NET WORK LISER

Discussion Paper Series

IZA DP No. 19019

October 2026

Understanding Migration Responses to Local Shocks

Kirill Borusyak
University of California, Berkeley

Rafael Dix-Carneiro
Duke University

Brian K. Kovak
Carnegie Mellon University
and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

ISSN: 2365-9793

Understanding Migration Responses to Local Shocks*

Abstract

A commonly used migration regression relates population changes to local labor demand shocks, often yielding small or statistically insignificant coefficients despite sizable effects on labor market outcomes. We argue that such findings need not imply that mobility is unresponsive to local economic conditions. Rather, conventional migration regressions can be biased, as they do not reflect how workers choose among competing locations. In a parsimonious equilibrium model, local population growth depends on differences between local shocks and those in relevant destinations; when these shocks are correlated, estimated responses are attenuated. Moreover, with industry-level shocks, standard shift-share measures misrepresent locations' effective exposure, leading to further attenuation. Using simulation experiments based on Brazilian and German administrative data, we show attenuation can exceed 50% and that conventional regression estimates can be highly imprecise. We propose theory-grounded empirical approaches that correct the biases, deliver estimates with a clear structural interpretation, and improve the precision of estimated mobility responses.

JEL classification

R23, J61, C21, F16

Keywords

internal migration, local labor markets, labor demand shocks, shift-share designs, spatial spillovers, sectoral mobility

Corresponding author

Brian K. Kovak
bkovak@cmu.edu

* *First draft:* April 2022. We thank David Autor, Dave Choi, David Dorn, Gordon Hanson, Felix Koenig, Camille Landais, and participants at numerous seminars and conferences for helpful comments, as well as Rodrigo Adão, Greg Howard, Isabela Manelici, and Sharon Traiberman for conference discussions. Borusyak acknowledges financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement Number 949995). Yazen Kashlan, Jacob Lefler, Ilaria Dal Barco, and Garrett Kent provided valuable research assistance. OpenAI's ChatGPT contributed valuable insights. Refine.ink was used to check for consistency and clarity. All errors are our own.

1 Introduction

Workers' interregional migration responses are critical in determining the impacts of local economic shocks. The welfare effects of a given shock depend upon whether workers can smooth adverse labor market outcomes by migrating (Yagan, 2014). The speed of migration adjustment affects how long the impacts of local shocks persist (Topel, 1986; Blanchard and Katz, 1992), and a lack of migration response may help explain long-lasting economic effects of changes in local labor demand (Dix-Carneiro and Kovak, 2017; Dix-Carneiro and Kovak, 2019; Autor et al., 2021). Mobility responses determine how local shocks influence interregional inequality (Topalova, 2010; Cadena and Kovak, 2016), and differences in mobility responses across demographic groups can drive between-group inequality as well (Bound and Holzer, 2000; Wozniak, 2010; Dix-Carneiro and Kovak, 2015). If migration frictions are particularly large, policy makers may consider programs to help workers relocate.

A voluminous literature spanning subfields of economics studies migration responses to local shocks using what we refer to as the "conventional migration regression." This regression relates 5- or 10-year local population growth to observed labor demand shocks, as in the following regression specification:

$$\hat{N}_r = \alpha + \beta \hat{x}_r^{\text{C}} + \varepsilon_r, \tag{1}$$

where $\hat{N}_r$ is the proportional change in population in region r, $\hat{x}_r^{\text{C}}$ (with the superscript C referring to the conventional regression) is an exogenous shock to local labor demand, often with a shift-share structure, and ε_r is an error term. In many cases, researchers find substantial effects of $\hat{x}_r^{\text{C}}$ on local economic outcomes such as wages and employment, but find estimates of β that are small and/or statistically indistinguishable from zero.[1] These results seem to imply a puzzle: despite substantial effects on workers' labor market outcomes, workers do not appear to adjust by relocating.[2] Prior

[1] Examples of papers finding substantial effects of local labor demand shocks on local economic outcomes but with estimates of β that are small and/or statistically indistinguishable from zero include Adão et al. (forthcoming); Autor et al. (2013); Autor et al. (2021); Autor et al. (2023) (except for those age 25-39); Bound and Holzer (2000) (for those with a high-school degree or less); Cadena and Kovak (2016); Choi et al. (2024); Dauth et al. (2014); Dix-Carneiro and Kovak (2017); Dix-Carneiro and Kovak (2019); Faber et al. (2021); Mian and Sufi (2014); Topalova (2010); and Yagan (2019). Additional papers estimating a version of (1) and finding substantial effects in at least some specifications include Albert and Monràs (2020); Albert et al. (2021); Bartik et al. (2019); Black et al. (2005); Boustan et al. (2020); Bustos et al. (2016); Foote et al. (2019); Foschi et al. (2025); Greenland et al. (2019); Hakobyan and McLaren (2016); McCaig (2011); McCaig et al. (2022); Monràs (2020); and Notowidigdo (2020).

[2] For example, Choi et al. (2024) echo a common sentiment in this literature: "... our result deepens the puzzle of why population does not appear to respond to these large, trade-driven employment shocks" (p.1541).

work has responded to these findings by inferring that migration costs are high, that workers are generally unresponsive to changing local economic conditions, or that interregional migration can be ignored in quantitative models.

In this paper, we characterize what the coefficient β in equation (1) measures and propose alternative empirical approaches that better capture the true population responses to observed and counterfactual labor demand shocks. To do so, we develop a parsimonious equilibrium model with mobility frictions across labor markets consisting of region-industry pairs. We interpret β through the lens of the model, contrasting it with the common interpretations invoked in prior work. We then assess model-consistent empirical approaches in simulation experiments based on Brazilian and German administrative data. We find that these approaches avoid biases that plague the conventional regression and outperform it in predicting the effects of observed and counterfactual shocks. Throughout our analysis, we presume that shocks are as-good-as-random, focusing on problems arising with the conventional migration regression even under this best-case scenario.

In the model, workers' decisions depend not only on economic conditions in their current market but also on those in other markets—their outside options. As a result, population growth in a given region r is governed not just by the relevant shocks to markets in that location but by the full constellation of shocks facing all markets. When shocks in the current region are positively correlated with shocks to workers' outside options in other regions, omitting the latter implies that β understates the true elasticity of population to local labor market conditions.

These problems are amplified when using shift-share labor demand measures based on industry-specific shocks, as in the vast majority of empirical work. Frictions in switching industries generate two additional sources of attenuation in this case. First, a worker tied to an industry is exposed to the same industry shock in both the initial region and potential destinations, so migration provides little escape from the shock. Second, we show that standard shift-share constructions mismeasure the relevant exposure to local shocks, compounding attenuation.

Building on this diagnosis, our model delivers tractable, theory-grounded empirical approaches to recover accurate population responses to labor demand shocks. These approaches require data on flows between regions and industries, local population changes, and the underlying shocks. The estimation problem reduces to recovering a single reduced-form parameter via nonlinear least squares (NLLS). With that parameter in hand, we can accurately predict the effects of observed or counter-

factual shocks. We also derive approximations to the model that yield linear regression specifications that can be readily implemented using standard econometric packages. These specifications incorporate information on both the direct shocks facing markets in a location and shocks facing markets elsewhere, reflecting the relevant outside option. In contrast to full "structural estimation" of the model, these reduced-form specifications estimate the relevant combination of structural parameters and have weaker data requirements (by not requiring market-level wage data). While our model-consistent approaches require data on pre-shock worker transitions across region-industry cells, which are often observable only in confidential administrative data, we examine further approximations that are feasible with typical publicly available data.

To assess the quantitative importance of the biases highlighted by the theory, we calibrate our model to real-world economic geography, industry structure, and mobility patterns across regions and industries using longitudinal administrative data on formally employed workers in Brazil and Germany. The main text focuses on Brazil and Appendix D provides results for Germany, with very similar findings in both contexts. Using the model, we simulate population responses to two types of shocks: purely regional shocks, which affect workers in a region equally irrespective of industry, and purely industry shocks, which affect workers in an industry equally irrespective of location, with the latter comprising the overwhelming majority of examples in this literature.[3] We then estimate the conventional migration regression on these simulated outcomes, allowing us to compare regression-based predictions against the true model-implied effects. We find that the conventional approach substantially understates population responses and poorly predicts the effects of labor demand shocks. These problems are most pronounced for industry-level shocks in shift-share designs, where the bias can exceed 50%. We find similar bias when simulating the impacts of real-world tariff changes from Brazil's early-1990s trade liberalization (Dix-Carneiro and Kovak, 2017) and import competition from China and Eastern Europe in Germany (Dauth et al. 2014; Dauth et al. 2021).

We proceed to study the performance of our proposed alternative methods vs. the conventional migration regression in realistic empirical settings. When industry shocks generate highly significant wage effects comparable to those following Brazil's trade liberalization (Dix-Carneiro and

[3]Even shocks that might appear purely regional can be industry-specific under our definition because they do not affect all local industries identically. For example, Bustos et al. (2016) use local soil and weather characteristics indicating suitability for genetically engineered soybeans, and Bartik et al. (2019) use geological features facilitating hydraulic fracturing to extract shale gas. Both variables can be thought of as shift-share exposure measures to a national single-industry shock.

Kovak, 2017), the model implies a cross-regional standard deviation of induced population changes of about 1.5%, which is sizable. In this case, the conventional regression yields estimates that are not only downward biased but also highly imprecise, mirroring the literature's findings of statistically insignificant migration-regression estimates. As a consequence of both misspecification and imprecision, the conventional regression makes poor predictions of the population changes in response to observed shocks, performing little better than predicting zero population change in every region.

Theory-grounded approaches exhibit vastly better performance. A model-consistent NLLS procedure yields estimates of population responsiveness to shocks that exhibit minimal bias and are more precise, regardless of whether the observed shocks are defined by region or industry. Conveniently, linear regressions based on a Taylor approximation to our nonlinear structural equation perform nearly as well as NLLS, making the approach straightforward to implement with standard econometric tools. When the magnitude of the observed industry shocks implies wage effects similar to those induced by the tariff shocks from the Brazilian trade liberalization of the 1990s, however, these approaches still produce statistically insignificant estimates. This does not imply a failure of the method; rather, it reflects that the shocks typically used in the literature are small relative to unobserved drivers of population change, even when those same shocks drive meaningful effects on wages. We therefore recommend using our approach to correct the quantitatively important biases identified above, while recognizing that low signal-to-noise environments may limit the precision of the resulting estimates.

When the full matrix of worker flows across regions and industries is unavailable to researchers, imputations based on more limited flow information do not match the performance of full-matrix approaches but can substantially outperform the conventional regression. Our results show that, while data on region-to-region migration flows are nearly sufficient to evaluate the effects of purely regional shocks, having *some* data on industry switching is indispensable for studying the migration effects of industry shocks in shift-share designs. Fortunately, data on industry-to-industry or even region-to-region-industry moves are commonly available from labor force surveys. Overall, these methods provide a practical path forward for empirical work in data-constrained settings.

In addition to informing empirical practice, our findings have further implications for models of local labor markets. Since conventional migration regression estimates are not informative about

the structural parameters governing mobility, they should not be used to guide modeling choices. Specifically, according to our framework, estimates of β close to zero may reflect the omission of correlated shocks to outside options rather than a true lack of mobility, and therefore do not justify omitting migration from models of local labor markets (e.g. Adão, 2025; Galle et al., 2023). Moreover, our approach elucidates how spatial migration networks lead to differences in the outside options facing workers in different regions. Some studies employ specifications arising from models that incorporate migration choice but have outside options that are common across markets and therefore do not capture the network structure of migration (e.g. Kleven et al., 2013; Agrawal and Foremny, 2019; Notowidigdo, 2020). These approaches are likely to experience issues similar to those we examine in the context of conventional migration regressions.

Our simple model of costly migration follows the tradition of static gravity models of location choice used to analyze both migration and commuting (e.g. Ahlfeldt et al. (2015), Morten and Oliveira (2018), Monte et al. (2018), and Amior and Manning (2019)). In these models, location choice probabilities take a logistic form across destinations, and bilateral costs of moving or commuting generate a network structure; most closely related in the migration context are the formulations by Tombe and Zhu (2019), Fan (2019), and Berkes et al. (2025). Our model also shares with Monte et al. (2018) the implication that local employment responses to labor demand shocks depend on a location's openness to outside labor markets, with heterogeneity driven by commuting shares in their setting and migration shares in ours. We depart from this literature by emphasizing the importance of mobility frictions across industries in addition to across regions, finding that omitting the industry dimension leads to quantitatively important biases.

Methodologically, our paper is related to Adão et al. (forthcoming), who study how trade linkages generate interregional spillovers in the employment and wage effects of increased Chinese imports in the U.S. Like them, we provide a model-consistent reduced-form equation for the effects of labor demand shocks in the presence of spatial spillovers; our analysis focuses on spillovers arising from migration while theirs focuses on spillovers arising from interregional trade.[4] Yet, the papers' objectives are quite distinct, with Adão et al. (forthcoming) seeking to estimate aggregate general equilibrium effects of a spatially heterogeneous labor demand shock, while our analysis focuses

[4]The full model in Adão et al. (forthcoming) features migration, but because it abstracts from bilateral migration costs it does not allow for migration network linkages that vary by location.

on interpreting and addressing the shortcomings of a commonly used reduced-form approach that generally ignores potential spillovers.

The paper proceeds as follows. Section 2 describes the model in which workers face costs of switching between region-industry markets and solves for the population responses to local shocks. Section 3 characterizes the conventional migration regression estimand and discusses problems of using the associated estimates to make within-sample or counterfactual predictions regarding the effects of shocks on regional populations. Section 4 develops empirical approaches to address these problems, including fully model-consistent approaches and alternatives with reduced data requirements. Section 5 presents the Brazilian data and descriptive statistics on worker transitions across regions and industries. Section 6 presents the quantitative results evaluating the conventional regression and our model-based approaches, and Section 7 concludes. Appendix A contains additional theoretical results. Derivations and proofs are relegated to Appendix B, and data construction is described in Appendix C. Appendix D presents quantitative results for Germany, confirming the conclusions of the Brazilian analysis.

2 Model of Mobility with Regional and Industry Frictions

2.1 Setup

This section presents a parsimonious model that yields clear and intuitive expressions for how labor demand shocks affect local populations in the presence of costly mobility. Consider a global economy consisting of R_w regions, a.k.a. locations, indexed by $r, o, d, k \in \mathcal{R}_w$ and I industries, a.k.a. sectors, indexed by $i, p, q \in \mathcal{I}$. For brevity, we will refer to region-industry combinations as "markets." Products, a.k.a. varieties, within each industry are differentiated by location of production, as in Armington (1969), Anderson and van Wincoop (2003), Head and Mayer (2014), and Adão et al. (2019), among many others. Product markets are frictionless, so the price of the industry-i good produced in region r, p_{ri}, is constant across locations of consumption. Homogeneous labor is the only input, and region-industry productivity is φ_{ri}.[5] Output and labor markets are perfectly competitive, so $p_{ri} = w_{ri}/\varphi_{ri}$, where w_{ri} is the wage in market r, i.

Individuals in all countries make optimal choices over consumption bundles q_{ri}. The representa-

[5] Appendix A.1 shows that allowing for agglomeration economies leads to isomorphic results.

tive consumer's utility from consumption is Cobb-Douglas across industries and constant elasticity of substitution (CES) across varieties within each industry:[6]

$$U = \sum_{i \in \mathcal{I}} \eta_i \log(Q_i), \quad \text{where } Q_i \equiv \left(\sum_{r \in \mathcal{R}_w} q_{ri}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{2}$$

$\sigma > 1$, and $\sum_i \eta_i = 1$. Normalizing world aggregate expenditure to one, the global demand for the variety of industry i produced in region r satisfies

$$q_{ri} = \eta_i p_{ri}^{-\sigma} P_i^{\sigma-1}, \tag{3}$$

where $P_i \equiv \left(\sum_{r \in \mathcal{R}_w} p_{ri}^{1-\sigma} \right)^{1/(1-\sigma)}$ is the CES exact price index for industry i.

A subset of R regions $\mathcal{R} \subset \mathcal{R}_w$ comprises the country of interest. This country is a small open economy in the sense that P_i is set on the world market and is exogenous to developments in the country of interest. Unless explicitly specified otherwise, sums over regions cover only the domestic ones. There is a mass N of domestic workers who are exogenously assigned to different markets, with L_{op}^0 workers originally in region o and industry p. They can switch to any other market of employment but face mobility frictions across both regions and industries. We assume that workers may move within the country of interest by choosing among markets $(r, i) \in \mathcal{R} \times \mathcal{I}$, but not internationally. Once the market is chosen, each worker inelastically supplies one unit of labor. The choice of the market for individual ω originally in market o, p maximizes the indirect utility of selecting destination market r, i, given by[7]

$$V_{\omega,ri} = \log\left(\frac{w_{ri}}{P}\right) - \log \tau_{op,ri} + \frac{1}{\theta} \epsilon_{\omega,ri}, \tag{4}$$

where $P \equiv \Pi_i \left(P_i/\eta_i\right)^{\eta_i}$ is the Cobb-Douglas exact price index across industries, $\log \tau_{op,ri}$ is the utility cost of moving from market o, p to r, i, and $\epsilon_{\omega,ri}$ is a taste shock following a type-I extreme-value (standard Gumbel) distribution that is iid across individuals and markets. The parameter $\theta > 0$ determines the strength of idiosyncratic preferences relative to the log real wage, with smaller values of θ implying stronger preference over markets and, as shown below, smaller employment

[6] Appendix A.2 adds a residential housing market to the model (simplified to have $I = 1$), showing very similar analytical results.

[7] Exogenous market-specific amenities can be incorporated without changing the analysis. Since our procedure does not rely on data on wages, w_{ri} can be reinterpreted as the amenity-adjusted wage. In addition, in unreported results we have shown, following Seo and Oh (2024), that our key results (Theorems 1 and 2, below) apply when moving costs exhibit arbitrary unobserved heterogeneity.

responses to wage changes across markets.

2.2 Market-Level Labor Demand

Demand for labor in a market is induced by the demand for the market's product variety. Given product demand in (3), the total revenue of producers in market r, i is

$$R_{ri} = \eta_i \left(\frac{w_{ri}/\varphi_{ri}}{P_i} \right)^{1-\sigma}. \tag{5}$$

Perfect competition and production with labor alone imply that market revenue is paid to workers in that market, so $w_{ri}L_{ri} = R_{ri}$, where L_{ri} is market-level employment. Plugging this into (5) and rearranging yields market-level labor demand:

$$L_{ri} = D_{ri} w_{ri}^{-\sigma}, \quad \text{where } D_{ri} \equiv \eta_i \left(\varphi_{ri} P_i\right)^{\sigma-1}. \tag{6}$$

Here, D_{ri} is a market demand shifter and σ is the labor demand elasticity, which is larger when product varieties across regions are more substitutable.

We consider a set of economic shocks that affect this demand shifter across markets in the country of interest. We denote the realized equilibrium objects that result from these shocks with a superscript 1, e.g. employment L_{ri}^1 and wages w_{ri}^1. Let hats represent log-changes relative to a no-shock counterfactual, i.e. $\hat{x} \equiv \log\left(x^1/x\right)$. Taking log-changes of (6) yields

$$\hat{L}_{ri} = \hat{D}_{ri} - \sigma \hat{w}_{ri}. \tag{7}$$

This expression describes how the quantity of labor demanded in market r, i responds to wage changes and demand shocks in that market. Letting bold symbols represent vectors with market-level elements in $\mathcal{R} \times \mathcal{I}$, such that $\mathbf{x} = (x_{ri})$ for any x (unless specified otherwise), this relationship can be expressed as

$$\widehat{\mathbf{L}} = \widehat{\mathbf{D}} - \sigma \widehat{\mathbf{w}}. \tag{8}$$

2.3 Market-Level Labor Supply

Given (4), the probability that an individual originating in market o, p chooses market r, i is given by

$$\pi_{op,ri} = \frac{(w_{ri}/\tau_{op,ri})^{\theta}}{\sum_{dq}(w_{dq}/\tau_{op,dq})^{\theta}}. \tag{9}$$

We refer to these choice probabilities as "out-migration shares," keeping in mind that they include transitions between markets representing different industries in the same region, so the transitions need not involve geographic migration. Equation (9) shows that θ is the elasticity reflecting how relative out-migration shares across destination markets respond to relative wage changes; a larger θ implies less variation in the idiosyncratic taste shocks and hence increased responsiveness to wage changes. Labor supply to market r, i is then determined by these out-migration shares and the initial distribution of workers across markets, L^0_{op}:

$$L_{ri} = \sum_{o,p} f_{op,ri} = \sum_{o,p} \pi_{op,ri} L^0_{op}, \tag{10}$$

where the flow $f_{op,ri} \equiv \pi_{op,ri} L^0_{op}$ is the number of people initially in market o, p who choose market r, i. While equations (9)–(10) hold at the no-shock equilibrium, analogous expressions relate employment L^1_{ri}, out-migration shares $\pi^1_{op,ri}$, and other variables in the presence of labor demand shocks.

Now consider a labor demand shock leading to changes in wages across (potentially all) markets, as in (7), while holding moving costs $\tau_{op,ri}$ fixed.[8] Moreover, assume these wage changes are small, such that the first-order approximation around the no-shock equilibrium is precise.[9] Totally log-differentiating (9) and (10) around the no-shock equilibrium yields

$$\hat{L}_{ri} = \sum_{o,p} \gamma_{op,ri} \hat{\pi}_{op,ri} = \theta \left(\hat{w}_{ri} - \sum_{o,p} \gamma_{op,ri} \sum_{d,q} \pi_{op,dq} \hat{w}_{dq} \right). \tag{11}$$

Here $\gamma_{op,ri} \equiv f_{op,ri}/L_{ri}$ is the share of those in market r, i who arrive from market o, p in the no-shock equilibrium, which we refer to as the "in-migration share." This expression describes how the

[8] Appendix A.4 generalizes our results to incorporate both labor demand shocks and labor supply shocks resulting from changes in $\tau_{op,ri}$. This does not affect the ensuing results, the only change being that the error term now includes unobserved shocks to labor demand and labor supply.

[9] We maintain this small-shock assumption in the main text for expositional clarity, but Appendix A.3 uses exact hat algebra following Dekle et al. (2008) to study exact responses to potentially large shocks. Our quantitative analysis confirms that the small-shock approximation is accurate even for shocks with scale similar to those facing Brazilian labor markets during the country's early 1990s trade liberalization.

quantity of labor supplied in market r, i responds to wage changes across all markets and can be represented in matrix form as

$$\widehat{\mathbf{L}} = \theta \left(\mathbb{I} - \Gamma' \Pi \right) \widehat{\mathbf{w}}, \tag{12}$$

where $\Gamma = (\gamma_{op,ri})$, $\Pi = (\pi_{op,ri})$, and $\mathbb{I}$ is the identity matrix.

2.4 Equilibrium Response to Shocks

Equating the labor demand and labor supply expressions in changes in (8) and (12) yields the equilibrium effects of labor demand shocks on employment and wages in region-industry markets:

$$\widehat{\mathbf{L}} = \Omega \widehat{\mathbf{D}}, \qquad \text{where } \Omega \equiv \mathbb{I} - \left(\mathbb{I} + \frac{\theta}{\sigma} \left(\mathbb{I} - \Gamma' \Pi \right) \right)^{-1}, \tag{13}$$

$$\widehat{\mathbf{w}} = \frac{1}{\sigma} (\mathbb{I} - \Omega) \widehat{\mathbf{D}}, \tag{14}$$

Equation (13) shows the reduced-form relationship between employment changes and demand shocks. Two practical implications are immediately evident. First, the employment change in a given market depends upon the shocks to *all* other markets. In fact, if all markets face the same shock so $\hat{D}_{ri} = \hat{d}\ \forall r, i$, then the employment in each market remains unchanged.[10] Second, the importance of each shock in changing a market's employment depends upon connections between the focal market and other markets, as reflected in Γ and Π, which enter Ω. This implies that Γ and Π are sufficient statistics for all relevant effects of the switching costs ($\tau_{op,ri}$) and region-industry productivity differences (φ_{ri}), so we will not need to estimate either so long as we can measure Γ and Π (as discussed in the next subsection).

Having assumed full employment, the market-level changes in employment in equation (13) can be aggregated across local industries to derive a prediction for the change in local population—the object of empirical interest in this paper. Let $N_r \equiv \sum_i L_{ri}$ be the population in region r, which in turn implies $\hat{N}_r = \sum_i \frac{L_{ri}}{N_r} \hat{L}_{ri}$. Then define an $R \times (RI)$ matrix of initial employment shares S such that $(S\mathbf{a})_r = \sum_i \frac{L_{ri}}{N_r} a_{ri}$. Therefore the $R \times 1$ vector of changes in local populations $\mathbf{N}$ is $\widehat{\mathbf{N}} = S\widehat{\mathbf{L}} = S\Omega\widehat{\mathbf{D}}$. Note that, consistent with our definition of changes as measured relative to the end-of-period quantities in the no-shock equilibrium, the shares in S are also measured in that way and therefore do not include a 0 superscript.

[10] Specifically, we use that $\Gamma'\iota = \Pi\iota = \iota$, where ι is a column vector of ones. Thus, $\left[\left(1 + \frac{\theta}{\sigma}\right)\mathbb{I} - \frac{\theta}{\sigma}\Gamma'\Pi\right]\iota = \iota$, $\left[\left(1 + \frac{\theta}{\sigma}\right)\mathbb{I} - \frac{\theta}{\sigma}\Gamma'\Pi\right]^{-1}\iota = \iota$, $\Omega\iota = 0$, and finally $\Omega\iota\hat{d} = 0$ for any scalar $\hat{d}$.

2.5 Linking to Observables

The expression relating employment changes to labor demand shocks in (13) depends upon the ratio of elasticities θ/σ, the demand shocks $\widehat{\mathbf{D}}$, and the migration share matrices Γ and Π. While we keep θ/σ as an unknown parameter, in this subsection we consider how to link the labor demand shocks and migration shares in the no-shock counterfactual to observable quantities. Specifically, we address the possibility that the demand shocks may include observed and unobserved components and explain how we estimate the counterfactual migration shares using pre-shock observations.

We allow for unobserved labor demand shocks by assuming that we do not directly observe the overall labor demand shocks but instead observe a demand shifter $\widehat{\mathbf{z}} = (\hat{z}_{ri})$ such that

$$\hat{D}_{ri} = \hat{z}_{ri} + \zeta_{1,ri}, \tag{15}$$

where $\boldsymbol{\zeta}_1 = (\zeta_{1,ri})$ captures unobserved shocks.[11] We assume that the observed demand shocks are as good as random with respect to unobserved shocks, such that

$$\mathbb{E}\left[\hat{z}_{ri} \mid \boldsymbol{\zeta}_1\right] = \mu \quad \forall r, i, \tag{16}$$

where μ is a constant, while allowing the shocks to be mutually correlated (cf. Borusyak et al., 2022, Assumption 1). In words, the shock to each market has the same mean conditional on all unobserved demand shocks. This assumption reflects the best-case empirical scenario, in which the observed shocks are not correlated with any unobserved confounders. The issues we highlight throughout the paper therefore arise even in this best-case scenario.

Given (15), we can rewrite (13) as

$$\widehat{\mathbf{L}} = \Omega\widehat{\mathbf{z}} + \boldsymbol{\zeta}, \tag{17}$$

where $\boldsymbol{\zeta} \equiv \Omega\boldsymbol{\zeta}_1 = (\zeta_{ri})$ is a random error term.

To operationalize the preceding expression, the researcher must measure Ω in (17), which in turn depends on Γ and Π, in the counterfactual setting without labor demand shocks. The ideal approach to estimating these counterfactual migration probabilities will depend on the context. In our case, we pursue the simple approach of using out-migration shares from the pre-shock period

[11] Equation (15) implies that we observe the demand shifter $\widehat{\mathbf{z}}$ in units corresponding to labor demand changes. In many empirical contexts, these units are unknown and one plausibly observes only $\psi\widehat{\mathbf{z}}$, for unknown $\psi > 0$. Appendix A.6 shows that this generally limits our ability to separately estimate θ/σ from ψ, but without affecting estimates of the effects of observed shocks on local populations.

to infer the counterfactual. Specifically, define time $t = -1$, 0, and 1, such that the labor demand shocks arrive between $t = 0$ and $t = 1$. Recall $\Pi \equiv (\pi_{op,ri})$ is the matrix of counterfactual out-migration shares between $t = 0$ and 1 in the absence of labor demand shocks. We assume that these *counterfactual* out-migration shares equal those *observed* in the prior period, i.e. $\Pi = \Pi^0$ where $\Pi^0 \equiv (\pi^0_{op,ri})$ is the matrix of observed out-migration shares between $t = -1$ and 0. Importantly, this assumption does not require the economy to be in steady state. Because Π^0 may reflect persistent population and employment trends, the no-shock counterfactual can accommodate continued changes in regional population and market-level employment.[12] Given these estimates and the initial population L^0_{op} in each market, we calculate the flow between markets o, p and r, i as $f_{op,ri} = L^0_{op}\pi^0_{op,ri}$, the resulting no-shock employment in each market $L_{ri} = \sum_{o,p} f_{op,ri}$, and the in-migration shares $\gamma_{op,ri} = f_{op,ri}/L_{ri}$. So, by assuming that the pre-shock out-migration shares would persist in the no-shock counterfactual, we can measure all relevant counterfactual quantities using observables.[13]

2.6 The Effect of Observable Shocks on Regional Populations

We now develop an approximation to the population changes in response to observable shocks that yields intuition for the ways in which migration connections between markets govern this relationship. The approximation ignores indirect connections between markets, so we refer to it as the "direct mobility" approximation. It is most accurate when migration matrices are diagonally dominant, so that the mass of individuals remaining in a market is large relative to the mass moving between any particular pair of markets. Specifically, write $\Gamma = \mathbb{I} + \Delta\Gamma$ and $\Pi = \mathbb{I} + \Delta\Pi$ and assume that $\Delta\Gamma'\Delta\Pi \approx \Delta\Gamma'\Delta\Gamma' \approx \Delta\Pi\Delta\Pi \approx \mathbf{0}$. Plugging this into (13) and then (17) yields

$$\hat{L}_{ri} \approx \frac{\theta}{\sigma}\left(\sum_{o,p}\gamma_{op,ri}(\hat{z}_{ri} - \hat{z}_{op}) + \sum_{d,q}\pi_{ri,dq}(\hat{z}_{ri} - \hat{z}_{dq})\right) + \zeta_{ri}. \tag{18}$$

[12] This approach to measuring Π could perform poorly if migration responses to prior shocks evolve over time, such that $\Pi \neq \Pi^0$. Reasonable strategies for measuring Π in that context could leverage additional historical migration data, for periods prior to $t = -1$.

[13] Having linked the right-hand side of (17) to observables, we briefly discuss the left-hand side, which represents the difference in market population at $t = 1$ between counterfactuals with and without labor demand shocks. In place of this unobserved quantity, researchers may employ the observed change in population between $t = 0$ and 1. The two are equivalent up to an additional error term representing employment growth in the no-shock counterfactual. Because this additional term is orthogonal to the observed labor demand shocks under the exogeneity assumption (16), our estimation procedures discussed below remain unaffected.

Equation (18) makes clear that *relative* shocks across markets drive employment responses. The first sum measures the inflow of workers induced by the shocks to move to ri from each origin op. This inflow is proportional to the no-shock migration rate, $\gamma_{op,ri}$, and the difference in shocks between the two locations. Similarly, the second sum measures the number of workers induced by the shocks to stay in ri rather than leaving for dq. In both cases, only marginal workers switch in response to small shocks, but the Gumbel distribution of idiosyncratic preferences links the mass of marginal workers to the observed total number of movers in the no-shock equilibrium.[14]

While (18) shows how shocks to all markets affect employment growth in market r, i, the object of our empirical interest is the change in population at the region level, aggregating across industries in each region, $\hat{N}_r = \sum_i \frac{L_{ri}}{N_r} \hat{L}_{ri}$. Applying the latter expression to (18), using $L_{ri} \approx L^0_{ri}$ under the direct mobility approximation,[15] and rearranging yields the following characterization:

Theorem 1. *Under the direct mobility approximation, the population change following a set of observable shocks* $\{\hat{z}_{ri}\}$ *satisfies*

$$\hat{N}_r \approx \frac{2\theta}{\sigma} \frac{M_r}{N_r} (\hat{x}^M_r - \hat{x}^M_{-r}) + \zeta_r, \tag{19}$$

$$\textit{where} \quad \hat{x}^M_r \equiv \sum_i \frac{F_{-r,ri}}{M_r} \hat{z}_{ri}, \quad \hat{x}^M_{-r} \equiv \sum_p \sum_{o \neq r} \frac{F_{op,r}}{M_r} \hat{z}_{op}, \tag{20}$$

and $\zeta_r = \sum_i \frac{L_{ri}}{N_r} \zeta_{ri}$. *Here* $\hat{x}^M_r$ *and* $\hat{x}^M_{-r}$ *capture the relevant weighted averages of shocks in region* r *and outside-option shocks in other regions, with weights governed by* $F_{-r,ri} = \sum_p \sum_{o \neq r} F_{op,ri}$, $F_{op,r} = \sum_i F_{op,ri}$, *and* $M_r \equiv \sum_{o \neq r} F_{or}$, *for* $F_{op,ri} \equiv \frac{1}{2}(f_{op,ri} + f_{ri,op})$ *and* $F_{or} \equiv \sum_p \sum_i F_{op,ri}$.

The "M" notation refers to the fact that the weights in constructing $\hat{x}^{\mathrm{M}}_r$ and $\hat{x}^{\mathrm{M}}_{-r}$ are based on the observed flows of movers between markets in r and in other regions. To understand the structure of $\hat{x}^{\mathrm{M}}_r$, consider why and how an increase in $\hat{z}_{ri}$ alone affects local population. Some workers whose best option was to stay in another region but who were nearly indifferent about moving to r, i would now move. Similarly, some workers who would otherwise move from r, i to another region would now stay in r, i. The density of these marginal movers is determined by how connected market r, i is

[14] Appendix A.7 formalizes this idea by considering a model extension with a general distribution of idiosyncratic preferences. We show that the density of marginal workers always governs employment responses to wage changes (extending (11)) and, when θ is small, also employment responses to labor demand shocks. The Gumbel assumption links these marginal densities to observed transition probabilities, which under our direct mobility approximation yield the weights in (18).

[15] While the direct mobility approximation implies that market-level employment does not change much in one period, $L_{ri} \approx L^0_{ri}$, it does *not* mean that the economy is close to a steady state, in the sense of $\lim_{t \to \infty} (\Pi')^t \mathbf{L} \approx \mathbf{L}$.

to other regions, as measured by the number of workers who move between r, i and markets in other regions (on average between the two directions), $F_{-r,ri}$, relative to the total number of migrants in and out of r, M_r.

Similarly, $\hat{x}^{\mathrm{M}}_{-r}$ averages the shocks in regions outside r — which we term "outside-option shocks" — that affect r's population. Those are the shocks in *destination* markets for workers in r who are close to indifferent between staying in r and leaving for a market in another location, and are the *origin* shocks for workers outside r who are on the margin of moving to r. The weights entering $\hat{x}^{\mathrm{M}}_{-r}$ in (20) are the shares of migrants between r and each connected market elsewhere, relative to the total number of movers in and out of r, $F_{op,r}/M_r$.

Putting these together, we see that in (19) population growth in region r depends on the ratio of the migration elasticity and labor demand elasticity, θ/σ, the local migration share, M_r/N_r, and the difference in appropriately measured average shocks facing local markets vs. those in migrant-connected markets, $(\hat{x}^{\mathrm{M}}_{r} - \hat{x}^{\mathrm{M}}_{-r})$. Only when local shocks differ from outside-option shocks in migrant-connected regions will the vector of shocks lead to changes in local population.

To derive further intuition, it is instructive to consider two special cases: shocks that vary only across regions and shocks that vary only across industries, with the latter used in nearly all empirical applications.

Purely Regional Shocks. Suppose economic shocks apply to all workers in the region, i.e., $\hat{z}_{ri} = \hat{z}_r \ \forall r, i$. In this case, the effect of shocks on local populations in (19) simplifies to

$$\hat{N}_r \approx \frac{2\theta}{\sigma}\frac{M_r}{N_r}\left(\hat{z}_r - \hat{z}_{-r}\right) + \zeta_r, \qquad \text{where } \hat{z}_{-r} \equiv \sum_{o \neq r} \frac{F_{or}}{M_r}\hat{z}_o. \tag{21}$$

This expression makes clear that local population growth depends not only upon the direct shock to that region, $\hat{z}_r$, but also upon shocks to other regions. Specifically, the outside-option shock $\hat{z}_{-r}$ is a weighted average of shocks to other regions, with more weight placed on regions with stronger migrant connections with r. Because in practice each region is connected to a very different mix of others, this outside-option shock will differ across regions as well. As we will see in Section 3, the relationship between direct shocks and relevant outside-option shocks is central to understanding and interpreting the conventional migration regression.

Purely Industry Shocks. Now consider labor demand shocks that affect workers in the same industry equally, regardless of their region, i.e. $\hat{z}_{ri} = \hat{z}_i \ \forall r, i$. As shown in (6), national industry shocks could arise from import competition (which affects $\hat{P}_i$), from changes in industry productivity that apply nationally ($\hat{\varphi}_{ri} = \hat{\varphi}_i$), or from global taste shocks ($\hat{\eta}_i$); the responses to all those shocks are equivalent under our model.

Now $\hat{x}_r^{\mathrm{M}}$ and $\hat{x}_{-r}^{\mathrm{M}}$ are both averages of the same shocks $\hat{z}_i$ with weights capturing two offsetting forces. Suppose labor demand in industry i goes up. On the one hand, this attracts workers to market r, i who would otherwise have been in another location. On the other hand, this also attracts some workers from r to work in industry i in other regions. The respective intensity of those forces depends on how connected market r, i is to other locations and how connected markets o, i for $o \neq r$ are to region r.

Two observations are in order. First, the standard shift-share measure of a region's exposure to national industry shocks, $\hat{x}_r^{\mathrm{C}} = \sum_i \frac{L_{ri}}{N_r} \hat{z}_i$, does not coincide with $\hat{x}_r^{\mathrm{M}}$. For example, consider an industry i in region r that is more connected to other regions than the employment-weighted regional average, i.e. $F_{-r,ri}/L_{ri} > M_r/N_r$. Within our model, this can happen because workers in that market face lower migration costs. This industry enters with weight $\frac{F_{-r,ri}}{M_r} > \frac{L_{ri}}{N_r}$: the shares $\frac{F_{-r,ri}}{M_r}$ reflect the composition of *movers* in and out of r rather than industry composition of *all* workers residing in r.

Second, because switching industries involves frictions, it can be more costly to do so than staying in the same industry, even when migrating to another region. In that case, market r, i will have particularly strong connections to markets o, i in other regions, $o \neq r$. Since both markets r, i and o, i are subject to the same shock $\hat{z}_i$, their contributions to $\hat{x}_r^{\mathrm{M}}$ and $\hat{x}_{-r}^{\mathrm{M}}$ in (20), respectively, offset each other. That is, if workers have trouble switching industry when they migrate, they will see little reason to migrate in response to national shocks to their industry.

To further illustrate this idea, consider an extreme scenario in which workers cannot switch industries, formalized by $\tau_{op,dq} = \infty$ for all o, d, and $p \neq q$. In that case pure industry shocks will not induce geographic migration at all. Because workers cannot change industries, they can only move from market r, i to markets $-r, i$. This means that $F_{-r,ri} = F_{(-r)i,r} = F_{(-r)i,ri}$, which implies $\hat{x}_r^{\mathrm{M}} = \hat{x}_{-r}^{\mathrm{M}}$ and $\hat{N}_r = 0$ in (19). This extreme example illustrates the more general point that industry-switching frictions reduce migration responses when labor demand shocks have an

industry-specific component.

Together, the results in this section highlight the central role of outside options in determining the effects of labor demand shocks on local populations when workers face costs of moving between markets. The relevant outside option for a given region depends upon the transition costs faced by workers, for which pre-shock transition flows serve as a sufficient statistic, and upon the nature of the shocks—for example, if they vary by region vs. industry. We take advantage of this characterization when proposing model-consistent empirical approaches in Section 4.

3 The Conventional Migration Regression

The model in the preceding section describes how local populations change in response to labor demand shocks across region-industry markets. Here, we investigate how to interpret the results of the conventional migration regression when the data generating process follows this model.

3.1 The Conventional Migration Regression: Defining β

We consider a conventional approach to estimating the impact of labor demand shocks on local population, via Ordinary Least Squares (OLS) regression

$$\hat{N}_r = \alpha + \beta \hat{x}_r^{\mathrm{C}} + \varepsilon_r, \tag{22}$$

with observations weighted by the local population N_r. The dependent variable is defined above as the proportional change in population in region r, i.e. $\hat{N}_r = \sum_i \frac{L_{ri}}{N_r} \hat{L}_{ri}$ (though the same issues arise when the dependent variable is the change in the total in-migration flow or the probability of out-migration). The independent variable is an employment-weighted average of shocks across industries within the region, $\hat{x}_r^{\mathrm{C}} \equiv \sum_i \frac{L_{ri}}{N_r} \hat{z}_{ri}$, where the superscript "C" stands for the conventional regression. In the case of purely regional shocks, this average is just the regional shock itself. In the case of purely industry shocks, this average becomes a shift-share variable (e.g., Borusyak et al., 2022). Given this definition, the corresponding least-squares estimator of β is

$$\hat{\beta} = \frac{\sum_r N_r \hat{N}_r \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right)}{\sum_r N_r \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right)^2}, \tag{23}$$

where $\bar{x} = \sum_r N_r \hat{x}_r^{\mathrm{C}} / N$.[16]

To assess whether (22) is a useful specification, we first need to understand what researchers want to learn from it, as evidenced by how they interpret the coefficient. There are several interpretations in the empirical literature, so for now we focus on the most common interpretation of β as the causal effect of a labor demand shock to a single region r on that region's population. Like any causal effect, this interpretation involves a counterfactual comparison: here, changing the economic conditions in r while the shocks to all other regions are zero or held fixed. We return to this and other interpretations in Section 3.3.

Given these desired interpretations, we examine what (22) is actually estimating. Rather than presuming that this equation is correctly specified, we assume that the data are generated according to the model of Section 2: equation (17) or its approximation (19). We also maintain the assumption that the $\{\hat{z}_{ri}\}$ are as-good-as-randomly assigned to markets with respect to all unobserved shocks to labor demand and supply, as in (16).[17] This assumption allows us to focus on problems that remain even in the best-case scenario in which unobserved shocks are uncorrelated with observed ones.

Since the estimate $\hat{\beta}$ has random variation due to observed and unobserved shocks, we focus on the *estimand* of regression (22), which we define as

$$\beta = \frac{\mathbb{E}\left[\sum_r N_r \hat{N}_r \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right)\right]}{\mathbb{E}\left[\sum_r N_r \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right)^2\right]}. \tag{24}$$

This expression replaces the numerator and denominator of $\hat{\beta}$ with their expectations taken over observed and unobserved shocks (while viewing the no-shock equilibrium variables as fixed). As we show in Appendix B.7, β in (24) approximates the probability limit of $\hat{\beta}$ in an asymptotic sequence of economies with growing numbers of regions, under appropriate regularity conditions. That is, in large samples where the impact of the random realizations of observed and unobserved shocks on $\hat{\beta}$ vanishes, we expect $\hat{\beta} \approx \beta$. Additionally, unlike $\mathbb{E}\left[\hat{\beta}\right]$, the formulation in (24) is analytically

[16] For tractability, the shares in constructing $\hat{x}_r^{\mathrm{C}}$ and the regression weights use end-of-period values in the no-shock counterfactual. In practice, researchers usually use beginning-of-period values, but these are likely similar, for example under the direct mobility approximation (see Appendix B.5). Results similar to those derived below, albeit more cumbersome, can be obtained for unweighted regressions used in some studies.

[17] Of course, bias will arise in the conventional migration regression when the shocks are correlated with unobserved determinants of population growth. See, for example, Greenland et al. (2019), who argue that the population regression in Autor et al. (2013) was confounded by pre-existing trends in population growth.

tractable, which allows for the convenient characterization presented in the following subsection.

3.2 Characterizing β

We now characterize β when the data were generated by the model described in Section 2. To make the analytical results more intuitive, we focus throughout on the direct mobility approximation to $\hat{N}_r$, shown in (19). To simplify the presentation, we also assume that the explanatory variable $\hat{x}_r^{\mathrm{C}}$ is homoskedastic; while this assumption is usually violated with shift-share variables resulting from industry-level shocks, the case of a heteroskedastic explanatory variable presented in Appendix A.5 is notationally more cumbersome without yielding additional intuition.

Let $M \equiv \sum_r M_r$ be the national number of geographic migrants, so M/N is the migrant share of the national population in the no-shock scenario. We also define the flows that would prevail if the initial and final distributions of workers across markets, L_{ri}^0 and L_{ri}, reflected an economy with no moving costs, i.e. if $\tau_{op,ri} = 1 \ \forall o, p, r, i$. We indicate the corresponding quantities with tildes. In this scenario, the probability of moving to any market equals the market's share of national population, regardless of the origin, such that the flow of movers would be $\widetilde{f}_{op,ri} = L_{op}^0 \cdot \frac{L_{ri}}{N}$, with $\widetilde{F}_{op,ri}$, $\widetilde{M}_r$, and $\widetilde{M}$ defined correspondingly.[18] We then have:

Theorem 2. *Suppose the data are generated by the direct mobility approximation to the model in (19) with $L_{ri} \approx L_{ri}^0$, that shocks are as good as random as in (16), and that the conventional regression's explanatory variable $\hat{x}_r^C$ is homoskedastic, such that* $\mathrm{Var}\left[\hat{x}_r^C\right] = v \ \forall r$. *Then*

$$\beta = \frac{2\theta}{\sigma} \cdot \frac{M/N}{\widetilde{M}/N} \cdot \frac{\rho^M - \rho}{1 - \widetilde{\rho}} \tag{25}$$

where $\rho^M \equiv \sum_r M_r \rho_r^M / M$ for $\rho_r^M \equiv \mathrm{Cov}\left[\hat{x}_r^C, \hat{x}_r^M\right]/v$; $\rho \equiv \sum_r M_r \rho_r / M$ for $\rho_r \equiv \mathrm{Cov}\left[\hat{x}_r^C, \hat{x}_{-r}^M\right]/v$; and terms with tildes refer to the costless-migration scenario, so $\widetilde{\rho} \equiv (\sum_r \widetilde{M}_r \widetilde{\rho}_r)/\widetilde{M}$ for $\widetilde{\rho}_r \equiv \mathrm{Cov}\left[\hat{x}_r^C, \widetilde{x}_{-r}^C\right]/v$ and $\widetilde{x}_{-r}^C \equiv \sum_{k \neq r} (\widetilde{F}_{rk}/\widetilde{M}_r)\hat{x}_k^C$.

This theorem shows that β can be decomposed as a product of three terms with distinct economic meanings. First, β depends on the structural elasticities: as discussed above, the true migration

[18] When $\tau_{op,ri} = 1 \ \forall o, p, r, i$, (9) implies that the share of individuals choosing market r, i does not depend on the source o, p: $\widetilde{\pi}_{op,ri} = w_{ri}^{\theta} / \sum_{dq} w_{dq}^{\theta} \equiv \widetilde{\pi}_{ri}$. In that case, $L_{ri} = \sum_{o,p} \widetilde{\pi}_{op,ri} L_{op}^0 = \widetilde{\pi}_{ri} N$, so $\widetilde{\pi}_{ri} = L_{ri}/N$, i.e. the probability of choosing market r, i equals its share of national population. This in turn implies that the flow of individuals from o, p choosing r, i is $\widetilde{f}_{op,ri} = \widetilde{\pi}_{op,ri} L_{op}^0 = \widetilde{\pi}_{ri} L_{op}^0 = L_{op}^0 L_{ri}/N$. We then have $\widetilde{F}_{op,ri} = \frac{1}{2}(\widetilde{f}_{op,ri} + \widetilde{f}_{ri,op})$, $\widetilde{F}_{or} = \sum_p \sum_i \widetilde{F}_{op,ri}$, $\widetilde{M}_r = \sum_{k \neq r} \widetilde{F}_{kr}$, and $\widetilde{M} = \sum_r \widetilde{M}_r$.

responses are larger when θ is large and/or σ is small, and β inherits those relationships. Second, β increases with the national migration share, M/N (while the term $\widetilde{M}/N$, capturing the same share without mobility costs, is generally close to one and can therefore be ignored in practice).[19] Third, β depends upon $\frac{\rho^{\mathrm{M}}-\rho}{1-\widetilde{\rho}}$, which we call the "attenuation factor."

The attenuation factor captures two reasons for attenuation due to the conventional regression's misspecification relative to (19). First, because the shocks to relevant outside-option markets are omitted from the conventional regression, the estimate is attenuated when shocks in strongly migrant-connected markets are positively correlated, as reflected in ρ (and more correlated than between all markets weighted by size, as reflected in $\widetilde{\rho}$). Second, the conventional migration regression suffers from additional misspecification by using $\hat{x}_r^{\mathrm{C}}$ to measure the region's direct exposure to shocks rather than the model-consistent $\hat{x}_r^{\mathrm{M}}$. Deviations of $\hat{x}_r^{\mathrm{M}}$ from $\hat{x}_r^{\mathrm{C}}$ generate measurement error and likely result in attenuation via smaller ρ^{M}. While this intuition is informal, and indeed it is not guaranteed that $\frac{\rho^{\mathrm{M}}-\rho}{1-\widetilde{\rho}} \leq 1$, we find that this condition holds in all our simulations in Section 6. Moreover, we now provide more precise intuition for both sources of attenuation in the special cases of purely regional shocks and purely industry shocks.

Purely Regional Shocks. When shocks vary only across regions, i.e. when $\hat{z}_{ri} = \hat{z}_r \ \forall r, i$, the decomposition of β in (25) simplifies to

$$\beta = \frac{2\theta}{\sigma} \cdot \frac{M/N}{\widetilde{M}/N} \cdot \frac{1-\rho}{1-\widetilde{\rho}}. \tag{26}$$

Here the attenuation factor simplifies in two ways. First, $\rho^{\mathrm{M}} = 1$, as the conventional regression correctly captures the regions' direct exposure to shocks. Attenuation only operates through the first channel, with $\rho = \sum_o \sum_{d \neq o} \frac{f_{od}}{M} \mathrm{Corr}\left[\hat{z}_o, \hat{z}_d\right]$ simplifying to the average correlation between shocks to pairs of distinct regions, weighted by migration flows between them. In the denominator, $\widetilde{\rho} = \sum_o \sum_{d \neq o} \frac{\widetilde{f}_{od}}{\widetilde{M}} \mathrm{Corr}\left[\hat{z}_o, \hat{z}_d\right]$ is a similar average shock correlation weighted by the migration flows in the no-migration-cost scenario, which depend only upon location size. Therefore, when shocks are particularly correlated between pairs of regions with migration flows that exceed those predicted by location size, we have $\rho > \widetilde{\rho}$, attenuating β.

[19] Specifically, letting $\widetilde{f}_{rr} \equiv \sum_p \sum_i \widetilde{f}_{rp,ri}$, $\frac{\widetilde{M}}{N} = 1 - \sum_r \frac{\widetilde{f}_{rr}}{N} = 1 - \sum_r \left(\frac{N_r}{N}\right)^2$ equals one minus the Herfindahl index of population across regions. The Herfindahl index is small when the number of regions is sufficiently large and population is not too concentrated in a small number of them, as in our data.

Intuitively, when shocks are more strongly correlated between migrant-connected regions, the shock to a given region and its relevant outside-option regions become more similar, and workers have a weaker incentive to migrate in response to those shocks, reducing the measured migration response captured in β. For example, if migration occurs primarily among geographically proximate regions and shocks are spatially correlated, workers have little incentive to move, and the conventional regression yields a relatively small value of β.

Purely Industry Shocks. Now consider shocks that vary only across industries, i.e. $\hat{z}_{ri} = \hat{z}_i \; \forall r, i$. As in the general case, differences between $\hat{x}^{\mathrm{M}}_r$, which places more weight on shocks to local industries with stronger migrant connections to outside locations, and $\hat{x}^{\mathrm{C}}_r$, which weights by local employment alone, lead to a lower ρ^{M} and more attenuation in β. The omission of the relevant outside-option for region r, $\hat{x}^{\mathrm{M}}_{-r}$, drives further attenuation the stronger is the correlation between this outside option and the relevant average of direct shocks, $\hat{x}^{\mathrm{M}}_r$. To make this latter comparison more intuitive, consider a scenario with minimal cross-industry mobility. In this case, workers moving between r and other regions will have similar industry compositions, and thus $\hat{x}^{\mathrm{M}}_r$ is similar to $\hat{x}^{\mathrm{M}}_{-r}$, implying that $\rho^{\mathrm{M}} \approx \rho$, and attenuation is severe.

This example reveals how a setting with industry frictions undermines the intuitive appeal of measuring local labor demand shocks using shift-share measures. With industry frictions, workers in different industries within a region face different shocks and have different outside options, leading to complex spillovers of the effects of shocks across industries and regions. Yet, all of these complexities are omitted by the conventional migration regression.

3.3 Interpreting β

We conclude this section by connecting the theoretical results back to the interpretations of the conventional regression estimates that we have encountered in the literature. We present stylized versions of such interpretations here, while Appendix A.8 provides examples from prior work.

Consider using β to measure the effect of a *counterfactual* vector of shocks, i.e. shocks distinct from those observed and used in estimating β. The most common example considers a shock to a single region, as in the claim *"The estimate of β implies that a shock of 0.1 to a given region would drive a $0.1 \cdot \beta$ increase in that region's population."* Theorem 2 makes it clear that this interpretation

is generally incorrect. From (21), the true response to a purely regional shock to a single region r, $\hat{z}_{ri} = \hat{z}_r \neq 0 \; \forall i$, is $\hat{N}_r \approx \frac{2\theta}{\sigma}\frac{M_r}{N_r}\hat{z}_r$. The conventional migration regression parameter β in (26) differs from this true effect in two ways. First, it depends upon the aggregate migration share, M/N, rather than the region-specific migration share, M_r/N_r. This is to be expected, since the single conventional regression parameter can only capture the average effect across regions. Nonetheless, omitting this source of heterogeneity reduces the conventional regression's explanatory power. Second, β includes the attenuation factor when the observed shocks used in estimation are correlated between markets with strong migrant connections or have an industry dimension, as in all empirical examples we have encountered. In this case, β will understate the average response to a counterfactual single-location shock. Analogous interpretations of β are similarly incorrect, e.g. the effect of changing a single region's shock from the 25th to the 75th percentile of the observed distribution or the effect of a shock to a single region while holding other shocks fixed at their observed values.

A more extreme interpretation treats β as reflecting population responsiveness to *any* counterfactual vector of shocks, as in the claim *"The small and statistically insignificant estimate of β implies that local populations are unresponsive to local labor demand shocks."* The analysis in the preceding paragraph makes clear that such a conclusion is equally unwarranted. Different constellations of shocks with different correlations among migration-connected markets can yield substantially different estimates of β without changing workers' true underlying responsiveness. It is therefore inappropriate to use small conventional migration regression estimates to justify omitting the migration mechanism from an analysis of the effects of labor demand shocks.

While β does not capture the effects of *counterfactual* shocks, perhaps it reflects the effect of the full vector of *observed* shocks, $\widehat{\mathbf{z}}$. An example of this interpretation appears in the claim *"The small and statistically insignificant estimate of β implies that local populations did not respond substantially to the observed shocks."* We now show that, in general, the conventional regression does not capture the impact of the full vector of observed shocks $\widehat{\mathbf{z}}$ (e.g., of the China shock overall). The *aggregate* effect of all shocks is mechanically zero, as internal migration cannot affect the national population. To give this class of interpretations a chance, we therefore consider the effect of all shocks on population in a *given* region r. According to Theorem 1, this effect is $\frac{2\theta}{\sigma}\frac{M_r}{N_r}(\hat{x}_r^{\mathrm{M}} - \hat{x}_{-r}^{\mathrm{M}})$, which does not coincide with $\beta \cdot \hat{x}_r^{\mathrm{C}}$ in all regions for any β, because $\hat{x}_r^{\mathrm{C}}$ is typically imperfectly correlated with $\hat{x}_r^{\mathrm{M}}$

and even less correlated with $\hat{x}^{\mathrm{M}}_{-r}$.[20] For example, consider a region in which $\hat{x}^{\mathrm{C}}_r$ is below the mean and similar to $\hat{x}^{\mathrm{M}}_r$. In this case, the conventional regression predicts population decline regardless of the magnitude of attenuation in β. However, assume that in this region $\hat{x}^{\mathrm{M}}_{-r}$ is even lower than $\hat{x}^{\mathrm{C}}_r$. In this case, local population will grow.

4 Alternative Empirical Approaches

Given the conventional migration regression's inability to accurately capture the effects of observed or counterfactual shocks on regional populations, in this section we describe alternative empirical approaches that integrate information on pre-shock migration patterns to estimate migration responses in a model-consistent way. Before describing these approaches, it is important to emphasize that we maintain the assumption that the researcher has available a set of local labor demand shocks that are as good as randomly assigned (as in (16)). Any endogeneity in the shocks will lead to inconsistency in both the conventional migration regression and our proposed approaches.

4.1 Model-Consistent Approaches

We begin with approaches assuming the researcher observes the full matrix of market-to-market transitions needed to calculate Π and Γ. All of these approaches are model-consistent, at least up to stated approximations when relevant.

Non-linear Least Squares: The model-implied relationship between local population changes and the vector of local labor demand shocks is given by equation (17) above. Aggregating this expression to the region level yields

$$\hat{N}_r = S_r' \Omega(\theta/\sigma)\widehat{\mathbf{z}} + \zeta_r, \qquad \text{where } \Omega(\lambda) \equiv \mathbb{I} - \left(\mathbb{I} + \lambda\left(\mathbb{I} - \Gamma'\Pi\right)\right)^{-1}, \tag{27}$$

where S_r' is a row of the initial employment share matrix S. Following Section 2.5, we observe the local population change $\hat{N}_r$, the employment shares S_r', the migration matrices Γ and Π, and the shocks $\widehat{\mathbf{z}}$, so we can estimate the parameter θ/σ using non-linear least squares (NLLS) with

[20] An edge-case exception is when industry switching costs are infinite. Both true migration effects and β are then zero, so the conventional regression correctly captures the lack of any migration in response to the national industry shocks. This is not the case in real data, where industry switches are observed (see Table 1, below).

population weights N_r:[21]

$$\widehat{\theta/\sigma} = \arg\min_{\lambda \geq 0} \sum_r \frac{N_r}{N} \left(\hat{N}_r - S_r' \Omega(\lambda) \widehat{\mathbf{z}} \right)^2. \tag{28}$$

In Appendix B.8 we provide mild conditions for identification, i.e. for θ/σ to uniquely solve the NLLS problem. Because national population may grow or contract over time, in practice we recommend including an intercept in equation (27).

Given the resulting estimate of θ/σ we can calculate shock-induced population changes in each region using the same relationship in (27), omitting the error term. This process can be used to predict internally valid effects of the observed $\widehat{\mathbf{z}}$ shocks and externally valid effects of counterfactual shocks. Note that these predictions can be made without separately identifying θ and σ; it is sufficient to estimate only their ratio, θ/σ. This approach differs from standard structural approaches, which separately identify the supply and demand elasticities θ and σ but require additional data, specifically market-level wages.

Under the assumptions of the model, this NLLS approach will yield consistent estimates without additional assumptions. However, under reasonable approximations of (27), researchers can estimate the effects of the vector of labor demand shocks using linear regression, avoiding the need for nonlinear optimization, as we discuss next.

Taylor Series Approximation: Consider the following first-order Taylor series approximation in $(\theta/\sigma)(\mathbb{I} - \Gamma'\Pi)$, in which the relationship between the shocks $\widehat{\mathbf{z}}$ and regional population growth $\hat{N}_r$ is linear in θ/σ:

$$\hat{N}_r \approx \frac{\theta}{\sigma} S_r' \left(\mathbb{I} - \Gamma'\Pi \right) \widehat{\mathbf{z}} + \zeta_r. \tag{29}$$

The quality of this approximation improves as either θ/σ or the share of movers approaches zero. Given this approximation, the researcher can calculate the vector $\mathbf{t} \equiv S\left(\mathbb{I} - \Gamma'\Pi\right)\widehat{\mathbf{z}}$, in which each element t_r corresponds to a region r. The researcher can then estimate the following linear regression,

[21] In general, estimation equations like (27) require a "recentering" adjustment to correctly leverage the as-good-as-random assignment of the shocks, $\mathbb{E}\left[\hat{z}_{ri} \mid \boldsymbol{\zeta}\right] = \mu$ for $\mu \neq 0$ (see the discussion in a similar context in Borusyak and Hull (2021, Appendix D.5)). However, it is easy to show that this issue does not arise in our setting because NLLS estimates are invariant to μ: demeaning the shock has no implications for predicted population changes. Recentering would be necessary, however, if the shocks were *conditionally* as-good-as-randomly assigned: e.g., for tariffs that apply in manufacturing industries only, but not for services.

in which $\hat{\alpha}_1$ estimates θ/σ:

$$\hat{N}_r = \alpha_0 + \alpha_1 t_r + \text{error}_r, \tag{30}$$

where the error term includes ζ_r and the approximation error. In Section 6.3 we will find that this approach performs well in practice.

Direct Mobility Approximation: A closely related approximation relies only upon the share of movers approaching zero, in the sense introduced in (18). Because this approach does not consider the additional possibility that θ/σ is small, the required assumption is stronger than for the Taylor series approximation. The researcher can then calculate the $R \times 1$ vector $\mathbf{m} \equiv S\left(\mathbb{I} - \Gamma' + \mathbb{I} - \Pi\right)\widehat{\mathbf{z}}$, which is the term in parentheses in (18) aggregated to the region level.[22] One can then estimate the following linear regression, in which $\hat{\delta}_1$ again estimates θ/σ:

$$\hat{N}_r = \delta_0 + \delta_1 m_r + \text{error}_r. \tag{31}$$

Bilateral Regression: Another linear-regression based alternative is to change the unit of analysis to the pair of regions and study the effects of observed shocks on changes in gross flows of workers between regions $o \neq d$, with pre-period flows f_{od} as weights (e.g. Sprung-Keyser et al. 2022). In Appendix B.9, we show that the following specification is consistent with the model under the direct mobility approximation in (19),

$$\hat{f}_{od} = \beta_0 + \beta_1 \left(\hat{x}^{\text{D}}_{od} - \hat{x}^{\text{O}}_{od}\right) + \text{error}_{od}, \tag{32}$$

$$\text{where} \quad \hat{x}^{\text{D}}_{od} \equiv \sum_q \frac{f_{o,dq}}{f_{od}} \hat{z}_{dq}, \quad \hat{x}^{\text{O}}_{od} \equiv \sum_p \frac{f_{op,d}}{f_{od}} \hat{z}_{op},$$

and where $\beta_1 = \theta/\sigma$. Like the measures in (19), the model-consistent shock measures in (32) in the pure industry shock case are weighted averages of industry shocks in the focal region (o or d), but the weights reflect bilateral migration connections between the two regions rather than connections to all other regions. Specifically, $\hat{x}^{\text{D}}_{od}$ averages shocks to industries in the destination d, placing more weight on industries receiving a larger share of movers from o to d. Similarly, $\hat{x}^{\text{O}}_{od}$ averages shocks to industries in the origin o, with more weight on industries sending a larger share of migrants from o to d. Therefore, the model-consistent shock measure associated with a given region will vary across

[22] We recommend using (18) rather than (19) in practice because the former does not involve the additional approximation $N^0_r \approx N_r$; see Appendix B.3.

bilateral partner regions. Note that by allowing the regression coefficients on $\hat{x}^{\mathrm{D}}_{od}$ and $\hat{x}^{\mathrm{O}}_{od}$ to differ, one can test the model-implied restriction that they differ by a factor of -1. Once β_1 is estimated, the researcher can make predictions about responses of regional populations to both observed and counterfactual shocks by aggregating the predicted changes in flows, as detailed in Appendix B.9.

4.2 Approaches with Reduced Transition-Data Demands

The preceding approaches all rely on observing transitions between region-industry markets in the pre-shock period to infer what these transitions would have been in the absence of shocks. In many countries, this level of detail is only available in confidential administrative data, which can be difficult or impossible to access. We therefore consider three alternative approaches to reducing these data demands, two that infer these transitions from more widely available data and one that presumes relevant shock variation is purely regional. In Section 6, we assess the quality of the resulting estimates in practice.

Imputation 1: Independent Regional and Industry Transitions. Our first approach to imputing market-to-market transitions, referred to as "imputation 1" in Section 6, only requires transition information at the region-to-region and industry-to-industry levels, which are widely available in many countries. The imputation procedure assumes that transitions in the region and industry dimensions are independent and therefore multiplicatively separable, such that the flow across markets is imputed as

$$f^{(\mathrm{imp1})}_{op,ri} = L_{ri} \frac{f_{o,r}}{N_r} \frac{f_{p,i}}{L_i},$$

where $f_{o,r}$ is the flow from region o to region r (irrespective of industry), $f_{p,i}$ is the flow from industry p to industry i (irrespective of region), and L_i is national employment in industry i. Because $f_{o,r}$ and $f_{p,i}$ are widely available, this method is feasible in many contexts, but may lead to inaccurate results by incorrectly measuring the relevant outside options by assuming that industry transition rates are identical across regions and vice versa.

Imputation 2: Industry-to-Industry and Region-to-Region-Industry Transitions. Our second imputation approach, referred to as "imputation 2" in Section 6, takes advantage of the fact that many labor force surveys record the current region, current industry, and the prior region.

This information allows one to calculate $f_{o,ri}$ in which region-to-region (o to r) transitions may differ based on current industry (i). Following Caliendo et al. (2019), we apportion these observed transitions to initial industries p based on observed industry-to-industry flows $f_{p,i}$. The probability a worker in industry i came from industry p is $f_{p,i}/L_i$, so market-to-market flows are imputed as

$$f_{op,ri}^{(\text{imp2})} = \frac{f_{p,i}}{L_i} f_{o,ri}.$$

These two imputation approaches substantially broaden the contexts in which researchers can implement the model-consistent approaches just described.[23]

Treating Shocks as Region-Specific. Another approach is to ignore the industry dimension by considering a simplified version of our model aggregating all industries into a single one. This approach is equivalent to assuming that workers face no frictions in switching industries. With this simplification, markets are regions, and therefore widely available pre-shock region-to-region flows are sufficient to measure the Γ and Π matrices. In place of the original region-by-industry shocks $\hat{z}_{ri}$, the analysis would plug in their shift-share aggregates $\hat{x}_r^{\text{C}} \equiv \sum_i \frac{L_{ri}}{N_r} \hat{z}_{ri}$. Then any of the theory-consistent methods from Section 4.1 can be applied using the regional flow matrices.

This approach is fully justified when the shocks are genuinely region-specific and the researcher chooses to use the direct mobility approximation: as shown in (21), the effects of region-specific shocks under the approximation depend only upon transitions between regions and not upon industry transitions. (Methods that do not rely on the direct mobility approximation, such as NLLS, strictly speaking require the full transition matrix, even with purely regional shocks.) However, this simplification is problematic when the true shocks have an industry component, even under low mobility. Equation (20) shows that both $\hat{x}_r^{\text{M}}$ and $\hat{x}_{-r}^{\text{M}}$ are not correctly summarized by standard employment-weighted shift-share variables. Therefore, when prior work uses other regions' shift-share aggregates in an attempt to address the omission of outside options in the conventional migration regression, substantial misspecification likely remains.[24]

[23] These two simple imputation approaches generally imply values for L_{ri}^0 that differ from those observed, but Section 6.3 finds that they perform reasonably well nonetheless.

[24] Examples include Albert et al. (2021), Faber et al. (2021), and Greenland et al. (2019), which predate our work and motivate their regional outside-option controls based on intuitive concerns about spatial spillovers. Autor et al. (2023) and Foschi et al. (2025) implemented these region-specific adjustments in response to an earlier version of this paper, which did not sufficiently emphasize the issues arising from industry mobility frictions. However, this approach is inappropriate in the context of our model when shocks incorporate industry variation.

5 Data and Descriptive Statistics

To assess the quantitative importance of the conceptual problems with the conventional migration regression that we have discussed in Section 3, we use data on observed worker transitions across regions and industries for Brazil and Germany. This section describes the data source and descriptive statistics for Brazil, with the same information for Germany presented in Appendix D.1.

5.1 Data

The model-consistent approaches described in Section 4 require information on worker transitions across region-industry markets, allowing us to calculate the transition matrices Γ and Π. We do so using administrative panel data that cover all formally employed workers in Brazil, allowing us to observe these detailed transitions across regions and industries.

We utilize data from the Relação Anual de Informações Sociais (RAIS) covering 1999 to 2010. This administrative dataset is a census of the Brazilian formal labor market that allows us to follow all formally employed workers across jobs in different industries and regions. Regions in our analysis are 486 time-consistent "microregions," which combine economically integrated contiguous municipalities (counties) with similar geographic and productive characteristics, following Dix-Carneiro and Kovak (2017). Industries are based on the "Subsetor IBGE" classification, which identifies 25 distinct industries including 12 in manufacturing. Our sample includes individuals age 18 to 64 with positive earnings in December of the relevant year. See Appendix C for details on the data source, sample, and variable construction.

The RAIS data allow us to observe formal employment in region-industry pairs ("markets") and transitions between them. Guided by our model, which does not include non-employment or informality, we use formal employment as a proxy for population.[25] For each pair of markets, we calculate migration flows in the no-shock counterfactual, $f_{op,dq}$, as observed 5-year transitions, averaged over intervals from 1999/2004 to 2005/2010 to reduce noise.

[25] A simple alternative we do not pursue would be to consider an extension of the model that treats informal jobs, non-employment, and unemployment as belonging to another fictitious sector that can be chosen by the workers and to use RAIS data to measure worker transitions between formal employment sectors and that fictitious one.

Table 1: Average Mobility Rates

Across	Mobility (%)
Regions	12.3
States	4.0
Industries	24.2
Regions and Industries	7.1

Notes: Percent of individuals with 5-year transitions between the cells defined in the first column, calculated from average 5-year flows between region-industry cells in RAIS data from 1999 to 2010. The last row measures the fraction of workers who change *both* region and industry in the same period.

5.2 Descriptive Evidence

We next provide summary statistics on the mobility rates and migration networks of the Brazilian economy, guided by the results in Sections 2 and 3, showing how they determine the population effects of labor demand shocks and the conventional regression's estimand β.

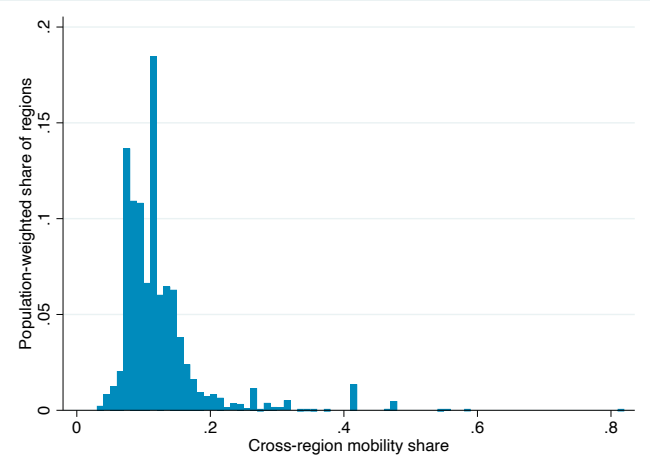
Table 1 reports worker mobility rates in different dimensions. The first row reports the national average regional migration rate (M/N), which plays a central role for β. Over a 5-year period, 12.3 percent of workers migrate across regions on average. The second row reports the probability of moving between the 27 Brazilian states (4.0 percent), showing that workers are more likely to move across regions within states than across states, which reflects a lower likelihood of long-distance moves. Despite the fact that we use a coarser classification of industries than regions in the RAIS data, industry transitions (third row) are substantially more likely than migration events (24.2 percent), and 7.1 percent of workers transition between both regions and industries in a given 5-year period.

Underlying these national average mobility rates is substantial heterogeneity in migration shares across regions. Figure 1 shows a histogram of M_r/N_r, the average number of migrants to and from r over a 5-year period expressed as a fraction of population N_r (with regions weighted by their population N_r). As discussed in Section 3.3, this heterogeneity in baseline migration intensity drives heterogeneity in the true effect of local labor demand shocks on local population, but is omitted from the estimate of β in the conventional migration regression.[26]

We next show that migration exhibits a network structure. The analyses in Section 2.6 emphasized the role of spillovers between connected regions and region-industry cells in driving migration

[26]We also find substantial heterogeneity across industries in the probability of switching industries, with rates ranging from 7% to 47%.

Figure 1: Geographic Mobility Rates Across Regions



Notes: The histogram reports, for each microregion r, the migration share M_r/N_r, calculated as one half the sum of the number of migrants to and from r as a fraction of the end-of-period population N_r, using 5-year flows between region-industry cells in RAIS data from 1999 to 2010. The histogram is weighted by each region's end-of-period population.

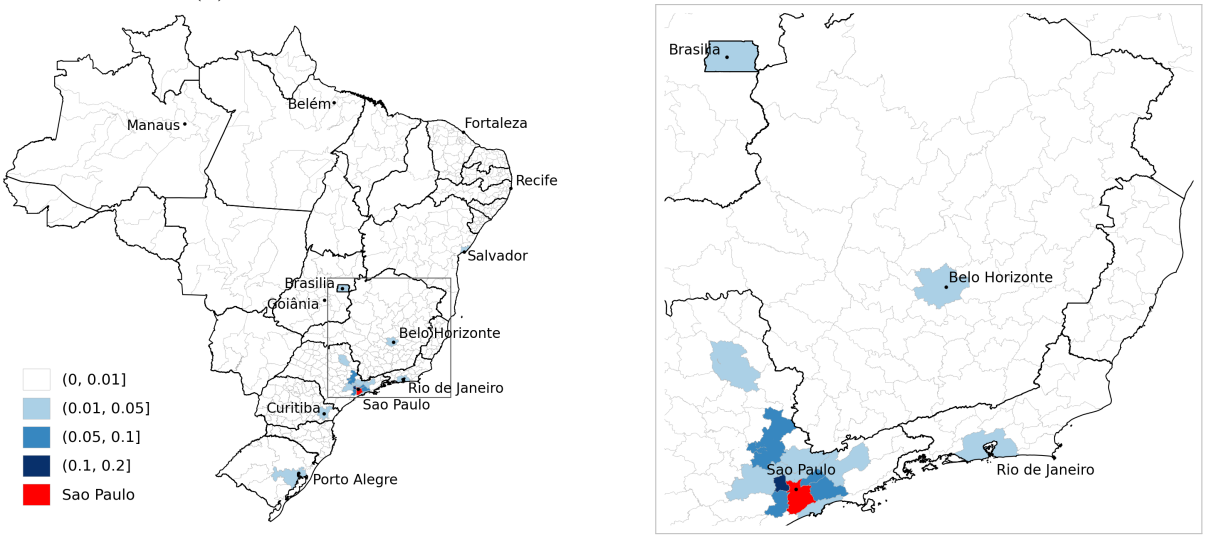
responses to labor demand shocks. These spillovers will be particularly important in practice if region or industry connections differ substantially for workers in different regions, e.g. due to migration costs increasing in distance.

Figure 2 illustrates the spatial concentration of migration connections between Brazilian regions. The map in panel (a) shows the destination distribution of migrants from São Paulo. The most important migrant destinations include neighboring regions in the same state, but also Rio de Janeiro and other cities in nearby states. Contrast this distribution with that of Rio de Janeiro in panel (b). While São Paulo is an important destination for migrants from Rio, their destination distribution differs from that of migrants originating in São Paulo, with many moving to neighboring locations within Rio de Janeiro State and to Manaus in the far northwest of Brazil. The differences between these two migrant distributions imply distinct outside options and therefore distinct spillovers for São Paulo and Rio when Brazilian markets face a vector of labor demand shocks.

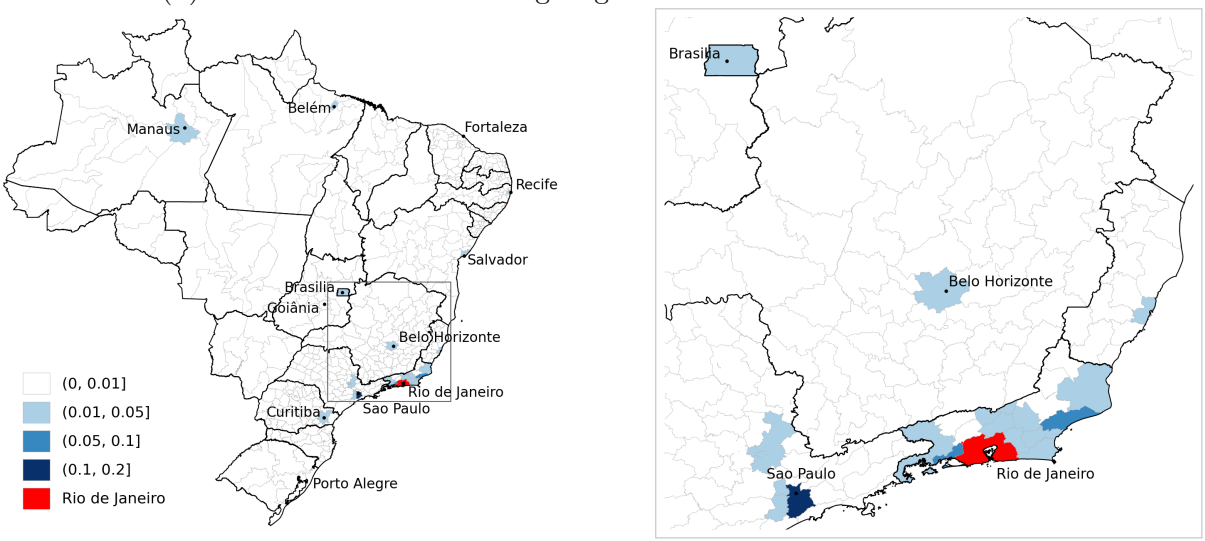
To confirm that these two relatively concentrated sets of migrant connections are not unique to São Paulo and Rio, we calculate the Herfindahl–Hirschman index $HHI_r = \sum_{k \neq r} (F_{kr}/M_r)^2$ measuring the average concentration of destinations and origins among those who migrated to or from a given region. The population-weighted average HHI_r across regions is 0.117 (s.d. 0.103), which is equivalent to equally-weighted connections to only 9 other regions. If migration were random, the latter value would equal the inverse HHI of regional populations, which is much higher,

Figure 2: Spatial Concentration of Migrant Connections

(a) Destination Shares Among Migrants from São Paulo



(b) Destination Shares Among Migrants from Rio de Janeiro



Notes: Both maps on the left plot the share of out-migrants from the relevant source region that choose each other region in Brazil, with darker blue regions capturing a larger share of out-migrants. These flows represent 5-year transitions averaged over 1999/2004 to 2005/2010, measured in RAIS data. Panel (a) plots the destination distribution for out-migrants from São Paulo, while panel (b) plots the same for Rio de Janeiro. The largest 12 cities by population are labeled. The maps on the right zoom in on the area around São Paulo and Rio de Janeiro.

at 29. Thus migrant networks in typical regions are highly concentrated, and workers in different regions may face substantial differences in spillovers from connected regions.

6 Quantitative Assessment of Migration Regressions

In this section we examine the bias and predictive inaccuracy of the conventional migration regression in the context of Brazilian geography, industrial structure, and baseline migration patterns. Based on the model in Section 2, we simulate equilibrium population responses to labor demand shocks that follow a variety of data generating processes (DGPs), including the real-world tariff shocks Brazilian industries faced during the country's 1990 trade liberalization. We then compare these true model-generated responses to those inferred from conventional migration regression estimates, the various alternative procedures developed in Section 4, and other specifications implemented in the prior literature. Parallel results for Germany appear in Appendix D.

6.1 Calibrating Model Parameters and Shock Distributions

To generate model-simulated population responses to labor demand shocks, we must first specify the distribution of these shocks and calibrate the relevant reduced-form parameter, θ/σ. We consider four data-generating processes for shocks. The first set of shocks consists of independent normal draws across regions, i.e. $\hat{z}_{ri} = \hat{z}_r$ and $\hat{z}_r \overset{iid}{\sim} \mathcal{N}(0, \sigma_R^2)$. We refer to these shocks as "iid across regions." The second set of shocks introduces spatial correlation by assigning the same shock to all locations (microregions) in the same Brazilian state. The state-level shocks are independently and normally distributed, such that $\hat{z}_{ri} = \hat{z}_{s(r)}$ and $\hat{z}_s \overset{iid}{\sim} \mathcal{N}(0, \sigma_S^2)$, where $s(r)$ is the state containing region r. We refer to these shocks as "iid across states." Both of these are pure regional shocks albeit with different correlation structures. Third, we consider pure industry shocks that are iid across industries: $\hat{z}_{ri} = \hat{z}_i$ and $\hat{z}_i \overset{iid}{\sim} \mathcal{N}(0, \sigma_I^2)$. For our initial results, we do not introduce unobserved shocks, which makes the shock variances σ_R^2, σ_S^2, and σ_I^2 immaterial; thus, we leave them undetermined until Section 6.3. The final set of shocks, which also vary across industries, are the actual tariff changes facing Brazilian industries during the country's 1990 trade liberalization.[27] These real-world shocks allow us to consider a realistic correlation structure across industries, as a complement to the iid process.

We draw on the literature to determine a realistic calibration of $\frac{\theta}{\sigma}$. For the migration elasticity, we use the median estimate from the 11 studies considered in Table B.1 of Head and Mayer (2021), $\theta = 1.63$. In the absence of estimates of the Armington substitution elasticity across sub-national

[27] Specifically, we let $\hat{z}_{ri} = \hat{z}_i = d\log(1+\tau_i)$, where τ_i is the tariff rate in tradable industry i, following Kovak (2013) and Dix-Carneiro and Kovak (2017). For non-tradable industries, we set $\hat{z}_{ri} = 0$.

Table 2: Bias in the Conventional Migration Regression

Shocks	Relative β	R^2
iid across regions	1.000	0.804
iid across states	0.423	0.833
iid across industries	0.511	0.551
Brazil industry tariff shock	0.470	0.636

Notes: For the first three rows, the table reports the median statistic across 601 simulations, while the last row reports each statistic for a single shock vector given by the Brazil tariff shock of the early 1990s. The first column rescales the resulting value of β relative to the median estimate under iid shocks across regions, while the second column reports the adjusted R^2. Simulations incorporate observed shocks only, without any unobserved shocks. See text for definitions of shock processes.

regions, we use standard estimates of the cross-country elasticity from Broda and Weinstein (2006) and Feenstra et al. (2018), both of which find $\sigma \approx 4$. We therefore set $\frac{\theta}{\sigma} = 0.4075$.

6.2 Bias in the Conventional Migration Regression

We first assess how misspecification affects results from the conventional migration regression, including the estimate of β and the goodness of fit. For each of the first three types of shocks described above, we generate 601 shock simulations. For the Brazilian trade-liberalization shocks, by contrast, we simulate the impact of the single observed vector of industry-specific tariff changes. Given the shocks and calibrated value of θ/σ, we generate changes in market-level employment as $\widehat{\mathbf{L}} = \Omega(\theta/\sigma)\widehat{\mathbf{z}}$ and changes in regional population as $S\widehat{\mathbf{L}}$.[28] Recall that Ω depends on the observed flow matrices Γ and Π, which we measure using the RAIS data (Section 5.1). For each set of shocks, we estimate the conventional regression using model-implied regional employment changes as the dependent variable and $\hat{x}_r^{\mathrm{C}} = \sum_i \frac{L_{ri}}{N_r}\hat{z}_{ri}$ as the independent variable, weighting the regression by population.

We collect the estimates of β and R^2 from each individual simulation and report their medians across simulations in Table 2. The "Relative β" column normalizes estimates of β by the estimate with iid shocks across regions. The first row reports that benchmark in which the relative β equals 1 by construction, corresponding to the fact that there is no attenuation (recall that in (26) $\rho = \widetilde{\rho} = 0$ and $\frac{1-\rho}{1-\widetilde{\rho}} = 1$). The misspecification of the conventional regression leads to a median R^2 across simulations of only 0.804, despite the absence of any unobserved shocks, such that model-consistent procedures would achieve $R^2 = 1$. This reduction in predictive performance reflects the fact that

[28] Note that we do not utilize the direct mobility approximation in the simulation, since we use (17) rather than (19). We do, however, maintain the assumption of small shocks; see Appendix A.3 that relaxes it.

β replaces the region-specific migration intensity $\frac{M_r}{N_r}$ with the national migration intensity $\frac{M}{N}$ in addition to omitting variation in $\hat{z}_{-r}$ (see equation (21)).

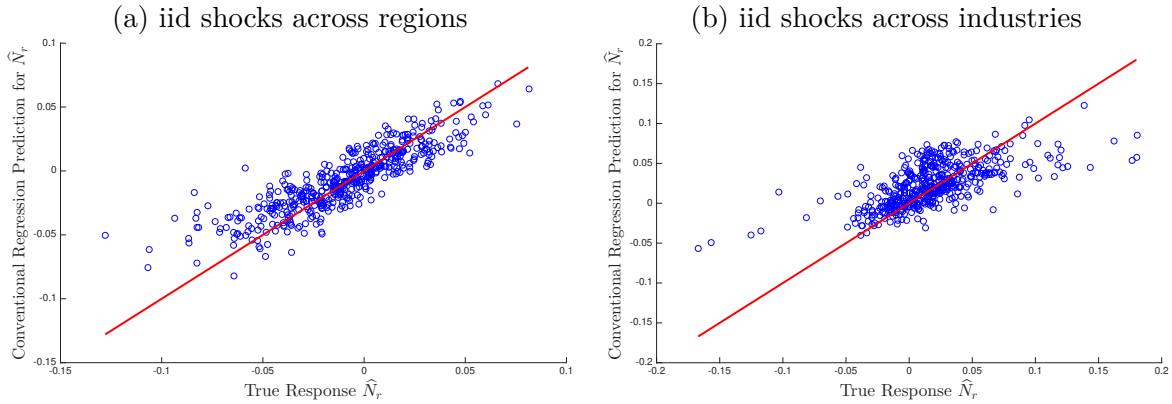
When purely regional shocks are particularly correlated between migrant-connected regions, ρ in (26) is greater than zero (and than $\widetilde{\rho}$) and the estimate of β is attenuated, biasing the estimates of population responses to counterfactual shocks. We can see this effect in the second row of Table 2, which presents results using iid shocks across states. The prevalence of intra-state migration (see Table 1) induces substantial correlation between migrant-connected regions when facing state-level shocks. As a result, the implied attenuation factor declines to 0.423 in the second row. Despite this substantial attenuation bias in β, the conventional regression yields a similar R^2 to the case with iid region shocks; while the variation explained by the conventional regression falls, with spatially correlated shocks the true variation in population is also lower.

The remaining two rows of Table 2 consider purely industry-specific shocks, which comprise the overwhelming majority of examples in the empirical literature (see footnote 3). The third row considers iid shocks across industries, and the fourth row uses the real-world liberalization shocks faced by Brazilian industries. In both cases the conventional regression estimates of β experience substantial attenuation: 0.511 and 0.470, respectively. This degree of attenuation is similar to that under iid state shocks. With iid industry shocks, however, attenuation arises from two sources: omitting shocks to the outside option and mismeasuring the relevant average shock to the location itself. As a result, the R^2 is substantially reduced with industry shocks, falling below 64 percent despite the fact that population growth depends solely upon the observed shocks. The results for iid industry shocks are similar to those for the actual tariff shocks, suggesting that drawing shocks iid is not implausibly simplistic.

Figure 3 visually presents the simulation results for a single representative draw. The y-axes measure predicted population changes based on the conventional migration regression estimates, $\hat{\beta}(\hat{x}_r^{\text{C}} - \bar{x})$, and the x-axes measure the true population change according to the model.[29] Each point represents a particular region. Panel (a) shows population changes when facing iid shocks across regions, while Panel (b) considers iid shocks across industries. If the regression perfectly captured the effect of the shocks on local populations, all of the points would lie along the 45-degree line, shown in red. The fact that many points are distant from the line implies that the regression-based

[29] In the predicted population change, $\bar{x}$ appears because of the regression intercept.

Figure 3: Conventional Regression Predictions vs. True Responses for Observed Shocks



Notes: Scatter plots compare the predicted population response to the observed labor demand shocks against the true model-based response for a single simulation. The simulation is chosen to be representative, i.e. to have $\hat{\beta}$ and R^2 most similar to their averages across many simulations. Each point represents a region. The y-axis is the predicted population response based on the conventional migration regression, $\hat{\beta}(\hat{x}_r^C - \bar{x})$. The x-axis is the true model-based population change. The left panel shows the result for iid shocks across regions, and the right panel shows the result for iid shocks across industries. 45-degree lines are shown in red.

predictions are quite inaccurate in many cases, consistent with the R^2 values in Table 2 being well below 1. This problem is substantially worse when shocks are industry-specific (in Panel (b)). In this example, many regions experience substantial population changes as a result of the shocks despite the relatively small predicted responses based on the conventional regression.

Together, the results in Table 2 and Figure 3 confirm the patterns predicted in our theoretical analysis and reveal substantial attenuation when shocks are correlated across migrant-connected regions. The low R^2 value for industry-level shocks also implies that the conventional regression yields relatively poor predictions of the effects of the observed shocks used to estimate β, even in a setting where migration is fully determined by the observed labor demand shocks.

6.3 Performance of Alternative Empirical Approaches

While the preceding analyses allow us to assess attenuation and misspecification in the conventional migration regression, it is unrealistic to assume that the researcher observes *all* shocks. We therefore introduce unobserved shocks, both for realism and to assess all of the approaches to measuring the effects of observed shocks on local populations proposed in Section 4.1. We assume the unobserved shocks to be orthogonal to the observed labor demand shocks as in (16) to maintain our focus on issues that arise even in the best-case scenario with exogenous shocks.

For each simulation, we generate the true employment changes at the market level as $\widehat{\mathbf{L}} =$

$\Omega(\theta/\sigma)\widehat{\mathbf{z}}+\boldsymbol{\zeta}$, where $\widehat{\mathbf{z}}$ is the vector of observed shocks and the error $\boldsymbol{\zeta}=(\zeta_{ri})$ captures the effects of unobserved shocks. This error consists of additive region, industry, and region-industry components, such that

$$\zeta_{ri}=\epsilon_r+\epsilon_i+\epsilon_{ri}-\bar{\epsilon}, \tag{33}$$

where $\bar{\epsilon}$ is a constant that ensures that the national employment change is zero. To ensure a realistic variance and correlation structure for this noise term, we take the observed 5-year changes in employment across markets and regress them on region fixed effects and industry fixed effects. We then measure the standard deviations of the region effects, the industry effects, and the market-level residual and average them across the 5-year windows in our RAIS data. Finally, we simulate ϵ_r, ϵ_i, and ϵ_{ri} as mean-zero normal random variables with the corresponding variances: $\epsilon_r \overset{iid}{\sim} \mathcal{N}(0, 0.050^2)$, $\epsilon_i \overset{iid}{\sim} \mathcal{N}(0, 0.082^2)$, and $\epsilon_{ri} \overset{iid}{\sim} \mathcal{N}(0, 0.154^2)$ (see Appendix C.2.1).

We focus on two of the DGPs for the observed $\widehat{\mathbf{z}}$ shocks from Table 2: iid across regions and iid across industries. Having specified the magnitude of variation in unobserved shocks, the magnitude of the observed shocks becomes meaningful. For comparability across shock types, we set the variance of the observed shocks to achieve a certain variation in the induced changes in population: "small" shocks target a standard deviation of 0.015 across regions (i.e., 1.5% of the population), "medium" shocks 0.025, and "large" shocks 0.035. In Panel A of Table 3, we report the standard deviations of underlying shocks that achieve these targets, along with the associated cross-regional variation in wage effects induced by the shocks (see Appendix C.3 for details).

The first two rows of Panel A show that to achieve the same variation in population effects, industry shocks must exhibit much more variation than region shocks. This is for two reasons. First, by averaging across industries, the shift-share variable $\hat{x}_r^{\mathrm{C}}$ has much less variation across regions than the underlying industry shocks. This force can be seen by considering the variation in the regional *wage* effects of the two types of shocks in the third row. Ten times larger industry shocks generate regional wage responses that are only 69% more variable. Second, per our theoretical analysis, shocks that induce correlations between the own local shocks and the outside options — such as pure industry shocks, even iid ones — induce weaker true population responses. Thus, iid industry shocks that generate 69% larger wage variation than iid regional shocks induce equal population responses.

Table 3: Estimates and Predictive Accuracy for Alternative Estimators

Shock size:	large		medium		small	
Shock type:	iid region	iid industry	iid region	iid industry	iid region	iid industry
Panel A: Shock magnitudes						
std. dev. of shocks	0.41	4.01	0.29	2.87	0.18	1.72
std. dev. of population effects ×100	3.50	3.50	2.50	2.50	1.50	1.50
std. dev. of wage effects ×100	9.10	15.40	6.50	11.00	3.90	6.60
Conventional wage regression	0.232	0.241	0.233	0.241	0.235	0.241
	(0.030)	(0.016)	(0.042)	(0.022)	(0.069)	(0.037)
Panel B: Conventional migration regression	0.085	0.042	0.086	0.043	0.089	0.043
	(0.032)	(0.021)	(0.044)	(0.027)	(0.073)	(0.044)
	{0.278}	{0.528}	{0.341}	{0.592}	{0.543}	{0.739}
Panel C: Model-consistent approaches						
NLLS	0.423	0.413	0.430	0.416	0.445	0.421
	(0.141)	(0.129)	(0.197)	(0.180)	(0.320)	(0.294)
	{0.043}	{0.034}	{0.084}	{0.068}	{0.235}	{0.200}
Taylor approx.	0.361	0.372	0.366	0.375	0.379	0.383
	(0.104)	(0.105)	(0.145)	(0.146)	(0.243)	(0.243)
	{0.042}	{0.050}	{0.080}	{0.086}	{0.219}	{0.210}
Direct mobility approx.	0.306	0.292	0.309	0.295	0.318	0.306
	(0.089)	(0.088)	(0.125)	(0.119)	(0.208)	(0.194)
	{0.056}	{0.180}	{0.094}	{0.216}	{0.235}	{0.324}
Bilateral regression	0.329	0.283	0.334	0.283	0.347	0.288
	(0.121)	(0.064)	(0.170)	(0.087)	(0.283)	(0.141)
	{0.126}	{0.234}	{0.215}	{0.263}	{0.519}	{0.335}
Panel D: Reduced transition-data demands						
Taylor, regional model	0.343	0.209	0.347	0.212	0.358	0.223
reg-to-reg	(0.098)	(0.077)	(0.137)	(0.100)	(0.229)	(0.159)
	{0.046}	{0.323}	{0.085}	{0.373}	{0.223}	{0.510}
Bilateral, regional shocks	0.329	0.207	0.334	0.211	0.347	0.217
	(0.121)	(0.095)	(0.170)	(0.129)	(0.283)	(0.212)
	{0.126}	{0.418}	{0.215}	{0.510}	{0.519}	{0.784}
Taylor imputation 1:	0.349	0.518	0.354	0.523	0.365	0.538
reg-to-reg and ind-to-ind	(0.100)	(0.150)	(0.140)	(0.205)	(0.234)	(0.337)
	{0.045}	{0.250}	{0.084}	{0.284}	{0.222}	{0.385}
Taylor imputation 2:	0.353	0.544	0.358	0.552	0.370	0.565
ind-to-ind and reg-to-reg-ind	(0.102)	(0.171)	(0.142)	(0.231)	(0.237)	(0.377)
	{0.044}	{0.222}	{0.083}	{0.260}	{0.223}	{0.387}

Notes: Panel A presents magnitude information for simulated shocks. The first row shows the standard deviation across regions or industries, corresponding to the shock variation, while the second and third rows report variation across regions. The fourth row reports estimates from a conventional wage regression relating changes in wages in regional markets to the employment-weighted average of shocks to local industries, $\hat{x}_r^{\mathrm{C}}$ (see Appendix C.3 for details on the third and fourth rows). Panel B shows estimates from the conventional migration regression. Panel C reports results from the model-consistent approaches to predicting the effect of observed shocks on regional populations discussed in Section 4.1. Panel D reports results from the approaches with reduced transition-data demands discussed in Section 4.2. Each column defines the shock process; see the text for details. In Panels B through D, numbers without parentheses or braces are median coefficient estimates across 601 simulations. Associated "standard errors" (standard deviations across simulations) are in parentheses. The relative mean squared errors, computed as the ratio of the median of (34) to the median of (35), are in curly braces.

Variation in wage effects reported in the table can help place our large, medium, and small shock magnitudes into context. The cross-regional standard deviation of the long-run wage effect in Dix-Carneiro and Kovak (2017) is approximately 6%, although this effect takes significantly longer

than five years to materialize. This effect is close to the magnitudes associated with the small industry shocks and the medium region shocks. For comparison, the standard deviation of wage effects induced by the China Shock is approximately 1% according to Autor et al. (2013) and even smaller according to Borusyak et al. (2022). Panel A also reports estimates from a simple regression of regional wage changes on the conventional regression's independent variable, $\hat{x}_r^{\mathrm{C}}$. Appendix C.3 describes the construction of the regional wage measure, which includes effects of both observed and unobserved shocks, following a procedure similar to the shock calibration for population changes. For all of the shocks considered in Table 3, these wage effects are highly statistically significant at conventional levels, indicating that even the small shocks are large enough to yield detectable effects on local labor market outcomes.

In Panels B and C of Table 3, we compare the performance of the conventional regression to that of the alternatives developed in Section 4.1. For each shock type and each estimation approach, we present the regression estimates and their standard errors (in parentheses) to assess biases and the ability of each estimator to statistically distinguish population responses from zero in the presence of unobserved shocks.[30] We calculate the estimates for each of $J = 601$ simulations and report the median. We also assess each approach's ability to capture population changes induced by the observed shocks by calculating the Mean Squared Error (MSE, in braces). For a given simulation j with simulated shock vector $\widehat{\mathbf{z}}_j$, let $\hat{N}_{r;j}^{\text{true}}$ be the true response to the observed shocks and let $\hat{N}_{r;j}^{\text{predicted}}$ be the predicted population growth from any estimation approach under consideration. We define

$$MSE_j = \sum_r \frac{N_r}{N} \left(\hat{N}_{r;j}^{\text{true}} - \hat{N}_{r;j}^{\text{predicted}} \right)^2, \tag{34}$$

which averages squared prediction errors across regions (with population weights). To help interpret the magnitude of this measure, we report its median value across simulations, normalized by the median MSE for an uninformative prediction, in which $\hat{N}_{r;j}^{\text{predicted}} = 0$:[31]

$$MSE_j^{\text{uninformative}} = \sum_r \frac{N_r}{N} \left(\hat{N}_{r;j}^{\text{true}} \right)^2, \tag{35}$$

[30] The "standard error" is computed as the true standard deviation across simulations. Any valid procedure for obtaining standard errors (such as exposure-robust standard errors of Borusyak et al. (2022)) provides an approximation for this object from the observed sample. Measures of statistical significance based on the median β estimate across simulations divided by this "standard error" should be interpreted appropriately.

[31] We use relative MSE rather than R^2 because the latter is not defined for NLLS, and its interpretation differs for bilateral regressions, estimated with a different unit of analysis.

which is equivalent to the variance of the true population changes driven by the shocks. Smaller values of the relative MSE imply more accurate measures of population changes induced by observed shocks, and a value of 1 implies no better prediction than assuming zero change in all regions.[32]

Panel B first examines the conventional migration regression, confirming its poor performance. As in Table 2, the coefficients with iid industry shocks are nearly half as large as those with iid region shocks, reflecting substantial attenuation bias with industry shocks.[33] Moreover, despite highly statistically significant wage effects for all simulated shocks, the conventional regression effect estimate is significant at conventional levels only for large shocks, whether region-specific or industry-specific. Conventional regression estimates are statistically insignificant even with medium industry shocks, which induce wage effects more than 50 percent larger than those in Dix-Carneiro and Kovak (2017). In other words, in our context, shocks would need to be implausibly large to generate a statistically significant coefficient estimate in the conventional migration regression. These patterns are consistent with many papers in the literature that find statistically insignificant conventional regression coefficients despite significant effects on local labor market outcomes. The relative MSE values are also quite large, in many cases exceeding 0.5, indicating poor performance capturing the actual shock-induced population changes.

Panel C then considers the model-consistent approaches described in Section 4.1, using the full market-to-market transition matrix. It is immediately clear that these approaches all perform far better than the conventional regression. The median parameter estimates are generally close to their theoretical target, $\theta/\sigma = 0.4075$, and the relative MSEs are far lower than those for the conventional regression.

We first estimate the model using nonlinear least squares, as in (28). The median coefficient estimates are all close to the target parameter, and, unlike in the conventional regression, the results are similar for iid region and iid industry shocks. The relative MSEs are also two to fifteen times smaller than those in the conventional regression. The estimates are statistically significant at the 5 percent level for large and medium shocks, but cannot be statistically distinguished from zero for

[32]Relative MSE values above one are theoretically possible if the method is either so poor at predicting population responses that it tends to get the sign wrong or is very noisy.

[33]The target magnitude for β is given by (25) in the case without attenuation, i.e. $\frac{2\theta}{\sigma} \cdot \frac{M}{N}$. Recall that in our calibration $\frac{\theta}{\sigma} = 0.4075$. From Table 1, $M/N = 0.123$, so the target magnitude is 0.100. The modest difference between this target and the conventional regression estimates for iid region shocks in Table 3 results from the approximations underlying (25) and the difference between the estimand in (24) and the expectation of (23) in finite samples.

small shocks. Thus, while our model-based approach improves power relative to the conventional regression, even a fully model-consistent analysis will struggle to identify effects on population growth when the observed shocks are small relative to unobserved drivers of population change.

The remaining model-consistent specifications rely on approximations to simplify the estimation procedure by yielding models that are linear in the parameters. First, the Taylor approximation in equation (30) is based on either low mobility or a small value of θ/σ. Quantitatively, this approach performs remarkably well across all analyses. The coefficients imply modest approximation bias relative to the target parameter, similarly for region- and industry-specific shocks. Even more impressively, the MSE values are close to those for NLLS. Given the ease of estimation, we view the Taylor approximation as a particularly favorable option for applied researchers.[34]

The direct mobility approximation in equation (31) imposes a stronger assumption than the Taylor approximation and experiences somewhat reduced performance as a result, both in terms of smaller point estimates and substantially larger relative MSE, particularly for industry shocks. Still, the direct mobility approximation performs vastly better than the conventional regression because, like the other model-consistent approaches, it uses model-consistent exposure measures, incorporates information on outside options, and implicitly accounts for heterogeneity in the local migration intensity (M_r/N_r).

One might expect the bilateral regression approach in (32) to yield superior results to the prior approaches because it has many more observations and more detailed information on migration flows. However, because the bilateral regression is derived based on the same direct mobility approximation as the prior estimator, it performs no better. Performance particularly degrades in terms of MSE for regional shocks.[35]

Together, we find that all the model-consistent approaches, including those based on approximations, exhibit vastly superior performance to the conventional regression in terms of MSE in nearly all cases and in terms of coefficient attenuation when facing industry shocks. Still, with small shocks even the model-consistent approaches struggle to measure the resulting population

[34] Another advantage of the Taylor approximation approach is robustness to the possibility that the observed shocks are a proportionally scaled version of the underlying demand shocks. As Appendix A.6 shows, we can implement a version of NLLS that incorporates a scaling factor $\psi > 0$ yielding similar population change predictions to those presented here but which is unable to separately identify ψ from θ/σ. In contrast, our specification based on the Taylor approximation directly identifies the response to appropriately scaled shocks.

[35] Appendix C.4 describes the simulation procedure for the bilateral regressions, and Appendix D provides intuition for why performance degrades more sharply for regional shocks than for industry shocks.

effects, with nearly all estimates statistically indistinguishable from zero. With medium shocks, the picture is more favorable, with many of the model-consistent approaches yielding statistically significant results. However, given that our medium shocks are quite large compared to those in prior literature, these findings emphasize the difficulty in estimating population effects, even for shocks with substantial effects on wages and other labor market outcomes.

While the strong performance of these model-consistent approaches is encouraging, they require that the analyst observes pre-shock transitions at the region-industry level, which in many contexts are unavailable or difficult to access. In Panel D of Table 3 we consider the approaches with reduced transition-data demands discussed in Section 4.2. Given the advantages of the Taylor approximation approach, we employ it here where possible.

We first consider the Taylor approximation ignoring industry variation, assuming that markets are purely regional (i.e., in a model with a single industry), and treating shift-share aggregates across industries as if they were regional shocks. Estimating this model requires transition information only at the region-to-region level, so it is feasible in most empirical contexts. Consistent with (21), for iid region shocks the results are similar to the Taylor approximation using full transition data between region-industry markets. In contrast, with iid industry shocks, the performance degrades sharply, with the MSE increasing by a factor of 2.4 to 6.5 compared to the fully model-consistent Taylor approximation for all shock sizes. The next row implements the bilateral regression analysis assuming purely regional markets and shocks, with similarly poor performance. In fact, with iid industry shocks, these bilateral MSEs approach those of the conventional regression. These findings reinforce one of our main conclusions: when shocks vary across industries, analyses incorporating only regional variation fail to accurately capture their effects on regional populations.

The last two rows of Table 3 assess the Taylor approximation when using region-industry transition matrices inferred from coarser (and more widely available) transition data. As discussed in Section 4.2, imputation 1 uses separate region-to-region and industry-to-industry data, assuming independent transitions across region and industry dimensions. While this approach exhibits larger MSEs than the version using full region-industry transition information, it still represents a substantial improvement over the conventional regression in all cases and avoids the attenuation experienced by the conventional regression with iid industry shocks. Imputation 2, which uses region-to-region-industry transitions and industry-to-industry transitions, exhibits slightly improved performance

with iid industry shocks, but is largely comparable to imputation 1. The results with both imputations make clear that it is essential to incorporate information on industry transitions when using shocks with industry variation, even when the most detailed data on market-to-market transitions are unavailable. In comparison to the conventional regression, the imputation-based Taylor approximations resolve the attenuation issue with iid industry shocks and much more accurately capture the shock-induced changes in regional populations.

In Appendix A.9, we further consider a variety of ad-hoc approaches to incorporating information on outside-option shocks to other regions motivated by those appearing in the prior literature. The regressions take the same form as the conventional regression in (22) with an additional independent variable aiming to capture shocks to outside-option markets. This outside-option control is a weighted average of $\hat{x}_k^{\mathrm{C}}$ in other regions (i.e. for $k \neq r$), with the weights differing by approach. The examples we study weight by pre-shock regional migration, the inverse of geographic distance, and both market size and distance, as in market access calculations. In all cases, these specifications perform far worse than the model-consistent approaches just discussed. Even when shocks are purely regional, these approaches are misspecified because they all omit regional heterogeneity in migration intensity, M_r/N_r, and because the distance- and market-access-based versions mismeasure relevant spillovers between regions given the model's assumptions. When shocks incorporate industry variation, the associated analyses perform particularly poorly because these averages entirely omit information on worker transitions across industries. In many cases, they yield MSEs that exceed even those of the conventional regression. These findings reemphasize the importance of incorporating regional heterogeneity in migration intensity and of incorporating information on industry transitions when shocks include industry variation.

7 Conclusion

This paper examines how to interpret conventional migration regression estimates that relate changes in local population to observed local labor demand shocks. This regression often yields small or insignificant coefficients, which suggests a puzzle: workers do not seem to migrate in response to such shocks, despite large effects on other outcomes, such as local wages. Using a simple equilibrium model with mobility costs across regions and industries, we show that the conventional

regression is misspecified. It omits shocks to relevant outside-option markets, omits important heterogeneity in effects, and, in the empirically relevant case of industry-level shocks, mismeasures the relevant average shock to the worker's current location. Estimates from the conventional migration regression are therefore biased, and acutely so when shocks are industry specific, as in the vast majority of empirical studies. We use Brazilian and German administrative data to show that these biases are quantitatively important and that predictions of population changes based on the conventional regression can be highly inaccurate, so researchers should be cautious in interpreting its estimates. We describe model-consistent empirical approaches with far superior performance, and show that similar approaches with reduced data demands can also perform far better than the conventional regression, as long as information on industry transitions is appropriately incorporated into analyses of industry-level shocks.

We implement these analyses in the context of a relatively parsimonious model, which yields tractability and allows us to focus on general problems with the conventional migration regression. Incorporating additional mechanisms will in some cases imply richer measures of shock exposure and outside options (cf. Appendix A.1 on agglomeration economies and Appendix A.2 on housing markets). Still, these richer models will admit model-consistent empirical approaches that parallel those we introduce here.

The insights from our analysis may have implications for other findings on worker mobility across space. First, while conventional regression estimates tend to suggest that highly educated workers move more in response to local economic shocks than less educated workers (e.g. Bound and Holzer 2000; Wozniak 2010), it is possible that the gap in coefficients across these two groups reflects differences in attenuation factors resulting from omitted outside-option shocks rather than differences in migration elasticities or rates. Second, a large literature has sought to explain steadily declining migration rates in the U.S. since the late 1980s (see Molloy et al. 2011 and Jia et al. 2023 for surveys). Our model suggests a novel alternative explanation: economic shocks may have become more correlated among migrant-connected regions, or there may be a growing share of the variation in economic shocks occurring at the level of national industries rather than regions.

Our focus on the network structure of worker reallocation can also bear fruit beyond the spatial context. For example, studies of the economic impacts of import competition at the industry level (e.g. Bernard et al. 2006 and Autor et al. 2014) may face similar issues due to spillovers in the

effects of shocks across industries. In fact, a recent paper by Kovak and Morrow (2025) explicitly draws on our proposed empirical approaches by using pre-shock industry transitions to construct outside-option tariff changes in the context of the Canada-US Free Trade Agreement. They find that tariff shocks to outside-option industries played an important role in driving Canadian workers' labor market outcomes.

References

Adão, Rodrigo (2025). "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil". *Working Paper*.

Adão, Rodrigo, Costas Arkolakis, and Federico Esposito (forthcoming). "General Equilibrium Effects in Space: Theory and Measurement". *American Economic Journal: Macroeconomics*.

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales (2019). "Shift-Share Designs: Theory and Inference". *The Quarterly Journal of Economics* 134(4): 1949–2010.

Agrawal, David R. and Dirk Foremny (2019). "Relocation of the Rich: Migration in Response to Top Tax Rate Changes from Spanish Reforms". *The Review of Economics and Statistics* 101(2): 214–232.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf (2015). "The Economics of Density: Evidence From the Berlin Wall". *Econometrica* 83(6): 2127–2189.

Albert, Christoph, Paula Bustos, and Jacopo Ponticelli (2021). "The effects of climate change on labor and capital reallocation". *Working Paper*.

Albert, Christoph and Joan Monràs (2020). "The Regional Impact of Economic Shocks: Why Immigration is Different from Import Competition". *Working paper*.

Amior, Michael and Alan Manning (2019). "Commuting, Migration and Local Joblessness". *CEP Discussion Paper No 1623* (October): 1–45.

Anderson, James E. and Eric van Wincoop (2003). "Gravity with Gravitas: A Solution to the Border Puzzle". *American Economic Review* 93(1): 170–192.

Armington, Paul S. (1969). "A Theory of Demand for Products Distinguished by Place of Production". *IMF Staff Papers* 16(1): 159–178.

Autor, David, David Dorn, and Gordon H. Hanson (2013). "The China Syndrome: Local Labor Market Impacts of Import Competition in the United States". *American Economic Review* 103(6): 2121–2168.

Autor, David, David Dorn, and Gordon H. Hanson (2021). "On the Persistence of the China Shock". *Brookings Papers on Economic Activity* Fall 2021: 381–447.

Autor, David, David Dorn, and Gordon H. Hanson (2023). *Trading Places: Mobility Responses of Native and Foreign-Born Adults to the China Trade Shock*. Working Paper 30904. National Bureau of Economic Research.

Autor, David, David Dorn, Gordon H. Hanson, and Jae Song (2014). "Trade Adjustment: Worker-Level Evidence". *The Quarterly Journal of Economics* 129(4): 1799–1860.

Bartik, Alexander W., Janet Currie, Michael Greenstone, and Christopher R. Knittel (2019). "The local economic and welfare consequences of hydraulic fracturing". *American Economic Journal: Applied Economics* 11(4): 105–155.

Berkes, Enrico, Ruben Gaetani, and Martí Mestieri (2025). "Technological Waves, Knowledge Diffusion, and Local Growth". *Journal of Political Economy Macroeconomics* 3(1): 75–121.

Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott (2006). "Survival of the best fit: Exposure to low-wage countries and the (uneven) growth of U.S. manufacturing plants". *Journal of International Economics* 68(1): 219–237.

Black, Dan, Terra McKinnish, and Seth Sanders (2005). "The economic impact of the coal boom and bust". *Economic Journal* 115(503): 449–476.

Blanchard, Olivier Jean and Lawrence F. Katz (1992). "Regional Evolutions". *Brookings Papers on Economic Activity* (1): 1–75.

Borusyak, Kirill and Peter Hull (2021). "Non-random exposure to exogenous shocks: Theory and applications". *National Bureau of Economic Research Working Paper*.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022). "Quasi-Experimental Shift-Share Research Designs". *Review of Economic Studies* 89(1): 181–213.

Bound, John and Harry J. Holzer (2000). "Demand shifts, population adjustments, and labor market outcomes during the 1980s". *Journal of Labor Economics* 18(1): 20–54.

Boustan, Leah, Matthew E. Kahn, Paul W. Rhode, and Maria Lucia Yanguas (2020). "The Effect of Natural Disasters on Economic Activity in US Counties: A Century of Data". *Journal of Urban Economics* 118(103257).

Broda, Christian and David E. Weinstein (2006). "Globalization and the Gains From Variety". *The Quarterly Journal of Economics* 121(2): 541–585.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli (2016). "Agricultural Productivity and Structural Transformation: Evidence from Brazil". *American Economic Review* 106(6): 1320–65.

Cadena, Brian C. and Brian K. Kovak (2016). "Immigrants equilibrate local labor markets: Evidence from the great recession". *American Economic Journal: Applied Economics* 8(1): 257–290.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro (2019). "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock". *Econometrica* 87(3): 741–835.

Choi, Jiwon, Ilyana Kuziemko, Ebonya Washington, and Gavin Wright (2024). "Local Economic and Political Effects of Trade Deals: Evidence from NAFTA". *American Economic Review* 114(6): 1540–1575.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2014). "The rise of the east and the far east: German labor markets and trade integration". *Journal of the European Economic Association* 12(6): 1643–1675.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2021). "Adjusting to Globalization in Germany". *Journal of Labor Economics* 39(1): 263–302.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008). "Global Rebalancing with Gravity: Measuring the Burden of Adjustment". *IMF Staff Papers* 55(3): 511–540.

Dix-Carneiro, Rafael and Brian K. Kovak (2015). "Trade liberalization and the skill premium: A local labor markets approach". *American Economic Review: Papers & Proceedings* 105(5): 551–557.

Dix-Carneiro, Rafael and Brian K. Kovak (2017). "Trade Liberalization and Regional Dynamics". *American Economic Review* 107(10): 2908–2946.

Dix-Carneiro, Rafael and Brian K. Kovak (2019). "Margins of Labor Market Adjustment to Trade". *Journal of International Economics* 117: 125–142.

Faber, Marius, Andrés Sarto, and Marco Tabellini (2021). "Local Shocks and Internal Migration: The Disparate Effects of Robots and Chinese Imports in the US". *IZA Discussion Paper No. 14623*.

Fan, Jingting (2019). "Internal Geography, Labor Mobility, and the Distributional Impacts of Trade". *American Economic Journal: Macroeconomics* 11(3): 252–88.

Feenstra, Robert, Philip Luck, Maurice Obstfeld, and Katheryn Russ (2018). "In Search of the Armington Elasticity". *Review of Economics and Statistics* 100(1): 135–150.

Foote, Andrew, Michel Grosz, and Ann Stevens (2019). "Locate Your Nearest Exit: Mass Layoffs and Local Labor Market Response". *Industrial and Labor Relations Review* 72(1): 101–126.

Foschi, Andrea, Christopher L. House, Christian Proebsting, and Linda L. Tesar (2025). "Should I Stay or Should I Go? The Response of Labor Migration to Economic Shocks". *Brookings Papers on Economic Activity* Spring 2025: 67–114.

Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi (2023). "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade". *The Review of Economic Studies* 90(1): 331–375.

Greenland, Andrew, John Lopresti, and Peter McHenry (2019). "Import Competition and Internal Migration". *The Review of Economics and Statistics* 101(1): 44–59.

Hakobyan, Shushanik and John McLaren (2016). "Looking for Local Labor Market Effects of NAFTA". *The Review of Economics and Statistics* 98(4): 728–741.

Head, Keith and Thierry Mayer (2014). "Chapter 3 - Gravity Equations: Workhorse, Toolkit, and Cookbook". In: *Handbook of International Economics*. Ed. by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff. Vol. 4. Handbook of International Economics. Elsevier, pp. 131–195.

Head, Keith and Thierry Mayer (2021). "The United States of Europe: A gravity model evaluation of the four freedoms". *Journal of Economic Perspectives* 35(2): 23–48.

Jia, Ning, Raven Molloy, Christopher Smith, and Abigail Wozniak (2023). "The Economics of Internal Migration: Advances and Policy Questions". *Journal of Economic Literature* 61(1): 144–80.

Kleven, Henrik Jacobsen, Camille Landais, and Emmanuel Saez (2013). "Taxation and International Migration of Superstars: Evidence from the European Football Market". *American Economic Review* 103(5): 1892–1924.

Kovak, Brian K and Peter M Morrow (2025). "The Long-Run Labour Market Effects of the Canada-U.S. Free Trade Agreement". *The Review of Economic Studies* 92(6): 4026–4058.

Kovak, Brian K. (2013). "Regional effects of trade reform: What is the correct measure of liberalization?" *American Economic Review* 103(5): 1960–1976.

McCaig, Brian (2011). "Exporting out of poverty: Provincial poverty in Vietnam and U.S. market access". *Journal of International Economics* 85(1): 102–113.

McCaig, Brian, Minh Nguyen, and Robert Kaestner (2022). "Export Expansion and Investment in Children's Human Capital: Evidence from the U.S.-Vietnam Bilateral Trade Agreement". *NBER Working Paper* (29830).

Mian, Atif and Amir Sufi (2014). "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82(6): 2197–2223.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak (2011). "Internal Migration in the United States". *Journal of Economic Perspectives* 25(3): 173–96.

Monràs, Joan (2020). "Economic Shocks and Internal Migration". *SSRN Working Paper*.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg (2018). "Commuting, Migration, and Local Employment Elasticities". *American Economic Review* 108(12): 3855–3890.

Morten, Melanie and Jacqueline Oliveira (2018). "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City". *NBER Working Paper* (22158).

Notowidigdo, Matthew J. (2020). "The Incidence of Local Labor Demand Shocks". *Journal of Labor Economics* 38(3): 687–725.

Seo, Jaeeun and Ryungha Oh (2024). "Sectoral Shocks and Labor Market Dynamics: A Sufficient Statistics Approach". *Working paper*.

Sprung-Keyser, Ben, Nathaniel Hendren, and Sonya Porter (2022). "The Radius of Economic Opportunity: Evidence from Migration and Local Labor Markets". *working paper*.

Tombe, Trevor and Xiaodong Zhu (2019). "Trade, Migration, and Productivity: A Quantitative Analysis of China". *American Economic Review* 109(5): 1843–72.

Topalova, Petia (2010). "Factor Immobility and Regional Impacts of Trade Liberalization : Evidence on Poverty from India". *American Economic Journal: Applied Economics* 2(4): 1–41.

Topel, Robert H. (1986). "Local labor markets". *Journal of Political Economy* 94(3): S111–S143.

Wozniak, Abigail (2010). "Are College Graduates More Responsive to Distant Labor Market Opportunities?" *Journal of Human Resources* 45(4): 944–970.

Yagan, Danny (2014). "Moving to Opportunity? Migratory Insurance over the Great Recession". *Working Paper*.

Yagan, Danny (2019). "Employment Hysteresis from the Great Recession". *Journal of Political Economy* 127(5): 2505–2558.

Online Appendix

A Additional Results

A.1 Agglomeration

In this appendix, we show that incorporating local agglomeration economies yields results that are isomorphic to those in the Section 2 model. Suppose productivity in each market increases in L_{ri} with constant elasticity $\lambda \in (0, 1/(\sigma-1))$, and denote $\tilde{\lambda} = (1-\lambda(\sigma-1))^{-1} \geq 1$. In this case, the expression for local labor demand is isomorphic to (6); the presence of agglomeration economies simply increases the labor demand elasticity.

$$\begin{aligned} L_{ri} &= \eta_i \left(\varphi_{ri} L_{ri}^{\lambda} P_i\right)^{\sigma-1} w_{ri}^{-\sigma} \\ &= \left(\eta_i \left(\varphi_{ri} P_i\right)^{(\sigma-1)}\right)^{\tilde{\lambda}} w_{ri}^{-\sigma\tilde{\lambda}}. \end{aligned}$$

A.2 Housing Market

In this appendix, we add a residential land market to the model. We follow Redding (2016) by assuming the quantity of local land is fixed, which introduces an additional congestion force in spatial equilibrium beyond that already generated in the baseline model by downward-sloping local labor demand. We simplify the model, however, to having $I = 1$ industry, as in Section 4.2 under "Treating Shocks as Region-Specific."[36]

Suppose that utility from consumption C and land (i.e., housing) H is Cobb-Douglas,

$$U = \left(\frac{C}{\alpha}\right)^{\alpha} \left(\frac{H}{1-\alpha}\right)^{1-\alpha}. \tag{A1}$$

Consumption is nationally traded at price P, and housing is priced locally at P_r^H. With fixed land supply $\bar{H}_r$, the market clearing condition for land is

$$\bar{H}_r = (1-\alpha) N_r \frac{v_r}{P_r^H}, \tag{A2}$$

where v_r is income per worker in region r, which includes both the wage w_r (varying only by region since $I = 1$) and land rents, assumed to be distributed to consumers with a lump-sum transfer, such that

$$v_r N_r = w_r N_r + (1-\alpha) v_r N_r, \quad \Longrightarrow \quad v_r = \frac{w_r}{\alpha}. \tag{A3}$$

From (A2) and (A3),

$$P_r^H = \frac{1-\alpha}{\alpha} \frac{w_r N_r}{\bar{H}_r}.$$

Thus, the indirect utility in region r (which replaces the real wage in (4)) is

$$U_r = \frac{v_r}{P^{\alpha} \left(P_r^H\right)^{1-\alpha}} = \frac{w_r \bar{H}_r^{1-\alpha}}{\alpha P^{\alpha} \left(\frac{1-\alpha}{\alpha}\right)^{1-\alpha} w_r^{1-\alpha} N_r^{1-\alpha}} = \left(\frac{w_r}{\alpha P}\right)^{\alpha} \left(\frac{\bar{H}_r}{(1-\alpha) N_r}\right)^{1-\alpha}. \tag{A4}$$

[36] Unlike in Section 4.2, here this simplification is made for tractability. Changes in the demand for housing depend on changes in the average regional wage. The latter reflects not only market-level wage changes but also the reallocation of workers across markets within region; see Appendix C.3 for further discussion of this issue.

Log-differentiating (A4) with fixed land supply,

$$\hat{U}_r = \alpha \hat{w}_r - (1-\alpha)\hat{N}_r - \alpha \hat{P}. \tag{A5}$$

Changes in labor supply now depend on differential changes in U_r rather than w_r:

$$\widehat{\mathbf{N}} = \theta \left(\mathbb{I} - \Gamma' \Pi\right) \widehat{\mathbf{U}} = \theta \left(\mathbb{I} - \Gamma' \Pi\right) \left(\alpha \widehat{\mathbf{w}} - (1-\alpha)\widehat{\mathbf{N}}\right), \tag{A6}$$

where the term $\alpha \hat{P}$ disappeared since it is the same in all regions.

Since firms do not use housing in this model, labor demand is the same as in the baseline model:

$$\widehat{\mathbf{L}} \equiv \widehat{\mathbf{N}} = \widehat{\mathbf{D}} - \sigma \widehat{\mathbf{w}}, \tag{A7}$$

The system (A6)–(A7) can be solved to describe the effect of local labor demand shocks on local populations:

$$\widehat{\mathbf{N}} = \left[\mathbb{I} + \left((1-\alpha)\sigma + \alpha\right) \frac{\theta}{\sigma} \left(\mathbb{I} - \Gamma' \Pi\right)\right]^{-1} \cdot \frac{\alpha\theta}{\sigma} \left(\mathbb{I} - \Gamma' \Pi\right) \widehat{\mathbf{D}}. \tag{A8}$$

Note that this expression depends on the structural parameters (θ, σ, α) via two combinations, $\alpha\theta/\sigma$ and $\left((1-\alpha)\sigma + \alpha\right)\frac{\theta}{\sigma}$, compared to the single reduced form parameter θ/σ in the baseline model. However, when we reintroduce the Taylor approximation in $\theta \to 0$, the resulting expression simplifies to exactly the form we found in the baseline model with pure region shocks, where the effect of local shocks depends only upon $\alpha\theta/\sigma$:

$$\widehat{\mathbf{N}} \approx \frac{\alpha\theta}{\sigma} \left(\mathbb{I} - \Gamma' \Pi\right) \widehat{\mathbf{D}}. \tag{A9}$$

It is easy to show that, as long as $\theta \left(\mathbb{I} - \Gamma' \Pi\right)$ is small, the direct-mobility approximation generalizes similarly. Because these expressions take the same form as in the baseline model, all of our results apply directly. The only change required is to appropriately calibrate the parameter $\alpha\theta/\sigma$ to account for the presence of housing demand.

A.3 Population Responses to Large Labor Demand Shocks

For expositional clarity and analytical tractability, the analysis in the main text studies small shocks relative to a no-shock counterfactual, such that the first-order approximation around the no-shock equilibrium is accurate (see Section 2.3). Here, we use exact hat algebra to determine how population responds to large shocks, following Dekle et al. (2008), allowing us to empirically assess how well the small-shock approximation performs in practice.

Let hats now denote exact log-changes. The labor demand equation (7) does not require the small-shock approximation and is therefore unchanged. The labor supply equation and the equilibrium characterization, however, do change. We have:

$$\begin{aligned}
\hat{L}_{ri} \equiv \log \frac{L^1_{ri}}{L_{ri}} &= \log \left(\frac{\sum_{o,p} L^0_{op} \frac{\left(w^1_{ri}/\tau_{op,ri}\right)^\theta}{\sum_{d,q}\left(w^1_{dq}/\tau_{op,dq}\right)^\theta}}{\sum_{o,p} L^0_{op} \frac{\left(w_{ri}/\tau_{op,ri}\right)^\theta}{\sum_{d,q}\left(w_{dq}/\tau_{op,dq}\right)^\theta}} \right) \\
&= \log \left(\sum_{o,p} \gamma_{op,ri} \frac{\exp \theta \hat{w}_{ri}}{\sum_{d,q} \left(w^1_{dq}/\tau_{op,dq}\right)^\theta / \sum_{d,q} \left(w_{dq}/\tau_{op,dq}\right)^\theta} \right)
\end{aligned}$$

$$= \theta \hat{w}_{ri} + \log \left(\sum_{o,p} \gamma_{op,ri} \left(\sum_{d,q} \pi_{op,dq} \exp \theta \hat{w}_{dq} \right)^{-1} \right)$$

$$= \frac{\theta}{\sigma} \left(\hat{D}_{ri} - \hat{L}_{ri} \right) + \log \left(\sum_{o,p} \gamma_{op,ri} \left(\sum_{d,q} \pi_{op,dq} \exp \frac{\theta}{\sigma} \left(\hat{D}_{dq} - \hat{L}_{dq} \right) \right)^{-1} \right), \quad \text{(A10)}$$

where the first line used the choice probabilities in (9), the second and third lines used standard exact hat algebra manipulations, and the last line plugged in (7). Here, (A10) is a system of nonlinear equations yielding $\widehat{\mathbf{L}}$ in response to any vector of labor demand shocks $\widehat{\mathbf{D}}$. It generalizes (13) in the sense that the log-linear approximation to (A10) gives $\widehat{\mathbf{L}} = \Omega \widehat{\mathbf{D}}$ as the solution. Aggregating into regional population changes,

$$\hat{N}_r = \log \frac{\sum_i L_{ri}^1}{\sum_i L_{ri}} = \log \sum_i \frac{L_{ri}}{N_r} \exp \hat{L}_{ri}. \quad \text{(A11)}$$

By solving (A10) numerically for a given shock vector $\widehat{\mathbf{D}}$ and plugging in the results to (A11), we can check the quality of the approximation in (13).

We now establish that, in our Brazilian data, the small-shock approximation is accurate with pure regional shocks of various sizes as well as "small" pure industry shocks from Table 3. Specifically, we simulate "large," "medium," and "small" pure region and pure industry shocks, as in Section 6.3 but without unobserved shocks.[37] We then generate the true population responses to these observed shocks, using (A10)–(A11). The first row of Table A1 regresses the resulting population changes on the small-shock approximation for them, as in (13).[38] The relative MSE reported in the table is essentially zero for pure regional shocks of all sizes we consider, implying that the small-shock approximation is innocuous in these scenarios. The approximation is less precise with pure industry shocks, with the relative MSE reaching 0.219 for the "large" shocks. This is because the industry shocks we consider are much larger than the regional shocks (see Panel A of Table 3). Yet, for the "small" shocks that are most comparable to the Brazilian trade liberalization (see Section 6.3), the relative MSE is only 0.058, suggesting that the small-shock approximation is sufficiently accurate and cannot materially hinder the quality of our model-consistent methods, as well as our conclusions about the problems with the conventional regression.

The remaining rows of Table A1 complement these findings by performing the core analyses from Table 3 but without the small-shock approximation. We perform model-consistent methods — NLLS, the Taylor approximation, and the direct-mobility approximation — and estimate the conventional regression. The patterns are consistent with those in the first row. As long as the shocks are purely regional of any size or "small" pure-industry, all model-consistent methods (derived using the small-shock approximation) are very accurate when the data are generated without the small-shock approximation. In contrast, the conventional migration regression makes relatively poor predictions even absent unobserved shocks.

[37] We do not include unobserved shocks in this simulation for the following reason. Without the small-shock approximation, population changes in (A10)–(A11) are not additively separable in the effects of observed and unobserved shocks. Thus, there is more than one natural definition of the effect of observed shocks: e.g., fixing the unobserved shocks at zero vs. at their realized value. We find it more informative to report the MSE when the estimand of interest is uniquely defined.

[38] This regression is infeasible, as constructing the true Ω requires knowing the true θ/σ. The results are similar with the feasible regressions described next.

Table A1: Precision of the Small-Shock Approximation

Shock size:	large		medium		small	
Shock type:	iid region	iid industry	iid region	iid industry	iid region	iid industry
Small-shock approx.	{0.000}	{0.219}	{0.000}	{0.133}	{0.000}	{0.058}
NLLS	{0.001}	{0.217}	{0.000}	{0.133}	{0.000}	{0.058}
Taylor approx.	{0.003}	{0.237}	{0.003}	{0.155}	{0.003}	{0.077}
Direct-mob. approx.	{0.018}	{0.338}	{0.017}	{0.264}	{0.017}	{0.194}
Conventional migration reg.	{0.189}	{0.494}	{0.192}	{0.458}	{0.192}	{0.439}

Notes: Each cell in this table reports the relative mean squared errors, computed as the ratio of the median of (34) across 601 simulations to the median of (35), for different regression specifications (as rows) and data-generating processes for the observed shocks (as columns). The true regional population responses on the left-hand side are generated using the exact hat algebra method in (A10)–(A11), with no unobserved shocks. See Appendix A.3 for specification details; the columns are as in Table 3.

A.4 Incorporating Labor Supply Shocks

In this appendix we show that the expressions from Section 2 extend naturally when the moving costs $\tau_{ri,dq}$, in addition to D_{ri}, are allowed to change, generating local labor supply shocks. In that case, (11) is replaced by

$$\begin{aligned}\hat{L}_{ri} &= \sum_{o,p} \gamma_{op,ri}\hat{\pi}_{op,ri} \\ &= \theta\left(\hat{w}_{ri} - \sum_{o,p}\gamma_{op,ri}\sum_{d,q}\pi_{op,dq}\hat{w}_{dq}\right) + \nu_{ri}\end{aligned}$$

where the labor supply shock is defined as

$$\nu_{ri} = -\theta\sum_{o,p}\gamma_{op,ri}\left(\hat{\tau}_{op,ri} - \sum_{d,q}\pi_{op,dq}\hat{\tau}_{op,dq}\right).$$

A destination market ri receives a large supply shock when moving costs into ri fall relative to moving costs from the same origins to alternative destinations. In vector form,

$$\widehat{\mathbf{L}} = \theta\left(\mathbb{I} - \Gamma'\Pi\right)\widehat{\mathbf{w}} + \boldsymbol{\nu} \tag{A12}$$

with $\boldsymbol{\nu} = \{\nu_{ri}\}$. Solving the system of demand and supply (i.e. combining with (8)), this implies $\widehat{\mathbf{L}} = \Omega\widehat{\mathbf{D}} + (\mathbb{I} - \Omega)\boldsymbol{\nu}$. Finally, isolating observed and unobserved shocks as in (15) yields

$$\widehat{\mathbf{L}} = \Omega\widehat{\mathbf{z}} + \boldsymbol{\zeta} + (\mathbb{I} - \Omega)\boldsymbol{\nu},$$

which generates a second, supply-driven, error term. The results of the paper continue to hold if the $\widehat{\mathbf{z}}$ shocks are exogenous in the sense of being mean-independent from both unobserved labor demand and labor supply shocks.

A.5 Characterizing β under Heteroskedasticity

We state a generalization of Theorem 2 which applies with heteroskedastic shocks:

Theorem A1. *Suppose the data are generated by the direct-mobility approximation to the model in (19) with $L_{ri} \approx L^0_{ri}$. Then*

$$\beta = \frac{2\theta}{\sigma} \cdot \frac{M^v/N}{\widetilde{M}^v/N} \cdot \frac{\rho^M - \rho}{1 - \widetilde{\rho}}. \tag{A13}$$

Here $\frac{M^v}{N} \equiv \frac{\sum_r M_r v_r}{\sum_r N_r v_r}$ is the national average migration share weighted by region-specific conventional regressor variances $v_r = \text{Var}[\hat{x}^C_r]$, $\rho^M \equiv (\sum_r M_r v_r \rho^M_r)/(\sum_r M_r v_r)$ for $\rho^M_r \equiv \text{Cov}[\hat{x}^C_r, \hat{x}^M_r]/v_r$, $\rho = (\sum_r M_r v_r \rho_r)/(\sum_r M_r v_r)$ for $\rho_r = \text{Cov}\left[\hat{x}^C_r, \hat{x}^M_{-r}\right]/v_r$. Similarly, $\frac{\widetilde{M}^v}{N} \equiv \frac{\sum_r \widetilde{M}_r v_r}{\sum_r N_r v_r}$, $\widetilde{\rho} = \frac{\sum_r \widetilde{M}_r v_r \widetilde{\rho}_r}{\sum_r \widetilde{M}_r v_r}$, and $\widetilde{\rho}_r = \text{Cov}\left[\hat{x}^C_r, \widetilde{x}^C_{-r}\right]/v_r$ for $\widetilde{x}^C_{-r} = \sum_{k \neq r}(\widetilde{F}_{rk}/\widetilde{M}_r)\hat{x}^C_k$. Moreover, when the $\hat{x}^C_r$ are homoskedastic, i.e. $v_r = v$ for all r, then ρ^M, ρ, and $\widetilde{\rho}$ coincide with the Theorem 2 definitions, $M^v = M$, and $\widetilde{M}^v = \widetilde{M}$.

A.6 Unknown Shock Scale

Suppose the researcher believes that

$$\hat{D}_{ri} = \psi \hat{z}_{ri} + \zeta_{1,ri}$$

with unknown $\psi > 0$ converting the units of $\hat{z}_{ri}$ into the units of labor demand. In that case, market-level employment changes follow

$$\widehat{\mathbf{L}} = \psi \Omega \widehat{\mathbf{z}} + \boldsymbol{\zeta}$$

and regional population changes satisfy

$$\hat{N}_r = \psi S_r' \Omega \widehat{\mathbf{z}} + \zeta_r.$$

This suggests an extended NLLS procedure:

$$\left(\hat{\psi}, \widehat{\theta/\sigma}\right) = \arg\min_{\psi,\lambda} \sum_r \frac{N_r}{N} \left(\hat{N}_r - \psi S_r' \Omega(\lambda) \widehat{\mathbf{z}}\right)^2$$

as well as its Taylor and low-mobility approximations.

However, ψ and θ/σ are only weakly identified separately. To see the issue most clearly, consider the analog of the Taylor approximation in (29):

$$\hat{N}_r \approx \psi \frac{\theta}{\sigma} S_r' \left(\mathbb{I} - \Gamma' \Pi\right) \widehat{\mathbf{z}} + \zeta_r.$$

Here ψ and θ/σ are entirely indistinguishable and we can only hope to estimate $\psi\theta/\sigma$. The same phenomenon arises under the direct-mobility approximation. However, conveniently, identifying them separately is not needed to estimate the effects of any observed or counterfactual shocks measured in the units of $\widehat{\mathbf{z}}$.

In the NLLS procedure, the nonlinearities in the model theoretically allow one to separate the two parameters. However, in unreported simulations we have found this identification to be very weak: while the estimates of $\psi\theta/\sigma$ are relatively stable, those of ψ and θ/σ separately are *extremely* noisy. But the flip side of this weak identification problem is that only $\psi\theta/\sigma$ matters for counterfactuals while the role of ψ and θ/σ separately is minimal — it does not limit our ability to make precise predictions for the population effects of labor demand shocks.

A.7 Model with Non-Gumbel Idiosyncratic Shocks

The role of marginal workers in our model does not rely on the Gumbel assumption. Suppressing the individual index ω for brevity, each individual nationally is characterized in our model by random variables $(O, \epsilon_{11}, \ldots, \epsilon_{RI})$, where O is the origin market[39] and

$$V_{ri} = \log\left(\frac{w_{ri}}{P}\right) - \log \tau_{O,ri} + \frac{1}{\theta}\epsilon_{ri}$$

is utility from choosing destination market ri. Unlike in the main analysis, we only require that $(\epsilon_{11}, \ldots, \epsilon_{RI})$ have a continuous joint density satisfying mild regularity conditions stated in Appendix B.10.

We describe marginal workers using their two preferred markets in the no-shock equilibrium. Let $A = \arg\max_{ri} V_{ri}$ be the preferred market, let $B = \arg\max_{ri \neq A} V_{ri}$ be the second-best market, and measure the gap between them in units of the taste shocks as

$$G = \theta(V_A - V_B) > 0.$$

For $dq \neq ri$, let

$$\kappa_{ri,dq} = \lim_{g \downarrow 0} \frac{\Pr(B = dq, G < g \mid A = ri)}{g}$$

measure the density of workers who are marginal between ri and dq, among those choosing ri.

We now show that the mass of marginal workers, $\kappa_{ri,dq}$, is key for two results in the model: it governs the labor supply response to wage changes (extending (11) to the general distribution of idiosyncratic preferences) and, when θ is small, it also governs the equilibrium employment effects of labor demand shocks (generalizing (29)). To see the first, let $A^1 = \arg\max_{ri}(V_{ri} + \widehat{w}_{ri})$ denote the preferred market after a vector of small wage changes, so that

$$\widehat{L}_{ri} = \frac{\Pr(A^1 = ri) - \Pr(A = ri)}{\Pr(A = ri)}.$$

We show that

$$\widehat{L}_{ri} = \theta \sum_{dq \neq ri} \kappa_{ri,dq}(\widehat{w}_{ri} - \widehat{w}_{dq}) \tag{A14}$$

(see Appendix B.10 for the proofs). Thus, the response to a relative wage change between two markets is proportional to the density of workers who are marginal between those markets.

The Gumbel assumption links these marginal densities to average transition probabilities in Π and Γ. Indeed, when the ϵ_{ri} are iid standard Gumbel (and independent of O), we show that

$$\kappa_{ri,dq} = \sum_{op} \gamma_{op,ri}\pi_{op,dq}, \qquad dq \neq ri, \tag{A15}$$

so (A14) reduces to (11).

Second, we characterize the equilibrium responses of market-level employment to labor demand

[39] This notation allows us to represent the choices made by individuals randomly drawn from the national population, rather than separately for each origin.

shocks when θ is small. [40] Combining labor supply with labor demand, we show that

$$\widehat{L}_{ri} \approx \frac{\theta}{\sigma} \sum_{dq \neq ri} \kappa_{ri,dq} (\widehat{D}_{ri} - \widehat{D}_{dq}). \tag{A16}$$

This characterization extends the outside-option logic of our model beyond the Gumbel distribution, showing again that the mass of marginal workers, $\kappa_{ri,dq}$, plays the central role.

Under Gumbel preferences, (A16) reduces to the Taylor approximation underlying (29). The direct-mobility approximation in Section 2.6 gives a complementary link to observed movers: with low mobility, (A15) implies

$$\kappa_{ri,dq} \approx \gamma_{dq,ri} + \pi_{ri,dq},$$

so the bilateral transition weights in (18) measure the same marginal densities that govern the general wage response in (A14), and the small-θ demand-shock response in (A16). This link from marginal to average mobility is the role of the Gumbel assumption highlighted in footnote 14.

A.8 Interpretations of β in Prior Literature

In this appendix, we present examples of prior work interpreting β in the conventional migration regression as capturing the effect of a counterfactual vector of shocks or as the effects of the observed vector of shocks themselves. As discussed in Section 3.3, both sets of interpretations are generally incorrect when workers face heterogeneous costs of transitioning between region-industry labor markets, as in our model. These interpretation issues are pervasive throughout the literature; the examples presented here were chosen for the clarity of the relevant quotes, not because the interpretations were more problematic than those in other papers. Our clarifications are added in brackets where appropriate.

A.8.1 Effects of Counterfactual Shocks

Starting with the effects of counterfactual shocks, we distinguish between two types of quotes. The first group of examples interprets β as capturing the effects of a counterfactual shock to a single location. In all cited cases, the independent variables are shift-share aggregates of industry-level shocks, as in $\hat{x}_r^{\mathrm{C}}$.

- Bound and Holzer (2000, p.41-42): "The estimates imply that a 10% shift in demand away from an area [i.e., their $\hat{x}_r^{\mathrm{C}}$] would lead to declines in population for all experience groups of about 5% among college graduates and nothing among those with high school or less..."
- Dauth et al. (2014, footnote 21): "… an increase in export (import) exposure by 1000 euro would increase (decrease) the regional population size by as little as 0.168 (0.165) per cent."
- Greenland et al. (2019, p.48): "Specifically, the point estimate implies that an interquartile increase in the NTR gap would have reduced a [commuting zone's] population growth by 0.015 log points ..."

[40] Formally, this involves considering a sequence of economies with $\theta \to 0$ that holds the observed data — specifically, the initial market employment vector and the no-shock transition structure — fixed by adjusting moving costs $\tau_{ri,dq}$ as needed. We implement this sequence by holding deterministic utility differences in taste-shock units fixed, so the $\kappa_{ri,dq}$ are fixed as well; in the Gumbel case, Π and, consequently, Γ are fixed. This approach is analogous to how, in a gravity trade model, Demidova et al. (2024) allow trade costs to adjust along a limiting sequence to prevent key trade shares from becoming degenerate.

In the second group of examples, the authors interpret a small and/or statistically insignificant estimate of β as implying a general lack of population responsiveness to local labor demand shocks, implicitly making inferences about the effects of *any* counterfactual vector of shocks.

- Autor et al. (2013, p.2141-42): "We find no robust evidence, however, that shocks to local manufacturing lead to substantial changes in population."
- Cadena and Kovak (2016, p.259): "Having established [using the conventional regression] that less-skilled Mexicans are highly geographically responsive to changes in labor market conditions while less-skilled natives are not ..."
- Dauth et al. (2014, p.1665-66): "Hence, in line with [the Autor et al. (2013)] results, we also find that the local adjustments to trade shocks do not seem to involve major population shifts across regions." [Note that this example implies a lack of responsiveness to all *trade* shocks rather than any local labor demand shocks.]
- Galle et al. (2017, footnote 13): "Our assumption of fixed groups applied to this setting implies no mobility across local labor markets and education categories. We view this as a reasonable assumption in light of existing literature that finds little evidence of trade exposure causing population shifts across local labor markets."[41]
- Adão (2025, footnote 7): "This paper abstracts from migration flows between regional labor markets in Brazil. This simplification is motivated by the empirical analysis in Section 4 [using the conventional regression] where I do not find evidence of migration responses following regional labor demand shocks."

A.8.2 Effects of Observed Shocks

The following quotes concern the effects of the observed shocks used to estimate the conventional migration regression rather than counterfactual shocks. Specifically, they interpret small and/or statistically insignificant estimates of β as implying that local populations did not respond substantially to the observed shocks.

- Dauth et al. (2014, p.1665): "In their analysis, [Autor et al. (2013)] emphasize sluggish adjustment of population across local labor markets. In other words, they find no evidence for population shifts in response to rising Chinese import competition in the United States."
- Dix-Carneiro and Kovak (2017, p.2923): "Table 3 rules out the interregional migration mechanism, showing that a region's working-age population did not respond to RTR_r" [Here, RTR_r is a regional shift-share aggregate of industry-level tariff changes.]
- Yagan (2019, p.2542): "The estimates indicate that individuals who experienced more severe Great Recession local shocks were insignificantly (0.073 percentage points) more likely to have moved after 2007 relative to the overall migration rate of 16 percent, suggesting that out-migration was not a major adjustment margin."

[41] Note that the published version in Galle et al. (2023) discusses the issue of mobility in more detail in their section 7.4, but the paper continues to omit the mobility channel from the quantitative analysis.

Table A2: Estimates and Predictive Accuracy for Ad-Hoc Spillover Adjustments

Shock size:	large		medium		small	
Shock type:	iid region	iid industry	iid region	iid industry	iid region	iid industry
$\hat{x}_r^{\mathrm{C}}$ minus migr.-wgt. avg., $\beta^D = -\beta^O$	0.082	0.053	0.083	0.053	0.087	0.054
	(0.030)	(0.022)	(0.042)	(0.029)	(0.069)	(0.047)
	{0.197}	{0.419}	{0.256}	{0.471}	{0.470}	{0.630}
Migration-weighted avg., β^D	0.081	0.052	0.083	0.052	0.087	0.053
	(0.030)	(0.022)	(0.042)	(0.029)	(0.070)	(0.046)
β^O	-0.086	-0.043	-0.087	-0.043	-0.091	-0.045
	(0.084)	(0.054)	(0.117)	(0.074)	(0.195)	(0.123)
	{0.256}	{0.538}	{0.382}	{0.691}	{0.820}	{1.240}
Distance-weighted avg., β^D	0.083	0.049	0.085	0.049	0.087	0.050
	(0.030)	(0.021)	(0.042)	(0.028)	(0.070)	(0.045)
β^O	-0.236	-0.040	-0.234	-0.038	-0.258	-0.037
	(0.593)	(0.090)	(0.822)	(0.125)	(1.361)	(0.205)
	{0.381}	{0.573}	{0.549}	{0.728}	{1.198}	{1.273}
Market access weighted avg., β^D	0.081	0.047	0.082	0.048	0.086	0.048
	(0.029)	(0.021)	(0.041)	(0.027)	(0.068)	(0.044)
β^O	-0.092	-0.041	-0.096	-0.042	-0.106	-0.046
	(0.137)	(0.079)	(0.192)	(0.110)	(0.318)	(0.181)
	{0.317}	{0.613}	{0.473}	{0.803}	{1.061}	{1.424}

Notes: Each row corresponds to a method of predicting the effect of observed shocks on regional population while using an ad-hoc method to incorporate outside-option shocks; see the text for details. Each column defines the shock process, described in the text. Numbers without parentheses or braces are median coefficient estimates across 601 simulations. Associated "standard errors" (standard deviations across simulations) are in parentheses. The relative mean squared errors, computed as the ratio of the median of (34) to the median of (35), are in curly braces.

A.9 Ad-hoc Empirical Approaches to Spillovers

In Table A2, we consider ad-hoc approaches to incorporating information on shocks to other regions that appear in the prior literature. These regressions take the following form

$$\hat{N}_r = \alpha + \beta^D \hat{x}_r^{\mathrm{C}} + \beta^O \hat{x}_{-r}^O + \mathrm{error}_r.$$

Here $\hat{x}_r^{\mathrm{C}}$ is as in the conventional migration regression, with associated regression coefficient β^D. The outside-option measure $\hat{x}_{-r}^O \equiv \sum_{k \neq r} \frac{\varphi_{kr}}{\sum_{o \neq r} \varphi_{or}} \hat{x}_k^{\mathrm{C}}$ is a weighted average of $\hat{x}_k^{\mathrm{C}}$ across other regions, with weights based on some measure $\varphi_{kr} \geq 0$ and associated regression coefficient β^O. In all cases, the analysis omits information on worker transitions across industries in constructing $\hat{x}_{-r}^O$, leading to particularly poor performance when shocks incorporate industry variation.

The most common approach in the prior literature is to measure the outside option shock as a migration-weighted average of shocks to other regions in which $\varphi_{kr} = F_{kr}$.[42] We first impose the restriction that $\beta^D = -\beta^O$, so the regressor is $\hat{x}_r^{\mathrm{C}} - \hat{x}_{-r}^O$. The MSE in this case is slightly improved relative to the conventional regression but much larger than the model-consistent approaches in Table 3, Panel C. For pure region shocks performance is still poor due to the omission of regional variation in migration intensity M_r/N_r, while for pure industry shocks misspecification of both the local and outside option shock measures generates additional problems. The second approach in Table A2 allows β^D and β^O to differ, simply adding $\hat{x}_{-r}^O$ as a second regressor. This approach performs quite poorly, with similar MSE to the conventional regression for large shocks, and much larger MSEs as the shock size decreases, reflecting the difficulty of separately estimating the effects of direct and indirect shocks.

[42] See, for example, Albert et al. (2021), Choi et al. (2024), and Foschi et al. (2025).

The final two approaches considered in Table A2 control for weighted averages of shocks to other regions that do not use data on region-to-region mobility. The first weights shocks to other regions based on the inverse geographic distance d_{kr}, so $\varphi_{kr} = d_{kr}^{-1}$.[43] The second weights shocks to other regions similar to how standard measures of market access are constructed, with weights decreasing in distance and increasing in local population, $\varphi_{kr} = d_{kr}^{-1.52} N_k$.[44] The distance elasticity of migration of -1.52 in this expression comes from a bilateral gravity equation estimated using the RAIS data with log symmetrized migration flow as the outcome and controlling for origin and destination fixed effects. Both approaches perform very poorly, even worse than the conventional regression. This is not surprising for industry shocks, given that the approaches do not incorporate any information on industry frictions, but they even underperform the conventional regression with regional shocks.

B Derivations and Proofs

B.1 Taylor Series Approximation (29)

When $\frac{\theta}{\sigma}$ is small or mobility is low, the largest eigenvalue of $\frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right)$ is small, such that

$$\left(\mathbb{I} + \frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right)\right)^{-1} = \mathbb{I} + \sum_{t=1}^{\infty}\left(-\frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right)\right)^{t} \approx \mathbb{I} - \frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right) \tag{B1}$$

and $\Omega \approx \frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right)$, which in turn implies $\widehat{\mathbf{L}} = \Omega\widehat{\mathbf{D}} \approx \frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right)\widehat{\mathbf{D}}$. Plugging in $\widehat{\mathbf{D}} = \widehat{\mathbf{z}} + \boldsymbol{\zeta}_1$ and aggregating markets into regions yields (29).

B.2 Direct-Mobility Approximation (18)

When $\Gamma = \mathbb{I} + \Delta\Gamma$ and $\Pi = \mathbb{I} + \Delta\Pi$ with $\Delta\Gamma'\Delta\Pi \approx 0$, we have

$$\mathbb{I} - \Gamma'\Pi = \mathbb{I} - \left(\mathbb{I} + \Delta\Gamma'\right)\left(\mathbb{I} + \Delta\Pi\right) \approx -\Delta\Gamma' - \Delta\Pi = \left(\mathbb{I} - \Gamma'\right) + \left(\mathbb{I} - \Pi\right).$$

By the Taylor series approximation from Appendix B.1,

$$\Omega \approx \frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right) \approx \frac{\theta}{\sigma}\left(\left(\mathbb{I} - \Gamma'\right) + \left(\mathbb{I} - \Pi\right)\right)$$

and

$$\widehat{\mathbf{L}} = \Omega\widehat{\mathbf{D}} \approx \frac{\theta}{\sigma}\left(\left(\mathbb{I} - \Gamma'\right) + \left(\mathbb{I} - \Pi\right)\right)\widehat{\mathbf{D}}.$$

[43] See Autor et al. (2023) and Greenland et al. (2019) for examples. In a different context focusing on trade-related spillovers, Adão et al. (forthcoming) use a similar approach, but with an exponent of -5 as opposed to -1.

[44] See Autor et al. (2021) for an example.

For a particular element ri,

$$\begin{aligned}\hat{L}_{ri} &\approx \frac{\theta}{\sigma}\left(\hat{D}_{ri} - \sum_{o,p}\gamma_{op,ri}\hat{D}_{op} + \hat{D}_{ri} - \sum_{d,q}\pi_{ri,dq}\hat{D}_{dq}\right) \\ &= \frac{\theta}{\sigma}\left(\sum_{o,p}\gamma_{op,ri}(\hat{D}_{ri} - \hat{D}_{op}) + \sum_{d,q}\pi_{ri,dq}(\hat{D}_{ri} - \hat{D}_{dq})\right). \end{aligned} \tag{B2}$$

Plugging in $\hat{D}_{ri} = \hat{z}_{ri} + \zeta_{1,ri}$ yields (18).

B.3 Theorem 1

We have

$$\begin{aligned}\hat{N}_r &= \sum_i \frac{L_{ri}}{N_r}\hat{L}_{ri} \\ &\approx \frac{\theta}{\sigma}\sum_i \frac{L_{ri}}{N_r}\sum_k\sum_p\left(\frac{f_{kp,ri}}{L_{ri}} + \frac{f_{ri,kp}}{L^0_{ri}}\right)(\hat{D}_{ri} - \hat{D}_{kp}) \\ &\approx \frac{2\theta}{\sigma}\sum_i\sum_k\sum_p \frac{F_{kp,ri}}{N_r}(\hat{D}_{ri} - \hat{D}_{kp}),\end{aligned}$$

where the second line follows by the direct-mobility approximation (B2) and the third line uses $L^0_{ri} \approx L_{ri}$ which again follows from the direct-mobility approximation since $\mathbf{L} = \Pi'\mathbf{L}^0 = \mathbf{L}^0 - (\mathbb{I} - \Pi')\mathbf{L}^0 \approx \mathbf{L}^0$. In the sum in the last expression, the terms corresponding to $k = r$ add up to zero:

$$\sum_i\sum_p \frac{F_{rp,ri}}{N_r}(\hat{D}_{ri} - \hat{D}_{rp}) = \sum_i \frac{F_{r,ri} - F_{ri,r}}{N_r}\hat{D}_{ri} = 0,$$

since $F_{r,ri} = F_{ri,r}$ for any r and i by definition of average bilateral flows, where $F_{r,ri} = \sum_p F_{rp,ri}$ and $F_{ri,r} = \sum_q F_{ri,rq}$. Thus,

$$\begin{aligned}\hat{N}_r &\approx \frac{2\theta}{\sigma}\sum_{k\neq r}\sum_i\sum_p \frac{F_{kp,ri}}{N_r}(\hat{z}_{ri} - \hat{z}_{kp}) + \zeta_r \\ &= \frac{2\theta}{\sigma}\left(\sum_{k\neq r}\sum_i \frac{F_{k,ri}}{N_r}\hat{z}_{ri} - \sum_{k\neq r}\sum_p \frac{F_{kp,r}}{N_r}\hat{z}_{kp}\right) + \zeta_r \\ &= \frac{2\theta}{\sigma}\frac{M_r}{N_r}\left(\hat{x}^{\mathrm{M}}_r - \hat{x}^{\mathrm{M}}_{-r}\right) + \zeta_r.\end{aligned}$$

This proves Theorem 1.

B.4 Equation (21)

Since $\hat{x}^{\mathrm{M}}_r$ is a weighted average of $\hat{z}_{ri}$ within the region, with pure regional shocks $\hat{x}^{\mathrm{M}}_r = \hat{z}_r$. Moreover,

$$\hat{x}^{\mathrm{M}}_{-r} = \sum_p\sum_{o\neq r}\frac{F_{op,r}}{M_r}\hat{z}_o = \sum_{o\neq r}\left(\sum_p \frac{F_{op,r}}{M_r}\right)\hat{z}_o = \sum_{o\neq r}\frac{F_{or}}{M_r}\hat{z}_o. \tag{B3}$$

B.5 Theorems 2 and A1

Since Theorem 2 is a special case of Theorem A1, we prove the latter.

Given that regression (22) includes the constant and our interest is in the slope parameter, we assume without loss of generality that the shocks have zero expectation, $\mu = 0$. We start with the numerator of β in (24), which can be represented as

$$\begin{aligned}
\mathbb{E}\left[\sum_r N_r \hat{N}_r \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right)\right] &= \mathbb{E}\left[\sum_r N_r \hat{N}_r \hat{x}_r^{\mathrm{C}}\right] \\
&= \mathbb{E}\left[\frac{2\theta}{\sigma}\sum_r M_r \hat{x}_r^{\mathrm{C}}\left(\hat{x}_r^{\mathrm{M}} - \hat{x}_{-r}^{\mathrm{M}}\right)\right] \\
&= \frac{2\theta}{\sigma}\sum_r M_r v_r \left(\rho_r^{\mathrm{M}} - \rho_r\right). \qquad \text{(B4)}
\end{aligned}$$

Here the first line used the fact that $\sum_r N_r \hat{N}_r \bar{x} = \bar{x}\sum_r N_r \hat{N}_r = \bar{x} N \hat{N} = 0$ because total national population is not allowed to change in the model. The second line plugged in (19) and used as-good-as-random assignment of $\widehat{\mathbf{z}}$, and the last line used the definitions of v_r, ρ_r^{M}, and ρ_r.

Now turn to the denominator of β. Note that $N_r \approx N_r^0$ under the direct-mobility approximation, and thus

$$\begin{aligned}
\widetilde{F}_{rk} &= \sum_{i,p} \frac{L_{ki}^0 L_{rp} + L_{rp}^0 L_{ki}}{2N} = \frac{N_k^0 N_r + N_r^0 N_k}{2N} \approx \frac{N_r N_k}{N}, \\
\widetilde{M}_r &= \sum_{k \neq r} \widetilde{F}_{rk} \approx \frac{N_r(N - N_r)}{N} = N_r\left(1 - \frac{N_r}{N}\right),
\end{aligned}$$

and thus

$$\widetilde{x}_{-r}^{\mathrm{C}} = \sum_{k \neq r} (\widetilde{F}_{rk}/\widetilde{M}_r)\hat{x}_k^{\mathrm{C}} = \sum_{k \neq r} \frac{N_k/N}{1 - N_r/N}\hat{x}_k^{\mathrm{C}}.$$

Since $\bar{x}$ is a weighted average of the conventional regressor in r and other regions, specifically $\bar{x} = \frac{N_r}{N}\hat{x}_r^{\mathrm{C}} + \left(1 - \frac{N_r}{N}\right)\widetilde{x}_{-r}^{\mathrm{C}}$, we have

$$\hat{x}_r^{\mathrm{C}} - \bar{x} = \left(1 - \frac{N_r}{N}\right)\left(\hat{x}_r^{\mathrm{C}} - \widetilde{x}_{-r}^{\mathrm{C}}\right). \qquad \text{(B5)}$$

Thus,

$$\begin{aligned}
\mathbb{E}\left[\sum_r N_r \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right)^2\right] &= \mathbb{E}\left[\sum_r N_r \hat{x}_r^{\mathrm{C}}\left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right)\right] \\
&= \mathbb{E}\left[\sum_r N_r \left(1 - \frac{N_r}{N}\right)\hat{x}_r^{\mathrm{C}}\left(\hat{x}_r^{\mathrm{C}} - \widetilde{x}_{-r}^{\mathrm{C}}\right)\right] \\
&= \mathbb{E}\left[\sum_r \widetilde{M}_r \hat{x}_r^{\mathrm{C}}\left(\hat{x}_r^{\mathrm{C}} - \widetilde{x}_{-r}^{\mathrm{C}}\right)\right] \\
&= \sum_r \widetilde{M}_r v_r \left(1 - \widetilde{\rho}_r\right). \qquad \text{(B6)}
\end{aligned}$$

Here the first line used $\sum_r N_r \bar{x} \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right) = \bar{x} \sum_r N_r \left(\hat{x}_r^{\mathrm{C}} - \bar{x}\right) = 0$ by definition of $\bar{x}$. The second line used (B5). The third line used the definition of $\widetilde{M}_r$, and the final line used the definition of $\widetilde{\rho}_r$ as well as $\mathbb{E}\left[\hat{x}_r^{\mathrm{C}}\right] = \mu = 0$.

Combining (B4) and (B6) yields

$$\begin{aligned}\beta &= \frac{2\theta}{\sigma} \frac{\sum_r M_r v_r \left(\rho_r^{\mathrm{M}} - \rho_r\right)}{\sum_r \widetilde{M}_r v_r \left(1 - \widetilde{\rho}_r\right)} \\ &= \frac{2\theta}{\sigma} \frac{\sum_r M_r v_r / \sum_r N_r v_r}{\sum_r \widetilde{M}_r v_r / \sum_r N_r v_r} \cdot \frac{\sum_r M_r v_r \left(\rho_r^{\mathrm{M}} - \rho_r\right) / \sum_r M_r v_r}{\sum_r \widetilde{M}_r v_r \left(1 - \widetilde{\rho}_r\right) / \sum_r \widetilde{M}_r v_r},\end{aligned}$$

which simplifies to (A13).

Finally, when the conventional regression shocks are homoskedastic (that is, $v_r = v$ for all r), all expressions immediately simplify to those in Theorem 2. We also note that, because $N_r \approx N_r^0$ under the direct-mobility approximation, the estimand is the same, up to the small approximation error, whether the regression is estimated with weights N_r, as we assumed, or N_r^0, justifying the argument in footnote 16.

B.6 Theorem 2 with Purely Regional Shocks

We now prove equation (26) for purely regional and homoskedastic shocks and the expressions for ρ and $\widetilde{\rho}$ below it.

With purely regional shocks, $\hat{x}_r^{\mathrm{M}} = \hat{z}_r$, such that $\rho_r^{\mathrm{M}} = 1$ and thus $\rho^{\mathrm{M}} = 1$, which implies (26). Next, we simplify ρ:

$$\begin{aligned}\rho &= \frac{1}{M} \sum_r M_r \rho_r = \frac{1}{Mv} \sum_r M_r \mathrm{Cov}\left[\hat{z}_r, \hat{x}_{-r}^{\mathrm{M}}\right] \\ &= \frac{1}{Mv} \sum_r \sum_{k \neq r} F_{kr} \mathrm{Cov}\left[\hat{z}_r, \hat{z}_k\right] \\ &= \frac{1}{Mv} \sum_r \sum_{k \neq r} f_{kr} \mathrm{Cov}\left[\hat{z}_r, \hat{z}_k\right] \\ &= \sum_r \sum_{k \neq r} \frac{f_{kr}}{M} \mathrm{Corr}\left[\hat{z}_r, \hat{z}_k\right],\end{aligned}$$

where the second line uses (B3), the third line uses the symmetry of both F_{kr} and $\mathrm{Cov}\left[\hat{z}_r, \hat{z}_k\right]$, and the last line uses homoskedasticity. Representing the denominator of β as in Appendix B.5 similarly justifies the analogous expression for $\widetilde{\rho}$.

B.7 $\hat{\beta}$ Is Close to β in Large Samples

In this appendix we provide the regularity conditions under which both $\hat{\beta}$ and β converge to the same number in an asymptotic sequence of economies with a growing number R of domestic regions. Under those conditions $\hat{\beta}$ and β are expected to be similar to each other when the observed number of regions is large.

We consider a general sequence of data-generating processes, indexed by the number of regions R, which jointly determine all relevant economic variables in equilibrium. Note this does not require iid sampling of regions which may be inappropriate in the setting we consider. With pure industry

shocks, I should also grow along the sequence to satisfy the assumptions (Borusyak et al., 2022). We have:

Proposition B1. *Suppose* $\sum_r \frac{N_r}{N} \hat{N}_r \left(\hat{x}_r^C - \bar{x}\right) \xrightarrow{p} C$ *and* $\sum_r \frac{N_r}{N} \left(\hat{x}_r^C - \bar{x}\right)^2 \xrightarrow{p} V > 0$ *as* $R \to \infty$*, and both sequences are uniformly integrable. Then* $\hat{\beta} \xrightarrow{p} \frac{C}{V}$ *and* $\beta \to \frac{C}{V}$*.*

Proof. The first statement follows by the continuous mapping theorem. For the second statement, note that uniform integrability implies $\mathbb{E}\left[\sum_r \frac{N_r}{N} \hat{N}_r \left(\hat{x}_r^C - \bar{x}\right)\right] \to C$ and $\mathbb{E}\left[\sum_r \frac{N_r}{N} \left(\hat{x}_r^C - \bar{x}\right)^2\right] \to V$. Thus, $\beta \to \frac{C}{V}$, as required. □

B.8 Identification with Model-Consistent Approaches

In this appendix, we provide conditions for identification of θ/σ in the model-consistent approaches of Section 4.1. We first characterize identification by NLLS for general market-level shocks. We then provide low-level conditions with the two shock structures emphasized in the paper. With purely regional shocks, we show that both NLLS and the Taylor approximation identify θ/σ under a very weak condition: that there is some geographic mobility. With purely industry shocks, we focus on the Taylor approximation and show that identification requires regional and industry transitions to interact, in a sense made precise below. We leave the proofs to the end of the appendix.

Proposition B2. *Let* $\Sigma_z \equiv \operatorname{Var}[\widehat{\mathbf{z}}]$*, which may be singular. Under the model and the shock-exogeneity assumption in (16),* $\lambda_0 = \theta/\sigma$ *minimizes the expectation of the NLLS objective in (28). It is the unique minimizer if and only if* $S\left(\Omega(\lambda) - \Omega(\lambda_0)\right)\Sigma_z^{1/2} \neq 0$ *for every* $\lambda \geq 0$ *such that* $\lambda \neq \lambda_0$*.*

Intuitively, identification requires that, for every incorrect value of λ, some direction along which the observed shocks vary generates different regional population responses under λ and λ_0.

Now consider the two main special cases. The result is particularly simple for regional shocks. Subject to non-degenerate regional shock variation, geographic mobility is both necessary and sufficient for identification: if workers can move across regions, relative regional shocks reveal their responsiveness to relative economic conditions.

Proposition B3. *Suppose as-good-as-random shocks* $\widehat{\mathbf{z}}$ *are purely regional,*

$$\widehat{\mathbf{z}} = C_R \widehat{\mathbf{z}}^R,$$

where C_R *is the* $RI \times R$ *matrix with* $(C_R)_{ri,o} = \mathbf{1}[r = o]$*, and suppose* $\operatorname{Var}\left[\widehat{\mathbf{z}}^R\right]$ *is positive definite. Suppose also that, from every origin market with positive initial employment, workers have positive probability of remaining in their origin region. Then:*

1. *Under the model* $\hat{N}_r = S_r' \Omega(\lambda_0)\widehat{\mathbf{z}} + \zeta_r$*,* θ/σ *is uniquely identified by NLLS if and only if* $M > 0$*;*
2. *If population changes are generated by the Taylor approximation in (29),* θ/σ *is uniquely identified by the regression in (30) if and only if* $M > 0$*.*

For purely industry shocks, we focus on the Taylor approximation, for which we are able to derive a tractable low-level condition for identification:

Proposition B4. *Suppose as-good-as-random shocks are purely industry-level,*

$$\widehat{\mathbf{z}} = C_I \widehat{\mathbf{z}}^I,$$

where C_I is the $RI \times I$ matrix with $(C_I)_{ri,q} = \mathbf{1}[i = q]$, and suppose $\mathrm{Var}\left[\widehat{\mathbf{z}}^I\right]$ is positive definite. If population changes are generated by the Taylor approximation in (29), θ/σ is identified by the regression in (30) if and only if $K_{ri} \neq 0$ for at least one market ri, where

$$K_{ri} = \sum_{op} L^0_{op} \left(\pi_{op,ri} - \pi_{op,r}\pi_{op,i}\right), \tag{B7}$$

and $\pi_{op,r} = \sum_q \pi_{op,rq}$ and $\pi_{op,i} = \sum_d \pi_{op,di}$ denote the probabilities that workers originating in market op choose region r and industry i, respectively.

Proposition B4 shows that identification with industry shocks requires a stronger condition than with regional shocks. The K_{ri} measure the extent to which workers' destination-region and destination-industry choices are related, conditional on their origin markets, after aggregation across origins. In particular, $K_{ri} = 0$ if destination region and destination industry are conditionally independent for every origin market. It is also zero if there is no geographic mobility or if there is no industry switching.[45] Thus, both geographic and industry mobility are necessary, but not sufficient, for identification from purely industry shocks: the two dimensions of worker transitions must interact in a way that does not cancel across origin markets.

Proof of Proposition B2. The expected NLLS objective function can be written as

$$\begin{aligned}
&\mathbb{E}\left[\sum_r \frac{N_r}{N}\left(\hat{N}_r - S_r'\Omega\left(\lambda\right)\widehat{\mathbf{z}}\right)^2\right] \\
&\quad= \sum_r \frac{N_r}{N}\mathbb{E}\left[\left(S_r'\Omega\left(\lambda_0\right)\widehat{\mathbf{z}} - S_r'\Omega\left(\lambda\right)\widehat{\mathbf{z}} + \zeta_r\right)^2\right] \\
&\quad= \sum_r \frac{N_r}{N}\mathbb{E}\left[\left(S_r'\Omega\left(\lambda_0\right)\left(\widehat{\mathbf{z}} - \mu\iota\right) - S_r'\Omega\left(\lambda\right)\left(\widehat{\mathbf{z}} - \mu\iota\right) + \zeta_r\right)^2\right] \\
&\quad= \sum_r \frac{N_r}{N}\mathrm{Var}\left[S_r'\left(\Omega\left(\lambda_0\right) - \Omega\left(\lambda\right)\right)\widehat{\mathbf{z}}\right] + \sum_r \frac{N_r}{N}\mathbb{E}\left[\zeta_r^2\right],
\end{aligned} \tag{B8}$$

where the second line plugs in the data generating process for $\hat{N}_r$, third line uses $\Omega\left(\lambda\right)\iota = 0$ (see footnotes 10 and 21), and the last line follows since, under $\mathbb{E}\left[\widehat{\mathbf{z}} - \mu\iota \mid \boldsymbol{\zeta}\right] = 0$, the cross-terms vanish by (16) applying the law of iterated expectations:

$$\mathbb{E}\left[\zeta_r \cdot S_r'\left(\Omega\left(\frac{\theta}{\sigma}\right) - \Omega\left(\lambda\right)\right)\left(\widehat{\mathbf{z}} - \mu\iota\right)\right] = 0$$

The second term in (B8) does not depend on λ, while the first is nonnegative and equals zero at $\lambda = \lambda_0$. Moreover,

$$\mathrm{Var}\left[S_r'\left(\Omega(\lambda_0) - \Omega(\lambda)\right)\widehat{\mathbf{z}}\right] = S_r'\left(\Omega(\lambda_0) - \Omega(\lambda)\right)\Sigma_z\left(\Omega(\lambda_0) - \Omega(\lambda)\right)' S_r.$$

Since $N_r > 0$ for every region, the first term in (B8) is zero if and only if

$$S\left(\Omega(\lambda_0) - \Omega(\lambda)\right)\Sigma_z^{1/2} = 0.$$

The result follows.

[45] With no geographic mobility, $\pi_{op,ri} = \mathbf{1}[r = o]\,\pi_{op,i}$ and $\pi_{op,r} = \mathbf{1}[r = o]$, such that all terms in (B7) vanish. The case of no industry mobility is analogous.

Proof of Proposition B3. We first establish the NLLS result. Let $D_N \equiv \text{diag}(N_r)$, $D \equiv \text{diag}(L_{ri})$, and $D^0 \equiv \text{diag}(L^0_{op})$. Because $\text{Var}\left[\widehat{\mathbf{z}}^R\right]$ is positive definite, Proposition B2 implies that NLLS identifies λ_0 if and only if

$$S\left(\Omega(\lambda) - \Omega(\lambda_0)\right) C_R \neq 0$$

for every $\lambda \neq \lambda_0$. To characterize this condition, define $A \equiv \mathbb{I} - \Gamma'\Pi$ and $B \equiv D^{1/2} A D^{-1/2}$. Because the definitions of Π and Γ imply $\Gamma' = D^{-1}\Pi' D^0$,

$$DA = D - D\Gamma'\Pi = D - \Pi' D^0 \Pi, \tag{B9}$$

which is symmetric and positive semidefinite. Thus, we have that $B = D^{-1/2}(DA)D^{-1/2}$ is symmetric and positive semidefinite, too, with $A = D^{-1/2} B D^{1/2}$. Given $\Omega(\lambda) = \mathbb{I} - (\mathbb{I} + \lambda A)^{-1}$, standard results in linear algebra imply

$$\begin{aligned}
\Omega(\lambda) - \Omega(\lambda_0) &= (\lambda - \lambda_0)\left(\mathbb{I} + \lambda A\right)^{-1} A \left(\mathbb{I} + \lambda_0 A\right)^{-1} \\
&= (\lambda - \lambda_0)\, D^{-1/2}\left(\mathbb{I} + \lambda B\right)^{-1} B \left(\mathbb{I} + \lambda_0 B\right)^{-1} D^{1/2} \\
&\equiv (\lambda - \lambda_0)\, D^{-1/2} \widetilde{B} D^{1/2}.
\end{aligned}$$

Finally, since S aggregates market-level outcomes to regions using employment shares,

$$S = D_N^{-1} C_R' D, \tag{B10}$$

and thus we have

$$S\left(\Omega(\lambda) - \Omega(\lambda_0)\right) C_R = (\lambda - \lambda_0)\, D_N^{-1} \cdot \left(D^{1/2} C_R\right)' \widetilde{B} \left(D^{1/2} C_R\right).$$

The middle matrix $\widetilde{B}$ is positive semi-definite and has the same null space as B. To see this, consider the eigendecomposition $B = U\,\text{diag}\,(b)\,U'$ for orthogonal U and $b = \{b_{ri}\} \geq 0$. Then $(\mathbb{I} + \lambda B)^{-1} = U\,\text{diag}\,(1/(1 + \lambda b_{ri}))\,U'$ and

$$\widetilde{B} = U\,\text{diag}\left(\frac{b_{ri}}{(1 + \lambda b_{ri})\,(1 + \lambda_0 b_{ri})}\right) U'.$$

Thus, since D_N^{-1} is non-singular, for $\lambda \neq \lambda_0$, $S\left(\Omega(\lambda) - \Omega(\lambda_0)\right) C_R = 0$ if and only if $0 = B D^{1/2} C_R = D^{1/2} A C_R$. Since $D^{1/2}$ is non-singular, this is further equivalent to

$$AC_R = 0. \tag{B11}$$

We next relate (B11) to geographic mobility. For any market-level vector $\mathbf{a}$, (B9) implies

$$\begin{aligned}
\mathbf{a}' DA\mathbf{a} &= \sum_{ri} L_{ri} a_{ri}^2 - \sum_{op} L^0_{op} \left(\sum_{ri} \pi_{op,ri} a_{ri}\right)^2 \\
&= \sum_{op} L^0_{op} \left[\sum_{ri} \pi_{op,ri} a_{ri}^2 - \left(\sum_{ri} \pi_{op,ri} a_{ri}\right)^2\right] \\
&\equiv \sum_{op} L^0_{op} \text{Var}_{\pi_{op}}\left[a_{ri}\right].
\end{aligned} \tag{B12}$$

Applying this expression to $\mathbf{a} = C_R\mathbf{c}$, for any regional vector $\mathbf{c} = (c_r)$, gives

$$\mathbf{c}'C_R'DAC_R\mathbf{c} = \sum_{op} L_{op}^0 \mathrm{Var}_{\pi_{op}}[c_r].$$

If there is no cross-regional mobility, the destination region is fixed at the origin region for every origin market, so the right-hand side is zero for every $\mathbf{c}$. Since DA is positive semidefinite, this implies $DAC_R = 0$ and in turn $AC_R = 0$. Conversely, suppose $M > 0$. Then there is some origin market op with positive probability of choosing a destination in a region $r \neq o$. By assumption, workers from that origin market also have positive probability of remaining in region o. Choosing $c_o = 1$ and $c_d = 0$ for $d \neq o$ makes the corresponding variance in (B12) strictly positive. Thus $C_R'DAC_R \neq 0$, and hence $AC_R \neq 0$.

For the Taylor approximation, the explanatory variable in (30) under purely regional shocks is

$$\mathbf{t} = SAC_R\widehat{\mathbf{z}}^R.$$

Since $\mathrm{Var}\left[\widehat{\mathbf{z}}^R\right]$ is positive definite, this regressor contains identifying variation if and only if

$$SAC_R \neq 0.$$

Using (B10),

$$SAC_R = D_N^{-1}C_R'DAC_R.$$

Since DA is positive semidefinite and D_N^{-1} and D are positive definite, $C_R'DAC_R = 0$ if and only if $AC_R = 0$, as before. Thus the Taylor approximation identifies θ/σ if and only if $M > 0$.[46]

Proof of Proposition B4. Under purely industry shocks, the regressor in the Taylor approximation is

$$\mathbf{t} = SAC_I\widehat{\mathbf{z}}^I.$$

Because $\mathrm{Var}\left[\widehat{\mathbf{z}}^I\right]$ is positive definite, $\mathbf{t}$ contains identifying variation if and only if

$$SAC_I \neq 0.$$

Let $K = C_R'DAC_I$. Using (B10),

$$SAC_I = D_N^{-1}C_R'DAC_I = D_N^{-1}K.$$

Since D_N is nonsingular, this is nonzero if and only if $K \neq 0$, as required.[47]

It remains to show that the elements of K are indeed K_{ri} from (B7). From (B9),

$$K = C_R'DC_I - C_R'\Pi'D^0\Pi C_I.$$

The (r, i) element of the first term is

$$(C_R'DC_I)_{ri} = L_{ri} = \sum_{op} L_{op}^0 \pi_{op,ri}.$$

[46] We note that including an intercept in the Taylor regression does not affect this conclusion since it is easy to verify that $\sum_r N_r t_r = 0$.

[47] As in the regional case, $\sum_r N_r t_r = 0$, so including an intercept does not alter the identification condition.

The corresponding element of the second term is

$$(C_R'\Pi' D^0 \Pi C_I)_{ri} = \sum_{op} L^0_{op} \left(\sum_q \pi_{op,rq} \right) \left(\sum_d \pi_{op,di} \right) = \sum_{op} L^0_{op} \pi_{op,r} \pi_{op,i}.$$

Subtracting gives (B7).

B.9 Bilateral Flow Regressions: Equation (32)

Using $f_{op,ri} \equiv \pi_{op,ri} L^0_{op}$ for $ri \neq op$ and totally log-differentiating (9), we have

$$\hat{f}_{op,ri} = \hat{\pi}_{op,ri} = \theta \left(\hat{w}_{ri} - \sum_{dq} \pi_{op,dq} \hat{w}_{dq} \right). \tag{B13}$$

Under the direct-mobility approximation, $\widehat{\mathbf{L}} = \Omega \widehat{\mathbf{D}} \approx \frac{\theta}{\sigma} (\mathbb{I} - \Gamma' \Pi) \widehat{\mathbf{D}}$ is small relative to $\widehat{\mathbf{D}}$. Thus, (14) yields $\widehat{\mathbf{w}} \approx (1/\sigma) \widehat{\mathbf{D}}$, which in turn implies

$$\hat{f}_{op,ri} \approx \frac{\theta}{\sigma} \left(\hat{D}_{ri} - \sum_{dq} \pi_{op,dq} \hat{D}_{dq} \right) \approx \frac{\theta}{\sigma} \left(\hat{D}_{ri} - \hat{D}_{op} \right).$$

Aggregating this to the level of region pairs,

$$\begin{aligned} \hat{f}_{or} &= \sum_{p,i} \frac{f_{op,ri}}{f_{or}} \hat{f}_{op,ri} = \frac{\theta}{\sigma} \sum_{p,i} \frac{f_{op,ri}}{f_{or}} \left(\hat{D}_{ri} - \hat{D}_{op} \right) \\ &= \frac{\theta}{\sigma} \left(\sum_i \frac{f_{o,ri}}{f_{or}} \hat{D}_{ri} - \sum_p \frac{f_{op,r}}{f_{or}} \hat{D}_{op} \right). \end{aligned}$$

Separating observed and unobserved labor demand shocks yields (32) with $\beta_1 = \theta/\sigma$.

We finally explain how to aggregate fitted values from the bilateral flow regression into a prediction of how regional population responded to the set of $\widehat{\mathbf{z}}$ shocks. The net change in the number of residents of r relative to the no-shock counterfactual equals the inflow minus the outflow:

$$N_r \hat{N}_r = \sum_{o \neq r} f_{or} \hat{f}_{or} - \sum_{d \neq r} f_{rd} \hat{f}_{rd}.$$

Thus,

$$\hat{N}_r = \sum_{o \neq r} \gamma_{or} \hat{f}_{or} - \sum_{d \neq r} \frac{f_{rd}}{N_r} \hat{f}_{rd} = \sum_{o \neq r} \gamma_{or} \hat{f}_{or} - \frac{N^0_r}{N_r} \sum_{d \neq r} \pi_{rd} \hat{f}_{rd}. \tag{B14}$$

While (B14) is written for the flow and population changes due to both observed and unobserved shocks, the analogous expression holds for the contribution of observed shocks only. Using the fitted values from the bilateral flow regression on the right-hand side (including the intercept), we get a prediction for the net population responses due to shocks, which we use, for instance, when evaluating the performance of bilateral regressions.

B.10 Proofs for Appendix A.7

Regularity conditions. We impose two conditions on near ties. First, for every pair of distinct markets ri, dq, the limits

$$\lim_{g\downarrow 0} \frac{\Pr(A = ri, B = dq, G < g)}{g} \quad \text{and} \quad \lim_{g\downarrow 0} \frac{\Pr(A = dq, B = ri, G < g)}{g}$$

exist, are finite, and coincide. This condition ensures that each pairwise indifference margin has a well-defined first-order density, independent of the side from which the margin is approached. Second, if $V_{(1)} \geq V_{(2)} \geq V_{(3)}$ denote the three highest utilities, we assume

$$\Pr\left(\theta(V_{(1)} - V_{(3)}) < g\right) = o(g) \qquad \text{as } g \downarrow 0.$$

This condition ensures that workers close to three or more markets, and hence near intersections of pairwise indifference margins, are negligible to first order. We also restrict attention to markets with $\Pr(A = ri) > 0$.

Proof of (A14). The first regularity condition implies

$$\frac{L_{ri}}{N}\kappa_{ri,dq} = \frac{L_{dq}}{N}\kappa_{dq,ri}, \qquad dq \neq ri. \tag{B15}$$

Indeed, both sides equal the common density of workers along the ri–dq indifference margin.

We first calculate the effect of changing one wage at a time. Increase $\log w_{dq}$ by $\delta > 0$, holding all other wages fixed. Because G is measured in taste-shock units, this increases the scaled utility of market dq by $\theta\delta$. For any $ri \neq dq$, workers leave ri for dq to first order when $A = ri$, $B = dq$, and $G < \theta\delta$. A worker for whom dq was initially third-best or lower can switch to dq only if the three highest utilities are within $\theta\delta$ in scaled units; by the second regularity condition, the probability of such events is $o(\delta)$. Therefore

$$\Pr(A^1 = ri) - \frac{L_{ri}}{N} = -\Pr(A = ri, B = dq, G < \theta\delta) + o(\delta) = -\theta\delta\,\frac{L_{ri}}{N}\kappa_{ri,dq} + o(\delta),$$

which implies

$$\frac{\partial \widehat{L}_{ri}}{\partial \widehat{w}_{dq}} = -\theta\kappa_{ri,dq}, \qquad dq \neq ri. \tag{B16}$$

The own-wage derivative follows from conservation of workers. Following an increase in $\log w_{ri}$, the increase in the probability of choosing ri equals, to first order, the sum of the probability losses in all other markets:

$$\Pr(A^1 = ri) - \frac{L_{ri}}{N} = \theta\delta \sum_{dq\neq ri} \frac{L_{dq}}{N}\kappa_{dq,ri} + o(\delta) = \theta\delta\,\frac{L_{ri}}{N} \sum_{dq\neq ri} \kappa_{ri,dq} + o(\delta),$$

where the last equality uses (B15). Hence

$$\frac{\partial \widehat{L}_{ri}}{\partial \widehat{w}_{ri}} = \theta \sum_{dq\neq ri} \kappa_{ri,dq}. \tag{B17}$$

Combining (B16)–(B17) for an arbitrary vector of small wage changes gives (A14).

Nesting the baseline model and proof of (A15). The result can be verified directly from the definition of $\kappa_{ri,dq}$ and the Gumbel density. A simpler argument follows by comparing the general labor-supply expression (A14) with the labor supply under Gumbel in (11). Equating them for any $\widehat{\mathbf{w}}$ implies (A15).

Proof of (A16). Along the sequence considered above, the no-shock choice structure is fixed as $\theta \to 0$, so the distribution of (A, B, G) and hence the coefficients $\kappa_{ri,dq}$ are fixed. Equation (A14) and labor demand in (8) imply

$$\widehat{L}_{ri} = \frac{\theta}{\sigma} \sum_{dq \neq ri} \kappa_{ri,dq} \left[(\widehat{D}_{ri} - \widehat{D}_{dq}) - (\widehat{L}_{ri} - \widehat{L}_{dq}) \right]. \tag{B18}$$

Because there are finitely many markets and the $\kappa_{ri,dq}$ are finite and fixed, (B18) implies $\widehat{\mathbf{L}} = O(\theta)$ for fixed $\widehat{\mathbf{D}}$ as $\theta \to 0$. The term involving employment differences on the right-hand side of (B18) is therefore $O(\theta^2)$. It follows that

$$\widehat{L}_{ri} = \frac{\theta}{\sigma} \sum_{dq \neq ri} \kappa_{ri,dq} (\widehat{D}_{ri} - \widehat{D}_{dq}) + O(\theta^2),$$

which implies (A16). The Gumbel special case follows by the argument above and the Taylor approximation in (29).

C Brazilian Data and Simulation Details

Our empirical work is based on *Relação Anual de Informações Sociais* (RAIS), an administrative employer-employee dataset maintained by the Brazilian Ministry of Labor. We consider data covering the period 1999–2010. RAIS constitutes a census of the Brazilian formal labor market and allows us to follow all formally employed workers longitudinally as they move across employers, industries, and geographic regions (De Negri et al., 2001; Saboia and Tolipan, 1985). Formal employment in RAIS includes workers with a signed work card (*carteira assinada*), which registers the worker's contract with the Ministry of Labor and entitles the worker to the protections and benefits established under the Brazilian labor legislation, as well as statutory public-sector employees, who are formally employed under a distinct legal regime and need not have a signed work card. RAIS excludes workers outside formal employment relationships, including the self-employed, informal workers without a formal labor contract, and individuals providing occasional services without an employment relationship. In addition, the dataset does not cover domestic workers or interns.

Restricting attention to formal employment is an important limitation in the Brazilian context, as informality rates exceeded 50 percent during our sample period (Dix-Carneiro and Kovak, 2019). Nevertheless, RAIS is the only nationwide longitudinal dataset with sufficient geographic and industrial detail to estimate worker mobility across local labor markets.

Geographic regions are defined using the Brazilian Statistical Agency's (IBGE) microregions, which group contiguous municipalities sharing similar geographic and economic characteristics into economically integrated local labor markets (IBGE, 2002). To ensure geographic consistency over time, we aggregate microregions following the procedure described in Dix-Carneiro and Kovak (2017), obtaining 486 time-consistent microregions spanning Brazil's 27 states. As documented in Dix-Carneiro and Kovak (2017), only 3.4 percent of individuals lived and worked in different microregions in the 2000 Population Census, implying that the workplace microregion provides an

excellent proxy for the worker's place of residence.

Industries are defined using the *Subsetor IBGE* classification reported in RAIS, which distinguishes 25 industries, including 12 manufacturing industries. Because a plant's industry classification may vary across years owing to reporting inconsistencies or reclassification, we assign each plant a single *Subsetor IBGE* code corresponding to the industry most frequently recorded for that plant over the full sample period, from 1986 (the first year of available data) through 2010 (the last). This procedure prevents spurious changes in industry classification from being counted as worker movements across industries. Combining the 486 microregions with the 25 industries yields 12,150 distinct region–industry markets.[48] When calculating observed Brazilian tariff changes used in our quantitative exercises, we follow Dix-Carneiro and Kovak (2017), with the exception that we aggregate tariffs to the *Subsetor IBGE* level (using simple averages).

Our sample definition and cleaning procedures closely follow Dix-Carneiro and Kovak (2017). Our sample consists of workers aged 18 to 64 with positive December earnings. RAIS reports both December earnings and average monthly earnings over the calendar year. We use December earnings to avoid distortions arising from seasonal fluctuations in employment and within-year inflation. We exclude observations with missing or inconsistent information, namely workers classified in "other/ignored" industries, individuals with inconsistent education records across simultaneous jobs, workers holding jobs in geographically implausible locations, and workers with clearly incorrect identifiers.[49] For workers holding multiple jobs in December, we retain the job with the highest December earnings and use it to assign the worker's industry and region.

C.1 Construction of Transition Matrices and Employment Stocks

The RAIS panel allows us to observe workers transitioning across all region–industry markets. Given the large number of markets (12,150), annual bilateral worker flows are often sparse. To reduce sampling noise, we construct transition matrices using five-year worker transitions and average observed flows across all overlapping five-year windows in the sample.

Specifically, let $f_{op,ri}^{t,t+5}$ denote the number of workers observed in market (o,p) in year t and in market (r,i) in year $t+5$, where $t = 1999, \ldots, 2005$. We rescale the flows in each year to match the national total from 1999, such that each 5-year interval contributes equally. We then define the average five-year worker flow as

$$f_{op,ri} = \frac{1}{7} \sum_{t=1999}^{2005} f_{op,ri}^{t,t+5}.$$

These average worker flows are also used to measure initial and final employment levels in each market, L_{op}^0 and L_{ri}. Combining these employment stocks with the worker flows $f_{op,ri}$, we obtain the out-migration and in-migration share matrices Π and Γ, as well as local industry composition.

Because employment levels are constructed from transition flows, we restrict attention to workers observed in both years t and $t+5$. As a result, workers who enter or exit formal employment between years t and $t+5$, as well as those who are otherwise not observed in both years, do not contribute to the measured employment stocks or transition matrices. This ensures that employment stocks are internally consistent with the observed transition flows.

[48] Starting in 1994, RAIS provides detailed five-digit industry codes that are approximately comparable to the International Standard Industrial Classification of All Economic Activities (ISIC). Exploiting this level of granularity, however, would be computationally infeasible, as it would require constructing and manipulating extremely large mobility matrices.

[49] RAIS identifies individuals using the Programa de Integração Social (PIS) number, which consists of 11 digits. We exclude individuals whose PIS identifiers consist entirely of zeros or contain fewer than 11 digits.

C.2 Calibration of Employment and Wage Noise Processes

The model requires estimates of the stochastic unobserved components governing employment and wage evolution across region–industry markets. We estimate these processes directly from the data using the same overlapping five-year windows.

C.2.1 Employment Process

For each five-year window starting in $t = 1999, \ldots, 2005$, we compute the five-year log employment growth in region r and industry i as the symmetric change

$$\Delta L_{ri}^{t} = \frac{L_{ri}^{t+5} - L_{ri}^{t}}{(L_{ri}^{t+5} + L_{ri}^{t})/2},$$

where $L_{op}^{t} = \sum_{ri} f_{op,ri}^{t,t+5}$ and $L_{ri}^{t+5} = \sum_{op} f_{op,ri}^{t,t+5}$. [50]

For each five-year window $t = 1999, \ldots, 2005$ separately, we then estimate

$$\Delta L_{ri}^{t} = \nu_{r}^{L,t} + \mu_{i}^{L,t} + \varepsilon_{ri}^{L,t},$$

where $\nu_{r}^{L,t}$ denotes region fixed effects, $\mu_{i}^{L,t}$ denotes industry fixed effects, and $\varepsilon_{ri}^{L,t}$ is the residual component. All regressions are estimated using employment weights $(L_{ri}^{t} + L_{ri}^{t+5})/2$.

Next, for each five-year window, we compute the employment-weighted standard deviations of the estimated region effects, industry effects, and residuals. The calibration parameters are obtained by averaging these standard deviations across the seven overlapping five-year windows. This procedure yields

$$\sigma_{\nu}^{L} = 0.050, \quad \sigma_{\mu}^{L} = 0.082, \quad \sigma_{\varepsilon}^{L} = 0.154$$

for the standard deviation of the unobserved pure region, pure industry, and idiosyncratic region-industry components, respectively.

C.2.2 Wage Process

To isolate region–industry wage components from worker composition effects, we first estimate separate Mincer regressions for each year $t = 1999, \ldots, 2010$:

$$\log w_{\omega}^{t} = \xi_{ri}^{t} + \beta^{t} Female_{\omega} + \sum_{k=2}^{5} \gamma_{k}^{t} Age_{\omega}^{k} + \sum_{m=2}^{9} \delta_{m}^{t} Educ_{\omega}^{m} + \eta_{\omega}^{t}, \tag{C1}$$

where w_{ω}^{t} denotes December earnings for worker ω employed in region r and industry i, $Female_{\omega}$ is a gender indicator, Age_{ω}^{k} denotes age-group indicators, $Educ_{\omega}^{m}$ denotes education-group indicators, and η_{ω}^{t} is a residual.[51] The estimated region–industry fixed effects ξ_{ri}^{t} capture composition-adjusted log wages.

[50] Note that, for consistency with Appendix C.1, we measure employment stocks by aggregating the flows of workers observed in both periods t and $t + 5$. We have verified that our conclusions are robust to measuring employment stocks directly from cross-sectional information.

[51] Since this analysis is year by year, rather than by 5-year window, we include all workers, whether or not they are observed 5 years before and after. The five age categories are: 18–24, 25–29, 30–39, 40–49, and 50–64. The nine education categories are: illiterate, 4th grade incomplete, 4th grade complete, 8th grade incomplete, 8th grade complete, high school incomplete, high school complete, college incomplete, and college complete or higher.

We then compute five-year changes in the region–industry log-wage components,

$$\Delta w_{ri}^{t} = \xi_{ri}^{t+5} - \xi_{ri}^{t},$$

and, for each five-year window, estimate

$$\Delta w_{ri}^{t} = \nu_{r}^{w,t} + \mu_{i}^{w,t} + \varepsilon_{ri}^{w,t},$$

using initial employment L_{ri}^{t} as weights.

For each five-year window, we compute the employment-weighted standard deviations of the estimated region effects, industry effects, and residuals. The calibration parameters are obtained by averaging these standard deviations across the seven overlapping five-year windows. This procedure yields

$$\sigma_{\nu}^{w} = 0.070, \quad \sigma_{\mu}^{w} = 0.048, \quad \sigma_{\varepsilon}^{w} = 0.094$$

for the standard deviation of the pure region, pure industry, and idiosyncratic region-industry components, respectively.

C.3 Wage Changes in Table 3

Row 3 of Table 3, Panel A, reports the median standard deviation (across simulations) of the regional wage effects of the $\widehat{\mathbf{z}}$ shocks. In each simulation, we compute these wage effects, denoted $\hat{w}_{r}^{\text{due to }\widehat{\mathbf{z}}}$, as follows. First, we compute market-level wage changes due to the $\widehat{\mathbf{z}}$ shocks using (14) (but excluding unobserved shocks), as

$$\widehat{\mathbf{w}}^{\text{due to }\widehat{\mathbf{z}}} = \frac{1}{\sigma}\left(\mathbb{I} - \Omega\right)\widehat{\mathbf{z}}.$$

Then we aggregate them from the level of markets to the regional level using employment weights,

$$\hat{w}_{r}^{\text{due to }\widehat{\mathbf{z}}} = S_{r}'\widehat{\mathbf{w}}^{\text{due to }\widehat{\mathbf{z}}}. \tag{C2}$$

We note that this object does not represent the change in the regional average wage $\bar{w}_r$ due to the $\widehat{\mathbf{z}}$ shocks. Suppressing "due to $\widehat{\mathbf{z}}$" for brevity, the latter satisfies, at the first order,

$$\hat{\bar{w}}_{r} \equiv \log\left(\frac{\sum_{i} w_{ri}^{1} L_{ri}^{1}}{\sum_{i} w_{ri} L_{ri}}\right) - \hat{N}_{r} = \sum_{i} \frac{w_{ri} L_{ri}}{\sum_{p} w_{rp} L_{rp}} \hat{w}_{ri} + \sum_{i} \left(\frac{w_{ri} L_{ri}}{\sum_{p} w_{rp} L_{rp}} - \frac{L_{ri}}{N_{r}}\right) \hat{L}_{ri}. \tag{C3}$$

This differs from (C2) in two ways. First, aggregation across industries involves payroll weights (i.e., proportional to $w_{ri}L_{ri}$) rather than employment weights. Second, the regional average wage can change purely from reallocating workers across industries with different wages within r, without changes in any w_{ri}, as captured by the second term in (C3). We use (C2) for its simplicity, particularly because it is implementable without wage data.

Row 4 of the same panel estimates a regression of wage changes due to both observed and unobserved shocks, $\hat{w}_r$, on $\hat{x}_{r}^{\text{C}}$. To parallel (33), we compute

$$\hat{w}_{r} = \hat{w}_{r}^{\text{due to }\widehat{\mathbf{z}}} + \zeta_{r}^{w},$$

where $\zeta_{r}^{w} = \sum_{i} \frac{L_{ri}}{N_r} \zeta_{ri}^{w}$ and

$$\zeta_{ri}^{w} = \epsilon_{r}^{w} + \epsilon_{i}^{w} + \epsilon_{ri}^{w}.$$

Here $\epsilon_{r}^{w}, \epsilon_{i}^{w}, \epsilon_{ri}^{w}$ are respectively mean-zero iid Gaussian regional, industry, and idiosyncratic com-

ponents. To ensure realistic simulations, their standard deviations are calibrated similarly to the noise components of employment changes; Appendix C.2.2 details the procedure.

C.4 Implementing Simulations Based on Bilateral Flow Regressions

To implement the bilateral flow regression (32) in our simulations, we need to obtain the outcome variable $\hat{f}_{or}$, the change in flows due to both observed and unobserved shocks, for each of the simulation draws. According to equation (B13), we can obtain $\hat{f}_{or}$ as a function of market-level wage changes:

$$\hat{f}_{or} = \sum_{p,i} \frac{f_{op,ri}}{f_{o,r}} \hat{f}_{op,ri} = \theta \sum_{p,i} \frac{f_{op,ri}}{f_{o,r}} \left(\hat{w}_{ri} - \sum_{dq} \pi_{op,dq} \hat{w}_{dq} \right). \tag{C4}$$

It would be natural to simulate wage changes using equation (14), while also including labor supply shocks (that are suppressed in our theory without loss of generality, see footnote 8, but allowed in our simulation design). The difficulty is that our simulation procedure calibrates the data-generating process for the error term in the market-level employment equation, $\boldsymbol{\zeta}$ in $\widehat{\mathbf{L}} = \Omega(\theta/\sigma)\widehat{\mathbf{z}} + \boldsymbol{\zeta}$, directly; see equation (33). Thus, we do not separately calibrate the process for $\zeta_{1,ri}$ in (15) or for the labor supply shocks. In the presence of both unobserved demand and supply shocks, $\widehat{\mathbf{w}}$ is not even uniquely determined by $\widehat{\mathbf{L}}$ (see equation (A12)).

To circumvent these problems, we adopt a simple solution by assuming that $\boldsymbol{\zeta}$ is fully generated by latent demand shocks, i.e. $\nu_{ri} = 0$. This assumption allows us to recover the implied wage changes by solving equation (12):

$$\widehat{\mathbf{L}} = \theta \left(\mathbb{I} - \Gamma' \Pi \right) \widehat{\mathbf{w}}.$$

The main complication is that the matrix $\mathbb{I} - \Gamma'\Pi$ is by construction singular. As a result, this linear system generally does not admit a solution $\widehat{\mathbf{w}}$ that exactly reproduces the simulated employment changes. We therefore recover wage changes by solving the least-squares problem

$$\widehat{\mathbf{w}} = \arg\min_{\widehat{\boldsymbol{\omega}}} \left\| \theta \left(\mathbb{I} - \Gamma' \Pi \right) \widehat{\boldsymbol{\omega}} - \widehat{\mathbf{L}} \right\|^2.$$

Plugging the solution into (C4) yields $\hat{f}_{or}$ for each simulation draw.

D Results with German Data

D.1 Data Sources and Cleaning

We replicate our core analysis with the confidential German employer-employee data provided by the Research Data Center at the Institute for Employment Research (IAB) for 1999–2008. Antoni et al. (2019) provide detailed data documentation. Our main dataset is the Sample of Integrated Labor Market Biographies (SIAB) which is a 2% panel of German workers, linking them over time and to the establishments where they work. We also use the Establishment History Panel (BHP), which provides information on those establishments. The main advantage of SIAB relative to RAIS is that informal work is relatively unimportant in the context of Germany; the main disadvantage is that we only observe 2% of the workers.[52]

While the data originally define regions as counties (*Kreise*), we aggregate them to $R = 141$ commuting zones (CZs) using the definition from Kosfeld and Werner (2012). This approach to

[52] A 50% sample also exists but data access to it is severely restricted.

defining the geographic boundary of local labor markets has become standard in the recent literature (Dorn and Levell, 2024). Most of the CZs are uniquely assigned to a German federal state (*Bundesland*); for the minority that cross a state boundary, we assign them to the state with the largest total employment in the CZ. Industries are originally defined by three-digit-level German classification of economic activities.[53] We aggregate them, mainly for computational reasons, into 23 two-digit manufacturing industries and 17 one-digit non-manufacturing sectors, with the total of $I = 40$.

The data cleaning procedure mostly follows that for our analysis of Brazil. In each year, we consider workers aged 18 to 64 who are employed as of June 30. For workers with multiple employment spells as of that day, we choose the one with the highest daily wage. We assign workers to the CZ and industry of their employer (rather than residential CZ). We drop workers for whom there is no valid establishment identifier, or where the region and industry are not specified for the establishment.

We construct transition flows $f_{op,ri}$ by aggregating the observed worker transitions over the five overlapping 5-year windows within our data. For disclosure reasons, we diverge slightly from the Brazil analysis by not normalizing flows in each 5-year window to the same national total.[54] The rest of the procedure for constructing transition matrices and employment stocks is identical to what Appendix C describes.

To evaluate the performance of the conventional migration regression, we consider two real industry shocks from the literature, to parallel the Brazilian trade liberalization shock in Table 2. Specifically, we follow Dauth et al. (2014) in considering the growth of German import penetration from China and Eastern European countries, measured in €1,000 per German worker (with zero values in all non-manufacturing industries). Dauth et al. (2014), like Autor et al. (2013), instrument German imports from China and Eastern Europe with imports into eight other destination countries (Australia, Canada, Japan, Norway, New Zealand, Sweden, Singapore, and the United Kingdom) from the same origins. We use the instrument version of the shock to parallel the reduced-form OLS regressions from the paper, as the instrument is considered exogenous. Borusyak et al. (2022) point out that, even if the shocks to manufacturing industries are exogenous, the resulting shift-share instrument is not because it is mechanically larger in manufacturing-heavy regions. As suggested in footnote 21, we adopt the simple recentering solution proposed in Borusyak and Hull (2023): to demean the shocks within manufacturing, while keeping them at zero in non-manufacturing. We construct the shocks following the replication package of Dauth et al. (2021b), with two exceptions. First, Dauth et al. (2021b) define the shocks by three-digit industry. We therefore aggregate both imports in the numerator and initial employment in the denominator to our 40 industries. Second, the two Dauth et al. papers use different time periods to define the shock. We use 2000–2010 as in Dauth et al. (2021a), such that the initial employment is measured in 2000, within our sample period.

We calibrate the variance components for employment and wages similarly to Appendix C.2. In the Mincer regression (C1), we control for a female indicator, all indicators of age in years, and professional training as the best available measure of education.[55] The rest of the procedure, e.g.

[53] Our choice of years, 1999–2008, is to ensure a consistent industry classification, which changes substantially in 2009.

[54] The disclosure procedure requires the number of workers underlying each statistic to be above a certain threshold. Aggregating worker flows without any renormalization ensures that the flows are integer-valued and equal to the underlying number of workers.

[55] Professional training is recorded in the variable `ausbildung_imp` (with `imp` for imputed), which has the following categories: "Lower secondary school qual. without vocational qual. [qualification]," "Lower secondary school qual. with vocational qual.," "Entrance qual. for FH or University without vocational qual.," "Entrance qual. for FH or University with vocational qual.," "University of Applied Sciences (FH)," and "University," as well as a missing value.

Table D1: Average Mobility Rates, German Data

Across	Mobility (%)
Regions	13.5
States	7.4
Industries	22.3
Regions and Industries	7.6

Notes: Same statistics as in Table 1, except using the German SIAB data described in Appendix D.1.

regression specifications, is identical to Appendix C.2. The calibration of the employment process yields:

$$\sigma_\nu^L = 0.022, \quad \sigma_\mu^L = 0.077, \quad \sigma_\varepsilon^L = 0.128.$$

The calibration for wages yields:

$$\sigma_\nu^w = 0.016, \quad \sigma_\mu^w = 0.034, \quad \sigma_\varepsilon^w = 0.065.$$

We note that these standard deviations are lower than in the Brazilian data.

D.2 Descriptive Statistics

Table D1 replicates Table 1 with German data. It shows that mobility rates across regions and industries are remarkably similar between the two countries (although they are not exactly comparable as the region and industry definitions are different). One difference is that the fraction of migration that crosses state lines is higher in Germany.

We also compute that the population-weighted average Herfindahl index of migration destinations is 0.086 (s.d. 0.049), which is equivalent to equally-weighted connections to 11.6 other regions out of 140 possibilities, with the overall effective number of regions (i.e., inverse Herfindahl of regional employment) of 62.7. The average Herfindahl index is similar to that of Brazil, although again the numbers are not directly comparable as Germany has fewer regions (but also much lower size inequality across regions, making the effective number of regions larger).

D.3 Results

Table D2 reports the median coefficient, relative to that with iid regional shocks, and R^2 of the conventional migration regression, replicating the analysis of Table 2 with German data. The key patterns are similar: attenuation is substantial with state-level or industry shocks, with the attenuation factor below 0.5 with iid industry shocks. Attenuation is somewhat weaker with the China import penetration shock (at 0.764) but strong with the Eastern European import penetration (at 0.553). The median R^2 of the regression is slightly larger than in Brazilian data but the conventional regression is still highly imperfect despite the lack of any unobserved shocks in this simulation. For instance, it is only 0.603 with iid industry shocks, only marginally higher than 0.551 in Brazilian data.

To evaluate the performance of alternative methods, Panel A of Table D3 calibrates the shock magnitudes in the same way as with Brazilian data, targeting the standard deviation of population effects to have the median across simulations of 0.035, 0.025, and 0.015 for "large", "medium", and "small" shocks, respectively. This yields calibrated variances of pure regional shocks very similar to those in Brazilian data, with similar wage effects, too. The standard deviation of pure industry shocks required to generate the same population response is almost twice as large in Germany,

Table D2: Bias in the Conventional Migration Regression, German Data

Shocks	Relative β	R^2
iid across regions	1.000	0.925
iid across states	0.652	0.846
iid across industries	0.446	0.603
China import penetration	0.764	0.749
Eastern European import penetration	0.553	0.725

Notes: Same statistics as in Table 2, except using the German SIAB data. See text for definitions of shock processes.

consistent with a more detailed industry classification we use (40 codes compared to 25 in Brazil), suggesting that shift-share regional aggregates of industry shocks have less variation. However, given the same population responses (and structural parameters), both datasets yield similar wage responses.

The other panels of Table D3 as well as Table D4 report the performance of alternative methods, for comparison with the conventional regression in Panel B. The patterns are qualitatively, and often quantitatively, similar to those reported in Tables 3 and A2 for Brazil. NLLS still resolves bias and exhibits excellent predictive performance. The Taylor approximation performs very well. The direct-mobility approximation vastly outperforms the conventional regression, particularly with regional shocks, as do methods with reduced transition-data demands. Ad hoc approaches still exhibit very poor performance in many cases, due to misspecification and very imprecise parameter estimates. Given these strong general similarities to the results for Brazil, we briefly focus on a few differences.

Consistent with lower variability in unobserved shocks, the levels of standard errors and mean-squared errors are smaller for most methods and shock magnitudes in Germany compared to Brazil. As a result, even the conventional regression is statistically significant at the 5-percent level for all shock types except small industry shocks. The MSE of the conventional regression is also much lower in Germany with regional shocks, but still large with iid industry shocks.

The bilateral regression exhibits modestly distinct performance in Germany than Brazil. In Brazil, the bilateral regression has larger MSE than the direct-mobility approximation for all shock types, while in Germany the bilateral MSE is smaller with industry shocks. A possible explanation for this pattern in Germany is that with industry shocks, bilateral regressions can predict gross mobility changes of the same sign for movers from region A to region B and from B to A — e.g., if in both directions workers move from the chemical industry to electronics, and one of those industries got a better national shock than the other. This signal is lost by aggregation into the direct-mobility unilateral specification, while the pure regional and cell-level components of the error terms do not wash out.

Finally, Table D5 verifies the precision of the small-shock approximation, analogously to Table A1 for Brazil. The approximation continues to be very precise with pure regional shocks of all sizes. Its precision deteriorates relative to the Brazilian case with pure industry shocks, as the underlying shock variances are larger (see Table D3, Panel A).

Table D3: Estimates and Predictive Accuracy for Alternative Estimators, German Data

Shock size:	large		medium		small	
Shock type:	iid region	iid industry	iid region	iid industry	iid region	iid industry
Panel A: Shock magnitudes						
std. dev. of shocks	0.38	7.54	0.27	5.39	0.16	3.23
std. dev. of population effects ×100	3.50	3.50	2.50	2.50	1.50	1.50
std. dev. of wage effects ×100	8.49	16.72	6.06	11.94	3.64	7.17
Conventional wage regression	0.228	0.240	0.228	0.240	0.228	0.240
	(0.008)	(0.006)	(0.011)	(0.007)	(0.019)	(0.012)
Panel B: Conventional migration regression	0.091	0.040	0.092	0.040	0.092	0.040
	(0.015)	(0.012)	(0.021)	(0.015)	(0.034)	(0.022)
	{0.093}	{0.437}	{0.106}	{0.461}	{0.146}	{0.526}
Panel C: Model-consistent approaches						
NLLS	0.408	0.405	0.408	0.404	0.408	0.402
	(0.073)	(0.069)	(0.102)	(0.096)	(0.171)	(0.158)
	{0.013}	{0.009}	{0.025}	{0.018}	{0.068}	{0.051}
Taylor approx.	0.361	0.370	0.361	0.371	0.362	0.371
	(0.057)	(0.058)	(0.079)	(0.081)	(0.132)	(0.135)
	{0.013}	{0.029}	{0.025}	{0.036}	{0.069}	{0.066}
Direct mobility approx.	0.327	0.288	0.327	0.287	0.327	0.288
	(0.051)	(0.054)	(0.072)	(0.072)	(0.119)	(0.115)
	{0.016}	{0.190}	{0.028}	{0.209}	{0.072}	{0.257}
Bilateral regression	0.330	0.294	0.331	0.294	0.332	0.294
	(0.054)	(0.011)	(0.075)	(0.013)	(0.125)	(0.019)
	{0.017}	{0.138}	{0.030}	{0.139}	{0.077}	{0.140}
Panel D: Reduced transition-data demands						
Taylor, regional model	0.354	0.187	0.355	0.188	0.355	0.188
reg-to-reg	(0.055)	(0.050)	(0.078)	(0.062)	(0.129)	(0.091)
	{0.014}	{0.346}	{0.026}	{0.364}	{0.070}	{0.420}
Bilateral, regional shocks	0.330	0.176	0.331	0.177	0.332	0.176
	(0.054)	(0.050)	(0.075)	(0.063)	(0.125)	(0.095)
	{0.017}	{0.343}	{0.030}	{0.362}	{0.077}	{0.431}
Taylor imputation 1:	0.358	0.508	0.358	0.506	0.359	0.510
reg-to-reg and ind-to-ind	(0.056)	(0.107)	(0.078)	(0.142)	(0.131)	(0.227)
	{0.013}	{0.270}	{0.026}	{0.288}	{0.069}	{0.333}
Taylor imputation 2:	0.359	0.577	0.359	0.570	0.360	0.573
ind-to-ind and reg-to-reg-ind	(0.056)	(0.117)	(0.079)	(0.154)	(0.131)	(0.247)
	{0.013}	{0.156}	{0.026}	{0.172}	{0.070}	{0.210}

Notes: Same statistics as in Table 3, except using the German SIAB data described in Appendix D.1.

Table D4: Estimates and Predictive Accuracy for Ad-Hoc Spillover Adjustments, German Data

Shock size:	large		medium		small	
Shock type:	iid region	iid industry	iid region	iid industry	iid region	iid industry
$\hat{x}_r^c$ minus migr.-wgt. avg., $\beta^D = -\beta^O$	0.089	0.047	0.089	0.047	0.089	0.047
	(0.014)	(0.013)	(0.020)	(0.016)	(0.033)	(0.024)
	{0.038}	{0.370}	{0.049}	{0.394}	{0.090}	{0.453}
Migration-weighted avg., β^D	0.090	0.047	0.090	0.047	0.090	0.047
	(0.015)	(0.013)	(0.021)	(0.016)	(0.035)	(0.024)
β^O	-0.083	-0.041	-0.085	-0.041	-0.088	-0.039
	(0.047)	(0.027)	(0.065)	(0.036)	(0.108)	(0.058)
	{0.052}	{0.384}	{0.079}	{0.427}	{0.179}	{0.543}
Distance-weighted avg., β^D	0.090	0.044	0.091	0.045	0.091	0.045
	(0.015)	(0.013)	(0.021)	(0.015)	(0.035)	(0.023)
β^O	-0.232	-0.070	-0.222	-0.071	-0.225	-0.070
	(0.249)	(0.079)	(0.342)	(0.105)	(0.563)	(0.170)
	{0.092}	{0.415}	{0.129}	{0.453}	{0.270}	{0.616}
Market access weighted avg., β^D	0.089	0.045	0.089	0.045	0.089	0.046
	(0.015)	(0.012)	(0.021)	(0.015)	(0.034)	(0.023)
β^O	-0.114	-0.051	-0.116	-0.052	-0.123	-0.050
	(0.086)	(0.043)	(0.120)	(0.057)	(0.200)	(0.092)
	{0.067}	{0.401}	{0.099}	{0.442}	{0.217}	{0.577}

Notes: Same statistics as in Table A2, except using the German SIAB data described in Appendix D.1. The distance elasticity of migration flows used for the last specification ("Market access weighted avg.") is -1.569, estimated as described in Appendix A.9.

Table D5: Precision of the Small-Shock Approximation, German Data

Shock size:	large		medium		small	
Shock type:	iid region	iid industry	iid region	iid industry	iid region	iid industry
Small-shock approx.	{0.000}	{0.455}	{0.000}	{0.355}	{0.000}	{0.196}
NLLS	{0.000}	{0.453}	{0.000}	{0.344}	{0.000}	{0.194}
Taylor approx.	{0.001}	{0.483}	{0.001}	{0.392}	{0.001}	{0.229}
Direct-mob. approx.	{0.003}	{0.577}	{0.003}	{0.493}	{0.003}	{0.364}
Conventional migration reg.	{0.075}	{0.576}	{0.075}	{0.485}	{0.075}	{0.432}

Notes: Same statistics as in Table A1, except using the German SIAB data described in Appendix D.1.

Appendix References

Adão, Rodrigo (2025). "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil". *Working Paper*.

Adão, Rodrigo, Costas Arkolakis, and Federico Esposito (forthcoming). "General Equilibrium Effects in Space: Theory and Measurement". *American Economic Journal: Macroeconomics*.

Albert, Christoph, Paula Bustos, and Jacopo Ponticelli (2021). "The effects of climate change on labor and capital reallocation". *Working Paper*.

Antoni, Manfred, Alexandra Schmucker, Stefan Seth, and Philipp vom Berge (2019). "Sample of Integrated Labour Market Biographies (SIAB) 1975-2017". *FDZ-Datenreport*.

Autor, David, David Dorn, and Gordon H. Hanson (2013). "The China Syndrome: Local Labor Market Impacts of Import Competition in the United States". *American Economic Review* 103(6): 2121–2168.

Autor, David, David Dorn, and Gordon H. Hanson (2021). "On the Persistence of the China Shock". *Brookings Papers on Economic Activity* Fall 2021: 381–447.

Autor, David, David Dorn, and Gordon H. Hanson (2023). *Trading Places: Mobility Responses of Native and Foreign-Born Adults to the China Trade Shock*. Working Paper 30904. National Bureau of Economic Research.

Borusyak, Kirill and Peter Hull (2023). "Non-Random Exposure to Exogenous Shocks". *Econometrica* 91(6): 2155–2185.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022). "Quasi-Experimental Shift-Share Research Designs". *Review of Economic Studies* 89(1): 181–213.

Bound, John and Harry J. Holzer (2000). "Demand shifts, population adjustments, and labor market outcomes during the 1980s". *Journal of Labor Economics* 18(1): 20–54.

Cadena, Brian C. and Brian K. Kovak (2016). "Immigrants equilibrate local labor markets: Evidence from the great recession". *American Economic Journal: Applied Economics* 8(1): 257–290.

Choi, Jiwon, Ilyana Kuziemko, Ebonya Washington, and Gavin Wright (2024). "Local Economic and Political Effects of Trade Deals: Evidence from NAFTA". *American Economic Review* 114(6): 1540–1575.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2014). "The rise of the east and the far east: German labor markets and trade integration". *Journal of the European Economic Association* 12(6): 1643–1675.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2021a). "Adjusting to Globalization in Germany". *Journal of Labor Economics* 39(1): 263–302.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner (2021b). "The Adjustment of Labor Markets to Robots". *Journal of the European Economic Association* 19(6): 3104–3153.

De Negri, João Alberto, Paulo Furtado de Castro, Natalia Ribeiro de Souza, and Jorge Saba Arbache (2001). "Mercado Formal de Trabalho: Comparação entre os Microdados da RAIS e da PNAD". *Texto Para Discussão* (No. 840): 29.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008). "Global Rebalancing with Gravity: Measuring the Burden of Adjustment". *IMF Staff Papers* 55(3): 511–540.

Demidova, Svetlana, Konstantin Kucheryavyy, Takumi Naito, and Andrés Rodríguez-Clare (2024). "The small open economy in a generalized gravity model". *Journal of International Economics* 152: 103997.

Dix-Carneiro, Rafael and Brian K. Kovak (2017). "Trade Liberalization and Regional Dynamics". *American Economic Review* 107(10): 2908–2946.

Dix-Carneiro, Rafael and Brian K. Kovak (2019). "Margins of Labor Market Adjustment to Trade". *Journal of International Economics* 117: 125–142.

Dorn, David and Peter Levell (2024). "Trade and inequality in Europe and the US". *Oxford Open Economics* 3(Supplement 1): i1042–i1068.

Foschi, Andrea, Christopher L. House, Christian Proebsting, and Linda L. Tesar (2025). "Should I Stay or Should I Go? The Response of Labor Migration to Economic Shocks". *Brookings Papers on Economic Activity* Spring 2025: 67–114.

Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi (2017). *Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade*. Working Paper 23737. National Bureau of Economic Research.

Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi (2023). "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade". *The Review of Economic Studies* 90(1): 331–375.

Greenland, Andrew, John Lopresti, and Peter McHenry (2019). "Import Competition and Internal Migration". *The Review of Economics and Statistics* 101(1): 44–59.

IBGE (2002). *Censo Demográfico 2000: Documentação dos Microdados da Amostra*. Instituto Brasileiro de Geografia e Estatística.

Kosfeld, Reinhold and Alexander Werner (2012). "Deutsche arbeitsmarktregionen–neuabgrenzung nach den kreisgebietsreformen 2007–2011". *Raumforschung und Raumordnung* 70(1): 49–64.

Redding, Stephen J. (2016). "Goods trade, factor mobility and welfare". *Journal of International Economics* 101: 148–167.

Saboia, Joao L.M. and Ricardo M.L. Tolipan (1985). *A Relação Anual de Informações Sociais (RAIS) e o mercado formal de trabalho no Brasil: uma nota*.

Yagan, Danny (2019). "Employment Hysteresis from the Great Recession". *Journal of Political Economy* 127(5): 2505–2558.