1.039	(0.866, 1.268)	1.544	(1.195, 2.812)
Czech Republic	0.151	(0.124, 0.186)	0.209	(0.163, 0.374)
Denmark	0.014	(0.012, 0.018)	0.034	(0.014, 0.076)
Estonia-Latvia	0.081	(0.067, 0.100)	0.043	(0.029, 0.085)
Finland	0.100	(0.083, 0.123)	0.154	(0.109, 0.279)
France	0.030	(0.025, 0.037)	0.057	(0.037, 0.125)
Germany	0.122	(0.101, 0.150)	0.201	(0.144, 0.347)
Greece	0.004	(0.003, 0.004)	0.018	(0.004, 0.061)
Hungary	0.214	(0.177, 0.264)	0.208	(0.178, 0.352)
India	0.126	(0.104, 0.155)	0.022	(-0.064, 0.101)
Indonesia	0.026	(0.022, 0.033)	-0.061	(-0.222, -0.004)
Ireland	0.135	(0.112, 0.167)	0.150	(0.128, 0.241)
Italy	0.008	(0.007, 0.010)	0.035	(0.012, 0.089)
Japan	0.095	(0.079, 0.117)	0.186	(0.120, 0.368)
Lithuania	0.065	(0.054, 0.079)	0.022	(-0.001, 0.052)
Mexico	0.121	(0.100, 0.150)	0.099	(0.086, 0.204)
Netherlands	0.043	(0.035, 0.053)	0.068	(0.042, 0.116)
Poland	0.086	(0.071, 0.107)	0.040	(0.026, 0.096)
Portugal	0.050	(0.042, 0.060)	0.055	(0.047, 0.093)
Republic of Korea	0.298	(0.248, 0.364)	0.399	(0.311, 0.654)
Rest of the World	0.293	(0.244, 0.358)	0.105	(-0.039, 0.246)
Romania	-0.005	(-0.006, -0.004)	-0.077	(-0.215, -0.029)
Russia	0.105	(0.087, 0.129)	0.103	(0.089, 0.157)
Slovak Republic	0.116	(0.096, 0.143)	0.120	(0.101, 0.207)
Slovenia	0.012	(0.009, 0.015)	0.020	(0.010, 0.045)
Spain	0.075	(0.062, 0.092)	0.112	(0.085, 0.213)
Sweden	0.076	(0.063, 0.094)	0.113	(0.085, 0.205)
Taiwan	0.695	(0.582, 0.843)	0.946	(0.743, 1.520)
Turkey	0.024	(0.020, 0.030)	0.019	(0.016, 0.042)
United Kingdom	0.014	(0.011, 0.017)	0.022	(0.013, 0.049)
United States	0.034	(0.028, 0.043)	0.071	(0.046, 0.136)

Notes. Estimates of welfare changes (computed as the minus of the equivalent variation) from replacing China's trade costs to all other countries in 2007 at their 1995 levels. "CES (standard gravity)" and "Mixed CES" report these welfare changes obtained using the factor demand system in Panels A and C, respectively, of Table 2. 95% confidence intervals computed using the bootstrap procedure documented in Appendix 2A.4.

Chapter 3

Services Trade and Labor Markets: The Spatial Unbundling of Worker Tasks

3.1 Introduction

An important feature of the recent globalization process was the deepening of international integration in services markets. As shown in Figure 3.1, the share of services in international trade has increased in the United States during the period of rapid trade expansion of 1980-2015. The growth in the flow of services has been especially strong among US exports: the share of services in total exports was 5.6% in 1980, 10.5% in 2000, and 18.0% in 2015. Simultaneously, the United States experienced divergent trends in labor market outcomes across high- and low-wage occupations employed in different sectors of the economy — for a review, see Acemoglu and Autor (2011). Motivated by these trends, I investigate the connection between the rise in services trade and changes in labor market outcomes both between- and within-sectors.

The starting point of my analysis is a neoclassical environment in which the production of a single final consumption good combines two intermediate inputs. Total production of intermediate goods corresponds to the sum of individual outputs of all sector employees, with a worker's output depending on her use of manual and cognitive tasks in the production process. The economy is populated by a continuum of workers endowed with heterogeneous bundles of cognitive and manual tasks. There is a technology of costly task exchange that allows workers to produce with a task bundle that differs from their task bundle endowment. In the model, trade in tasks is costless within a production facility, implying that sectoral output only depends on the total amount of tasks supplied by sector employees. However, trade in tasks is costly between sectors and countries. These costs endogenously determine the extent of between-sector trade in tasks both within a country ("outsourcing") and between countries ("offshoring").

In this environment, Section 3.2 analyzes, conditional on good prices, the labor market equilibrium under the restriction on the set of feasible sectoral task allocations imposed by the bundling of task endowments at the worker-level. Whenever sectoral task intensities are sufficiently different, I show that the equilibrium exhibits different task prices in the two sectors of the economy.[1] In this case, workers self-select into sectors based on the cognitive-manual ratio of their task endowment: those with a low ratio are employed in the manual-intensive sector, and those with a high ratio are employed in the cognitive-intensive sector. Such an equilibrium is sustained with a higher relative price of cognitive tasks in the cognitive-intensive sector.

Using this simple model, the rest of my paper offers a comprehensive exploration of how changes in trading costs, either in goods or tasks, affects labor market outcomes between- and within-sectors in different countries. To this extent, Section 3.3 considers two countries, Home and Foreign, with Home being a small economy abundant in workers with a high endowment ratio of cognitive-to-manual tasks. I then analyze a move from autarky to free trade in intermediate goods, which, from Home's perspective, is equivalent to an increase in the relative price of the cognitive-intensive intermediate good. Such a shock triggers an employment expansion in the

[1] The environment is isomorphic to that in Dornbusch, Fischer, and Samuelson (1980). In this paper, task price equalization depends on the task endowment of the continuum of workers to be allocated across sectors. In Dornbusch, Fischer, and Samuelson (1980), factor price equalization depends on the factor intensity of the continuum of goods to be allocated across countries.

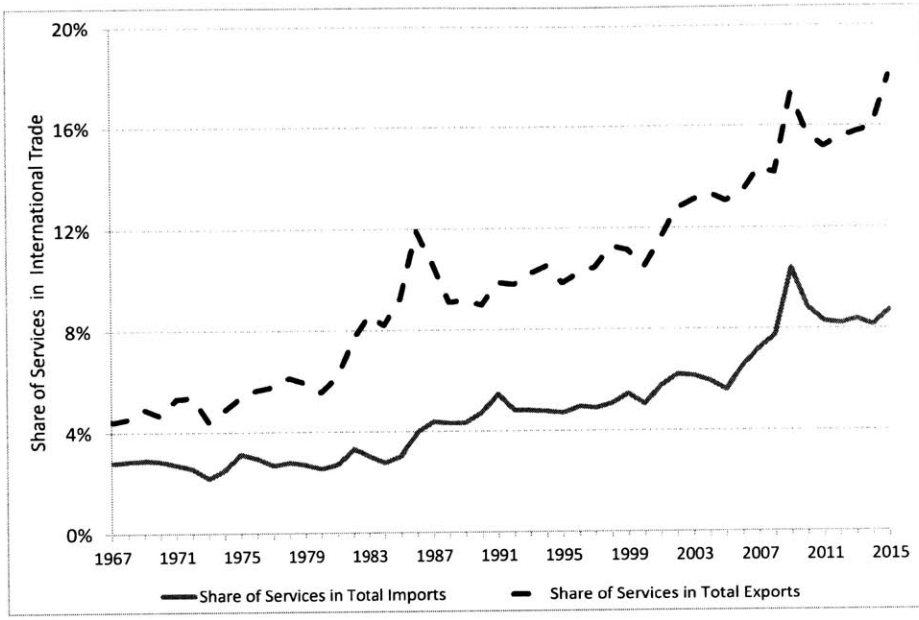

Figure 3.1: International Trade in Services, United States 1967-2015

Note. International trade of goods and services by product extracted from the National Accounts (Table 4.2.5). Services exclude transport, travel, and intellectual property rights.

cognitive-intensive sector that, due to the initial worker sorting, causes the price of cognitive-intensive tasks to rise in both sectors. Compared to Foreign, the free trade equilibrium in the Home country entails less dispersion in task prices across sectors. Specifically, the relative price of cognitive tasks is lower in Home's cognitive-intensive sector and higher in Home's manual-intensive sector.

Starting from the free trade equilibrium, Section 3.4 investigates the effect on Home's labor market of spatially unbundling the task output of its workers. The unbundling technology consists of the remote transmission of cognitive tasks at a fixed utility cost. Trade in tasks, therefore, captures the great unbundling phenomenon described by Baldwin (2006): international competition shifts from firms and sectors in different countries to individual workers performing similar tasks in different sectors and countries. This analysis proceeds in two steps.

First, I consider the labor market consequences of domestic trade in cognitive tasks. Taking advantage of the lower price of cognitive tasks in the manual-intensive sector, the cognitive-intensive sector "outsources" part of its demand for the cognitive task input. The magnitude of between-sector task transactions depends on the mass of marginal workers for whom the difference in sectoral task prices surpasses the utility cost. In this context, I show that adoption of the unbundling technology triggers between-sector convergence in task prices: the price of cognitive tasks falls in the cognitive-intensive sector and rises in the manual-intensive sector.

Second, I consider the effect of international trade in cognitive tasks. Since Foreign has a higher cognitive task price in the cognitive-intensive sector than Home, the unbundling technology induces Home workers with the highest cognitive-manual task ratio to sell their cognitive task output to producers at the Foreign country. Such a process increases between-sector task price dispersion in the Home country, pushing task prices closer to their corresponding levels in the Foreign country. This offshoring process triggers a reduction in Home's exports of the cognitive-intensive good that is compensated by the new income of high-end occupations selling their services abroad. This change in the pattern of international trade is consistent with the recent increase in the participation of services in exports observed in the United States (Figure 3.1).

My paper is mainly related to three strands of the literature. First, this paper draws on Heckman and Scheinkman's (1987) idea that workers cannot unbundle their skills. Although Heckman and Scheinkman (1987) mainly investigate task price equalization across sectors, I focus on the spatial unbundling of workers' task output as a force that endogenously generates outsourcing and offshoring. In this sense, my paper is closely related to the literature analyzing the labor market consequences of trade in tasks arising from changes in transmission costs — see Antràs, Garicano, and Rossi-Hansberg (2006), Grossman and Rossi-Hansberg (2008b, 2012), and Acemoglu and Autor (2011). While these papers study the allocation of different tasks across locations, my framework analyzes task unbundling at the worker level. In this environment, I establish that trade in tasks contributes to task price convergence in case of domestic outsourcing or task price divergence in case of international offshoring. These changes affect both between- and within-sector wage inequality.

This paper is also related to the growing empirical literature analyzing the labor market effects of offshoring — for a review, see Hummels, Munch, and Xiang (2016).

Several papers have documented that offshoring has a significant positive impact on the average wage and the skill wage premium across industries, occupations, and firms.[2] More recently, Hummels, Jørgensen, Munch, and Xiang (2014) use Danish employer-employee data to investigate the impact of offshoring on wages within each job-spell. They find that offshoring tends to increase high-skilled wages and decrease low-skilled wages. The magnitudes of these effects vary depending on workers' task specialization. My paper provides a framework to analyze the differential impact of offshoring on sector employees specialized in the production of different tasks, rationalizing the rich pattern of between- and within-sector wage responses documented in Hummels, Jørgensen, Munch, and Xiang (2014).

Finally, my paper is related to the literature studying the impact of production fragmentation on factor prices across countries. I follow Deardorff (2001) by modelling task unbundling as an innovation that allows the spatial fragmentation of the production process. In this sense, the impact of offshoring on international task price convergence is related to Helpman's (1984) result that international trade in factors increases the likelihood of factor price equalization across countries.

This chapter is organized as follows. Section 3.2 describes the labor market equilibrium conditional on the price of intermediate goods. Section 3.3 relates cross-country differences in workforce composition to cross-country differences in sectoral task prices when intermediate goods are freely traded. Section 3.4 introduces between-sector trade in tasks both within a country ("outsourcing") and between countries ("offshoring"). Section 3.5 offers some concluding remarks.

[2]Using cross-industry regressions, Feenstra and Hanson (1997) and Hsieh and Woo (2005) document a positive impact of offshoring on the relative demand for skilled workers respectively in Mexico and Hong Kong. In the United States, Amiti and Wei (2005) show that offshoring had a positive impact on productivity and a weak impact on employment across manufacturing industries. Using cross-firm exposure to tariff changes, Amiti and Davis (2012) find a positive impact of a fall in offshoring costs on the wage of employees in Indonesian import-using firms. Finally, Ebenstein, Harrison, McMillan, and Phillips (2009), Criscuolo and Garicano (2010) and Liu and Trefler (2011) analyze employment and wages in occupations differentially exposed to shocks in offshoring costs.

3.2 Model

This section presents a model where production of intermediate goods requires workers to perform manual and cognitive tasks. The novel feature of the model is the restriction regarding the cost of remotely trading tasks. In the first part of the paper, I assume that tasks can only be traded within the production facility of each intermediate good. As discussed below, this assumption implies that each sector's production function only depends on the total amount of tasks supplied by sector employees. In the second part of the paper, I introduce the possibility of between-sector trade in tasks both within a country ("outsourcing") and between countries ("offshoring").

3.2.1 Environment

Consider a world economy with two countries: the Home country, $c = H$, and the Foreign country, $c = F$. Each country is populated by L^c workers that inelastically supply one unit of labor and consume a homogeneous final good. Assume that Home is small relative to the world economy, $L^H \equiv 1 < L^F$, and, therefore, Home does not affect world prices.

Workers. Each worker is associated with an occupations u that determines her ability to produce cognitive and manual tasks. The average worker in occupation u supplies $\bar{T}_C(u) \equiv s(u) \cdot a(u)$ units of the cognitive task and $\bar{T}_M(u) \equiv \max\{(1 - s(u)) \cdot a(u);\ 0\}$ of the manual task. Thus, $s(u)$ is the occupation intensity in cognitive skills, and $a(u)$ is the average absolute advantage of workers in occupation u. Without loss of generality, assume that $s(u)$ is strictly increasing in u, and that workers are uniformly distributed across occupations.

To introduce a motivation for international trade, I allow countries to differ with respect to their occupational composition. In particular, I assume that Home is relatively abundant in cognitive-intensive occupations: $u \in U[0, \bar{u}^c]$ such that $\bar{u}^F < \bar{u}^H$. To simplify the analysis, Home's additional occupations are fully specialized in cognitive tasks: $s(\bar{u}^F) = 1$ so that $s(u) > 1$ for all $u \in (\bar{u}^F, \bar{u}^H]$.

Production. There is a single consumption good, C^c, whose production requires two intermediate inputs: $C^c = G(Q_1^c, Q_2^c)$ where $G(.)$ is increasing, concave, differentiable, and homogeneous of degree one. Production of intermediate goods combines cognitive and manual tasks supplied by sector employees. For each intermediate product $j = 1, 2$, the production function is given by

$$Q_j^c \equiv \int_{u \in \mathcal{U}_j^c} f_j \left[T_{jC}(u), T_{jM}(u) \right] du \tag{3.1}$$

where $f(.)$ is increasing, concave, differentiable, and homogeneous of degree one.

This production technology implies that the output of worker u depends on the task bundle used in production, $\left(T_{jC}(u), T_{jM}(u)\right)$. Without trade in tasks, workers have to produce with their own task endowment, reducing the environment to a standard Roy model, as in Heckman and Honoré (1990), with the productivity of worker u in sector j given by $f_j\left[\bar{T}_C(u), \bar{T}_M(u)\right]$. However, trade in tasks opens the possibility of workers achieving higher productivity by exchanging their task output in the marketplace. That is, the task bundle used in production, $\left(T_{jC}(u), T_{jM}(u)\right)$, is not necessarily equal to the worker's task endowment, $\left(\bar{T}_C(u), \bar{T}_M(u)\right)$. In light of this discussion, I now turn to the characterization of the optimal production decision whenever tasks can be freely traded among employees within each sector.

3.2.2 Sectoral Task Demand

Let us start by solving the sector's optimal allocation of tasks to workers conditional on the sector's employee set. To this extent, I assume that workers produce fixed bundle of tasks, $\left(\bar{T}_C(u), \bar{T}_M(u)\right)$, but that their task output can be freely traded within the production facility. Therefore, worker's final output $f_j\left[T_{jC}(u), T_{jM}(u)\right]$ may be produced with a different task bundle. In this environment, the optimal task allocation problem in sector j is given by

$$F_j\left(\mathcal{U}_j^c\right) \equiv \max_{\{T_{jC}(u),T_{jM}(u)\}_{u\in\mathcal{U}_j^c}} \int_{u\in\mathcal{U}_j^c} f_j\left[T_{jC}(u), T_{jM}(u)\right] du$$

such that

$$\int_{u\in\mathcal{U}_j^c} T_{jt}(u)du = \int_{u\in\mathcal{U}_j^c} \bar{T}_t(u)du.$$

Because $F_j(.)$ is differentiable and homogeneous of degree one, the first-order condition immediately implies that

$$\frac{T_{jC}(u)}{T_{jM}(u)} = \frac{\int_{u\in\mathcal{U}_j^c} \bar{T}_C(u)}{\int_{u\in\mathcal{U}_j^c} \bar{T}_M(u)} \quad \text{for all} \quad u \in \mathcal{U}_j^c.$$

Intuitively, costless task trade within each sector allows workers to exchange their task outputs in order to achieve the task intensity that maximizes their final output. Substituting this constant task ratio into the production function in (3.1), we establish the following result.

Claim 3. *Suppose trade in tasks is costless among workers within a sector. Then, the production function in sector j depends only on the aggregate amount of tasks supplied by sector employees, $\mathcal{U}_j^c$:*

$$F_j\left(\mathcal{U}_j^c\right) = f_j\left(T_{jC}, T_{jM}\right) \quad \textit{where} \quad T_{jt} \equiv \int_{u\in\mathcal{U}_j^c} \bar{T}_t(u)du \tag{3.2}$$

and $f_j(.)$ is the production function in (3.1).

As a corollary of this result, the task output of each employee has the same value for producers in sector j, allowing us to focus on the sectoral task demand. Conditional on a potentially sector-specific task price, profit maximization yields the following familiar conditions:

$$T_{jt} = a_{jt}(w_{jC}^c, w_{jM}^c)\cdot Q_j^c \quad \text{and} \quad w_{jC}^c\cdot a_{jC}(w_{jC}^c, w_{jM}^c) + w_{jM}^c\cdot a_{jM}(w_{jC}^c, w_{jM}^c) = p_j^c \tag{3.3}$$

where, in country c, w_{jt}^c is the price of task t in the intermediate sector j and p_j^c is the price of the intermediate good j.

In order to obtain analytical predictions regarding international trade flows, I assume that production of intermediate good 1 is relatively more intensive in the cognitive task compared to production of intermediate good 2. To be more precise, I impose the standard assumption that

$$\frac{a_{1C}(w)}{a_{1M}(w)} > \frac{a_{2C}(w)}{a_{2M}(w)} \quad \text{for all} \quad w \equiv (w_C, w_M). \tag{3.4}$$

Finally, profit maximization by final good producers implies that

$$Q_j^c = \alpha_j\left(p_1^c, p_2^c\right) \cdot C^c \quad \text{and} \quad p_1^c \cdot \alpha_1\left(p_1^c, p_2^c\right) + p_1^c \cdot \alpha_2\left(p_1^c, p_2^c\right) = 1, \tag{3.5}$$

where the price of the final good is normalized to one.

3.2.3 Sectoral Task Supply

Now let us turn to the characterization of the sectoral task supply in two cases: (i) equal sector task prices, and (ii) different sector-specific task prices. To this end, it is useful to define the task supply of workers in occupations $[0, u]$ as

$$T(u) \equiv (T_C(u), T_M(u)) = \left(\int_0^u s(\tilde{u}) \cdot a(\tilde{u})\, d\tilde{u}, \int_0^u ((1 - s(\tilde{u})) \cdot a(\tilde{u})\, d\tilde{u}\right) \tag{3.6}$$

where, by construction, $T_C(u)$ and $T_M(u)$ are increasing and continuously differentiable in $u \in (0, \bar{u}^c)$.

The sectoral allocation of workers is determined by the worker's optimal choice of sector of employment. Notice that, by the discussion above, worker u earns w_{jt}^c in return for each unit of task t supplied to sector j. Assuming that workers self-select into the sector yielding the highest income, worker u solves the following problem:

$$W^c(u) \equiv \max_{j \in \{1,2\}} W^c(u; j) \quad \text{s.t.} \quad W^c(u; j) \equiv \left[w_{jC}^c \cdot s(u) + (1 - s(u)) w_{jM}^c\right] \cdot a(u).$$

In this environment, consider an equilibrium with task price equalization across sectors: $w^c_{1C} = w^c_{2C}$ and $w^c_{1M} = w^c_{2M}$. In this case, workers are indifferent between the two sectors of the economy and, therefore, any sectoral allocation of workers is feasible. This implies that the task supply vector $(T_{1C}, T_{1M}, T_{2C}, T_{2M})$ is feasible if, and only if,

$$(T_{1C}, T_{1M}) + (T_{2C}, T_{2M}) = (T_C(\bar{u}^c), T_M(\bar{u}^c))$$

such that

$$(T_{1C}, T_{1M}) \in \left\{T_M \in [0, T_M(\bar{u}^c)], T_C \in [T_C(\underline{u}), T_C(\bar{u}^c) - T_C(\bar{u})] : T_M = T_M(\underline{u}) = T_M(\bar{u}^c) - T_M(\bar{u})\right\}$$

where the existence of $(\underline{u}, \bar{u})$ is guaranteed by the continuity of $(T_C(u), T_M(u))$.

Given the occupational distribution, the lens-shaped set in Figure 3.2 represents the set of feasible task supply allocation with task price equalization. Starting from O_1, the lower limit of the lens-shaped set is $T(u)$ determined by the allocation of the occupations with low cognitive intensity, $[0, u]$, to sector 1. Analogously, the upper limit is $T(\bar{u}^c) - T(u)$ determined by the allocation of the occupations with high cognitive intensity, $[u, \bar{u}^c]$, to sector 1. By changing the sectoral allocation of workers, it is possible to achieve any task supply vector in the lens-shaped set. For instance, point A can be achieved by allocating a share $O_1A'/O_1T(u^*)$ of all workers in occupations $[0, u^*]$ to sector 1, and a share $A''T(u^*)/O_2T(u^*)$ of all workers in occupations $[u^*, \bar{u}^c]$ to sector 1. Such an allocation implies that sector 1 receives the task supply vectors O_1A' and O_1A'' that, combined, deliver point A. By moving u^*, it is straight forward to verify that there are other sectoral allocation of workers that deliver the task supply in A.

Second, consider an equilibrium without task price equalization across sectors. Because of assumption (3.4), sector 1 is intensive in the cognitive task which implies that the relevant departure from task price equalization entails $\frac{w^c_{1C}}{w^c_{1M}} > \frac{w^c_{2C}}{w^c_{2M}}$. Thus, sector 1 employs workers in occupations with task intensity above $u^{c,*}$, implying that

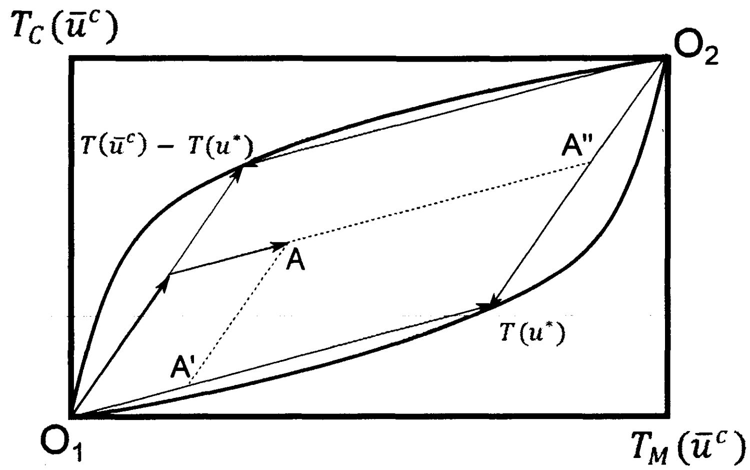

Figure 3.2: Task bundling restriction and the feasible set of sectoral task supply bundles.

sector task supply is given by

$$T_1^c = T(\bar{u}^c) - T(u^{c,*}), \quad T_2^c = T(u^{c,*}), \quad \text{and} \quad x^c \equiv \frac{s(u^{c,*})}{1 - s(u^{c,*})} = \frac{w_{2M}^c - w_{1M}^c}{w_{1C}^c - w_{2C}^c} \tag{3.7}$$

where, as discussed below, an equilibrium with positive employment in both sectors entails $w_{1C}^c > w_{2C}^c$ and $w_{1M}^c < w_{2M}^c$.[3]

3.2.4 Task Market Clearing

Any competitive equilibrium requires task market clearing with sectoral task demand described in equation (3.3) and sectoral task supply described in Section 3.2.3. In the rest of the paper, I restrict the analysis to equilibria without task price equalization across sectors. As illustrated in point E of Figure 3.3, such an equilibrium occurs whenever production of the two intermediate goods have sufficiently different task intensities. To see this, let us ignore for a moment the bundling restriction, allowing

[3]Given task prices, the selection of workers to sectors follows from the solution of the following problem:

$$J^{c,*}(u) = arg \max_{j \in \{1,2\}} W^c(u;j) = (w_{1M}^c + u(w_{1C}^c - w_{1M}^c)) \left(\frac{w_{2M}^c + u(w_{2C}^c - w_{2M}^c)}{w_{1M}^c + u(w_{1C}^c - w_{1M}^c)} \right)^{j-1} \cdot a(u).$$

Thus, $sign \frac{\partial^2 \log W^c(u;j)}{\partial j \partial u} = sign[w_{2C}^c w_{1M}^c - w_{1C}^c w_{2M}^c] < 0$ and, therefore, $J^{c,*}(u)$ is decreasing in u.

workers to freely to sell their task outputs to different employers. This would lead to an equilibrium allocation inside the area formed by the isovalue curves passing through point E — precisely, the point of tangency of the isovalue curves. In this case, the sectoral task intensities are so different that the two sectors, if faced with the same task prices, have task demands that are not feasible under the restrictions imposed by the bundling of tasks at the worker-level. In order to both sectors exhibit similar relative task demands that are feasible in equilibrium, sector 1 has to face a higher relative price of the cognitive task than sector 2. Conditional on good prices, the zero profit condition implies this can only be achieved if $w^c_{1C} > w^c_{2C}$ and $w^c_{1M} < w^c_{2M}$. Appendix 3A.1.1 provides a formal proof of this claim.

Consistent with the discussion above, I focus on the case of extreme sectoral task intensities that lead to equilibria without task price equalization.[4] In particular, I assume that the two sectors have sufficiently different task intensities such that the following condition holds in equilibrium:

$$\phi^c_{1C} \geq \frac{w^c_{1C}(w^c_{2M} - w^c_{1M})}{w^c_{1C}w^c_{2M} - w^c_{2C}w^c_{1M}} > \frac{w^c_{2C}(w^c_{2M} - w^c_{1M})}{w^c_{1C}w^c_{2M} - w^c_{2C}w^c_{1M}} \geq \phi^c_{2C} \tag{3.8}$$

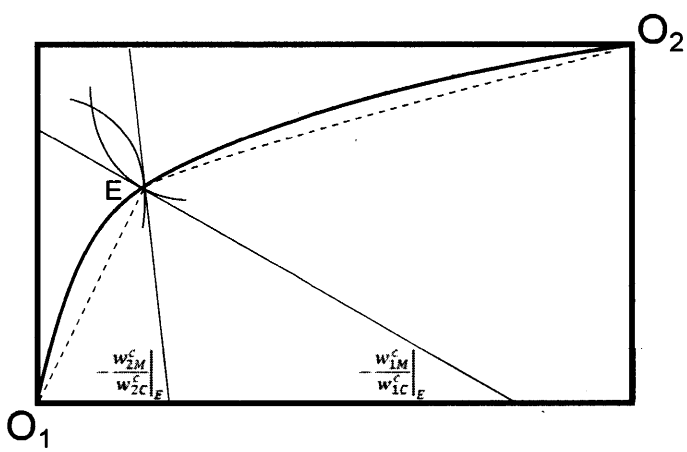

Figure 3.3: Competitite equilibrium without task price equalization.

[4] In the context of hedonic regressions, Heckman and Scheinkman (1987) and Jones (1988) analyze the fundamental conditions in the economy that guarantee task price equalization. In this paper, I abstract from a formal analysis of these conditions to focus on comparative statics in an environment without task price equalization.

where ϕ^c_{jC} is the share of cognitive tasks in the total production cost of industry j of county c.[5]

Given the price of intermediate goods, condition (3.8) guarantees that an increase in the relative price of cognitive tasks in either sector leads to an employment expansion in the manual-intensive sector. To see this, notice that the threshold in (3.7) and the zero profit condition in (3.3) yield

$$\begin{aligned} \hat{x}^c + \widehat{\frac{p^c_1}{p^c_2}} = & \left[\phi^c_{1C} - \frac{w^c_{1C}(w^c_{2M} - w^c_{1M})}{w^c_{1C}w^c_{2M} - w^c_{2C}w^c_{1M}}\right] \cdot \frac{\hat{w}^c_{1C} - \hat{w}^c_{1M}}{w^c_{1C} - w^c_{2C}} \\ & + \left[\frac{w^c_{2C}(w^c_{2M} - w^c_{1M})}{w^c_{1C}w^c_{2M} - w^c_{2C}w^c_{1M}} - \phi^c_{2C}\right] \cdot \frac{\hat{w}^c_{2C} - \hat{w}^c_{2M}}{w^c_{1C} - w^c_{2C}} \end{aligned} \tag{3.9}$$

where a hat denotes the log-change of the variable.

The equilibrium relation (3.9) plays a central role in remaining of the paper. To gain intuition for it, let us consider the effect of an increase in sector 1's relative price of cognitive tasks, $\hat{w}^c_{1C} - \hat{w}^c_{1M}$, while holding good prices constant. By the zero profit condition, such an increase is associated with an increase in sector 1's cognitive task price, $\hat{w}^c_{1C} = (1 - \phi^c_{1C})(\hat{w}^c_{1C} - \hat{w}^c_{1M})$, and a decline in sector 1's manual task price, $\hat{w}^c_{1M} = -\phi^c_{1C}(\hat{w}^c_{1C} - \hat{w}^c_{1M})$. These changes have opposite effects on the employment decision of workers in the marginal occupation. With a strong enough cognitive intensity of sector 1, the rise in w^c_{1C} is weaker than the fall in w^c_{1M}, pushing marginal workers away from sector 1. The reverse argument holds for the manual-intensive sector 2.

To complete the characterization of the equilibrium without task price equalization, consider the log-linear version of the task market clearing condition:

$$h^c \cdot \hat{u}^c + g^c_1 \cdot \hat{x}^c = -\eta^c_1 \cdot (\hat{w}^c_{1C} - \hat{w}^c_{1M}) \quad \text{and} \quad g^c_2 \cdot \hat{x}^c = -\eta^c_2 \cdot (\hat{w}^c_{2C} - \hat{w}^c_{2M}), \tag{3.10}$$

where g^c_j is the positive elasticity of the sectoral relative task supply to the sectoral

[5]The existence of equilibria satisfying condition (3.8) is exemplified by the case of Cobb-Douglas production technology. First, if $\phi_{1C} = 1$ and $\phi_{1C} = 0$, then $w^c_{1M} = 0$ and $w^c_{2C} = 0$ and condition (3.8) holds trivially. Second, if $\phi_{1C} \in (0,1)$ and $\phi_{1C} = 0$, then the condition is equivalent to $w^c_{1M} \geq \phi_{1M}$. In numerical simulations with $a(u) = 1$, I attest that this condition is satisfied for all values $\phi_{1C} \in (0,1)$.

allocation of occupations, h^c is the positive elasticity of sector 1's cognitive task supply to the country's occupational composition, and η_j^c is the elasticity of substitution between cognitive and manual tasks in sector j.

Together with equation (3.9), the market clearing condition in (3.10) determines task price responses to changes in intermediate good prices. Because sector 1 employs occupations with a relatively higher ratio of cognitive-manual task endowment, the departure of marginal workers from sector 1 triggers an increase in the relative supply of cognitive tasks in both sectors (i.e., $g_j > 0$). Thus, market clearing requires a decline in the relative price of cognitive tasks, $\hat{w}_{jC}^c - \hat{w}_{jM}^c < 0$, whenever employment in the manual-intensive sector expands, $\hat{x}^c > 0$.

3.3 Trade in Intermediate Goods

This section analyzes the effect of international trade in intermediate goods on Home's labor market if task prices are not equalized in the two sectors. Toward this end, I start by describing cross-country differences in the autarky equilibrium. Such differences determine the effect of trade opening on Home's relative good prices and, consequently, on Home's relative task prices.

3.3.1 Competitive Equilibrium: Autarky

Without international trade, domestic markets for intermediate goods have to clear with demand given by (3.5). Combining this condition with sectoral task supply in (3.7), one obtains the following condition:

$$\frac{T_M(\bar{u}^F) - T_M(u^{c,*})}{\bar{T}_M(u^{c,*})} = \frac{\alpha_1(p_1^c/p_2^c)}{\alpha_2(p_1^c/p_2^c)} \frac{a_{1M}^c(w_{1C}^c/w_{1M}^c)}{a_{2M}^c(w_{2C}^c/w_{2M}^c)}. \tag{3.11}$$

As shown in Appendix 3A.1.2, equation (3.11) commands a negative relation between employment in the cognitive-intensive sector and the relative price of the cognitive-intensive good. Intuitively, the final good sector requires a lower relative price in order to absorb the relative output produced by the additional employees in the cognitive-intensive sector. As a consequence of this relation, Proposition 4 establishes cross-country differences in the autarky equilibrium.

Proposition 4. *In the autarky equilibrium, the Home country has a lower relative price of the cognitive-intensive, p_1^c/p_2^c, than the Foreign country.*

Proof. See Appendix 3A.1.2.

To understand Proposition 4, recall that, compared to Foreign, Home is abundant in high-end occupations that self-select into the cognitive-intensive sector. Thus, given identical good prices, the two countries have the same relative good demand, but Home has a larger relative production of the cognitive-intensive good. Since this cannot be an equilibrium, the good market clearing condition in (3.11) implies that Home must have a lower relative price of the cognitive-intensive good.

3.3.2 Competitive Equilibrium: Free Trade in Intermediate Goods

Starting from the autarky equilibrium, I now analyze the labor market consequences of allowing intermediate goods to be traded internationally. By Proposition 4, Home has a lower relative price of the cognitive-intensive intermediate good in the autarky equilibrium, $p_1^H/p_2^H < p_1^F/p_2^F$. Given the simplifying assumption that Home is a small economy, its integration to the world economy does not affect the relative price of intermediate goods in the Foreign country. However, producers in the Home country are able to sell their output in the world market, causing the relative price of the cognitive-intensive intermediate good in the Home country to increase to the level observed in the Foreign country.

Using the equilibrium relations (3.9) and (3.10), one obtains the effect of the rise in the relative price of the cognitive-intensive good on Home's labor market. Such an effect is summarized in Proposition 5.

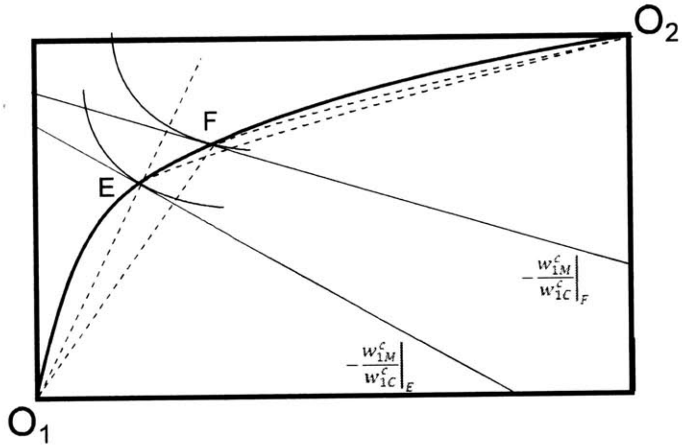

Figure 3.4: Autarky equilibrium (E) and Free trade in intermediate goods (F) – Home country.

Proposition 5. *In the Home country, the increase in the relative price of the cognitive intensive good caused by its integration to the world economy triggers (i) an employment expansion in the cognitive-intensive sector; and (ii) an increase in the relative price of the cognitive task in both sectors.*

Proof. See Appendix 3A.1.2.

Market integration triggers an increase in the relative price of the cognitive-intensive good in the Home country, creating incentives for Home producers to expand production in sector 1. In order to do so, sector 1 has to absorb part of the occupations originally employed in sector 2. As illustrated in Figure 3.4, these sector-switchers supply a lower relative output of the cognitive task than the original employees of sector 1, causing a reduction in the relative amount of the cognitive task in sector 1. Similarly, the employment reduction in sector 2 also causes a reduction in the sector's relative amount of the cognitive task, because sector-switchers supply a higher relative output of cognitive tasks than those staying in sector 2. This requires the cognitive task to become relatively more expensive in both sectors — the shallower isocost curve sustaining the equilibrium task allocation F.

Lastly, I use relations (3.9) and (3.10) to compare labor market outcomes across countries in the free trade equilibrium.

Proposition 6. *Compared to the Foreign country, the free trade equilibrium in the Home country entails (i) lower employment in the manual-intensive sector,* $u^{*,H} < u^{*,F}$*; and (ii) less dispersion in task prices,*

$$w_{2C}^{F} < w_{2C}^{H} < w_{1C}^{H} < w_{1C}^{F} \quad \textit{and} \quad w_{1M}^{F} < w_{1M}^{H} < w_{2M}^{H} < w_{2M}^{F}. \tag{3.12}$$

Proof. See Appendix 3A.1.2.

The law of comparative advantage holds in this environment: the lower autarky price of the cognitive-intensive good in the Home country implies that Home is an exporter of the cognitive-intensive good in the free trade equilibrium. In the free trade equilibrium, task prices cannot be equalized across countries. To see this, suppose they are identical. In this case, both countries have identical sectoral relative task demands, but, because of its relative abundance in high-end occupations, Home's sector 1 has a higher relative amount of the cognitive task. Since this violates task market clearing, Home must have a lower relative price of cognitive tasks in sector 1 than Foreign. Such a lower relative task price implies an increase in sector 1's relative task demand for cognitive tasks and, simultaneously, a reduction in sector 1's relative supply of cognitive tasks by reducing the employment threshold (see equation (3.9)). Finally, the lower employment threshold reduces the relative supply of cognitive tasks in Home's sector 2, requiring the relative price of cognitive tasks to be higher at Home's sector 2. Given this discussion, the task price ordering follows immediately from the zero profit condition with identical intermediate good prices in the two countries.

3.4 Trade in Cognitive Tasks

In this section, I introduce the possibility of remote transmission of cognitive tasks while restricting manual tasks to be performed at the production sight. This shock is motivated by the idea that recent innovations in information technology significantly reduced the cost of transmitting the output of cognitive tasks like data analysis and problem solving. As we shall see below, such a shock triggers domestic between-sector trade in tasks ("outsourcing") and international trade in tasks ("offshoring"). To simplify the analysis, I assume that the technology is only introduced at the Home country, so that Foreign prices of goods and tasks remain constant.

3.4.1 Domestic Trade in Cognitive Tasks

Starting from the free trade equilibrium described in Section 3.3.2, I introduce costly transmission of cognitive tasks between sectors in the Home country. Any worker can sell remotely its cognitive task output by paying a lump utility cost γ_D. This implies that workers in occupation $u \in (\underline{D}^H, \bar{D}^H)$ engage in outsourcing. That is, workers in these occupations sell their manual task output to the manual-intensive sector and their cognitive task output to the cognitive-intensive sector such that

$$(\tilde{w}_{1C}^H - \tilde{w}_{2C}^H) \cdot s(\underline{D}^H) = \gamma_D \quad \text{and} \quad (\tilde{w}_{2M}^H - \tilde{w}_{2M}^H) \cdot (1 - s(\bar{D}^H)) = \gamma_D, \tag{3.13}$$

and $\tilde{w}_{jt}^H$ is the price of task t in sector j in the equilibrium with between-sector task trade.

To get an intuition for condition (3.13), consider the incentives to engage in remote task trading. For a worker in occupation u employed in sector 2, the benefit of remotely selling one unit of the cognitive task to sector 1 is the between-sector price differential, $(\tilde{w}_{1C}^H - \tilde{w}_{2C}^H) \cdot s(u)$. Since this benefit is increasing in the workers cognitive task output, it exceeds the utility cost γ_D for those in occupations with skill intensity above $\underline{D}^H$. Similarly, workers engaging in remote task trade must be better off than if employed in sector 1. In this case, they abdicate the higher return for their manual tasks, $(\tilde{w}_{2M}^H - \tilde{w}_{2M}^H) \cdot (1 - s(u))$, but they also save on the cost γ_D. Since the benefit is decreasing in the cognitive intensity of the occupation, it exceeds the utility cost γ_D for those in occupations with skill intensity below $\bar{D}^H$.

Notice that a high enough transaction cost prevents between-sector trade in tasks. Specifically, for any $\gamma_D \geq \bar{\gamma}_D \equiv (w_{1C}^H - w_{2C}^H) \cdot s(u^{H,*}) = (w_{2M}^H - w_{2M}^H) \cdot (1 - s(u^{H,*}))$, workers do not want to remotely trade their task output in the initial free trade equilibrium. For any $\gamma_D < \bar{\gamma}_D$, there is a positive mass of workers that engage in outsourcing. They take advantage of the between-sector differential in task prices by selling their manual task output to the manual-intensive sector and their cognitive task output to the cognitive intensive sector. Proposition 7 analyzes the effect of this technology on Home's labor market outcomes for a given level of world good prices.

Proposition 7. *For any $\gamma_D \in [0, \bar{\gamma}^D)$, a reduction in domestic transmission costs triggers (i) an increase in outsourcing in the Home country; and (ii) a reduction in between-sector task price differentials such that*

$$w^H_{2C} < \tilde{w}^H_{2C} \leq \tilde{w}^H_{1C} < w^H_{1C} \quad \textit{and} \quad w^H_{1M} < \tilde{w}^H_{1M} \leq \tilde{w}^H_{2M} < w^H_{2M}.$$

Proof. See Appendix 3A.1.3 for a proof with similar sector sizes in the initial equilibrium.

The intuition behind Proposition 7 is conveyed by the introduction of costless remote transmission of cognitive tasks, $\gamma_D = 0$. From condition (3.13), it is straight forward to see that any between-sector differential in task prices leads all workers to supply manual tasks to the manual-intensive sector and cognitive tasks to the cognitive-intensive sector. Since the production technology in both sectors requires positive amounts of the two tasks, there is task price equalization across sectors. Conditional on international good prices, this can only occur with task prices in a level in between their original levels in sectors 1 and 2. To see this, consider the unitary isovalue curves represented in Figure 3.5. With costless trade in tasks, these unitary isovalue curves have to be sustained by an intermediate level of the relative task prices that is common for the two sectors. Moreover, constant good prices imply that this can only be achieved with task price convergence.

This price convergence is a consequence of the reallocation of task supply across sectors implied by workers engaging in outsourcing. That is, workers in occupations $u \in (\underline{D}^H, \bar{D}^H)$ sell their task output to different sectors. In the case of costless trade

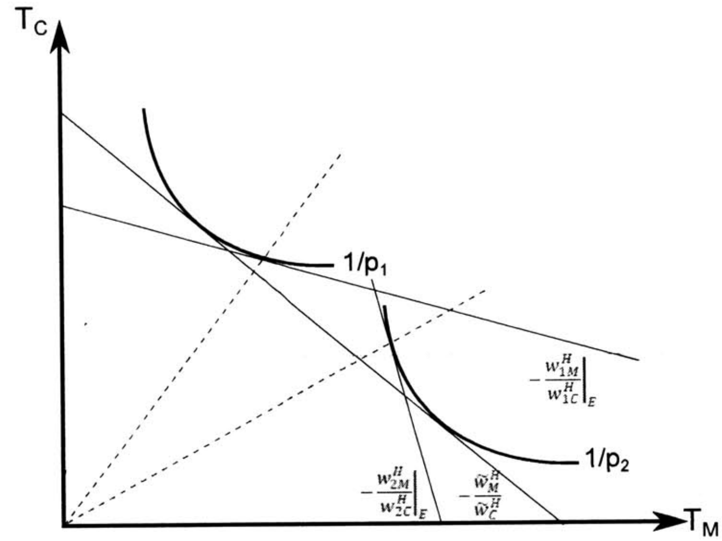

Figure 3.5: Task price convergence and Domestic trade in cognitive tasks – Home country.

$\gamma_D = 0$, the equilibrium task allocation is illustrated by point B in Figure 3.6. Because workers have to perform manual tasks in the production facility, point E represents the sectoral allocation of occupations determined by the sectoral division of manual tasks. Thus, vector EB represents the excess of cognitive tasks that sector 2 supplies to sector 1: Producers in sector 1 "outsource" part of the cognitive task output of their employees to producers in sector 2. Similarly to production fragmentation in Helpman (1984), outsourcing in this model facilitates task price equalization across sectors within a country.

3.4.2 International Trade in Cognitive Tasks

Lastly, let us analyze the labor market consequences of introducing costly international transmission of cognitive tasks in the free trade equilibrium of Section 3.3.2. Specifically, assume that workers can sell their cognitive task output internationally at a utility cost γ_I. Because of the higher price of cognitive tasks at the Foreign country (Proposition 6), Home workers in highly cognitive intensive occupations have an incentive to engage in "offshoring".

To see this, notice that, at a cost γ_I, workers can obtain a higher payment for their

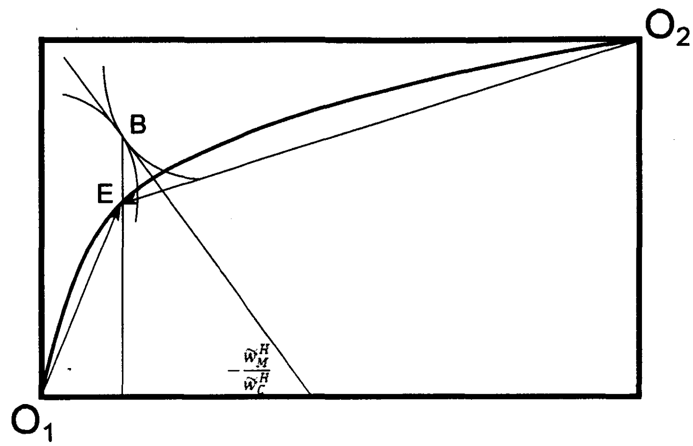

Figure 3.6: Endogenous task outsourcing and Domestic trade in cognitive tasks – Home country.

cognitive task output, $(w_{1C}^F - \tilde{w}_{1C}^H) \cdot s(u)$. As above, a high enough transaction cost prevents international trade in tasks: for any $\gamma_I \geq \bar{\gamma}_I \equiv (w_{1C}^F - \tilde{w}_{1C}^H) \cdot s(\bar{u}^H)$, all workers prefer to sell their cognitive task output for a lower price at the domestic market. For levels γ_I slightly below $\bar{\gamma}_I$, those engaging in offshoring only produce cognitive tasks and, therefore, workers in occupations $u \in (\bar{I}^H, \bar{u}^H)$ sell their cognitive task output to producers in the cognitive-intensive sector in the Foreign country. For simplicity, let us restrict the analysis to values of γ_I sufficiently close to $\bar{\gamma}_I$ such that $\bar{I}^H > \bar{u}^F$ and

$$(w_{1C}^F - \tilde{w}_{1C}^H) \cdot s(\bar{I}^H) = \gamma_I. \tag{3.14}$$

In this sense, "offshoring" is effectively a reduction in Home's supply of high-end occupations. For $\gamma_I < \bar{\gamma}_I$, international trade in tasks causes the range of occupations supplying tasks at Home to change from $[0, \bar{u}^H]$ to $[0, \bar{I}^H]$ with $\bar{u}^H > \bar{I}^H$. Proposition 8 analyzes the effect of a reduction in the cost of internationally trading tasks at a neighbourhood of the prohibitive level $\bar{\gamma}^I$.

Proposition 8. *For γ^I sufficiently close to $\bar{\gamma}^I$, a reduction in international transmission costs triggers (i) an increase in offshoring in the Home country, $\bar{I}^H < \bar{u}^H$; (ii) lower employment in the cognitive-intensive sector; and (iii) an increase in between-*

sector task price differentials,

$$\tilde{w}_{2C}^{H} < w_{2C}^{H} < w_{1C}^{H} < \tilde{w}_{1C}^{H} \quad \text{and} \quad \tilde{w}_{1M}^{H} < w_{1M}^{H} < w_{2M}^{H} < \tilde{w}_{2M}^{H}. \tag{3.15}$$

Proof. Appendix 3A.1.4.

Intuitively, offshoring generates excess relative demand for cognitive tasks in Home's sector 1. In order to eliminate this excess relative demand, sector 1 experiences an increase in the relative price of cognitive tasks accompanied by an employment contraction. In sector 2, the employment expansion increases the relative amount of the cognitive task, requiring a lower relative price of cognitive tasks. These effects translate into more disperse task prices that are closer to those observed at Foreign.

As a direct consequence of the lower employment in sector 1, offshoring reduces Home's exports of the cognitive-intensive good. This reduction is compensated by the income of high-end occupations selling their services in the Foreign country. Thus, following a reduction in the cost of internationally trading tasks, the model predicts that Home's exports move from cognitive-intensive goods to cognitive-intensive services. Such a change in the trade pattern is consistent with the recent trends for the US economy shown in Figure 3.1.

3.5 Conclusion

This paper starts from one observation: the rise in services trade in the recent globalization episode was accompanied by substantial changes in the wage structure in different countries. Motivated by these trends, I developed a theoretical framework where services trade arise from the spatial unbundling of workers' task output. A decline in the transmission cost of cognitive tasks endogenously determine the extent of between-sector trade in tasks both within a country ("outsourcing") and between countries ("offshoring"). Considering a small economy abundant in workers with high productivity in cognitive tasks, I show that outsourcing triggers a reduction in between-sector task price differentials. In contrast, by allowing workers to sell their cognitive task output at a higher price in the world market, offshoring triggers an increase in between-sector task price differentials. Finally, the task endowment heterogeneity of sector employees implies that such shocks affect both between- and within-sector wage inequality.

In this paper, I investigated the effect of unbundling workers' task output on changes in both the flow and the price of tasks across sectors and countries. Movements in wage differentials depend crucially on the correlation between workers' cognitive-to-manual endowment ratio and their overall productivity. Thus, precise predictions regarding responses in between- and within-sector wage inequality require knowledge of this correlation among employees in each sector.

Lastly, my theoretical framework provides a set of testable predictions regarding the labor market consequences of declines in the cost of trading services in a two-sector environment. However, testing these predictions entails two potential challenges. First, it is necessary to recover the unobserved sectoral task prices, which requires panel data containing information on the cognitive content of workers' task endowments. By mapping worker types to occupations, the procedure in Autor, Levy, and Murnane (2003) yields task content variables from detailed surveys on time use across occupations. Second, testing the model's predictions demand a quasi-experiment that replicates the task unbundling shock studied in Section 3.4. As in Akerman, Gaarder, and Mogstad (2015), the timing of broadband availability across regional labor markets constitutes such an experiment.

3A Appendix

3A.1 Proofs

3A.1.1 Equilibrium without task price equalization

Let us start from the equilibrium without the bundling restriction with positive production in both sectors. Given product prices, task prices are determined by the zero profit condition:

$$\tilde{w}_C^c \cdot a_{jC}(\tilde{w}_C^c, \tilde{w}_M^c) + \tilde{w}_M^c \cdot a_{jM}(\tilde{w}_C^c, \tilde{w}_M^c) = p_j^c \quad \text{for} \quad j = 1, 2.$$

In equilibrium, production is determined by the task market clearing condition:

$$a_{1t}(\tilde{w}_C^c, \tilde{w}_M^c)Q_1^c + a_{2t}(\tilde{w}_C^c, \tilde{w}_M^c)Q_2^c = \bar{T}_t(\bar{u}^c) \quad \text{for} \quad t = C, M.$$

Thus, sector 1's task allocation is $T_{1t}^c = a_{1t}(\tilde{w}_C^c, \tilde{w}_M^c) \cdot Q_1^c$. Define $\bar{u}$ as the solution of $a_{1M}(\tilde{w}_C^c, \tilde{w}_M^c)Q_1^c = T_M(\bar{u}^c) - T_M(\bar{u})$. If $a_{1C}(\tilde{w}_C, \tilde{w}_M^c)Q_1^c > T_C(\bar{u}^c) - T_C(\bar{u})$, then this equilibrium does not satisfy the bundling restriction described in Section 3.2.3. In order to satisfy this restriction, sector 1's relative demand for the cognitive task has to decrease, which requires a higher relative price of the cognitive task. By a similar argument, sector 2's relative demand for the cognitive task has to increase, which requires a lower relative price of the cognitive task. Thus,

$$\frac{w_{1C}^c}{w_{1M}^c} > \frac{\tilde{w}_C^c}{\tilde{w}_M^c} > \frac{w_{2C}^c}{w_{2M}^c}.$$

Given good prices, this inequality in terms of relative task prices implies that $w_{1C}^c > \tilde{w}_C^c > w_C^c$ and $w_{1M}^c < \tilde{w}_M^c < w_{2M}^c$.

3A.1.2 Proofs of Propositions 4, 5, and 6

Consider an equilibrium without task price equalization satisfying condition (3.8). In all proofs, $\omega_j^c \equiv w_{jC}^c/w_{jM}^c$ is the relative price of the cognitive task in sector j; $\omega_M^c \equiv w_{2M}^c/w_{1M}^c$ is the relative price of the manual task in sector 2; and $p^c \equiv p_1^c/p_2^c$ is the relative price of cognitive-intensive intermediate 1.

For any $\bar{u}^c > \bar{u}^F$, task market clearing in sectors 1 and 2 imply, respectively, that

$$\frac{T_C(\bar{u}^c) - T_C(u^{c,*})}{T_M(\bar{u}^F) - T_M(u^{c,*})} = \frac{a_{1C}(\omega_1^c)}{a_{1M}(\omega_1^c)} \quad \text{and} \quad \frac{T_C(u^{c,*})}{T_M(u^{c,*})} = \frac{a_{2C}(\omega_2^c)}{a_{2M}(\omega_2^c)} \tag{3.16}$$

where $u^{c,*}$ follows from the indifference condition of workers,

$$x^c \equiv \frac{s(u^{c,*})}{1 - s(u^{c,*})} = \frac{\omega_M^c - 1}{\omega_1^c - \omega_2^c \omega_M^c}. \tag{3.17}$$

In any equilibrium, the ratio of the zero profit conditions in the two sectors implies that

$$\frac{\omega_1^c \cdot a_{1C}(\omega_1^c) + a_{1M}(\omega_1^c)}{\omega_2^c \cdot a_{2C}(\omega_2^c) + a_{2M}(\omega_2^c)} \frac{1}{\omega_M^c} = p^c \tag{3.18}$$

Define the elasticity of sectoral relative task supply to the sectoral allocation of occupations:

$$\tilde{g}_1(u) \equiv \frac{\partial \log \frac{T_C(\bar{u}^c) - T_C(u)}{T_M(\bar{u}^F) - T_M(u)}}{\partial \log u} > 0; \quad \tilde{g}_2(u) \equiv \frac{\partial \log \frac{T_C(u)}{T_M(u)}}{\partial \log u} > 0; \quad \text{and} \quad \tilde{g}_3(u) \equiv \frac{\partial \log \frac{T_M(\bar{u}^F) - T_M(u)}{T_M(u)}}{\partial \log u} < 0.$$

The sign of these elasticities follow directly from the definition of $(T_C(u), T_M(u))$ in (3.6). Also, notice that $\tilde{g}_1(u) + \tilde{g}_3(u) < 0$ because $T_C(u)$ and $T_M(u)$ are increasing in u. Lastly, define the elasticity of sector 1's cognitive task supply with respect to the country's occupational composition:

$$h^c \equiv \frac{\partial \log(T_C(\bar{u}^c) - T_C(u^{c,*}))}{\partial \log \bar{u}^c} > 0.$$

Define $g_j^c \equiv \tilde{g}_j(u^{c,*})(1 - s(u^{c,*}))/\varepsilon_s(u^{c,*})$ where $\varepsilon_s(.) \equiv \frac{\partial \log s(u)}{\partial \log u} > 0$. Equations

(3.16)- (3.18) imply that

$$h^c \cdot \hat{u}^c + g_1^c \cdot \hat{x}^c = -\eta_1^c \cdot \hat{\omega}_1^c; \tag{3.19}$$

$$g_2^c \cdot \hat{x}^c = -\eta_2^c \cdot \hat{\omega}_2^c; \tag{3.20}$$

$$\hat{p}^c + \hat{\omega}_M^c = \phi_{1C}^c \hat{\omega}_1^c - \phi_{2C}^c \hat{\omega}_2^c; \tag{3.21}$$

$$\hat{x}^c = -\frac{w_{1C}^c}{w_{1C} - w_{2C}} \hat{\omega}_1^c + \frac{w_{2C}^c}{w_{1C} - w_{2C}} \hat{\omega}_2^c + \frac{w_{1C}^c w_{2M}^c - w_{2C}^c w_{1M}^c}{(w_{2M}^c - w_{1M}^c)(w_{1C}^c - w_{2C}^c)} \hat{\omega}_M^c \tag{3.22}$$

where $\eta_j^c \equiv -\frac{\log(a_{jC}(\omega_j^c)/a_{jM}(\omega_j^c))}{\partial \log \omega_j^c} > 0$ is the elasticity of substitution in sector j at price ω_j^c.

Substituting (3.19)-(3.21) into (3.22),

$$\frac{h^c}{\eta_1^c} \left(\Omega^c \phi_{1C}^c - w_{1C}^c\right) \hat{u}^c + \Omega^c \hat{p}^c = -\left[\frac{g_1^c}{\eta_1^c} \left(\Omega^c \phi_{1C}^c - w_{1C}^c\right) + \frac{g_2^c}{\eta_2^c} \left(w_{2C}^c - \phi_{2C}^c \Omega^c\right) + \left(w_{1C}^c - w_{2C}^c\right)\right] \hat{x}^c \tag{3.23}$$

where $\Omega^c \equiv \frac{w_{1C}^c w_{2M}^c - w_{2C}^c w_{1M}^c}{w_{2M}^c - w_{1M}^c} > 0$ in any equilibrium without task price equalization.

Under condition (3.8), the term in brackets on the right-hand side is positive, so the relative employment share in sector 2 decreases ($\hat{x}^c < 0$) if the relative price of the cognitive-intensive intermediate increases ($\hat{p}^c > 0$) or the range of cognitive intensive occupations increases ($\hat{\bar{u}}^c > 0$). Notice that equation (3.23) follows from the task market clearing conditions implied by the sectoral choice of workers and the profit maximization problem of intermediate good producers. Hence, it holds whenever prices are exogenously determined in the world market or endogenously determined in the local market.

Proof of Proposition 4

Define $\rho^c \equiv -\frac{\partial \log \alpha_1(p^c)/\alpha_2(p^c)}{\partial \log p^c} > 0$. The combination of the log-linearized version of this equation (3.11) with (3.19)-(3.21) yields

$$\left[\phi_{2C}^c g_2^c - (g_3^c + \phi_{1C}^c g_1^c)\right] \hat{x}^c = \rho \hat{p}^c + h^c \phi_{1C}^c \hat{\bar{u}}^c$$

where the term in brackets is positive since $g_1 > 0$, $g_2 > 0$, and $g_3 + g_1 < 0$.

Applying this expression into equation (3.23),

$$\hat{p}^c = -\frac{h^c}{\eta_1^c} \frac{(\Omega^c \phi_{1C}^c - w_{1C}^c) + \eta_1^c A \cdot \phi_{1C}^c}{\rho^c A + \Omega^c} \cdot \hat{\bar{u}}^c \tag{3.24}$$

where

$$A \equiv \left[\frac{g_1^c}{\eta_1^c}(\Omega^c \phi_{1C}^c - w_{1C}^c) + \frac{g_2^c}{\eta_2^c}(w_{2C}^c - \phi_{2C}^c \Omega^c) + (w_{1C}^c - w_{2C}^c)\right] / \left[\phi_{2C}^c g_2^c - (g_3^c + \phi_{1C}^c g_1^c)\right].$$

Notice that condition (3.8) guarantees that $A > 0$ and $\Omega^c \phi_{1C}^c > w_{1C}^c$. Hence, Proposition 4 follows directly from equations (3.24) whenever $\bar{u}^H > \bar{u}^F$. □

Proof of Proposition 5

Relative to the autarky equilibrium, the Home country experiences an increase in the relative price of the cognitive-intensive task, $\hat{p}^H > 0$. Thus, Part (i) of Proposition 5 follows immediately from (3.23) with $\hat{\bar{u}}^c = 0$. By equations (3.19)-(3.20), $\hat{x}^H < 0$ implies that $\hat{\omega}_j > 0$, which establishes Part (ii) of Proposition 5. □

Proof of Proposition 6

In the free trade equilibrium, one can obtain the differences in labor market outcomes between the Home and the Forein country by setting $\hat{p}^c = 0$ and $\hat{\bar{u}}^H > 0$ in equations (3.19)-(3.23). To establish Part (i), notice that equation (3.23) immediately yields $x^H < x^F$. Using equation (3.20), we immediately get that $\hat{\omega}_2^H > 0$. In addition, the combination of equations (3.23) yields (3.19):

$$\hat{\omega}_1^c = -\frac{h^c}{\eta_1^c} \frac{\frac{g_2^c}{\eta_2^c}(w_{2C}^c - \phi_{2C}^c \Omega^c) + (w_{1C}^c - w_{2C}^c)}{\left[\frac{g_1^c}{\eta_1^c}(\Omega^c \phi_{1C}^c - w_{1C}^c) + \frac{g_2^c}{\eta_2^c}(w_{2C}^c - \phi_{2C}^c \Omega^c) + (w_{1C}^c - w_{2C}^c)\right]} \cdot \hat{\bar{u}}^c < 0. \tag{3.25}$$

Thus, $\bar{u}^H > \bar{u}^F$ implies that $\omega_1^H < \omega_1^F$ and $\omega_2^H > \omega_2^F$. To compare task prices, notice that the zero profit condition implies

$$\phi^c_{jC}\hat{\omega}^c_{jC} + \hat{\omega}^c_{jM} = \hat{p}^c_j. \tag{3.26}$$

By setting $\hat{p}^c_2 = 0$ in equation (3.26), we conclude that $\hat{\omega}^c_2 > 0$ can only occur if $\hat{w}^c_{2C} > 0$ and $\hat{w}^c_{1M} < 0$. Thus, $w^H_{2C} > w^F_{2C}$ and $w^H_{2M} < w^F_{2M}$. Using the same rationale, one concludes that $w^H_{1C} < w^F_{1C}$ and $w^H_{1M} > w^F_{1M}$. □

3A.1.3 Proof of Propositions 7

Given the task supply decision of workers, the ratio between the task market clearing conditions is

$$\frac{\bar{T}_C(\bar{u}^H) - \bar{T}_C(\underline{D}^H)}{\bar{T}_M(\bar{u}^F) - \bar{T}_M(\bar{D}^H)} = \frac{a_{1C}(\omega^c_1)}{a_{1M}(\omega^c_1)} \quad \text{and} \quad \frac{\bar{T}_C(\underline{D}^H)}{\bar{T}_M(\bar{D}^H)} = \frac{a_{2C}(\omega^c_2)}{a_{2M}(\omega^c_2)}. \tag{3.27}$$

The equilibrium is characterized by equation (3.27) together with the zero profit conditions in (3.3), and the thresholds in (3.13). The log-linearization of equation (3.27) implies that

$$-h_{1C}(\underline{D}^H) \cdot \hat{\underline{D}}^H + h_{1M}(\bar{D}^H) \cdot \hat{\bar{D}}^H = -\eta^c_1 \left(\hat{w}^c_{1C} - \hat{w}^c_{1M}\right) \tag{3.28}$$

$$h_{2C}(\underline{D}^H) \cdot \hat{\underline{D}}^H - h_{2M}(\bar{D}^H) \cdot \hat{\bar{D}}^H = -\eta^c_2 \left(\hat{w}^c_{2C} - \hat{w}^c_{2M}\right), \tag{3.29}$$

$g_{1C} \equiv \frac{\partial \log \int_u^{\bar{u}^c} s(u)a(u)du}{\partial \log \bar{u}^c}$, $h_{1C}(u) \equiv -\frac{\partial \log \int_u^{\bar{u}^c} s(u)a(u)du}{\partial \log u}$, $h_{1M}(u) \equiv -\frac{\partial \log \int_u^{\bar{u}^F} (1-s(u))a(u)du}{\partial \log u}$, $h_{2C}(u) \equiv \frac{\partial \log \int_0^u s(u)a(u)du}{\partial \log u}$, and $h_{2M}(u) \equiv \frac{\partial \log \int_0^u (1-s(u))a(u)du}{\partial \log u}$.

Combining these relations with the log-linear versions of (3.3) and (3.13), one obtains the following equilibrium relations:

$$\left[\frac{h_{1C}}{\varepsilon_C}\frac{\phi^c_{1M}}{\phi^c_{1C}}\frac{w^c_{1C}}{w^c_{1C} - w^c_{2C}} + \frac{h_{1M}}{\varepsilon_M}\frac{w_{1M}}{w^c_{2M} - w^c_{1M}} + \frac{\eta^c_1}{\phi^c_{1C}}\right]\hat{w}^c_{1M}$$
$$-\left[\frac{h_{1C}}{\varepsilon_C}\frac{\phi^c_{2M}}{\phi^c_{2C}}\frac{w^c_{2C}}{w^c_{1C} - w^c_{2C}} + \frac{h_{1M}}{\varepsilon_M}\frac{w^c_{2M}}{w^c_{2M} - w^c_{1M}}\right]\hat{w}^c_{2M} = -\left(\frac{h_{1C}}{\varepsilon_C} + \frac{h_{1M}}{\varepsilon_M}\right)\hat{\gamma}_D$$

$$\left[\frac{h_{2C}}{\varepsilon_L}\frac{\phi_{1M}^c}{\phi_{1C}^c}\frac{w_{1C}^c}{w_{1C}^c - w_{2C}^c} + \frac{h_{2M}}{\varepsilon_H}\frac{w_{1M}}{w_{2M}^c - w_{1M}^c}\right]\hat{w}_{1M}^c$$
$$-\left[\frac{h_{2C}}{\varepsilon_L}\frac{\phi_{2M}^c}{\phi_{2C}^c}\frac{w_{2C}^c}{w_{1C}^c - w_{2C}^c} + \frac{h_{2M}}{\varepsilon_H}\frac{w_{2M}^c}{w_{2M}^c - w_{1M}^c} + \frac{\eta_2^c}{\phi_{2C}^c}\right]\hat{w}_{2M}^c = -\left(\frac{h_{2C}}{\varepsilon_L} + \frac{h_{2M}}{\varepsilon_H}\right)\hat{\gamma}_D$$

where $\varepsilon_C(u) \equiv \frac{\partial \log s(u)}{\partial u} > 0$ and $\varepsilon_M(u) \equiv -\frac{\partial \log(1-s(u))}{\partial u} > 0$.

Let us write this system as

$$A_1\hat{w}_{1M} - A_2\hat{w}_{2M}^c = -a_1\hat{\gamma}$$
$$B_1\hat{w}_{1M} - B_2\hat{w}_{2M}^c = -b_1\hat{\gamma}$$

Thus,

$$\hat{w}_{1M}^c = \left(\frac{b_1A_2 - a_1B_2}{A_1B_2 - B_1A_2}\right)\hat{\gamma}_D; \tag{3.30}$$

$$\hat{w}_{2M}^c = \left(\frac{b_1A_1 - a_1B_1}{A_1B_2 - B_1A_2}\right)\hat{\gamma}_D \tag{3.31}$$

To get the sign on these coefficients, notice that, by definition, $h_{1t}T_{1t} = h_{2t}T_{2t}$. If sectors have similar sizes initially, then $A_1 > B_1$, $A_2 < B_2$, and $a_1 \approx b_1$. Proposition 7 follows from the fact that, by equations (3.30)-(3.31), $\hat{w}_{1M}^H > 0$ and $\hat{w}_{2M}^H < 0$ if $\hat{\gamma}_D < 0$. □

3A.1.4 Proof of Propositions 8

Denoting $\varepsilon^H \equiv \frac{d\log s(u)}{dlogu}\big|_{u=\bar{I}^H}$, the log-linearization of (3.14) implies

$$\varepsilon^H\hat{\bar{I}}^H - \frac{w_{1C}^H\phi_{1M}^H}{w_{1C}^F - w_{1C}^H}\hat{\omega}_1^H = \hat{\gamma}_I.$$

Since all steps leading to equation (3.25) are still valid, we can write

$$\left(\varepsilon^H + \frac{w_{1C}^H\phi_{1M}^H}{w_{1C}^F - w_{1C}^H}\frac{h^H}{\eta_1^H}\frac{\frac{g_2^H}{\eta_2^H}\left(w_{2C}^H - \phi_{2C}^H\Omega^H\right) + \left(w_{1C}^H - w_{2C}^H\right)}{\left[\frac{g_1^H}{\eta_1^H}\left(\Omega^H\phi_{1C}^H - w_{1C}^H\right) + \frac{g_2^H}{\eta_2^H}\left(w_{2C}^H - \phi_{2C}^H\Omega^H\right) + \left(w_{1C}^H - w_{2C}^H\right)\right]}\right)\hat{I}^H = \hat{\gamma}_I$$

Thus, $\hat{\gamma}^I < 0$ implies $\hat{I}^H < 0$. Proposition 8 follows from the expressions in Appendix 3A.1.2. □

Bibliography

D. Acemoglu and D. H. Autor. Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4:1043–1171, 2011.

A. Akerman, I. Gaarder, and M. Mogstad. The skill complementarity of broadband internet, 2015.

T. Allen, C. Arkolakis, and Y. Takahashi. Universal gravity. *Northwestern working paper*, 2014.

F. Alvarez and R. E. Lucas. General equilibrium analysis of the eaton-kortum model of international trade. *Journal of Monetary Economics*, 54(6):1726–1768, 2007.

M. Amiti and D. R. Davis. Trade, firms, and wages: Theory and evidence. *The Review of economic studies*, 79(1):1–36, 2012.

M. Amiti and S.-J. Wei. *Service Offshoring, Productivity, and Employment: Evidence from the United States (EPub)*. Number 5-238. International Monetary Fund, 2005.

J. E. Anderson and E. Van Wincoop. Gravity with gravitas: A solution to the border puzzle. *The American Economic Review*, 93(1):170–192, 2003.

J. E. Anderson and E. Van Wincoop. Trade costs. *Journal of Economic Literature*, 42(3):691–751, 2004.

P. Antràs, L. Garicano, and E. Rossi-Hansberg. Offshoring in a knowledge economy*. *The Quarterly journal of economics*, 121(1):31–77, 2006.

C. Arkolakis, A. Costinot, D. Donaldson, and A. Rodriguez-Clare. The elusive pro-competitive effects of trade. *mimeo*, 2012a.

C. Arkolakis, A. Costinot, and A. Rodríguez-Clare. New trade models, same old gains? *American Economic Review*, 102(1):94–130, 2012b.

K. Arrow and F. Hahn. *General competitive analysis*. San Francisco: Holden-Day, 1971.

K. Arrow, H. Chenery, B. Minhas, and R. Solow. Capital-labor substitution and economic efficiency. *Review of Economics and Statistics*, 3:225–250, 1961.

E. Artuç, S. Chaudhuri, and J. McLaren. Trade Shocks and Labor Adjustment: A Structural Empirical Approach. *The American Economic Review*, pages 1008–1045, 2010.

D. H. Autor, F. Levy, and R. J. Murnane. The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, pages 1279–1333, 2003.

D. H. Autor, L. F. Katz, and M. S. Kearney. Trends in US wage inequality: Revising the revisionists. *The Review of Economics and Statistics*, 90(2):300–323, 2008.

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

D. H. Autor, D. Dorn, G. H. Hanson, and J. Song. Trade Adjustment: Worker-Level Evidence. *The Quarterly Journal of Economics*, 129(4):1799–1860, 2014.

R. Baldwin. Globalisation: the great unbundling (s). *Economic Council of Finland*, 20(3):5–47, 2006.

R. E. Baldwin and A. J. Venables. Regional economic integration. In G. M. Grossman and K. Rogoff, editors, *Handbook of International Economics*, volume 3. New York: Elsevier, 1995.

M. Bas, T. Mayer, and M. Thoenig. From micro to macro: Demand, supply, and heterogeneity in the trade elasticity. *mimeo*, 2015.

E. Berman, J. Bound, and S. Machin. Implications of Skill-Biased Technological Change: International Evidence. *The Quarterly Journal of Economics*, 113(4): 1245–1279, 1998.

A. B. Bernard and J. B. Jensen. Exceptional exporter performance: Cause, effect, or both. *Journal of International Economics*, 47:1–25, 1999.

S. Berry. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262, 1994.

S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

S. T. Berry and P. A. Haile. Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797, 2014.

S. T. Berry, A. Gandhi, and P. A. Haile. Connected substitutes and invertibility of demand. *Econometrica*, 81(5):2087–2111, 2013.

G. J. Borjas. Self-Selection and the Earnings of Immigrants. *The American Economic Review*, 77(4):531–553, 1987.

H. P. Bowen, E. E. Leamer, and L. Sveikauskas. Multicountry, multifactor tests of the factor abundance theory. *American Economic Review*, 77(5):791–809, 1987.

T. F. Bresnahan and R. J. Gordon. *The economics of new goods*. Chicago: University of Chicago Press, 2008.

A. Burstein and G. Gopinath. International prices and exchange rates. In G. Gopinath, E. Helpman, and K. Rogoff, editors, *Handbook of International Economics*, volume 4. New York: Elsevier, 2013.

A. Burstein and J. Vogel. Globalization, technology, and the skill premium. 2010. Manuscript, Columbia University and UCLA.

A. Burstein and J. Vogel. International trade, technology, and the skill premium. *Manuscript, Columbia University and UCLA*, 2015.

A. Burstein, J. C. Neves, and S. Rebelo. Distribution costs and real exchange-rate dynamics during exchange-rate-based stabilizations. *Journal of Monetary Economics*, 50(6):1189–1214, 2003.

A. Burstein, E. Morales, and J. Vogel. Accounting for Changes in Between-Group Inequality. Technical report, Working Paper, 2015.

L. Caliendo and F. Parro. Estimates of the trade and welfare effects of NAFTA. *Review of Economic Studies*, 82(1):1–44, 2015.

L. Caliendo, F. Parro, E. Rossi-Hansberg, and P.-D. Sarte. The impact of regional and sectoral productivity changes on the us economy. Technical report, National Bureau of Economic Research, 2014.

L. Caliendo, M. Dvorkin, and F. Parro. The Impact of Trade on Labor Market Dynamics. Technical report, National Bureau of Economic Research, 2015.

A. C. Cameron and D. L. Miller. Robust inference with clustered data. UC Davis working paper, 2010.

I. A. Canay, A. Santos, and A. M. Shaikh. On the testability of identification in some nonparametric models with endogeneity. *Econometrica*, 81(6):2535–2559, 2013.

P. Chiappori and I. Ekeland. *The Economics and Mathematics of Aggregation: Formal Models of Efficient Group Behavior*. Boston: Now, 2009.

F. Costa, J. Garred, and J. P. Pessoa. Winners and Losers from a Commodities-for-Manufactures Trade Boom. Technical report, Centre for Economic Performance, LSE, 2014.

A. Costinot and A. Rodríguez-Clare. Trade theory with numbers: Quantifying the consequences of globalization. In G. Gopinath, E. Helpman, and K. Rogoff, editors, *Handbook of International Economics*, volume 4. New York: Elsevier, 2013.

A. Costinot and A. Rodríguez-Clare. Trade theory with numbers: Quantifying the consequences of globalization. In *In: Helpman, E.(Ed.), Handbook of international economics*. Citeseer, 2013.

A. Costinot and J. Vogel. Beyond Ricardo: Assignment Models in International Trade. Technical report, National Bureau of Economic Research, 2014.

A. Costinot and J. E. Vogel. Matching and Inequality in the World Economy. *Journal of Political Economy*, 118(4):747–786, 2010.

A. Costinot, D. Donaldson, and I. Komunjer. What goods do countries trade? a quantitative exploration of ricardo's ideas. *Review of Economic Studies*, 79(2): 581–608, 2012a.

A. Costinot, D. Donaldson, and C. Smith. Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from a 9 million-field partition of the earth. *Journal of Political Economy, Forthcoming*, 2012b.

A. Costinot, D. Donaldson, and C. Smith. Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from a 1.7 million fields around the world. *Journal of Political Economy*. In press, 2015.

C. Criscuolo and L. Garicano. Offshoring and wage inequality: Using occupational licensing as a shifter of offshoring costs. *The American Economic Review*, 100(2): 439–443, 2010.

D. R. Davis and D. Weinstein. An account of global factor trade. *American Economic Review*, 91(5):1423–1453, 2001.

A. V. Deardorff. Fragmentation in simple trade models. *The North American Journal of Economics and Finance*, 12(2):121–137, 2001.

A. V. Deardorff and R. W. Staiger. An interpretation of the factor content of trade. *Journal of International Economics*, 24(1):93–107, 1988.

R. Dekle, J. Eaton, and S. Kortum. Unbalanced trade. *The American Economic Review*, pages 351–355, 2007.

R. Dekle, J. Eaton, and S. Kortum. Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers*, 55(3):511–540, 2008.

R. Dix-Carneiro. Trade liberalization and labor market dynamics. *Econometrica*, 82 (3):825–885, 2014.

R. Dix-Carneiro and B. K. Kovak. Trade reform and regional dynamics: evidence from 25 years of brazilian matched employer-employee data. Technical report, National Bureau of Economic Research, 2015a.

R. Dix-Carneiro and B. K. Kovak. Trade liberalization and the skill premium: A local labor markets approach. *American Economic Review*, 105(5):551–57, 2015b.

A. K. Dixit and V. Norman. *Theory of International Trade*. Cambridge University Press, 1980.

R. Dornbusch, S. Fischer, and P. A. Samuelson. Heckscher-ohlin trade theory with a continuum of goods. *The Quarterly Journal of Economics*, pages 203–224, 1980.

J. Eaton and S. Kortum. Technology, geography and trade. *Econometrica*, 70(5): 1741–1779, 2002.

A. Ebenstein, A. Harrison, M. McMillan, and S. Phillips. Estimating the impact of trade and offshoring on American workers using the current population surveys. *Review of Economics and Statistics*, (0), 2009.

P. Fajgelbaum and A. Khandelwal. Measuring the unequal gains from trade. *UCLA working paper*, 2014.

R. C. Feenstra and G. H. Hanson. Foreign direct investment and relative wages: Evidence from Mexico's maquiladoras. *Journal of international economics*, 42(3): 371–393, 1997.

F. H. Ferreira, S. Firpo, and J. Messina. A more level playing field? Explaining the decline in earnings inequality in Brazil, 1995-2012. 2014.

E. French and C. Taber. Identification of Models of the Labor Market. *Handbook of Labor Economics*, 4:537–617, 2011.

J. A. Frias, D. S. Kaplan, and E. Verhoogen. Exports and within-plant wage distributions: Evidence from mexico. *American Economic Review*, 102(3):435–40, 2012.

S. Galle, A. Rodriguez-Clare, and M. Yi. Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade. *UC Berkeley*, mimeo, 2015.

N. Gaston and D. Trefler. The labour market consequences of the Canada-US Free Trade Agreement. *Canadian Journal of Economics*, 30(1):18–41, 1997.

P. K. Goldberg. Product differentiation and oligopoly in international markets: The case of the u.s. automobile industry. *Econometrica*, 63(4):891–951, 1995.

P. K. Goldberg and N. Pavcnik. Distributional Effects of Globalization in Developing Countries. *Journal of Economic Literature*, 45(1):39–82, 2007.

G. M. Grossman. Partially mobile capital: A general approach to two-sector trade theory. *Journal of International Economics*, 15(1):1–17, 1983.

G. M. Grossman. Heterogeneous workers and international trade. *Review of World Economics*, 149(2):211–245, 2013.

G. M. Grossman and E. Rossi-Hansberg. Trading tasks: A simple theory of offshoring. *American Economic Review*, 98(5):1978–1997, 2008a.

G. M. Grossman and E. Rossi-Hansberg. Trading tasks: A simple theory of offshoring. *American Economic Review*, 98(5):1978–1997, 2008b.

G. M. Grossman and E. Rossi-Hansberg. Task trade between similar countries. *Econometrica*, 80(2):593–629, 2012.

G. H. Hanson. The Rise of Middle Kingdoms: Emerging Economies in Global Trade. *The Journal of Economic Perspectives*, 26(2):41–63, 2012.

G. H. Hanson and R. Robertson. China and the manufacturing exports of other developing countries. In R. C. Feenstra and S.-J. Wei, editors, *China's Growing Role in World Trade*. University of Chicago Press, 2010.

G. H. Hanson, N. Lind, and M.-A. Muendler. The dynamics of comparative advantage. *mimeo UCSD*, 2014.

J. Hausman. Exact consumer's surplus and deadweight loss. *American Economic Review*, 71:662–676, 1981.

J. Hausman and W. Newey. Nonparametric estimation of exact consumers surplus and deadweight loss. *Econometrica*, 63(6):1445–1476, 1995.

K. Head and T. Mayer. Gravity equations: Toolkit, cookbook, workhorse. In G. Gopinath, E. Helpman, and K. Rogoff, editors, *Handbook of International Economics*, volume 4. New York: Elsevier, 2013.

K. Head and J. Ries. Increasing returns versus national product differentiation as an explanation for the pattern of u.s.-canada trade. *American Economic Review*, 91 (4):858–876, 2001.

J. Heckman and J. Scheinkman. The importance of bundling in a Gorman-Lancaster model of earnings. *The Review of Economic Studies*, 54(2):243–255, 1987.

J. J. Heckman and B. E. Honoré. The Empirical Content of the Roy Model. *Econometrica*, pages 1121–1149, 1990.

J. J. Heckman and G. Sedlacek. Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market. *The Journal of Political Economy*, 93(6):1077, 1985.

E. Helpman. Solutions of general equilibrium problems for a trading world. *Econometrica*, 44(3):547–559, 1976.

E. Helpman. A simple theory of international trade with multinational corporations. *The Journal of Political Economy*, pages 451–471, 1984.

E. Helpman, O. Itskhoki, and S. Redding. Inequality and unemployment in a global economy. *Econometrica*, 78(4):1239–1283, 2010.

E. Helpman, O. Itskhoki, M.-A. Muendler, and S. J. Redding. Trade and inequality: From theory to estimation. Technical report, National Bureau of Economic Research, 2015.

T. Hertel. Global applied general equilibrium analysis using the global trade analysis project framework. In P. B. Dixon and D. W. Jorgenson, editors, *Handbook of Computable General Equilibrium Modeling*, volume 1. New York: Elsevier, 2013.

T. W. Hertel, R. A. McDougall, B. G. Narayanan, and A. H. Aguiar. Behavioral parameters. In B. G. Narayanan, A. H. Aguiar, and R. A. McDougall, editors, *Global Trade, Assistance, and Production: The GTAP 8 Data Base*. Center for Global Trade Analysis, Purdue University, 2012.

P. Howitt. Gross substitutability with multi-valued excess demand functions. *Econometrica*, 48(6):1567–1573, 1980.

C.-T. Hsieh and R. Ossa. A global view of productivity growth in china and india. *NBER Working Paper*, 2011.

C.-T. Hsieh and K. T. Woo. The impact of outsourcing to china on hong kong's labor market. *The American Economic Review*, 95(5):1673–1687, 2005.

C.-T. Hsieh, E. Hurst, C. I. Jones, and P. J. Klenow. The Allocation of Talent and US Economic Growth. Technical report, National Bureau of Economic Research, 2013.

D. Hummels, R. Jørgensen, J. Munch, and C. Xiang. The wage effects of offshoring: Evidence from danish matched worker-firm data. *American Economic Review*, 104 (6):1597–1629, 2014.

D. Hummels, J. R. Munch, and C. Xiang. Offshoring and labor markets. Technical report, National Bureau of Economic Research, 2016.

D. S. Jacks, C. M. Meissner, and D. Novy. Trade booms, trade busts, and trade costs. *Journal of International Economics*, 83(2):185–201, 2011.

R. C. Johnson and G. Noguera. Accounting for intermediates: Production sharing and trade in value added. *Journal of International Economics*, 86(2):224–236, 2012.

L. E. Jones. The characteristics model, hedonic prices, and the clientele effect. *The Journal of Political Economy*, pages 551–567, 1988.

R. W. Jones. The structure of simple general equilibrium models. *The Journal of Political Economy*, pages 557–572, 1965.

R. W. Jones. Income distribution and effective protection in a multicommodity trade model. *Journal of Economic Theory*, 11(1):1–15, 1975.

G. Kambourov. Labour market regulations and the sectoral reallocation of workers: The case of trade reforms. *The Review of Economic Studies*, 76(4):1321–1358, 2009.

L. F. Katz and K. M. Murphy. Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78, 1992.

B. K. Kovak. Regional Effects of Trade Reform: What is the Correct Measure of Liberalization? *The American Economic Review*, 103(5):1960–1976, 2013.

P. Krugman. Technology, trade and factor prices. *Journal of International Economics*, 50:51–71, 2000.

D. Lagakos and M. E. Waugh. Selection, Agriculture, and Cross-Country Productivity Differences. *The American Economic Review*, 103(2):948–980, 2013.

R. Z. Lawrence and M. J. Slaughter. International trade and American wages in the 1980s: giant sucking sound or small hiccup? *Brookings papers on economic activity. Microeconomics*, pages 161–226, 1993.

E. E. Leamer. What's the use of factor contents? *Journal of International Economics*, 50:17–49, 2000.

E. E. Leamer. Tantalus on the road to asymptopia. *Journal of Economic Perspectives*, 24(2):31–46, 2010.

W. Leontief. Domestic production and foreign trade; the american capital position re-examined. *Proceedings of the American philosophical Society*, pages 332–349, 1953.

A. A. Levchenko and J. Zhang. The evolution of comparative advantage: Measurement and welfare implications. Michigan working paper, 2011.

R. Liu and D. Trefler. A sorted tale of globalization: white collar jobs and the rise of service offshoring. Technical report, National Bureau of Economic Research, 2011.

C. F. Manski. *Partial Identification of Probability Distributions*. New York: Springer-Verlag, 2003.

A. Mas-Colell. On the uniqueness of equilibrium once again. In W. Barnett, B. Cornet, C. d'Aspremont, J. Gabszewicz, and A. Mas-Colell, editors, *Proceedings of the sixth international symposium in economic theory and econometrics*. Cambridge: Cambridge University Press, 1991.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

K. Matsuyama. Beyond icebergs: Towards a theory of biased globalization. *Review of Economic Studies*, 74:237–253, 2007.

R. L. Matzkin. Nonparametric Identification. *Handbook of Econometrics*, 6:5307–5368, 2007.

J. Meade. *A Geometry of International Trade.* London: George Allen and Unwin, 1952.

N. A. Menezes-Filho and M.-A. Muendler. Labor reallocation in response to trade reform. Technical report, National Bureau of Economic Research, 2011.

P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

C. B. Mulligan and Y. Rubinstein. Selection, investment, and women's relative wages over time. *The Quarterly Journal of Economics*, 123(3):1061–1110, 2008.

M. Mussa. Imperfect factor mobility and the distribution of income. *Journal of International Economics*, 12(1):125–141, 1982.

P. Neary and A. Schweinberger. Factor content functions and the theory of international trade. *Review of Economic Studies*, 53(3):421–432, 1986.

A. Nevo. A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics and Management Strategy*, 9(4):513–548, 2000.

A. Nevo. Empirical Models of Consumer Behavior. *Annual Review of Economics*, 3 (1):51–75, 09 2011.

W. Newey and J. Powell. Instrumental variable estimation in nonparametric models. *Econometrica*, 71(5):1565–1578, 2003a.

W. K. Newey. Nonparametric instrumental variables estimation. *MIT working paper*, 2013a.

W. K. Newey. Nonparametric Instrumental Variables Estimation. *The American Economic Review*, 103(3):550–556, 2013b.

W. K. Newey and J. L. Powell. Instrumental Variable Estimation of Nonparametric Models. *Econometrica*, 71(5):1565–1578, 2003b.

D. Novy. International trade without CES: Estimating translog gravity. *Journal of International Economics*, 89(2):271–282, 2013.

E. Oberfield and D. Raval. Micro data and macro technology. Princeton working paper, 2014.

F. Ohnsorge and D. Trefler. Sorting It Out: International Trade with Heterogeneous Workers. *Journal of Political Economy*, 115(5):868–892, 2007.

R. A. Pollak, R. C. Sickles, and T. J. Wales. The ces-translog: Specification and estimation of a new cost function. *Review of Economics and Statistics*, 66(4):602–607, 1984.

T. Rader. *Theory of General Economic Equilibrium.* Academic Press, 1972.

N. Ramondo and A. Rodríguez-Clare. Trade, multinational production, and the gains from openness. *Journal of Political Economy*, 121(2):273–322, 2013.

A. L. Revenga. Exporting jobs? The impact of import competition on employment and wages in US manufacturing. *The Quarterly Journal of Economics*, 107(1): 255–284, 1992.

A. D. Roy. Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, 3(2):pp. 135–146, 1951.

H. E. Scarf and C. A. Wilson. Uniqueness of equilibrium in the multi-country ricardo model. In T. J. Kehoe, T. N. Srinivasan, and J. Whalley, editors, *Frontiers in Applied General Equilibrium Modeling*. Cambridge: Cambridge University Press-Cambridge University Press, 2005.

P. K. Schott. Across-product versus within-product specialization in international trade. *Quarterly Journal of Economics*, 119(2):647–678, 2004.

J. Shapiro. Trade, co2, and the environment. *Yale working paper*, 2012.

J. H. W. Stock, James H. and M. Yogo. A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business and Economic Statistics*, 20(4), 2002.

W. Stolper and P. A. Samuelson. Protection and Real Wages. *Review of Economic Studies*, 9(1):58–73, 1941.

F. Taylor. *The guidance of production in a socialist state*. Minneapolis: University of Minnesota Press, 1938.

P. Topalova. Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India. *American Economic Journal: Applied Economics*, 2 (4):1–41, 2010.

D. Trefler. International factor price differences: Leontief was right! *Journal of Political Economy*, 101(6):961–987, 1993.

D. Trefler. The case of the missing trade and other mysteries. *American Economic Review*, 85(5):1029–46, 1995.

E. A. Verhoogen. Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector. *The Quarterly Journal of Economics*, 123(2):489–530, 2008.

R. Wacziarg and J. S. Wallack. Trade liberalization and intersectoral labor movements. *Journal of International Economics*, 64(2):411–439, 2004.

C. A. Wilson. On the general structure of ricardian models with a continuum of goods: Applications to growth, tariff theory, and technical change. *Econometrica*, 48(7):1675–1702, 1980.

A. Woodland. Direct and indirect trade utility functions. *Review of Economic Studies*, 47:907–926, 1980.

K.-M. Yi. Can vertical specialization explain the growth of world trade? *Journal of Political Economy*, 111(1):52–102, 2003.

A. Young. Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services. *The American Economic Review*, 104(11):3635 3667, 2014.