

Trade and Inequality: From Theory to Estimation*

Elhanan Helpman
Harvard University
and CIFAR

Oleg Itskhoki
Princeton University
and NBER

Marc-Andreas
Muendler
UC San Diego
and NBER

Stephen Redding
Princeton University
and NBER

February 12, 2014

Abstract

While neoclassical theory emphasizes the impact of trade on wage inequality *between* occupations and sectors, more recent theories of firm heterogeneity point to the impact of trade on wage dispersion *within* occupations and sectors. Using linked employer-employee data for Brazil, we show that much of overall wage inequality arises within sector-occupations and for workers with similar observable characteristics; this within component is driven by wage dispersion between firms; and wage dispersion between firms is related to firm employment size and trade participation. We then extend the heterogeneous-firm model of trade and inequality from Helpman, Itskhoki, and Redding (2010) and estimate it with Brazilian data. We show that the estimated model provides a close approximation to the observed distribution of wages and employment. We use the estimated model to undertake counterfactuals, in which we find sizable effects of trade on wage inequality.

Key words: Wage Inequality, International Trade

JEL classification: F12, F16, E24

*We thank the National Science Foundation for financial support. We also thank SECEX/MEDIC and Ricardo Markwald at FUNCEX Rio de Janeiro for sharing firm trade data. We are very grateful to Sam Bazzi, Lorenzo Casaburi, Itzik Fadlon, Adam Guren, Eduardo Morales and Jesse Schreger for excellent research assistance. We thank Mark Aguiar, Carlos Henrique Corseuil, Kerem Coşar, Angus Deaton, Thibault Fally, Cecilia Fielor, Leo Feler, Penny Goldberg, Gita Gopinath, Gene Grossman, Bo Honoré, Guido Imbens, Larry Katz, Francis Kramarz, Rasmus Lentz, Brian McCaig, Marc Melitz, Naercio Menezes-Filho, Eduardo Morales, Ulrich Müller, Andriy Norets, Richard Rogerson, Esteban Rossi-Hansberg, Andres Santos, Mine Senses, Robert Shimer, Chris Sims, John Van Reenen, Jon Vogel, Yoto Yotov and seminar participants at the Administrative Data for Public Policy Conference Singapore, Bayreuth, CESifo, IAE Workshop at UAB, CEPR, Columbia, EITI, FGV Rio de Janeiro, Georgia, Harvard, Humboldt Berlin, IAS, IGC Trade Programme Meeting at Columbia, JHU-SAIS, Maryland, Munich, NBER, Notre Dame, Princeton, SED, Syracuse, Toronto, the Philadelphia Fed Trade Workshop, UBC Vancouver, UC San Diego, Wisconsin, World Bank and WTO for helpful comments and suggestions. The usual disclaimer applies.

1 Introduction

The field of international trade has undergone a transformation in the last decade, as attention has shifted to heterogeneous firms as drivers of foreign trade. Until recently, however, research on the labor market effects of international trade has been heavily influenced by the Heckscher-Ohlin and Specific Factors models, which provide predictions about relative wages across skill groups, occupations and sectors. In contrast to the predictions of those theories, empirical studies find increased wage inequality in both developed and developing countries, growing residual wage dispersion among workers with similar observed characteristics, and increased wage dispersion across plants and firms *within* sectors. In part due to this disconnect, previous studies have concluded that the contribution of trade to growing wage inequality is modest at best.

This paper argues that these apparently discordant empirical findings are in fact consistent with a trade-based explanation for wage inequality, but one rooted in recent models of firm heterogeneity rather than neoclassical trade theories. For this purpose we develop a theoretical model of firm wages, employment and export status. We develop a methodology for estimating this model and illustrate with Brazilian data how the estimated model can be used to quantify the contribution of trade to the cross-section dispersion of wages between firms and workers.

To motivate our theoretical model, we first provide evidence on a number of stylized facts about wage inequality. These stylized facts combine approaches from different parts of the trade and labor literatures to provide an integrated view of the sources of wage inequality in Brazil. First, much of overall wage inequality occurs *within sectors and occupations* rather than between sectors and occupations. Second, a large share of this wage inequality within sectors and occupations is driven by wage inequality *between* rather than within firms. Third, both of these findings are robust to controlling for observed worker characteristics, suggesting that this wage inequality between firms within sector-occupations is *residual* wage inequality. These features of the data motivate our theoretical model's focus on wage inequality between firms for workers with similar observed characteristics.

We measure the between-firm component of wage inequality by including a firm-occupation-year fixed effect in a Mincer regression for log worker wages on controls for observed worker characteristics. This *firm wage component* includes both wage premia for workers with identical characteristics and unobserved differences in workforce composition across firms. Our analysis focuses on this overall wage component because recent theories of firm heterogeneity emphasize both sources of wage differences across firms. We estimate the firm wage component separately for each sector-occupation-year, because these theories emphasize that the firm wage component can vary across sectors, occupations and time. We find a strong relationship between the firm wage component and trade participation: exporters are on average larger and pay higher wages than non-exporters. While these exporter premia are robust features of the data, there is overlap in the exporter and non-exporter employment and wage distributions, so that some non-exporters are larger and pay higher wages than some exporters.

To account for these features of the data, we extend the theoretical framework of Helpman,

Itskhoki, and Redding (2010) to include two additional sources of heterogeneity across firms besides productivity: the cost of screening workers and the size of the fixed cost of exporting. Heterogeneous screening costs allow for variation in wages across firms after controlling for their employment size and export status, while idiosyncratic exporting costs allow some small low-wage firms to profitably export and some large high-wage firms to serve only the domestic market. We use the theoretical framework to derive an econometric model of firm employment, wages and export status. This econometric model explains positive exporter premia for employment and wages and predicts imperfect correlations between firm employment, wages and export status. It also highlights that the exporter wage premium depends on both the *selection* into exporting of more productive firms that pay higher wages and the increase in firm wages because of the greater *market access* of exporters.

We derive the closed-form solution for the econometric model’s likelihood function and estimate its parameters using maximum likelihood. We show that the parameterized model provides a good fit to the data, both for first and second moments of wages, employment and export status and for the distributions of wages and employment across firms and workers. We show that trade participation is important for the model’s fit, which deteriorates substantially when we shut down the market access effects of exporting. We use the estimated model to undertake counterfactuals, in which we find sizable effects of trade on wage inequality, comparable in magnitude to the inequality movements observed in our sample. Additionally, we provide bounds on the effect of trade on wage inequality in a class of econometric models consistent with the size and exporter wage premia observed in the cross-section of Brazilian manufacturing firms.

Our paper is related to a number of strands of research. Methodologically, our work connects most closely with the wider literature quantifying models of international trade and labor markets: Coşar, Guner, and Tybout (2011) fit and calibrate a model of search and matching frictions in which convex hiring costs induce wage variation across firms; Coşar (2010), Dix-Carneiro (2013) and Kambourov (2009) examine the role of sector-specific human capital in influencing the response of the economy to trade liberalization; Das, Roberts, and Tybout (2007) use a Bayesian approach to structurally estimate market entry costs from firm turnover; Irarrazabal, Moxnes, and Opromolla (2011) examine the determinants of trade versus multinational activity; Eaton, Kortum, Kramarz, and Sampogaro (2013) explore the labor market effects of trade liberalization in a model of firm-to-firm trade; Egger, Egger, and Kreickemeier (2011) provide evidence on the wage inequality predictions of a model of firm heterogeneity and fair wages.

As discussed above, several empirical studies have suggested that the Heckscher-Ohlin and Specific Factors models—as conventionally interpreted—provide at best an incomplete explanation for observed wage inequality. First, changes in the relative returns to observed measures of skills (e.g., education and experience) and changes in sectoral wage premia account for a limited share of change in overall wage inequality, leaving a substantial role for residual wage inequality.¹ Second, the Stolper-Samuelson theorem predicts a rise in the relative skilled wage

¹For developed country evidence, see for example Autor, Katz, and Kearney (2008), Juhn, Murphy, and Pierce (1993), and Lemieux (2006). For developing country evidence, see for example Attanasio, Goldberg, and Pavcnik

in skill-abundant countries and a fall in the relative skilled wage in unskilled-abundant countries in response to trade liberalization. Yet wage inequality rises following trade liberalization in both developed and developing countries (e.g., Goldberg and Pavcnik 2007).² Third, much of the change in the relative demand for skilled and unskilled workers in developed countries has occurred within sectors and occupations rather than across sectors and occupations (e.g., Katz and Murphy 1992 and Berman, Bound, and Griliches 1994). Fourth, while wage dispersion between plants and firms is an empirically-important source of wage inequality (e.g., Davis and Haltiwanger 1991 and Faggio, Salvanes, and Van Reenen 2010), neoclassical trade theory is not able to elucidate it. Each of these features of the data can be explained within the class of new trade models based on firm heterogeneity.

Models of firm heterogeneity suggest two sets of reasons for wage variation across firms. One line of research assumes competitive labor markets, so that all workers with the same characteristics are paid the same wage, but wages vary across firms as a result of differences in workforce composition (see for example Yeaple 2005, Verhoogen 2008, Bustos 2011, Burstein and Vogel 2012, Monte 2011 and Sampson 2012). Another line of research introduces labor market frictions so that workers with the same characteristics can be paid different wages by different firms. For example, efficiency or fair wages can result in wage variation across firms when the wage that induces worker effort, or is perceived to be fair, varies with the revenue of the firm (see for example Egger and Kreickemeier 2009, Amiti and Davis 2012 and Davis and Harrigan 2011). Furthermore, search and matching frictions and the resulting bargaining over the surplus from production can induce wages to vary across firms (see for example Davidson, Matusz, and Shevchenko 2008).³ Helpman, Itskhoki, and Redding (2010; HIR henceforth) develop a model with frictional assortative matching of workers to firms, and hence featuring both mechanisms for wage variation across firms. We estimate the extended HIR model, for which we show that the reduced-form representation takes a log linear selection form. We conjecture that such a reduced form can be derived from other models of heterogeneous firms, and hence it is potentially more general than the particular microfoundations we rely on. Nonetheless, by providing microfoundations for our econometric model, we ensure that the estimated model is consistent with individual optimization and equilibrium, obtain theoretical restrictions, and are able to undertake model-based counterfactuals.

Our paper also relates to the existing literature estimating search-theoretic equilibrium models of the labor market, including Burdett and Mortensen (1998), Cahuc, Postel-Vinay, and Robin (2006), Postel-Vinay and Robin (2002), and Postel-Vinay and Thuron (2010). The focus of this literature is typically on employment and wage dynamics. In contrast, the focus of our

(2004), Ferreira, Leite, and Wai-Poi (2010), Goldberg and Pavcnik (2005), Gonzaga, Menezes-Filho, and Terra (2006), and Menezes-Filho, Muendler, and Ramey (2008).

²Increasing wage inequality in both developed and developing countries can be explained by re-interpreting the Stolper-Samuelson Theorem as applying within sectors as when production stages are offshored (see for example Feenstra and Hanson (1996), Feenstra and Hanson (1999) and Trefler and Zhu (2005)).

³Search and matching frictions may also influence income inequality through unemployment, as in Davidson and Matusz (2010), Felbermayr, Prat, and Schmerer (2011), and Helpman and Itskhoki (2010).

analysis is the cross-section variation in employment and wages across firms. More specifically, we study how international trade affects the employment and wage distribution across firms, which requires a substantially richer product market structure (e.g., as in Melitz 2003) than what is typically assumed in the labor-macro literature (i.e., perfect substitutes under perfect competition). Furthermore, firm size is at the core of the mechanism that we explore, and hence we need to depart from the assumptions in the labor-macro literature which cast the analysis in terms of jobs rather than firms, as in this literature firm boundaries are often not well-defined. As a result, we choose to abstract from the labor market dynamics that are central for the labor-macro literature, in order to embed a rich frictional labor market into a model of international trade in differentiated product markets that we use to analyze the cross-sectional relationship between firm employment, wages and export participation.

Finally, our work is also related to empirical research using plant and firm data that has found substantial differences in wages and employment between exporters and non-exporters following Bernard and Jensen (1995, 1997). More recent research using linked employer-employee datasets has sought to determine the sources of the exporter wage premium, including Schank, Schnabel, and Wagner (2007), Munch and Skaksen (2008), Frías, Kaplan, and Verhoogen (2009), Davidson, Heyman, Matusz, Sjöholm, and Zhu (2011), Krishna, Poole, and Senses (2011), and Baumgarten (2011). This literature typically makes the assumption that the matching of workers to firms is random after controlling for firm fixed effects, worker fixed effects and time-varying worker observables. Under this assumption, these empirical studies typically find that the exporter wage premium is composed of both unobserved differences in workforce composition and wage premia for workers with identical characteristics, with the relative importance of these two forces varying across studies. Since we focus on the overall firm wage component, including both these sources of differences in wages across firms, we need not rely on the assumption of conditional random matching of workers to firms which is violated in our theoretical framework. Furthermore, while the above empirical studies typically estimate a time-invariant wage fixed effect for each firm, a key feature of our approach is that the firm component of wages can change over time with firm revenue (e.g., with both changes in market conditions and firm entry into export markets). In contrast to these empirical studies, we use our theoretical framework to derive an econometric model that jointly determines firm employment, wages and export status and can be estimated using exclusively the cross-section of the data.

The remainder of the paper is structured as follows. In Section 2, we introduce our data and some background information. In Section 3, we present some stylized facts about wage inequality in Brazil. Motivated by these findings, Section 4 develops a heterogeneous-firm model of trade and inequality, derives an econometric model that we estimate using these data, and reports our counterfactuals for the impact of trade liberalization on wage inequality. Section 5 concludes. A supplementary web appendix contains detailed derivations, description of the data, and additional results.⁴

⁴Access at http://www.princeton.edu/~itskhoki/papers/TradeInequalityEvidence_appendix.pdf.

2 Data and Background

Our main dataset is a linked employer-employee dataset for Brazil from 1986-1998, which we briefly describe here and discuss in further detail in the web appendix. The source for these administrative data is the *Relação Anual de Informações Sociais* (RAIS) database of the Brazilian Ministry of Labor. By law, all formally-registered firms are required to report information each year on each worker employed by the firm, as recorded in RAIS. The data contain a unique identifier for each worker, which remains with the worker throughout his or her work history as well as the tax identifier of the worker's employer.

We focus on the manufacturing sector, because manufacturing goods are typically tradable and there is substantial heterogeneity across sectors, occupations and firms within manufacturing. Therefore this sector provides a suitable testing ground for traditional and heterogeneous firm theories of international trade. Manufacturing is also an important source of employment in Brazil, accounting in 1990 and 1998 for around 23 and 19 percent of total employment (formal and informal) respectively. Our data cover all manufacturing firms and workers in the formal sector, which Goldberg and Pavcnik (2003) estimates accounts for around 84 percent of manufacturing employment.

Our annual earnings measure is a worker's mean monthly wage, averaging the worker's wage payments over the course of a worker's employment spell during a calendar year.⁵ For every worker with employment during a calendar year, we keep the worker's last recorded job spell and, if there are multiple spells spanning into the final month of the year, the highest-paid job spell (randomly dropping ties). Therefore our definition of firm employment is the count of employees whose employment spell at the firm is their final (highest-paid) job of the year.

We undertake our analysis at the firm rather than the plant level, because recent theories of firm heterogeneity and trade are concerned with firms, and wage and exporting decisions are arguably firm based. For our baseline sample we focus on firms with five or more employees, because we analyze wage variation within and across firms, and the behavior of firms with a handful of employees may be heavily influenced by idiosyncratic factors. But we find a similar pattern of results using the universe of firms. Our baseline sample includes an average of 6.38 million workers and 92,513 firms in each year.

Each worker is classified in each year by her or his occupation. In our baseline empirical analysis, we use five standard occupational categories that are closely related to skill groups, described in Table 1 which also reports the employment shares of each occupation and the mean log wage in each occupation relative to manufacturing overall. In robustness checks, we also make use of the more disaggregated *Classificação Brasileira de Ocupações (CBO)* definition of occupations, which breaks down manufacturing into around 350 occupations.

Each firm is classified in each year by its main sector according to a classification compiled by

⁵Wages are reported as multiples of the minimum wage, which implies that inflation that raises the wages of all workers by the same proportion leaves this measure of relative wages unchanged. Empirically, we find a smooth left tail of the wage distribution in manufacturing, which suggests that the minimum wage is not strongly binding in manufacturing during our sample period. RAIS does not report hours, overtime, investment or physical capital.

Table 1: Occupation Employment Shares and Relative Mean Log Wages, 1994

CBO	Occupation	Employment share (percent)	Relative mean log wage
1	Professional and Managerial	7.2	1.12
2	Skilled White Collar	10.8	0.38
3	Unskilled White Collar	8.8	0.07
4	Skilled Blue Collar	63.1	-0.14
5	Unskilled Blue Collar	10.0	-0.39

Note: Share in total formal manufacturing-sector employment; log wage minus average log wage in formal manufacturing sector.

Table 2: Sectoral Employment Shares and Relative Mean Log Wages, 1994

IBGE	Sector	Emplmnt share (percent)	Relative mean log wage	Exporter share (percent)	
				Firms	Emplmnt
2	Non-metallic Minerals	4.6	-0.21	4.7	34.6
3	Metallic Products	10.3	0.31	9.9	57.6
4	Mach., Equip. and Instruments	5.9	0.48	25.4	71.8
5	Electrical & Telecomm. Equip.	4.3	0.41	19.9	70.9
6	Transport Equip.	6.0	0.73	13.6	75.3
7	Wood & Furniture	6.9	-0.51	8.0	39.7
8	Paper & Printing	5.5	0.20	4.8	37.0
9	Rubber, Tobacco, Leather, etc.	5.1	-0.05	12.8	56.7
10	Chemical & Pharm. Products	9.4	0.31	15.6	56.8
11	Apparel & Textiles	15.1	-0.34	4.8	42.7
12	Footwear	5.4	-0.44	16.8	72.3
13	Food, Beverages & Alcohol	21.3	-0.18	4.1	42.2
	All Manufacturing Sectors	100	0.00	9.0	51.8

Note: Share in total formal manufacturing-sector employment; log wage minus average log wage in formal manufacturing sector; share of firms that export; employment share of exporters.

the *Instituto Brasileiro de Geografia e Estatística (IBGE)*, which disaggregates manufacturing into twelve sectors roughly corresponding to two-digit International Standard Industrial Classification (ISIC) sectors. Sectoral employment shares and the mean log wage in each sector relative to the overall manufacturing mean log wage are reported in Table 2. From 1994 onwards, firms are classified according to the more finely-detailed *National Classification of Economic Activities (CNAE)*, which breaks down manufacturing into over 250 industries. In robustness checks we use this more detailed classification when it is available. In the web appendix, we report the concordance between our baseline five occupations and twelve sectors and the more disaggregated occupations and sectors.

From Tables 1 and 2, there is substantial variation in average wages across both occupations and sectors. Skilled White Collar workers are paid on average 52 and 77 log points over Skilled and Unskilled Blue Collar workers respectively, which correspond to wage premia of roughly

68 and 116 percent respectively. Machinery and equipment sectors 4–6 pay an average wage premium of around 72 percent compared to the typical manufacturing wage, while furniture and footwear sectors (7 and 12) pay on average less than two thirds of the typical manufacturing wage. Therefore both occupations and sectors are consequential for wages, leaving open the possibility that between-sector and between-occupation effects could be important for the evolution of overall wage inequality. We provide evidence on the extent to which this is the case below.

RAIS also reports information on worker educational attainment. Our choice of educational classification is guided by the existing labor economics literature, including Autor, Katz, and Krueger (1998) and Katz and Autor (1999). In our baseline specification, we distinguish the following four categories: (i) Less than High School, (ii) High School, (iii) Some College, and (iv) College Degree. We also report the results of a robustness test using nine more disaggregated educational categories.⁶ During our sample period, the shares of workers with some college education or a college degree are relatively constant, while the share of workers with high-school education rises by around 10 percentage points. In addition to these data on educational attainment, RAIS also reports information on age and gender for each worker. Finally, we construct a measure of a worker’s tenure with a firm based on the number of months for which the worker has been employed by the firm.

We combine the linked employer-employee data from RAIS with trade transactions data from *Secretaria de Comércio Exterior* (SECEX) that are available from 1986-1998. These trade transactions data report for each export and import customs shipment the tax identifier of the firm, the product exported and the destination country served. We merge the trade transactions and linked employer-employee data using the tax identifier of the firm. As shown in Table 2, exporters account for a much larger share of employment than the number of firms: the fraction of exporters ranges from 4.1 to 25.4 percent, while the exporter share of employment ranges from 34.6 to 75.3 percent. Since exporters account for a disproportionate share of employment, differences in wages between exporters and non-exporters can have disproportionate effects on the distribution of wages across *workers*.

Our sample period includes changes in both trade and labor market policies in Brazil. Tariffs are lowered in 1988 and further reduced between 1990 and 1993, whereas non-tariff barriers are dropped by presidential decree in January 1990. Following this trade liberalization, the share of exporting firms nearly doubles between 1990 and 1993, and their employment share increases by around 10 percentage points.⁷ In contrast, following Brazil’s real exchange rate appreciation of 1995, both the share of firms that export and the employment share of exporters decline by around the same magnitude. In 1988, there was also a reform of the labor market. Finally, the late 1980s and early 1990s witnessed some industrial policy initiatives, which were mostly

⁶The nine more disaggregated categories are: Illiterate, some primary, complete primary, some middle, complete middle, some high, complete high, some college, and complete college.

⁷For an in-depth discussion of trade liberalization in Brazil, see for example Kume, Piani, and Souza (2003). The changes in the exporter employment share discussed above are reflected in a similar pattern of aggregate manufacturing exports, as shown in the web appendix.

applied on an industry-wide basis.⁸

Our theoretical model is concerned with the cross-section distribution of wages and employment. Therefore we estimate it using data across firms and workers for a given year. We show that we find consistent predictions for the effects of trade on wage inequality across years both before and after the reforms.

3 Stylized Facts

In this section, we combine different approaches from the trade and labor literatures to develop a set of stylized facts on wage inequality in Brazil. We present a sequence of variance decompositions that quantify the relative importance of alternative possible sources of wage inequality. In each year, we decompose overall wage inequality (T) into a within (W) and a between component (B) as follows:

$$\begin{aligned}
 T_t &= W_t + B_t \\
 T_t &= \frac{1}{N_t} \sum_{\ell} \sum_{i \in \ell} (w_{it} - \bar{w}_t)^2, \\
 W_t &= \frac{1}{N_t} \sum_{\ell} \sum_{i \in \ell} (w_{it} - \bar{w}_{\ell t})^2, \\
 B_t &= \frac{1}{N_t} \sum_{\ell} N_{\ell t} (\bar{w}_{\ell t} - \bar{w}_t)^2,
 \end{aligned} \tag{1}$$

where workers are indexed by i and time by t ; ℓ denotes sector, occupation or sector-occupation cells depending on the specification; N_t and $N_{\ell t}$ denote the overall number of workers and the number of workers within cell ℓ ; w_{it} , $\bar{w}_{\ell t}$ and \bar{w}_t are the log worker wage, the average log wage within cell ℓ and the overall average log wage. The use of the log wage ensures that the decomposition is not sensitive to the choice of units for wages and facilitates the inclusion of controls for observable worker characteristics.

When undertaking this decomposition, we report results for the level of wage inequality for 1994, because this year is after trade liberalization and before the major appreciation of the real. We report growth results for 1986-1995, because this is the period over which the growth in wage inequality in Brazilian manufacturing occurs. We find a similar pattern of results for different years, as shown in the figures in subsections D2-D4 of the web appendix, which displays the evolution of overall wage inequality and its components for each year in our sample.

3.1 Within versus between sectors and occupations

We start by decomposing overall wage inequality into within and between components using the decomposition (1) for sector, occupation and sector-occupation cells, building upon Davis and

⁸The main elements of the 1988 labor market reform were a reduction of the maximum working hours per week from 48 to 44, an increase in the minimum overtime premium from 20 percent to 50 percent, and a reduction in the maximum number of hours in a continuous shift from 8 to 6 hours, among other institutional changes. Among the industrial policy initiatives, some tax exemptions differentially benefited small firms while foreign-exchange restrictions and special import regimes tended to favor select large-scale firms until 1990.

Table 3: Contribution of the Within Component to Log Wage Inequality

	Level (percent)	Change (percent)
	1994	1986–95
A. Main Period		
Within occupation	82	92
Within sector	83	73
Within sector-occupation	68	66
Within detailed-occupation	61	60
Within sector–detailed-occupation	56	54
B. Late Period	1994	1994–98
Within detailed-sector–detailed-occupation	47	141

Note: Each cell in the table reports the contribution of the within component to total log wage inequality. The unreported between component is 100 percent minus the reported within component. The within component exceeds 100 percent when the between component moves in the opposite direction partially offsetting its effect.

Haltiwanger (1991) and Katz and Murphy (1992).⁹

In Panel A of Table 3, we report the contribution of each within component to the level and growth of overall wage inequality. Although the contribution of the within component inevitably falls as one considers more and more disaggregated categories, it accounts for 82, 83 and 68 percent of the level of overall wage inequality for occupations, sectors and sector-occupations respectively (first column). Similarly, the majority of the growth in the variance of log wages of around 17.4 percent (corresponding to a 8.3 percent increase in the standard deviation of log wages) is explained by wage inequality within occupations, sectors and sector-occupations (second column).

Fact 1 *The within sector-occupation component of wage inequality accounts for the majority of both the level and growth of wage inequality in Brazil between 1986 and 1995.*

While our baseline results use the IBGE classification of twelve manufacturing sectors and five occupations, the importance of the within component is robust to the use of alternative more detailed definitions of sectors and occupations. In Panel A of Table 3, we report results using detailed occupation cells based on more than 300 occupations in the CBO classification (fourth row) and using sector-detailed-occupation cells defined using IBGE sectors and CBO occupations (fifth row). As a further robustness check, Panel B of Table 3 reports results for a later time period 1994–98, for which the more finely-detailed CNAE sector classification is available. Here we consider detailed-sector-detailed-occupation cells based on around 350 CNAE sectors and more than 250 CBO occupations. While some occupations do not exist in some sectors, there are still around 40,000 sector-occupation cells in this specification. Yet we continue to find that the within component accounts for around half of the level and most of the growth of wage inequality.

⁹See also Barth, Bryson, Davis, and Freeman (2011) for additional evidence using U.S. plant-level data and Faggio, Salvanes, and Van Reenen (2010) for evidence using U.K. firm-level data.

Neoclassical theories of international trade emphasize wage inequality between different types of workers (Heckscher-Ohlin model) or sectors (Specific Factors model). Our findings suggest that this focus on the between component abstracts from an important potential channel through which trade can affect wage inequality. Of course, our results do not rule out the possibility that Heckscher-Ohlin and Specific-Factors forces play a role in the wage distribution. As shown in Feenstra and Hanson (1996, 1999) and Trefler and Zhu (2005), the Stolper-Samuelson Theorem can be re-interpreted as applying at a more disaggregated level within sectors and occupations such as production stages. But these neoclassical theories emphasize *dissimilarities* across sectors and occupations, and if their mechanisms are the dominant influences on the growth of wage inequality, we would expect to observe a substantial between-component for grossly-different occupations and sectors (e.g., Managers versus Unskilled Blue-collar workers and Textiles versus Chemicals and Pharmaceuticals). Yet the within component accounts for a substantial proportion of overall wage inequality and continues to do so even when we consider around 40,000 disaggregated sector-occupations. Therefore, while the forces highlighted by neoclassical trade theory may be active, there appear to be other important mechanisms that are also at work.

3.2 Worker observables and residual wage inequality

We now examine whether the contribution of the within-sector-occupation component of wage inequality is robust to controlling for observed worker characteristics, building on the recent literature on residual wage inequality including Autor, Katz, and Kearney (2008), Juhn, Murphy, and Pierce (1993), and Lemieux (2006). To control for worker observables, we estimate the following OLS Mincer regression for log worker wages:

$$w_{it} = z'_{it}\vartheta_t + \nu_{it}, \quad (2)$$

where i still denotes workers, z_{it} is a vector of observable worker characteristics, ϑ_t is a vector of returns to worker observables, and ν_{it} is a residual. We estimate this Mincer regression for each year separately, allowing the coefficients on worker observables (ϑ_t) to change over time to capture changes in the rate of return to these characteristics. We control for worker observables nonparametrically by including indicator variables for the following categories: education (we use our baseline four categories discussed above and report a robustness test using the nine more disaggregated categories), age (using the categories 10-14, 15-17, 18-24, 25-29, 30-39, 40-49, 50-64, 65+), quintiles of experience (tenure) at the firm, and gender.

The empirical specification (2) serves as a conditioning exercise, which allows us to decompose the variation in log wages into the component correlated with worker observables and the orthogonal residual component:

$$T_t = \text{var}(w_{it}) = \text{var}(z'_{it}\hat{\vartheta}_t) + \text{var}(\hat{\nu}_{it}), \quad (3)$$

where the hat denotes the estimated value from regression (2). We refer to $\text{var}(\hat{\nu}_{it})$ as *residual*

Table 4: Worker Observables and Residual Log Wage Inequality

	Level (percent) 1994	Change (percent) 1986–95
Residual wage inequality	59	49
— within sector-occupation	89	90

Note: The first row decomposes the level and growth of overall log wage inequality into the contributions of worker observables and residual (within-group) wage inequality using (2) and (3). The unreported contribution of worker observables equals 100 percent minus the reported contribution of residual wage inequality. The second row reports the within sector-occupation component of residual wage inequality.

wage inequality. We can further decompose residual wage inequality into its within and between components using sector, occupation or sector-occupation cells, by applying (1) to the estimated residuals $\hat{\nu}_{it}$.

Table 4 reports the results of the variance decomposition (3). In the first row, we find that the worker observables and residual components make roughly equal contributions towards both the level (1990) and growth (1986–1995) of overall wage inequality.¹⁰ In the second row, we decompose the level and growth of residual wage inequality into its within and between sector-occupation components. We find that the within sector-occupation component dominates, explaining around 90 percent of both level and growth of the residual wage inequality.¹¹ Comparing the results in Tables 4 and 3, the within sector-occupation component is more important for residual wage inequality than for overall wage inequality, which is consistent with the fact that much of the variation in worker observables is between sector-occupation cells. This enhanced dominance of the within-sector-occupation component after controlling for worker observables suggests that the majority of residual wage inequality is a within sector-occupation phenomenon.

Note that residual wage inequality is measured relative to the worker characteristics included in the regression (2). In principle, there can be other unmeasured worker characteristics that matter for wages and that are observed by the firm but are uncorrelated with the worker characteristics available in our data. To the extent that this is the case, the contribution of worker characteristics could be larger than estimated here. On the other hand, the decomposition (3) projects all variation in wages that is correlated with the included worker characteristics on worker observables. Therefore, if the firm component of wages is correlated with these worker characteristics, some of its contribution to wage variation can be attributed to worker observables. In the next subsection, we use a different decomposition, in which we explicitly control for firm fixed effects and find a smaller contribution from worker observables towards overall

¹⁰Consistent with Attanasio, Goldberg, and Pavcnik (2004) and Menezes-Filho, Muendler, and Ramey (2008), we find an increase in the estimated returns to education and experience (tenure) in the Mincer log wage equation, as reported in the web appendix.

¹¹We find a similar pattern of results using the nine more disaggregated education categories: for example, the residual component accounts for 57 percent of the level (1994) of wage inequality and 46 percent of the growth (1986–95) of wage inequality. Around 90 percent of both the level and growth of this residual wage inequality is again explained by the within sector-occupation component.

Table 5: Regional Robustness

	OVERALL INEQUALITY		RESIDUAL INEQUALITY	
	Level	Change	Level	Change
	1994	1986–95	1994	1986–95
Within sector-occupation	68	66	89	90
Within sector-occupation, São Paulo	64	49	89	71
Within sector-occupation-state	58	38	76	56
Within sector-occupation-meso	54	30	72	49

Note: All entries are in percent. The first line duplicates the baseline results from Table 3 (overall inequality) and 4 (residual inequality). The second line reports the same decomposition for the state of São Paulo. The last two lines report the within component using sector-occupation-region cells, where regions are first 27 states and second 136 meso regions.

wage inequality. Keeping these caveats in mind, we state:

Fact 2 *Residual wage inequality is at least as important as worker observables in explaining the overall level and growth of wage inequality in Brazil from 1986-1995. Most of the level and growth of residual wage inequality is within-sector-occupation.*

One potential concern is that regional differences in wages could drive wage inequality within sector-occupations for workers with similar observed characteristics.¹² In Table 5, we demonstrate the robustness of our results to controlling for region. In the first row, we restate our baseline results. In the second row, we report results for the state of São Paulo, which accounts for around 45 percent of formal manufacturing employment in our sample. In the third and fourth rows, we report results using sector-occupation-region cells instead of sector-occupation cells, where we define regions in terms of either 27 states or 136 meso regions. These specifications abstract from any variation in wages across workers within sector-occupations that occurs between regions. Nonetheless, in each specification, we continue to find that a sizeable fraction of wage inequality is a within phenomenon. This is particularly notable for residual wage inequality, where the within component still accounts for over two thirds of the level and around half of the growth of residual inequality even for the detailed meso-regions.¹³

Another potential concern is that our findings for wage inequality could be influenced by changes in workforce composition if, for example, residual wage inequality is more prevalent within certain skill groups. We follow Lemieux (2006) to address this concern by holding workforce composition across cells constant at its beginning of the sample values, and find the same pattern of residual wage inequality, as discussed further in the web appendix.

¹²For empirical evidence of wage variation across Brazilian states, see for example Fally, Paillacar, and Terra (2010) and Kovak (2011).

¹³While we find that the majority of residual wage inequality is within sector-occupation-region, we also estimated the Mincer wage regression (2) including meso region fixed effects and sector fixed effects, where the sector fixed effects capture the between-sector component of wage inequality after controlling for worker observables and region. As shown in Figure D8 in the web appendix, these estimated sector fixed effects are relatively stable over time, and do not account for the rising and declining pattern of wage inequality over time.

3.3 Between versus within-firm wage inequality

We now decompose wage inequality within sectors and occupations into the contributions of within-firm and between-firm components. Here we build on the recent literature in labor economics that has used linked employer-employee data to estimate firm wage components, including Abowd, Kramarz, and Margolis (1999), Abowd, Kramarz, Margolis, and Troske (2001), Menezes-Filho, Muendler, and Ramey (2008) and Lazear and Shaw (2009).

For each sector-occupation-year cell, we decompose wage inequality across workers in that cell into within and between-firm components. To do so, we regress log worker wages on firm fixed effects and observable worker characteristics for each sector-occupation-year separately:¹⁴

$$w_{it} = z'_{it}\vartheta_{\ell t} + \psi_{j\ell t} + \nu_{it}, \quad (4)$$

where i again indexes workers, j indexes firms, and ℓ indexes sector-occupation cells; we normalize the firm-occupation-year fixed effects $\psi_{j\ell t}$ to sum to zero for each sector-occupation-year, which implies that the regression constant is separately identified (and we absorb it into the worker observables component, $z'_{it}\vartheta_{\ell t}$); we allow the coefficients ($\vartheta_{\ell t}$) on observed worker characteristics (z'_{it}) to differ across sector-occupations ℓ and time t to capture variation in their rate of return; and ν_{it} is a stochastic error.¹⁵

Although $\psi_{j\ell t}$ is a firm fixed effect, the regression (4) is estimated for each sector-occupation-year, and therefore this firm-occupation-year fixed effect varies over time and across occupations. To emphasize this difference from time-invariant firm fixed effects, we refer to $\psi_{j\ell t}$ as a *firm wage component*. We allow the firm wage component to change over time, because theories of heterogeneous firms and trade such as Helpman, Itskhoki, and Redding (2010) emphasize that firm wages vary with firm revenue (e.g., as firms enter and exit export markets). Similarly, we allow the firm wage component to differ across occupations because these theories imply that the sensitivity of firm wages to firm revenue can differ across occupations. We also consider a restricted version of equation (4) excluding the controls for worker observables (so that $z'_{it}\vartheta_{\ell t}$ consists solely of the regression constant). We distinguish between our estimates of $\psi_{j\ell t}$ with and without the controls for worker observables by using the terms conditional and unconditional firm wage components respectively ($\psi_{j\ell t}^C$ and $\psi_{j\ell t}^U$).

We use the estimated firm-occupation-year fixed effects controlling for worker observables ($\hat{\psi}_{j\ell t}^C$) as our baseline measure of the firm component of wages in the econometric model below.

¹⁴While we estimate a separate regression for each sector-occupation-year, we could equivalently pool sector-occupation-year observations and allow all coefficients to vary by sector-occupation-year. With this log linear regression specification, these two procedures generate the same estimated firm-occupation-year fixed effects.

¹⁵In Table 6, we treat the firm-occupation-year fixed effects ($\hat{\psi}_{j\ell t}$) as data. In the model developed below, we make the theoretical assumption that the firm observes these wage components and that the model is about these wage components, which can be therefore taken as data in its estimation. In contrast, without this theoretical assumption, $\hat{\psi}_{j\ell t}$ should be interpreted as estimates, in which case the variance of these estimates equals their true variance plus the variance of a sampling error that depends on the average number of workers employed by a firm. Since this average is around 70 workers in our data, the resulting correction for the variance of the sampling error is small, as discussed further in the web appendix.

These firm-occupation-year fixed effects capture both firm wage premia for workers with identical characteristics and unobserved differences in workforce composition across firms (including average match effects). The theoretical literature on heterogeneous firms and labor markets considers both these sources of wage differences across firms, and our objective is to quantify the overall contribution of the firm component to wage inequality, rather than sorting out further its different components. Our baseline specification allows the firm-occupation-year fixed effects ($\hat{\psi}_{j\ell t}^C$) to be correlated with worker observables, as will be the case, for example, if there is assortative matching on worker observables across firms.

Using the estimates from (4), we decompose wage inequality within each sector-occupation-year into the following four terms:

$$\text{var}(w_{it}) = \text{var}(z'_{it}\hat{\vartheta}_{\ell t}) + \text{var}(\hat{\psi}_{j\ell t}^C) + 2\text{cov}(z'_{it}\hat{\vartheta}_{\ell t}, \hat{\psi}_{j\ell t}^C) + \text{var}(\hat{\nu}_{it}). \quad (5)$$

These four terms are: (1) worker observables; (2) the between-firm component (firm-occupation year fixed effects); (3) the covariance between worker observables and the firm component; (4) the within-firm component (residual), which by construction is orthogonal to the other terms. In the restricted version of equation (4) excluding the controls for worker observables, the decomposition (5) includes only the between-firm and within-firm components.

We summarize the aggregate results from these decompositions as the employment-weighted average of the results for each sector-occupation-year cell. These aggregate results capture the average importance of the between-firm and within-firm components in accounting for wage variation within sector-occupation-year cells. Note that these results do not capture average differences in wages between occupations within firms, because the firm-occupation-year fixed effects have a mean of zero for each sector-occupation-year.

In the first two columns of Table 6, we report the results for the unconditional firm wage component ($\hat{\psi}_{j\ell t}^U$). We find that between and within-firm wage inequality make roughly equal contributions to the level of wage inequality within sector-occupations (first column, top panel). In contrast, the growth of wage inequality within sector-occupations is almost entirely explained by wage inequality between firms (second column, top panel).

In the final two columns of Table 6, we summarize the results for the conditional firm wage component ($\hat{\psi}_{j\ell t}^C$). As shown in the third column, we find that the between-firm and within-firm (residual) components account for roughly equal amounts of the level of wage inequality within sector-occupations (39 and 37 percent respectively). Of the other two components, worker observables account for around one eighth, and the covariance between worker observables and the firm component of wages accounts for the remaining one tenth. In contrast, in the fourth column, changes in between-firm wage dispersion account for most (86 percent) of the growth in wage inequality within sector-occupations. The next largest contribution (around one quarter) comes from an increased correlation between worker observables and the firm wage component (consistent with increased assortative matching on worker observables). Changes in the residual

Table 6: Decomposition of Log Wage Inequality *within* Sector-Occupations

	UNCONDITIONAL FIRM WAGE COMPONENT, $\hat{\psi}_{jlt}^U$		CONDITIONAL FIRM WAGE COMPONENT, $\hat{\psi}_{jlt}^C$	
	Level 1994	Change 1986–1995	Level 1994	Change 1986–1995
Between-firm wage inequality	55	115	39	86
Within-firm wage inequality	45	–15	37	–11
Worker observables			13	2
Covar observables–firm effects			11	24
Between-firm wage inequality				
— Between-firm within meso region	64	52	58	51

Note: All entries are in percent. Decomposition of the level and growth of wage inequality within sector-occupations (employment-weighted average of the results for each sector-occupation). The decomposition in the first two columns corresponds to the unconditional firm wage component that does not control for worker observables. The decomposition in the last two columns corresponds to the conditional firm wage component that controls for worker observables. Figures may not sum exactly to 100 percent due to rounding.

within-firm wage dispersion make a small negative contribution.¹⁶

We find a similar pattern of results for other years and definitions of sectors and occupations. And although for brevity we concentrate on aggregate results, the same pattern is pervasive across sectors and occupations. Note that the contribution of worker observables is smaller in Table 6 than in Table 4, because we now control for firm-occupation-year fixed effects and focus on wage inequality within-sector occupations.¹⁷

Since location is a fixed characteristic of the firm within a sector-occupation-year cell, the Mincer regression (4) cannot be augmented with region fixed effects, because these are perfectly colinear with the firm-occupation-year fixed effects. However, in the bottom panel of Table 6, we decompose the variation in the firm-occupation-year fixed effects (between-firm wage inequality) into variation within and between meso regions. Although this specification is conservative, because the within-meso-region component abstracts from any differences in firm wages across meso regions, more than one half of between-firm wage inequality within sector-occupations

¹⁶We find similar results using firm-year rather than firm-occupation-year fixed effects. For example, estimating the regression (4) for each sector-year and implementing the decomposition (5), we find the following contributions to the within-sector component of wage inequality in 1994: worker observables (21 percent); between-firm component (29 percent); covariance (13 percent); within-firm component (37 percent). Over time, the between-firm component accounts for 76 percent of the growth in wage inequality within sectors from 1986–1995. This similarity of the results using firm-year fixed effects is consistent with our normalization of the firm-occupation-year fixed effects to sum to zero for each sector-occupation-year, which implies that they do not capture average differences in wages between occupations within firms.

¹⁷The reduced contribution from worker observables is explained by two differences from the specification in Section 3.2. First, the Mincer regression (4) is estimated separately for each sector-occupation-year, instead of pooling observations across sectors and occupations for each year. Since much of the variation in worker observables occurs across sectors and occupations (sorting across sectors and occupations), this generates a smaller contribution from worker observables. Second, the Mincer regression (4) includes firm-occupation-year fixed effects, which can be correlated with worker observables (sorting across firms). Once we allow for this correlation, we again find a smaller contribution from worker observables. Empirically, these two differences are of roughly equal importance in explaining the reduction in the contribution from worker observables.

occurs within meso regions.

Fact 3 *Between-firm and within-firm dispersion make roughly equal contributions to the level of wage inequality within sector-occupations, but the growth of wage inequality within sector-occupations is largely accounted for by between-firm wage dispersion.*

This importance of between rather than within firm variation points towards theories of heterogeneity between firms as the relevant framework for understanding both the overall wage distribution and the wage distribution across workers with similar observed characteristics.

3.4 Size and exporter wage premia

We now examine the relationship between the firm wage component and firm employment and export status, building on the empirical trade literature following Bernard and Jensen (1995, 1997). We first construct a measure of firm wages in each year by aggregating our firm-occupation-year wage components from the previous subsection to the firm-year level using employment weights. We next regress both these firm wage components on firm employment and export status for each year:

$$\psi_{jt} = \lambda_{\ell t}^o + \lambda_t^s h_{jt} + \lambda_t^x \iota_t + \nu_{jt}, \quad (6)$$

where we again index firms by j ; ℓ now denotes sectors; h_{jt} is log firm employment; $\iota_{jt} \in \{0, 1\}$ is a dummy for whether a firm exports; ν_{jt} is the residual; $\lambda_{\ell t}^o$ is a sector-time fixed effect; λ_t^s is the employment size wage premium and λ_t^x is the exporter wage premium, where we allow both of these premia to vary over time.

In Table 7, we report the results for 1994 for both measures of the firm wage component. Consistent with the existing labor and international trade literatures (e.g., Oi and Idson 1999, Bernard and Jensen 1995, 1997), we find positive and statistically significant premia for employment size and export status. Using the firm wage component controlling for worker observables ($\hat{\psi}_{jt}^C$), we find a size premium of $\hat{\lambda}^s = 0.104$ and an exporter premium of $\hat{\lambda}^x = 0.168$.¹⁸ In this reduced-form specification, the exporter wage premium does not have a causal interpretation, because it captures both the non-random selection of high-wage firms into exporting (beyond that captured by firm size) and the impact of exporting on the wage paid by a given firm. In contrast, our structural model below separates out these two components of the exporter wage premium by modeling a firm's endogenous decision to export.

Although the employment size and exporter wage premia are statistically significant in both specifications, the correlation between firm wages, employment and export status is imperfect. After netting out the sector fixed effects, the within R -squared is around 0.15. This pattern of results suggests that there is a systematic component of firm wages (related to firm size and export status) and an idiosyncratic component. While the R -squared of the reduced-form

¹⁸Augmenting regression (6) with firm employment growth has little effect on either the estimated size and exporter wage premia or on the regression fit. In Helpman, Itskhoki, Muendler, and Redding (2012), we show that exporter wage premia are also observed in a panel data specification including firm fixed effects.

Table 7: Size and Exporter Wage Premia, 1994

	UNCONDITIONAL FIRM WAGE COMPONENT, $\hat{\psi}_{jt}^U$	CONDITIONAL FIRM WAGE COMPONENT, $\hat{\psi}_{jt}^C$
Firm Employment Size	0.122*** (0.010)	0.104*** (0.009)
Firm Export Status	0.262*** (0.042)	0.168*** (0.024)
Sector Fixed Effects	yes	yes
Within R-squared	0.165	0.130
Observations	91,410	91,410

Note: Parameter estimates from the cross-section specification (6); *** denotes statistical significance at the 1 percent level; standard errors are heteroscedasticity robust.

regressions in Table 7 suggests that the idiosyncratic component is large relative to the systematic component, this does not rule out changes in the systematic component having economically-meaningful effects on wage inequality. Indeed, changes in the systematic component shift the entire wage distribution and hence can have a substantial effect on overall wage dispersion. Furthermore, the regressions in Table 7 are run at the level of the firm, and hence the estimated exporter wage premium applies to all workers employed by exporting firms. Since exporters account for a large share of employment, the average differences in wages between exporters and non-exporters have a disproportionate effect on the distribution of wages across workers. The next section quantifies the effect of trade on inequality using an estimated model that captures both the systematic and idiosyncratic components of firm wages and hence reproduces the cross-section relationship between firm wages, employment and export status in these regressions.

Fact 4 *Larger firms on average pay higher wages. Controlling for size, exporters on average pay higher wages than non-exporters. Nonetheless, controlling for size and export status, the remaining variation in wage across firms is substantial.*

Taken together, the findings of this section have established a number of key stylized facts that point towards the relevance of recent theories of wage inequality based on firm heterogeneity. Within-sector-occupation inequality accounts for much of overall wage inequality in Brazil. Most of this within-sector-occupation dispersion is residual wage inequality. Furthermore, between-firm variation in wages accounts for most of this residual wage inequality within sector-occupations. Finally, wage variation across firms exhibits robust employment size and exporter wage premia.

4 Structural Model

Guided by the empirical findings in the previous section, we now develop and estimate an extension of the Helpman, Itskhoki, and Redding (2010, HIR henceforth) model. In the HIR

model, wages vary between firms within sector-occupations for workers with similar observed characteristics, and these between-firm differences in wages are correlated with firm employment and export status.¹⁹ Our stylized facts suggest that these are essential ingredients for a model to capture the empirical patterns of inequality and its relationship with trade openness. In what follows we first describe and generalize the HIR model; we then develop a method for estimating this extended model; and lastly we apply the model to the Brazilian data and conduct counterfactuals.

4.1 Theoretical framework

We begin by briefly describing the theoretical framework of HIR, emphasizing the modifications we make in order to take the model to the data. The economy consists of many sectors, some or all of which manufacture differentiated products. The model’s predictions for wages and employment across firms within each differentiated sector hold regardless of general equilibrium effects. Therefore we focus on variation across firms and workers within one such differentiated sector. We are concerned with cross-section dispersion in wages and employment across firms in a given year and develop a static model to characterize such cross-section dispersion.

Within the sector there are a large number of monopolistically competitive firms, each supplying a distinct horizontally-differentiated variety. Demand functions for varieties emanate from constant elasticity of substitution (CES) preferences. As a result, a firm’s revenue in market m (domestic or foreign) can be expressed in terms of its output supplied to this market (Y_m) and a demand shifter (A_m):

$$R_m = A_m Y_m^\beta, \quad m \in \{d, x\},$$

where d denotes the domestic market and x the export market. The demand shifter A_m depends on aggregate sectoral expenditure and the sectoral price index in market m . Since every firm is small relative to the sector, the firm takes this demand shifter as given. The parameter $\beta \in (0, 1)$ controls the elasticity of substitution between varieties.

In order to export, a firm has to incur a fixed cost $e^\varepsilon F_x$, where ε is firm-specific and F_x is common to all firms in the sector. In addition, there are iceberg variable trade costs: $\tau > 1$ units of a variety have to be exported for one unit to arrive in the foreign market. An exporting firm allocates its output between the domestic and export market to maximize revenue. As a result, the firm’s revenue ($R = R_d + R_x$) can be expressed as a function of its output ($Y = Y_d + Y_x$),

¹⁹We focus our econometric analysis on firm exporting rather than firm importing. While the mechanism linking trade and wage inequality in our theoretical model is driven by firm export-market participation as in Melitz (2003), the model can also be extended to capture firm selection into importing as in Amiti and Davis (2012). To the extent that firm importing increases productivity and raises revenue per worker, it results in a similar *importer* wage premium, and our methodology could be applied to this other dimension of firm selection. In practice, firm exporting and importing are strongly positively correlated in the cross section, and hence in our estimation we capture most of the overall effect of firm trade participation.

the demand shifter in the domestic market, and a *market access* variable (Υ_x):

$$R = [1 + \iota(\Upsilon_x - 1)]^{1-\beta} A_d Y^\beta, \quad \text{where} \quad \Upsilon_x = 1 + \tau^{\frac{-\beta}{1-\beta}} \left(\frac{A_x}{A_d} \right)^{\frac{1}{1-\beta}} \quad (7)$$

and ι is an indicator variable, equal to one when the firm exports and equal to zero otherwise. The revenue of a non-exporter is $R = A_d Y^\beta$, while the revenue of an exporter is $R = \Upsilon_x^{1-\beta} A_d Y^\beta$. The firm revenue premium from exporting ($\Upsilon_x^{1-\beta}$) is decreasing in the variable trade cost parameter (τ) and increasing in the foreign demand shifter relative to the domestic demand shifter (A_x/A_d).

We assume that firm output (Y) depends on firm productivity (θ), the measure of workers hired by the firm (H), and the average ability of these workers (\bar{a}):

$$Y = e^\theta H^\gamma \bar{a}, \quad 0 < \gamma < 1. \quad (8)$$

HIR show that this production function can be derived from human capital complementarities (e.g., production takes place in teams and the productivity of a worker depends on the average productivity of her team), or from a model of a managerial time constraint (e.g., a manager with a fixed amount of time who needs to allocate some time to every worker, as in Rosen 1982). Importantly, the production technology (8) exhibits complementarity between the firm's productivity and average worker ability.

Firms and workers are matched in a labor market that exhibits search and matching frictions of the Diamond-Mortensen-Pissarides type. A firm bears a search cost bN in order to randomly match with N workers. The hiring cost b is endogenously determined by the tightness of the labor market and is taken as given by each firm in the sector. In our econometric model, labor market tightness is absorbed in the constants of the estimation equations. For this reason we do not elaborate these details below, and the interested reader can find them in HIR.

Workers are heterogenous in their ability, a , which is drawn from a Pareto distribution $G(a) = 1 - (a_{\min}/a)^k$ for $a \geq a_{\min} > 0$ and $k > 1$. We assume that both firms and workers are *ex ante* equally unaware of the realizations for ability and only know the underlying distribution. In our static model, the worker ability draws admit two interpretations: they can be inherent unobserved worker characteristics or match-specific productivities, as well as any combination of the two. This modeling assumption is consistent with our empirical focus on the component of wages for workers with similar observable characteristics. In HIR we extend the framework to explicitly account for observable worker heterogeneity and different occupations, but here we do not attempt to explain this additional dimension of wage variation and keep the analysis focused on residual wage dispersion within occupations.

Although a firm cannot observe the individual abilities of its N workers, it can invest resources in screening in order to obtain a signal of these abilities. By choosing an ability threshold a_c , a firm can identify workers with abilities below a_c , but it cannot identify the precise ability of each worker. Screening costs increase with the ability threshold and equal $e^{-\eta} C(a_c)^\delta / \delta$, where

η is firm specific while δ and C are common to all firms. We assume $\delta > k$, which ensures a positive equilibrium size-wage premium, as found empirically in the previous section. The incentive to screen workers results from the complementarity of firm productivity and worker abilities in the production function (8), and we show that the more productive firms choose to be more selective in the labor market. Therefore, higher-ability workers are more likely to end up employed by more productive firms, and the model features imperfect (noisy) assortative matching on unobservables in the labor market.

The timing of decisions is as follows. Each firm in a given sector learns its idiosyncratic draw $(\theta, \eta, \varepsilon)$, corresponding to productivity, screening costs, and fixed export costs respectively. Given this triplet, the firm chooses whether to serve only the domestic market or to also export.²⁰ Each firm pays the search costs and matches with the chosen number of workers. After matching, each firm chooses its screening threshold and hires the workers with abilities above this threshold. Therefore, a firm that has searched for N workers and has chosen the ability cutoff a_c hires

$$H = N(1 - G(a_c)) = N \left(\frac{a_{\min}}{a_c} \right)^k \quad (9)$$

workers whose expected ability is

$$\bar{a} = \mathbb{E}\{a | a \geq a_c\} = \frac{k}{k-1} a_c, \quad (10)$$

by the properties of the Pareto distribution. Neither the firm nor its hired workers have information on the abilities of individual workers beyond the fact that they are above the cutoff a_c .

After the firm has paid all the costs—exporting, search and screening—it engages in multilateral bargaining with its H workers over wages, as in Stole and Zwiebel (1996). HIR show that the outcome of this bargaining game is a wage rate

$$W = \frac{\beta\gamma}{1 + \beta\gamma} \frac{R}{H},$$

so that the wage bill is a fixed fraction of firm revenue. Workers who have not been matched with firms, or whose abilities have fallen below their firm's threshold, become unemployed and are not observed in our data.

Anticipating this bargaining outcome, a firm maximizes its profits by choosing the number of workers to match with (N), the screening threshold (a_c), and whether to export:

$$\Pi = \max_{N, a_c, \iota \in \{0,1\}} \left\{ \frac{1}{1 + \beta\gamma} R(N, a_c, \iota) - bN - \frac{C e^{-\eta}}{\delta} (a_c)^\delta - \iota F_x e^\varepsilon \right\},$$

where the revenue function $R(N, a_c, \iota)$ is defined by (7)–(10). HIR show that the solution to

²⁰All firms serve the domestic market since we assume no associated fixed costs. In our empirical implementation, we condition on firm entry into production and analyze a firm's decision to serve the export market and its choices of employment and wages. Therefore we do not model the firm's entry decision here. Similarly, we do not explicitly characterize workers' decisions to search for employment in a given sector, and refer the reader to HIR.

this problem yields (see (S16) in the online supplement to HIR, and the web appendix):

$$R = \kappa_r [1 + \iota(\Upsilon_x - 1)]^{\frac{1-\beta}{\Gamma}} (e^\theta)^{\frac{\beta}{\Gamma}} (e^\eta)^{\frac{\beta(1-\gamma k)}{\delta\Gamma}}, \quad (11)$$

$$H = \kappa_h [1 + \iota(\Upsilon_x - 1)]^{\frac{(1-\beta)(1-k/\delta)}{\Gamma}} (e^\theta)^{\frac{\beta(1-k/\delta)}{\Gamma}} (e^\eta)^{\frac{\beta(1-\gamma k)(1-k/\delta)}{\delta\Gamma} - \frac{k}{\delta}}, \quad (12)$$

$$W = \kappa_w [1 + \iota(\Upsilon_x - 1)]^{\frac{k(1-\beta)}{\delta\Gamma}} (e^\theta)^{\frac{\beta k}{\delta\Gamma}} (e^\eta)^{\frac{k}{\delta} \left(1 + \frac{\beta(1-\gamma k)}{\delta\Gamma}\right)}, \quad (13)$$

and a firm chooses to export in addition to serving the domestic market if and only if:

$$\kappa_\pi \left(\Upsilon_x^{\frac{1-\beta}{\Gamma}} - 1 \right) (e^\theta)^{\frac{\beta}{\Gamma}} (e^\eta)^{\frac{\beta(1-\gamma k)}{\delta\Gamma}} \geq F_x e^\varepsilon, \quad (14)$$

where $\Gamma \equiv 1 - \beta\gamma - \beta(1 - \gamma k)/\delta > 0$ is a derived parameter and the κ_s 's ($s = r, h, w, \pi$) are combinations of variables and parameters that are common to all firms in the sector. Condition (14) requires that the additional profits from exporting net of the fixed exporting cost are positive, and derives from the fact that operational profits are a constant fraction of revenues. When condition (14) holds, $\iota = 1$; otherwise $\iota = 0$.

Equations (12)–(14) describe firm employment, wages, and export participation. This model features two additional sources of firm heterogeneity that do not exist in HIR: heterogeneity in fixed export costs and in screening costs. Without heterogeneous export-market entry costs, a firm's revenue and wage bill would perfectly predict its export status. This prediction is inconsistent with the data, in which there is considerable overlap in the wage and employment distributions between non-exporters and exporters. Some small low-wage firms export in the data, but nonetheless, exporters are on average larger and pay higher wages, consistent with the *selection* and *market access* forces in the model. Without heterogeneous screening costs, employment and wages are perfectly correlated across firms, whereas in the data this correlation is imperfect. Incorporating these two additional sources of heterogeneity enables the model to match the empirical cross-sectional distribution of firm employment, wages and export status.

Our theoretical model predicts that firms with higher productivity θ hire more workers, are more likely to export, and pay higher wages.²¹ Firms with higher screening productivity η are both more selective in the labor market and more profitable, and hence pay higher wages and are more likely to export. However, the effect of screening cost draws on firm employment is more subtle because of two opposing forces. Lower screening costs raise a firm's profitability and result in a larger scale of operation (i.e., increase the number of matches N), but also increase a firm's selectivity in the labor market (reduce the ratio of hires H/N). On net, the effect of lower screening costs on employment is negative.

²¹In this model with Stole-Zwiebel bargaining, equilibrium wages are equalized with the firm's outside option to replace a worker, since the outside option for all workers is unemployment. Firms that are more selective in the labor market have workforces that are more costly to replace and hence end up paying higher wages. Due to complementarity in production, more productive firms choose to be both larger and more selective, and hence pay higher wages. Through this mechanism, exporters are larger and pay higher wages than non-exporters.

The model features a two-way relationship between exporting and firm characteristics. On the one hand, there is a *selection effect*, whereby firms with high productivity have large employment and high wages and are more likely to find it profitable to export. On the other hand, there is a *market access effect*, whereby exporting feeds back into higher firm employment and wages. Access to the foreign market requires a larger scale of production, which is complementary with greater selectivity in the labor market. Hence, exporters have workforces of greater average ability and pay higher wages. In the theoretical literature following Melitz (2003), these two effects are typically not separated because firm productivity perfectly predicts export status. Our framework emphasizes the distinction between these two effects. Using the estimated model, we show that the magnitude of the trade effects on wage inequality depends on the relationship between these two forces.

As discussed in the Introduction, there exist other models of heterogeneous firms with competitive and frictional labor markets—including competitive assignment, search-and-matching, fair-wage and efficiency-wage models—that can generate a firm-size wage premium and exporter premia for wages and employment. We conjecture that our reduced-form specification (12)–(14) captures to a first approximation the predictions for wages, employment and export status of a variety of heterogeneous firm models.²² However, by providing explicit microfoundations for these equations, we ensure that they are consistent with individual optimization and equilibrium, obtain theoretical restrictions on parameters, and are able to undertake model-based counterfactuals.

4.2 Econometric model

Having described the theoretical structure of the model, we now adapt it for empirical estimation. Taking logarithms in (12)–(14) and assuming that the shocks (θ, η, ϵ) are jointly normally distributed, we obtain the following log linear selection model for employment, wages and export status:

$$\begin{cases} h &= \alpha_h + \mu_h \iota + u, \\ w &= \alpha_w + \mu_w \iota + \zeta u + v, \\ \iota &= \mathbb{I}\{z \geq f\}, \end{cases} \quad (u, v, z)' \sim \mathcal{N}\left(\mathbf{0}, \begin{pmatrix} \sigma_u^2 & 0 & \rho_u \sigma_u \\ 0 & \sigma_v^2 & \rho_v \sigma_v \\ \rho_u \sigma_u & \rho_v \sigma_v & 1 \end{pmatrix}\right), \quad (15)$$

where h and w are the natural logarithms of employment and wages respectively and $\mathbb{I}\{\cdot\}$ denotes an indicator function. The reduced-form shocks (u, v, z) are functions of the underlying structural shocks (θ, η, ϵ) as defined in the appendix. The mean-zero normalization for all shocks, the unit-variance normalization for the export-participation shock z and the orthogonality nor-

²²It is important to note that not every model that predicts variation in firm wages, size and export status is capable of explaining the data. For example, if we were to replace the shock to screening in our theoretical model with a shock to the bargaining power of a firm relative to its workers, the resulting correlations would not be consistent with the data. Indeed, the high wage of a firm in this case is a signal of its low bargaining power, which in turn reduces its profitability and hence diminishes its chances to export. In other words, controlling for employment, exporters should be relatively low-wage firms, reflecting their high bargaining power and hence high profitability. Similar logic is likely to rule out the importance of heterogeneity in monitoring ability across firms.

malization for the shocks to employment and wages (u, v) are all without loss of generality. The coefficients μ_h and μ_w capture the *market access* effects of trade on employment and wages. The correlations (ρ_u, ρ_v) between (u, v) and z capture the *selection effects* of high employment and wage firms into exporting.

Our econometric model (15) is the reduced-form of the structural model, and consequently the estimated coefficients, $\Theta \equiv \{\alpha_h, \alpha_w, \zeta, \sigma_u, \sigma_v, \rho_u, \rho_v, \mu_h, \mu_w, f\}$, are reduced-form functions of the parameters of our theoretical model and variables such as trade and labor market costs that are likely to change over time. A reduction in variable trade costs raises firm revenue through improved market access ($\Upsilon_x^{(1-\beta)/\Gamma}$), and hence increases the market access effects for firm employment and wages (μ_h, μ_w) and reduces the export threshold (f). Both the market access effects (μ_h, μ_w) and the export threshold (f) also depend on relative market demands in the domestic and export market (A_x/A_d), and hence can rise and fall over time with fluctuations in the real exchange rate. Since these coefficients of the econometric model are likely to change over time, we estimate the model using cross-section data on firms in each given year.

In the appendix we provide the closed-form relationships between the reduced-form and the structural parameters. Not all structural parameters are identified. Nonetheless, we show that the reduced-form coefficients provide sufficient statistics for the impact of trade on the wage distribution. In other words, the estimated model offers enough information to undertake counterfactuals for the impact of reductions in fixed and variable trade costs on wage inequality, as shown below. Specifically, variable trade costs τ affect only (μ_h, μ_w, f) , through their impact on the market access variable Υ_x , while the fixed exporting cost F_x affects only f .²³ Below we use these relationships between the reduced-form and structural parameters to quantify the effects of reductions in τ and F_x on wage inequality.

We estimate the model using Maximum Likelihood. Our econometric model (15) takes a form similar to a Tobit Type 5 model in Amemiya (1985) or a regression model with endogenous switching in Maddala (1983). A unit of observation in the model is a firm j , and each observation is a triplet of firm log employment, log wages and binary export status, $x_j = (h_j, w_j, \iota_j)'$. The likelihood function of the data $\mathcal{L}(\Theta|\{x_j\}) \equiv \prod_j \mathbb{P}_\Theta \{x_j\}$ admits a simple closed-form expression (see the appendix):

$$\mathbb{P}_\Theta \{x_j\} = \frac{1}{\sigma_u} \phi(\hat{u}_j) \frac{1}{\sigma_v} \phi(\hat{v}_j) \left[\Phi \left(\frac{f - \rho_u \hat{u}_j - \rho_v \hat{v}_j}{\sqrt{1 - \rho_u^2 - \rho_v^2}} \right) \right]^{1-\iota_j} \left[1 - \Phi \left(\frac{f - \rho_u \hat{u}_j - \rho_v \hat{v}_j}{\sqrt{1 - \rho_u^2 - \rho_v^2}} \right) \right]^{\iota_j}, \quad (16)$$

where $\hat{u}_j \equiv (h_j - \alpha_h - \mu_h \iota_j)/\sigma_u$, and $\hat{v}_j \equiv [(w_j - \alpha_w - \mu_w \iota_j) - \zeta \sigma_u \hat{u}_j]/\sigma_v$. The functions $\phi(\cdot)$ and $\Phi(\cdot)$ in (16) are respectively the density and cumulative distribution functions of a standard normal. This simple expression for the density of the data x_j is intuitive: the first two terms reflect the likelihood of the continuous distribution of shocks which result in the observed

²³We have $f = (\alpha_\pi + \log F_x - \log [\Upsilon_x^{(1-\beta)/\Gamma} - 1])/\sigma$, $\mu_h = (1 - k/\delta) \log \Upsilon_x^{(1-\beta)/\Gamma}$ and $\mu_w = (k/\delta) \log \Upsilon_x^{(1-\beta)/\Gamma}$ where Υ_x is defined in (7), and α_π and σ are defined in the appendix. In the theoretical model, $\Upsilon_x > 1$, $0 < k/\delta < 1$, and $\Gamma > 0$, which implies $\mu_h, \mu_w > 0$. Note that $\exp\{\mu_h + \mu_w\} = \Upsilon_x^{(1-\beta)/\Gamma}$, which allows us to use the estimates of (μ_h, μ_w) to isolate the effect of variable trade costs on the export threshold f .

employment and wages, while the last two terms are a standard Probit likelihood for binary export status given the employment and wages of the firm.

Our econometric framework models the effect of exporting on wage and employment outcomes as well as selection into exporting. In selection models of this form, one approach to separately identifying the parameters of the outcome and selection equations is to find an instrument that satisfies the exclusion restriction of only affecting the wage and employment outcomes through export status. Such firm-level instruments are inherently hard to find, because most firm-level variables are endogenously determined by analogy with employment and wages, and hence cannot be used as instruments. We therefore identify the coefficients of the outcome and selection equations using structural identifying assumptions. In particular, we impose the theoretical restriction that the market access effects are positive ($\mu_h, \mu_w > 0$) and invoke a covariance condition on the structural shocks ($\sigma_{\theta\eta} = 0$) as is common in the structural econometrics literature following Koopmans (1949), Fisher (1966) and Wolpin (2013). This orthogonality condition excludes a mechanical relationship between wages and employment based on an arbitrary correlation between the productivity and screening shocks in the wage and employment equations. We show that these structural assumptions identify the parameters of the model and the relative importance of market access effects (μ_h, μ_w) and selection effects (ρ_u, ρ_v). Later we relax the orthogonality condition and impose only the theoretical restriction that the market access effects are positive ($\mu_h, \mu_w > 0$). Under these weaker assumptions, we provide lower and upper bounds to the relative importance of market access and selection effects, and hence to the impact of trade on wage inequality.

Our assumption of orthogonality between the structural shocks θ and η implies two additional inequalities on μ_w/μ_h (see the appendix). We summarize these restrictions in:

$$\mu_h, \mu_w \geq 0, \quad \zeta \leq \mu_w/\mu_h \leq \zeta + \frac{\sigma_v^2}{(1 + \zeta)\sigma_u^2}. \quad (17)$$

The interpretation of these inequalities is as follows. In the model without screening shocks ($\eta \equiv 0$), employment and wages are perfectly correlated, which implies $\zeta = \mu_w/\mu_h = k/(\delta - k)$. The second equality ($\mu_w/\mu_h = k/(\delta - k)$) does not depend on the presence of screening shocks and follows from the general definition of μ_h and μ_w (see footnote 23). The first equality ($\zeta = \mu_w/\mu_h$) becomes an inequality when screening shocks are introduced, because ζ controls the covariance between employment and wages within the groups of non-exporters and exporters, and this covariance becomes weaker with the importance of screening shocks. This explains the role of ζ as the lower bound in (17). At the same time, the difference between μ_w/μ_h and ζ is bounded above by the relative dispersion of the screening and productivity shocks, which explains the upper bound in (17).

Our structural approach contrasts with two alternative identification strategies that have been used in the related empirical literature on trade and labor markets. First, one can rely on time-series shocks to trade costs and export demand (such as MERCOSUR or exchange rate devaluation, as used for example in Verhoogen 2008). While this identification strategy is

attractive in some empirical settings, it is not well suited to our cross-sectional model of wages, employment and export status. Such time-series shocks to trade costs and export demand affect both the market access coefficients (μ_h and μ_w) and selection into exporting (f). Therefore they do not satisfy the exclusion restriction of only affecting wages and employment through selection into exporting. Furthermore, these time-series shocks relate to the dynamics of employment and wages over time, whereas our model is concerned with the contribution of firm selection into export markets towards cross-section wage inequality. Second, following Abowd, Kramarz, and Margolis (1999), another reduced-form literature has used linked employer-employee datasets to seek to separately identify firm, worker and match wage effects. To do so, this literature typically relies on the assumption that the mobility of workers across firms is random after controlling for firm fixed effects, worker fixed effects and time-varying worker observables. While this identification strategy also has attractions, it is again not well suited to our application. Our model features (noisy) assortative matching of workers on unobservables across firms. Therefore the assumption that the time-varying unobserved component of wages is uncorrelated with movements of workers between firms is unlikely to be satisfied in our setting.

4.3 Estimation and model fit

We now turn to the empirical estimation of the model. We first discuss the coefficient estimates and the model fit. We later use the estimated model to undertake model-based counterfactuals and evaluate the contribution of trade to wage inequality. Finally, we close with a discussion of a number of robustness and sensitivity checks, which in particular provide bounds for the effect of trade on wage inequality.

As for our stylized facts above, we pick 1994 as the baseline year for our estimation, because this year is after trade liberalization and before the major appreciation of the Real in 1995. We also estimate the model for all other years and find very similar results, both qualitatively and quantitatively. The data for the baseline estimation consists of firm export status, employment and wages. Our baseline firm wage variable is the firm wage component (ψ_{jt}^C) from the Mincer regression (4), consistent with our focus on residual wage inequality within occupations.

Table 8: Coefficient estimates, 1994

	COEFFICIENT	STD ERROR
μ_h	1.992	0.019
μ_w	0.197	0.022
ρ_u	0.023	0.004
ρ_v	0.199	0.024
f	1.341	0.006

Note: Number of observations: 91,410 (firms). Maximum likelihood estimates and robust (sandwich-form) asymptotic standard errors (see the web appendix).

We start by discussing the coefficients Θ . The key coefficients of interest for the effects of trade on wage inequality are the market access coefficients (μ_h, μ_w), the selection correlation

coefficients (ρ_u, ρ_v) , and the export threshold (f) . In Table 8, we report the estimated values of these coefficients and their standard errors for our baseline year, and the estimates for all other coefficients and for all other years can be found in the web appendix. As shown in the table, we find positive market access and selection effects. Therefore exporting raises both the employment and wages of a given firm $(\mu_h, \mu_w > 0)$ and there is selection of high employment and high wage firms into exporting $(\rho_u, \rho_v > 0)$. The export threshold f captures the fraction of non-exporters, which is equal to $\Phi(f)$ in the model. Small standard errors reported in the second column of Table 8 reflect that all coefficients are estimated precisely given the large size of the sample. In the robustness section below we discuss the values of these parameters under alternative estimation assumptions, as well as contrast the estimated market access coefficients with the reduced-form exporter wage premium.

In the appendix, we show the evolution of these and all other coefficients for each year of our sample. While the selection correlations are relatively stable over time, the market access coefficients and export threshold experience some dynamics, which is consistent with changes in variable trade costs and relative foreign market demand.²⁴ The intercepts α_h and α_w fit the means of firm log employment and wages, and among other things depend on equilibrium variables common to all firms such as the domestic demand shifter and labor market costs. The coefficient ζ controls the covariance between firm employment and wages, and is relatively stable over time. Finally, σ_u accounts for the residual variance of employment, and σ_v for the residual variance of wages conditional on employment and export status. Since employment and wages are not very closely correlated, changes over time in σ_v account for a large part of the fluctuation in wage dispersion. Nonetheless, as we show below, other model parameters also play an important role in shaping wage inequality in both the cross-section and for counterfactual changes in model parameters such as trade costs.

We next examine the model's fit. In Table 9 we report moments in the data and in an artificial dataset simulated using the estimated model. We focus on the first and second moments of the firm employment and wage distributions, both unconditional and conditional on firm export status. These moments provide a good characterization of the overall joint distribution of firm employment, wages and export status, so that an efficient GMM procedure based on the subset of moments reported in Table 9 recovers parameter estimates close to our maximum-likelihood estimates (see the appendix).

Table 9 shows that the model matches all first moments, both conditional and unconditional, as well as the unconditional second moments. The fit of the model is worse for the conditional second moments, in particular for the standard deviations of firm employment and wages among exporters. Indeed, the model does not allow the standard deviations of wages and employment

²⁴Menezes-Filho and Muendler (2011) find net employment reductions in comparative-advantage industries and at exporters. Consistent with these results, we find a decline in the market access effect for employment (μ_h) over time, which reflects in part exchange rate appreciation. Despite these time-series changes, the cross-section data used in our estimation exhibit a substantial positive market access effect for employment (μ_h) in each year in our sample. As a result, we find similar predicted impacts of trade on cross-section wage dispersion for each year in our sample.

Table 9: Firm moments, 1994

	DATA			MODEL		
	All firms	Non-exp.	Exporters	All firms	Non-exp.	Exporters
Mean h	2.96	2.78	4.82	2.96	2.78	4.83
Mean w	-0.33	-0.37	-0.01	-0.33	-0.37	0.00
Std deviation h	1.20	1.00	1.46	1.20	1.05	1.05
Std deviation w	0.43	0.43	0.38	0.43	0.42	0.42
Correlation h & w	0.33	0.24	0.32	0.32	0.25	0.24
Fraction of exporters	9.0%			9.0%		

Note: h is firm log employment and w is firm log wage (conditional firm wage component, ψ_{jt}^C , from (4) used as data in estimation).

Table 10: Moments of worker wage dispersion, 1994

	DATA	MODEL
Std deviation	0.42	0.46
— non-exporters	0.42	0.42
— exporters	0.35	0.42
Gini coefficient	0.23	0.25
90/10-ratio	2.95	3.23
— 90/50	1.63	1.80
— 50/10	1.81	1.80

Note: Each worker is assigned the wage of the firm (conditional firm wage component from 4) to construct the distribution of wages across workers. 90/10-ratio is the ratio of the wages in the 90th percentile and the 10th percentile in the wage distribution (and similar for 90/50 and 50/10).

to differ across exporters and non-exporters, while the data exhibit such variation.

We next examine the model’s ability to fit the moments of the wage distribution across workers. Consistent with the model, we calculate the worker-level moments by assigning the firm wage (firm wage component from (4)) to each worker employed by the firm. Table 10 shows the model’s fit for moments capturing worker wage dispersion—the standard deviation of log wages, Gini coefficient and percentile ratios. The model overpredicts wage dispersion in the upper tail and among exporters, while matching it closely in the lower tail and among non-exporters. Although these moments are complex non-linear transformations of the firm employment and wage distributions that are not targeted directly in the estimation, we find that the model matches these moments relatively closely. Furthermore, the quality of the fit is similar across the different measures of wage inequality. We thus proceed with the remainder of the analysis by using the standard deviation of worker log wages as our main inequality measure, but the results are similar for the other measures of wage inequality.

In Figure 1, we examine the ability of the model to fit the entire distributions of observed employment and wages. In the first panel of the figure, we display kernel densities for firm employment and wages, for all firms and for exporters and non-exporters separately. In the second panel of the figure, we display kernel densities for the distribution of wages across workers,

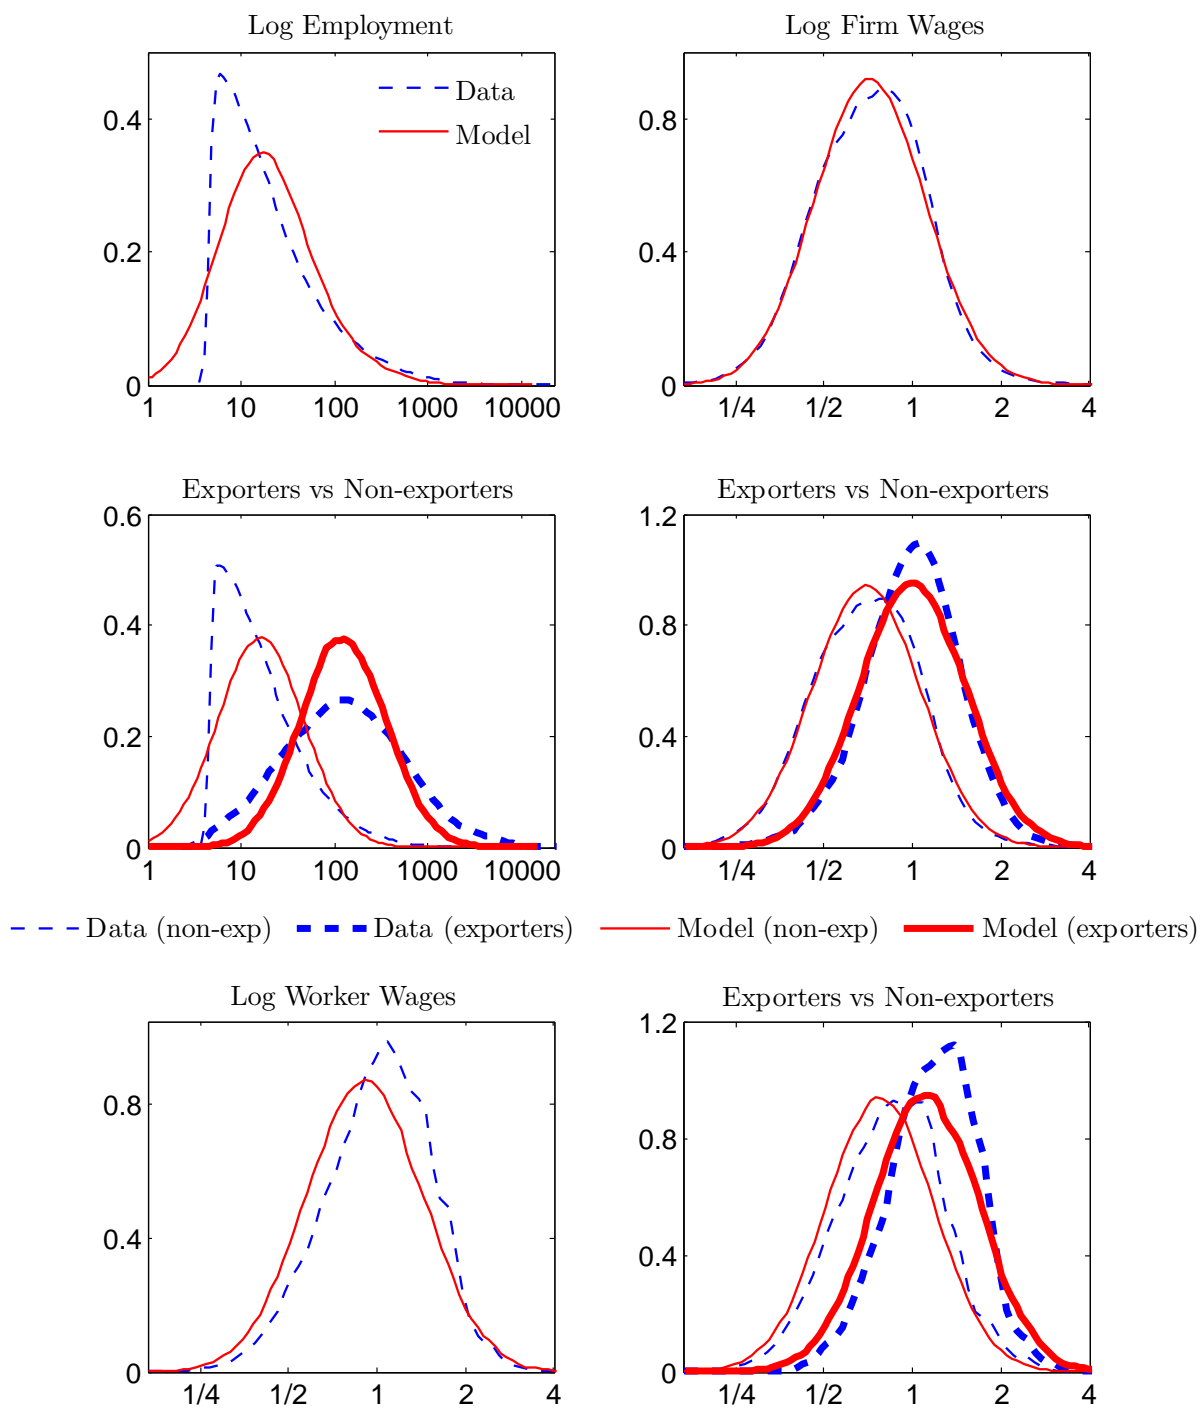


Figure 1: Kernel densities for firm and worker employment and wages, 1994

Note: Log firm employment (scale: number of workers) and log firm and worker wages (scale: multiples of sample average log wage).

Table 11: Employment and exporter wage premia, 1994

	DATA	MODEL
Employment premium	0.10	0.10
Exporter premium	0.16	0.16
<i>R</i> -squared	0.11	0.11

Note: Coefficients and *R*-squared from the regression of firm log wages (firm wage component from (4)) on firm log employment and export status. To ensure the comparability of the results in the data and model, these regressions exclude industry fixed effects (which explains the small difference in estimates from Table 7).

both for all workers and for workers employed by exporters and non-exporters separately. We find that the model is overall successful in fitting these distributions, and in particular captures both the wide overlap in the employment and wage distributions across exporters and non-exporters, as well as the noticeable rightward shift in the employment and wage distributions of exporters relative to non-exporters.²⁵

Finally, we examine the model’s ability to fit the cross-sectional relationship between firm wages, employment and export status in the reduced-form regressions in Table 7. In Table 11, we compare the coefficients and *R*-squared in these regressions in the data to those in the simulated dataset from the estimated model (we again use the firm wage component from (4)). To ensure the comparability of the results in the data and model, these regressions exclude industry fixed effects (which explains the small difference in the estimates between Tables 7 and Table 11). The model matches the employment-size and exporter premia as well as the overall fit of the regressions. In both the data and model, larger firms pay higher wages (with an elasticity of 10 percent) and exporters pay higher wages conditional on their employment size (by 16 percent). In both cases, this relationship is noisy, and hence firm size and export status explain only around 11 percent of the variation in wages. This cross-sectional relationship between firm size, export status and wages is at the core of the trade-and-inequality mechanism that we emphasize in this paper, and hence the ability of the model to replicate this empirical relationship is an important specification check.

Having established the ability of the model to match the empirical patterns of employment, wages and exporting, we next use the estimated model to undertake counterfactuals.

4.4 Counterfactuals

Using the estimated model, we carry out two counterfactual exercises. While we again focus on our benchmark year of 1994, results are similar for all years in our sample. In both counterfactuals we hold all parameters of the model constant apart from trade costs. In the first exercise, we only change the fixed cost of exporting F_x by gradually lowering it from high levels at which no firm exports to low levels at which all firms export. These changes in the fixed cost of ex-

²⁵One noticeable failure in the fit of the distributions is that the employment distribution in the data is more skewed than the log-normal distribution assumed in the structural model. As a result, the model underpredicts the employment share of the exporters, despite matching exactly the fraction of exporting firms.

porting influence wage inequality through the export threshold f in the equation for selection into export markets in (15). In the second exercise, we only change the variable (iceberg) trade cost τ , again lowering it gradually from prohibitively high levels to low levels at which most firms find it optimal to export. These changes in the variable trade cost affect wage inequality through both the export threshold f in the selection equation and the market access premia $\{\mu_w, \mu_h\}$ in the wage and employment equations in (15).²⁶

We report the results of these counterfactuals in the two panels of Figure 2, where we plot the standard deviation of log worker wages against the exporter employment share. For ease of interpretation, the inequality measure in each panel is normalized by its counterfactual value in autarky (equal to 0.43). We vary trade costs and trace out counterfactual wage inequality for all values of the exporter employment share in between zero and one. Therefore we compare the actual level of wage inequality in 1994 to both the counterfactual autarky value and counterfactual values in open economy equilibria in which some or all firms export.

As shown in Figure 2, both counterfactual exercises emphasize a hump-shape relationship between wage inequality and trade openness. This inverted-U-shape pattern is a key theoretical result in the HIR model, in which it is obtained under substantially more stylized assumptions (no heterogeneity in screening or fixed exporting costs and a Pareto distribution of productivity). Therefore, we now confirm that this theoretical conclusion also holds in a substantially richer quantitative model capable of capturing the salient features of the employment and wage distributions for the cross-section of firms.

Quantitatively, we find that the wage inequality predicted by the model for 1994 (corresponding to the red dot in both panels of Figure 2) is 7.6 percent above the counterfactual level of inequality in autarky.²⁷ Interestingly, this corresponds roughly to the peak of inequality when we vary the fixed costs of exporting in the left panel of Figure 2. Therefore, starting from the level of inequality in 1994, further reductions in fixed exporting costs do not lead to additional increases in inequality. This is because at this level of trade openness almost half of the Brazilian manufacturing labor force is employed by exporting firms, and hence further increases in trade participation make the distribution of wages only more equal.

This is not the case, however, for the variation in variable trade costs shown in the right panel of Figure 2. The reason is that reductions in variable trade costs affect wage inequality through both an increase in the extensive margin of export participation and an increase in the intensive margin of the exporter premia for inframarginal exporters. In this case, the peak

²⁶Using the structure of the model, a reduction in variable trade costs corresponds to a proportional increase in the market access coefficients μ_h and μ_w , as well as a reduction in the exporting threshold f in (15). We provide the details of the design of this counterfactual in the appendix. Note that, in the general equilibrium of the model, changes in f_x and τ typically lead to changes in variables common to all firms that are captured by the intercepts α_h and α_w in equation (15). However, the measures of wage inequality that we consider are not sensitive to the values of these intercepts that are common to all firms, and hence our counterfactuals are robust to these general equilibrium considerations (HIR provide further discussion of this point).

²⁷The autarky counterfactual is the most straightforward to calculate, as when no firms export ($\rho_u, \rho_v, \mu_h, \mu_w$) are inconsequential and only the values of $(\zeta, \sigma_u, \sigma_v)$ are required to simulate the joint distribution of employment and wages up to the intercepts (α_h, α_w) , which are common across firms and do not affect the inequality calculations.




Figure 2: Counterfactuals: (a) Fixed export costs and (b) Variable trade cost

Note: The figure plots counterfactual standard deviation of log worker wages (scaled by its counterfactual autarky level equal to 0.43) against exporter employment share in the simulated model. The parameters of the model are held constant at their 1994 estimated values, with the exception of fixed exporting cost F_x in the left panel and variable trade cost τ in the right panel (which are varied from high to low values to cover the full range of exporter employment share). The red dot denotes the points corresponding to the 1994 estimated model.

of inequality is 9.6 percent above the autarky level, and it is reached when over 60 percent of workers are employed by exporters. Finally, around the estimated level of trade openness in 1994 in the right panel of Figure 2, further reductions in variable trade costs that lead to a 10 percentage point increase in the exporter employment share result in about a 2 percent increase in wage inequality.

To put these numbers into perspective, the increase in the relevant component of wage inequality in Brazil between 1986 and 1995 amounted to an increase of around 20 percent in the standard deviation of log worker wages.²⁸ Export participation rises in the early 1990s before reaching a peak in 1993 and declining thereafter. Over these years, the exporter employment share increased by around 10 percentage points, while wage inequality rose by around 4 percent. Our counterfactual exercises have the caveats that we focus solely on the mechanism of firm-based variation in wages and do not consider other possible influences on inequality. Nonetheless, the results of these counterfactuals suggest that the trade-channel in the model can generate sizeable movements in wage inequality.

Bounds of inequality effects We now examine the robustness of our quantitative conclusions. We begin by relaxing the assumption that θ and η are orthogonal. Imposing only the theoretical restriction that the market access effects are positive ($\mu_h, \mu_w > 0$), we provide lower and upper bounds to the relative importance of market access and selection effects, and hence to the impact of trade on wage inequality.

The selection effect suggests that more productive, higher-wage-firms are more likely to

²⁸We quote here the increase in the between-firm within-sector-occupation component of worker wage inequality which is the focus of our model. The increase in overall wage inequality was more modest: the overall standard deviation of log worker wages increased by 8 percent.

Table 12: Alternative specifications

	Estimated coefficients				GMM objective
	μ_h	μ_w	ρ_u	ρ_v	
Theoretical restriction (17)	1.99	0.20	0.02	0.20	0.038
No market access effect ($\mu_w = 0$)	1.92	—	0.06	0.42	0.027
No selection effect ($\rho_v = 0$)	2.04	0.36	0.00	—	0.043
No trade effects ($\mu_h = \mu_w = 0$)	—	—	0.60	0.18	0.149

Note: Theoretical restriction corresponds to our maximum likelihood estimates under the theoretical restriction (17) from Table 8. No-market-access, no-selection, and no-trade-effects lines correspond to maximum likelihood estimates under the restrictions that $\mu_w = 0$, $\rho_v = 0$ and $\mu_h = \mu_w = 0$ respectively. The last column reports the square root of the GMM objective based on eleven first and second moments of employment and wages for exporters and non-exporters (as reported in Table 9) using the diagonal of the efficient weighting matrix (see the text and appendix for details).

select into exporting, and in the model it is captured by the selection correlations ρ_u and ρ_v . In contrast, the market access effect suggests that access to foreign markets induces firms to increase both their employment and wages, and in the model it is captured by the coefficients μ_h and μ_w . The trade effects on inequality are stronger, the stronger is the market access channel relative to the selection channel.²⁹

We now use these properties of the econometric model (15) to provide lower and upper bounds on the importance of the market access and selection effects consistent with the theoretical restriction that $\mu_h, \mu_w \geq 0$. Specifically, we estimate a version of the model in which we shut down the market access effect in the wage equation (setting $\mu_w = 0$ and estimating μ_h) and also a version of the model in which we shut down the wage-equation selection effect (setting $\rho_v = 0$ and estimating ρ_u). We expect the former model to correspond to the lower bound on the effects of trade on inequality and the latter model to correspond to the upper bound.

The first three rows of Table 12 compare our baseline estimates with those for the two limiting cases of the model—with no market access and no selection effects, respectively. The estimated parameters for the wage equation, not surprisingly, vary a lot across the three cases.³⁰ In particular, the market access coefficient μ_w equals 0.2 in our baseline specification, which is about half way between the value of zero under the no-market-access specification and the value of 0.36 under the no-selection specification. These estimates of the market access effect of exporting on wages (μ_w) can be contrasted with the reduced-form estimate of the exporter wage premium of 0.16 in Table 11. This contrast emphasizes that the reduced-form coefficient by itself is insufficient to determine the market access effect, implying that back-of-the-envelope calculations based on the reduced-form estimates are unlikely to be reliable.

Similarly, the selection correlation ρ_v equals 0.20 in our baseline specification, which is again

²⁹Indeed, if the data can be fully explained by the selection mechanism with absolutely no market access effect ($\mu_h = \mu_w = 0$), trade participation merely reflects selection and does not lead to any changes in the wage or employment distributions.

³⁰At the same time, the estimated parameters for the employment equation are not sensitive across specifications: the selection correlation ρ_u is always close to zero and the market access coefficient μ_h is around two.

about half-way between zero under no selection and 0.42 under no market access effects. Therefore, our baseline parameter estimates incorporate both selection and market access forces, and lie squarely within the two bounds without being close to either extreme.

In the last column of Table 12, we report a measure of the model’s fit across specifications, as captured by the square root of a GMM objective function based on the eleven first and second moments for exporters and non-exporters reported in Table 9.³¹ We find that our maximum likelihood estimates under our identifying assumption that θ and η are orthogonal fit the data well, with a value for the GMM objective function equal to 0.038, which implies a cumulative discrepancy between the moments in the model and in the data equal to 3.8 percent of the sample standard deviation of the moments. The no-market-access and no-selection estimates have a similar level of fit with GMM objectives of the same order of magnitude. Notably, all three specifications reproduce the empirical employment-size and exporter wage premia (as in Table 11). Therefore, by considering the no-selection and no-market-access bounds, we do not sacrifice the ability of the model to account for the observed correlation structure between firm employment, export status and wages.

Finally, the last row of Table 12 considers an additional specification, in which we completely shut down the effects of trade on employment and wage distributions by setting both $\mu_h = 0$ and $\mu_w = 0$. Unsurprisingly, this specification generates large estimated selection correlations ρ_u and ρ_v , which enables even this restricted version of the model to reproduce the pattern of empirical correlations of employment, export status and wages. Nonetheless, this specification falls short of matching the average differences in employment size and wages between exporters and non-exporters, and as a result, its fit (as reflected in the GMM objective) is an order of magnitude worse than that of the other three specifications discussed above. Therefore, the data reject this specification in which there are no market access effects of trade on employment and wages. In other words, the data viewed through the prism of our econometric model suggests that trade participation is an important determinant of employment and wage variation across firms.

We use the estimated coefficients from these lower and upper bounds to examine the sensitivity of the model’s quantitative conclusions to the relative importance of the selection and market access forces. In Table 13, the first column of Panel A reports the percentage changes in inequality between the fitted 1994 cross-section and the counterfactual of autarky. We report this comparison for both our baseline estimates and for the no-market-access and no-selection bounds. Our baseline estimates predict 7.6 percent higher wage inequality in the 1994 cross-section than under autarky (as in Figure 2). In comparison, the no-market-access and no-selection bounds predict an increase in wage inequality of 3.3 and 8.3 percent respectively. Therefore, our baseline

³¹Specifically, the GMM objective function we use is the sum of the squared moment conditions scaled by the respective standard deviation of the moment in the data, as discussed further in the appendix. Since we use an overidentified set of moments, the GMM objective is separated from zero even when we use the unconstrained GMM procedure to estimate the model. Specifically, for the 1994 sample in Table 12, the efficient GMM procedure results in the square root of the GMM objective equal to 0.024 which is the comparison benchmark for the last column of the table.

Table 13: Trade effects on inequality

	1994 over autarky	Peak over autarky
PANEL A: INEQUALITY BOUNDS		
Theoretical restriction (17)	7.6%	7.8%
Lower bound (no market access, $\mu_w = 0$)	3.3%	5.0%
Upper bound (no selection, $\rho_v = 0$)	8.3%	8.8%
PANEL B: ROBUSTNESS AND EXTENSIONS		
Sector-region heterogeneity	8.8%	—
Multi-destination A	14.0%	—
Multi-destination B	15.9%	—

Note: Percent changes in inequality (standard deviation of log worker wages) relative to autarky. The first column computes the change in inequality between the predicted values for 1994 and the autarky counterfactual. The second column computes the peak change in inequality, going from autarky to the counterfactual value of f resulting in maximal inequality given the other parameters (i.e., corresponding to the peak in the left panel of Figure 2). The sector-region dummies specification allows the parameters (α_h, α_w, f) to differ across sector-regions (12 sectors \times 136 meso regions). The multi-destination specifications distinguish between exporters serving different numbers of destination markets (see the text for further discussion).

estimates fall within the two bounds, closer to the no-selection bound. It is worth discussing why the no-market-access bound still produces sizeable inequality effects, despite no market access effects of exporting on wages ($\mu_w = 0$). The reason is that there are positive market access effects of exporting on employment ($\mu_h > 0$), so that reductions in trade costs reallocate employment towards larger firms that pay higher wages, inducing an increase in wage inequality.

The second column of Panel A reports an alternative quantification of the bounds on inequality effects by comparing the maximum inequality effects across the three specifications. Specifically, we calculate the counterfactual change between autarky and peak inequality in the open economy.³² We find smaller differences in peak wage inequality across the three specifications: our baseline specification implies peak inequality 7.8 percent above autarky, while the corresponding lower (no-market-access) bound is 5 percent and the corresponding upper (no-selection) bound is 8.8 percent. The reason for this tighter interval is that peak inequality is attained for different levels of trade openness across specifications. Without market access effects, the peak inequality is attained already when 4 percent of firms export, while without selection effects the peak does not happen until 11.5 percent of firms export. As a point of comparison, the fraction of exporters in 1994 is 9 percent, and this roughly corresponds to the peak level of inequality under our baseline estimates.

Therefore, imposing only the theoretical restriction $\mu_h, \mu_w > 0$, the no-market-access and no-selection bounds define a fairly tight interval of the possible effects of trade on wage inequality. Our baseline specification lies approximately in the center of this interval.

³²For this counterfactual, we vary the export threshold f (and hence the fraction of exporting firms) holding other parameters constant (as in the left panel of Figure 2) to find the maximum level of inequality.

Further robustness and extensions We close this section by considering an additional robustness test and extension. One potential concern is that employment and wages could vary systematically across industries and regions in ways that are correlated with export status. To address this concern, we extend our econometric model to allow the constants of the employment and wage equations (α_h, α_w) and the export threshold f to vary by sector-region.³³ In this specification, we find similar market access and selection effects $(\mu_h, \mu_w, \rho_u, \rho_v)$ and similar predicted impacts of trade on wage inequality as in our baseline estimates. As shown in the first row of Panel B of Table 13, predicted wage inequality in 1994 in the model with sector-region heterogeneity is 8.8 percent above the counterfactual autarky level, slightly above the value for our baseline estimates. This suggests that our findings are robust to unobserved heterogeneity across sectors and regions.

More generally, we find an impact of trade on wage inequality that is in line with our baseline estimates across a range of additional robustness tests. For example, when we re-estimate the model for 1990, the predicted increase in wage inequality relative to autarky is 7.4 percent. Additionally, we re-estimate the model using the unconditional firm wage component instead of the conditional firm wage component from the Mincer regression (4), and find a slightly larger counterfactual increase in inequality over autarky equal to 8.8%. This is the case because a portion of the exporter wage premium is accounted for by sorting of workers on observables across firms, which we net out from our baseline analysis focused on residual wage dispersion within occupations. Indeed, the model estimation using the unconditional wage component results in both a slightly larger market access effect μ_w and a slightly larger selection correlation ρ_v .

Finally, we consider an extension in which we refine our definition of exporting. The distinction between exporters and non-exporters potentially abstracts from heterogeneity among exporters. In particular, some Brazilian firms export only to Argentina, whereas others export to multiple destinations, some of which may be remote. In a generalization of the theoretical model to incorporate multiple destinations, each additional export market increases firm market access and hence employment and wages (that is, μ_h and μ_w depend on the set of export markets served), and firms that export to more destinations are on average more productive, larger and pay higher wages (through the selection correlations ρ_u and ρ_v). Since exporting to more destinations raises employment and wages, this generalization is likely to magnify the impact of trade on wage inequality. While a complete treatment of this generalization is outside the scope of this paper, the appendix derives the likelihood function for a simple extension of our econometric model to incorporate multiple destinations. In this extension, we split exporters into three mutually-exclusive bins based on their number of destination markets and estimate separate values of $(\mu_{h,j}, \mu_{w,j})$ for each bin j , as well as separate exporting cutoffs f_j . We report results for two alternative definitions of bins. In specification A, we distinguish between firms exporting to only one destination, 2–5 destinations, and 6 and more destinations. In specifica-

³³Specifically, we estimate separate values for (α_h, α_w, f) for each sector-region bin, which allows the model to perfectly match mean employment and wages, as well as export participation, within each bin. We do so for the 12 sectors and 136 meso-regions, amounting to 1,632 sector-region bins.

tion B, we consider firms exporting to 5 and fewer countries, 6–24 countries, and 25 and more countries.

The last two rows of Table 13 report the results from these multi-destination extensions. We indeed find substantially larger inequality effects: the fitted 1994 cross-section has a level of inequality 14 to 16 percent above the counterfactual autarky level, which is around twice as large as in our baseline estimates. Furthermore, the introduction of a finer partitioning of firms by export status implies greater scope for further increases in wage inequality beyond the levels achieved in Brazil in 1994. As trade costs are reduced further, there is a reallocation of employment not only from non-exporters to exporters but also towards exporters serving more destination markets that are larger and pay higher wages.

5 Conclusion

Using linked employer-employee data for Brazil, we provide evidence on between-firm differences in wages as a mechanism for trade to affect wage inequality in recent theories of heterogeneous firms. We begin by developing a set of stylized facts that provide support for this mechanism. We find that around two thirds of overall wage inequality occurs within sector-occupations. Most of this within-sector-occupation inequality is residual wage inequality. Between-firm wage dispersion accounts for a substantial proportion of this residual wage inequality within sectors and occupations. These between-firm differences in wages are systematically but imperfectly related to trade participation: exporters on average pay higher wages than non-exporters even after controlling for employment.

Guided by these stylized facts, we extend the heterogeneous-firm model of trade and inequality from Helpman, Itskhoki, and Redding (2010) and estimate it using the Brazilian data. This extended model incorporates three dimensions of firm heterogeneity—productivity, fixed exporting costs and worker screening—each of which is central to matching the data. We use the theory to derive an econometric model for wages, employment and export status that features two channels through which trade affects wage inequality: a market access effect (exporting raises the employment and wages of a given firm) and a selection effect (exporting firms are on average larger and pay higher wages than other firms). We use maximum likelihood to estimate the model under our identifying assumptions. We also relax these identifying assumptions to provide lower and upper bounds to the impact of trade on wage inequality for the cases of no-market-access and no-selection effects.

We show that the estimated model matches first and second moments of wages and employment for exporters and non-exporters and approximates the observed distributions of wages and employment for both firms and workers. Our estimates imply sizeable effects of trade on wage inequality, with the opening of the closed economy to trade raising wage inequality by around 10 percent. The estimated model implies a non-monotonic relationship between wage inequality and trade openness, where trade liberalization at first raises and later reduces wage inequality, confirming the theoretical prediction of Helpman, Itskhoki, and Redding (2010).

Although trade expands the set of opportunities for all firms and workers, only some firms find it profitable to take advantage of these opportunities, which is the mechanism driving trade's effect on wage inequality in our model. Enriching the model to introduce a finer partitioning of trading opportunities (e.g., distinguishing between multiple destination markets) magnifies trade's effect on wage inequality through this mechanism.

References

- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67(2), 251–333.
- ABOWD, J. M., F. KRAMARZ, D. N. MARGOLIS, AND K. R. TROSKE (2001): “The Relative Importance of Employer and Employee Effects on Compensation: A Comparison of France and the United States,” *Journal of the Japanese and International Economies*, 15(4), 419–436.
- AMEMIYA, T. (1985): *Advanced Econometrics*. Harvard University Press, Cambridge, Mass.
- AMITI, M., AND D. R. DAVIS (2012): “Trade, Firms, and Wages: Theory and Evidence,” *Review of Economic Studies*, 79(1), 1–36, forthcoming.
- ATTANASIO, O., P. K. GOLDBERG, AND N. PAVCNİK (2004): “Trade Reforms and Wage Inequality in Colombia,” *Journal of Development Economics*, 74(2), 331–366.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2008): “Trends in U.S. Wage Inequality: Revising the Revisionists,” *Review of Economics and Statistics*, 90(2), 300–323.
- BARTH, E., A. BRYSON, J. C. DAVIS, AND R. FREEMAN (2011): “The Contribution of Dispersion across Plants to the Increase in US Earnings Dispersion,” Presented at the EEA-ESEM meetings, unpublished manuscript.
- BAUMGARTEN, D. (2011): “Exporters and the Rise in Wage Inequality: Evidence from German Linked Employer-Employee Data,” RWI, unpublished manuscript.
- BERMAN, E., J. BOUND, AND Z. GRILICHES (1994): “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures,” *Quarterly Journal of Economics*, 109(2), 367–97.
- BERNARD, A. B., AND J. B. JENSEN (1995): “Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987,” *Brookings Papers on Economic Activity: Microeconomics*, 1995(1), 67–112.
- (1997): “Exporters, Skill Upgrading, and the Wage Gap,” *Journal of International Economics*, 42(1-2), 3–31.
- BURDETT, K., AND D. T. MORTENSEN (1998): “Wage Differentials, Employer Size, and Unemployment,” *International Economic Review*, 39(2), 257–273.
- BURSTEIN, A., AND J. VOGEL (2012): “Globalization, Technology, and the Skill Premium: A Quantitative Analysis,” *Columbia University, mimeograph*.
- BUSTOS, P. (2011): “Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms,” *American Economic Review*, 101(1), 30440.
- CAHUC, P., F. POSTEL-VINAY, AND J. M. ROBIN (2006): “Wage Bargaining with On-the-Job Search: Theory and Evidence,” *Econometrica*, 74(2), 323–64.
- COŞAR, A. K. (2010): “Adjusting to Trade Liberalization: Reallocation and Labor Market Policies,” University of Chicago Booth School of Business, unpublished manuscript.
- COŞAR, A. K., N. GUNER, AND J. TYBOUT (2011): “Firm Dynamics, Job Turnover, and Wage Distributions in an Open Economy,” *IMDEA Working Paper*, 2011/06.
- DAS, S., M. J. ROBERTS, AND J. R. TYBOUT (2007): “Market Entry Costs, Producer Heterogeneity, and Export Dynamics,” *Econometrica*, 75(3), 837–873.
- DAVIDSON, C., F. HEYMAN, S. MATUSZ, F. SJÖHOLM, AND S. C. ZHU (2011): “Globalization and Imperfect Labor Market Sorting,” *Research Institute of Industrial Economics Working Paper*, 856, Stockholm.

- DAVIDSON, C., AND S. J. MATUSZ (2010): *International Trade with Equilibrium Unemployment*. Princeton University Press, Princeton, NJ.
- DAVIDSON, C., S. J. MATUSZ, AND A. SHEVCHENKO (2008): “Globalization and Firm Level Adjustment with Imperfect Labor Markets,” *Journal of International Economics*, 75(2), 295–309.
- DAVIS, D. R., AND J. HARRIGAN (2011): “Good Jobs, Bad Jobs, and Trade Liberalization,” *Journal of International Economics*, 84(1), 26–36.
- DAVIS, S. J., AND J. C. HALTIWANGER (1991): “Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-86,” *Brookings Papers on Economic Activity: Microeconomics*, 1991(1), 115–80.
- DIX-CARNEIRO, R. (2013): “Trade Liberalization and Labor Market Dynamics,” Duke University, Unpublished Manuscript.
- EATON, J., S. KORTUM, F. KRAMARZ, AND R. SAMPOGARO (2013): “Firm-to-Firm Trade: Imports, Exports and the Labor Market,” Yale University, unpublished manuscript.
- EGGER, H., P. EGGER, AND U. KREICKEMEIER (2011): “Trade, Wages, and Profits,” University of Tübingen, unpublished manuscript.
- EGGER, H., AND U. KREICKEMEIER (2009): “Firm Heterogeneity and the Labor Market Effects of Trade Liberalization,” *International Economic Review*, 50(1), 187–216.
- FAGGIO, G., K. G. SALVANES, AND J. VAN REENEN (2010): “The Evolution of Inequality in Productivity and Wages: Panel Data Evidence,” *Industrial and Corporate Change*, 19(6), 1919–1951.
- FALLY, T., R. PAILLACAR, AND C. TERRA (2010): “Economic Geography and Wages in Brazil: Evidence from Micro-Data,” *Journal of Development Economics*, 91(1), 155–168.
- FEENSTRA, R. C., AND G. H. HANSON (1996): “Foreign Investment, Outsourcing and Relative Wages,” in *The Political Economy of Trade Policy: Papers in Honor of Jagdish Baghwati*, ed. by R. C. Feenstra, G. M. Grossman, and D. A. Irwin, chap. 6, pp. 89–128. MIT Press, Cambridge, Mass.
- (1999): “The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979-1990,” *Quarterly Journal of Economics*, 114(3), 907–40.
- FELBERMAYR, G., J. PRAT, AND H.-J. SCHMERER (2011): “Globalization and Labor Market Outcomes: Wage Bargaining, Search Frictions, and Firm Heterogeneity,” *Journal of Economic Theory*, 146(1), 39–73.
- FERREIRA, F. H. G., P. G. LEITE, AND M. WAI-POI (2010): “Trade Liberalization, Employment Flows and Wage Inequality in Brazil,” in *The Poor under Globalization in Asia, Latin America and Africa*, ed. by M. Nissanke, and E. Thorbecke, UNU-Wider Studies in Development Economics, chap. 8, pp. 199–254. Oxford University Press, Oxford.
- FISHER, F. (1966): *The Identification Problem in Econometrics*. McGraw-Hill, New York.
- FRÍAS, J. A., D. S. KAPLAN, AND E. A. VERHOOGEN (2009): “Exports and Wage Premia: Evidence from Mexican Employer-Employee Data,” Columbia University, unpublished manuscript.
- GOLDBERG, P. K., AND N. PAVCNİK (2003): “The Response of the Informal Sector to Trade Liberalization,” *Journal of Development Economics*, 72(2), 463–96.
- (2005): “Trade, Wages, and the Political Economy of Trade Protection: Evidence from the Colombian Trade Reforms,” *Journal of International Economics*, 66(1), 75–105.
- (2007): “Distributional Effects of Globalization in Developing Countries,” *Journal of Economic Literature*, 45(1), 39–82.
- GONZAGA, G., N. A. MENEZES-FILHO, AND M. C. TERRA (2006): “Trade Liberalization and the

- Evolution of Skill Earnings Differentials in Brazil,” *Journal of International Economics*, 68(2), 345–67.
- HELPMAN, E., AND O. ITSKHOKI (2010): “Labour Market Rigidities, Trade and Unemployment,” *Review of Economic Studies*, 77(3), 11001137.
- HELPMAN, E., O. ITSKHOKI, M. MUENDLER, AND S. REDDING (2012): “Trade and Inequality: From Theory to Estimation,” *NBER Working Paper*, 17991.
- HELPMAN, E., O. ITSKHOKI, AND S. REDDING (2010): “Inequality and Unemployment in a Global Economy,” *Econometrica*, 78(4), 1239–1283.
- IRARRAZABAL, A., A. MOXNES, AND L. D. OPROMOLLA (2011): “The Margins of Multinational Production and the Role of Intra-firm Trade,” Dartmouth College, unpublished manuscript.
- JUHN, C., K. M. MURPHY, AND B. PIERCE (1993): “Wage Inequality and the Rise in Returns to Skill,” *Journal of Political Economy*, 101(3), 410–442.
- KAMBOUROV, G. (2009): “Labour Market Regulations and the Sectoral Reallocation of Workers: The Case of Trade Reforms,” *Review of Economic Studies*, 76(4), 1321–1358.
- KATZ, L. F., AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963–1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 107(1), 35–78.
- KOOPMANS, T. C. (1949): “Identification Problems in Economic Model Construction,” *Econometrica*, 17(2), 125–44.
- KOVAK, B. K. (2011): “Local Labor Market Effects of Trade Policy: Evidence from Brazilian Liberalization,” Carnegie Mellon University, unpublished manuscript.
- KRISHNA, P., J. P. POOLE, AND M. Z. SENSES (2011): “Trade Liberalization, Firm Heterogeneity, and Wages: New Evidence from Matched Employer-Employee Data,” *World Bank Policy Research Working Paper*, 5711.
- KUME, H., G. PIANI, AND C. F. B. D. SOUZA (2003): “A Política Brasileira de Importação no Período 1987–98: Descrição e Avaliação,” in *A abertura comercial brasileira nos anos 1990: Impactos sobre emprego e salários*, ed. by C. H. Corseuil, and H. Kume, chap. 1, pp. 9–37. MTE and IPEA, Rio de Janeiro.
- LAZEAR, E. P., AND K. L. SHAW (2009): “Wage Structure, Raises and Mobility: An Introduction to International Comparisons of the Structure of Wages Within and Across Firms,” in *The Structure of Wages: An International Comparison*, ed. by E. P. Lazear, and K. L. Shaw, chap. 1. University of Chicago Press, Chicago.
- LEMIEUX, T. (2006): “Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?,” *American Economic Review*, 96(3), 461–498.
- MADDALA, G. S. (1983): *Limited-dependent and qualitative variables in econometrics*, vol. 3 of *Econometric Society Monographs*. Cambridge University Press, Cambridge, UK.
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71(6), 1695–1725.
- MENEZES-FILHO, N. A., AND M.-A. MUENDLER (2011): “Labor Reallocation in Response to Trade Reform,” *NBER Working Paper*, 17372.
- MENEZES-FILHO, N. A., M.-A. MUENDLER, AND G. RAMEY (2008): “The Structure of Worker Compensation in Brazil, with a Comparison to France and the United States,” *Review of Economics and Statistics*, 90(2), 324–346.
- MONTE, F. (2011): “Skill Bias, Trade, and Wage Dispersion,” *Journal of International Economics*, 83(2),

202–218.

- MUNCH, J. R., AND J. R. SKAKSEN (2008): “Human Capital and Wages in Exporting Firms,” *Journal of International Economics*, 75(2), 363–372.
- OI, W. Y., AND T. L. IDSON (1999): “Firm Size and Wages,” in *Handbook of labor economics*, ed. by O. Ashenfelter, and D. Card, vol. 3B, pp. 2165–2214. Elsevier Science, North-Holland, Amsterdam, New York and Oxford.
- POSTEL-VINAY, F., AND J.-M. ROBIN (2002): “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 70(6), 2295–2350.
- POSTEL-VINAY, F., AND H. THURON (2010): “On-the-Job Search, Productivity Shocks and the Individual Earnings Process,” *International Economic Review*, 51(3), 599–629.
- ROSEN, S. (1982): “Authority, Control, and the Distribution of Earnings,” *Bell Journal of Economics*, 13(2), 311–323.
- SAMPSON, T. (2012): “Selection into Trade and Wage Inequality,” London School of Economics, unpublished manuscript.
- SCHANK, T., C. SCHNABEL, AND J. WAGNER (2007): “Do Exporters Really Pay Higher Wages? First Evidence from German Linked Employer-Employee Data,” *Journal of International Economics*, 72(1), 52–74.
- STOLE, L. A., AND J. ZWIEBEL (1996): “Intra-firm Bargaining under Non-binding Contracts,” *Review of Economic Studies*, 63(3), 375–410.
- TREFLER, D., AND S. C. ZHU (2005): “Trade and Inequality in Developing Countries: A General Equilibrium Analysis,” *Journal of International Economics*, 65(1), 21–48.
- VERHOOGEN, E. A. (2008): “Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector,” *Quarterly Journal of Economics*, 123(2), 489–530.
- WOLPIN, K. I. (2013): *The Limits of Inference without Theory*, Tjalling C. Koopmans Memorial Lectures. MIT Press, Cambridge MA.
- YEAPLE, S. R. (2005): “A Simple Model of Firm Heterogeneity, International Trade, and Wages,” *Journal of International Economics*, 65(1), 1–20.