Understanding Migration Responses to Local Shocks*

Kirill Borusyak†, Rafael Dix-Carneiro‡, and Brian K. Kovak§

June 2023

Abstract

We examine how to interpret estimates from a commonly used migration regression relating changes in local population to exogenous local labor demand shocks. Using a simple model of local labor markets with mobility costs, we find that most conclusions drawn from migration regression estimates are likely to be substantially misleading. Intuitively, the conventional migration regression is misspecified due to the bilateral nature of location choices. Workers choose where to live based not only on the shock to their current location, but also on the shocks to potential alternative locations, which are omitted from the regression. Analytical results and simulations based on Brazilian data show that conventional migration regression estimates are inaccurate for the local population effects of either shocks to individual locations or all observed shocks taken together and often substantially understate the amount of worker reallocation driven by observed shocks. These problems are particularly acute when workers face industry switching costs in addition to geographic mobility costs. Simple alternative approaches leveraging the model's structure exhibit far better performance.

*We thank David Autor, Dave Choi, David Dorn, Gordon Hanson, Felix Koenig, Camille Landais, and participants at numerous seminars and conferences for helpful comments. 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). All errors are our own.

†Department of Economics, University College London (k.borusyak@ucl.ac.uk)

‡Department of Economics, Duke University (rafael.dix.carneiro@duke.edu)

§Heinz College, Carnegie Mellon University (bkovak@cmu.edu)

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, 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.[1]

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 local population growth to observed labor demand shocks, as in the following regression specification:

$$\hat{L}_\ell = \alpha + \beta \hat{z}_\ell + \varepsilon_\ell, \tag{1}$$

where $\hat{L}_\ell$ is the proportional change in population in location ℓ, $\hat{z}_\ell$ is an observed local labor demand shock, and ε_ℓ is the error term. In many cases, researchers find substantial effects of the shocks on local economic outcomes such as wages and employment, but find estimates of β that are small and/or statistically indistinguishable from zero.[2] These results seem to imply a puzzle: despite substantial effects of labor demand shocks on workers' outcomes, they do not appear to adjust by relocating.[3] Prior work has responded to these findings by inferring that migration costs are high,

[1]For example, the U.S. Trade Adjustment Assistance program includes funding for distant job search activities and an allowance for relocation costs (Hyman, 2018).

[2]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 Bound and Holzer (2000) (for those with a high-school degree or less); Topalova (2010); Autor et al. (2013); Dauth et al. (2014); Mian and Sufi (2014); Cadena and Kovak (2016); Dix-Carneiro and Kovak (2017); Dix-Carneiro and Kovak (2019); Yagan (2019); Autor et al. (2021) (except for those age 25-39); Choi et al. (2021); and Faber et al. (2021) (for Chinese import competition), among others. Additional papers estimating a version of (1) and finding substantial effects in at least some specifications include Black et al. (2005); Wozniak (2010); McCaig (2011); Bustos et al. (2016); Hakobyan and McLaren (2016); Bartik et al. (2019); Foote et al. (2019); Greenland et al. (2019); Albert and Monràs (2020); Boustan et al. (2020); Monràs (2020); Notowidigdo (2020); and Albert et al. (2021), among others.

[3]For example, Choi et al. (2021) echo a common sentiment in this literature: "This finding deepens the puzzle raised in recent papers that find no or limited migration response to large, negative local employment shocks."

that workers are generally unresponsive to changing local economic conditions, or that interregional migration can be ignored in quantitative models.

In this paper, we seek to resolve this puzzle by showing that most common interpretations of β in (1) can yield misleading conclusions. In particular, the estimate of β can be close to zero even when the observed shocks led to substantial spatial reallocation and workers are highly responsive to local economic conditions. These problems arise even when the observed shock is "exogenous," in the sense of being as good as random with respect to unobserved local labor demand and supply shocks.[4]

The core intuition emerging from our analysis is that the conventional migration regression is misspecified due to the bilateral nature of location choices. When workers consider economic conditions in deciding whether and where to move, their decisions depend upon both the shock to their current location and the shocks facing potential alternative locations. Since workers face different migration costs across origin-destination pairs, the most important potential alternatives differ across locations. By omitting the shocks facing these relevant alternative locations, the conventional migration regression in (1) is misspecified.

This misspecification leads to two problems. First, when the omitted shocks to relevant alternative locations are correlated with the shock to the current location, the conventional migration regression suffers from omitted-variable bias. For example, even if workers are very responsive to differences in labor demand conditions, there will be little incentive to migrate when workers' current and potential alternative locations face very similar labor demand shocks. In this case, the conventional migration regression estimate β will be close to zero even when the true migration elasticity with respect to local economic conditions is high. Second, the predicted values from the conventional regression fail to capture the true effects of the observed shocks on local populations because it omits variation in the outside options faced by potential migrants. Therefore, the conventional regression yields misleading conclusions about both worker mobility and the effects of observed shocks on local populations.

We analyze these interpretation problems in the context of a static model of local labor markets in which workers face mobility costs and have idiosyncratic preferences for living in different locations. The model is deliberately simple, focusing on costly migration while omitting features

[4] Equation (19) defines "as good as random" in this context.

such as capital markets or forward-looking behavior.[5] The simplicity of the model allows us to derive an intuitive approximation for the effect of a given vector of local labor demand shocks on local population growth, capturing the intuition just described. The shocks' effect on a location's population depends on the direct shock to the location minus the migration-weighted average shock to other locations, which is a sufficient statistic for relevant spillovers.

We use this model to derive a novel decomposition for β in the conventional migration regression. We show that β increases in i) the ratio of migration and labor demand elasticities, ii) the baseline share of migrants in the national population, and iii) an "attenuation factor" that is below one when shocks are particularly positively correlated between regions with large migration flows. This decomposition allows us to contrast various common interpretations of β against true model-based responses to observed or counterfactual shocks. Examples include interpreting β as the effect of a unit shock to a single location or interpreting $\beta(\hat{z}_k - \hat{z}_\ell)$ as the difference in the effects of shocks on locations k and ℓ. We show that these and similar interpretations provide poor predictions for the true effects on local populations, particularly when the shocks used to estimate β are correlated across migrant-connected locations and thus attenuation is severe.[6]

In contrast, we find that these issues in population regressions do not imply similar problems when studying the effects of local shocks on local *wage* changes. When the conventional regression estimate of β is close to zero, wage regression estimates will experience minimal confounding from migration, irrespective of the mechanism driving the small estimate of β. Intuitively, because labor demand shocks have a direct effect on local wages and an indirect effect through migration spillovers, when $\beta \approx 0$ this indirect effect is approximately uncorrelated with the direct effect and therefore does not substantially bias the estimate of the local shock's effect on the local wage. So, even when population regressions yield misleading interpretations, associated wage analyses may still yield valid conclusions.

The full version of the model incorporates mobility costs across both regions and industries, and shows that the interpretation problems in the conventional migration regression are exacerbated

[5] As discussed in Section 2 and Appendices A.14 and A.15, we find very similar results when incorporating inelastic housing supply and local agglomeration economies.

[6] We show that β is correctly interpreted as the *difference* in the *average* effect of *all* shocks across locations facing higher vs. lower *direct* shocks. While correct, this interpretation is quite limited. Even when β is small, the shocks may have led to substantial spatial reallocation because locations facing similar direct shocks may exhibit substantial variation in shocks to their migrant-connected locations.

in this more realistic setting. Many papers estimating (1) analyze local labor demand shocks with a shift-share structure in which typically $\hat{z}_\ell = \sum_n \frac{L^0_{\ell n}}{L^0_\ell} \hat{x}_n$, where $\hat{x}_n$ is a national labor demand shock facing industry n, and $L^0_{\ell n}/L^0_\ell$ is industry n's initial share of employment in location ℓ.[7] If industry switching costs are large, workers see diminished benefit to moving across locations because they primarily face the shock to their industry no matter where they choose to live. When labor demand shocks have an industry component, the presence of industry switching frictions thus reduces migration beyond what one would observe in a setting with regional frictions alone, further attenuating β. Moreover, with industry frictions, workers in different industries face different shocks and have different outside options, leading to complex spillovers across industries and regions, all of which are omitted by the conventional migration regression. As a result, our quantitative analysis finds that the conventional regression misses about half of the true population reallocation in response to national industry-level labor demand shocks.

Given the failure of the conventional migration regression to provide accurate predictions of the effects of observed or counterfactual local labor demand shocks, we propose alternative approaches leveraging the model-based expression that relates the vector of local shocks to population changes in each location. With observed labor demand shocks and data on the pre-shock migration flows between locations or industry-location pairs (depending on whether the shocks have an industry component), one can use nonlinear least squares (NLLS) to estimate the relevant parameter and predict the effects of observed or counterfactual labor demand shocks. As an alternative, we also show how an approximation to the model implies a model-consistent version of the conventional migration regression, which incorporates an appropriate migration-weighted average of shocks to other locations (or location-industries). Similarly, the same approximation admits a regression at the location pair level, relating the change in migration flow to the difference in shocks facing each location in the pair. While these procedures are specific to our model, the same approach can be applied to extended models that include additional economic forces.

To assess the quantitative importance of our insights, we study the effects of simulated shocks in the context of real-world economic geography and migration patterns. We base our simulations

[7] Of the papers estimating conventional migration regressions listed in footnote 2, the following use shift-share labor demand shock measures: Topalova (2010); Wozniak (2010); McCaig (2011); Autor et al. (2013); Dauth et al. (2014); Cadena and Kovak (2016); Hakobyan and McLaren (2016); Dix-Carneiro and Kovak (2017); Dix-Carneiro and Kovak (2019); Greenland et al. (2019); Albert and Monràs (2020); Boustan et al. (2020); Notowidigdo (2020); Autor et al. (2021); Choi et al. (2021); and Faber et al. (2021).

on longitudinal administrative data covering formally employed workers in Brazil, which allow us to observe worker transitions across locations and industries. Due to data limitations, we use formal employment as a proxy for local population in this exercise. We simulate labor demand shocks following a variety of data-generating processes and generate model-implied local population changes when workers face migration frictions alone or both migration and industry frictions. The resulting population changes are then used to estimate the conventional migration regression, the results of which we compare to the true model-based effects of the simulated labor demand shocks.

The results of this simulation exercise confirm the practical quantitative importance of the concerns raised by the theoretical analysis. The regression-based estimates poorly predict the effects of the simulated shocks themselves, and estimates of the effects of shocks to individual locations are even less accurate. These problems are largest when analyzing industry-level shocks via shift-share regressions in the presence of location and industry frictions.

To assess the performance of our model-consistent NLLS procedure we introduce random variation from unobservable factors into the simulation procedure and vary the magnitude of this residual variation.[8] For all magnitudes of the unobserved component, the predictive performance of this NLLS procedure far outperforms that of the conventional migration regression, which often performs as poorly as an uninformative prediction of zero migration response. Our findings further highlight the importance of accounting for observed connections between industries in addition to locations and the value of leveraging data on employment changes at the location-industry level when studying the effects of industry shocks.

Our results inform the many literatures estimating the conventional migration regression in (1), which span development, international, labor, resource, and urban economics.[9] We emphasize interpretation challenges resulting from cross-location spillovers that violate the stable unit treatment value assumption (SUTVA). A small number of papers in this literature anticipated this spillover issue and introduced *ad-hoc* controls in an effort to address it. While studying the effects of Chinese import competition on U.S. local population changes, Greenland et al. (2019) control for the weighted average of shocks to other locations, with weights proportional to the squared inverse distance. In studying the effects of droughts on regional labor and capital reallocation in Brazil,

[8]We do not implement the NLLS procedure in the simulations without additional residual variation because it would fit perfectly by construction.

[9]See footnote 2 for references.

Albert et al. (2021) include a migration-weighted average of drought measures in other locations.[10] Our analysis refines and provides a theoretical rationalization for these approaches to dealing with cross-location spillovers.

Our findings have further implications for models of local labor markets. Since conventional migration regression estimates are not informative about underlying mobility parameters, the estimates should not be used to guide modeling choices. Specifically, estimates of β that are close to zero may simply reflect correlated shocks rather than a lack of mobility, so they do not justify omitting migration from models of local labor markets (e.g. Adão, 2016; Galle et al., forthcoming). Moreover, our approach elucidates how spatial migration networks lead to differences in the outside options facing workers in different locations. Studies that do not estimate the conventional regression but instead employ model-consistent specifications arising from models that assume away spillovers through migration networks are likely to experience issues similar to those we examine in the context of conventional migration regressions (e.g. Kleven et al., 2013; Agrawal and Foremny, 2019; Notowidigdo, 2020).

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); 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) and Fan (2019). On the labor demand side, the standard multi-sector Armington (1969) model that we employ has been used by Anderson and van Wincoop (2003), Head and Mayer (2014), and Adão et al. (2019), among many others.

Methodologically, our paper is related to Adão et al. (2021), 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 a local labor demand shock in presence of spatial spillovers; our analysis focuses on spillovers from migration

[10] Spillovers via migration have also been considered contexts other than local population growth. Bertoli and Moraga (2013) consider spillovers in international migration flows between countries. Arthi et al. (2019) consider migration spillovers when studying the effects of the 19^{th} century cotton famine on local mortality rates in the UK. When studying effects of trade liberalization on the non-agriculture share of local employment in China, Tian et al. (2021) include both the direct tariff shock facing a location and a migration-weighted average of tariff changes facing other locations. Imbert et al. (2021) control for an inverse-distance-weighted average of shocks to other locations in their study of how in-migration affects firms in China.

while theirs focuses on interregional trade.[11] Yet, the papers' objectives are quite distinct, with Adão et al. (2021) seeking to estimate aggregate general equilibrium effects of a spatially heterogeneous shock to labor demand and our analysis informing a commonly used reduced-form approach that generally ignores potential spillovers.

The paper proceeds as follows. Section 2 describes the baseline model of local labor markets in which workers face costs of moving across locations, solves the model, and relates model-based expressions to observable quantities. Section 3 defines and characterizes the conventional migration regression coefficient and discusses problems when using the associated estimates to make within-sample or counterfactual predictions. It also describes model-consistent estimation procedures. Section 4 presents the full model including frictions in switching both locations and industries. Section 5 presents the data and descriptive statistics on worker transitions across locations and industries. Section 6 presents the simulation-based quantitative results, and Section 7 concludes.

2 Baseline Model of Mobility with Regional Frictions

2.1 Setup

This section presents a stylized model that yields clear and intuitive expressions for how local labor demand shocks affect local populations in the presence of costly mobility. Although this framework omits a number of potentially relevant mechanisms, including forward-looking behavior and capital markets, it highlights general issues with standard regression-based approaches to understanding how changes in local populations respond to local shocks. In Appendices A.14 and A.15, we show that very similar results emerge when incorporating inelastic housing supply and local agglomeration economies.[12] Here, we introduce a baseline version of the model in which workers face mobility costs across locations but are freely mobile across industries. Section 4 studies a more realistic setting with frictions in both dimensions.

Consider a small open economy consisting of multiple sectors indexed by $n \in \mathcal{N}$ produced in

[11] The extended model of Adão et al. (2021) features migration, but because it lacks migration costs it does not allow for migration network linkages that vary by location.

[12] Specifically, Appendix A.14 shows that the model presented here is isomorphic to baseline model presented here under the low-mobility approximation introduced below in (17). Appendix A.15 shows that the results are isomorphic to those presented here for agglomeration economies in which productivity increases in (i) local industry employment or (ii) overall local employment across industries.

locations indexed by $\ell \in \mathscr{L}$. Products within a sector 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 a sector-n good produced in location ℓ, $p_{\ell n}$, is constant across destinations. Homogeneous labor is the only input, and sector-location productivity is $\varphi_{\ell n}$. Output and labor markets are perfectly competitive, so $p_{\ell n} = w_\ell/\varphi_{\ell n}$, where w_ℓ is the wage in location ℓ.

Individuals make optimal choices over consumption bundles and their location of employment and face frictions in moving across locations. We assume that workers may migrate within the country of interest by choosing among locations $\ell \in \mathscr{L}$, but not internationally. Unless explicitly specified otherwise, sums over locations cover only places in the country of interest. Each of L workers inelastically supplies one unit of labor in their chosen location, so changes in employment are equivalent to changes in population. Utility from consumption is Cobb-Douglas across sectors and CES across location-specific varieties within sector:

$$U = \sum_n \eta_n \ln(Q_n), \quad \text{where } Q_n \equiv \left(\sum_{\ell \in \mathscr{L}_w} q_{\ell n}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{2}$$

and $\sum_n \eta_n = 1$. Note that Q_n aggregates over location-specific varieties produced around the world ($\ell \in \mathscr{L}_w$), including those within the country of interest ($\mathscr{L} \subset \mathscr{L}_w$). Normalizing world aggregate expenditure to one, the aggregate optimal demand bundle satisfies

$$q_{\ell n} = p_{\ell n}^{-\sigma} P_n^{\sigma-1} \eta_n, \tag{3}$$

where $P_n \equiv \left(\sum_{\ell \in \mathscr{L}_w} p_{\ell n}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ is the CES exact price index for sector n, including varieties produced globally. The small-country assumption implies that P_n is set on the world market and is exogenous to developments in the country of interest.

The initial distribution of individuals across locations is taken as given, with L_o^0 workers in location o, and the indirect utility of choosing destination location ℓ for individual i initially living in o is[13]

$$V_{io\ell} = \ln\left(\frac{w_\ell}{P}\right) - \ln \tau_{o\ell} + \frac{1}{\theta}\epsilon_{i\ell}, \tag{4}$$

where $P \equiv (\Pi_n P_n^{\eta_n}) / (\Pi_n \eta_n^{\eta_n})$ is the Cobb-Douglas exact price index across sectors, $\ln \tau_{o\ell}$ is the utility cost of moving from o to ℓ, and $\epsilon_{i\ell}$ is a taste shock following a type-I extreme-value (standard Gumbel) distribution that is i.i.d. across individuals and locations. The parameter θ determines the strength of location preferences relative to the log real wage, with smaller values of θ implying stronger location preference and larger frictions across locations.

2.2 Regional Labor Demand

Given perfectly competitive markets, $p_{\ell n} = w_\ell/\varphi_{\ell n}$, and revenue to sector-$n$ producers in location ℓ is

$$R_{\ell n} = \eta_n \left(\frac{w_\ell/\varphi_{\ell n}}{P_n}\right)^{1-\sigma}. \tag{5}$$

Perfect competition and regional trade balance imply that the regional wagebill equals total regional revenue across sectors:

$$w_\ell L_\ell = \sum_n R_{\ell n}. \tag{6}$$

Plugging in (5) and rearranging yields regional labor demand:

$$L_\ell = D_\ell w_\ell^{-\sigma}, \quad \text{where } D_\ell \equiv \sum_n \eta_n \left(\varphi_{\ell n} P_n\right)^{\sigma-1}. \tag{7}$$

The term D_ℓ is a regional demand shifter affected by changes in productivity $\varphi_{\ell n}$ or industry price indexes P_n. Equation (7) shows that σ is the labor demand elasticity, which is larger when product varieties across locations are more substitutable.

We consider an economic shock that affects this demand shifter across locations in the country of interest.[14] Let hats represent proportional changes relative to a no-shock counterfactual, and

[13] Exogenous local amenities can be incorporated without changing the analysis. Since we do not use data on wages, w_d can be reinterpreted as the amenity-adjusted wage.

[14] For example, a change to the vector of sectoral price indexes yields a shift-share structure in $\hat{D}_\ell$. Differentiating the definition of D_ℓ in (7) yields

$$\hat{D}_\ell = (\sigma - 1) \sum_n \frac{\eta_n (\varphi_{\ell n} P_n)^{\sigma-1} w_\ell^{-\sigma}}{D_\ell w_\ell^{-\sigma}} \hat{P}_n = (\sigma - 1) \sum_n \frac{L_{\ell n}}{L_\ell} \hat{P}_n.$$

assume small changes so $\hat{x} \equiv dx/x = d\ln x$.[15] The total derivative of (7) is then

$$\hat{L}_\ell = \hat{D}_\ell - \sigma\hat{w}_\ell. \tag{8}$$

This expression describes how the quantity of labor demanded in location ℓ responds to local wage changes and local labor demand shocks, and can be represented in matrix notation as

$$\hat{\mathbf{L}} = \hat{\mathbf{D}} - \sigma\hat{\mathbf{w}}. \tag{9}$$

2.3 Regional Labor Supply

Given the assumptions in (4), the probability that an individual in location o chooses to live in location ℓ is given by $\pi_{o\ell}$, where

$$\pi_{o\ell} = \frac{(w_\ell/\tau_{o\ell})^\theta}{\sum_d (w_d/\tau_{od})^\theta}. \tag{10}$$

We refer to these location-choice probabilities as "out-migration shares." Equation (10) shows that θ is the migration elasticity reflecting how relative out-migration shares across destination locations respond to relative wage changes; a larger θ implies less variation in the idiosyncratic taste shocks and hence increased responsiveness to wage differences. Labor supply to location ℓ is then determined by these out-migration shares and the initial distribution of workers across locations, L_o^0:

$$L_\ell = \sum_o f_{o\ell} = \sum_o \pi_{o\ell} L_o^0, \tag{11}$$

where $f_{o\ell} \equiv \pi_{o\ell} L_o^0$ is the number of people initially in location o who choose location ℓ.

Now consider a labor demand shock leading to small changes in wages across (potentially all) locations, as in (8), while holding moving costs $\tau_{o\ell}$ fixed.[16] Totally differentiating (10) and (11)

[15] We maintain this small-shock assumption in the main text for expositional clarity, but Appendix A.12 uses exact hat algebra following Dekle et al. (2008) to study exact responses to potentially large shocks. In results available upon request, we confirm that the small-shock approximation is accurate even for shocks that are far larger than those facing Brazilian labor markets during the country's early 1990s trade liberalization.

[16] Appendix A.1 generalizes our results to incorporate both labor demand shocks and labor supply shocks resulting from changes in $\tau_{o\ell}$. This does not affect the ensuing results.

yields[17]

$$\hat{L}_\ell = \sum_o \gamma_{o\ell}\hat{\pi}_{o\ell} = \theta\left(\hat{w}_\ell - \sum_o \gamma_{o\ell}\sum_d \pi_{od}\hat{w}_d\right). \tag{12}$$

Here $\gamma_{o\ell} \equiv f_{o\ell}/L_\ell$ is the share of those in location ℓ who arrived from location o, which we refer to as the "in-migration share."[18] This expression describes how the quantity of labor supplied in location ℓ responds to wage changes across all destinations and can be written in matrix notation as

$$\hat{\mathbf{L}} = \theta\left(I - \Gamma'\Pi\right)\hat{\mathbf{w}}, \tag{13}$$

where $\Gamma = (\gamma_{o\ell})$ is the matrix of in-migration shares and $\Pi = (\pi_{o\ell})$ is the matrix of out-migration shares.

2.4 Equilibrium Response to Shocks

Equating the labor demand and labor supply expressions in changes in (9) and (13) yields the equilibrium effects of regional labor demand shocks on regional population and wages:

$$\hat{\mathbf{L}} = \Omega\hat{\mathbf{D}}, \tag{14}$$

$$\hat{\mathbf{w}} = \frac{1}{\sigma}(\mathbb{I} - \Omega)\hat{\mathbf{D}}, \tag{15}$$

$$\Omega \equiv \left(\mathbb{I} - \left(\mathbb{I} + \frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right)\right)^{-1}\right). \tag{16}$$

Equation (14) shows the relationship between population changes and local labor demand shocks, i.e. the relationship whose empirical investigation we seek to inform. Two practical implications are immediately evident. First, the population change in a given location depends upon the shocks to *all* other locations. In fact, if all locations face the same shock so $\hat{D}_\ell = \hat{d}\ \forall\ell$, then the population in each location remains unchanged.[19] Second, the importance of each shock in changing local population depends upon migrant connections between the focal location and other locations, as reflected in Γ and Π, which enter Ω. This implies that Γ and Π are sufficient statistics for all relevant

[17] Note that, in a slight abuse of notation to avoid clutter, $\pi_{o\ell}$ in (11) is the out-migration share in the presence of labor demand shocks, while the version in (12) is the out-migration share in the absence of the labor demand shocks, such that $\hat{\pi}_{o\ell}$ is the proportional difference between these two sets of migration shares.

[18] The labor supply expression in (12) is similar to appendix equation (F.3) in Berkes et al. (2021).

[19] 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 $\hat{d}$.

effects of the moving costs ($\tau_{o\ell}$) and sector-location productivity differences ($\varphi_{\ell n}$), so we will not need to estimate either as long as we can measure Γ and Π (as discussed in the next subsection).

To gain intuition for the ways in which these migrant connections influence the effects of shocks on connected regions, we consider an approximation in which mobility is low enough that the effects of indirect connections become unimportant. 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 (16) and then (14) yields[20]

$$\hat{L}_\ell \approx \frac{\theta}{\sigma}\left(\sum_o \gamma_{o\ell}(\hat{D}_\ell - \hat{D}_o) + \sum_d \pi_{\ell d}(\hat{D}_\ell - \hat{D}_d)\right). \tag{17}$$

The effect of a shock to location ℓ on the population in ℓ depends upon how the local shock compares to shocks facing all other locations, with more weight placed upon shocks in the location's typical migrant sources (captured by $\gamma_{o\ell}$) and its typical migrant destinations (captured by $\pi_{\ell d}$). This expression also shows that when the migration elasticity θ is larger, more individuals choose to migrate in response to a given vector of shocks. Similarly, when the local labor demand elasticity σ is smaller, a given vector of demand shifts leads to larger population changes.

2.5 Linking to Observables

The expressions relating population changes to labor demand shocks in (14) and (17) depend upon the ratio of elasticities θ/σ, the demand shocks $\hat{\mathbf{D}}$, and the migration share matrices Γ and Π, which refer to a counterfactual without labor demand shocks. While we keep θ/σ as an unknown parameter, in this subsection we consider how to link the labor demand shocks and counterfactual migration shares 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 $\hat{\mathbf{D}} = (\hat{D}_\ell)$, but instead observe a demand shifter $\hat{\mathbf{z}} = (\hat{z}_\ell)$ such that

$$\hat{D}_\ell = \hat{z}_\ell + \zeta_{1\ell}, \tag{18}$$

[20] Appendix A.2 derives (17).

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

$$\mathbb{E}\left[\hat{z}_\ell \mid \boldsymbol{\zeta}_1\right] = \mu \quad \forall \ell, \tag{19}$$

where μ is a constant, while allowing the shocks to be mutually correlated (c.f. Borusyak et al. (2022), Assumption 1). In words, the shock to each location has the same mean conditional on all unobserved demand shocks. This assumption allows us to emphasize issues that arise even in the best-case empirical scenario, when the observed shocks are not confounded by unobservables.[21]

Given (18), we can rewrite (14) as

$$\hat{\mathbf{L}} = \Omega\hat{\mathbf{z}} + \boldsymbol{\zeta}_2, \tag{20}$$

where $\boldsymbol{\zeta}_2 \equiv \Omega\boldsymbol{\zeta}_1 = (\zeta_{2\ell})$ is a random error term. Similarly isolating the impacts of observed and unobserved shocks in (17) and rearranging terms, we arrive at an approximation for the population responses to shocks that will be particularly useful for our analytical results:

$$\hat{L}_\ell \approx \frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}\left(\hat{z}_\ell - \hat{z}_{-\ell}\right) + \zeta_{2\ell}, \qquad \text{where } \hat{z}_{-\ell} \equiv \sum_{k\neq\ell}\frac{F_{k\ell}}{M_\ell}\hat{z}_k, \tag{21}$$

$F_{k\ell} \equiv \frac{1}{2}(f_{k\ell} + f_{\ell k})$ is the average of migration flows between k and ℓ, and $M_\ell \equiv \sum_{k\neq\ell} F_{k\ell}$ is the average number of migrants to and from ℓ.[22] Equation (21) shows that local population responds to the gap between the local shock, $\hat{z}_\ell$, and the migration-weighted average of shocks to other locations, $\hat{z}_{-\ell}$. This response is stronger when the migration elasticity θ is larger and when the labor demand elasticity σ is smaller; it is also stronger in locations with a larger migration share, M_ℓ/L_ℓ.

To empirically operationalize the preceding two expressions, we must estimate the terms referring to the counterfactual setting without labor demand shocks: Ω in (20), which depends on Γ and Π, and M_ℓ, L_ℓ, and $F_{k\ell}$ in (21). We do so by assuming that the out-migration shares from the pre-shock period would have persisted in the absence of labor demand shocks.

To formalize this idea, define time $t = -1$, 0, and 1, such that the labor demand shocks arrive

[21] As with any observational research design, the conventional migration regression will yield inconsistent estimates when the shocks are correlated with unobserved determinants of regional 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.

[22] Appendix A.3 provides the derivation.

between $t = 0$ and $t = 1$, consistent with our earlier notation L_ℓ^0 for the observed population before the arrival of the shock. Let $\Pi \equiv (\pi_{o\ell})$ be 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 $\Pi^0 \equiv (\pi_{o\ell}^0)$ is the matrix of observed out-migration shares between $t = -1$ and 0. In other words, we presume that the existing migration patterns would continue absent observed and unobserved shocks at $t = 1$. This approach allows us to measure Π using pre-shock data while accounting for the possibility of persistent population trends across periods. Given these estimates and the initial population in each location, we calculate the remaining quantities in the no-shock counterfactual that appear in (20) and (21), including Γ, M_ℓ, L_ℓ, and $F_{k\ell}$.[23] 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.[24]

3 Conventional Regression and Model-Consistent Procedures

The model in the preceding section describes how local populations change in response to labor demand shocks across locations. In this section, we investigate how to interpret the results of the conventional migration regression when the data generating process follows the model just described.

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

$$\hat{L}_\ell = \alpha + \beta \hat{z}_\ell + \varepsilon_\ell, \tag{22}$$

[23] We follow the definitions of these variables to calculate their values in the no-shock counterfactual: $f_{k\ell} = L_k^0 \pi_{k\ell}^0$, $F_{k\ell} = \frac{1}{2}(f_{k\ell} + f_{\ell k})$, $M_\ell = \sum_{k \neq \ell} F_{k\ell}$, $L_\ell = \sum_k f_{k\ell}$, and $\gamma_{k\ell} = f_{k\ell}/L_\ell$ with $\Gamma = (\gamma_{k\ell})$.

[24] Having linked the right-hand side of (20) and (21) to observables, we briefly discuss the left-hand side, which represents the difference in local population in $t = 1$ between counterfactuals with and without labor demand shocks. In place of this unobserved quantity, researchers employ the observed change in population between $t = 0$ and 1. As shown in Appendix A.4, the two are equivalent up to an additional error term representing population growth in the no-shock counterfactual. Because this additional term is orthogonal to the observed labor demand shocks under the exogeneity assumption (19), our identification and estimation procedures discussed below remain unaffected.

with observations weighted by the pre-shock population L_ℓ^0. The corresponding OLS estimator of β is

$$\hat{\beta} = \frac{\sum_\ell L_\ell^0 \hat{L}_\ell \left(\hat{z}_\ell - \bar{z}\right)}{\sum_\ell L_\ell^0 \left(\hat{z}_\ell - \bar{z}\right)^2} \tag{23}$$

where $\bar{z} = \sum_\ell L_\ell^0 \hat{z}_\ell / \sum_\ell L_\ell^0$.[25]

We do not assume that equation (22) is correctly specified. Instead, we assume that the data on $\hat{L}_\ell$ are generated according to the model of Section 2: equation (20) or its approximation (21). We also maintain the assumption that the $(\hat{z}_\ell)$ are as-good-as-randomly assigned to regions with respect to all unobserved shocks to labor demand and supply, as in (19).

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

$$\beta = \frac{\mathbb{E}\left[\sum_\ell L_\ell^0 \hat{L}_\ell \left(\hat{z}_\ell - \bar{z}\right)\right]}{\mathbb{E}\left[\sum_\ell L_\ell^0 \left(\hat{z}_\ell - \bar{z}\right)^2\right]}. \tag{24}$$

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 as fixed). This formulation has a number of benefits. First, it applies in finite samples of regions and thus allows us to avoid random sampling assumptions, which are inappropriate in typical applications including all regions of the country. Second, as we show in Appendix A.5, 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 unobserved shocks on $\hat{\beta}$ vanishes, we expect $\hat{\beta} \approx \beta$. Third, unlike $\mathbb{E}[\hat{\beta}]$, the formulation in (24) is analytically 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 results intuitively interpretable, we focus throughout on the low-mobility approximation to $\hat{L}_\ell$, shown in (21).

To understand our characterization, it is helpful to define some notation. Let $M \equiv \sum_{o \neq d} f_{od}$ be

[25] Results similar to those derived below, albeit more cumbersome, can be obtained for unweighted regressions used in some studies.

the national total number of migrants in the country of interest, so M/L is the migrant share of the population. We then define similar quantities that would prevail if the observed local populations reflected an economy with no migration costs, i.e. if $\tau_{od} = 1\ \forall o, d$. Without migration costs, the probability of moving to any destination will equal the destination's share of national population, regardless of the origin.[26] Define $\widetilde{f}_{od} \equiv L_o^0 \cdot \frac{L_d}{L}$ as the number of migrants from o to d and $\widetilde{M} \equiv \sum_{o \neq d} \widetilde{f}_{od}$ as the national total number of migrants in this costless-migration setting. The following theorem then characterizes β estimated with data generated by the model.

Theorem 1. *Suppose the data are generated by the low-mobility approximation to the baseline model, in which workers face mobility costs across locations but are freely mobile across industries (21), and the shocks are homoskedastic, i.e.* $\mathrm{Var}[\hat{z}_\ell]$ *is the same for all* ℓ*. Then*

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

where $\rho \equiv \sum_{o \neq d} \frac{f_{od}}{M} \mathrm{Corr}[\hat{z}_o, \hat{z}_d]$ *is the average correlation between shocks to pairs of distinct regions weighted by migration flows, and* $\widetilde{\rho} = \sum_{o \neq d} \frac{\widetilde{f}_{od}}{\widetilde{M}} \mathrm{Corr}[\hat{z}_o, \hat{z}_d]$ *is a similar shock correlation weighted by the flows in the no-migration-cost scenario.*

Appendix A.6 proves this result, along with its generalization to heteroskedastic shocks.

Equation (25) shows that β can be viewed as a product of several factors with distinct economic meanings. First, β depends on the structural *elasticities*: as discussed above, true migration responses are larger when θ is large or σ is small, and β inherits those relationships. Second, β increases with the national migration *share* M/L (the term $\widetilde{M}/L$ is generally close to one and can therefore be ignored in practice).[27] Third, β is smaller when $\rho > \widetilde{\rho}$; that is, if shocks are particularly positively correlated between regions with large migration flows (i.e. flows that exceed those one would expect based solely on location sizes)—for instance, if migration is more likely at close geographic distances and shocks exhibit spatial correlation. We call the term $\frac{1-\rho}{1-\widetilde{\rho}}$ the *attenuation*

[26]When $\tau_{\ell d} = 1\ \forall \ell, d$, (10) implies that the share of individuals choosing destination d does not depend on source ℓ: $\pi_{\ell d} = w_d^\theta / \sum_{d'} w_{d'}^\theta \equiv \pi_d$. In that case, $L_d = \sum_\ell \pi_{\ell d} L_\ell^0 = \pi_d L$, so $\pi_d = L_d/L$, i.e. the probability of choosing d equals d's share of national population. This in turn implies that the flow of individuals from ℓ choosing d is $\pi_{\ell d} L_\ell^0 = \pi_d L_\ell^0 = L_\ell^0 L_d / L$.

[27]Specifically, $\frac{\widetilde{M}}{L} = 1 - \sum_\ell \frac{\widetilde{f}_{\ell\ell}}{L} = 1 - \sum_\ell \left(\frac{L_\ell}{L}\right)^2$ 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.

factor: it reflects how the mutual correlation of shocks along the migration network attenuates the OLS coefficient. When shocks to migrant-connected locations are strongly correlated, individuals have minimal incentive to migrate, so we observe minimal migration response and estimate a relatively small (attenuated) value of β.

3.3 What Do We Learn from β?

The expression for β in Theorem 1 reveals how one can and cannot interpret estimates from conventional migration regressions. Here, we will consider various potential interpretations and compare them to correct interpretations in light of the model. To be precise about the causal effects of interest, we define $\hat{\mathbf{L}}(\hat{\mathbf{z}}) \equiv \Omega\hat{\mathbf{z}} + \boldsymbol{\zeta}_2$, the model-predicted population growth in ℓ under shock vector $\hat{\mathbf{z}}$, as in (20).[28] $\hat{L}_\ell(\hat{\mathbf{z}})$ depends on the entire shock vector $\hat{\mathbf{z}}$, not just the direct shock to ℓ. To analyze counterfactual shocks, we denote a vector $\hat{\mathbf{z}}$ whose ℓ^{th} element has been changed from $\hat{z}_\ell$ to the scalar value a by $(\hat{\mathbf{z}} \overset{\ell}{\hookleftarrow} a) \equiv (\hat{z}_1, ..., \hat{z}_{\ell-1}, a, \hat{z}_{\ell+1}, ..., \hat{z}_R)$.

Effect of a shock to a single location. Researchers often use β to infer the causal effect of a labor demand shock to a single location on that location's population. A common approach interprets β as the effect of a unit labor demand shock to location ℓ on ℓ's population, implicitly fixing other locations' labor demand shocks at zero. Equation (21) shows that the true causal effect is $\hat{L}_\ell(\mathbf{0} \overset{\ell}{\hookleftarrow} 1) - \hat{L}_\ell(\mathbf{0}) \approx \frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}$. The estimate of β in (25) differs from this model-implied population growth in two important ways. First, it misses location-specific heterogeneity in the effect of a local shock; the model-implied population growth depends on the location-specific migration share $\frac{M_\ell}{L_\ell}$, while the regression estimate depends on the national migration share $\frac{M}{L}$ (recall $\widetilde{\frac{M}{L}} \approx 1$). As we will see in Section 5.2, local migration shares often vary substantially, so this heterogeneity can be important in practice.[29]

Second, the estimate of β includes the attenuation factor $\frac{1-\rho}{1-\bar{\rho}}$ resulting from the correlation in shocks among locations with strong migrant connections. Because of this attenuation factor, a researcher may estimate a small value of β that is close to zero even when a shock to any single location would actually lead to a substantial migration response. Consistent with this analysis,

[28] Note that $\hat{L}_\ell(\hat{\mathbf{z}})$ includes the effect of unobserved shocks, so $\hat{L}_\ell(\mathbf{0}) = \zeta_{2\ell}$.

[29] This problem would not arise if the outcome variable of the regression was the net population change divided by the gross migration flow M_ℓ rather than divided by the overall population L_ℓ.

Section 6 shows that β substantially understates the true effect of a unit shock to a single location in a realistic empirical setting.

The same analysis and conclusions apply to other interpretations of β involving the effect of a shock to a single location. Examples include i) interpreting $\beta \hat{z}_\ell$ as the effect of the observed shock to ℓ, holding all other locations' shocks fixed at their observed values ($\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} 0)$), ii) interpreting $\beta(\hat{z}_k - \hat{z}_\ell)$ as the difference in population growth caused by the difference in observed shocks across locations k and ℓ ($\hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} \hat{z}_k) - \hat{L}_\ell(\hat{\mathbf{z}})$), and iii) interpreting $\beta(\hat{z}^{p75} - \hat{z}^{p25})$ as the effect of changing the shock to location ℓ from the 25^{th} to the 75^{th} percentile of the shock distribution ($\hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} \hat{z}^{p75}) - \hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} \hat{z}^{p25})$). All of these cases suffer from the same problems as the unit shock discussed above: the regression-based prediction omits local heterogeneity in migration intensity and suffers from attenuation resulting from correlated shocks across migrant-connected locations.[30]

Effect of shocks to all locations. The preceding discussion makes clear that the conventional migration regression does not capture the effects of shocks to individual locations. But one might still think it could capture the effect of the entire observed shock vector $\hat{\mathbf{z}}$ on population growth in each location ℓ, i.e. $\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\mathbf{0}) \approx \frac{2\theta}{\sigma} \frac{M_\ell}{L_\ell} (\hat{z}_\ell - \hat{z}_{-\ell})$.

It is important to note that the (population-weighted) average effect of any shock vector $\hat{\mathbf{z}}$ is zero by construction, since national population is assumed to be unaffected by the shock.[31] A meaningful interpretation of β must therefore involve comparisons across locations. To see why, consider using $\beta \hat{z}_\ell$ to predict the shock vector's effect on population growth in each location ℓ. In many empirical contexts, the shocks to all locations have the same sign, so $\beta \hat{z}_\ell$ would incorrectly imply either population growth in all locations or population declines in all locations. A more sensible approach attributes to the local shock only its effect relative to the mean, $\beta\,(\hat{z}_\ell - \bar{z})$. This approach ensures that the predicted effects of the shocks average to zero across locations and is consistent with the inclusion of an intercept term α in the conventional regression (22). It also yields the best (mean

[30] In case i) the true effect is $\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} 0) \approx \frac{2\theta}{\sigma} \frac{M_\ell}{L_\ell} \hat{z}_\ell$, while the regression-based prediction is $\beta \hat{z}_\ell = \frac{2\theta}{\sigma} \cdot \frac{M/L}{\widetilde{M}/L} \cdot \frac{1-\rho}{1-\tilde{\rho}} \hat{z}_\ell$. In case ii) the true effect on location ℓ's population of changing ℓ's shock from $\hat{z}_\ell$ to $\hat{z}_k$, fixing other locations' shocks at their observed values is $\hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} \hat{z}_k) - \hat{L}_\ell(\hat{\mathbf{z}}) \approx \frac{2\theta}{\sigma} \frac{M_\ell}{L_\ell} (\hat{z}_k - \hat{z}_\ell)$, while the regression-based prediction is $\beta(\hat{z}_k - \hat{z}_\ell) = \frac{2\theta}{\sigma} \cdot \frac{M/L}{\widetilde{M}/L} \cdot \frac{1-\rho}{1-\tilde{\rho}} (\hat{z}_k - \hat{z}_\ell)$. In case iii) the true effect is $\hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} \hat{z}^{p75}) - \hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\leftarrowtail} \hat{z}^{p25}) \approx \frac{2\theta}{\sigma} \frac{M_\ell}{L_\ell} (\hat{z}^{p75} - \hat{z}^{p25})$, while the regression-based prediction is $\beta(\hat{z}^{p75} - \hat{z}^{p25}) = \frac{2\theta}{\sigma} \cdot \frac{M/L}{\widetilde{M}/L} \cdot \frac{1-\rho}{1-\tilde{\rho}} (\hat{z}^{p75} - \hat{z}^{p25})$.

[31] This assumption can be relaxed by allowing for population growth or international migration.

squared error (MSE) minimizing) prediction of the expected population growth in location ℓ among linear functions of $(\hat{z}_\ell - \bar{z})$.[32]

However, using $\hat{z}_\ell - \bar{z}$ alone to predict local population growth is likely to yield very poor predictions. To see why, use the decomposition of β in (25) to write

$$\beta\left(\hat{z}_\ell - \bar{z}\right) = \frac{2\theta}{\sigma} \cdot \frac{M/L}{\widetilde{M/L}} \left(\left(\hat{z}_\ell - \bar{z}\right) - \frac{\rho - \widetilde{\rho}}{1 - \widetilde{\rho}} \cdot \left(\hat{z}_\ell - \bar{z}\right) \right). \tag{26}$$

Compare this prediction to (21), the model's expected population growth in ℓ when facing the observed vector of demeaned shocks: $\frac{2\theta}{\sigma} \cdot \frac{M_\ell}{L_\ell} \left(\left(\hat{z}_\ell - \bar{z}\right) - \left(\hat{z}_{-\ell} - \bar{z}\right)\right)$. The estimate $\beta\left(\hat{z}_\ell - \bar{z}\right)$ omits two key sources of variation. As before, it omits heterogeneity based on the local migration intensity, $\frac{M_\ell}{L_\ell}$. More importantly, it replaces the shocks to other migration-connected locations, $(\hat{z}_{-\ell} - \bar{z})$, with the scaled local shock relative to the mean $\frac{\rho-\widetilde{\rho}}{1-\widetilde{\rho}}\left(\hat{z}_\ell - \bar{z}\right)$. Shocks to other locations are important drivers of local population change according to the model and can differ substantially among locations facing similar direct shocks, but this information is omitted from (26). Because the fitted values from the population regression only use information on the direct shocks to the location, they cannot account for the extent to which other migrant sources and destinations faced similar or different shocks. As we will see in Section 5.2, the omission of these cross-location spillover effects can be quantitatively important.

Rather than attempting to use β to predict the effect of $\hat{\mathbf{z}}$ on one location's population growth, an alternative approach would interpret $\beta(\hat{z}_k - \hat{z}_\ell)$ as the difference in the effects of all shocks between particular locations k and ℓ: $(\hat{L}_k(\hat{\mathbf{z}}) - \hat{L}_k(\mathbf{0})) - (\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\mathbf{0}))$.[33] This interpretation suffers from similar problems to those already discussed: it omits heterogeneity in migration intensity and fails to account for the particular differences between migrant-connected shocks $\hat{z}_{-k}$ and $\hat{z}_{-\ell}$.[34]

Yet, as with any regression coefficient estimate, there is a valid interpretation of β in (22);

[32] See Appendix A.7 for a proof.

[33] Note the distinction between this interpretation focusing on the effect of *all* shocks as opposed to interpretation ii) in footnote 30, which considers the effect of changing location $\ell's$ shock from $\hat{z}_\ell$ to $\hat{z}_k$, while holding all other locations' shocks fixed: $\hat{L}_\ell(\hat{\mathbf{z}} \overset{\ell}{\hookleftarrow} \hat{z}_k) - \hat{L}_\ell(\hat{\mathbf{z}})$.

[34] The difference in the effects of all shocks $\hat{\mathbf{z}}$ between locations k and ℓ is

$$\left(\hat{L}_k(\hat{\mathbf{z}}) - \hat{L}_k(\mathbf{0})\right) - \left(\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\mathbf{0})\right) \approx \frac{2\theta}{\sigma}\left[\left(\frac{M_k}{L_k}\hat{z}_k - \frac{M_\ell}{L_\ell}\hat{z}_\ell\right) - \left(\frac{M_k}{L_k}\hat{z}_{-k} - \frac{M_\ell}{L_\ell}\hat{z}_{-\ell}\right)\right].$$

$\beta(\hat{z}_k - \hat{z}_\ell)$ omits heterogeneity based on initial migration intensity and captures the differences in $\hat{z}_{-k}$ and $\hat{z}_{-\ell}$ using only $(\hat{z}_k - \hat{z}_\ell)$.

it measures the difference in average effects of *all* shocks between locations facing different *direct* shocks. The most straightforward way to demonstrate this interpretation is to consider the following regression, which relates the effect of all shocks on population growth in location ℓ to location ℓ's direct shock:

$$\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\mathbf{0}) = a + b\hat{z}_\ell + e_\ell. \tag{27}$$

Here b cannot be estimated directly because the dependent variable is not observed. However, the only difference between the specification in (27) and the conventional migration regression in (22) is on the left side. The dependent variable in (27) is the difference in counterfactual population growth with vs. without the observed shocks, $\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\mathbf{0})$, while the dependent variable in (22) is the observed population change over time, $\hat{L}_\ell$. As shown in Appendix A.4, the difference in these two measures is the population growth in the no-shock counterfactual plus the population effect of unobserved shocks. Because this difference is orthogonal to the observed labor demand shocks under the exogeneity assumption (19), the estimands β and b are equal. We can therefore interpret an estimate of β as we would interpret b: the difference in the average effect of *all* shocks between places facing high vs. low *direct* shocks, divided by the corresponding difference in direct shocks.[35] Importantly, even this limited interpretation only holds for the shocks used to estimate β. Counterfactual shocks will generally induce different correlation structures across migrant-connected regions that will not be captured by β estimated using the realized shocks.

To summarize, an estimate of β from the conventional migration regression in (22) does not yield accurate predictions regarding the causal effects of observed or counterfactual shocks because it ignores information on shocks to migrant-connected locations and omits heterogeneity based on local migration intensity. The estimate of β may be close to zero even when the observed shock led to substantial spatial reallocation, when workers are highly responsive to local economic conditions,

[35] To formalize this point, split the sample of locations into those with above-average shocks, $\ell|\hat{z}_\ell \geq \overline{z}$, and those with below-average shocks, $\ell|\hat{z}_\ell < \overline{z}$. One can then express β in terms of the difference in the average effect of all shocks between these two groups of locations:

$$\beta = \Lambda \mathbb{E}\left[\sum_{\ell|\hat{z}_\ell \geq \overline{z}} v_\ell(\hat{L}_\ell(\mathbf{z}) - \hat{L}_\ell(\mathbf{0})) - \sum_{\ell|\hat{z}_\ell < \overline{z}} v_\ell(\hat{L}_\ell(\mathbf{z}) - \hat{L}_\ell(\mathbf{0}))\right],$$

where $v_\ell \equiv \frac{L_\ell^0|\hat{z}_\ell - \overline{z}|}{\mathbb{E}\left[\sum_{k|\hat{z}_k \geq \overline{z}} L_k^0|\hat{z}_k - \overline{z}|\right]}$ and $\Lambda \equiv \frac{\mathbb{E}\left[\sum_{\ell|\hat{z}_\ell \geq \overline{z}} L_\ell^0(\hat{z}_\ell - \overline{z})\right]}{\mathbb{E}\left[\sum_\ell L_\ell^0(\hat{z}_\ell - \overline{z})^2\right]}$. Each sum in this expression for β is the weighted-average of the effects of all shocks in the set of places with either above-average or below-average shocks, with weights given by v_ℓ, which add to one in each group of locations. β is the difference in these average effects scaled by Λ, which places β on a per-unit-of-$\hat{z}$ scale.

and when a counterfactual shock to a single location would drive substantial migration. We learn little about interregional mobility costs or migration elasticities from an estimate of β and should therefore not use results from conventional migration regressions to justify model restrictions or to calibrate model parameters.

Confounding migration in wage regressions. Some studies estimate the conventional migration regression in (22) not to draw conclusions about the effects of shocks on local populations, but to assess whether interregional migration is likely to substantially affect estimates of the effect of local shocks on local economic outcomes such as wages. Appendix A.8 examines this approach in light of the model, finding that when the conventional regression estimate of β is close to zero, wage regression estimates will experience minimal confounding from migration, irrespective of the mechanism driving the small estimate of β. Intuitively, because labor demand shocks have a direct effect on local wages and an indirect effect through migration spillovers, when $\beta \approx 0$ this indirect effect is approximately uncorrelated with the direct effect and therefore does not substantially bias the estimate of the local shock's effect on the local wage.

3.4 Model-Consistent Specifications

We now describe alternative empirical approaches that integrate information on regional shocks and pre-shock migration patterns to estimate migration responses in a way that is consistent with the model.

The model-implied relationship between local population changes and the vector of local labor demand shocks is given by equation (20) above:

$$\hat{\mathbf{L}} = \Omega(\theta/\sigma)\hat{\mathbf{z}} + \boldsymbol{\zeta}_2, \qquad \text{where } \Omega\left(\theta/\sigma\right) \equiv \mathbb{I} - \left(\mathbb{I} + \frac{\theta}{\sigma}\left(\mathbb{I} - \Gamma'\Pi\right)\right)^{-1}. \tag{28}$$

Following Section 2.5, we observe the population change $\hat{\mathbf{L}}$, the migration matrices Γ and Π, and the shocks $\hat{\mathbf{z}}$, so we can estimate the parameter θ/σ using non-linear least squares (NLLS).[36] Formally, in Appendix A.10 we show that whenever $\hat{\mathbf{z}}$ has a non-degenerate variance-covariance matrix, θ/σ

[36]We recommend an intercept is included in equation (28) in practice to capture exogenous changes in national population (due to population growth or international migration) and to increase estimation efficiency.

uniquely solves the NLLS problem[37]

$$\min_{\lambda} \mathbb{E}\left[\left(\hat{\mathbf{L}} - \Omega(\lambda)\hat{\mathbf{z}}\right)' \left(\hat{\mathbf{L}} - \Omega(\lambda)\hat{\mathbf{z}}\right)\right]. \tag{29}$$

Given the resulting estimate of θ/σ we can calculate shock-induced population changes in each location using the same relationship in (28), omitting the error term. This process can be used to predict internally valid effects of the observed shocks used to estimate θ/σ in (28) and externally valid effects of counterfactual shocks, under the model's assumptions.

An alternative approach relies on the low-mobility approximation in (21) to estimate a specification like the conventional migration regression in (22), but using the model-consistent independent variable, $\frac{M_\ell}{L_\ell}(\hat{z}_\ell - \hat{z}_{-\ell})$. This measure captures local heterogeneity in migration intensity (M_ℓ/L_ℓ) and incorporates information on shocks to other locations with strong migrant connections ($\hat{z}_{-\ell}$).[38] While relying on the same information as the NLLS procedure just discussed, this approach maintains a specification that is linear in the parameters and can be estimated by OLS.

Another alternative is to change the unit of analysis to the location pair and study the effects of observed shocks changes in migration flows between pairs of locations.[39] In Appendix A.13, we show that the following specification is consistent with the model under the low-mobility approximation in (21),

$$\hat{f}_{o\ell} = \alpha + \beta_1 \hat{z}_\ell + \beta_2 \hat{z}_o + \varepsilon_{o\ell}, \tag{30}$$

where $f_{o\ell}$ is the migration flow from origin o to destination ℓ where $o \neq \ell$, and the regression coefficients $\beta_1 = -\beta_2 = \theta/\sigma$. One can also recover model-consistent estimates of the local population changes driven by the shocks $\hat{\mathbf{z}}$ by aggregating the changes in migration flows estimated using (30)

[37]Borusyak and Hull (2021, Appendix D.5) show in a similar setting that a "recentering" adjustment is generally necessary for the appropriate moment condition to hold under the as-good-as-random assignment of the shocks, $\mathbb{E}[\hat{z}_\ell \mid \boldsymbol{\zeta}_1] = \mu$. We show in Appendix A.9 that this issue does not arise in our setting because NLLS estimates are invariant to μ: demeaning the shock has no implications for population changes. Recentering would be necessary, however, if the shocks were *conditionally* as-good-as-randomly assigned: e.g. higher in regions with a higher manufacturing employment share, as in many empirical settings (e.g. Autor et al. (2013)).

[38]In Section 4, we generalize the model to include both location and industry switching costs and define a parallel linear specification capturing the effects of both shocks to ℓ and those in other locations in equation (33).

[39]See Sprung-Keyser et al. (2022) for an example of this approach.

up to the single-location level:

$$\hat{L}_\ell = \sum_{k \neq \ell} \gamma_{k\ell} \hat{f}_{k\ell} - \sum_{k \neq \ell} \pi_{\ell k} \hat{f}_{\ell k} \approx \frac{2\theta}{\sigma} \cdot \frac{M_\ell^0}{L_\ell^0} \cdot (\hat{z}_\ell - \hat{z}_{-\ell}) + \zeta_{2\ell} \tag{31}$$

This approach again relies on the same information as NLLS but can be estimated by OLS.

Two caveats should be kept in mind here. First, we assume that the researcher has available a set of shocks to local labor demand that are as good as randomly assigned (as in (19)) with which to estimate θ/σ. Any endogeneity in the shocks will lead to inconsistency in both the conventional migration regression and our proposed approaches. Second, our model is purposefully stylized in order to focus on the conventional migration regression's failure to account for shocks to other migrant-connected locations. The model does not include other mechanisms such as agglomeration economies, housing, capital markets, and forward-looking behavior. Instead, our relatively parsimonious model focuses on the cross-location spillover problem faced by conventional migration regressions.

4 Full Model with Regional and Industry Frictions

The baseline model described in the preceding section allows for frictions in moving across locations, but assumes costless mobility across industries within a given location. Here, we examine a more realistic setting in which workers face costly mobility across both industries and locations. We do so by applying the same analysis at the location-industry level. In particular, the indirect utility in (4) can be written identically, but with the subscript o referring to the worker's initial location-industry pair and ℓ referring to the location-industry pair that the worker might choose. It will be helpful to make this distinction explicit by defining the initial location and industry as o and p and the new location and industry as d and q. To clarify the distinction between prior and current location-industry information, we separate the subscripts with a comma. Workers face a moving cost $\tau_{op,dq}$ when switching from location-industry pair op to pair dq, and θ determines the responsiveness of labor supply to wage differences between one location-industry pair and another. Note that, for simplicity, this setup assumes that the same parameter θ applies to both the location and industry dimensions.

The model extends directly to the location-industry context, and we again make use of the low-mobility approximation to derive an intuitive expression relating labor demand shocks to population changes at the location-industry level.[40] Applying (17) and (18) in that setting yields

$$\hat{L}_{\ell n} \approx \frac{\theta}{\sigma}\left(\sum_{o\in\mathscr{L}}\sum_{p\in\mathscr{N}}\gamma_{op,\ell n}(\hat{z}_{\ell n}-\hat{z}_{op})+\sum_{d\in\mathscr{L}}\sum_{q\in\mathscr{N}}\pi_{\ell n,dq}(\hat{z}_{\ell n}-\hat{z}_{dq})\right)+\zeta_{\ell n}, \tag{32}$$

where $\gamma_{op,\ell n} = f_{op,\ell n}/L_{\ell n}$ is the share of workers in location-industry cell ℓn who came from cell op and $\pi_{\ell n,dq} = f_{\ell n,dq}/L^0_{\ell n}$ is the share of workers in cell ℓn who went to cell dq (both in the no-shock counterfactual), the $\hat{z}_{\ell n}$ are observable labor demand shocks in cell ℓn, and $\boldsymbol{\zeta} = (\zeta_{\ell n})$ is an error term such that $\mathbb{E}[\hat{z}_{\ell n} \mid \boldsymbol{\zeta}] = \mu,\ \forall \ell, n$, as in (19).

This model allows us to interpret the conventional location-level migration regression when the data reflect a setting with frictions across both industries and locations. The conventional regression remains identical to equation (22), where we aggregate across industries to yield the location's average labor demand shock $\hat{z}_\ell \equiv \sum_n \frac{L^0_{\ell n}}{L^0_\ell}\hat{z}_{\ell n}$ and the location's change in population $\hat{L}_\ell = \sum_n \frac{L_{\ell n}}{L_\ell}\hat{L}_{\ell n}$. Applying the latter to (32) and rearranging yields

$$\hat{L}_\ell \approx \frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}(\hat{z}^{\text{mov}}_\ell - \hat{z}^{\text{mov}}_{-\ell}) + \zeta_\ell, \tag{33}$$

$$\text{where}\quad \hat{z}^{\text{mov}}_\ell \equiv \sum_n \frac{F_{-\ell,\ell n}}{M_\ell}\hat{z}_{\ell n},\quad \hat{z}^{\text{mov}}_{-\ell} \equiv \sum_p\sum_{o\neq\ell}\frac{F_{op,\ell}}{M_\ell}\hat{z}_{op}, \tag{34}$$

and $\zeta_\ell = \sum_n \frac{L_{\ell n}}{L_n}\zeta_{\ell n}$ (see Appendix A.11 for the proof). In this expression, $F_{-\ell,\ell n} \equiv \sum_p\sum_{o\neq\ell}F_{op,\ell n}$ is the migration flow between location-industry ℓn and other locations (recall that $F_{op,\ell n} \equiv \frac{1}{2}(f_{op,\ell n} + f_{\ell n,op})$ is the average of flows in both directions), and $F_{op,\ell} \equiv \sum_n F_{op,\ell n}$ is the flow between location-industry op and location ℓ, all referring to the no-shock counterfactual. The terms $\hat{z}^{\text{mov}}_\ell$ and $\hat{z}^{\text{mov}}_{-\ell}$ reflect shocks facing industries in location ℓ and those in other locations, respectively. Both are weighted averages with weights depending on location ℓ's migrant connections to other locations, hence the superscripts referring to "movers."

Equation (33) shows that the local population change depends upon how local shocks compare to outside-option shocks to other locations, as in the baseline model. In fact, in the case of a

[40] Note that the low-mobility approximation may perform more poorly in the location-industry context, since there may be substantial mobility across industries within location. However, our quantitative investigation in Section 6 suggests this is not the case in our context.

purely regional shock, i.e. $\hat{z}_{\ell n} = \hat{x}_\ell \ \forall \ell, n$, the two models imply identical population responses under the low-mobility approximation, since $\hat{z}_\ell^{\text{mov}} = \hat{z}_\ell$ and $\hat{z}_{-\ell}^{\text{mov}} = \hat{z}_{-\ell}$. With regional shocks, (33) is therefore equivalent to (21). In other words, when the shock is purely regional, the presence of industry frictions does not alter migration behavior, and all of the problems of interpretation discussed in Section 3.3 and the solutions described in Section 3.4 apply to the model with location and industry frictions.

Now consider shocks that vary across industries. In a regional migration analysis, researchers typically incorporate industry-level shocks using a shift-share structure, in which the average shock facing workers in location ℓ is an employment-weighted average of industry shocks: $\hat{z}_\ell = \sum_n \frac{L^0_{\ell n}}{L^0_\ell} \hat{x}_n$. When switching industries is costless, as in the baseline model of Section 2, this shift-share measure captures the regional labor demand shock facing all workers in location ℓ, i.e. the shift-share measure *is* the regional labor demand shock when there are no industry switching frictions. In contrast, when industry switching frictions are present, as in the full model developed in this section, workers in the same location but in different industries experience different labor demand changes and have different incentives to migrate. We therefore find distinct migration behavior in models with and without industry frictions when workers face industry-level labor demand shocks.

We can see this distinction in (33) with purely industry-level shocks, i.e. $\hat{z}_{\ell n} = \hat{x}_n \ \forall \ell, n$. The term for shocks facing location ℓ becomes $\hat{z}_\ell^{\text{mov}} = \sum_n \frac{F_{-\ell,\ell n}}{M_\ell} \hat{x}_n$, which places more weight on shocks to local industries with stronger migrant connections to outside locations. These outside migrant connections are stronger for industries in which migration costs are relatively low, so a larger fraction of workers in these industries are close to indifferent between locations, and a positive (negative) shock to that industry will lead to more net in-migration (out-migration) among these marginal workers. Conversely, the outside-option term becomes $\hat{z}_{-\ell}^{\text{mov}} \equiv \sum_p \frac{F_{(-\ell)p,\ell}}{M_\ell} \hat{x}_p$, where $F_{(-\ell)p,\ell}$ is the migrant flow between all industries in location ℓ and industry p in other locations. This migrant connection is stronger when workers in ℓ can transition into industry p in other locations relatively easily, so a more positive (negative) shock to that industry will lead to more net out-migration from (net in-migration to) location ℓ.

To illustrate the implications of industry frictions for migration behavior, 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, although regional shocks will induce migration, pure industry shocks will not. Because workers cannot change industries, those migrating to or from ℓ will have the same industry mix as those in ℓ. This means that $F_{-\ell,\ell n} = F_{(-\ell)n,\ell} = F_{(-\ell)n,\ell n}$, which implies that $\hat{z}_\ell^{\text{mov}} = \hat{z}_{-\ell}^{\text{mov}}$ and $\hat{L}_\ell = 0$ in (33). When workers cannot switch industries, they face the same industry-level shock irrespective of their location and thus have no incentive to migrate in response to pure industry shocks.

This extreme example illustrates the more general point that industry-switching frictions reduce migration responses when labor demand shocks have an national industry-specific component. Intuitively, if switching industries is difficult, workers see less benefit in moving to a location with a more favorably affected industry mix because they would still face much of their original industry's shock in the new location.

We can formalize this intuition by characterizing β, the conventional migration regression estimand in (24), in the context of the full model with location and industry mobility costs. Note that we lead with the case of heteroskedastic shocks because shift-share variables are inherently heteroskedastic.

Theorem 2. *Suppose the data are generated by the low-mobility approximation to the full model in which workers face mobility costs across both locations and industries (33). Then*

$$\beta = \frac{2\theta}{\sigma} \cdot \frac{M^v/L}{\widetilde{M}^v/L} \cdot \frac{\rho^{mov} - \rho}{1 - \widetilde{\rho}} \tag{35}$$

where $v_\ell \equiv \operatorname{Var}[\hat{z}_\ell]$ is the variance of the local shock to ℓ, $M^v/L \equiv (\sum_\ell M_\ell v_\ell)/(\sum_\ell L_\ell v_\ell)$ is the shock-variance weighted national average migration share, and $\rho \equiv (\sum_\ell M_\ell v_\ell \rho_\ell)/(\sum_\ell M_\ell v_\ell)$ for $\rho_\ell \equiv \operatorname{Cov}\left[\hat{z}_\ell, \hat{z}_{-\ell}^{mov}\right]/v_\ell$. As in Theorem 1, tildes refer to the costless-migration setting, so $\widetilde{M}^v/L \equiv (\sum_\ell \widetilde{M}_\ell v_\ell)/(\sum_\ell L_\ell v_\ell)$ and $\widetilde{\rho} \equiv (\sum_\ell \widetilde{M}_\ell v_\ell \widetilde{\rho}_\ell)/(\sum_\ell \widetilde{M}_\ell v_\ell)$ for $\widetilde{\rho}_\ell \equiv \operatorname{Cov}[\hat{z}_\ell, \widetilde{z}_{-\ell}]/v_\ell$ and $\tilde{z}_{-\ell} \equiv \sum_{d \neq \ell} (\frac{1}{2}(\tilde{f}_{\ell d} + \tilde{f}_{d\ell})/\widetilde{M}_\ell)\hat{z}_d$. Finally, $\rho^{mov} \equiv (\sum_\ell M_\ell v_\ell \rho_\ell^{mov})/(\sum_\ell M_\ell v_\ell)$, and $\rho_\ell^{mov} \equiv \operatorname{Cov}\left[\hat{z}_\ell^{mov}, \hat{z}_\ell\right]/v_\ell$.

The proof is given in Appendix A.11. As in the baseline model, β can be multiplicatively decomposed into the ratio of structural elasticities, a term depending on the national migration share (which differs from Theorem 1 in the baseline model only due to heteroskedasticity), and the attenuation factor $\frac{\rho^{\text{mov}} - \rho}{1 - \widetilde{\rho}}$. This attenuation factor captures two reasons for attenuation in the conventional migration regression estimate. As in the baseline model, ρ is larger when the shock to ℓ, $\hat{z}_\ell$, is more strongly correlated with the shocks to migrant-connected locations, $\hat{z}_{-\ell}^{\text{mov}}$. As this correlation grows, β exhibits more attenuation. In the presence of industry switching frictions, the conventional regression suffers from additional misspecification by using $\hat{z}_\ell$ as the direct shock measure rather than the model-consistent $\hat{z}_\ell^{\text{mov}}$. The more $\hat{z}_\ell^{\text{mov}}$ deviates from $\hat{z}_\ell$, the smaller is ρ^{mov} and the more β exhibits attenuation. Returning to the example of industry-level shocks with limited cross-industry mobility, workers moving between ℓ and other regions will have similar industry composition, and thus $\hat{z}_\ell^{\text{mov}}$ is similar to $\hat{z}_{-\ell}^{\text{mov}}$. This implies that $\rho^{\text{mov}} \approx \rho$, and attenuation is severe.

This discussion reveals how industry frictions undermine the intuitive appeal of measuring local labor demand shocks using shift-share measures. Only when industry switching frictions are absent does the industry shift-share represent a true regional labor demand shock facing all local workers. With industry frictions, workers in different industries 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, exacerbating the biases examined in Section 3.3.

5 Data and Descriptive Statistics

To assess the quantitative importance of the conceptual problems with the conventional migration regression that we have discussed in Sections 2 through 4, we use data on observed worker transitions across locations and industries. This section describes the data source and presents descriptive statistics relevant to our quantitative analysis.

5.1 Data

The NLLS approach described in Section 3.4 and the model simulations presented in Section 6 require information on worker transitions across location-industry pairs, 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 locations and industries. Subsequent revisions will present results using similar data in other countries.

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 locations (De Negri et al., 2001; Saboia and Tolipan, 1985).[41] Locations in our analysis are based on the "microregion" definition of the Brazilian Statistical Agency (IBGE, 2002), which combines economically integrated contiguous municipalities (counties) with similar geographic and productive characteristics. We aggregate microregions slightly to ensure consistent boundaries over time, following Dix-Carneiro and Kovak (2017), which yields 486 time-consistent microregion locations in Brazil's 27 states.[42] Industries are based on the "Subsetor IBGE" classification the RAIS dataset uses to identify the industry associated with each employer. This classification 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.[43] We also drop individuals with missing or inconsistent information, including "other/ignored" industries, contradictory education levels across jobs in the same year, simultaneously holding jobs in very distant geographic areas, etc. If a worker holds multiple jobs, we select the job with the highest December earnings and use it to assign the worker's location and industry of employment.

The RAIS data allow us to observe formal employment in location-industry cells and transitions between them, but do not provide any information on non-employed or informally employed individuals. In our model, which does not include non-employment or informality, population and formal employment are indistinguishable.[44] We therefore use formal employment as a proxy for population and use RAIS data to calculate formally-employed worker flows between location-industry pairs, omitting transitions into or out of formal employment. Because we use detailed geographic and industry information, workers may transition between 12,150 = 486 $\times$ 25 location-industry cells. We calculate 5-year transitions and average observed flows for each pair of location-industry cells from 1999/2004 to 2005/2010 to reduce noise. We use these averages to calculate migration flows,

[41] Formality is defined as having a signed work card (*carteira assinada*), which registers the worker's contract with the Ministry of Labor and gives them the right to benefits and protections under the legal employment system. The restriction to formal workers is important in the Brazilian context, as informality rates measured in the Brazilian Census exceed 50% during our sample period (Dix-Carneiro and Kovak, 2019).

[42] As Dix-Carneiro and Kovak (2017) report using Census data, only 3.4 percent of individuals lived and worked in different microregions in the year 2000, so the microregion of employment is a quite accurate measure of both employment and residence location.

[43] RAIS reports earnings for December and average monthly earnings during employed months in the reference year. We use December earnings to ensure that our results are not influenced by seasonal variation or month-to-month inflation.

[44] A simple extension of the model would treat informal jobs and unemployment as belonging to another fictitious sector that can be chosen by the workers.

employment levels, and employment changes both at the location and location-industry levels.[45]

5.2 Descriptive Evidence

We now investigate aspects of the transition matrices, Γ and Π, that are relevant to the population effects of local labor demand shocks. Section 3.3 shows that the conventional migration regression estimate β omits two important elements that drive the population effects of local labor demand shocks in the model: heterogeneity in local migration shares and differences in shocks to migrant-connected locations. This section therefore presents summary statistics on mobility rates and the concentration of migrant connections across locations and industries.

The conventional migration regression estimate β depends upon the national average mobility rate (M/L), which is shown in the first row of Table 1: 12.2 percent of workers migrate across locations (microregions) over a 5-year period 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 locations within states than across states, which reflects a lower likelihood of long-distance moves. Industry transitions (third row) are substantially more likely than migration events (23.5 percent), and 6.8 percent of workers transition between both locations and industries in a given 5-year period.

Underlying these national average mobility rates is substantial heterogeneity in migration shares across locations. Figure 1 shows a histogram of M_ℓ/L_ℓ, the share of individuals in ℓ who move to a different location in a typical year (weighted by the region's initial employment). 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 (as in (21)), but this heterogeneity is omitted from the estimate of β in the conventional migration regression (as in (25)).[46]

We next show that migration exhibits a network structure. The analyses in Sections 3.3 and 4 emphasized the role of spillovers between connected locations and location-industry cells in driving migration responses to labor demand shocks. These spillovers will be particularly important in practice if location or industry connections differ substantially for workers in different locations or

[45] Specifically, define $f_{op,dq}^{t,t+5}$ as the observed number of workers moving from cell op to cell dq between years t and $t+5$. Then calculate the average across years as $\overline{f}_{op,dq} \equiv \frac{1}{7}\sum_{t=1999}^{2005} f_{op,dq}^{t,t+5}$. These average flows are then used to measure the initial and final employment levels, local industry compositions, and transition matrices.

[46] We also find substantial heterogeneity across industries in the probability of switching industries, with rates ranging from 4.5% to 45%.

Table 1: Average Mobility Rates, %

Across	Mobility (%)
Locations	12.2
States	4.0
Industries	23.5
Locations and Industries	6.8

Notes: Percent of individuals with 5-year transitions between the cells defined in the first column, calculated from average 5-year flows between location-industry cells in RAIS data from 1999 to 2010 as described in footnote 45. The last row measures the fraction of workers who change *both* location and industry in the same year.

industries, e.g. due to migration costs increasing in distance. To assess the degree to which these connections are concentrated, we calculate the Herfindahl–Hirschman index (HHI) measuring the average concentration of destinations and origins among those who migrated to or from a given region:

$$HHI_\ell = \sum_{k \neq \ell} \left(\frac{\frac{1}{2} \left(\pi_{\ell k} + \gamma_{k \ell} \right)}{1 - \frac{1}{2} \left(\pi_{\ell \ell} + \gamma_{\ell \ell} \right)} \right)^2 . \tag{36}$$

We also calculate a parallel expression for those who switched to or from industry n, HHI_n. The distributions of these concentration measures are shown in Figure 2. The average HHI_ℓ across locations is 0.116 (s.d. 0.102), which is equivalent to equally-weighted connections to only 9 other locations out of 485 possibilities. Industry transitions are relatively more dispersed, with an average HHI_n across industries of 0.109 (s.d. 0.050), which is equivalent to equally-weighted connections to 9 other industries out of 24 possibilities. Figure 2 also documents the variation in these concentration measures across both locations and industries. The substantial concentration in connections and its heterogeneity across locations suggest that workers in different locations may face substantial differences in spillovers from connected locations and industries.

6 Quantitative Assessment of Migration Regressions

In this section we examine the predictive accuracy of the conventional migration regression in the context of Brazilian geography, industrial structure, and baseline migration patterns. We simulate migration behavior based on the baseline model with frictions across locations only and on the full model with frictions across both locations and industries. We consider population responses to labor demand shocks that follow a variety of data generating processes, including the real-world

Figure 1: Geographic Mobility Rates Across Regions

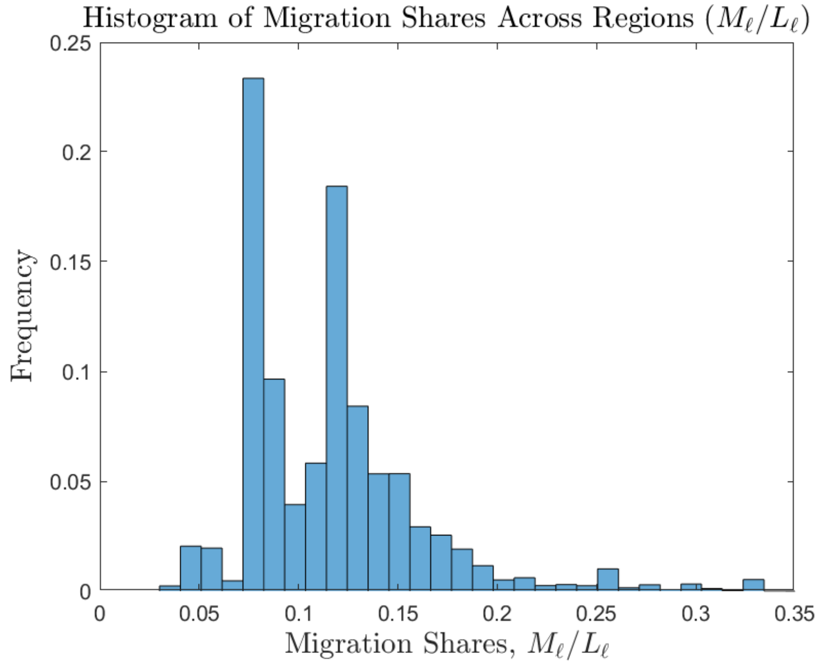

Notes: Percent of individuals migrating across locations (microregions), calculated from average 5-year flows between location-industry cells in RAIS data from 1999 to 2010 as described in footnote 45. The histogram is weighted by each location's initial population.

tariff shocks Brazilian industries faced during the country's 1990 trade liberalization, and compare these true model-generated responses to those inferred from conventional migration regression estimates. We further compare the conventional regression with the model-consistent NLLS and other regression-based procedures discussed in Section 3.4.

6.1 Simulation Procedure and Calibration

To generate model-simulated migration behavior, we must first specify the distribution of labor demand shocks and calibrate the relevant parameter ratio, θ/σ. We consider four data-generating processes for shocks. The first set of shocks consists of independent standard normal draws across locations, i.e. $\hat{x}_\ell \sim N(0,1)$ and $\hat{z}_{\ell n} = \hat{x}_\ell$. We refer to these shocks as "*iid* across locations." 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 $\hat{x}_s \sim N(0, \sigma_2^2)$ and $\hat{z}_{\ell n} = \hat{x}_{s(\ell)}$, where $s(\ell)$ is the state containing location ℓ. We refer to these shocks as "*iid* across states." Third, we generate independently and normally

Figure 2: Concentration of Locations and Industries Among Switchers

Panel (a): Across Locations Panel (b): Across Industries

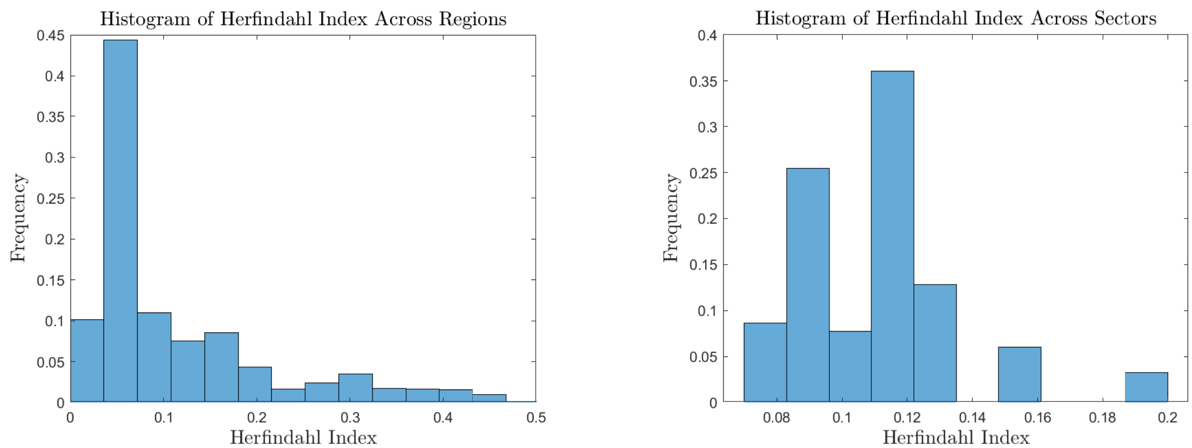

Notes: Panel (a) shows the distribution of HHI indexes measuring the concentration of sources and destinations for those switching locations, as defined in (36), weighted by the initial regional employment. Panel (b) similarly shows the distribution of HHI indexes across industries.

distributed shocks across industries, such that $\hat{x}_n \sim N(0, \sigma_3^2)$, and $\hat{z}_{\ell n} = \hat{x}_n$. We refer to these shocks as "*iid* across industries." Finally, we use the actual tariff changes facing Brazilian industries during the country's 1990 trade liberalization to calculate $\hat{x}_n$. These real-world shocks allow us to assess the results for shocks with a realistic correlation structure across industries, as a complement to the *iid* process.[47] For all four data generating processes, we calculate the regional shock $\hat{z}_\ell$ from the region-industry shocks $\hat{z}_{\ell n}$ using the weighted-average measure $\hat{z}_\ell = \sum_n \frac{L^0_{\ell n}}{L^0_\ell} \hat{z}_{\ell n}$. To ensure comparable magnitudes across the shocks, we scale each shock measure so that the expected cross-location variance of shocks is equal to one.[48]

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 locations, 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 $\frac{\theta}{\sigma} = 0.4075$.

Given the relevant vector of shocks and calibrated value of θ/σ, we generate the change in employment based on $\hat{\mathbf{L}} = \Omega \hat{\mathbf{z}}$ (i.e. (20) without the error term reflecting unobserved shocks, which

[47] Specifically, we let $\hat{x}_n = d\ln(1 + \tau_n)$, using the tariff measures from Kovak (2013) and Dix-Carneiro and Kovak (2017) and then rescale the resulting shift-share shocks to have variance equal to one across locations.

[48] Specifically, $\sigma_2 = 1.070$, $\sigma_3 = 5.454$, and the Brazilian tariff shocks are scaled by a factor 52.64.

we return to below).[49] Recall that Ω depends on the observed flow matrices Γ and Π, which we measure using the RAIS data (Section 5.1). To investigate the effects of industry switching frictions, we generate versions of $\hat{\mathbf{L}}$ from both the baseline model with regional frictions alone (Section 2) and the full version with industry-region frictions (Section 4), using the same structural parameters.

We estimate the conventional location-level migration regression using the model-implied employment changes as the dependent variable and the direct regional shocks as the independent variable, weighting the regression by the initial population. We then save the estimate of β and R^2 for each simulation. For ease of interpretation, we report the estimate of β relative to the setting with *iid* shocks across locations, since in that case the conventional regression does not suffer from any attenuation. We also calculate what we call the "reallocation index," defined as the population reallocated by the shocks based on the conventional regression's prediction, relative to the true population that was reallocated by the shocks:

$$\text{Reallocation index} \equiv \frac{\frac{1}{2}\sum_\ell L_\ell \left|\hat{\beta}\left(\hat{z}_\ell - \overline{z}\right)\right|}{\frac{1}{2}\sum_\ell L_\ell \left|\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\mathbf{0})\right|}. \tag{37}$$

In this expression, $\hat{\beta}(\hat{z}_\ell - \overline{z})$ is the proportional population change in ℓ due to observed shocks as predicted by conventional migration regression (see Section 3.3), while $\hat{L}_\ell(\hat{\mathbf{z}}) - \hat{L}_\ell(\mathbf{0})$ is the true population change (subtracting $\hat{L}_\ell(\mathbf{0})$ removes the effect of any unobserved shocks, which becomes relevant in Section 6.3). The absolute values avoid having population increases in one location cancel out decreases elsewhere (which is necessary on average with a fixed population), and the $\frac{1}{2}$ terms avoid double-counting these shifts across locations. We collect these statistics for each simulation and then repeat the entire process 500 times and report averages across the simulations, unless specified otherwise.

6.2 Assessing the Conventional Migration Regression

Table 2 presents simulation results assessing the performance of the conventional migration regression. Panel A reports conventional migration regression estimates when the data are generated from the model with location frictions only (as in Section 2). The first row reports estimates when

[49] Note that we do not impose the low-mobility assumption in the simulation, since we use (20) rather than (21). We do, however, maintain the assumption of small shocks, such that $\hat{x} \equiv dx/x = d\ln x$.

Table 2: Conventional Migration Regression Performance

Shocks	Relative β	Reallocation Index	R^2
Panel A: Model with location frictions only			
iid across locations	1.000	0.963	0.796
iid across states	0.419	0.984	0.824
iid across industries	0.810	0.904	0.733
industry tariff shock	0.794	0.864	0.717
Panel B: Model with location and industry frictions			
iid across locations	1.000	0.963	0.805
iid across states	0.426	0.985	0.828
iid across industries	0.520	0.800	0.552
industry tariff shock	0.468	0.789	0.625

Notes: Table reports conventional migration regression estimates β relative to the estimate under *iid* shocks across locations, the reallocation index defined in (37), and associated R^2, averaged across 500 simulations. See text for definitions of shock processes. Panel A simulates local employment changes using the model with cross-location frictions only (Section 2), while Panel B simulates employment changes using the full model with location-industry frictions (Section 4).

shocks are *iid* across locations, in which case there will be no attenuation, so $\rho = 0$ and $\frac{1-\rho}{1-\tilde{\rho}} = 1$. Since we report the estimate of β relative to this case, the relative estimate is 1 by definition. The reallocation index is close to 1, indicating that the conventional regression captures nearly all of the true population reallocation driven by the shocks. Yet, the misspecification of the conventional regression leads to an average R^2 across simulations of only 0.796, despite the absence of any unobserved shocks, such that model-consistent procedures would achieve $R^2 = 1$. This relatively poor predictive performance reflects the fact that β omits location-specific migration intensity $\frac{M_\ell}{L_\ell}$ and predicts the spillover effects from migrant-connected locations $(\hat{z}_{-\ell} - \overline{z})$ using only the direct shock facing the location $(\hat{z}_\ell - \overline{z})$, as discussed in Section 3.3.

When shocks are correlated between migrant-connected locations, ρ in (25) is greater than zero and the estimate of β falls, 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 locations when facing state-level shocks. As a result, the implied attenuation factor declines to 0.419 in the second row. Industry shift-share shocks in the third row similarly induce spatial correlation, but much less so than with state-level shocks, so the β estimate is only moderately below that for the *iid* location shocks. The real-world shocks facing Brazilian industries during trade liberalization yield similar results to those in the prior row, suggesting that the *iid* shocks across industries yield realistic implications.

While the model with location frictions alone is helpful for understanding the potential problems interpreting the conventional migration regression estimate of β, the model with both industry and location frictions is more realistic and exhibits stronger attenuation when shocks have an industry component. The simulation results from this full model appear in panel B of Table 2. When shocks are *iid* across locations, the conventional regression performs nearly identically to the case with cross-region frictions in panel A in terms of implied population reallocation and R^2. This close similarity reflects the finding in Section 4 that the two models are identical under the low-mobility approximation. However, when we introduce industry-level shocks in the last two rows of panel B, the estimates are much smaller in the full model than in the baseline model, with respective attenuation factors of 0.520 and 0.468, reflecting the additional attenuation induced by industry shocks in the presence of industry frictions. The conventional regression also understates the extent to which the shocks reallocated population between locations in response to industry-level shocks by around 20 percent due to the omission of variation across locations in the shocks to migrant-connected location-industry pairs and the misspecification of the effects of direct shocks (i.e. using $\hat{z}_\ell$ rather than $\hat{z}_\ell^{\text{mov}}$). With industry-level shocks, the regression fit is poor, with R^2 values below 63 percent, despite the fact that population growth depends solely upon the observed shocks.

Figure 3 visually presents the simulation results for a single representative simulation draw. The y-axes measure predicted population changes based on the conventional migration regression estimates, $\beta(\hat{z}_\ell - \bar{z})$, and the x-axes measure the true population change in the full model with industry-region frictions. Each point represents a particular location. Panel (a) shows population changes when facing *iid* shocks across locations, while Panel (b) considers *iid* shocks across industries. Note the difference in x-axis scale between the two panels, showing that there is less true reallocation in Panel (b) with industry shocks. This difference confirms the theoretical conclusion from Section 4 that industry-switching frictions reduce true migration responses when labor demand shocks have an industry-specific component.

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 predictions are quite inaccurate in many cases, consistent

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

Panel (a): *iid* shocks across locations　　Panel (b): *iid* shocks across industries

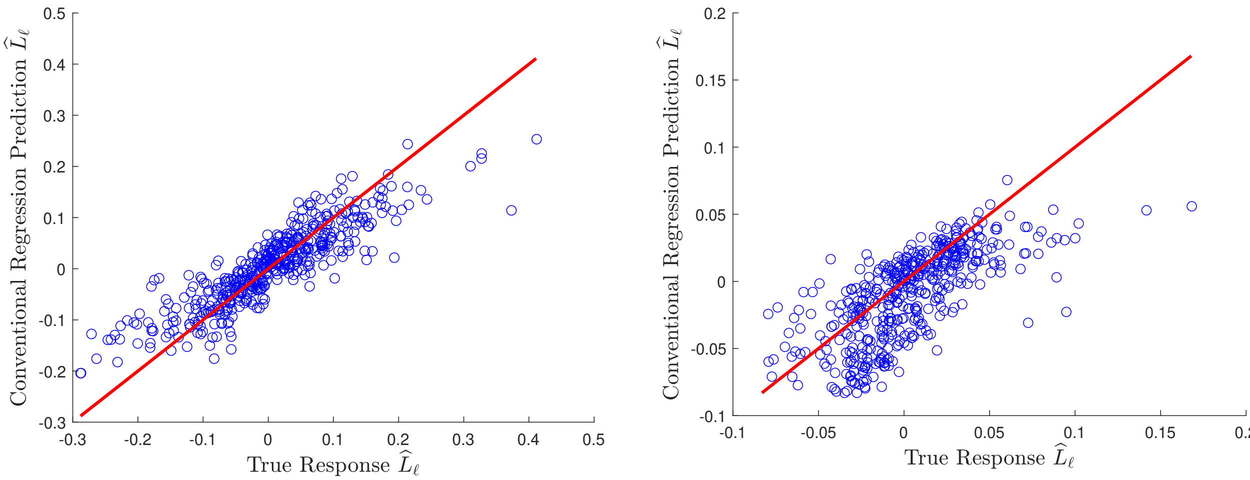

Notes: Scatter plots compare the predicted population response to the observed labor demand shocks against the true model-based response for a single representative simulation of the model with location-industry frictions (Section 4). Each point represents a location. The y-axis is the predicted population response based on the conventional migration regression, $\beta(\hat{z}_\ell - \bar{z})$. The x-axis is the true model-based population change in the model with location-industry frictions. The left panel shows the result for *iid* shocks across locations, and the right panel shows the result for *iid* shocks across industries. 45-degree lines are shown in red.

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. As a result, the regression-based predictions capture on average only 80 percent of the true population reallocation caused by the labor demand shocks, based on the reallocation index reported in the penultimate row of Table 2.

Together, the results in Table 2 and Figure 3 confirm the patterns predicted in our theoretical analysis and reveal quantitatively substantial attenuation when shocks are correlated across migrant-connected locations. 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 determined by the labor demand shocks alone, without additional noise. In other words, the conventional migration regression performs poorly in generating internally valid estimates of the effects of observed shocks.

Figure 4 illustrates that the estimate of β also cannot be interpreted as the effect of a counterfactual unit labor demand shock to a single location on that location's population. The regression-based prediction for this counterfactual is $\beta \times 1$, so the counterfactual prediction is the constant β for

Figure 4: Conventional Regression Predictions vs. True Responses for Counterfactual Unit Shocks

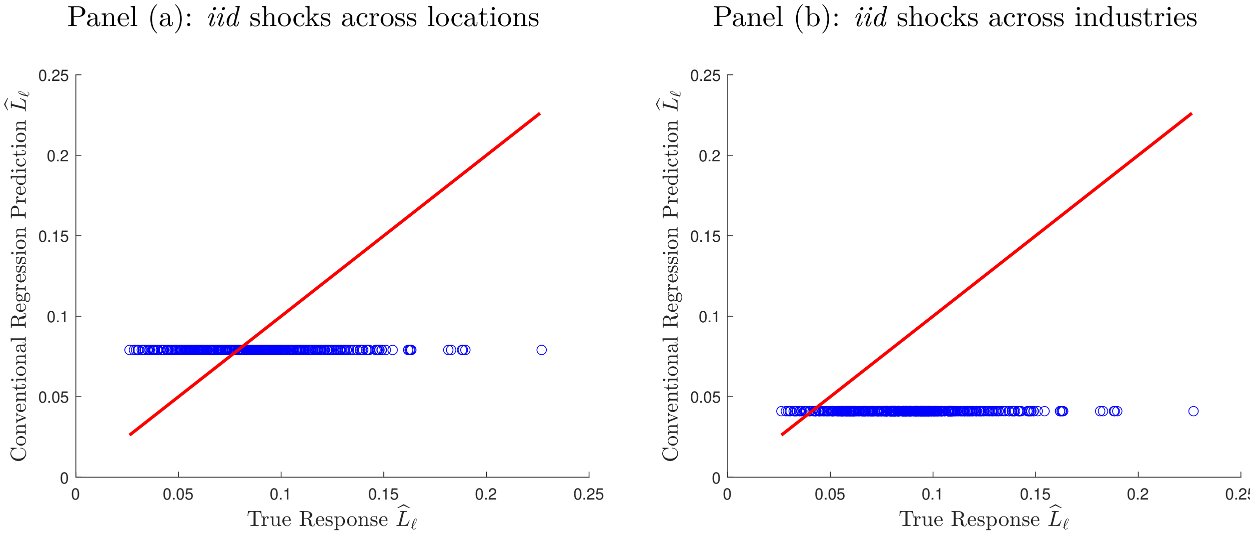

Notes: Scatter plots compare the predicted population response to a counterfactual unit shock to a single location against the true model-based response. Each point represents a location. The y-axis is the predicted population response based on the conventional migration regression estimate, β, averaged across simulations. The x-axis is the true model-based population change in the model with location-industry frictions (Section 4) when the relevant location faces a unit shock. The left panel shows the result when β is estimated using *iid* shocks across locations, and the right panel shows the result when β is estimated using *iid* shocks across industries. 45-degree lines are shown in red.

all locations. The figure shows the true model-based population change when individually shocking each location and plots the responses against the estimate of β. Panel (a) considers the average β estimate using *iid* regional shocks in the model with location-industry frictions. While the regression estimate based on *iid* regional shocks matches the mean response across locations, it omits significant heterogeneity in the magnitude of the true counterfactual responses in each location. The situation is worse when β is estimated using a regional shift-share measure based on *iid* industry shocks in Panel (b). Not only does the estimate omit all heterogeneity, but it is subject to severe attenuation due to the shock correlation across locations and the industry-switching frictions (as seen in the last two rows of Table 2). As a result, the regression-based counterfactual estimate substantially understates the true effect in nearly every location.

6.3 Introducing Noise

The preceding analyses determine the true population changes based on the model's structure and simulated labor demand shocks alone. In that setting, the model fits the data perfectly by construction, so we were unable to assess the quality of estimates based on the proposed NLLS procedure introduced in Section 3.4. Here, we add noise from unobserved changes in labor demand

and supply to the baseline simulations of Section 6.2. This addition allows us to assess the relative performance of the conventional migration regression vs. the proposed model-consistent procedures, and more closely mimics the data observed in a typical study of migration responses to local shocks. We assume the noise to be orthogonal to the observed labor demand shocks as in (19) to focus on issues that arise even in the best-case scenario with exogenous shocks.

We generate the true population changes at the location-industry level using the full model with location-industry frictions as $\hat{\mathbf{L}} = \Omega\hat{\mathbf{z}} + \boldsymbol{\zeta}$, where $\hat{\mathbf{z}}$ is the vector of observed shocks and $\boldsymbol{\zeta}$ captures the effects of unobservable shocks. The unobserved shocks in $\boldsymbol{\zeta} = (\zeta_{\ell n})$ consist of additive location, industry, and location-industry components, such that

$$\zeta_{\ell n} = \epsilon_\ell + \epsilon_n + \epsilon_{\ell n}.$$

To ensure a realistic correlation structure for this noise term, we take the observed 5-year changes in employment in each location-industry cell and regress them on location fixed effects and industry fixed effects, which allows us to measure the location-level variation, industry-level variation, and residual location-industry variation, which we use to calibrate the variance of each noise component.[50]

We focus on *iid* shocks across locations and *iid* shocks across industries and compare the in-sample fit of five methods for estimating the population effects of the simulated labor demand shocks. We first estimate β from the conventional regression (22) and, like before, take $\beta\left(\hat{z}_\ell - \bar{z}\right)$ as the predicted population growth in location ℓ. Second, we estimate θ/σ using the NLLS specification (28) based on the correctly specified model allowing for location-industry mobility frictions and estimated using location-industry observations. We then use that parameter estimate to generate predicted population changes in response to the observed shocks. While this approach has the benefit of perfectly matching the data generating process, its data requirements are substantial.

[50] Specifically, for each 5-year transition matrix from the historical data (without averaging across years), we estimate $\hat{L}_{\ell n} = \epsilon_\ell + \epsilon_n + \epsilon_{\ell n}$, where $\hat{L}_{\ell n}$ is the proportional change in the number of workers in location ℓ and industry n (averaged across years of our data), ϵ_ℓ are location fixed effects, ϵ_n are industry fixed effects, and $\epsilon_{\ell n}$ is the error term. We then estimate variances of each component and average them across the years. The employment-weighted standard deviation of the estimated location terms is 0.09, of the estimate industry terms is 0.08, and of the estimated residual is 0.18. We then simulate normally-distributed noise components as $\epsilon_\ell \sim \mathcal{N}(0, 0.09^2)$, $\epsilon_n \sim \mathcal{N}(0, 0.08^2)$, and $\epsilon_{\ell n} \sim \mathcal{N}(0, 0.18^2)$. Finally, we demean the resulting overall noise term $\zeta_{\ell n}$ so that $\sum_{\ell,n} \frac{L_{\ell n}}{L}\zeta_{\ell n} = 0$ in each simulation. In unreported results we additionally vary the magnitude of the unobserved shocks, finding qualitatively similar results to those presented here.

Specifically, in order to calculate the transition matrices Γ and Π in (28), one must have pre-shock data on transitions across location-industry pairs. Because such data are unavailable in some contexts, our third approach implements the NLLS procedure but uses a misspecified model that only allows for frictions across locations, as in the baseline model of Section 2. The advantage of this specification is that it is feasible without data on mobility across location-industry cells; it only requires pre-shock flows across locations, which are much more readily available in many countries.

We then consider two approaches that use simple linear regressions, as in the conventional migration regression, but incorporate information on potential spillovers from shocks to other locations and industries. Our fourth approach estimates a model-consistent regression based on the low-mobility approximation to the model, in which the dependent variable is the local population growth, $\hat{L}_\ell$, and the independent variable is $\frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}(\hat{z}_\ell^{\text{mov}} - \hat{z}_{-\ell}^{\text{mov}})$, following equation (33). Finally, the fifth approach seeks to reduce the data requirements even further than the misspecified NLLS approach by estimating the conventional regression with an additional control for shocks to other locations, weighted by size and distance:

$$\hat{L}_\ell = \alpha + \beta\hat{z}_\ell + \nu\hat{z}_{-\ell}^{\text{dist}} + \varepsilon_\ell, \qquad \text{where } \hat{z}_{-\ell}^{\text{dist}} \equiv \frac{\sum_{k\neq\ell}\frac{L_k}{(dist_{\ell k})^{\varsigma}}\widehat{z}_k}{\sum_{k\neq\ell}\frac{L_k}{(dist_{\ell k})^{\varsigma}}} \tag{38}$$

and $\varsigma = -1.25$ is the distance elasticity of migration, estimated using a standard gravity equation. This approach is similar to those used by Greenland et al. (2019) and Autor et al. (2021), and only requires information on location size and distance rather than pre-shock transitions. This approach allows us to assess whether distance and location size capture the information on transition costs reflected in pre-shock migration flows; if so, this approach will perform well, at least for the case of location-specific shocks.

We implement this simulation and estimation procedure across $S = 500$ simulations and use the Mean Square Error (MSE) as the measure of in-sample predictive quality for the response of regional population changes to shocks. For a given simulation s, let $\hat{L}_{\ell;s}^{\text{true}} \equiv \hat{L}_\ell(\hat{\mathbf{z}}_s) - \hat{L}_\ell(\mathbf{0})$ be the true response to shocks and let $\hat{L}_{\ell;s}^{\text{predicted}}$ be the predicted effect on local population growth using one of the five methods just described.

$$MSE = \frac{1}{S}\sum_{s=1}^{S}\left[\sum_{\ell}\frac{L_\ell}{L}\left(\hat{L}_{\ell;s}^{\text{true}} - \hat{L}_{\ell;s}^{\text{predicted}}\right)^2\right], \tag{39}$$

Table 3: MSE of Conventional Regression and NLLS Predictions

Prediction method	MSE, relative to uninformative prediction of $\hat{L}_\ell = 0$	
	iid location shocks (1)	*iid* industry shocks (2)
Conventional regression	0.244	0.514
NLLS, region-industry level	0.026	0.001
NLLS, misspecified regional model	0.033	0.358
OLS, low-mobility approximation	0.049	0.177
Distance-weighted average of $\hat{z}_k$	0.229	0.666

Notes: Each row corresponds to a method of predicting the effect of observed shocks on regional population, while each column defines the shock process; see the text for details. Each cell reports the ratio of the mean squared error (MSE) for a given prediction method to the MSE from the uninformative prediction, computed via (39) and (40), respectively. All simulations are based on the model with location-industry frictions (Section 4).

which averages squared prediction errors both across locations (with population weights) and across the simulations. We rescale this MSE measure relative to its value for the uninformed prediction, $\hat{L}_{\ell;s}^{\text{predicted}} = 0$, i.e. by the variance of the true population changes driven by the shocks, averaged across simulations:

$$MSE_{\text{uninformative}} = \frac{1}{S} \sum_{s=1}^{S} \left[\sum_{\ell=1}^{R} \frac{L_\ell}{L} \left(\hat{L}_{\ell;s}^{\text{true}} \right)^2 \right]. \tag{40}$$

Table 3 reports the results for our five methods and for *iid* location and *iid* industry shock processes. The first row confirms that predictions by the conventional regression are poor: the MSE of its predictions is 24 percent of the uninformative prediction for *iid* location shocks and 51 percent for *iid* industry shocks. The second row confirms that the NLLS procedure based on the true model performs extremely well, with prediction errors approaching zero for both regional and industry-level shocks. Things are quite different, however, in the third row when using NLLS based on the misspecified model with regional mobility costs but costless industry mobility. When shocks are truly location-specific, this approach performs extremely well — nearly as well as NLLS based upon the correctly specified model. In contrast, when shocks are industry-specific, this approach performs terribly, with relative MSE approaching that of the conventional regression. This finding shows that, while the data requirements to implement the baseline model with regional mobility costs are quite modest (region transitions only), its performance is very poor when the data reflect a situation with costs of moving across industries and an industrial component to the shocks.

The fourth row of Table 3 is based on predictions from the location-level regression following equation (33). This approach is consistent with the model under the low-mobility approximation, so this row can be seen as assessing the quality of that approximation in the Brazilian context. With *iid* location shocks, this approach performs very well, with prediction error only 4.9 percent that of the uninformative prediction. The performance is a bit worse with *iid* industry shocks, but the relative MSE in this case is still only one third that of the conventional migration regression and much better than the misspecified NLLS approach. Finally, the fifth row of Table 3 implements the approach controlling for the distance-weighted average of shocks to other locations, described in (38). As is clear in the table, this approach performs poorly even with *iid* regional shocks, approaching the MSE of the conventional regression without any attempt to address potential spillovers. This result implies that distance and location size do not sufficiently capture the information on heterogeneous mobility costs reflected in pre-shock transitions across locations.

Taken together, these results yield two lessons. First, accounting for the spillovers via the model-consistent NLLS procedure allows one to capture the population responses of local labor demand shocks even in the presence of noise in the data, which the conventional regression is not capable of. Second, ignoring the industry dimension of the data is relatively innocuous when the shocks are purely region-specific. However, when the shocks are industry-driven, as in with shift-share shocks commonly used in the literature, it is crucial to take the structure of industry mobility into account, in using the appropriate model as well as estimating the structural parameter at the disaggregated location-industry level.

7 Conclusion

This paper has examined how to interpret conventional migration regression estimates relating changes in local population or employment to observed local labor demand shocks. These estimates often yield the seemingly puzzling finding that workers do not migrate in response to local labor demand shocks despite large effects on other outcomes, such as local wages. Using a simple model of local labor markets with mobility costs, we find that this puzzling conclusion may be unwarranted.

Analytical results and quantitative simulations show that the conventional migration regression provides inaccurate estimates of both the local population effects of shocks to individual locations and the within-sample effects of the entire vector of observed shocks (a distinction that is often overlooked in the literature). These issues are particularly severe when the local shocks are constructed as shift-share aggregates of national industry shocks, which is a common approach. These findings suggest that population may have responded substantially to the shocks studied in prior work, even when the associated conventional migration regression estimates were small.

Going forward, we recommend that researchers adopt methods that account for spillovers of shocks between connected locations, as in our proposed NLLS and linear regression procedures. Yet, to implement these and related analyses, the data demands are significant. Our analysis using Brazilian data shows the substantial practical benefit of observing transitions between relatively detailed location-industry cells to successfully capture the interrelationships in a large and diverse economy.

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 spatially correlated 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 spatially correlated, or there may be a growing share of the variation in economic shocks occurring at the level of national industries rather than locations. Our focus on the network structure of worker reallocation can also bear fruit beyond the spatial context. For example, to the best of our knowledge, studies of the economic impacts of import competition at the industry level (e.g. Bernard et al. 2006 and Autor et al. 2014) have not considered the potential effects of spillovers in shocks across industries. Yet, since the costs of switching industries are likely to be very heterogeneous across different industry pairs, similar interpretation challenges to those we emphasize in the spatial context may arise in studies using cross-industry variation.

References

Adão, Rodrigo, "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil," *Working Paper*, 2016.

_ , Costas Arkolakis, and Federico Esposito, "General Equilibrium Effects in Space: Theory and Measurement," *Working Paper*, 2021.

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *The Quarterly Journal of Economics*, 08 2019, *134* (4), 1949–2010.

Agrawal, David R. and Dirk Foremny, "Relocation of the Rich: Migration in Response to Top Tax Rate Changes from Spanish Reforms," *The Review of Economics and Statistics*, 05 2019, *101* (2), 214–232.

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

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

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

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

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

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

Arthi, V, B Beach, and W Hanlon, "Recessions, Mortality, and Migration Bias: Evidence from the Lancashire Cotton Famine," 2019, (1552692).

Autor, David, David Dorn, and Gordon Hanson, "On the Persistence of the China Shock," *NBER Working Paper*, 2021, (29401).

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

_ , _ , _ , and Jae Song, " Trade Adjustment: Worker-Level Evidence *," *The Quarterly Journal of Economics*, 09 2014, *129* (4), 1799–1860.

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

Berkes, Enrico, Ruben Gaetani, and Martí Mestieri, "Cities and Technological Waves," *CEPR Discussion Paper*, 2021, (DP16794).

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

Bertoli, Simone and Jesús Fernández-Huertas Moraga, "Multilateral resistance to migration," *Journal of Development Economics*, 2013, *102*, 79–100.

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

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

Borusyak, Kirill and Peter Hull, "Non-Random Exposure to Exogenous Shocks: Theory and Applications," *Mimeo*, 2021.

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

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

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

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

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

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

Choi, Jiwon, Ilyana Kuziemko, Ebonya Washington, and Gavin Wright, "Local Economic and Political Effects of Trade Deals: Evidence from NAFTA," *NBER Working Paper*, 2021, (29525).

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

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

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

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

— **and** — , "Trade Liberalization and Regional Dynamics," *American Economic Review*, 2017, *107* (10), 2908–2946.

— **and** — , "Margins of Labor Market Adjustment to Trade," *Journal of International Economics*, 2019, *117*, 125–142.

Faber, Marius, Andrés Sarto, and Marco Tabellini, "Local Shocks and Internal Migration: The Disparate Effects of Robots and Chinese Imports in the US," 2021, (14623).

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

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

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

Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi, "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *Review of Economic Studies*, forthcoming.

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

Hakobyan, Shushanik and John McLaren, "Looking for Local Labor Market Effects of NAFTA," *The Review of Economics and Statistics*, 10 2016, *98* (4), 728–741.

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

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

Hyman, Benjamin G., "Can Displaced Labor Be Retrained? Evidence from Quasi-Random Assignment to Trade Adjustment Assistance," *Working paper*, 2018.

IBGE, *Censo Demográfico 2000: Documentação das Microdados da Amostra*, Instituto Brasileiro de Geografia e Estatística, 2002.

Imbert, Clement, Marlon Seror, Yifan Zhang, and Yanos Zylberberg, "Migrants and Firms: Evidence from China," *working paper*, 2021.

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

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

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

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

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

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

Monràs, Joan, "Economic Shocks and Internal Migration," *Ssrn*, 2020.

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

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

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

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

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

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

Tian, Yuan, Junjie Xia, and Rudai Yang, "Trade-Induced Urbanization and the Making of Modern Agriculture," *working paper*, 2021.

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

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

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

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

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

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

A Theory Appendix

A.1 Incorporating Labor Supply Shocks

In this appendix we show that the expressions from Section 2 extend naturally when the moving costs $\tau_{\ell d}$, in addition to D_ℓ, are allowed to change, generating local labor supply shocks. In that case, (12) is replaced by

$$\begin{aligned}\hat{L}_d &= \sum_\ell \gamma_{\ell d}\hat{\pi}_{\ell d} \\ &= \theta\left(\hat{w}_d - \sum_\ell \gamma_{\ell d}\sum_{d'}\pi_{\ell d'}\hat{w}_{d'}\right) + \nu_d\end{aligned}$$

where the labor supply shock is defined as

$$\nu_d = -\theta\sum_\ell \gamma_{\ell d}\left(\hat{\tau}_{\ell d} - \sum_{d'}\pi_{\ell d'}\hat{\tau}_{\ell d'}\right).$$

The supply shock is high when moving costs to d decline, relative to the moving costs from the same origins to other destinations. In vector form, $\hat{\mathbf{L}} = \theta\left(I - \Gamma'\Pi\right)\hat{\mathbf{w}} + \boldsymbol{\nu}$ with $\boldsymbol{\nu} = \{\nu_d\}$. Solving the system of demand and supply (i.e. using (9)), this implies

$$\hat{\mathbf{L}} = \Omega\hat{\mathbf{D}} + (I - \Omega)\boldsymbol{\nu}.$$

Finally, (18) yields

$$\hat{\mathbf{L}} = \Omega\hat{\mathbf{z}} + \boldsymbol{\zeta}_2 + (I - \Omega)\boldsymbol{\nu},$$

which generates a second, supply-driven, error term. The results of the paper apply if the $\hat{\mathbf{z}}$ shocks are mean-independent from both unobserved labor demand and labor supply shocks.

A.2 Low-Mobility Approximation Derivation

Plugging $\Gamma = I + \Delta\Gamma$ and $\Pi = I + \Delta\Pi$ into (16) and assuming $\Delta\Gamma'\Delta\Pi \approx 0$ yields

$$\Omega \approx I - \left(I - \frac{\theta}{\sigma}\left(\Delta\Gamma' + \Delta\Pi\right)\right)^{-1}.$$

Let $A \equiv \frac{\theta}{\sigma}\left(\Delta\Gamma' + \Delta\Pi\right)$ such that $\Omega = I - (I - A)^{-1}$. It is straightforward to verify that $(I - A)^{-1} = I + A + A^2 + A^3 + \dots$ and therefore $\Omega = -\left(A + A^2 + A^3 + \dots\right)$. Moreover, under the low-mobility assumption, $A + A^2 + A^3 + \dots \approx A$ since $\Delta\Gamma'\Delta\Pi \approx \Delta\Gamma'\Delta\Gamma' \approx \Delta\Pi\Delta\Pi \approx \mathbf{0}$.We thus have

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

and

$$\hat{\mathbf{L}} = \frac{\theta}{\sigma}\left(I - \Gamma' + I - \Pi\right)\hat{\mathbf{D}},$$

which is (17) in matrix notation.

A.3 Derivation of Equation (21)

From (17) we have:

$$\begin{aligned}
\hat{L}_\ell &\approx \frac{2\theta}{\sigma}\sum_{k\neq\ell}\frac{\gamma_{k\ell}+\pi_{\ell k}}{2}\left(\hat{D}_\ell-\hat{D}_k\right).\\
&= \frac{2\theta}{\sigma}\sum_{k\neq\ell}\frac{f_{k\ell}/L_\ell+f_{\ell k}/L_\ell^0}{2}\left(\hat{D}_\ell-\hat{D}_k\right)\\
&\approx \frac{2\theta}{\sigma}\sum_{k\neq\ell}\frac{F_{k\ell}}{L_\ell}\left(\hat{D}_\ell-\hat{D}_k\right)\\
&= \frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}\left(\hat{D}_\ell-\sum_{k\neq\ell}\frac{F_{k\ell}}{M_\ell}\hat{D}_k\right)\\
&= \frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}\left(\hat{z}_\ell+\zeta_{1\ell}-\sum_{k\neq\ell}\frac{F_{k\ell}}{M_\ell}\left(\hat{z}_k+\zeta_{1k}\right)\right)\\
&= \frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}\left(\hat{z}_\ell-\hat{z}_{-\ell}\right)+\zeta_{2\ell}.
\end{aligned}$$

where the first line cancels the terms for the same origin and destination, the third line uses that under the low-mobility approximation $L_\ell^0 \approx L_\ell$, and the other lines rearrange terms and use the definitions of $F_{k\ell}$, M_ℓ, and $\hat{z}_{-\ell}$.

A.4 Analyzing Observed Population Changes

Let $\mathbf{L}^0$ be the exogenously given vector of initial populations at $t = 0$, and $\mathbf{L}^1(\hat{\mathbf{z}}, \zeta_1)$ be the vector of populations at $t = 1$ as a function of observed and unobserved labor demand shocks.[51] Then the observed population change between $t = 0$ and $t = 1$ is $\log \mathbf{L}^1(\hat{\mathbf{z}}, \zeta_1) - \log \mathbf{L}^0$. We first can compare it to the population change relative to the no-shock counterfactual, as in footnote 24: by equation (20),

$$\log \mathbf{L}^1(\hat{\mathbf{z}}, \zeta_1) - \log \mathbf{L}^1(\mathbf{0}, \mathbf{0}) \equiv \hat{\mathbf{L}} = \Omega \cdot (\hat{\mathbf{z}} + \zeta_1).$$

This counterfactual differs from the observed change because of the secular trend in population,

$$\log \mathbf{L}^1(\mathbf{0}, \mathbf{0}) - \log \mathbf{L}^0 = \log\left(\Pi\mathbf{L}^0\right) - \log \mathbf{L}^0,$$

which is non-stochastic in our framework. Thus, it is orthogonal from $\hat{\mathbf{z}}$ by the as-good-as-randomness of the observed shocks.

The same argument applies to the causal impact of the observed shocks on population, $\log \mathbf{L}^1(\hat{\mathbf{z}}, \zeta_1) - \log \mathbf{L}^1(\mathbf{0}, \zeta_1)$, as in Section 3.3. It differs from the observed population change by

$$\log \mathbf{L}^1(\mathbf{0}, \zeta_1) - \log \mathbf{L}^0 = \log\left(\Pi\mathbf{L}^0\right) - \log \mathbf{L}^0 + \Omega\zeta_1,$$

which captures both the secular trend and the effects of (small) unobserved labor demand shocks, and is again orthogonal to $\hat{\mathbf{z}}$.

[51] Labor supply shocks which can be incorporated as in Appendix A.1.

A.5 $\hat{\beta}$ 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 of domestic regions R. 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 R, which jointly determine all relevant economic variables. We have:

Proposition A1. *Suppose* $\sum_\ell \frac{L_\ell^0}{L^0}\hat{L}_\ell\left(\hat{z}_\ell-\bar{z}\right)\xrightarrow{p}C$ *and* $\sum_\ell \frac{L_\ell^0}{L^0}\left(\hat{z}_\ell-\bar{z}\right)^2\xrightarrow{p}V>0$ *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 $\mathbb{E}\left[\sum_\ell \frac{L_\ell^0}{L^0}\hat{L}_\ell\left(\hat{z}_\ell-\bar{z}\right)\right]\to C$ and $\mathbb{E}\left[\sum_\ell \frac{L_\ell^0}{L^0}\left(\hat{z}_\ell-\bar{z}\right)^2\right]\to V$. Thus, $\beta\to\frac{C}{V}$, as required. □

A.6 Generalization and Proof of Theorem 1

We prove a generalization of Theorem 1 which applies beyond homoskedastic shocks:

Theorem A1. *Suppose the data are generated by the low-mobility approximation (21) to the baseline model. Then*

$$\beta=\frac{2\theta}{\sigma}\cdot\frac{M^v/L}{\widetilde{M}^v/L}\cdot\frac{1-\rho}{1-\tilde{\rho}}. \tag{41}$$

Here $\frac{M^v}{L}\equiv\frac{\sum_\ell M_\ell v_\ell}{\sum_\ell L_\ell v_\ell}$ *is the national average migration share weighted by shock variances* $v_\ell=\mathrm{Var}\left[\hat{z}_\ell\right]$*, and* $\rho=\frac{\sum_\ell M_\ell v_\ell\rho_\ell}{\sum_\ell M_\ell v_\ell}\in[-1,1]$ *for* $\rho_\ell=\frac{\mathrm{Cov}[\hat{z}_\ell,\hat{z}_{-\ell}]}{v_\ell}$*. Similarly,* $\frac{\widetilde{M}^v}{L}\equiv\frac{\sum_\ell \widetilde{M}_\ell v_\ell}{\sum_\ell L_\ell v_\ell}$*,* $\tilde{\rho}=\frac{\sum_\ell \widetilde{M}_\ell v_\ell\tilde{\rho}_\ell}{\sum_\ell \widetilde{M}_\ell v_\ell}\in[-1,1]$*, and* $\tilde{\rho}_\ell=\frac{\mathrm{Cov}[\hat{z}_\ell,\tilde{z}_{-\ell}]}{v_\ell}$ *for* $\tilde{z}_{-\ell}=\sum_{d\neq\ell}\frac{(\tilde{f}^0_{\ell d}+\tilde{f}^0_{d\ell})/2}{\widetilde{M}^0_\ell}\hat{z}_d$*. Moreover, when shocks are homoskedastic, i.e.* $v_\ell=v$ *for all* ℓ*, then* ρ *and* $\tilde{\rho}$ *coincide with the Theorem 1 definitions,* $M^v=M$*, and* $\widetilde{M}^v=\widetilde{M}$*.*

Proof. 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, $\mu=0$. Moreover, under the low-mobility approximation, $L_\ell\approx L_\ell^0$, and weighting the regression by the no-shock counterfactual employment levels L_ℓ is equivalent to weighting it by pre-shock employment L_ℓ^0. We therefore consider L_ℓ weights, which are more convenient for analytical results, and similarly redefine $\bar{z}=\sum_\ell \frac{L_\ell}{L}\hat{z}_\ell$.

We start with the numerator of β in (24), which can be represented as

$$\begin{aligned}\mathbb{E}\left[\sum_\ell L_\ell\hat{L}_\ell\left(\hat{z}_\ell-\bar{z}\right)\right]&=\mathbb{E}\left[\sum_\ell L_\ell\hat{L}_\ell\hat{z}_\ell\right]\\&=\mathbb{E}\left[\frac{2\theta}{\sigma}\sum_\ell M_\ell\hat{z}_\ell\left(\hat{z}_\ell-\hat{z}_{-\ell}\right)\right]\\&=\frac{2\theta}{\sigma}\sum_\ell M_\ell v_\ell\left(1-\rho_\ell\right).\end{aligned} \tag{42}$$

Here the first line used the fact that $\sum_\ell L_\ell\hat{L}_\ell\bar{z}=\bar{z}\sum_\ell L_\ell\hat{L}_\ell=\bar{z}\hat{L}=0$ because total national population is not allowed to change in the model. The second line plugged in (21) and used as-good-as-random assignment of $\hat{\mathbf{z}}$, and the last line used the definitions of v_ℓ and ρ_ℓ.

Now turn to the denominator of β. Since $\bar{z}$ is a weighted average of the shock to ℓ and other regions, specifically $\bar{z} = \frac{L_\ell}{L}\hat{z}_\ell + \left(1 - \frac{L_\ell}{L}\right)\widetilde{z}_{-\ell}$, we have

$$\hat{z}_\ell - \bar{z} = \left(1 - \frac{L_\ell}{L}\right)\left(\hat{z}_\ell - \widetilde{z}_{-\ell}\right). \tag{43}$$

Thus,

$$\begin{aligned}
\mathbb{E}\left[\sum_\ell L_\ell \left(\hat{z}_\ell - \bar{z}\right)^2\right] &= \mathbb{E}\left[\sum_\ell L_\ell \hat{z}_\ell \left(\hat{z}_\ell - \bar{z}\right)\right] \\
&= \mathbb{E}\left[\sum_\ell L_\ell \left(1 - \frac{L_\ell}{L}\right)\hat{z}_\ell \left(\hat{z}_\ell - \widetilde{z}_{-\ell}\right)\right] \\
&= \mathbb{E}\left[\sum_\ell \widetilde{M}_\ell \hat{z}_\ell \left(\hat{z}_\ell - \widetilde{z}_{-\ell}\right)\right] \\
&= \sum_\ell \widetilde{M}_\ell v_\ell \left(1 - \widetilde{\rho}_\ell\right).
\end{aligned} \tag{44}$$

Here the first line used $\sum_\ell L_\ell \bar{z}\left(\hat{z}_\ell - \bar{z}\right) = \bar{z}\sum_\ell L_\ell\left(\hat{z}_\ell - \bar{z}\right) = 0$ by definition of $\bar{z}$. The second line used (43). The third line used the definition of $\widetilde{M}_\ell$, and the final line used the definition of $\widetilde{\rho}_\ell$.

Combining (42) and (44) yields

$$\begin{aligned}
\beta &= \frac{2\theta}{\sigma}\frac{\sum_\ell M_\ell v_\ell \left(1 - \rho_\ell\right)}{\sum_\ell \widetilde{M}_\ell v_\ell \left(1 - \widetilde{\rho}_\ell\right)} \\
&= \frac{2\theta}{\sigma}\frac{\sum_\ell M_\ell v_\ell / \sum_\ell L_\ell v_\ell}{\sum_\ell \widetilde{M}_\ell v_\ell / \sum_\ell L_\ell v_\ell} \cdot \frac{\sum_\ell M_\ell v_\ell \left(1 - \rho_\ell\right) / \sum_\ell M_\ell v_\ell}{\sum_\ell \widetilde{M}_\ell v_\ell \left(1 - \widetilde{\rho}_\ell\right) / \sum_\ell \widetilde{M}_\ell v_\ell},
\end{aligned}$$

which implies (41).

To show that $\rho \in [-1, 1]$, we write

$$\rho = \frac{\sum_\ell M_\ell \mathrm{Cov}\left[\hat{z}_\ell, \hat{z}_{-\ell}\right]}{\sum_\ell M_\ell v_\ell} = \frac{\sum_{k \neq \ell} F_{k\ell}\mathrm{Cov}\left[\hat{z}_\ell, \hat{z}_k\right]}{\sum_\ell M_\ell v_\ell} = \frac{\sum_{k \neq \ell} f_{k\ell}\mathrm{Cov}\left[\hat{z}_\ell, \hat{z}_k\right]}{\sum_{k \neq \ell} f_{k\ell}\frac{\mathrm{Var}[\hat{z}_\ell] + \mathrm{Var}[\hat{z}_k]}{2}}.$$

We therefore have $\rho \leq 1$ because

$$\sum_{k \neq \ell} f_{k\ell}\frac{\mathrm{Var}\left[\hat{z}_\ell\right] + \mathrm{Var}\left[\hat{z}_k\right]}{2} - \sum_{k \neq \ell} f_{k\ell}\mathrm{Cov}\left[\hat{z}_\ell, \hat{z}_k\right] = \frac{1}{2}\sum_{k \neq \ell} f_{k\ell}\mathrm{Var}\left[\hat{z}_\ell - \hat{z}_k\right] \geq 0$$

and $\rho \geq -1$ because

$$\sum_{k \neq \ell} f_{k\ell}\frac{\mathrm{Var}\left[\hat{z}_\ell\right] + \mathrm{Var}\left[\hat{z}_k\right]}{2} + \sum_{k \neq \ell} f_{k\ell}\mathrm{Cov}\left[\hat{z}_\ell, \hat{z}_k\right] = \frac{1}{2}\sum_{k \neq \ell} f_{k\ell}\mathrm{Var}\left[\hat{z}_\ell + \hat{z}_k\right] \geq 0.$$

The argument for $\tilde{\rho} \in [-1, 1]$ follows similarly.

Finally, suppose shocks are homoskedastic; that is, $v_\ell = v$ for all ℓ. This immediately implies

$M^v/L = M/L$ and $\widetilde{M}^v/L = \widetilde{M}/L$. Moreover, the numerator of β equals

$$\begin{aligned}
\frac{2\theta}{\sigma} vM(1-\rho) &= \frac{2\theta}{\sigma}\sum_\ell v_\ell M_\ell(1-\rho_\ell) \\
&= \frac{2\theta}{\sigma}\sum_\ell \left(vM_\ell - \sum_{k\neq\ell} F_{k\ell}\mathrm{Cov}\left[\hat{z}_k,\hat{z}_\ell\right]\right) \\
&= \frac{2\theta}{\sigma} v\sum_\ell \left(M_\ell - \sum_{k\neq\ell} f_{k\ell}\frac{\mathrm{Cov}\left[\hat{z}_k,\hat{z}_\ell\right]}{v}\right) \\
&= \frac{2\theta}{\sigma} vM\left(1 - \sum_{k\neq\ell} \frac{f_{k\ell}}{M}\mathrm{Corr}\left[\hat{z}_k,\hat{z}_\ell\right]\right),
\end{aligned}$$

such that the Theorem 1 formula for ρ applies. The denominator of β analogously justifies the Theorem 1 formula for $\tilde{\rho}$. □

A.7 Proof $\beta(\hat{z}_\ell - \bar{z})$ Yields MSE-Minimizing Predicted Population Growth

Proposition. *$\beta(z_\ell - \bar{z})$ is the best linear predictor of $\hat{L}_\ell$ among all $b(z_\ell - \bar{z})$, i.e. after demeaning the shocks, in the sense that it solves*

$$\min_b \mathbb{E}\left[\sum_\ell L_\ell\left(\hat{L}_\ell - b(z_\ell - \bar{z})\right)^2\right]. \tag{45}$$

Proof. The first order condition for (45) yields

$$\begin{aligned}
\mathbb{E}\left[\sum_\ell L_\ell\left(\hat{L}_\ell - b(z_\ell - \bar{z})\right)\cdot(z_\ell - \bar{z})\right] &= 0 \\
\mathbb{E}\left[\sum_\ell L_\ell \hat{L}_\ell (z_\ell - \bar{z})\right] &= b\cdot\mathbb{E}\left[\sum_\ell L_\ell (z_\ell - \bar{z})^2\right],
\end{aligned}$$

which implies $b = \beta$. □

A.8 Relationship Between $\beta \approx 0$ and Local Economic Outcomes

In this appendix we show in the context of our baseline model of Section 2 that the regression of local wage changes on the $\hat{\boldsymbol{z}}$ shock is approximately unbiased if and only if β from the migration regression is small, regardless of the reason why that it is small.[52]

By (8) and (18),

$$\hat{w}_\ell = \frac{1}{\sigma}\left(\hat{D}_\ell - \hat{L}_\ell\right) = \frac{1}{\sigma}\left(\hat{z}_\ell + \zeta_{1\ell} - \hat{L}_\ell\right).$$

Therefore, a conventional regression of $\hat{w}_\ell$ on the regional shock $\hat{z}_\ell$ would produce an estimand $\frac{1}{\sigma}(1-\beta)$. This estimand is close to $\frac{1}{\sigma}$ whenever $\beta \approx 0$, regardless of the causes, described in

[52]In a model extension (available by request) which allows for the choice of non-employment we establish a similar result for regressions in which the outcome variable is local employment change, which is now distinct from the change in local population.

Theorem 1. Note, however, that in some locations fitted values from the wage regression, $\frac{1}{\sigma}(1-\beta)\hat{z}_\ell$, may be far from the true effect, $\frac{1}{\sigma}\left(\hat{z}_\ell - \hat{L}_\ell\right)$, because migration responses $\hat{L}_\ell$ can be large in some regions. Moreover, the regression-based predicted effect of a counterfactual shock to a specific location ℓ, $\frac{1}{\sigma}(1-\beta)$, may differ substantially from the true effect of $\frac{1}{\sigma}\left(1 - \frac{2\theta}{\sigma}\frac{M_\ell}{L_\ell}\right)$, for locations with high migration shares.

A.9 Invariance to Recentering

In this appendix we show that recentering as in Borusyak and Hull (2021) is not needed for the NLLS problem (29) when the shocks are unconditionally as-good-as-randomly assigned to all regions, and how it can be performed with conditional as-good-as-random assignment.

Let $\Omega_\ell(\lambda)$ be the ℓth row of $\Omega(\lambda)$, such that $\hat{L}_\ell = \Omega_\ell(\lambda)\hat{\mathbf{z}} + \zeta_{2\ell}$. As we have shown in footnote 19, $\Omega(\lambda)\iota = 0$ (for any λ), and thus $\frac{\partial\Omega}{\partial\lambda}\iota = 0$ as well, where $\frac{\partial\Omega}{\partial\lambda}$ is a matrix of element-by-element derivatives of $\Omega(\lambda)$.

The first order condition corresponding to the NLLS problem (29) is

$$\mathbb{E}\left[\sum_\ell \left(\hat{L}_\ell - \Omega_\ell(\lambda)\hat{\mathbf{z}}\right)\cdot\frac{\partial\Omega_\ell}{\partial\lambda}\hat{\mathbf{z}}\right] = 0. \tag{46}$$

It holds at the true $\lambda = \frac{\theta}{\sigma}$ if and only if

$$\mathbb{E}\left[\sum_\ell \zeta_{2\ell}\cdot\frac{\partial\Omega_\ell}{\partial\lambda}\hat{\mathbf{z}}\right] = 0.$$

When $\mathbb{E}\left[\hat{z}_\ell \mid \boldsymbol{\zeta}_2\right] = \mu$ for all ℓ, i.e. in the unconditional as-good-as-random assignment case, this holds because, by the law of iterated expectations,

$$\mathbb{E}\left[\sum_\ell \zeta_{2\ell}\cdot\frac{\partial\Omega_\ell}{\partial\lambda}\hat{\mathbf{z}}\right] = \mathbb{E}\left[\sum_\ell \zeta_{2\ell}\cdot\frac{\partial\Omega_\ell}{\partial\lambda}\mu\iota\right] = 0.$$

This situation contrasts with *conditionally* as-good-as-random shocks, formalized by $\mathbb{E}\left[\hat{z}_\ell \mid \boldsymbol{\zeta}_2\right] = \widetilde{\mu}' q_\ell$, where q_ℓ is a vector of observables. This weaker assumption is appropriate, for instance, when $\hat{z}_\ell$ is a shift-share variable based on some shifters (e.g. the growth of import penetration from China) which only happen in manufacturing industries, with q_ℓ including the regional share of manufacturing employment in addition to the intercept (Borusyak et al., 2022). In that case the NLLS moment condition (46) need not hold, and its recentered version should be used (Borusyak and Hull, 2021):

$$\mathbb{E}\left[\sum_\ell \left(\hat{L}_\ell - \Omega_\ell(\lambda)\hat{\mathbf{z}}\right)\cdot\frac{\partial\Omega_\ell}{\partial\lambda}\left(\hat{\mathbf{z}} - Q\widetilde{\mu}\right)\right] = 0,$$

where matrix Q collects the q_ℓ, and $\widetilde{\mu}$ can be estimated in a zeroth step.

A.10 Non-Linear Least Squares Identification

In this appendix we prove that $\frac{\theta}{\sigma}$ uniquely solves the minimization problem (29) whenever the variance-covariance matrix of $\hat{\mathbf{z}}$ is non-degenerate (i.e., positive definite).

The NLLS objective function (in expectation) can be written as

$$\begin{aligned}
&\mathbb{E}\left[\left(\hat{\mathbf{L}}-\Omega\left(\lambda\right)\hat{\mathbf{z}}\right)'\left(\hat{\mathbf{L}}-\Omega\left(\lambda\right)\hat{\mathbf{z}}\right)\right]\\
&\qquad=\mathbb{E}\left[\left(\Omega\left(\frac{\theta}{\sigma}\right)\hat{\mathbf{z}}-\Omega\left(\lambda\right)\hat{\mathbf{z}}+\boldsymbol{\zeta}_2\right)'\left(\Omega\left(\frac{\theta}{\sigma}\right)\hat{\mathbf{z}}-\Omega\left(\lambda\right)\hat{\mathbf{z}}+\boldsymbol{\zeta}_2\right)\right]\\
&\qquad=\mathbb{E}\left[\left(\Omega\left(\frac{\theta}{\sigma}\right)\left(\hat{\mathbf{z}}-\mu\iota\right)-\Omega\left(\lambda\right)\left(\hat{\mathbf{z}}-\mu\iota\right)+\boldsymbol{\zeta}_2\right)'\left(\Omega\left(\frac{\theta}{\sigma}\right)\left(\hat{\mathbf{z}}-\mu\iota\right)-\Omega\left(\lambda\right)\left(\hat{\mathbf{z}}-\mu\iota\right)+\boldsymbol{\zeta}_2\right)\right]\\
&\qquad=\operatorname{Var}\left[\left(\Omega\left(\frac{\theta}{\sigma}\right)-\Omega\left(\lambda\right)\right)\hat{\mathbf{z}}\right]+\mathbb{E}\left[\boldsymbol{\zeta}_2'\boldsymbol{\zeta}_2\right],
\end{aligned} \tag{47}$$

where the third line uses $\Omega\left(\lambda\right)\iota=0$ (see footnote 19 and Appendix A.9) and the last line follows since, under $\mathbb{E}\left[\hat{\mathbf{z}}-\mu\iota\mid\boldsymbol{\zeta}_2\right]=0$, the cross-terms

$$\mathbb{E}\left[\boldsymbol{\zeta}_2'\left(\Omega\left(\frac{\theta}{\sigma}\right)-\Omega\left(\lambda\right)\right)\left(\hat{\mathbf{z}}-\mu\iota\right)\right]=0$$

by the law of iterated expectations.

We now show that $\frac{\theta}{\sigma}$ uniquely minimizes (47). The second term in (47) does not depend on λ. The first term is zero when $\lambda=\frac{\theta}{\sigma}$ and positive whenever $\Omega\left(\lambda\right)\neq\Omega\left(\frac{\theta}{\sigma}\right)$ since the variance-covariance matrix of $\hat{\mathbf{z}}$ is non-degenerate. However, $\Omega\left(\lambda\right)=\Omega\left(\frac{\theta}{\sigma}\right)$ implies $\left(I+\lambda\left(I-\Gamma'\Pi\right)\right)^{-1}=\left(I+\frac{\theta}{\sigma}\left(I-\Gamma'\Pi\right)\right)^{-1}$ and in turn $\lambda\left(I-\Gamma'\Pi\right)=\frac{\theta}{\sigma}\left(I-\Gamma'\Pi\right)$, which is impossible for $\lambda\neq\frac{\theta}{\sigma}$ as long as there is any mobility in the no-shock equilibrium, i.e. $\Gamma'\Pi\neq I$.

A.11 Proofs for the Full Model

Here we derive equation (33) and prove Theorem 2. Analogously to the proofs for the baseline model in Appendices A.3 and A.6, we use the fact that under the low-mobility approximation, $L^0_{\ell n}\approx L_{\ell n}$, $L^0_\ell\approx L_\ell$, and that it is irrelevant whether regional observations in the regression and region-by-industry observations when defining $\hat{z}_\ell$ are weighted by the pre-shock or no-shock (counterfactual) employment levels.

From (32)

$$\begin{aligned}
\hat{L}_\ell&=\sum_{n\in\mathcal{N}}\frac{L_{\ell n}}{L_\ell}\hat{L}_{\ell n}\\
&\approx\frac{\theta}{\sigma}\sum_{n\in\mathcal{N}}\frac{L_{\ell n}}{L_\ell}\sum_{k\in\mathscr{L}}\sum_{p\in\mathcal{N}}\left(\frac{f_{kp,\ell n}}{L_{\ell n}}+\frac{f_{\ell n,kp}}{L^0_{\ell n}}\right)\left(\hat{z}_{\ell n}-\hat{z}_{kp}\right)+\zeta_\ell\\
&\approx\frac{2\theta}{\sigma}\sum_{n\in\mathcal{N}}\sum_{k\in\mathscr{L}}\sum_{p\in\mathcal{N}}\frac{F_{kp,\ell n}}{L_\ell}\left(\hat{z}_{\ell n}-\hat{z}_{kp}\right)+\zeta_\ell.
\end{aligned}$$

In this sum, the terms corresponding to $\ell=k$ add up to zero:

$$\sum_{n\in\mathcal{N}}\sum_{p\in\mathcal{N}}\frac{F_{\ell p,\ell n}}{L_\ell}\left(\hat{z}_{\ell n}-\hat{z}_{\ell p}\right)=\sum_{n}\frac{F_{\ell,\ell n}-F_{\ell n,\ell}}{L_\ell}\hat{z}_{\ell n}=0,$$

since $F_{\ell,\ell n} = F_{\ell n,\ell}$ for any ℓ and n by definition of average bilateral flows. Thus,

$$\begin{aligned}
\hat{L}_\ell &\approx \frac{2\theta}{\sigma} \sum_{k \neq \ell} \sum_{n \in \mathcal{N}} \sum_{p \in \mathcal{N}} \frac{F_{kp,\ell n}}{L_\ell} \left(\hat{z}_{\ell n} - \hat{z}_{kp}\right) + \zeta_\ell \\
&= \frac{2\theta}{\sigma} \left(\sum_{k \neq \ell} \sum_{n \in \mathcal{N}} \frac{F_{k,\ell n}}{L_\ell} \hat{z}_{\ell n} - \sum_{k \neq \ell} \sum_{p \in \mathcal{N}} \frac{F_{kp,\ell}}{L_\ell} \hat{z}_{kp} \right) + \zeta_\ell \\
&= \frac{2\theta}{\sigma} \frac{M_\ell}{L_\ell} \left(\hat{z}_\ell^{\text{mov}} - \hat{z}_{-\ell}^{\text{mov}} \right) + \zeta_\ell,
\end{aligned}$$

proving (33).

From (33), the proof of Theorem 2 follows similarly to the proof of Theorem A1. Specifically, the denominator of β is exactly the same:

$$\mathbb{E}\left[\sum_\ell L_\ell \left(\hat{z}_\ell - \bar{z}\right)^2 \right] = \sum_\ell \widetilde{M}_\ell v_\ell \left(1 - \widetilde{\rho}_\ell\right).$$

Assuming without loss of generality that $\mathbb{E}\left[\hat{z}_\ell\right] = \mu = 0$, the numerator can further be expressed as

$$\mathbb{E}\left[\sum_\ell L_\ell \hat{L}_\ell \hat{z}_\ell \right] = \mathbb{E}\left[\frac{2\theta}{\sigma} \sum_\ell M_\ell \hat{z}_\ell \left(\hat{z}_\ell^{\text{mov}} - \hat{z}_{-\ell}^{\text{mov}} \right) \right] = \frac{2\theta}{\sigma} \sum_\ell M_\ell v_\ell \left(\rho_\ell^{\text{mov}} - \rho_\ell \right).$$

Putting the two together,

$$\beta = \frac{2\theta}{\sigma} \frac{\sum_\ell M_\ell v_\ell \left(\rho_\ell^{\text{mov}} - \rho_\ell \right)}{\sum_\ell \widetilde{M}_\ell v_\ell \left(1 - \widetilde{\rho}_\ell\right)} = \frac{2\theta}{\sigma} \cdot \frac{M^v / L}{\widetilde{M}^v / L} \cdot \frac{\rho^{\text{mov}} - \rho}{1 - \widetilde{\rho}},$$

as required by the theorem.

A.12 Allowing for Large Shocks

For expositional clarity and analytical tractability, the analysis in the main text studies small shocks relative to a no-shock counterfactual, such that $\hat{x} \equiv dx/x = d \ln x$ (see Section 2.2), i.e. we consider a first-order approximation to the effects of arbitrary shocks. Here, we use exact hat algebra to determine how population responds to large shocks, following Dekle et al. (2008), allowing us to empirically assess empirically how well the first-order approximation performs in practice.

Let hats now denote exact log-changes. The labor demand equation is unchanged:

$$\hat{L}_\ell = \hat{D}_\ell - \sigma \hat{w}_\ell. \tag{48}$$

The labor supply equation, however, does change. Let superscripts 0 and 1 refer to two counterfactual situations.

$$\begin{aligned}
\hat{L}_\ell \equiv \log \frac{L_\ell^1}{L_\ell^0} &= \log \frac{\sum_o L_o \frac{\left(w_\ell^1/\tau_{o\ell}\right)^\theta}{\sum_d \left(w_d^1/\tau_{od}\right)^\theta}}{\sum_o L_o \frac{\left(w_\ell^0/\tau_{o\ell}\right)^\theta}{\sum_d \left(w_d^0/\tau_{od}\right)^\theta}} \\
&= \log \sum_o \gamma_{o\ell} \frac{\exp \theta \hat{w}_\ell}{\sum_d \left(w_d^1/\tau_{od}\right)^\theta / \sum_d \left(w_d^0/\tau_{od}\right)^\theta} \\
&= \theta \hat{w}_\ell + \log \sum_o \gamma_{o\ell} \left(\sum_d \pi_{od} \exp \theta \hat{w}_d \right)^{-1} \\
&= \frac{\theta}{\sigma} \left(\hat{D}_\ell - \hat{L}_\ell \right) + \log \sum_o \gamma_{o\ell} \left(\sum_d \pi_{od} \exp \frac{\theta}{\sigma} \left(\hat{D}_d - \hat{L}_d \right) \right)^{-1},
\end{aligned} \tag{49}$$

where the first line used the choice probabilities in (10), the second and third lines used standard exact hat algebra manipulations, and the last line plugged in (48). Here, (49) is a system of nonlinear equations yielding $\hat{\mathbf{L}}$ in response to any $\hat{\mathbf{D}}$ whose log-linear approximation is $\hat{\mathbf{L}} = \Omega \hat{\mathbf{D}}$ as in (14) in the main text. By solving (49) numerically for a given shock vector $\hat{\mathbf{D}}$, we can check the quality of the approximation in (14).

In results available upon request, we implement this comparison using simulated shocks and find that the first-order approximation yields quantitatively accurate results even for shocks that are substantially larger than those faced by Brazilian labor markets during the country's early 1990s trade liberalization. For example, the approximation error is very small for shocks with a standard deviation of 0.7, while the shocks in Brazil's liberalization had a standard deviation of 0.064.

A.13 Bilateral Flow Regressions

We first show that bilateral flow regressions as in (30) are consistent with the model in Section 2, under the low-mobility approximation in (21). Using $f_{o\ell} \equiv \pi_{o\ell} L_o^0$ for $\ell \neq o$ and differentiating (10), we have

$$\hat{f}_{o\ell} = \hat{\pi}_{o\ell} = \theta \left(\hat{w}_\ell - \sum_d \pi_{od} \hat{w}_d \right).$$

Plugging in (15) and imposing the low-mobility approximation in which $\Pi, \Gamma \approx \mathbb{I}$ and hence $\hat{\mathbf{w}} \approx (1/\sigma)\hat{\mathbf{D}}$ yields

$$\hat{f}_{o\ell} \approx \frac{\theta}{\sigma} \left(\hat{D}_\ell - \sum_d \pi_{od} \hat{D}_d \right) \approx \frac{\theta}{\sigma} \left(\hat{D}_\ell - \hat{D}_o \right).$$

Now separate observed and unobserved labor demand shocks by plugging in (18).

$$\hat{f}_{o\ell} \approx \frac{\theta}{\sigma} (\hat{z}_\ell - \hat{z}_o) + \zeta_{o\ell}, \tag{50}$$

where $\zeta_{o\ell}$ is an error term reflecting unobserved shocks. This expression corresponds to the pair-level flow regression in (30).

Given estimates from the bilateral flow regression, it is straightforward to aggregate up to the location level to answer the same question motivating the conventional migration regression: How did observed shocks affect the population in each location ℓ? First, note that $\hat{L}_\ell = \sum_o \gamma_{o\ell} \hat{f}_{o\ell} =$

$\gamma_{\ell\ell}\hat{f}_{\ell\ell} + \sum_{o\neq\ell}\gamma_{o\ell}\hat{f}_{o\ell}$. Since the initial population is exogenously fixed, we also have $\sum_d \pi_{\ell d}\hat{f}_{\ell d} = 0$ and thus $\gamma_{\ell\ell}\hat{f}_{\ell\ell} = \pi_{\ell\ell}\hat{f}_{\ell\ell} = -\sum_{d\neq\ell}\pi_{\ell d}\hat{f}_{\ell d}$. Plugging in $\gamma_{\ell\ell}\hat{f}_{\ell\ell}$ from here yields

$$\hat{L}_\ell = \sum_{k\neq\ell}\gamma_{k\ell}\hat{f}_{k\ell} - \sum_{k\neq\ell}\pi_{\ell k}\hat{f}_{\ell k}.$$

Plugging in (50) yields

$$\begin{aligned}\hat{L}_\ell &\approx \frac{\theta}{\sigma}\sum_{k\neq\ell}\left(\gamma_{k\ell}+\pi_{\ell k}\right)\left(\hat{z}_\ell - \hat{z}_k\right) + \zeta_{2\ell} \\ &\approx \frac{2\theta}{\sigma}\frac{M_\ell^0}{L_\ell^0}\left(\hat{z}_\ell - \hat{z}_{-\ell}\right) + \zeta_{2\ell},\end{aligned}$$

which corresponds precisely to (21) in the main text.

A.14 Housing Market

In this appendix, we consider the addition of a residential land market to the model with regional mobility costs in Section 2. We follow Redding (2016) by assuming the quantity of local land is fixed, which creates an additional congestion force in spatial equilibrium beyond that already induced in the baseline model by downward-sloping local labor demand. Suppose that utility from consumption C and land H is Cobb-Douglas,

$$U = \left(\frac{C}{\alpha}\right)^\alpha \left(\frac{H}{1-\alpha}\right)^{1-\alpha}. \tag{51}$$

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

$$\bar{H}_\ell = (1-\alpha)\, L_\ell \frac{v_\ell}{r_\ell}, \tag{52}$$

where v_ℓ is income, which includes both the wage and land rents, assumed to be distributed to consumers with a lump-sum transfer, such that

$$\begin{aligned} v_\ell L_\ell &= w_\ell L_\ell + (1-\alpha)v_\ell L_\ell \\ v_\ell &= \frac{w_\ell}{\alpha}. \end{aligned} \tag{53}$$

From (52) and (53),

$$r_\ell = \frac{1-\alpha}{\alpha}\frac{w_\ell L_\ell}{\bar{H}_\ell}.$$

Thus, the indirect utility in region ℓ is

$$V_\ell = \frac{v_\ell}{P^\alpha r_\ell^{1-\alpha}} = \frac{w_\ell \bar{H}_\ell^{1-\alpha}}{\alpha P^\alpha \left(\frac{1-\alpha}{\alpha}\right)^{1-\alpha} w_\ell^{1-\alpha} L_\ell^{1-\alpha}} = \left(\frac{w_\ell}{\alpha P}\right)^\alpha \left(\frac{\bar{H}_\ell}{(1-\alpha)L_\ell}\right)^{1-\alpha}. \tag{54}$$

Differentiating (54) with fixed land supply,

$$\hat{V}_\ell = \alpha\hat{w}_\ell - (1-\alpha)\hat{L}_\ell - \alpha\hat{P}. \tag{55}$$

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

$$\hat{\mathbf{L}} = \hat{\mathbf{D}} - \sigma \hat{\mathbf{w}}, \tag{56}$$

Changes in labor supply now depend on differential changes in V_ℓ rather than w_ℓ:

$$\hat{\mathbf{L}} = \theta \left(I - \Gamma' \Pi \right) \hat{\mathbf{V}} = \theta \left(I - \Gamma' \Pi \right) \left(\alpha \hat{\mathbf{w}} - (1-\alpha) \hat{\mathbf{L}} \right), \tag{57}$$

where the term $\alpha \hat{P}$ disappeared since it is the same in all regions. The system (56)-(57) can be solved to describe the effect of local labor demand shocks on local populations:

$$\hat{\mathbf{L}} = \left[I + ((1-\alpha)\sigma + \alpha) \frac{\theta}{\sigma} \left(I - \Gamma' \Pi \right) \right]^{-1} \cdot \frac{\alpha \theta}{\sigma} \left(I - \Gamma' \Pi \right) \hat{\mathbf{D}}. \tag{58}$$

Note that this expression depends independently upon all three structural parameters (θ, σ, α), unlike the baseline model, which depends only upon the single reduced form parameter θ/σ. However, when we reintroduce the low-mobility approximation, the resulting expression simplifies to exactly the form we found in the baseline model, where the effect of local shocks depends only upon $\alpha\theta/\sigma$:

$$\hat{\mathbf{L}} \approx \frac{\alpha \theta}{\sigma} \left(I - \Gamma' \Pi \right) \hat{\mathbf{D}}. \tag{59}$$

Because this expression takes the same form as in the baseline model, all of our result apply directly. The only change required is to appropriately calibrate the parameter $\alpha\theta/\sigma$ to account for the presence of housing demand.

A.15 Agglomeration

In this appendix, we show that two standard ways of incorporating local agglomeration economies yield results that are isomorphic to those in the model with regional mobility costs in Section 2.

Approach 1: Productivity increases in $L_{\ell n}$ with elasticity ρ (and denoting $\tilde{\rho} = (1 - \rho(\sigma - 1))^{-1} \geq 1$). In this case, the expression for local labor demand is isomorphic to (7); the presence of agglomeration economies simply increases the labor demand elasticity.

$$\begin{aligned} L_{\ell n} &= \alpha_n \left(\varphi_{\ell n} L_{\ell n}^{\rho} P_n \right)^{\sigma - 1} w_\ell^{-\sigma} \\ &= \left(\alpha_n \left(\varphi_{\ell n} P_n \right)^{(\sigma-1)} \right)^{\tilde{\rho}} w_\ell^{-\sigma \tilde{\rho}} \end{aligned}$$

and

$$L_\ell = \sum_n L_{\ell n} = D_\ell w_\ell^{-\sigma \tilde{\rho}}.$$

Approach 2: Productivity increases with L_ℓ with elasticity ρ. Once again, the expression for local labor demand is isomorphic to (7), and the presence of agglomeration economies increases the labor demand elasticity.

$$L_{\ell n} = \alpha_n \left(\varphi_{\ell n} L_\ell^{\rho} P_n \right)^{\sigma - 1} w_\ell^{-\sigma}$$

and

$$
\begin{aligned}
L_\ell = \sum_n L_{\ell n} &= \sum_n \alpha_n \left(\varphi_{\ell n} P_n\right)^{\sigma-1} \cdot L_\ell^{\rho(\sigma-1)} w_\ell^{-\sigma} \\
&= \left(\sum_n \alpha_n \left(\varphi_{\ell n} P_n\right)^{\sigma-1}\right)^{\tilde{\rho}} w_\ell^{-\sigma\tilde{\rho}} = D_\ell w_\ell^{-\sigma\tilde{\rho}}.
\end{aligned}
$$