

Migration frictions, earnings differentials, and spatial misallocation: Evidence from Thailand

Ashish Shenoy

Department of Agricultural and Resource Economics

University of California, Davis

1 Shields Avenue

Davis, CA 95616

530-752-0824

shenoy@ucdavis.edu

August 2023

Abstract

This paper uses revealed-preference location decisions of workers in Thailand to quantify the disutility of labor migration, characterize the migration contribution to labor supply elasticity, and estimate the effect of migration frictions on spatial earnings differentials and labor misallocation. I structurally estimate a spatial equilibrium model using commodity prices as instruments for local earnings to overcome endogeneity and selection, and to identify the net present value returns to potential migration. Estimation employs a novel strategy using maximum likelihood to accommodate measurement error and choice-based sampling. I find migration contributes 9.5 percentage points to labor supply elasticity at the extensive margin, which is 25–50 percent as large as existing intensive-margin estimates among non-movers. The disutility from migration is 1.0–1.2 times annual earnings; alleviating this friction would induce a quarter of the population to relocate and lower spatial earnings variation by 20 percent. However, gains would be realized primarily in non-wage utility with a modest 3 percent increase in national product, suggesting migration frictions play a limited role relative to preference heterogeneity in productive misallocation.

Keywords: migration, labor misallocation, spatial equilibrium, Thailand

JEL Codes: O15, J61, R23

1 Introduction

Around the world there exist large differences in earnings across space, both within and between countries, that persist even after adjusting for local prices and human capital (Acemoglu and Dell, 2010). Such variation highlights the role played by location in both labor supply and demand. On the demand side, if firms were perfectly mobile, they would relocate to where labor was the cheapest. Similarly on the supply side, a perfectly mobile labor force would migrate toward higher wages. Geographic mismatch between demand and supply gives rise to earnings variation across markets. The mobility frictions that generate this mismatch may lower aggregate productivity through the misallocation of labor and can increase exposure to local shocks by concentrating economic impacts within a smaller geographic area.

In this paper I quantify the economic importance of mobility frictions across provinces in Thailand in the 1980's and 1990's. Like many countries, Thailand has substantial geographic variation in labor earnings, with the standard deviation across its 73 provinces at around half of the median. I analyze the labor supply contribution to this variation through the lens of a spatial equilibrium model of worker location choice (see Rosen, 1979; Roback, 1982; Moretti, 2011). In the model, the migration elasticity to differences in earnings across space is mediated by variation in location-specific amenities common to all workers, idiosyncratic worker preferences for place of residence, and migration frictions raising the cost of relocation. I use the model to characterize the net contribution of migration to local labor supply elasticity.

Model parameters are structurally estimated from the revealed preference migration response to changes in local market conditions. Data comes from annual cross-sectional labor force surveys conducted by the government of Thailand. From 1985 to 2000, surveys included questions about history of residence from which I construct a pseudo-panel of annual province-to-province migration flows. Data on gross bidirectional flows, as opposed to net flows or gross population changes, allow for separate identification of location-specific preferences and mobility costs, both of which may mute the migration elasticity to local earnings. The net migration response to a change in earnings reflects the importance of earnings relative to workers' geographic preferences. Conditional on net flows, high gross flows would indicate location-specific preference heterogeneity independent of location history, while low gross flows would suggest workers perceive a disutility from leaving their current place of residence.¹ I define the latter to be a mobility friction, and simulate counterfactual scenarios to investigate how much such frictions sustain spatial earnings gaps and depress aggregate productivity.

The main empirical innovation of this paper is to estimate model parameters using international commodity prices to isolate exogenous earnings variation caused by shifts in local labor demand.

¹Estimation reveals the perceived or anticipated disutility of migration as the choice to migrate is made before the disutility is realized. The actual experience may be better or worse than expected.

This identification strategy resolves three distinct challenges. First, instrumenting for local labor demand overcomes the standard problem of labor supply endogeneity. In spatial equilibrium, endogeneity arises when a correlated preference shock drives workers to or from a given destination, shifting the local labor supply curve and inducing a spurious negative relationship between migration and earnings. My empirical strategy addresses this concern by instrumenting for local earnings using global commodity prices interacted with local sensitivity to commodities, thereby leveraging exogenous, demand-based variation in the local earnings process.

Second, this empirical strategy separates changes in the return to migration from selection on unobservables. Any revealed-preference empirical strategy relies on observed earnings to infer the potential return to migration. However, recent empirical work has established the possibility that much spatial earnings variation can be attributed to worker selection (Young, 2013; Hamory et al., 2021) rather than local productivity. If workers sort geographically according to unobserved traits, then observed earnings levels may not reflect the true potential income one might earn from a move. To avoid selection bias, I exploit variation caused by fluctuations in local labor demand and verify that the resulting earnings changes are not driven by differential worker selection. Therefore, estimates from this paper are accurately based on changes in the potential return to migration rather than on unobserved characteristics of the population at the destination.

Third, the commodity-based instrument enables results to be quantified in currency units. In general, location choice is a long-term decision with benefits realized over time. Current earnings levels do not convey expectations about the future, making it difficult to calculate the expected return to migration in net present value terms. However, the portion of earnings identified by commodity instruments inherits the same time series properties as the commodity price series themselves. Therefore, estimation leverages earnings variation for which the expected net present value is readily computed, allowing other parameter estimates to be interpreted relative to monetary returns. I verify in the reduced form that for the same size contemporaneous shock to earnings, there is a greater migration response when the shock is generated by a more permanent commodity series. This differential response indicates that labor markets incorporate information about the expected future value of current earnings shocks, allowing other preference parameters to be scaled into meaningful units comparable to a dollar of net-present-value earnings.

This final feature expands on existing work using commodity price shocks as a source of identifying variation in local labor demand. A number of other studies use commodity prices to estimate the effect of changes in earnings on outcomes such as occupation selection (e.g. Young, 2014; James, 2015; Allcott and Keniston, 2017) and civil conflict (e.g. Dube and Vargas, 2013; Berman et al., 2017; McGuirk and Burke, 2020). I extend this literature by leveraging differences in the time series properties of different commodity prices as a proxy for expectations about the future value for shocks of comparable size.

This paper also introduces a methodological innovation to overcome two common weaknesses

found in the data. First, migration rates are measured with noise. Since the sample size is small relative to the number of potential migration channels, there is high sampling variance and zero observed migration along the majority of province-to-province channels. This issue also arises in retail data with fine or uncommon product categories (e.g [Gandhi et al., 2023](#)) as well as trade data with errors in measurement or reporting ([Feenstra et al., 2005](#); [UNCTAD, 2012](#)). Second, because sampling is stratified by province, sampling frequencies are endogenous to workers' residence and therefore to their migration decision. Such choice-based sampling arises whenever surveys stratify by outcome or oversample rare outcomes for statistical power ([Cosslett, 1981](#)). I show that naive model estimation that ignores these factors would generate inaccurate estimates of the importance of mobility frictions relative to earnings.

To preserve consistency, I derive a formula for the data-generating process that comprises both the spatial choice model as well as the survey sampling design. This strategy selects parameters to maximize the joint likelihood of each survey respondent having made their observed migration choice from their place of origin *and* of having been surveyed at their current destination conditional on that choice. Estimation requires imposing the constraint that aggregate migration flows in and out of each province must sum to the actual change in population for that province. The estimator in this paper is a generalization of the choice-based maximum likelihood approach proposed by [Manski and Lerman \(1977\)](#) that supplements individual choice data with aggregate choice probabilities. While the Thai data does not accurately measure aggregate migration probabilities along any specific province-to-province channel, the net change in each province's population reflects the aggregate sum across all migration channels. I show that this procedure has an equivalent interpretation that the set of survey respondents in each province represent the outcome of a draw from a multinomial distribution of potential origin provinces, with probabilities governed by migration rates and population sizes.

Results show the perceived disutility of migration to be substantial. The average utility cost of migrating is equivalent to 1.0–1.2 times annual earnings in net-present-value terms. Furthermore, the variation across individuals in their preferences over local non-wage amenities is on the same order of magnitude as the observed spatial variation in earnings. These two factors combine to limit the role of the labor market in efficient reallocation. The size of the estimated utility penalty to migration relative to annual earnings is consistent with comparable estimates from Brazil ([Morten and Oliveira, 2023](#)), and smaller but on the same order as the estimated cost in the U.S. ([Kennan and Walker, 2011](#)). A parallel avenue of research on the effects of trade finds the cost of changing sector, rather than changing location, to be significantly greater ([Artuç et al., 2010](#); [Dix-Carneiro, 2014](#)). It is unclear in general whether worker behavior is primarily governed by migration frictions with sector choice being a consequence of the location decision, or by sectoral frictions with location choice conforming to sector of occupation.

Despite the high effective cost of relocation, migrants play an important role in local labor

supply. The average elasticity to destination earnings of migration along a given province-to-province channel is near unity. Summing across all possible origins, migration contributes up to 9.5 percentage points to local labor supply elasticity. This figure is between 25 and 50 percent as large as existing estimates of labor supply elasticity among nonmovers in this period (see [Singh et al., 1986](#); [Bauer et al., 1988](#)).

Barriers to migration contribute to spatial variation in earnings. In a world with a perfectly mobile labor force where earnings exactly offset local non-wage amenities, a quarter of the Thai labor force would live in a different province and the standard deviation of earnings in Thailand in 2000 would be 20 percent lower than its actual level. Alternatively, earnings variation would need to be over twice as large to maintain the population distribution in 2000 if there were no barriers to migration. These two quantifications suggest that costly migration plays an important role in constraining the population and sustaining spatial earnings gaps.

While migration frictions limit the mobility of Thai workers, their effects on aggregate productivity are modest. With perfectly elastic local labor demand, eliminating barriers to migration would have raised aggregate earnings in the year 2000 by only 3 percent. The small size of the increase is due to the role that non-wage components of the utility function play in location decisions. Local amenities are negatively correlated with observed earnings, and therefore act as a countervailing force attracting workers to areas of low labor productivity. Therefore, while lowering mobility frictions would induce a high degree of relocation, much of the gains would be realized in non-wage utility and would do little to alleviate productive misallocation over space.

Research on labor misallocation and spatial earnings disparities is motivated by empirical evidence on the returns to migration. Cross-border studies using both randomized variation from visa lotteries ([McKenzie et al., 2010](#); [Clemens and Postel, 2017](#); [Mobarak et al., 2023](#)) as well as observational evidence ([Clemens, 2013](#); [Hendricks and Schoellman, 2017](#); [Clemens et al., 2019](#)) demonstrate that migrants who travel internationally enjoy higher earnings at their destination. Within-country, a small number of experimental evaluations find comparably high returns ([Bryan et al., 2014](#); [Baseler, 2023](#)). Panel data from a broader range of settings confirms that the labor return to domestic migration is typically positive, though migrants do not capture the full gap in observed earnings between their origin and destination because some of the gap may be driven by labor market sorting ([Beegle et al., 2011](#); [Hamory et al., 2021](#); [Alvarez, 2020](#); [Lagakos et al., 2020](#)).

Positive returns to migration reflect the potential for economic gains from worker relocation. This study quantifies the magnitude of potential gains and associated mobility frictions preventing those gains using plausibly exogenous variation in the returns to migration. My estimate of the productivity loss from labor misallocation within Thailand is consistent with comparable research in Brazil ([Morten and Oliveira, 2023](#)) and Indonesia ([Bryan and Morten, 2019](#)). These authors similarly discover that heterogeneity in workers' location-specific preferences continues to drive a wedge between location choice and labor market returns even in the absence of mobility frictions. In

contrast, [Desmet et al. \(2018\)](#) forecast much larger aggregate gains from migration across national borders.

My work more generally relates to labor market transformation in the process of economic development. I study a period of rapid development in Thailand with real per capita GDP growing by six percent annually. Accompanying growth, the agricultural fraction of the national labor force shrunk from almost two thirds to only 40 percent by the year 2000. This pattern of movement out of agriculture is echoed in growth trajectories worldwide ([Herrendorf et al., 2014](#)). [Gollin et al. \(2013\)](#) document a pervasive gap between agricultural and non-agricultural productivity that widens at the low end of the income scale, further evidencing a fundamental link between development and labor force composition. Such sectoral shifts are strongly tied to geography as different industries have different patterns of land use [Eckert and Peters \(2022\)](#). My findings reveal that mobility frictions may impede the pace of urbanization and structural transformation, and thereby sustain agricultural productivity gaps, along the path of national development.

Section 2 describes the available data in more detail and provides some descriptive facts about the period and sample of study. Section 3 introduces the empirical model of location choice, describes the commodity-based instruments used for estimation, and presents reduced-form evidence of their validity. Next, Section 4 details the maximum likelihood estimation procedure, and Section 5 reports results.

2 Labor Market Data Description

Data primarily comes from the Thai labor force survey (LFS), an annual, nationally representative cross-sectional survey conducted by the National Statistics Office of Thailand. The LFS covers roughly 0.3 percent of the population per year from 1985–1993, and doubles to around 0.6 percent per year thereafter. This paper restricts analysis to men² between the ages of 16 and 60, excluding those out of the labor force for disability or schooling. Roughly one quarter of respondents meet these criteria. In the first four years of study my sample includes roughly 36,000 respondents per year in 73 provinces. From 1989–1993 this number expands to around 50,000, and from 1994 onward the sample size is roughly 88,000 per year.

The LFS contains information on demographics and labor market outcomes. From 1985–2000, it also asked where respondents resided previously, up to five years back, which I use to construct yearly province-to-province migration flows. This measure captures substantially more migration than standard census questions about place of birth. For instance, imputing an annualized migration rate from place of birth in the 1990 and 2000 Thai censuses would yield a 20–40 percent underestimate in the population of study. Regrettably, the LFS only records respondents’ most

²Women are excluded due to low labor force participation, preventing calculation of their earnings and returns to migration.

recent move and lacks information on prior residence history.

LFS data reveal large, persistent spatial variation in earnings. In 2000, the standard deviation of province average earnings was around 45 percent of the mean, and the correlation between a province's earnings level in 2000 and 1985 is 0.73. Migrants unsurprisingly respond to this variation, so that a province's earnings level is a significant predictor of its population growth. Figure 1 plots the positive relationship between population growth and average earnings by province. Regression confirms that a log point increase in average earnings in 1985 corresponds to a 19 percent rise in population over the period of study.

On average, five percent of the working age population of men in Thailand move every year, and this value remains steady over the period of study. Despite the net directionality towards high earnings, gross migration between provinces is roughly five times the net flow, indicating substantial travel in the reverse direction. Migration rates in general decrease with distance: migrants travel only 83 percent as far as they would had they selected a destination at random. 30–40 percent of migration within the country is to or from Bangkok, whose population is an order of magnitude greater than any other province.

Summary statistics for demographic characteristics of migrants and the full sample are presented in Table 1. In general, migrants tend to be younger, work longer hours, and earn less, in part due to lower employment rates. Otherwise they look similar demographically to the general population. Notably, the period of study is characterized by universally low levels of education. Only a quarter of the population has completed primary education and only ten percent has completed secondary education, with little difference between migrants and non-migrants.

One distinguishing characteristic of migrants is sector of employment. Relative to the rest of the population, migrants are much less likely to be employed in agriculture, and instead tend to concentrate in manufacturing and construction. Among non-migrants just over half the population works in agriculture; this number drops to 40 percent among those who have migrated within the last year, some of whom are likely seasonal agricultural workers. Among those who migrated between 1 and 8 years ago, only 30 percent work in agriculture.

3 Empirical Model

I construct a static model of location choice with heterogeneous preferences following [Moretti \(2011\)](#). The model consists of forward-looking agents who each start in an initial location and choose a destination according to their preferences over expected earnings, location-specific non-wage amenities, and a disutility of migration. To identify expected earnings, I generate instruments by interacting local sensitivity to commodity prices with global price fluctuations. Earnings variation induced by these instruments inherits its time series characteristics from them, and therefore has a known expected future value. Reduced-form evidence shows that instruments reflect changes

to the potential earnings a worker would receive were they to locate in a province, and that workers are sensitive to both contemporaneous earnings as well as expected future earnings.

3.1 Location Choice

Let the economy contain a continuum of workers indexed by n distributed across J discrete provinces, with each province constituting a labor market. At each time t , each worker starts in an initial market and chooses one destination, which may be the same as the starting point. Workers then supply one unit of labor inelastically to the destination market and earn the local wage. The timing is such that in each period workers observe the prevailing wage everywhere, make a location decision, and then earn income.

Let the utility worker n would derive from selecting province j at time t , conditional on currently residing in province i , be given by

$$V_{njt} = \mathbb{E}_t[Y_{njt}] - c_{ijt} + A_j + \epsilon_{njt} \quad (1)$$

where Y_{njt} is the net present value of earnings at the destination, c_{ijt} is the disutility of moving from i to j , and $A_j + \epsilon_{njt}$ represent the non-earnings utility derived from living at destination j . The implicit utility coefficient on earnings is one, which provides a scale to translate migration costs and amenities into dollar equivalents. Each worker chooses the province that maximizes their period utility.

Preferences terms $A_j + \epsilon_{njt}$ capture the non-pecuniary value of living in a location, commonly referred to as amenities. These may include preferences over food, climate, culture, or any other location-specific aspect that delivers utility. Because individuals value local amenities differently, the non-wage utility is modeled an individual-location-specific random variable. It can be decomposed into province-specific terms A_j that describe the average amenity level of a province and idiosyncratic preferences ϵ_{njt} that describe each individual's tastes. The variance of this idiosyncratic term reflects the heterogeneity of tastes in the population relative to earnings utility. Heterogeneity mutes the migration response to earnings differences as non-wage amenities take on greater importance.

The disutility from migrating c_{ijt} is referred to as a “mobility friction” or “migration cost”. It encompasses the monetary cost of transportation as well as any psychic costs, barriers related to establishing new financial or social networks, and any other non-monetary disutility that migration may generate. It also embeds any prior sorting according to fixed, idiosyncratic locational preferences. Migration costs also diminish earnings-based relocation as workers need greater compensation to justify a move.

The key theoretical distinction between migration costs and earnings- or amenity-based preferences is history dependence. Individuals generally perceive a utility difference when deciding

between two provinces. Some of this gap is symmetric across all agents in the economy, but the portion that depends on current residence is considered a migration cost paid only if a location choice requires changing one’s place of residence. Mathematically, the cost of moving from province i to j is $c_{ijt} = V_{njjt} - V_{nijt}$. Earnings and amenities in i and j cancel out of the difference and therefore do not enter into the cost term.

The utility specification and subsequent estimation omit local prices and abstract from higher moments of the income distribution beyond expected earnings due to data limitations. I discuss these omissions in Appendix B and present evidence that they do not bias results. Note that estimation will recover perceived utility parameters prior to migrating; it is possible but unobservable in data that the realized cost of migration differs from what agents expect *ex ante*.

3.2 Measuring Potential Earnings

The earnings term $\mathbb{E}_t[Y_{njt}]$ in the model reflects the net present value of expected future income agent n would earn were they to locate in province j . Observed earnings in data may be a poor proxy for this counterfactual for two reasons. First, measured earnings are a function of both province-specific productivity and individual-specific characteristics. To the extent that there is selection across labor markets—either on individual characteristics or individual-location match quality—observed earnings differences may not reflect the true income an agent would receive upon moving. Second, observed income is not informative about future expectations. Workers may anticipate different streams of lifetime income from provinces with the same current earnings level based on the nature of that earnings. Expectations are further complicated by the fact that agents may subsequently migrate again, and provinces provide different option value to remigration.

I reconcile the discrepancy between counterfactual lifetime earnings in the model and observed current earnings in the data using instrumental variables. Given province-level productivity shifters z_{jt}^k indexed by k , we can decompose individual potential earnings in a single year as

$$y_{njt} = \sum_k \kappa^k z_{jt}^k + \mu_{njt} \tag{2}$$

where κ quantifies the transmission of local productivity into worker earnings, and μ_{njt} includes both residual province-level productivity as well as any worker-specific skill or location match quality. As long as shocks to z^k are orthogonal to how workers sort across provinces (i.e. the realized distribution of μ_{njt}), then they serve as valid instruments for productivity. Hence, consistent estimates of κ can be recovered from a regression of (2).

This relationship can also be projected forward in time, so that the identified portion of local productivity inherits the same time-series properties as the underlying instrument. To relate productivity to expected earnings, we require a second identifying assumption that shocks to z_{jt}^k are uncorrelated with the option value a province provides to remigration. Under this assumption,

such option value appears in the expectation of future μ , and expected future values of $\kappa^k z^k$ remain valid proxies for future earnings.³ This identifying assumption circumvents the need to explicitly model the dynamic location choice such as in [Caliendo et al. \(2019\)](#); [Balboni \(2021\)](#).

Instrumenting for local productivity has the additional benefit of allowing a causal interpretation of the migration response to changes in local earnings. In general, earnings are endogenous because migrants contribute to local labor supply. In the spatial model, endogeneity arises if time-varying shocks to the preference term ϵ_{njt} are correlated across agents. In this case, we would measure a spurious negative relationship between migration and local earnings driven by the correlated supply shock. Instrumenting for earnings using local productivity shifters isolates demand-driven variation in earnings exogenous to supply.

3.3 Instrument Construction

I generate instruments for earnings using global commodity price series. Instruments vary at the province-year level according to each province’s sensitivity to each commodity price. The expected sensitivity to a price shock in a province is computed as the weighted average of each industry’s sensitivity to the shock weighted by local industry intensity.

To construct instruments, I first compute industry sensitivity by regressing earnings in industry ℓ on the price of commodity k in the time series at the national level

$$y_{nt} = \omega^{k,\ell} \tilde{p}_t^k + \varepsilon_{nt}$$

independently for each commodity and industry at the 2-digit level. For the price shock \tilde{p} I use deviation from trend to isolate the unanticipated component of the series, which represents the innovation in that period.

Second, I derive each province’s sensitivity to a commodity as the weighted average of industry sensitivities, weighted by labor force composition. Formally, the sensitivity of province j to commodity k , given 2-digit industries indexed by ℓ , is $\sum_{\ell} \omega^{k,\ell} s_{j0}^{\ell}$. In this expression, s_{j0}^{ℓ} is the share of the labor force of province j employed in industry ℓ in the base year of 1985.

Third, I compute a commodity-province-year instrument value by interacting the cross-sectional sensitivity variation with time-series price variation. Formally,

$$z_{jt}^k = \tilde{p}_t^k \sum_{\ell} \omega^{k,\ell} s_{j0}^{\ell} \tag{3}$$

Generating a province-level shock as the sum of local industry-level shocks is motivated by an underlying assumption of integrated local labor markets. A shock that drives up earnings in one industry must attract workers from other local industries until the local wage equilibrates at a new,

³Empirically, the (unknown) future behavior of μ becomes a nuisance parameter to be estimated.

higher level. Crucially, this integration allows labor markets to be characterized by a unified local wage.

Instrument construction follows a Bartik-style procedure taking a weighted average over industry composition (Bartik, 1991). Instruments based on commodity shocks employ two sources of variation that are orthogonal to individual preferences and therefore plausibly exogenous. The first source of variation comes from time-series changes in global commodity prices. Since Thailand is small relative to the global economy, it is unlikely that prices are affected by local conditions in any individual province.

Second, cross-sectional variation in industry composition is drawn from a reference period that precedes analysis, and therefore is unlikely to have been influenced by subsequent preference-based migration. Furthermore, fixed characteristics of a province that may be correlated with industry composition are absorbed by the amenity term A_j , meaning that violations of exclusion would have to come from time-varying province characteristics correlated with both industry composition and migration. Combining these two sources of plausibly exogenous variation generates valid instruments for local earnings. Any endogenous correlation between y and ϵ must operate through the uninstrumented portion of earnings, represented in (2) by μ .

3.3.1 Selection of Commodity Series

Instruments for local earnings derive from prices of major Thai manufacturing imports, which are inputs into production and therefore influence productivity. I consider every commodity that represented over 1 percent of Thai imports in 1995. This criterion selects oil, cotton, wood, iron, aluminum, and copper. Oil and petroleum products in 1995 represented 6.5 percent of national imports. I use the average spot price of Dated Brent, West Texas Intermediate, and Dubai Fateh price series for analysis. Cotton also represents a large input in the manufacturing sector during the period of study, comprising 1.1 percent of imports in 1995 while clothing and textile products made up 6.5 percent of exports. Prices are derived from the A Index, CIF at Liverpool. Finally, in 1995 wood and lumber represented around 2 percent of Thai imports, while furniture and other wood products comprised a comparable fraction of exports. For analysis I use the Japanese import price for Malaysian Meranti. In addition to these three commodities, I consider iron, aluminum, and copper. Iron accounted for 4.5 percent of imports in 1995, and the latter two metals made up just over 1 percent each. For iron I use the Chinese iron ore import price at Tianjin port; for the other metals I use the London Metals Exchange CIF spot price. Finally, I try a composite index consisting of aluminum, copper, iron ore, lead, nickel, tin, and zinc compiled by the World Bank.

I first verify that each commodity instrument induces variation in provincial earnings. Regression follows almost directly from (2). For a given instrument, suppressing the k index, I run

$$y_{jt} = \kappa z_{jt} + \gamma_j + \gamma_t + \mu_{jt} \tag{4}$$

where each observation is a province-year, κ is the effect of a commodity price shock on local earnings, and γ s are fixed effects. Price series are rescaled so that they each represent comparable shocks to earnings.

Table 2 presents results from the first-stage regression of earnings using oil, cotton, and wood. Regression confirms that each instrument has a significant impact on earnings that persists after controlling for the other instruments.⁴ Unfortunately, the remaining four price series have no predictive power for local earnings, and are henceforth dropped from analysis. Appendix C presents first-stage results using these instruments as well as robustness tests verifying that the main results of this paper are not sensitive to their inclusion.

As discussed above, it is important that instruments identify changes in the potential earnings a prospective migrant would receive and not selection on underlying worker characteristics. Without panel data, it is impossible to directly study unobservable characteristics. Nevertheless, Appendix C provides two pieces of suggestive evidence that instruments identify productivity shocks. First, the first-stage regression is robust to controlling for observable age and education, both of which affect earnings. Second, the regression is insensitive to dropping recent migrants from the province average. Differential selection can occur through new workers entering the market or existing workers departing. By dropping migrants, I verify that the former channel plays no role in first-stage results; the latter is unobservable in my data. These two pieces of evidence suggest commodity-based instruments indeed identify productivity-based labor demand shocks.

3.4 Shock Permanence

To characterize the net present value of identified earnings variation, it only remains to estimate the permanence of shocks to commodity prices. To do so I model each series as the sum of a random walk and a stochastic component, a decomposition common in literature on household earnings (e.g. [Blundell et al., 2008](#)). Although the realized permanence of any individual shock remains unobservable under this decomposition, expected permanence is computed from the relative variance of each component. This modeling technique allows for simple estimation as well as a tractable characterization of earnings utility.

Formally, let each price series, suppressing k indices, evolve as

$$\begin{aligned} p_t &= R_t + s_t \\ R_t &= R_{t-1} + r_t \end{aligned} \tag{5}$$

where r_t and s_t are independent, normally distributed shocks drawn from stationary distributions. With normality, the expected permanence of an observed shock is proportional to the variance of

⁴Since these represent industry inputs, an increase in the price leads to a decrease in productivity, and therefore each series has a negative rescaling factor.

the permanent component. We can define a permanence parameter ρ as

$$\mathbb{E}_t[p_{t+\tau}|p_{t-1}, p_t] = p_{t-1} + \frac{\sigma_r^2}{\sigma_r^2 + \sigma_s^2} \Delta p_t \equiv p_{t-1} + \rho \Delta p_t$$

where Δ denotes differences over time and σ^2 denotes variance. This expectation does not depend on distance into the future τ because of the random walk nature of R_t . As a result, the utility valuation of future earnings collapses to a single term that is insensitive to functional forms that vary the relative weights placed on each future period, such as β - δ style hyperbolic discounting or finite time horizons, except through their effect on the total weight placed on the future relative to the present. (See Appendix B for a derivation.)

The variance of permanent and transitory shocks to each price series is computed empirically from first differences. Note

$$\Delta p_t = r_t + \Delta s_t$$

Because (p_t, p_{t-1}) are jointly normally distributed, this expression implies the moment conditions

$$\begin{aligned} \text{var}(\Delta p_t) &= \sigma_r^2 + 2\sigma_s^2 \\ \text{cov}(\Delta p_t, \Delta p_{t-1}) &= -\sigma_s^2 \end{aligned} \tag{6}$$

For each commodity, (6) estimates the expected contribution of permanent and temporary factors to a change in the price. Observed price shocks have anticipated permanent and temporary components in these ratios.

Table 3 presents the long-run permanence of the three commodity price series through 2000. Crude oil prices appear to be the most permanent: over 98 percent of a typical shock persists into the future. Cotton prices have a similarly high permanent fraction of 91.5 percent. In contrast, permanent and transitory shocks contribute nearly equally to wood prices, meaning roughly half of a typical shock dissipates by the following year.

Given heterogeneous permanence, each instrument corresponds to a different net present value for the same contemporaneous change in earnings. The full discounted value of a shock is largest in absolute terms when driven by the price of crude oil and smallest when driven by the price of wood. If agents are forward-looking, theory predicts that the migration response to earnings shocks induced by these commodities will be proportionately sized.

3.5 Reduced-Form Migration Results

Identification of the choice model relies on workers being sensitive to the anticipated future value of a contemporary shock. Note that this condition does not require every agent to fully understand

the times series properties of every commodity price and its transmission into earnings. Instead, market mechanisms may lead agents to respond appropriately to income variation with differing permanence. For instance, it would be sufficient if local businesses sent signals in their employment offers, or if workers associated greater income permanence with industries and markets more sensitive to commodities of greater price permanence.

I present reduced-form evidence that workers are sensitive to the expected future component of earnings in their migration choice. To do so, I compare two-stage least squares estimates of the elasticity of migration to earnings identified with instruments of differing permanence. Since a more permanent commodity price corresponds to a higher net present value of earnings given the same size contemporaneous shock, the model predicts two-stage least squares using a more permanent instrument will reveal a larger migration response.

I run regressions of the form

$$m_{ijt} = \lambda^y y_{jt} + \lambda^d d_{ij} + \lambda^x y_{jt} \times d_{ij} + \gamma_j + \gamma_{it} + \varepsilon_{ijt} \quad (7)$$

where the unit of observation is a migration channel from origin i to destination j in year t . m_{ijt} is a dummy for observing migrants along the channel⁵, and the regressor of interest, y_{jt} , measures log earnings at the destination after controlling for age and education. Distance between provinces, d_{ij} , is included as a control as well as its interaction with earnings, and γ s represent fixed effects. With origin-year fixed effects, using destination income as a regressor is identical to using the earnings gap between origin and destination. Migration and earnings are seasonally adjusted as discussed in Appendix A.

Results confirm that the migration response to earnings shocks induced by oil and cotton prices is larger than that driven by wood prices. OLS regression of (7) in the first column of Table 4 indicates a one log point increase in earnings at a given location is associated with a 16.21 percent greater chance of observing migrants to that destination along a given migration channel, though this result lacks a causal interpretation for reasons described above. Importantly, columns 2 through 4 estimate the causal migration response to earnings separately using each commodity instrument. Labor migration is more likely to be observed following a more permanent shock, with the regression coefficient being almost three times greater when using the more permanent instruments of crude oil and cotton relative to the more temporary instrument of wood. The final column combines all three instruments and estimates a migration response between the extremes.⁶ Reduced form results are consistent with markets incorporating beliefs about future earnings, justifying the revealed-preference estimation in the next section.

⁵Alternate parameterizations of migration prevalence are discussed in Appendix A.

⁶Appendix Table S1 verifies that results are robust to excluding Bangkok and to omitting the post-financial crisis years of 1998–2000.

4 Structural Estimation

In this section I discuss estimation of the choice model using revealed-preference migration decisions. Mobility parameters of interest are the disutility of migration and variance in idiosyncratic tastes, scaled into dollar terms from the net present value of an earnings shock. Estimation is complicated by two features of the Thai LFS common to many data sources. First, measured migration rates are noisy, with zero migrants observed along many province-to-province channels. Second, sampling is stratified by current province, which is an endogenous outcome of migration. I present a novel implementation of an estimation strategy that explicitly models data sampling subject to aggregate changes in population. I show this method has an equivalent interpretation treating each province-level survey as a random draw from a multinomial distribution with probabilities governed by migration rates.

4.1 Assumptions for Estimation

The value function from the choice model in (1) can be written as

$$V_{nijt} = y_{jt} + \beta \hat{y}_{jt} - c_{ijt} + A_j + \sum_{\tau=1}^T \delta^\tau \mathbb{E}_t[\mu_{njt+\tau}] + \epsilon_{njt}$$

where y_{jt} represents current earnings, $\hat{y}_{jt} = \sum_k \rho^k \kappa^k z_{jt}^k$ represents the identified portion of expected future earnings, μ_{njt} represents the unexplained portion of earnings, and $\beta = \sum_{\tau=1}^T \delta^\tau$ represents the discount value on the future. I make three assumptions to relate the choice model to data.

Assumption 1: Parameterization of expected future earnings.

Let the expectation over the unexplained portion of future earnings be linear:

$$\sum_{\tau=1}^T \delta^\tau \mathbb{E}_t[\mu_{njt+\tau}] = \rho^\mu \mu_{jt}$$

This assumption admits many functional forms including the stochastic/random walk form above or a stationary AR1 process. It requires multiplicative separability between μ_{njt} and any sequence of discounting.

Assumption 2: Linear control function for endogeneity.

Given the exclusion restriction $\epsilon_{njt} \perp z_{jt}^k$, endogeneity between idiosyncratic preferences ϵ_{njt} and earnings y_{jt} must enter through the uninstrumented portion of earnings. Following [Rivers and Vuong \(1988\)](#), we can write

$$\epsilon_{njt} = f(\mu_{nt}) + e_{njt}$$

where e_{njt} is i.i.d. and uncorrelated with other variables. In estimation, I let $f(\cdot)$ be linear.

Assumptions 1 and 2 together imply that μ_{njt} enters the value function linearly; relaxing either assumption admits other functional forms as well. The utility expression becomes

$$V_{njt} = y_{jt} + \beta \hat{y}_{jt} - c_{ijt} + A_j + \phi \mu_{jt} + e_{njt} \equiv U_{jt} - c_{ijt} + e_{njt} \quad (8)$$

where ϕ is a parameter to be estimated that incorporates both the expected permanence of μ as well as endogeneity through the control function. Without further identifying variation, it is impossible to decompose the relative importance of these factors. (8) represents the utility expression to be estimated.

Assumption 3: Distribution of idiosyncratic preferences.

Let e_{njt} follow an i.i.d. extreme value distribution. This assumption gives a familiar closed-form solution for migration probabilities

$$\begin{aligned} m_{ijt} &= \mathbb{P}[e_{njt} - e_{nj't} \geq U_{j't} - c_{ij't} - U_{jt} + c_{ijt}, \forall j' \in \{1, \dots, J\}] \\ &= \frac{\exp(U_{jt} - c_{ijt})}{\sum_{j'} \exp(U_{j't} - c_{ij't})} \end{aligned} \quad (9)$$

where m_{ijt} represents the fraction of the population of province i that moves to j in year t .

Given infinite data, all parameters in (8) could be estimated. However, due to data limitations, estimation is unstable along dimensions that lack sufficient variation. Therefore, I make two further simplifications.

Assumption 4: Parameterization of migration disutility.

Let the migration disutility consist of fixed and distance-based components

$$c_{ijt} = \mathbf{1}\{i \neq j\}C + \eta d_{ij} \quad (10)$$

where d_{ij} is the distance between markets i and j . With enough observations, migration costs are nonparametrically identified up to bilateral symmetry $c_{ijt} = c_{jit}$. However, with the small sample size relative to the number of possible migration channels, there are no observed migrants along many channels. This fact would drive the estimated friction along such channels to ∞ and eliminate their informativeness. It is unlikely that some migrations cause infinite disutility, and indeed unlikely that there were truly no migrants along channels with observed zeroes, so a parametric assumption disciplines estimation.

Assumption 5: Calibration of discounting.

Fix the value placed on future earnings in (8) to

$$\beta = \sum_{\tau=1}^{\infty} (0.82\delta)^{\tau} \tag{11}$$

with a fixed discount rate δ . In principle, β is identified from variation in the permanent and transitory components of earnings shocks. Unfortunately, instruments do not provide sufficient variation for this, a problem common to identification of discount rates in studies of dynamic choice (see [Frederick et al., 2002](#); [Magnac and Thesmar, 2002](#)). With so few instruments, two of which have high permanence, the present and future components of earnings are largely collinear, making estimation of β unstable. To calibrate future valuation, I fix the effective discount rate to be a combination of time preference and the likelihood of re-migration. In what follows, I present results using the average re-migration rate of 18.8 percent found in data combined with discount preferences of $\delta = 0.97, 0.95,$ and 0.9 . With greater variation in the permanence of earnings instruments, the methodology in this paper could recover the valuation workers place on the future.

Under Assumptions 1–5, exogenous parameters to be estimated in (8) and (10) are $\{C, \eta, A_j, \phi, \sigma_e\}$. Distance d_{ij} in (10) is computed using province centroids, and earnings y_{jt} come from the LFS. Expected future earnings \hat{y}_{jt} are constructed from instrument values, rescaled to represent a common earnings shock and multiplied by their expected permanence. Finally, the endogenous component of earnings, μ_{jt} , is estimated as the residual from the first stage regression of y_{jt} on z_{jt}^k in (4).⁷

The model is fit to workers’ observed changes in location, taking the initial population distribution as given. Location-specific preferences are separately identified from migration costs using the difference between gross and net flows between provinces. The net migration response to a change in net present value of expected income scales workers’ geographic preferences to the utility value of earnings. Conditional on net migration, high gross flows between provinces would indicate geographic preferences are not tied to current location, reflecting low mobility costs and higher idiosyncratic location-specific preferences. In contrast, low gross flows would indicate mobility frictions tether workers to their current place of residence.

4.2 Maximum Likelihood Estimation

Conditional on the exogenous variables in (8) and (10), parameters map to a unique set of predicted province-to-province migration probabilities \hat{m}_{ijt} . Estimation searches over the parameter space to fit \hat{m}_{ijt} to observed migration patterns.

One common approach follows [Berry et al. \(2004\)](#) in taking logs of (9) to generate a system

⁷This procedure implicitly lets counterfactual μ_{njt} be a linear function of measured average μ_{jt} .

of linear equations to estimate by regression (e.g. [Diamond, 2016](#)). This method-of-moments minimizes the distance between the model-predicted $\log \hat{m}_{ijt}$ and observed \log migration rates. Such an approach would be inappropriate in this setting because there is sampling error in observed migration probabilities and, in particular, zero observed migration along many channels. Missing channels and sampling error more generally are not random; measured migration depends on both the true migration rate as well as origin and destination province sizes.

This issue also arises in, e.g., retail data with infrequently purchased or granular product categories ([Gandhi et al., 2023](#)), or in trade data with measurement or reporting errors ([Feenstra et al., 2005](#); [UNCTAD, 2012](#)). [Gandhi et al. \(2023\)](#) show estimation bias grows with the probability of observing zeros in the sample and propose a set-identified adjustment to [Berry et al. \(2004\)](#) to recover utility parameters over the range of possible bias. I introduce a different approach that imposes constraints based on aggregate population changes; this would be analogous to using aggregate sales data or import/export quantities in other settings.

An alternative for estimation using small samples is maximum likelihood on individual choice data. This strategy maximizes the joint likelihood that each sampled individual is observed living in their origin in the prior year and in their destination in the current year. The likelihood function can be written as

$$\mathcal{L} = \prod_{nt} \mathbb{P}[i_{nt}, j_{nt}] = \prod_{nt} \mathbb{P}[i_{nt}] \mathbb{P}[j_{nt} | i_{nt}]$$

where i_{nt} and j_{nt} denote individual n 's origin and destination provinces, respectively, in year t . When individuals are randomly sampled from their home location, for instance in a census extract or a prospective sample drawn prior to migration, the origin term does not depend on parameters to be estimated and therefore drops out of the maximization. The second term is simply the destination choice probability, so estimation reduces to maximizing the joint likelihood of each observed migration (or non-migration).

$$\max \mathcal{L} \propto \max \prod_{nt} \mathbb{P}[j_{nt} | i_{nt}] = \max \prod_{nt} \hat{m}_{i_{nt}j_{nt}} \quad (12)$$

This maximization is subject to the constraint that probabilities be non-negative and sum to one, which is met by construction in [\(9\)](#).

In contrast, when sampling is stratified by current location, as in the Thai LFS, [\(12\)](#) no longer returns consistent parameter estimates. Inconsistency arises because sampling probabilities are endogenous to migration rates, an issue described as choice-based sampling by [Manski and Lerman \(1977\)](#). The probability of observing an individual from a given origin location i now depends on model parameters as the migration choice affects the likelihood of being sampled. Similarly $\mathbb{P}[j|i]$ is no longer equal to m_{ijt} because the sample is not a random draw from the past population of

province i . This concern is present in any data stratified by outcome, such as when rare outcomes are oversampled for accuracy.

With choice-based sampling, the likelihood function can be rewritten as

$$\mathcal{L} = \prod_{nt} \mathbb{P}[i_{nt}, j_{nt}] = \prod_{nt} \mathbb{P}[j_{nt}] \mathbb{P}[i_{nt}|j_{nt}]$$

where $\mathbb{P}[j_{nt}]$, the survey frequency in province j at time t , is independent of migration probabilities. The second term in the likelihood function is the probability that a surveyed resident of j came from i . This probability can be expressed as the number of migrants from i in j over the total population of j , which gives the maximand

$$\max \mathcal{L} = \max \prod_{nt} \mathbb{P}[i_{nt}|j_{nt}] = \max \prod_{nt} \frac{N_{i_{nt}t-1} \hat{m}_{i_{nt}j_{nt}t}}{N_{j_{nt}t}} \quad (13)$$

where N_{jt} measures the total population of province j at time t .

Population sizes can be found in external data that don't depend on parameters, so it may appear that the maximands in (12) and (13) are identical. The difference is that estimated probabilities $\mathbb{P}[i_{nt}|j_{nt}]$ are no longer guaranteed to sum to one; this requirement must be imposed as a constraint. Rearranging terms, this set of constraints can be written as

$$N_{jt} = N_{jt-1} - \sum_{j' \neq j} \hat{m}_{jj't} N_{j't-1} + \sum_{i \neq j} \hat{m}_{ijt} N_{it-1} \quad \forall j, t \quad (14)$$

which can be interpreted as a law of motion for population. The population of province j in year t must equal the population in year $t - 1$ minus the number of out-migrants plus the number of in-migrants. Maximizing (13) subject to (14) is a generalization of the choice-based estimator proposed by [Manski and Lerman \(1977\)](#), where total population is a sum of province-specific aggregate choice probabilities. In general, consistency would not be possible without some form of aggregate choice data.

This estimator has an equivalent interpretation that treats each province-year of the LFS as a unit of observation. Each province-year survey is a random draw of n_{jt} individuals from the population of j . Every individual belongs to a bin according to their prior province, with bin frequencies equal to the fraction that migrated from that province $\mathbb{P}_t[i|j] = N_{it-1} m_{ijt} / N_{jt}$. Thus, each province-year survey can be considered a multinomial random variable distributed as

$$\{n_{ijt}\} \sim \mathcal{M} \left[n_{jt}, \left\{ \frac{N_{it-1} m_{ijt}}{N_{jt}} \right\} \right] = \frac{n_{jt}!}{\prod_i n_{ijt}!} \prod_i \left(\frac{N_{it-1} m_{ijt}}{N_{jt}} \right)^{n_{ijt}}$$

where n_{ijt} is the number of individuals observed in province j at time t hailing from province i . The factorial terms in this expression are independent of parameters and therefore drop out of

maximization. Taking a joint probability across survey province-years yields a maximand exactly proportional to (13). Maximization is again subject to the constraint that probabilities must sum to one, meaning (14) remains in place. Interpretation of the maximum likelihood estimate at either the individual or the survey-stratification-unit level produces identical estimators.

5 Results

Results reveal barriers to migration to be on the order of one year’s earnings, but migration remains an important contributor to labor supply elasticity. Eliminating mobility frictions would lower cross-province earnings variation by 20 percent and lead a quarter of the population to relocate. However, the returns would be primarily realized in non-wage utility as production would only increase by 3 percent.

5.1 Mobility Frictions

Estimation scales parameters relative to the utility value of earnings. For reference, average annual earnings in Thailand over the period are \$2,225 in 2015 USD. The standard deviation across provinces is around 45 percent of the mean. Were this variation to remain stable over time, the standard deviation in net present value terms according to (11) at the discount rate of 0.95 would be \$4,702 in 2015 USD. I benchmark estimation results against these reference values. The full range of point estimates are presented in Table 5.

Results indicate mobility frictions play a substantial role in blunting the utility returns to migration. I find the average cost among all possible migrations to be between 1.00 and 1.20 times average annual earnings.⁸ At the preferred discount rate of 0.95, the implied utility cost of migration is 1.14 times average earnings, or 54 percent of the standard deviation of lifetime earnings utility across provinces. Breaking apart this value, 56 percent of the disutility is in C , a fixed penalty for any relocation, and the remaining 44 percent varies with distance and makes farther away locations less attractive. The high fixed component may be because displacement itself is more unpleasant than transit, social networks must be rebuilt no matter how far one travels, or workers are already well-sorted according to tastes.

Heterogeneity in idiosyncratic preferences expands the wedge between local earnings and the utility returns to migration. The standard deviation of e_{njt} is 1.06–1.12 times the standard deviation of lifetime earnings utility. Since this portion of preference is uncorrelated with earnings and distance, it inspires many moves that do not necessarily have high market returns.

Table 6 demonstrates the importance of accounting for the duration of earnings and the sampling procedure. The first column reproduces main parameter estimates. The next two columns present estimates under assumptions that earnings innovations are fully transitory and permanent. These

⁸The average cost over realized migrations is only slightly smaller, ranging from 0.98 to 1.17 times annual earnings.

assumptions bracket the possible valuations that may be inferred from an observed income shock. Parameter estimates range from 25 to 125 percent of the main estimate. The breadth of this range highlights the role expected future earnings play in assessing frictions; assuming full permanence is closer to the truth because the identifying variation has a high degree of permanence.

The final two columns of Table 6 present results under estimation that fails to account for measurement error and endogenous sampling. Column 4 uses a method-of-moments treating observed migration rates as precise measures of the truth. This method overestimates migration frictions, especially as they relate to distance traveled, and underestimates idiosyncratic preference heterogeneity. Bias arises because it treats zero observed migration along a channel to be the truth rather than a random draw from a distribution with positive weight at zero.

The fifth column of Table 6 presents results from maximum likelihood estimation of (12) that uses sampling weights but ignores choice-based sampling. The estimated utility parameters are uniformly smaller than the truth because the estimator places more weight on evidence from smaller provinces. Smaller provinces tend to have lower earnings and see a net outflow of migrants over the period of study, so this strategy overstates the importance of earnings relative to other factors.

5.2 Labor Supply Elasticity

I next compute the migration contribution to labor supply elasticity implied by model parameters. The elasticity of migration between any two provinces can be derived analytically from (9) as

$$\mathcal{E}_{m,y} \equiv \frac{\partial m_{ijt}}{\partial y_{jt}} \frac{y_{jt}}{m_{ijt}} = \sigma_e^{-1} y_{jt} (1 - m_{ijt}) \quad (15)$$

This calculation focuses exclusively on the extensive margin of labor market choice and abstracts from intensive-margin variation in labor hours or days. Model-generated elasticities avoid measurement issues discussed above that can confound reduced-form estimation. Results in this section use a discount rate of 0.95.

Parameter values indicate the average elasticity of migration along any province-to-province channel is close to unity. A one percent increase in earnings at a destination raises immigration from an origin by 0.87 percent on average.⁹

Summing across provinces yields the net migration contribution to local labor supply elasticity. On average, the migration response to a local earnings shock adds between 0.095, in the case of a fully permanent shock, and 0.02, with a fully transitory shock, to local labor supply elasticity. Of this, roughly half (0.051 in the case of a permanent shock) comes from variation in new arrivals, and the remainder is attributed to current residents deciding (not) to relocate.

The migration contribution to labor supply elasticity is economically meaningful compared to

⁹This percentage is assessed on a very small base; on average only 0.075 percent of the population of a province relocates to a given destination in a typical year.

intensive-margin elasticities among non-movers. A permanent one percent increase in earnings would cause a province’s male labor force to grow by 0.095 percent. By comparison, the few existing estimates of male labor supply elasticity in Thailand around the period of study range from 0.21 to 0.44 (Singh et al., 1986; Bauer et al., 1988), and comparable elasticities are estimated for agricultural households in other Asian countries (Singh et al., 1986). Thus, internal migration adds anywhere from 25 to 50 percent to local labor supply elasticity.

5.3 Long-Term Earnings Differentials

I simulate three counterfactual scenarios to quantify the economic importance of mobility frictions. The ideal exercise would be to calculate the counterfactual earnings and population distribution were there no migration costs. Unfortunately, this calculation requires credible estimates of local labor demand in addition to migration elasticities. As workers move in the spatial model, the model equilibrates both because destination earnings fall, lowering the location’s attractiveness, and because new migrants are drawn from farther down the preference distribution. I present simulation results based on different assumptions about the relative size of these two forces.

First, I consider the case where earnings exactly offset differences in local amenities. This counterfactual equalizes location-specific average utility without mobility frictions so that location choice would be determined only by workers’ idiosyncratic preferences. In this counterfactual world the correlation between earnings and amenities¹⁰ is exactly -1 ; in reality the coefficient from a regression of earnings on amenities is -1.12 , significant at the 1 percent level with a regression R^2 of 0.89, revealing excess earnings variation. The standard deviation of earnings falls by just under 20 percent in this counterfactual scenario, indicating that local amenities explain 81–84 percent of the spatial earnings variation in Thailand and mobility frictions account for the rest. Figure 2 depicts observed and counterfactual amenities-based earnings levels side-by-side.

Second, I calculate the earnings levels that would sustain the year 2000 population distribution without migration costs. This counterfactual corresponds to a world with free mobility and perfectly inelastic labor demand at existing population levels.¹¹ Counterfactual earnings are mapped next to actual earnings in Figure 3. The standard deviation of earnings would be 2.3 times greater were migration barriers removed in a world with perfectly inelastic labor demand. The change is predictably driven by increased earnings in larger provinces accompanied by decreases in smaller ones: regressing the difference between counterfactual and actual earnings on province population yields a positive coefficient significant at the 1 percent level with a regression R^2 of 0.85. This

¹⁰Appendix B details how amenity values are computed from estimated parameters.

¹¹I apply the normalization that average national earnings remain unchanged. I drop the province of Bangkok for this exercise because it is an order of magnitude larger than the rest. The model attributes its size to its high initial population that faces barriers to leaving, rather than to excessive amenities. Therefore, the earnings needed to justify the population of Bangkok without barriers to mobility would dwarf all other provinces and overwhelm any meaningful variation.

counterfactual exercise suggests that barriers to mobility play a significant role in maintaining the size of larger provinces.

Third, I consider the opposite extreme of perfectly elastic local labor demand around the year 2000 earnings levels. In this scenario, 26.5 percent of the male working-age population relocates from their province of residence. Of this quantity, 9.3 percentage points consist of workers leaving Bangkok while the remaining 17.3 percentage points are migrations around the rest of the country. Maps of actual and counterfactual province population are given in Figure 4. Relocation is substantially greater than the annual level of 5 percent observed in the data, again suggesting that the disutility of migration deters much desired relocation.

Despite the high level of migration in this counterfactual scenario, national income only grows by 3 percent. This modest productivity effect comes about because labor reallocation has two offsetting motivators. On net, workers move out of provinces with medium levels of earnings and amenities toward both high-income, low-amenity markets and low-income, high-amenity ones.¹² As a result, a substantial fraction of moves in this scenario deliver positive utility but negative earnings returns. Allowing for downward-sloping local labor demand would have an ambiguous effect on this calculation that depends on whether earnings fall faster in high-productivity or low-productivity locations. In either case, the productivity effects of labor reallocation are dampened by the inverse relationship between earnings and non-wage amenities.

6 Conclusion

In this paper I quantify the disutility of migration between provinces in Thailand and evaluate its role in labor misallocation and spatial earnings differentials. I estimate a spatial equilibrium model of worker location choice using revealed-preference migration decisions from 1985 to 2000. The model is identified from exogenous variation in local labor demand based on fluctuations in global commodity prices interacted with local industry exposure to those commodities. The time series characteristics of commodity prices translate observed earnings variation into expected net present value. For estimation I jointly model location choice and data sampling to overcome inconsistency introduced by measurement error and endogenous stratification.

A differential migration response based on earnings permanence is consistent with related findings that households have a higher propensity to consume out of more permanent income shocks (e.g. Paxson, 1992). A large body of work uses this fact to investigate the degree to which households are insured against earnings fluctuations. Labor markets can act as an ex ante counterpart to ex post consumption smoothing. The more households adjust their labor supply in response to demand shocks, the less fluctuation in earnings the local market will experience (Yang, 2004;

¹²I again restrict to reallocation outside of Bangkok. Including Bangkok, the model actually predicts a 5 percent decrease in national product as there is a large flow out of this productive, low-amenity metropolis.

Jayachandran, 2006). Since migration is one channel of adjustment, migration barriers and the associated labor supply elasticity are informative about household insurance against local shocks.

Results of this study indicate that migrating generates a disutility of 1.0–1.2 times average annual earnings. Despite the high effective cost of relocation, migrants play an important role in the labor force, adding up to 9.5 percentage points to local labor supply elasticity. Eliminating barriers to mobility could lower spatial earnings variation by 20 percent and lead up to a quarter of the labor force to relocate. However, this relocation would raise national product by only 3 percent as many moves would be motivated by amenities and locational preferences. Together, these numbers indicate that mobility frictions play a role in sustaining spatial earnings disparities, but the highest returns to lowering frictions would come in the form of non-wage utility rather than labor market productivity.

References


- Acemoglu, Daron and Melissa Dell**, “Productivity Differences Between and Within Countries,” *American Economic Journal: Macroeconomics*, 2010, 2 (1), 169–188.
- Allcott, Hunt and Daniel Keniston**, “Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America,” *The Review of Economic Studies*, 2017, 85 (2), 695–731.
- Alvarez, Jorge A.**, “The Agricultural Wage Gap: Evidence from Brazilian Micro-data,” *American Economic Journal: Macroeconomics*, 2020, 12 (1), 153–73.
- Artuç, Erhan, Shubham Chaudhuri, and John McLaren**, “Trade Shocks and Labor Adjustment: A Structural Empirical Approach,” *American Economic Review*, 2010, 100 (3), 1008–45.
- Balboni, Claire**, “In Harm’s Way? Infrastructure Investments and the Persistence of Coastal Cities,” 2021. Unpublished manuscript.
- Bartik, Timothy J.**, *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute for Employment Research, 1991.
- Baseler, Travis**, “Hidden Income and the Perceived Returns to Migration,” *American Economics Journal: Applied Economics*, 2023.
- Bauer, John, Naohiro Ogawa, and Nipon Poapongsakorn**, “Forecasts of Labor Force and Wages for Thailand,” HOMES Research Report 5, East-West Population Institute 1988.
- Beegle, Kathleen, Joachim De Weerd, and Stefan Dercon**, “Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey,” *Review of Economics and Statistics*, 2011, 93 (3), 1010–1033.
- Berman, Nicolas, Mathieu Couttenier, Dominic Rohner, and Mathias Thoenig**, “This Mine Is Mine! How Minerals Fuel Conflicts in Africa,” *American Economic Review*, 2017, 107 (6), 1564–1610.
- Berry, Steven, James Levinsohn, and Ariel Pakes**, “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, 2004, 112 (1), 68–105.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, “Consumption Inequality and Partial Insurance,” *American Economic Review*, 2008, 98 (5), 1887–1921.
- Bryan, Gharad and Melanie Morten**, “The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia,” *Journal of Political Economy*, 2019, 127 (5), 2229–68.
- , **Shyamal Chowdhury, and Ahmed Mushfiq Mobarak**, “Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh,” *Econometrica*, 2014, 82 (5), 1671–1748.
- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro**, “Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock,” *Econometrica*, 2019, 87 (3), 741–835.

- Clemens, Michael A.**, “Why Do Programmers Earn More in Houston than Hyderabad? Evidence from Randomized Processing of US Visas,” *American Economic Review*, 2013, *103* (3), 198–202.
- **and Hannah Postel**, “Temporary work visas as US-Haiti development cooperation: a preliminary impact evaluation,” *IZA Journal of Labor and Development*, 2017, *6* (4).
- , **Claudio E. Montenegro, and Lant Pritchett**, “The Place Premium: Bounding the Price Equivalent of Migration Barriers,” *The Review of Economics and Statistics*, 2019, *101* (2), 201–213.
- Cosslett, Stephen R.**, “Efficient Estimation of Discrete-Choice Models,” in Charles F. Manski and Daniel L. McFadden, eds., *Structural Analysis of Discrete Data and Econometric Applications*, Cambridge: The MIT Press, 1981, pp. 51–111.
- Desmet, Klaus, Dávid Krisztián Nagy, and Esteban Rossi-Hansberg**, “The Geography of Development,” *Journal of Political Economy*, 2018, *126* (3), 903–83.
- Diamond, Rebecca**, “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000,” *American Economic Review*, 2016, *106* (3), 479–524.
- Dix-Carneiro, Rafael**, “Trade Liberalization and Labor Market Dynamics,” *Econometrica*, 2014, *82* (3), 825–885.
- Dube, Oeindrila and Juan F. Vargas**, “Commodity Price Shocks and Civil Conflict: Evidence from Colombia,” *The Review of Economic Studies*, 2013, *4* (285), 1384–1421.
- Eckert, Fabian and Michael Peters**, “Spatial Structural Change,” NBER Working Papers 30489, National Bureau of Economic Research 2022.
- Feenstra, Robert C., Robert E. Lipsey, Haiyan Deng, Alyson C. Ma, and Hengyong Mo**, “World Trade Flows: 1962-2000,” NBER Working Papers 11040, National Bureau of Economic Research 2005.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue**, “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 2002, *40* (2), 351–401.
- Gandhi, Amit, Zhentong Lu, and Xiaoxia Shi**, “Estimating demand for differentiated products with zeroes in market share data,” *Quantitative Economics*, 2023, *14* (2), 381–418.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh**, “The Agricultural Productivity Gap,” *Quarterly Journal of Economics*, 2013, *129* (2), 939–993.
- Hamory, Joan, Marieke Kleemans, Nicholas Li, and Edward Miguel**, “Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata,” *Journal of the European Economic Association*, 2021, *19* (3), 1522–1555.
- Hendricks, Lutz and Todd Schoellman**, “Human Capital and Development Accounting: New Evidence from Wage Gains at Migration,” *Quarterly Journal of Economics*, 2017, *133* (2), 665–700.

- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi**, “Growth and Structural Transformation,” in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 2, Elsevier, 2014, chapter 6, pp. 855–941.
- James, Alexander**, “The resource curse: A statistical mirage?,” *Journal of Development Economics*, 2015, 114 (C), 55–63.
- Jayachandran, Seema**, “Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries,” *Journal of Political Economy*, 2006, 114 (3), 538–575.
- Kennan, John and James R. Walker**, “The Effect of Expected Income on Individual Migration Decisions,” *Econometrica*, 2011, 79 (1), 211–251.
- Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E. Waugh**, “Migration costs and observational returns to migration in the developing world,” *Journal of Monetary Economics*, 2020, 113 (C), 138–154.
- Magnac, Thierry and David Thesmar**, “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 2002, 70 (2), 801–816.
- Manski, Charles F. and Steven R. Lerman**, “The Estimation of Choice Probabilities from Choice Based Samples,” *Econometrica*, 1977, 45 (8), 1977–1988.
- McGuirk, Eoin and Marshall Burke**, “The Economic Origins of Conflict in Africa,” *Journal of Political Economy*, 2020, 128 (10), 3940–3997.
- McKenzie, David, Steven Stillman, and John Gibson**, “How Important is Selection? Experimental vs. Non-Experimental Measures of the Income Gains from Migration,” *Journal of the European Economic Association*, 2010, 8 (4), 913–945.
- Mobarak, Ahmed Mushfiq, Iffath Anwar Sharif, and Maheshwor Shrestha**, “Migration and Remittances: The Impacts of a Government Intermediated International Migration Program,” *American Economic Journal: Applied Economics*, 2023.
- Moretti, Enrico**, “Local Labor Markets,” in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, 2011, chapter 14, pp. 1237–1313.
- Morten, Melanie and Jaqueline Oliveira**, “Paving the Way to Development: Costly Migration and Labor Market Integration,” NBER Working Papers 22158, National Bureau of Economic Research 2023.
- Paxson, Christina H.**, “Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand,” *American Economic Review*, 1992, 82 (1), 15–33.
- Rivers, Douglas and Quang H. Vuong**, “Limited information estimators and exogeneity tests for simultaneous probit models,” *Journal of Econometrics*, 1988, 39 (3), 347–366.
- Roback, Jennifer**, “Wages, Rents, and the Quality of Life,” *Journal of Political Economy*, 1982, 90 (6), 1257–1278.

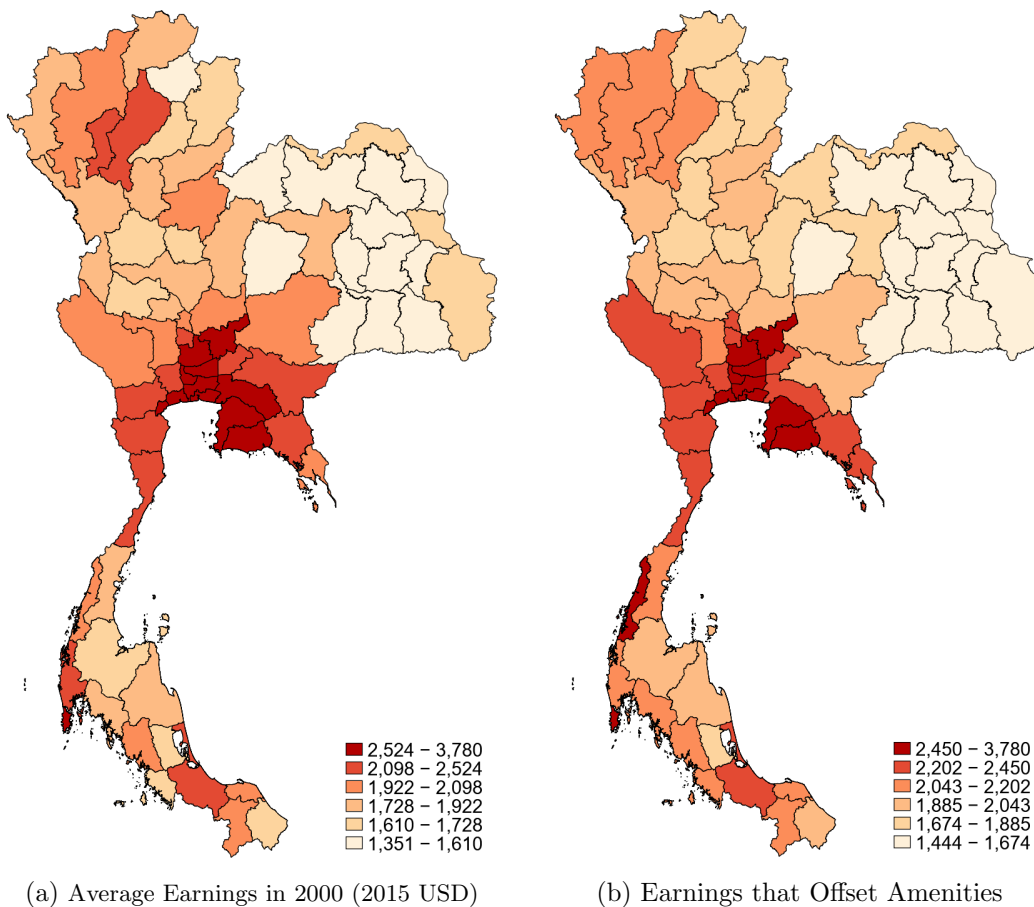
- Rosen, Sherwin**, “Wage-based indexes of urban quality of life,” in Peter N. Miezkowski and Mahlon R. Straszheim, eds., *Current Issues in Urban Economics*, Johns Hopkins University Press, 1979, pp. 74–104.
- Singh, Inderjit, Lyn Squire, and John Strauss**, “A Survey of Agricultural Household Models: Recent Findings and Policy Implications,” *World Bank Economic Review*, 1986, 1 (1), 149–179.
- UNCTAD**, *A Practical Guide to Trade Policy Analysis*, United Nations Conference on Trade and Development, 2012.
- Yang, Liu**, “Unequal Provinces but Equal Families? An Analysis of Inequality and Migration in Thailand,” Unpublished manuscript, University of Chicago 2004.
- Young, Alwyn**, “Inequality, the Urban-Rural Gap, and Migration,” *Quarterly Journal of Economics*, 2013, 128 (4), 1727–1785.
- , “Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services,” *American Economic Review*, 2014, 104 (11), 3635–67.

Figure 1: Population Growth versus Average Earnings



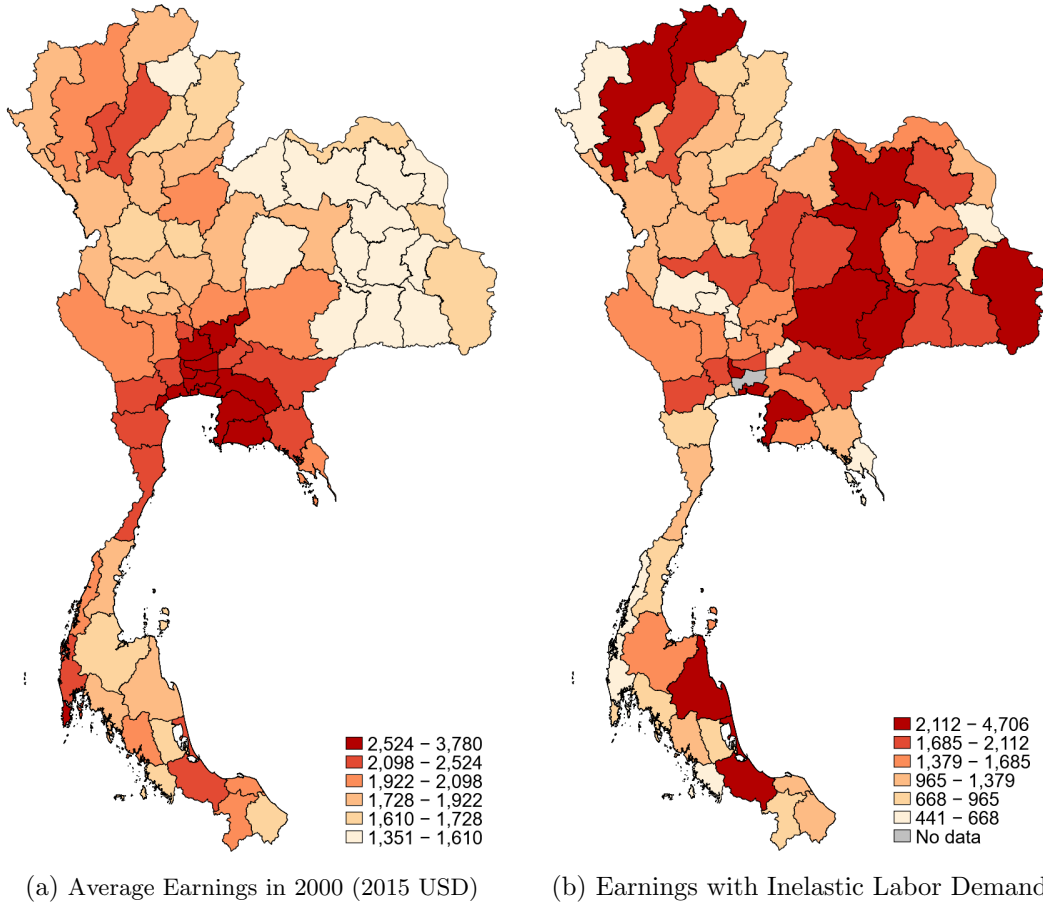
Notes: Scatter plot of population growth from 1985–2000 versus log average earnings in 1985 by province. The slope of the regression line is 18.8 and significant at the 1% level.

Figure 2: 2000 Observed Earnings and Counterfactual Earnings



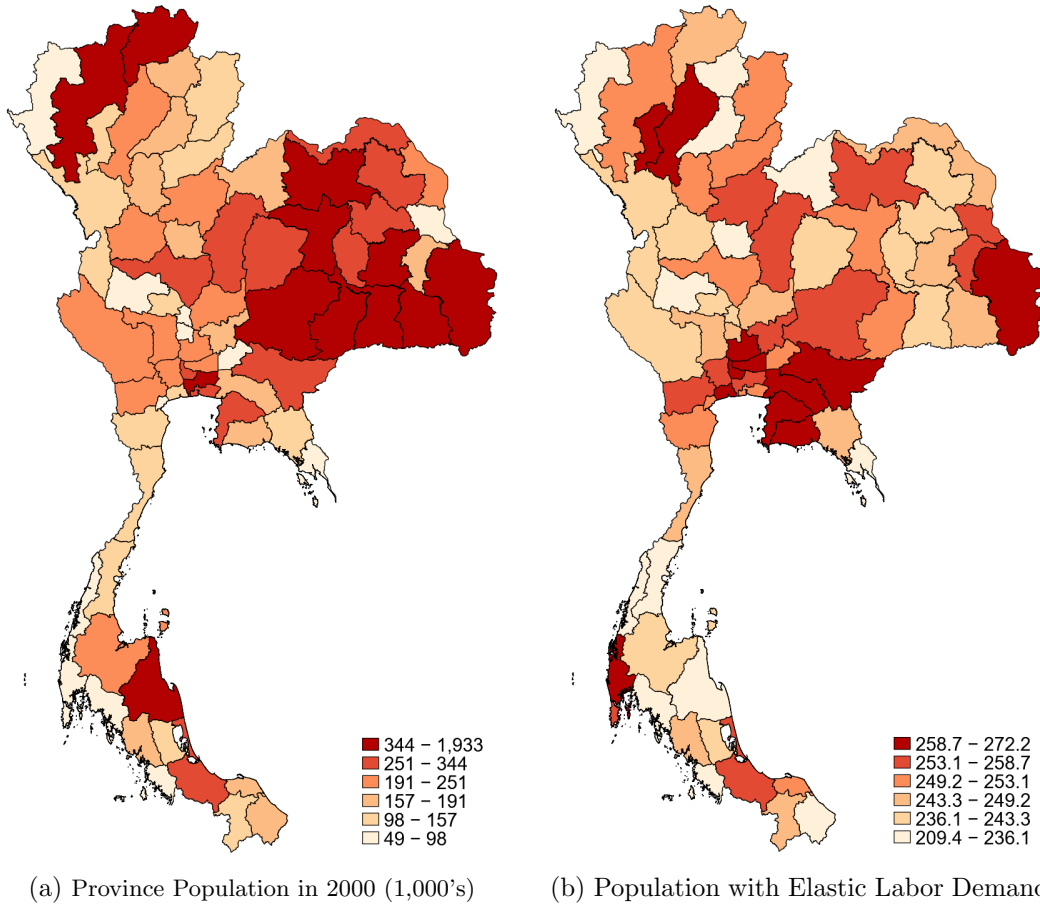
Notes: Observed earnings in the year 2000 and earnings required to perfectly offset province amenity levels, both in 2015 USD. The standard deviation of counterfactual earnings is 18 percent lower. Regressing actual earnings on amenity levels yields a coefficient of -1.12 significant at the 1% level with a regression R^2 of 0.89.

Figure 3: 2000 Observed Earnings and Counterfactual Earnings



Notes: Observed earnings in the year 2000 and earnings required to maintain the same population distribution without barriers to mobility, excluding Bangkok. The standard deviation of counterfactual earnings is 2.27 times greater. Regressing the difference between counterfactual and actual earnings on province population size yields a positive coefficient significant at the 1% level with a regression R^2 of 0.85.

Figure 4: 2000 Observed Population and Counterfactual Population



Notes: Observed population distribution in the year 2000 and the counterfactual population distribution with perfect mobility holding local earnings constant. 26.7 percent of the population relocates under the counterfactual distribution. Population change in the counterfactual has a positive quadratic relationship to province size, meaning large and small provinces grow while medium provinces shrink, significant at the 1% level with a regression R^2 of 0.19.

Table 1: Summary Statistics

	All	Migrants
Age	35.0 (12.0)	29.14 (9.66)
Household size	5.0 (2.0)	4.5 (2.8)
Employment	0.94 (0.23)	0.90 (0.31)
Hours (cond. on work)	52.7 (13.9)	54.4 (13.5)
1985 Earnings (2015 USD)	1,109 (1,954)	1,121 (2,034)
2000 Earnings (2015 USD)	2,041 (2,833)	1,761 (2,247)
At least primary education	0.24 (0.43)	0.26 (0.44)
At least secondary education	0.11 (0.31)	0.09 (0.29)
N	993,227	54,776

Notes: Averages weighted by inverse sampling frequency. Standard deviations in parentheses.

Table 2: Earnings Responsiveness to Normalized Price Shocks

Outcome:	Local Earnings			
Crude Oil	1.00*** (0.24)			0.64*** (0.22)
Cotton		1.00*** (0.23)		0.49* (0.26)
Wood			1.00*** (0.23)	0.60*** (0.21)
Fixed Effects:				
Province	X	X	X	X
Year	X	X	X	X
R Squared	0.93	0.93	0.93	0.94
Observations	1,095	1,095	1,095	1,095
Rescale Factor	-0.016	-0.034	-0.22	
Partial F-Stat	16.8	19.3	19.5	15.1

Notes: First-stage regression of province earnings on province exposure to commodity price shocks. Shock magnitude is rescaled to represent a consistent magnitude for all commodities. Robust standard errors clustered by province in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Permanence of Commodity Prices

	Crude Oil	Cotton	Wood
Fraction Permanent	0.983 (0.21)	0.915 (0.24)	0.512 (0.40)
var(Perm.)	0.36	0.45	0.02
var(Temp.)	0.03	0.08	0.02

Notes: The long-run permanence of various commodity series, running from 1960 through 2000. Data from the World Bank's databank. Bootstrapped standard errors in parentheses.

Table 4: Migration Response to Local Earnings

Outcome:	Migration between Provinces				
Instrument:	OLS	Crude Oil	Cotton	Wood	Combined
Dest. Income	16.21*** (2.57)	47.74*** (15.15)	42.49** (19.94)	16.35 (15.91)	30.64** (13.05)
Distance	18.36*** (3.40)	31.89*** (5.73)	33.08*** (6.77)	26.00*** (5.25)	28.60*** (4.92)
Income×Distance	-3.10*** (0.46)	-4.90*** (0.77)	-5.06*** (0.90)	-4.11*** (0.70)	-4.46*** (0.66)
Don't Move	67.59*** (2.23)	67.44*** (2.19)	67.42*** (2.17)	67.50*** (2.19)	67.47*** (2.18)
Fixed Effects:					
Destination	X	X	X	X	X
Origin×Year	X	X	X	X	X
Population Controls	X	X	X	X	X
R Squared	0.27	0.27	0.27	0.27	0.27
Observations	79,935	79,935	79,935	79,935	79,935

Notes: OLS and IV regressions of a dummy for any observed migration on earnings at destination. Robust standard errors clustered by destination province in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5: ML Estimated Parameters

<i>Calibrated Parameters</i>			
Discount Rate (δ)	0.90	0.95	0.97
Remigration Rate	0.19	0.19	0.19
Time Horizon (T)	∞	∞	∞
<i>Estimated Coefficients</i>			
Fixed Cost (C)	-4.28	-4.85	-5.15
	[-5.45, -3.17]	[-6.19, -3.58]	[-6.56, -3.77]
Distance (η)	-0.61	-0.70	-0.74
	[-0.78, -0.46]	[-0.89, -0.52]	[-0.94, -0.55]
St. Dev. Tastes (σ_e)	1.02	1.15	1.22
	[0.75, 1.29]	[0.85, 1.46]	[0.90, 1.55]
<i>Size relative to earnings</i>			
Migration Disutility/ Avg. Earnings	1.00	1.14	1.20
St. Dev. Tastes/ St. Dev. Earnings	1.12	1.07	1.06

Notes: Estimated structural parameters for a range of discount parameters. All values are in terms of log earnings. Bootstrapped 95% confidence intervals in square braces. The likelihood ratio index for all specifications, relative to a model in which all parameters are 0, is 0.99996. The bottom panel shows cost parameters relative to average earnings and standard deviation of tastes relative to standard deviation of earnings. The bottom row presents the rescaled ratio if observed earnings were perfectly permanent.

Table 6: Parameter Estimates under Alternate Specifications

	Model Estimate	Shock Duration		Alternate Method	
		Transitory	Permanent	GMM	Logit
<i>Discount Rate = 0.90</i>					
Fixed Cost	-4.28 [-5.45, -3.17]	-1.12	-5.24	-4.02	-1.94
Distance	-0.61 [-0.78, -0.46]	-0.16	0.75	-1.01	-0.28
St. Dev. Tastes	1.02 [0.75, 1.29]	0.27	1.24	0.73	0.46
<i>Discount Rate = 0.95</i>					
Fixed Cost	-4.85 [-6.19, -3.58]	-1.12	-5.99	-4.66	-2.21
Distance	-0.70 [-0.89, -0.52]	-0.16	-0.86	-1.17	-0.31
St. Dev. Tastes	1.15 [0.85, 1.46]	0.26	1.42	0.85	0.52
<i>Discount Rate = 0.97</i>					
Fixed Cost	-5.15 [-6.56, -3.77]	-1.12	-6.37	-5.06	-2.36
Distance	-0.74 [-0.94, -0.55]	-0.16	-0.91	-1.27	-0.33
St. Dev. Tastes	1.22 [0.90, 1.55]	0.27	1.51	0.92	0.55

Notes: Estimated parameters with alternate specifications. Columns 2 and 3 assume observed earnings are fully transitory and fully permanent, respectively. Column 4 contains parameter estimates obtained using GMM to match predicted migration to observed migration ignoring measurement error. Column 5 contains parameter estimates obtained using multinomial logit ignoring choice-based sampling.

Acknowledgements and Disclosure Statement

This work was supported in part by a National Science Foundation Graduate Research Fellowship. I declare I have no conflicts of interest.

I am grateful to Abhijit Banerjee, Esther Dufo, and Rob Townsend for their excellent guidance. I also thank Stephen Boucher, Michael Carter, Michael Stepner, and workshop participants at Johns Hopkins, MIT, UC Davis, UC San Diego, and Yale for helpful feedback and advice. I benefited greatly from the help of Jiaqiang Chen and researchers at the University of the Thai Chamber of Commerce in obtaining the data and from excellent research assistance by Yulu Pi. Any mistakes that remain are my own.

Email: shenoy@ucdavis.edu

Supplementary Appendix

For Online Publication Only

A Reduced-Form Outcomes

A.1 Seasonality

Seasonality affects both migration and earnings in the data. Among migrants, there are significantly more migrants who have lived in a location for less than one year in the third-quarter rounds of survey than in the first quarter. This pattern corresponds with the agricultural season in most of the country, during which time many people temporarily stay with relatives in rural areas to help with farm labor. The average ratio of migrants observed in the third quarter relative to the first quarter is strongly correlated with the fraction of a province engaged in agriculture, supporting this pattern. No such seasonality is observed among migrants who have lived in the province for one year or more, suggesting that the swell in migrants in the third quarter is largely temporary seasonal work.

I address seasonality in migration by adjusting migration rates based on the destination and quarter of survey. For each province, I compute the ratio of zero-year migrants to one-year migrants from the previous year in each quarter. I treat the ratio in the first quarter as the true ratio, and an excess or deficit in the third quarter relative to the first to represent seasonal migration. I thus deflate the observed migration flows to a given destination in the third quarter this ratio of ratios. Annual migration is computed as the average of adjusted first-quarter-equivalent migration in each included round of the survey for a given year.

The seasonal pattern in the labor market also generates cyclicity in earnings. Earnings are around two percent higher during the third quarter across all sectors of the labor force. However, this pattern masks different underlying mechanisms by industry. In the lower human capital sectors of agriculture, construction, and retail trade, the wage increase persists even after controlling for years of education. This suggests that higher demand relative to supply boosts wages across the board for these sectors during the thick part of the labor cycle. However, in higher human capital sectors such as manufacturing, utilities, and the public sector, the seasonal pattern in earnings disappears after controlling for education. These sectors also shrink as a portion of the labor force during the thick season. These two facts together suggest that seasonality is driven by occupation selection: in the third quarter the lower end of the earnings distribution switches sectors to meet seasonal demand.

Given the varying drivers of seasonality by sector, I deseason earnings by regressing earnings on a dummy for third quarter, year dummies, and dummies for age and education bin separately for each sector. I then partial out the seasonal component for all analysis. Earnings are also Winsorized at the 99th percentile.

I consider labor market returns in terms of total earnings rather than the hourly wage rate because a significant portion of the population draws income from either family agriculture or self-owned businesses for which hours worked is not well defined. Furthermore, hours are only reported for a single week of work; imputing wage rates from this data would lead to significant measurement error in hours. In contrast, income is computed over a full month and reported in significantly more detail, making it a more accurate measure of labor market returns.

A.2 Measure of Migration

Estimation suffers from high noise due to the sparseness of the survey. To isolate individual local markets, I define a labor market to be a province, of which there are 73 in the country. However, this means that the number of people surveyed is small relative to the number of migrants and total possible migration channels. As a result, many migration flows are missed and coded as 0, which causes difficulty in estimation.

In the data, there are no observed migrants from a given origin to destination in the vast majority of cases. In only 16 percent of origin-destination-year cells is any migration present in the survey. The majority of zeros observed in the data are likely not generated by an underlying truth of no migration, but are more likely caused by the fact that the survey is sparse relative to the number of migrants. This conjecture is supported by the fact that the number of cells with any observed migration is around 10 percent per year in the earlier years with smaller surveys (and only 5 percent in the two years with only one survey round), and climbs to over 30 percent in the later, significantly larger surveys. At the same time, the 95th and 99th percentiles of migration share remain stable over time; only lower percentiles fill in as the survey size increases. Taken together, these two facts suggest that many small migration flows are missed entirely. Much of this paper deals with avoiding bias from the large number of zeroes.

The ideal reduced-form measure of migration would be the percent of the population of i that relocates to j at time t . Were this measured accurately, it would readily generate an estimate of elasticity. However, the sparseness of the survey means it is

The percent migration measure is formally constructed, suppressing time subscripts for simplicity, as

$$m_{ij t} = \frac{\sum_{j_n=j, i_n=i} \mathbb{P}_j^{-1}}{N_i}$$

where n indexes individuals surveyed; i_n and j_n are the individuals' previous and current provinces, respectively; \mathbb{P}_j is the sampling probability of an individual from province j ; and N_i is the total population size of province i .

Taking expectations under the assumption that sampling frequency is independent of the probability of being a migrant,

$$\mathbb{E}[m_{ij t}] = \frac{\sum_{j_n=j} \mathbb{P}_j^{-1} \frac{m_{ij t}^* N_i}{N_j}}{N_i} = \frac{m_{ij t}^*}{N_j} \sum_{j_n=j} \mathbb{P}_j^{-1} = m_{ij t}^*$$

because the probability a given resident of j hails from i is equal to the population fraction $m_{ij t}^* N_i / N_j$ for the true migration rate $m_{ij t}^*$. The last equality follows because the sampling probability in a given province is the sample size divided by the population size. Therefore, the measure constructed this way is an unbiased estimate of the migration probability.

The variance of this measure is

$$\text{Var}[m_{ij t}] = (N_i)^{-2} \sum_{j(n)=j} \mathbb{P}_j^{-2} \frac{m_{ij t}^* N_i}{N_j} \left(1 - \frac{m_{ij t}^* N_i}{N_j}\right) = m_{ij t}^* \left(\frac{N_j}{N_i} - m_{ij t}^*\right) \sum_{j_n=j} (\mathbb{P}_j)^{-2}$$

For small values of m^* this variance is increasing in m^* , the population of the destination province, and the variance of survey probabilities. It is decreasing in the population of the origin province and the size of the

survey. Taking these two expressions together, we can rewrite $m_{ijt} = m_{ijt}^* + \zeta_{ijt}$ for some mean-0 ζ that is independent across observations but not identically distributed.

Using the relationship given by (7), the OLS estimator becomes

$$\hat{\beta} = (X'X)^{-1}X'm = \beta + (X'X)^{-1}X'(\varepsilon + \zeta)$$

with a comparable derivation for 2SLS. This final expression reveals two facts about this regression, also true after instrumenting. First, the sampling method introduces heteroskedasticity, even when ε is homoskedastic, because the variance of ζ depends on sampling and the parameters.

Second, and more worrisome, estimation of β is generally not median unbiased. Although $\hat{\beta} = \beta$ in expectation, the error term ζ is a recentered sum of Bernoulli variables; in the extreme case of constant sampling probabilities it is Binomial. As the sample gets large relative to m^* , it converges to a normal distribution. However, with small samples, in particular when $\frac{m_{ijt}^* N_i}{N_j} < 1$, the majority of the probability mass is on an observation of 0 migrants and the median lies well below the mean. Thus the majority of estimates $\hat{\beta}$ will be below the true value of β .

The median bias is apparent in the IV regression: coefficients follow roughly the same ordering as with the dummy regression, but estimation is noisier and several point estimates are actually below 0. Results are presented in Table S2. To limit the effect of bias in the reduced form, I collapse the outcome measure into a dummy for any observed migration. This greatly lowers the variance of the measure and brings the mean closer to the median, at the cost that the coefficient no longer has a ready interpretation.

B Modeling Assumptions

B.1 Uncertainty in Individual Earnings

In the model as written, workers only have preferences over their expected earnings at each possible destination; higher order moments do not appear in the utility function. Notably, the model abstracts from labor market risk. In reality, worker earnings may vary due to uncertainty in the availability of employment or type of job. Although this type of idiosyncratic risk may be quantitatively significant, it is difficult to capture empirically due to both measurement and identification challenges. However, I provide some evidence that it is a second order concern in relation to average earnings.

Labor market risk only matters in the empirical exercise to the extent that it varies with both province and year. Any persistent component of province-level earnings risk will enter into the province fixed effect or amenity term for estimation. Similarly, any time-varying uncertainty that is consistent throughout the country will be subsumed into a time fixed effect and will not affect the location decision. Unfortunately, the residual uncertainty cannot be measured, nor is there sufficient exogenous variation to separately identify it.

Labor market uncertainty appears in the data as variation in earnings. However, in a repeated cross section, uncertainty cannot be separated from heterogeneity in worker unobservable characteristics. It is impossible to determine whether an individual with low earnings is facing an unlucky year or is a persistent low earner. Panel data would allow for measurement of province-level uncertainty by controlling for worker fixed effects. Were a measure with such variation possible, it still suffers from potential endogeneity and would need to be identified with exogenous instruments.

Empirically, the effect of risk on migration appears to be small relative to expected earnings. As a proxy for market risk, I take the income variance within a given province-year cell. This measure is a composite of idiosyncratic risk and idiosyncratic variation in unobserved earnings potential. Incorporating this measure into the IV version of (7), with results presented in Table S1, shows that risk is likely unimportant in two ways. First, the point estimate indicates that the migration effect of a standard deviation change in earnings is an order of magnitude smaller than a standard deviation change in level. Second, the point estimates on the effect of earnings levels are largely unchanged when variance enters the regression. These two facts suggest that estimation of the model without explicitly incorporating risk will be consistent because, to the extent that market uncertainty enters migration considerations, it seems to have little effect on the decision with respect to earnings.

B.2 Local Prices

Typical formulations of the spatial equilibrium model include local prices in the utility function. Their inclusion is motivated by the fact that prices mediate the utility of earnings and may be endogenous to population through local markets such as real estate. This paper omits prices from the empirical model due to lack of reliable localized price data. This omission may threaten estimation if prices are affected by the commodity instruments used for identification. I investigate whether this is a concern using data from the SocioEconomic Survey (SES). The SES is a biennial household survey that includes household demographics, earnings, and expenditures. Province earnings levels reported in the SES closely track those in the LFS, suggesting the samples are comparable.

I compute average housing prices from the SES and include them as the dependent variable in the first stage regression (4). Results are presented in Table S3. First-stage coefficients on local housing prices are uniformly smaller than coefficients on earnings, and not statistically significant. From this exercise, I cannot reject that local prices are unaffected by the commodity instruments, justifying their use as identifying variation for earnings.

B.3 Expected Future Utility

By modeling shocks as the sum of a transitory and permanent component, the weight placed on the expected future value of a contemporaneous earnings shock can be expressed as a single parameter representing the sum of all future discount rates $\sum_{\tau=1}^T \delta^\tau$. This is because the permanent component of each series enters as a random walk so that

$$\begin{aligned} \mathbb{E}_t[\tilde{p}_{t+\tau}^k] &= \mathbb{E}_t[\tilde{p}_{t+\tau'}^k] \quad \forall \tau, \tau' > 0 \\ \implies \sum_{\tau=1}^T \delta^\tau \mathbb{E}_t[w_{jt+\tau}] &= \left(\sum_k \rho^k \kappa^k z_{jt}^k + \bar{w}_j \right) \sum_{\tau=1}^T \delta^\tau + \sum_{\tau=1}^T \delta^\tau \mathbb{E}_t[\mu_{jt+\tau}] \end{aligned}$$

That is, all future periods have the same expected price level given the present information set, which means that in expectation all future periods have the same predictable component of earnings. Because of this uniformity, the forward-lookingness term $\sum_{\tau=1}^T \delta^\tau$ can be collapsed into a single parameter; variation in T or δ do not interact with differences in expected earnings in different periods. As a result, the maximum likelihood estimate of the cost of migration is not sensitive to alternate functional forms that vary the relative

weights placed on each future period, such as β - δ style hyperbolic discounting or finite time horizons, except through their effect on the total weight placed on the future relative to the present.

B.4 Computation of Amenities

I back out the local amenity values from the province-specific estimates \hat{A}_j by adjusting for average earnings and discounting. In estimation, \hat{A}_j is the sum of local amenities A_j and province average earnings \bar{y}_j in net present value terms. The latter is estimated consistently as a byproduct of the first stage regression (4) and can therefore be subtracted from \hat{A}_j . Since current earnings are included in the value function in full, the NPV value of earnings in \hat{A}_j excludes the present. Amenities, like all other parameters, are estimated in present-value terms and must be adjusted accordingly. Formally, the full adjustment can be written as

$$\hat{A}_j = \frac{\beta}{1-\beta}\bar{y}_j + \frac{1}{1-\beta}\alpha_j$$

for a discount factor β , average earnings level \bar{y}_j , and a single-year amenity value α_j .

C Instrument Validity

C.1 Commodity Instrument Selection

First-stage regressions of local earnings on iron, aluminum, copper, and a metals composite from (4) indicate that these commodities do not sufficiently influence local earnings. Table S4 reports first-stage estimates for these four price series. No series is a significant predictor of local earnings, nor are all four jointly significant. Due to the lack of a first stage, I drop these instruments from analysis.

Inclusion of these instruments does not substantially change any estimates. Column 2 of Table S1 presents results from the IV version of (7) using the four insignificant price series as additional instruments. Due to their lack of identifying variation, their inclusion does not affect the main estimates. Table S5 maximum likelihood estimates with extra instruments included in addition to the primary three instruments, using a discount rate of 0.95. Again the inclusion of insignificant instruments does not qualitatively affect the outcome.

C.2 Selection on Unobservables

The empirical exercise relies on assuming that an observed change in the average earnings within a province is reflective of a change in workers' earnings potential in that province. However, the observed earnings level is a function of both underlying province productivity as well as unobserved characteristics of the population in the province. Only the underlying productivity affects workers' earnings potential in the province; changes in population characteristics should not affect an individual worker's location choice. Therefore, to accurately interpret the empirical result, the instruments must identify earnings changes driven by province productivity rather than by changes in the population.

To more precisely characterize the desired variation, consider a simple model where the potential earnings of a workern in province j can be decomposed into a province-specific productivity \mathcal{Y} and a worker-specific

piece ξ . Suppressing time subscripts, let

$$y_{nj} = \mathcal{Y}_j + \xi_{nj}$$

Note that ξ can include both the worker’s portable skills that affect earnings in every province as well as worker-province match quality. Averaging across workers, the observed earnings level is

$$\tilde{y}_j = \frac{1}{n} \sum_{n \in j} y_{nj} = \mathcal{Y}_j + \frac{1}{n} \sum_{n \in j} \xi_{nj}$$

Clearly, an observed change in \tilde{y}_j may be caused by either a change in productivity \mathcal{Y}_j or by a change in the population of j that affects the distribution of x_{inj} among residents. However only changes to \mathcal{Y}_j translate into changes in potential earnings y_{nj} .

Panel data would allow for the inclusion of individual fixed effects in the first-stage regression. Adding fixed effects would control for unobserved individual characteristics ξ_{nj} among non-movers and verify that all remaining variation is due to underlying productivity. Lacking panel data, I cannot separately quantify local productivity and the distribution of individual characteristics. Instead I present two pieces of evidence that suggest the instruments identify earnings variation that reflects underlying productivity.

First, the first stage regression is insensitive to controlling for observable worker characteristics. In the main specification, province earnings are computed as the average of worker residual earnings after controlling for age and education. Since these two observable factors are known to influence earnings, the main specification controls for them to avoid any selection along these dimensions. Table S6 presents results from the first stage regression in (4) without controlling for these observable characteristics. The results, while large, are similar to those of Table 2. In other words, the identifying variation for the IV strategy is not driven by selection on observables, a fact that can hopefully be extrapolated to unobservables.

Second, the the first stage regression is insensitive to dropping migrants. Changes in the distribution of unobservables ξ_{nj} may be caused by new workers entering the population, i.e. immigration, or by old workers leaving the population, i.e. outmigration. Excluding both of these groups would leave a stable subpopulation with a stable distribution of x_i ; earnings changes for this subpopulation will be driven by local productivity rather than individual characteristics. It is impossible to identify future outmigrants in retrospective data, but recent immigrants can be excluded. Table S7 presents first stage regression results dropping recent migrants from the province earnings level. The results here are almost identical to those in Table 2, suggesting that the iv strategy does not rely on variation in the distribution of ξ caused by selective immigration. It remains possible that the results are driven by selective outmigration, but each outmigrant from a province is an immigrant in a different province. It is unlikely that migrant departures would be selected on unobservables while their destinations are not.

Migration Response to Local Earnings					
	IV	More Instr.	Pre-1998	No Bangkok	Var(Y_{jt})
Dest. Income	30.64** (13.05)	25.79** (11.83)	24.34*** (9.330)	18.00 (11.86)	28.88** (12.20)
Distance	28.60*** (4.918)	29.04*** (4.892)	27.30*** (4.274)	32.66*** (4.504)	29.05*** (4.530)
Income×Distance	-4.460*** (0.658)	-4.518*** (0.655)	-4.246*** (0.569)	-5.005*** (0.603)	-4.520*** (0.607)
Don't Move	67.47*** (2.183)	67.47*** (2.183)	69.94*** (2.177)	68.89*** (1.635)	67.47*** (2.186)
Income Variance					-0.0599 (0.0927)
Fixed Effects:					
Destination	X	X	X	X	X
Origin×Year	X	X	X	X	X
Population Controls	X	X	X	X	X
R Squared	0.271	0.272	0.278	0.227	0.271
Observations	79935	79935	63948	77760	79935

Table S1: Column 1 reproduces the IV results from Table 4. Column 2 presents IV results after including prices for aluminum, copper, iron ore, and an international metals index as additional instruments. Columns 3 and 4 present results from the IV regression in (7) dropping years 1998 and onward and excluding the province of Bangkok, respectively. Column 5 adds province-level earnings variance as an additional regressor. Robust standard errors clustered by destination province in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Commodity Instruments: Migration Percent					
	OLS	Crude Oil	Cotton	Wood	Combined
Dest. Income	-0.0748** (0.0349)	-0.0931 (0.150)	0.206 (0.165)	-0.0786 (0.0891)	-0.0176 (0.0849)
Distance	-0.265*** (0.0838)	-0.188* (0.100)	-0.136 (0.107)	-0.0932 (0.0716)	-0.137* (0.0787)
Income×Distance	0.0320*** (0.0111)	0.0217 (0.0132)	0.0148 (0.0142)	0.00914 (0.00941)	0.0150 (0.0103)
Don't Move	95.53*** (0.335)	95.53*** (0.330)	95.53*** (0.330)	95.53*** (0.330)	95.53*** (0.330)
Fixed Effects:					
Destination	X	X	X	X	X
Origin×Year	X	X	X	X	X
R Squared	0.998	0.998	0.998	0.998	0.998
Observations	79935	79935	79935	79935	79935

Table S2: IV regressions of the portion of a origin province that migrates to a destination on earnings at destination. Robust standard errors clustered by destination province in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Housing Price Responsiveness to Commodity Shocks

Crude Oil	0.523 (0.583)			0.501 (0.561)
Cotton		0.500 (0.803)		0.261 (0.897)
Wood			0.436 (0.759)	0.356 (0.803)
Fixed Effects:				
Province	X	X	X	X
Year	X	X	X	X
R Squared	0.893	0.892	0.892	0.893
Observations	584	584	584	584

Table S3: First-stage regression of local housing prices on province exposure to commodity price shocks using data from the SocioEconomic Survey (SES). Shock magnitude is rescaled to represent a consistent one-unit earnings shock for all commodities. Relative to the effect on earnings, coefficients for housing prices are smaller and insignificant. Robust standard errors clustered by province in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Earnings Responsiveness to Commodity Price Shocks					
Aluminum	1.000 (46.08)				-98.13 (112.1)
Copper		1.000 (1.380)			1.097 (1.576)
Iron Ore			1.000 (208.8)		150.5 (176.9)
Metals Index				1.000 (2.405)	4.275 (6.484)
Fixed Effects:					
Province	X	X	X	X	X
Year	X	X	X	X	X
R Squared	0.932	0.932	0.932	0.932	0.932
Observations	1095	1095	1095	1095	1095
Rescale Factor	0.000650	0.0287	0.0000452	0.00810	
Partial F	0.000471	0.525	0.0000229	0.173	0.500

Table S4: First-stage regression of province earnings on province exposure to commodity price shocks. Shock magnitude is rescaled to represent a consistent magnitude for all commodities. Robust standard errors clustered by province in parentheses. *** p<0.01, ** p<0.05, * p<0.1

ML Estimates with Extra Instruments Included

<i>Discount Rate = 0.90</i>			
Fixed Cost	-4.28	-4.56	-4.31
	[-5.45, -3.17]		
Distance	-0.61	-0.65	-0.62
	[-0.78, -0.46]		
St. Dev. Tastes	1.02	1.08	1.02
	[0.75, 1.29]		
<hr/>			
<i>Discount Rate = 0.95</i>			
Fixed Cost	-4.85	-5.15	-4.88
	[-6.19, -3.58]		
Distance	-0.70	-0.74	-0.70
	[-0.89, -0.52]		
St. Dev. Tastes	1.15	1.22	1.16
	[0.85, 1.46]		
<hr/>			
<i>Discount Rate = 0.97</i>			
Fixed Cost	-5.15	-5.47	-5.17
	[-6.56, -3.77]		
Distance	-0.74	-0.78	-0.74
	[-0.94, -0.55]		
St. Dev. Tastes	1.22	1.30	1.23
	[0.90, 1.55]		
<hr/>			
Additional Included Instruments:			
Iron Ore		X	
Aluminum		X	
Copper		X	
Metals Index			X

Table S5: Maximum likelihood estimates with other instruments included. Adding insignificant instruments does not significantly change any results.

Non-Residualized Earnings Responsiveness to Commodity Shocks

Crude Oil	2.202*** (0.374)			1.813*** (0.324)
Cotton		1.177*** (0.333)		0.251 (0.328)
Wood			1.561*** (0.388)	0.767** (0.359)
Fixed Effects:				
Province	X	X	X	X
Year	X	X	X	X
R Squared	0.906	0.902	0.904	0.907
Observations	1095	1095	1095	1095
Partial F	34.72	12.51	16.17	16.59

Table S6: First-stage regression of province earnings on province exposure to commodity price shocks without controlling for age and education. Shock magnitude is rescaled to represent a consistent magnitude for all commodities. Robust standard errors clustered by province in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Non-Migrant Earnings Responsiveness to Commodity Shocks

Crude Oil	1.077*** (0.266)			0.742*** (0.246)
Cotton		0.992*** (0.237)		0.479* (0.274)
Wood			0.995*** (0.242)	0.561** (0.219)
Fixed Effects:				
Province	X	X	X	X
Year	X	X	X	X
R Squared	0.931	0.930	0.931	0.932
Observations	1095	1095	1095	1095
Rescale Factor	-0.0155	-0.0340	-0.222	
Partial F	16.37	17.47	16.85	15.87

Table S7: First-stage regression of province earnings on province exposure to commodity price shocks excluding earnings of recent migrants. Shock magnitude is rescaled to represent a consistent magnitude for all commodities. Robust standard errors clustered by province in parentheses. *** p<0.01, ** p<0.05, * p<0.1