A Mapping of Labor Mobility Costs in Developing Countries

Erhan Artuç

Daniel Lederman

Guido Porto

The World Bank
Poverty Reduction and Economic Management Network
Latin America and the Caribbean Region
Development Research Group
August 2013

Abstract

Estimates of labor mobility costs are needed to assess the responses of employment and wages to trade shocks when factor adjustment is costly. Available methods to estimate those costs rely on panel data, which are seldom available in developing countries. The authors propose a method to estimate mobility costs using readily

obtainable data worldwide. The estimator matches the changes in observed sectoral employment allocations with the predicted allocations from a model of costly labor adjustment. This paper estimates a world map of labor mobility costs and uses those estimates to explore the response of labor markets to trade policy.

This paper is a joint product of International Trade Department, Poverty Reduction and Economic Management Network; the Office of the Chief Economist, Latin America and the Caribbean Region; and the Trade and International Integration Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The authors may be contacted at eartuc@worldbank.org and dlederman@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

A Mapping of Labor Mobility Costs in Developing Countries*

Erhan Artuç [†]	Daniel Lederman [‡]	Guido Porto [§]
DECTI	LAC	Dept. of Economics
The World Bank	The World Bank	UNLP

JEL CODES: F16, D58, J2, J6.

Key Words: **World Labor Mobility Costs, Wage Dynamics, Employment Dynamics**

---

*We want to thank Irene Brambilla for detailed comments. David Jaume provided very valuable research assistance. Comments from seminar participants at Universidad de La Plata and Universidad de San Andres are greatly appreciated. We are also grateful to M. Haddad, B. Hoekman, and A. Mattoo for support with the project. This paper is part of the Trade and Labor Adjustment Costs program led by the International Trade Department of the World Bank (PRMTR).

[†]Development Research Group, The World Bank.

[‡]Latin America and the Caribbean, The World Bank.

[§]Departamento de Economia, Universidad Nacional de La Plata.

1 Introduction

Labor market frictions, such as moving costs, firing-hiring costs, or sector-specific skills, make labor adjustment typically costly.¹ In this setting, a trade shock will only induce a gradual response of wages and employment and this pattern of sluggish labor adjustment has important welfare implications.² The assessment of these labor market responses requires estimates of the costs of labor mobility, but these estimates are seldom available in developing countries. In this paper, our aim is to create a map of estimates of labor mobility costs across the developing world and to use these estimates to explore labor market responses to trade shocks.

We set up a dynamic model of sectoral employment choices and we estimate it for a large sample of developing countries. We adopt the labor adjustment analytical framework of Artuç, Chaudhuri and McLaren (2010), where workers can move across sectors (e.g., in response to wage differences) at a cost. This cost has a common and a idiosyncratic component. The common component captures the average mobility cost of a labor market friction, while the idiosyncratic cost captures worker-specific costs. The parameters governing these costs can only be estimated with panel data, which are hard to find in developing countries.³ To overcome this limitation, we adapt the model and we propose a minimum distance estimator that requires only a time series of cross-sections of sectoral employment and wages—more easily obtainable data. Without the gross flows from the panels, we use net flows to identify the common mobility cost by matching the response of those flows to observed wage differences. We also need to impose a normalization of the idiosyncratic costs. This normalization turns out to be opportune because our model allows for utility compensating differentials across sectors and also because our estimates are robust to small departures from this normalization. In the end, we generate a robust hierarchy of countries based on mobility costs

---

¹Labor immobility is documented in Wacziarg and Wallack (2004), who show little inter-industry flows after liberalization across countries, and Muendler (2010) and Menezes-Filho and Muendler (2011), who show that the absorption of displaced workers from de-protected industries in Brazil was very slow. Labor immobility is also indirectly suggested by the presence of wage differentials, created in part by tariff protection (Attanasio, Goldberg and Pavcnik, 2004; Goldberg and Pavcnik, 2005; Galiani and Porto, 2010).

²The estimation of the impacts of trade liberalization in the presence of imperfect labor mobility is a major ongoing theme in the recent trade literature. Structural models of the dynamics of costly labor adjustment following trade policy and trade shocks include Artuç, Chaudhuri and McLaren (2008, 2010), Coşar (2013), Coşar, Guner and Tybout (2011), Davidson and Matusz (2000; 2004a; 2004b; 2006a; 2006b, 2010), Dix-Carneiro (2013), and Kambourov (2009).

³Panel data provide information on the level of the gross employment flows (which identify the average mobility cost) and on the responsiveness of those flows to the observed wage differentials (which identifies the idiosyncratic component).

caused by labor market frictions. This allows us to assess the responses to trade shock in the presence of costly labor adjustment in a wide array of countries.

We use the United Nations Industrial Development Organization (UNIDO) database, which provides information on labor allocations and wages in manufactures, to estimate a map of the labor mobility costs for 22 developed countries and 25 developing countries. We estimate large costs of labor mobility. On average, the labor mobility costs in developing countries are equivalent to 4.93 times the annual wage. In developed countries, the mobility costs are 2.41—much lower, as expected. The highest costs are estimated in South Asia (5.45), Latin America (5.34), Eastern Europe and Central Asia (4.96), Middle East and North Africa (4.40), Sub-Saharan Africa (4.26), and East Asia and Pacific (3.03). Labor mobility costs are negatively correlated with per capita GDP and positively correlated with poverty rates. They are also inversely correlated with tertiary educational attainments and schooling quality, but are uncorrelated with primary and secondary education enrolment. Finally, mobility costs positively correlate with other frictions, distortions and constraints in the economy.

To illustrate how our estimates of labor mobility costs can be used for policy analysis, we run simulations of the labor market responses to trade liberalization. For each developing country, we separately explore the impacts of a hypothetical decrease in the prices of Food and Beverages and Textiles (due to tariff cuts, for instance). The magnitude of the labor mobility costs matters for the responses of these economies to such a trade shock. Typically, countries only reach close to the steady state after 6 years and the higher the mobility costs are, the longer this transition takes. This imperfect adjustment is costly. We estimate measures of trade adjustment costs and these estimates vary widely across countries. On average, the costs of adjustment to a trade shock in the food sector can be twice as high as the actual gains from trade, while the costs of adjustment to a trade shock in the textiles sector can be equivalent to about 60 percent of the gains from trade.

The rest of the paper is organized as follows. In section 2, we introduce the structural model of labor mobility costs and, in section 3, we discuss the estimation algorithm. The mapping of the estimates of the labor mobility costs is in section 4. Section 5 presents the simulations of the impacts of trade shocks. Finally, section 6 concludes.

2 A Model of Labor Mobility Costs

Our model of labor mobility costs is based on Artuç, Chaudhuri and McLaren (2010). There are N sectors in the economy, M manufacturing sectors and one non-manufacturing sector.⁴ At a given time period, each agent is employed in a sector and earns the sectoral market wage. At the end of each time period t , the agent chooses a sector of employment for the next period, $t + 1$. If the utility differential is larger than the cost of moving, workers move. This determines a new vector of equilibrium labor allocations. We can then estimate the key moving cost parameters by matching the employment predictions of the model with the employment allocations observed in the data.

A worker employed in sector i at time t earns the current sector specific wage w_t^i and enjoys a sector specific (utility) effect denoted by η^i . These η^i can be interpreted as compensating differentials across sectors. Both w_t^i and η^i are common to all workers in a given sector so that there is no worker heterogeneity.⁵ The agent observes both w and η , but only w is observed in the data.

At the end of each time period t , the agent chooses the next period sector of employment based on the expected stream of future wages and on the moving costs. The cost of choosing alternative j for agent l who is currently in sector i is $C + \varepsilon_t^{j,l}$. The “moving cost” has two components, a deterministic part, C , common to all agents, and a random part, $\varepsilon_t^{j,l}$, specific to agent l . All agents are identical except for their individual moving cost shock $\varepsilon_t^{j,l}$, and their current sector. Hence, the state of each agent can be summarized with his/her sector i . We assume that $C = 0$ if agents stay in their current sector. At the end of time t , the random component of the “moving cost,” $\varepsilon_t^{j,l}$, is revealed.

Agents are risk neutral, have rational expectations and a common discount factor $\beta < 1$. Let $U_t^{i,l}$ be the present discounted choice-specific utility of agent l currently employed in sector i . Let V_{t+1}^j be the expected value of U , conditional on the vector of idiosyncratic shocks, $\varepsilon_t^{j,l}$. The Bellman equation is

$$(1) \quad U_t^{i,l} = w_t^i + \eta^i + \max_j \left\{ \beta E_t V_{t+1}^j - C - \varepsilon_t^{j,l} \right\}.$$

---

⁴In other settings, this “residual” sector could also include unemployment or informality. See our discussion below.

⁵Since, as we explain below, we work only with aggregate data, this is an unavoidable assumption. Dix-Carneiro (2013) introduces worker heterogeneity in a related structural model of labor mobility costs.

We now need to solve the model to compute the equilibrium flows of workers across sectors. This solution delivers employment allocations for all sectors i and periods t , and we can thus recover the structural parameters by matching the employment solution of the model with the employment levels observed in the data, our task in section 3. To find the solution, take expectations of (1) with respect to agent specific shocks to get

$$(2) \quad V_t^i = w_t^i + \eta^i + E_t \max_j \left\{ \beta V_{t+1}^j - C - \varepsilon_t^{j,l} \right\}.$$

Dropping the agent superscript l for notational convenience, we can rearrange the value function as

$$V_t^i = w_t^i + \eta^i + \beta E_t V_{t+1}^i + E_t \max_j \{ \varepsilon_t^j + \bar{\varepsilon}_t^i \},$$

where

$$\bar{\varepsilon}_t^i = [\beta E_t V_{t+1}^j - \beta E_t V_{t+1}^i] - C.$$

Then, the choice specific values can be written as

$$(3) \quad V_t^i = w_t^i + \eta^i + \beta E_t V_{t+1}^i + \Omega_t^i.$$

Here, the option value Ω_t^i is equal to

$$\Omega_t^i = \sum_{j=1}^N \int_{-\infty}^{\infty} (\varepsilon^j + \bar{\varepsilon}_t^{ij}) f(\varepsilon^j) \prod_{k \neq j} F(\varepsilon^j + \bar{\varepsilon}_t^{ij} - \bar{\varepsilon}_t^{ik}) d\varepsilon^j,$$

where $F(\varepsilon)$ is the cumulative distribution function and $f(\varepsilon)$ is the probability density function of the moving cost shocks. The option value, Ω_t^i , is the additional utility generated by the possibility to change sectors in the future (and thus to enjoy potential wage differentials). As the moving cost C increases, the option value decreases, and it diminishes to zero when the moving cost goes to infinity.

In principle, the model can be solved for any distributional assumption on $F(\varepsilon)$. As it is standard

in discrete choice model, we assume that ε is distributed iid extreme value type I with location parameter $-\nu\gamma$, scale parameter ν , and cdf $F(\varepsilon) = \exp(-\exp(-\varepsilon/\nu - \gamma))$, where $E(\varepsilon) = 0$, $Var(\varepsilon) = \pi^2\nu^2/6$ and γ is the Euler's constant. This assumption allows us to have an analytical solution for the option value, which becomes $\Omega_t^i = \nu \log \sum_k \exp((E_t V_{t+1}^k - E_t V_{t+1}^i - C_t^k) \frac{1}{\nu})$.⁶ This analytical solution simplifies the Bellman equation and makes it tractable.⁷

We can now derive the employment allocations implied by the model. Let m_t^{ij} be the ratio of agents who switch from sector i to sector j . This can be interpreted as gross flows from i to j , or the probability of choosing j conditional on i . The total number of agents moving from i to j is equal to $y_t^{ij} = L_t^i m_t^{ij}$, where L_t^i is the number of agents who are in i at time t . Under the extreme value distributional assumption, the gross flow m_t^{ij} can be written as

$$(4) \quad m_t^{ij} = \frac{\exp\left(\left(E_t V_{t+1}^j - E_t V_{t+1}^i - C_t^j\right) \frac{1}{\nu}\right)}{\sum_{k=1}^N \exp\left(\left(E_t V_{t+1}^k - E_t V_{t+1}^i - C_t^k\right) \frac{1}{\nu}\right)}.$$

Finally, the allocation of labor between sectors is given by:

$$(5) \quad L_{t+1}^j = \sum_{k \neq j} m_t^{kj} L_t^k + m_t^{jj} L_t^j.$$

There are four key pieces in the model: the common mobility costs, C , the variance of the idiosyncratic mobility costs, ν , the compensating differentials, η^i , and the wage differentials. At each time period, workers compare the costs and benefits of changing employment sector. The benefits are given by the utility differentials, that is the wage differentials net of the compensating differentials. Workers move when the utility differentials are larger than the mobility costs, inclusive of the idiosyncratic benefits. In the end, given the utility differentials, the flow of workers across sectors depends on C and ν , which are the parameters that we want to estimate.

---

⁶See Artu (2012) for the derivation.

⁷See McFadden (1973)

3 Estimation

Our goal in this paper is to create the most comprehensive map of estimates of labor mobility costs possible. Artuç, Chaudhuri and McLaren (2010) derive estimators of these mobility costs based on panel data. The panels are used to build measures of gross employment flows across sectors that, together with utility compensating differentials, identify C , ν , and η^i . In many countries, where the needed panel data is not available, this approach cannot be implemented. For many of these countries, there is availability of time series of sector-level employment and wages, for example in the UNIDO database. In this section, we derive estimators of the mobility costs based on this readily available data.

3.1 The UNIDO Data

It is convenient to begin with a description of the data, because the type of data available imposes some restrictions on the estimation algorithm. In the analysis, we use INDSTAT4, the UNIDO Industrial Statistics Database for the period 1990-2008. The dataset provides information on number of establishments, number of employees, wages, output, value added, and gross fixed capital formation. For our purposes, we need information on employment and wages for the estimation and also on value added for the simulations (see section 5 below). From these data, we build series for the wage streams w_t^i and for the labor allocations L_t^i , for sector i at time t . For the estimation, we aggregate the data into eight major sectors, namely Metals & Minerals, Chemicals & Petroleum Products, Machinery, Food & Beverages, Wood Products, Textiles & Clothing, Miscellaneous Equipments, Motor Vehicles. The UNIDO data has a good coverage of the manufacturing sector but does not cover the non-manufacturing sector. To overcome this limitation, we use national account data to construct measures of labor allocations in the non-manufacturing sector, which we label as sector 1. Note that we do not observe wages for the non-manufacturing sector. In the end, our data comprises time series of employment allocations for both non-manufactures and manufactures, L_t^1 and L_t^i , and wages for manufactures, w_t^i . Wages (utilities) in the the non-manufacturing sector will be calibrated from the data.

3.2 The Estimator

We propose a simulation estimator where we compare the labor allocations simulated with our structural model with the labor allocations observed in the data. Concretely, we define a minimum distance estimator that matches changes in employment allocations for all the manufacturing sectors across time:

$$(6) \quad \hat{C} = \arg \min_C \sum_{t=1}^{T-1} \sum_{i=2}^N \omega_t^i \left(\left(\tilde{L}_{t+1}^i(C; \eta, \mathbf{u}^1) - \tilde{L}_t^i(C; \eta, \mathbf{u}^1) \right) - (L_{t+1}^i - L_t^i) \right)^2,$$

where $\tilde{L}_t^j$ are the employment prediction of the model and ω_t^i are weights used for efficiency. Identification depends on the response of employment allocations to wage differences in the data. In (6), we account for aggregate shocks and for sector-specific fixed effects. However, since we only have aggregate employment and wage series, we cannot control for observed and unobserved heterogeneity. In principle, observed heterogeneity can be dealt with by matching employment and wage series conditional of those characteristics (skills, gender, age). Consequently, our estimates capture a reduced-form association between labor flows and average sectoral wage differences which we interpret as reduced-form estimates of average labor market frictions in the economy. Unobserved heterogeneity can create potential biases (Dix-Carneiro, 2013; Lee and Wolpin, 2006). Self-selection can make the expected wage in a sector to be a poor estimate of the wage differentials for workers employed in other sectors. For example, if the wage actually offered to a worker in sector i is lower than the average observed wage used by our estimator, then labor re-allocation can be small even with large average wage differences. As a result, our estimate of C can be spuriously large. With aggregate data such as UNIDO's, these potential biases are unavoidable.

To implement the estimator, we start with guesses for the mobility costs, C/ν , the compensating differentials for the manufacturing sectors, η^i , and the utility differentials for the non-manufacturing sector, u_t^1 .⁸ Given these guesses, we solve the model with backward iteration.⁹ We first calculate

---

⁸Given the data on employment allocation, and thus the lack of data on gross flows, we can only identify C/ν , the ratio of the common mobility cost C and the variance of the utility shocks ν , but we cannot separately identify these parameters. In the Appendix, we show that this limitation turns out not to be relevant for our purposes.

⁹This requires a finite time horizon assumption, so that the optimization problem ends at time T , and an assumption about expectations, namely $E_t w_\tau = w_\tau$ for any $\tau \geq t$, so that there is no aggregate uncertainty. This restriction allows us to calculate expected values V from observed wages, rather than from expected wages.

the values V_t^i, V_t^1 backwards using (3).¹⁰ We then calculate the gross flows, m_t^{ij} and, with them, we predict the next period labor allocation in sector j , $\tilde{L}_{t+1}^j$,

$$\tilde{L}_{t+1}^j = \sum_{i=1}^K \tilde{L}_t^i m_t^{ij},$$

where $\tilde{L}_1^i = L_1^i$ for $t = 1$. These predictions are then compared with the data and the guesses are updated until convergence. To achieve efficiency, the model is estimated in two steps. We first use the identity matrix as the weighting matrix and we then plug in the residuals from this step in ω_t^i .

We should note here that, in order to simplify the numerical search given the large number of parameters and the limited data, we find the utility differentials u_t^1 and the compensating differentials η^i that simultaneously solve the following system of equations:

$$(7) \quad \tilde{L}_t^1(\mathbf{u}^1, \eta; C/\nu) = L_t^1,$$

for all t , and, for each sector i ,

$$(8) \quad \frac{1}{T} \sum_t \tilde{L}_t^i(\mathbf{u}^1, \eta; C/\nu) = \frac{1}{T} \sum_t L_t^i,$$

where T is the number of time periods in the data. This is a calibration procedure where we match u_t^1 and η such that the predictions of the model match both the whole time series of non-manufacturing sector labor allocations and the average employment for each sector i . The vectors $\mathbf{u}^1$ and η are re-calibrated in each step of the minimization search.¹¹

The variance can be computed numerically. Given the solution $\hat{C}$, let $\hat{\Lambda}_{t+1}^i = \hat{L}_{t+1}^i(\hat{C}; \eta, \mathbf{u}^1) - \hat{L}_t^i(\hat{C}; \eta, \mathbf{u}^1)$ be the predicted employment changes and let $\Lambda_{t+1}^i = L_{t+1}^i - L_t^i$ be the observed employment changes. Define

$$(9) \quad \hat{h}_{t+1}^i = \frac{\partial \hat{\Lambda}_{t+1}^i}{\partial C} e_{t+1}^i,$$

¹⁰Note that, for time T , the values are simply equal to the instantaneous utility, $V_T^i = w_T^i + \eta_i$ and $V_T^1 = w_T^1 + \eta_T^1$.

¹¹Our estimator has a flavor similar to the estimator in Goolsbee and Petrin (2004) in the sense that they match micro-data to recover the parameters and aggregate moments to recover the fixed-effects. The key difference is that Goolsbee and Petrin have much more detailed disaggregated data and thus can estimate rather than calibrate the fixed-effects, as we do here.

where $e_{t+1}^i = \Lambda_{t+1}^i - \hat{\Lambda}_{t+1}^i(C)$ is the residual. Then, the variance of $\hat{C}$ is:

$$(10) \quad \hat{V}(\hat{C}) = \frac{1}{n} \left[\frac{1}{n} \sum_{i,t} \frac{\partial \hat{h}_{t+1}^i}{\partial C} \right]^{-1} \left(\frac{1}{n} \sum_{i,t} (\hat{h}_{t+1}^i)^2 \right) \left[\frac{1}{n} \sum_{i,t} \frac{\partial \hat{h}_{t+1}^i}{\partial C} \right]^{-1},$$

where n is the total number of observations, summing across sectors i and time t . Note that the calculation of the variance requires numerical estimates of the first and second derivatives of $\hat{\Lambda}$.

4 A Mapping of World Mobility Costs

4.1 Descriptive Statistics

Table 1 shows the estimates of the labor mobility costs for 47 countries around the world and Figure 1 introduces our labor mobility cost map. Table 2 reports averages for different groups of countries. For all the countries in the world, the average C is 3.75. This means that, when moving sectors, workers face a common utility cost that is equivalent to 3.75 times the annual average wage in the economy. In general, developing countries show much higher C than developed countries. On average, the mobility cost is 2.41 for developed countries and 4.93 for developing countries, more than twice as large. The lowest costs are estimated in Singapore, the United States, and Japan; the highest costs are estimated in Bangladesh, Ethiopia, Turkey, Azerbaijan, and Peru. Figure 2 shows the distribution of C . As expected, the density for developing countries lies to the right of the one for developed countries. In addition, there is a much larger dispersion in C across developing countries. The density for developed countries is, in contrast, much more concentrated.

To explore differences in C , we report averages for countries by region and by income levels. The lowest labor mobility costs are in North America, 1.65, and in Western Europe, 2.61. In developing countries, the highest average C s are estimated for South Asia (5.45), Latin America (5.34), Eastern Europe & Central Asia (4.96), Middle East & North Africa (4.40), Sub-Saharan Africa (4.26), and East Asia & Pacific (3.03). When countries are grouped by income level, the estimated costs are 2.40 in High income OECD countries, and 2.55 in High income, non-OECD countries. The highest average C , at 6.81, is in Low income countries. The mobility cost in Upper middle income countries

is 4.84 and in Lower middle income countries, 4.68.

4.2 Correlates

We now explore some correlates of the labor mobility costs. We do this by plotting simple bivariate non-parametric regressions between various country characteristics and the size of the estimated mobility costs. We organize this description around four groups of correlates: country well-being, features of the labor market, education, and other constraints. No causality is implied by this analysis, only simple correlations. Data on these correlates are from the World Development Indicators, and they represent averages for the period 1995-2007.

Figure 3 describes the correlation between the mobility costs and some measures of well-being. Richer countries, in terms of per capita GDP, tend to show lower mobility costs (top left panel). However, there is no obvious correlation between the growth rate of per capita GDP and the size of those costs (top right panel). We don't observe any statistical correlation between C and inequality (measured by the Gini coefficient), but there is a very strong positive correlation with both the poverty head-count and the poverty gap.

In Figure 4, we plot C against the structure of employment in agriculture, industry, and services. The correlation is strongly negative with agriculture and strongly positive with industry and services. Furthermore, the top panel of Figure 5 documents a positive correlation between the mobility costs and self-employment, while the bottom panel reveals a positive correlation with the proportion of individuals in vulnerable employment conditions. The implications are that countries that are more highly specialized in non-primary sectors tend to show lower labor mobility costs, as do countries with large self-employment and low job quality.

Figure 6 describes the correlation with educational variables. The top panels and the bottom-left panel show that the labor mobility costs are inversely correlated with the educational attainment of the labor force. This can be more clearly seen in the bottom-left panel, where the negative correlation of C with the share of the labor force with tertiary education is evident. In the bottom-right panel, we plot C against the pupil to teacher ratio in secondary education. The graph shows a positive correlation, thus suggesting the countries with lower education quality (higher pupil-teacher

ratio) tend to also show higher labor mobility costs.

Finally, we plot in Figure 7 the correlations with various other indicators of constraints and distortions. It is not surprising to find an overall positive correlation between the mobility costs and these indicators. For example, labor mobility costs are positively associated with constraints such as business start-up costs, firing costs, procedures to enforce a contract, days to exports, days to import, and time required to start a business. This means that labor market rigidities are more prevalent in countries where other types of rigidities and distortions are also present.

5 The Estimates at Work: Simulating Labor Market Responses

Our estimates of C have a high descriptive value, especially for developing countries. They can be used to characterize and assess differences in labor market frictions across countries. In this section, we illustrate how to use those same estimates to simulate labor market responses to a trade shock and to derive measures of trade adjustment costs. Given the limitations of our data, we work with simple simulations where we shock the price of a sector. We focus on trade shocks to Food and Beverages, but also show simulations to shocks in Textiles. We assume that the price of these goods exogenously and unexpectedly decline by 30 percent and we run independent simulations for all the countries in our sample.¹² All the results that follow should be taken as indications of potential impacts of trade reforms and as a simple illustration of the uses of our estimates of C .

To simulate the economy, we need to add more structure into our model. We specify production and demand functions and we calibrate the initial steady state of the economy. When the economy is hit by the trade shock, we solve for the transition path to the new steady state. Details on the structure of the simulations and the algorithms to find the solutions are in Artuç, Chaudhuri and McLaren (2008, 2010).

On the demand side, we assume Cobb-Douglas preferences

$$(11) \quad u = \prod_g x_g^{\theta_g},$$

---

¹²We do not attempt to develop a global model of trade adjustment. This is, however, doable, if we impose heavy structure to a global model such as Hoekman and Olarreaga (2008).

where x_g is the consumption of good g and θ_g is its share of total expenditure. To be consistent across the paper, we work with nine goods, eight traded goods and the non-traded residual sector. We represent those preferences with data on budget shares compiled by the International Comparison Program.¹³

Production functions are also assumed to be Cobb-Douglas

$$(12) \quad Q_g = A_g L_g^{\alpha_g} K_g^{1-\alpha_g},$$

where Q_g is physical output of good g , A_g is a technology parameter, and L_g and K_g are labor and capital respectively. Labor is imperfectly mobile because, as before, workers can move across sectors after paying the moving costs C . Instead, capital is assumed to be fixed as in Artu, Chaudhuri, and McLaren (2010).¹⁴ The parameters α_g are approximated with the share of wage bill in value added at a sectoral level. Assuming each sector pays a wage equal to the marginal product of labor, we then solve for the technology parameters (including differences in capital)

$$(13) \quad \tilde{A}_g = \frac{1}{\alpha_g} w_g L_g^{1-\alpha_g},$$

where $\tilde{A}_g = A_g K_g^{1-\alpha_g}$. Note that there is an important difference in the treatment of the traded and residual sectors. For the traded sectors, the UNIDO data include wages and employment and thus we can easily recover $\tilde{A}_g$. For the residual sector, we only observe L_g . For the purpose of the initial calibration, we thus set the wage to the average wage of the economy.

We focus first on shocks to Food and Beverages. We report the responses of employment allocations and wages for each of the 25 developing countries in our sample in Figures 8 to 15. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of the affected sector (Food and Beverages) are plotted with a thick line and the responses of the residual sector, with a medium-thick line. To simplify the presentation of the results, we aggregate all the remaining traded sectors into one. These responses are plotted with a thin line. To streamline the exposition, all our results are presented as proportional changes relative to the initial steady

---

¹³Details can be found at http://siteresources.worldbank.org/ICPEXT/Resources/ICP_2011.html

¹⁴See Artu, Bet, Brambilla, and Porto (2012) for simulations with imperfect capital mobility.

state. The transitional dynamics of each country are interesting and revealing in themselves. But rather than attempting to describe all these dynamics, we prepared a typology of responses that are prevalent in our sample.

In all countries, on impact, the real wage in Food and Beverages declines. The decrease in food prices causes a loss of profitability in the sector that translates one to one to nominal wages. The decline in the price index (CPI) is proportional to the food share (which is less than 1). Real wages increase in the rest of the economy, both in the residual sector and in other manufactures, because of the increase in purchasing power (with constant nominal wages and a lower CPI). There are sizeable differences in these initial responses because the weight of food in the price index varies across countries.

The resulting changes in intersectoral wage differentials create incentives for workers to move away from the food sector. The real wage in the sector thus gradually starts increasing. It is remarkable that, in most cases, real wages actually recover and are in fact higher in the new steady state than in the initial steady state. There are only four exceptions, namely Latvia, Romania, Costa Rica, and South Africa. There are, however, significant differences in the time it takes to recover. In Azerbaijan, Russia, or India, for example, the recovery occurs in only 2 years (after the initial wage decline—year 3 of the transition). In Peru, in contrast, it takes 12 years, and in Turkey, 10 years.

Employment in the shocked sector declines, as expected, because firms shrink. It is noteworthy that food workers flow to other traded sectors rather than to the residual sector. In fact, in no country does the residual sector grow. It is not easy to track down why this happens, because the result follows from a combination of the large size of the residual sector, which thus only slightly reacts, of both the initial wage and the compensation differentials, and of the wage responses themselves. The analysis of wage reactions shows that both wages increase (as already pointed out) and that the increase in the wage of the other traded sectors is always larger than the increase in the wage of the residual sector. Only India shows similar responses in these two wages after the second year of the transition. In addition, the real wage in other traded sectors first increases, but then declines. However, in no country are the responses reverted (so real wages in the new

steady state are always higher than in the initial steady state). In the residual, non-traded sector, the initial increase in real wage roughly perpetuates during the whole transition. This is so for all countries and it is the consequence of the size of the sector. Since the residual sector is very large, compared with other traded sectors in these economies, the (often low) inflow of workers does not affect equilibrium wages to a large extent.

Overall, thus, our findings suggest sluggish responses of the labor market, especially in the affected sector, due to labor mobility costs. This can also be seen by computing, for each country, the number of years needed to converge to within 95 percent of the new steady state level of real wages in Food and Beverages.¹⁵ Results are reported in Figure 16. Worldwide, the average convergence speed is 4 years, but it is slower in developing countries (5.44 years) than in developed countries (2.5 years). There is significant variation in speed. In Lithuania, for instance, it takes 12 years to reach 95 percent of the steady state; in Turkey and Azerbaijan, 11 years; in Peru, 10. By contrast, it takes only two years in Latvia, South Africa, Romania, Costa Rica, and Senegal. The convergence speed is increasing with C . For mobility costs C of up to around 4, the convergence speed is constant at 2 years. This includes most developed countries. For C higher than 4, the convergence speed steeply increases with the mobility costs.

5.1 The Gains from Trade and Trade Adjustment Costs

We can also use our model to estimate measures of gains from trade and of trade adjustment costs in the affected sector, Food and Beverages. As our measure of welfare, we use the workers' values, given by V_t in equation (3), which is the present discounted utility for a (random) worker in the food sector at time t . Note that our model generates bilateral flows of workers between sectors during the transition and during the steady state—the difference being that wages change in the transition but the wage differentials are constant in the steady state. In consequence, when we refer to trade adjustment costs for workers in food and beverages, we are making a statement about a random worker that may, or may not switch sectors.

Let V_0 and V_∞ be the welfare of a worker in Food and Beverages in the pre-shock and post-shock steady states. Let V be the present discounted value of the utility of a worker in the food sector

---

¹⁵Other variables, such as employment, produce similar results.

along the transition. The actual gains from trade are given by

$$(14) \quad G = V - V_0.$$

The potential gains from trade are instead given by

$$(15) \quad PG = V_\infty - V_0.$$

The differences between the potential and the actual gains from trade are caused by the costs of labor mobility, which prevent the economy from instantaneously reaching the new steady state. As in Davidson and Matusz (2010), this allows us to estimate Trade Adjustment Costs, TAC, as:

$$(16) \quad TAC = V_\infty - V.$$

To help in interpreting these measures, Figure 17 plots a hypothetical scenario where the pre-shock and post-shock values V_0 V_∞ are the present discounted value of v_0 and v_t , respectively. The scenario features gains from trade ($V_\infty > V_0$) and a transition path that shows an initial decline in welfare and a later recovery. The potential gains from trade are $PG = A + C$, the actual gains are $G = C - B$ and $TAC = A + B$.

Results are reported in Table 3. Column 1 reproduces the level of mobility costs C . In columns 2 and 3, we show the potential and the actual gains from trade as a share of the value in the initial steady state. The potential gains from trade are always positive. All countries stand to gain from lower food prices although the magnitudes vary widely, from a very low 0.55 percent in the United States to almost 20 percent in Azerbaijan. The reason why a decline in the price of food and beverages causes potential welfare gains is twofold. As shown above, nominal wages in F&B decrease and this represents a welfare loss. However, lower food prices raise the real wage, not only in the food sector but also in other sectors of the economy. In the model, the intertemporal welfare of a worker V has two components, the real wage and the option value associated with the probability of future labor choices (possibly implying a switch of sectors). Thus, even though the

real wage of a food worker may decline following the trade shock, the increase in the real wage in the rest of the economy provides an option value that help raise overall welfare. Our results suggest that, for all countries, the new steady state level of intertemporal welfare is actually higher than the pre-shock welfare. Among developing countries, the potential gains from trade are equivalent to 7.5 percent of initial welfare.

The actual gains from trade also tend to be positive. They are also quite large: for developing countries, the gains from trade, as a share of initial welfare, are 5.2 percent. This is surprising because of the initial drop in the real wage that we documented in Figures 8 to 15—which means that, along the transition, welfare first declines but then recovers. The results from our simulations suggest that, in 45 out of 47 countries, the future gains from lower food prices more than compensate the short-run losses (in terms of Figure 17, area C is larger than area A). Among developing countries, only Peru would lose from lower prices. This is because the initial drop in welfare is very pronounced and the recovery takes a long time. The actual gains from trade are smaller than the potential gains from trade because of the costs of labor mobility. In countries with low C , such as the United States or Singapore, the differences are negligible. In countries with high C , such as Peru, Azerbaijan, or Turkey, the differences are sizeable.

In all countries, the trade adjustment costs TAC are positive. This can be seen in column 4 of Table 3, which shows TAC as a share of initial values. In low- C countries, and in fact in most developed countries, these costs are very low. As expected, developing countries with higher C face much larger TAC . In Peru, for instance, TAC are equivalent to 7.25 percent of the initial welfare. In Turkey, Azerbaijan, and Lithuania TAC s are equivalent to 5.92, 5.79 and 5.18 percent of initial welfare, respectively. In Figure 18, note that the level of TAC increases with C . The average TAC for developing countries is 2.3 percent.

In the literature, TAC is typically reported as a share of the total gains from trade (Davidson and Matusz, 2010). This measure gives a sense of the gains from trade that are forgone due to the costs of labor mobility. We report TAC as a share of the potential and actual gains from trade in columns 5 and 6. It is noteworthy that the costs of trade adjustment can represent a very large fraction of the gains from trade. In countries where there are actual losses from trade (Peru and

Denmark), the ratio of TAC to potential gains is actually greater than 1 (because area B is greater than area C). In all the remaining countries, TAC are bounded by PG . The share of TAC in PG varies a lot, from as high as 91.79 in Turkey and 93.87 in Chile to as low as 3.46 in South Africa and 3.5 in Japan. The average for developing countries is 33.

The ratio of TAC to the actual gains G can vary widely. In South Africa and Japan, the ratios are 3.58 and 3.62, respectively. These economies, among others, can quickly enjoy the gains from trade because the costs of adjusting labor, C , are relatively low. In other cases, TAC are huge. In Turkey, TAC can be more than 11 times higher than the gains from trade, and in Chile more than 5 times as high. The average for developing countries is 101.7, so TAC and G are similar, but the median is 27. The size of TAC , in these cases, is dictated by the various parameters of the model, not only by C . For instance, the ratios are roughly the same in Bangladesh, a country with one of the highest C , as in Egypt, a country with the average C for developing countries. The ratios depend on both TAC and G . A large mobility costs makes TAC large and G small, so the ratio tends to be high in these cases. But a large share of food prices in the CPI will create a large G , even for a given TAC , and thus a lower relative importance of trade adjustment costs.

To complete the illustration of some of the uses of C , we end with an analysis of the gains from trade and the trade adjustment costs following a 30 percent decline in the price of textiles.¹⁶ In the UNIDO data, the textile sector tends to be larger than the food sector (in terms of employment). In turn, the weight of textiles and clothing in the consumer price index is much smaller than the weight of food. As a result, there will be larger losses from wage responses in the textile sector, and lower gains from CPI changes. In fact, we find welfare losses across most countries (Table 4). In column 2, the potential gains from trade are negative for most developing countries (with the exception of Chile, Ecuador, Russia, South Africa, and Latvia); in column 3, actual losses from trade are found all over the developing world, except in South Africa and Latvia.

TAC s, as a share of initial welfare, are in column 4. As in the case of Food and Beverages, we find worldwide positive trade adjustment costs in textiles, even when the simulations show welfare costs. This is an interesting result of our model. It is often argued that, when there are losses from trade as in the case of textiles, factor adjustment costs can actually protect workers in the

---

¹⁶A full set of simulation results as in Figures 8 to 15 can be found in the online appendix to this paper.

short-run. In other words, when labor is imperfectly mobile, the gradual adjustment to the steady state may ameliorate the short-run losses associated with the loss of protection. In our model, the short-run welfare loss of a textile worker overshoots the post-shock steady state welfare level so that the short-run losses are actually larger than the long-run losses. Figure 19 presents the case of Bangladesh. This implies positive TAC s. In other words, textile workers are worse off under mobility costs, even if they would lose in the frictionless model.

To end, columns 5 and 6 show that the ratio of TAC to the gains (losses) from trade vary widely across countries. This, as argued above, is not only due to C but also to the whole mixture of factors captured by the parameters of our model, such as CPI textiles shares, the share of labor in textiles production and so on.

6 Conclusions

The premise of our paper is that, in the presence of labor market frictions, trade shocks can have distinct dynamic effects on wages and employment. To explore this idea, we proposed an estimator of labor mobility costs and we built a map of those costs across the developing world. In line with the literature, we estimate large costs of labor mobility, especially for developing countries. These costs, however, vary a lot across countries. They are negatively correlated with various measures of development (per capita GDP, educational attainments) and positively correlated with other frictions, distortions and constraints in the economy.

Our estimates can be used to assess policies. We combined the structural model with the estimated labor mobility costs to simulate the responses of labor markets to trade shocks in developing countries. These simulations allowed us to illustrate the interplay between labor market frictions and trade shocks and to quantify the gains from trade and the trade adjustment costs. We find that transitions are long, 6 years on average, and longer in countries with higher mobility costs. This creates trade adjustment costs that are typically large, and larger for countries with higher costs.

We think about our estimates of labor mobility costs as a useful tool for policy analysis. We show here that our model and estimator work well with readily available aggregate data. Importantly, if more detailed data could be compiled, the estimation and the simulation results can be improved to

account for workers heterogeneity (by for example estimating different costs for skilled and unskilled workers), to incorporate informality, or to better deal with the non-manufacturing sector.

References

Attanasio, O., P. Golberg, and N. Pavcnik (2004). “Trade Reforms and Income Inequality in Colombia,” *Journal of Development Economics*, vol. 74, pp. 331-366.

Artuç, E. (2012). “A Tractable Method for Estimating Dynamic Discrete Models with Policy Shocks,” mimeo, World Bank.

Artuç, E. and J. McLaren (2010). “A Structural Empirical Approach to Trade Shocks and Labor Adjustment: An Application to Turkey,” in G. Porto and B. Hoekman (2010).

Artuç, E. and J. McLaren (2012). “Trade Policy and Wage Inequality: A Structural Analysis with Occupational and Sectoral Mobility,” mimeo World Bank and University of Virginia.

Artuç, E., S. Chaudhuri, and J. McLaren (2008). “Delay and Dynamics in Labor Market Adjustment: Simulation Results,” *Journal of International Economics*, 75(1): 1-13.

Artuç, E., S. Chaudhuri, and J. McLaren (2010). “Trade Shocks and Labor Adjustment: A Structural Empirical Approach,” *American Economic Review*, 100(3), pp. 1008–45.

Artuç, E., G. Bet, I. Brambilla, and G. Porto (2012). “Trade Shocks, Firm-Level Investment Inaction, and Labor Market Responses,” mimeo Universidad de La Plata.

Coşar, A. K. (2013). “Adjusting to Trade Liberalization: Reallocation and Labor Market Policies,” mimeo University of Chicago.

Coşar, A. K., N. Guner, and J. Tybout (2011). “Firm Dynamics, Job Turnover, and Wage Distributions in an Open Economy,” NBER Working Papers 16326, National Bureau of Economic Research, September.

Davidson, C. and S. J. Matusz (2000). “Globalization and Labour-Market Adjustment: How Fast and At What Cost?” *Oxford Review of Economic Policy*, 16: 42-56.

- Davidson, C. and S. J. Matusz (2004a). "An Overlapping-generations Model of Escape Clause Protection," *Review of International Economics*, 12: 749-768.
- Davidson, C. and S. J. Matusz (2004b). *International Trade and Labor Markets: Theory, Evidence, and Policy Implications*, W.E. Upjohn Institute for Employment Research.
- Davidson, C. and S. J. Matusz (2006a). "Long-Run Lunacy, Short-Run Sanity: a Simple Model of Trade with Labor Market Turnover," *Review of International Economics*, 14: 261-276.
- Davidson, C. and S. J. Matusz (2006b). "Trade Liberalization and Compensation," *International Economic Review*," 47: 723-747.
- Davidson, C. and S. J. Matusz (2010). "Modeling, Measuring, and Compensating the Adjustment Costs Associated with Trade Reforms," in G. Porto and B. Hoekman (2010).
- Dix-Carneiro, R. (2013). "Trade Liberalization and Labor Market Dynamics," University of Maryland.
- Galiani, S. and G. Porto (2010). "Trends in Tariff Reforms and in the Structure of Wages," *The Review of Economics and Statistics*, vol. 92, No. 3, pp. 482-494.
- Goldberg, P. and N. Pavcnik (2005). "Trade, Wages, and the Political Economy of Trade Protection: Evidence from the Colombian Trade Reforms," *Journal of International Economics*, vol. 66, pp. 75-105.
- Goolsbee, A. and A. Petrin (2004). "The Consumer Gains from Direct Broadcast Satellites and The Competition with Cable TV," *Econometrica*, Vol. 72, No. 2, pp. 351-381.
- Hoekman, B. and M. Olarreaga (2008). *Global Trade and Poor Nations*. Washington, DC: Brookings Institution Press.
- Kambourov, G. (2009). "Labour Market Regulations and the Sectoral Reallocation of Workers: The Case of Trade Reforms," *Review of Economic Studies*, 76(4), pp. 1321-1358.
- Lee, D. and K. Wolpin (2006). "Inter-sectoral Labor Mobility and the Growth of the Service Sector," *Econometrica*, Vol. 74, No. 1, pp. 1-46.

- Menezes-Filho, N. and M. Muendler (2011). “Labor Reallocation in Response to Trade Reform,” CESifo Working Paper, March.
- McFadden, Daniel (1973). “Conditional Logit Analysis of Qualitative Choice Behavior,” in P. Zarembka (ed.) *Frontiers in Econometrics*, New York, Academic Press.
- McFadden, D. (1989). “A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration,” *Econometrica*, 57(5), pp. 995–1026.
- Muendler, M. (2010). “Trade Reform, Employment Allocation and Worker Flows,” in G. Porto and B. Hoekman (2010).
- Porto, G. and B. Hoekman (2010) (Eds). *Trade Adjustment Costs in Developing Countries: Impacts, Determinants and Policy Responses*, CEPR and World Bank, Washington DC.
- Wacziarg, R. and J. S. Wallack (2004). “Trade Liberalization and Intersectoral Labour Movements,” *Journal of International Economics* 64, pp. 411-439.

APPENDIX: Assessment

In this Appendix, we assess our estimates of C and the implied simulation results. There are two issues to address. First, the estimation algorithm is a mixture of a calibration exercise and a non-linear estimator. We want to show that the calibration embedded in our algorithm is not biasing the estimates of the mobility costs C . Second, since we do not observe labor flows, we can only estimate C/ν . We want to investigate the implications of this limitation both in terms of the cross-country differences in C (section 4) and of the simulation results for a given country (section 5). We tackle these issues sequentially.

Monte Carlo Experiments

To assess the role of the calibration in the estimation of C , we perform the following Monte Carlo experiment. We assume values for the parameters of the economy, namely $C = 3$ and $\nu = 1$ (we also adopt values for η^i and u_t^1). Given these parameters and various shocks to the economy (ε), we simulate data using the structural model. These data mimic the data observed in UNIDO, that is, a time series of labor allocations across sectors and equilibrium wages. We simulate the economy for 15 years (a time series comparable with that in the data). Using this simulated data, we estimate C using our estimation algorithm (which embeds the calibration of η^i and u_t^1). We repeat this experiment 100 times.¹⁷ In each replication, we get estimates of C and its variance. In the experiment, the average value of C is 3.06, with a standard deviation of 0.20. For the variance, the estimated standard error of C is 0.14 (with a standard deviation of 0.13). Even with few replication, thus, the experiment suggests our algorithm works quite well.

C versus C/ν

A concern with our data is that we can only estimate the ratio of the mobility costs, C , to the variance of the utility shocks, ν . This is a consequence of having time series data rather than panel data, which thus prevents us from observing bilateral gross employment flows across sectors. This data limitation makes it impossible to separately identify the two parameters. To see why, recall

---

¹⁷We run only 100 replications in our Monte Carlo because each iteration takes a long computing time.

the moment condition derived by Artuç, Chaudhuri and McLaren (2010):

$$(17) \quad E_t\left[\frac{\beta}{\nu}(\tilde{w}_{t+1}^k - \tilde{w}_{t+1}^j) + \beta(\ln m_{t+1}^{jk} - \ln m_{t+1}^{kk}) - \frac{(1-\beta)}{\nu}C^{jk} - (\ln m_t^{jk} - \ln m_t^{jj})\right] = 0,$$

where $\tilde{w}_t^i$ is the real wage in sector i at time t and β is the discount factor. From this, we infer that the level of employment flows across sectors identifies C/ν (given β) and that the correlation between the response of the flows m to the expected wage difference identifies $1/\nu$ (given β).

Our estimator compares wage and employment fluctuations. If wage differences are large but workers move little (i.e., the net flows are small), then we estimate a high C , given our normalization of ν . However, a large ν can also account for small employment responses to wage fluctuations. The variance of the utility shock, ν , indicates how much workers value wage shocks vis-à-vis idiosyncratic utility shocks (net of the utility compensating differentials). A high ν implies a high variance for these shocks and thus a lower weight on wage shocks. In other words, when ν is large, workers care a lot more about the utility shocks than about wage shocks, so that the relative importance of wage differences decreases. As a result, the employment response to a given wage differential will be small. Note that, with gross flows, this is not a problem because a large ν implies very large gross flows and we can use this information to pin down the separate role of C and ν . Large gross flows imply that workers move from sector i to sector j , and from j to i at the same time, something that is not compatible with a high C .¹⁸

In practice, this limitation is not a source of great concern. Since we are normalizing ν , we are in fact creating a hierarchy of countries based on the labor mobility costs C , which is the average, common, cost of various labor market frictions. Parameter ν is instead related to the idiosyncratic cost of those frictions. We argue that, while there is, a priori, little reason to believe that ν will differ across countries, there are numerous important reasons to believe C will. The correlations uncovered in Figures 3 to 7 are a clear indication of this. The ranking of countries based on our estimates is an arguably reasonable ranking of labor market frictions and this ranking will thus only

---

¹⁸Note that, even with good panel data (meaning long panels with many sectors), while the identification of C is straightforward, the identification of ν is always difficult. Complications in the identification of ν arise if, as it is often the case, there is little variation in observed aggregate wages. Worker heterogeneity may also affect identification. For instance, there might be large changes in individual wages, but if these changes go in different directions, the aggregate wage differentials used to identify ν may be artificially small (Dix-Carneiro, 2011).

fail if we believe that workers in, say, Peru, care less about wage shocks (and more about utility shocks) than workers in Costa Rica do. There might be isolated cases where such a scenario can arise, but this is unlikely to be a systematic problem for our hierarchy.

There is still the issue of the implications of normalizing ν . Even if ν is the same across countries, it could be different from 1. How different can it be and if so, can this affect the estimation of C ? In other words, can the ranking change under a different normalization? First, $\nu \approx 1$ is a reasonable approximation. Artuç (2012) indeed argues that the most plausible estimates of ν using very rich U.S. panel data are actually small, with $\nu = 1$ being in fact a good approximation to the true value. Moreover, there is little variation in ν across specifications. In consequence, our normalization actually provides reasonable estimates for the level of C as well. Second, we can estimate C for alternative values of ν in a small neighborhood of 1, say in the interval $[0.90, 1.1]$. In Table A1, we report the baseline C (column 1), and the estimated C for $\nu = 0.9$ (column 2) and for $\nu = 1.1$ (column 3). As it can be seen, these alternative estimates are typically within 5 to 10 percent of the baseline C . Moreover, the correlation between them, at 0.86 for low ν and 0.99 for high ν , is very high so that the ranking of countries in terms of the mobility costs is robust.

A different concern is whether the normalization of ν affects the results from the simulations. For a given C , if the true ν is higher than 1, then the response of the employment flows to the utility shocks relative to the wage shocks would be underestimated. This means that the response of workers to the wage differential created by a trade shock would be overestimated. In the end, the simulations would produce a more responsive labor market than it really is. If the true ν is instead lower than 1, then the simulations would produce a more sluggish labor market. The evidence in Artuç (2012) based on U.S. data suggests ν could be, if anything, a bit lower than 1. Labor markets would thus be slightly more responsive than what our simulations suggest.

We can assess this directly by re-running the simulations using the estimates of C and ν in the limits of the interval for ν , $[0.90, 1.10]$.¹⁹ We did this for all countries in the sample and results are in an online appendix. Here, we show differences in the dynamics of real wages in Food and Beverages for six countries: the top two countries in terms of C , Peru and Azerbaijan, two countries

---

¹⁹Note that we are re-estimating C for the different values of ν and, as a result, the simulations mix the effects of both changes.

with average C , Egypt and Bolivia, and the bottom two countries in terms of C , South Africa and Latvia. The results are displayed in Figure A1. As it can be seen, for $\nu = 0.9$ or $\nu = 1.1$, the dynamics of the system, given the estimated C/ν , are to a very large extent the same as those for $\nu = 1$. A similar exercise with exactly the same conclusions can be found in Artuç, Chaudhuri, and McLaren (2008). As a consequence, even though the economies will behave differently under different ν , the differences are, both quantitatively and qualitatively, almost irrelevant for our analysis.

Table 1
Labor Mobility Costs

	Mobility Costs	Standard Error
	C	
Peru	7.94	1.938
Azerbaijan	7.81	8.218
Turkey	7.39	0.279
Ethiopia	7.06	0.193
Bangladesh	6.55	0.724
Indonesia	6.20	0.227
Lithuania	6.07	0.619
Chile	5.72	0.683
Ecuador	5.53	1.045
Bulgaria	5.47	0.301
France	5.33	0.646
Denmark	5.08	1.956
Egypt	4.95	1.075
Bolivia	4.93	0.549
Mongolia	4.88	0.538
Russia	4.56	0.680
Iran	4.52	0.462
Georgia	4.42	0.377
Syria	4.40	0.963
India	4.35	0.274
Mauritania	4.33	0.622
Jordan	4.13	0.318
Oman	4.01	0.382
Senegal	3.68	0.087
Poland	2.99	1.198
Czech Republic	2.97	0.698
Greece	2.72	0.099
Belgium	2.57	2.749
Costa Rica	2.56	0.731
Austria	2.46	0.337
Romania	2.40	0.338
Portugal	2.17	0.375
Germany	2.16	0.522
Canada	2.13	0.830
Sweden	2.04	0.644
Finland	2.00	0.399
South Africa	1.95	0.702
Slovakia	1.91	0.630
Norway	1.88	0.366
Spain	1.86	0.050
Ireland	1.81	0.651
Great Britain	1.75	1.260
Korea	1.57	0.770
Latvia	1.55	0.633
Japan	1.42	2.844
United States	1.16	0.496
Singapore	1.09	0.767

Notes: Estimates of labor mobility costs C using UNIDO data.

Table 2
Labor Mobility Costs
Descriptive Statistics

	Observations	Mean	Standard Error
	(1)	(2)	(3)
All countries	47	3.75	1.93
Developed	22	2.41	1.11
Developing	25	4.93	1.72
By region			
Western Europe	16	2.61	1.09
North America	2	1.65	0.69
Eastern Europe & Central Asia	8	4.96	2.21
South Asia	2	5.45	1.56
Latin America & Caribbean	5	5.34	1.93
East Asia & Pacific	5	3.03	2.34
Middle East & North Africa	5	4.40	0.37
Sub-Saharan Africa	4	4.26	2.12
By Income			
High income: OECD	20	2.40	1.07
High income: non-OECD	2	2.55	2.06
Upper middle income	14	4.83	2.13
Lower middle income	9	4.68	0.69
Low income	2	6.81	0.36

Note: Average of labor mobility costs C for different groups of countries.

Table 3
Gains from Trade and Trade Adjustment Costs
Food & Beverages

	<i>C</i>	Gains from Trade		Trade Adjustment Costs		
		Potential	Actual	Initial Value	Potential Gain	Actual Gain
Peru	7.94	4.32	-2.93	7.25	167.75	-247.60
Azerbaijan	7.81	19.58	13.79	5.79	29.57	41.99
Turkey	7.39	6.45	0.53	5.92	91.79	1117.84
Ethiopia	7.06	8.46	5.44	3.02	35.73	55.58
Bangladesh	6.55	15.04	11.14	3.90	25.95	35.04
Indonesia	6.20	13.15	9.15	4.00	30.41	43.70
Lithuania	6.07	8.19	3.01	5.18	63.23	171.94
Chile	5.72	3.62	0.58	3.04	83.97	523.82
Ecuador	5.53	4.83	2.39	2.45	50.66	102.69
Bulgaria	5.47	6.02	3.05	2.97	49.28	97.18
France	5.33	3.40	0.79	2.61	76.78	330.74
Denmark	5.08	2.26	-0.28	2.54	112.63	-891.90
Egypt	4.95	7.47	5.64	1.82	24.42	32.31
Bolivia	4.93	5.44	4.00	1.44	26.51	36.07
Mongolia	4.88	5.71	4.40	1.31	23.02	29.91
Russia	4.56	8.94	7.20	1.74	19.43	24.12
Iran	4.52	5.81	4.71	1.09	18.81	23.17
Georgia	4.42	5.52	4.59	0.93	16.87	20.29
Syria	4.40	8.93	7.58	1.35	15.14	17.84
India	4.35	8.45	7.76	0.69	8.17	8.90
Mauritania	4.33	18.43	16.63	1.81	9.81	10.87
Jordan	4.13	6.32	5.46	0.87	13.72	15.90
Oman	4.01	5.04	4.21	0.83	16.51	19.78
Senegal	3.68	4.52	4.24	0.28	6.20	6.61
Poland	2.99	4.57	4.15	0.43	9.32	10.28
Czech Republic	2.97	4.61	4.18	0.43	9.24	10.18
Greece	2.72	2.02	1.80	0.22	10.89	12.23
Belgium	2.57	2.01	1.83	0.18	9.02	9.92
Costa Rica	2.56	2.26	2.03	0.23	10.26	11.43
Austria	2.46	1.69	1.55	0.15	8.71	9.54
Romania	2.40	4.76	4.55	0.21	4.41	4.62
Portugal	2.17	2.47	2.29	0.17	6.99	7.52
Germany	2.16	1.72	1.65	0.07	4.35	4.54
Canada	2.13	1.31	1.19	0.13	9.53	10.53
Sweden	2.04	1.79	1.67	0.12	6.72	7.20
Finland	2.00	1.89	1.77	0.12	6.23	6.64
South Africa	1.95	2.52	2.43	0.09	3.46	3.58
Slovakia	1.91	2.70	2.58	0.12	4.27	4.46
Norway	1.88	1.45	1.36	0.09	6.39	6.83
Spain	1.86	1.54	1.45	0.09	6.01	6.39
Ireland	1.81	0.57	0.45	0.13	21.82	27.91
Great Britain	1.75	1.02	0.93	0.09	8.45	9.23
Korea	1.57	1.78	1.70	0.08	4.29	4.49
Latvia	1.55	2.18	2.07	0.11	4.86	5.11
Japan	1.42	1.64	1.58	0.06	3.50	3.62
United States	1.16	0.55	0.50	0.06	10.46	11.68
Singapore	1.09	0.78	0.71	0.07	9.39	10.37

Note: Gains from trade and trade adjustment costs following a decrease of 30% in the price of food and beverages.

Table 4
Gains from Trade and Trade Adjustment Costs
Textiles

	<i>C</i>	Gains from Trade		Trade Adjustment Costs		
		Potential	Actual	Initial Value	Potential Gain	Actual Gain
Peru	7.94	-5.14	-6.97	1.83	35.66	26.28
Azerbaijan	7.81	-1.68	-3.24	1.57	93.17	48.23
Turkey	7.39	-3.28	-6.63	3.35	102.33	50.58
Ethiopia	7.06	-3.16	-6.78	3.62	114.58	53.40
Bangladesh	6.55	-8.87	-11.46	2.59	29.25	22.63
Indonesia	6.20	-2.58	-6.26	3.68	142.93	58.84
Lithuania	6.07	-5.95	-8.26	2.31	38.86	27.98
Chile	5.72	0.78	-1.76	2.55	324.51	144.54
Ecuador	5.53	0.39	-1.40	1.79	462.64	127.58
Bulgaria	5.47	-2.13	-4.44	2.31	108.26	51.98
France	5.33	-0.39	-2.31	1.92	485.96	82.93
Denmark	5.08	-0.42	-1.75	1.33	315.51	75.93
Egypt	4.95	-1.35	-2.94	1.58	117.09	53.94
Bolivia	4.93	-0.52	-1.49	0.97	185.57	64.98
Mongolia	4.88	-3.15	-4.72	1.57	49.84	33.26
Russia	4.56	0.07	-0.75	0.82	1233.93	108.82
Iran	4.52	-0.37	-1.35	0.99	267.00	72.75
Georgia	4.42	-0.80	-0.97	0.17	21.43	17.65
Syria	4.40	-2.41	-3.82	1.41	58.49	36.90
India	4.35	-0.70	-1.34	0.64	92.05	47.93
Mauritania	4.33	-10.62	-11.54	0.92	8.68	7.99
Jordan	4.13	-0.16	-0.80	0.63	386.28	79.44
Oman	4.01	0.79	0.43	0.36	45.79	84.46
Senegal	3.68	-0.01	-0.12	0.11	1343.20	93.07
Poland	2.99	-0.56	-0.93	0.37	65.05	39.41
Czech Republic	2.97	-0.03	-0.26	0.23	804.45	88.94
Greece	2.72	0.33	0.17	0.16	49.16	96.68
Belgium	2.57	0.19	0.07	0.11	60.72	154.58
Costa Rica	2.56	-1.07	-1.09	0.02	1.95	1.91
Austria	2.46	0.56	0.47	0.09	15.98	19.02
Romania	2.40	-0.95	-1.09	0.14	15.30	13.27
Portugal	2.17	-0.32	-0.43	0.11	34.79	25.81
Germany	2.16	0.37	0.33	0.05	12.27	13.99
Canada	2.13	0.18	0.13	0.06	31.10	45.14
Sweden	2.04	0.37	0.30	0.07	18.67	22.95
Finland	2.00	0.21	0.15	0.06	28.31	39.50
South Africa	1.95	0.14	0.13	0.01	4.58	4.80
Slovakia	1.91	-0.17	-0.21	0.04	23.72	19.17
Norway	1.88	0.31	0.27	0.04	12.20	13.90
Spain	1.86	0.20	0.16	0.05	23.52	30.75
Ireland	1.81	0.24	0.20	0.04	17.12	20.65
Great Britain	1.75	0.30	0.26	0.04	13.99	16.27
Korea	1.57	-0.11	-0.17	0.06	56.56	36.13
Latvia	1.55	0.10	0.04	0.07	64.72	183.48
Japan	1.42	0.07	0.05	0.03	36.32	57.04
United States	1.16	0.16	0.14	0.03	17.13	20.67
Singapore	1.09	0.09	0.07	0.02	25.20	33.68

Note: Gains from trade and trade adjustment costs following a decrease of 30% in the price of textiles and clothing.

Table A1
Robustness to ν

	C	C with	
		$\nu = 0.9$	$\nu = 1.1$
Peru	7.94	7.38	8.52
Azerbaijan	7.81	7.11	8.52
Turkey	7.39	6.66	8.09
Ethiopia	7.06	6.59	7.21
Bangladesh	6.55	5.91	7.2
Indonesia	6.2	5.74	6.56
Lithuania	6.07	5.5	6.65
Chile	5.72	5.26	6.14
Ecuador	5.53	5.21	5.69
Bulgaria	5.47	5.04	5.87
France	5.33	4.98	5.58
Denmark	5.08	5.03	4.84
Egypt	4.95	4.83	5.01
Bolivia	4.93	4.7	5.07
Mongolia	4.88	4.53	5.21
Russia	4.56	4.29	4.75
Iran	4.52	4.32	4.72
Georgia	4.42	4.17	4.65
Syria	4.4	4.42	3.98
India	4.35	4.21	4.73
Mauritania	4.33	4.18	4.48
Jordan	4.13	3.84	4.37
Oman	4.01	3.77	4.24
Senegal	3.68	9.6	3.95
Poland	2.99	2.8	3.22
Czech Republic	2.97	2.83	3.15
Greece	2.72	2.47	3
Belgium	2.57	2.86	2.27
Costa Rica	2.56	2.57	2.58
Austria	2.46	2.28	2.67
Romania	2.4	2.25	2.58
Portugal	2.17	1.94	2.4
Germany	2.16	2.02	2.3
Canada	2.13	1.98	2.3
Sweden	2.04	1.95	2.14
Finland	2	1.8	2.22
South Africa	1.95	2.51	1.84
Slovakia	1.91	1.74	2.08
Norway	1.88	1.78	2.02
Spain	1.86	1.79	1.97
Ireland	1.81	1.61	2.01
Great Britain	1.75	1.73	1.8
Korea	1.57	1.47	1.67
Latvia	1.55	6.97	1.83
Japan	1.42	1.24	1.6
United States	1.16	1.06	1.26
Singapore	1.09	0.95	1.24

Note: Gains from trade and trade adjustment costs following a decrease of 30% in the price of textiles and clothing.

Figure 1
A Mapping of Labor Mobility Costs

Notes: Estimates of labor mobility costs C using UNIDO data.

Figure 2
Density of Mobility Costs

Notes: Estimates of labor mobility costs C using UNIDO data. Non-parametric density estimators for developed and developing countries.

Figure 3
Correlates of Mobility Costs
Country Well-Being

Notes: Correlations of estimates of labor mobility costs C with per capita GDP, per capita GDP growth rate, the Gini coefficient, the poverty gap, and the poverty head count.

Figure 4
Correlates of Mobility Costs
Features of Labor Markets

Notes: Correlations of estimates of labor mobility costs C with share of employment in agriculture, industry, and services.

Figure 5
Correlates of Mobility Costs
Features of Labor Markets

Notes: Correlations of estimates of labor mobility costs C with the share of self-employed individuals in the labor force and with the share of vulnerable employment.

Figure 6
Correlates of Mobility Costs
Education

Notes: Correlations of estimates of labor mobility costs C with the share of the labor force with primary education, the share of the labor force with secondary education, the share of the labor force with tertiary education, and with the pupil to teacher ratio in secondary school.

Figure 7
Correlates of Mobility Costs
Other Constraints

Notes: Correlations of estimates of labor mobility costs C with firing costs, the number of procedures needed to enforce a contract, the time required to start a business, the time required to export, and the time to import.

Figure 8
Impacts of Trade Shock
Eastern Europe and Central Asia

A) Azerbaijan

B) Bulgaria

C) Georgia

D) Lithuania

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 9
Impacts of Trade Shock
Eastern Europe and Central Asia

A) Latvia

B) Romania

C) Russia

D) Turkey

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 10
Impacts of Trade Shock
South Asia

A) Bangladesh

B) India

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 11
Impacts of Trade Shock
Latin America

A) Bolivia

B) Chile

C) Costa Rica

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 12
Impacts of Trade Shock
Latin America

A) Ecuador

B) Peru

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 13
Impacts of Trade Shock
East Asia and Pacific

A) Indonesia

B) Mongolia

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 14
Impacts of Trade Shock
Middle East and North Africa

A) Egypt

B) Iran

C) Jordan

D) Syria

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 15
Impacts of Trade Shock
Sub-Saharan Africa

A) Ethiopia

B) Mauritania

C) Senegal

D) South Africa

Note: Simulations of responses to a 30 percent price decline in Food and Beverages. Each graph shows six responses using solid lines for wages and dashed lines for employment. The responses of Food and Beverages are plotted with a thick line, the responses of the residual sector, with a medium-thick line, and the responses of the remaining traded sectors, with a thin line. The graphs show proportional changes relative to the initial steady state.

Figure 16
Convergence to Steady State
Number of Years to Within 95%

Notes: convergence.

Figure 17
Trade Adjustment Costs

Notes: The gains from trade and trade adjustment costs. The potential gains from trade are $PG = A + C$, the actual gains are $G = C - B$ and the trade adjustment costs are $TAC = A + B$.

Figure 18
Trade Adjustment Costs
and Labor Mobility Costs

Notes: Correlations of estimates of labor mobility costs C with estimates of Trade Adjustment Costs.

Figure 19
Overshooting in Welfare Losses
Textiles in Bangladesh

Notes: simulation results following a 30% price decline in Textiles in Bangladesh. The short run welfare loss overshoots the post-shock steady state level of welfare, thus creating positive trade adjustment costs even in the presence of losses from trade.

Figure A1
Robustness to ν

A) Peru

B) Azerbaijan

C) Egypt

D) Bolivia

E) South Africa

F) Latvia

Note: Simulations of responses to a 30 percent price decline in Food and Beverages for $\nu = 1$, $\nu = 0.9$ and $\nu = 1.1$.