Local Adjustment to Immigrant-driven Labor Supply Shocks

Joan Monras*

Universitat Pompeu Fabra, CREI, Barcelona GSE, and CEPR

July 13, 2020

Abstract

When comparing high- to low-immigrant locations, a large literature documents small effects of immigration on labor market outcomes over ten year horizons. At the same time, the literature has also documented short-run negative effects of immigrant-driven labor supply shocks, at least for some groups of native workers. Taken together, these results suggests that there are mechanisms in place that help local economies recover from the short-run effects of immigrant shocks. This paper introduces a small open city spatial equilibrium model that allows to decompose the local adjustment through various channels. I document that by 1990 Miami had fully recovered from the Mariel Boatlift shock and I use the model and new estimates on internal migration patterns to show that around 50 percent of this recovery was due to internal migration.

Key Words: International and internal migration, local shocks, local labor demand elasticity, technology adoption.

JEL Classification: F22, J20, J30

*Correspondence: jm3364@gmail.com. All errors are mine.

1 Introduction

A large part of the literature evaluates the labor market effect of immigration by comparing changes in local labor market outcomes between high- and low-immigrant locations over ten year horizons using Census data. To identify the causal effect of immigration most of this literature uses the immigrant networks instrument. Using this strategy most papers report very small effects on outcomes such as wages and employment rates, both when looking at particular types of native workers, or when focusing on the "average" native worker (Altonji and Card, 1991; Borjas, 2003; Borjas et al., 1996; Card, 2009; Cortes, 2008).

At the same time, the literature has also documented short-run negative effects of immigrant-driven labor supply shocks at least for some groups of native workers (Borjas, 2017; Monras, 2020). Taken together, these results suggests that there are mechanisms in place that help local economies recover from the short-run effects of immigrant shocks.

Previous literature has investigated some of the mechanisms that may help local economies absorb immigrant shocks. For example, Lewis (2012) uses the Census of Manufactures to assess whether plants facing immigrant-driven increases in the number of high-school dropouts adopt fewer machines per worker. His estimates suggest that a 1 percentage point increase in the fraction of high-school dropouts hired in a plant leads to a decline in plant-level machinery adoption of about 6 percent. Similarly, Clemens et al. (2018) explain the lack of employment gains for natives when the Bracero program was removed by the patterns of technology adoption in response to immigrant shocks. As explained in Lewis (2012) and Lewis (2013), the adoption of forms of capital that substitute low-skilled labor tends to attenuate the effect that immigrant-driven changes in the skill mix have on the returns to relative skills.

Other papers have instead focused on how immigrant-induced changes in factor mix can be absorbed in the context of multisector economies. In open economy models, it is sufficient to expand the sectors using more intensively the type of labor brought by immigrant inflows. These type of models require intense across sector relocations which are typically not found in the data both in the US and in other countries. See for example the pioneering work by Hanson and Slaughter (2002) and Lewis (2003), and the papers by Dustmann and Glitz (2015) and Gonzalez and Ortega (2010)

A few recent papers in the literature, have emphasized that internal relocation may be behind the fast absorption of immigrants into local labor markets. For instance, Monras (2020) suggests that internal migration responded to the unexpectedly large inflow of Mexican immigrants following the Mexico Peso crisis of the 90s, while Amior (2020) provides systematic evidence that internal migration plays a crucial role. His estimates suggest that internal migration may account for the full adjustment. This recent literature seems to contradict earlier accounts of the role of internal migration in dissipating immigrant-driven labor supply shocks (Card and DiNardo, 2000).

Hence, providing a framework for thinking about the relative importance of internal migration and other factors in dissipating immigrant-induced labor supply shocks may be helpful for the advancement of the literature. This is the main contribution of this paper.

In the first part of the paper, I introduce a spatial equilibrium model that provides a structural, yet simple, framework to quantify the importance of different factors in dissipating immigrant shocks. The model represents a "small open" city with two sectors. The first sector produces a tradable good using labor and other factors. The second sector, produces a non-tradable good that satisfies the local demand for housing combining land and the tradable good as inputs. The model incorporates the key elements that help to analyze the effect of immigration on local welfare, measured through the indirect utility function, while taking into account local labor and housing markets.

Under the assumption that in the short-run local technologies, factors, and internal migration does not adjust, and hence adjustment comes through factor and rental price changes, and that in the long-run indirect utility needs to recover the pre-shock level for the economy to return to the spatial equilibrium, I propose a method to translate short-run estimates on the effect of immigration on wages and longer-run estimates on internal migration, to evaluate the relative importance of the different adjustment mechanisms. The intuition behind the exercise is simple. If during a certain period of time, say a decade, we see little changes in factor and rental prices and a lot of internal migration then factors other than internal migration such as technology adoption were perhaps not very important in dissipating wage shocks. If instead we see little internal migration and yet wages, housing prices (and hence indirect utility) recover pre-shock levels then the role of factors other than internal migration such as technology adoption is likely to be much more important.

In the second part of the paper, I study the Mariel Boatlift episode through the lenses of the model. I first document that wages in Miami declined relative to a number of control groups in the first few years after the shock, as has been shown in previous literature. Second, I show that by 1990 wage and rental price changes in Miami were similar to the rest of the US during the decade – not always emphasized in previous papers. Hence, over longer time horizons, I show that wages and rental prices in Miami of *all types* of workers were similar to those in the rest of the country, despite the large inflow of low-skilled immigrants at the beginning of the decade and the evidence pointing at short-run low-skilled workers' wage declines and rental prices increase, as documented in Saiz (2003). This evidence suggests that by 1990 Miami was back into equilibrium and had fully absorbed the immigrant-driven labor supply shock.

Next, I document that while the share of low-skilled workers increased one to one with the inflow of Cuban immigrants in the early 1980s, by the end of the decade it had increased by only .6 low-skilled workers for each low-skilled Cuban immigrant. More precisely, I document that the share of low-skilled workers increased on impact with the Mariel shock, stayed high until 1985, and then declined until 1990 although it remained higher than it was in 1980. The beginning of the decline in the share of low-skilled workers living in Miami coincides with the period when short-run wage effects are estimated to be larger, suggesting that internal migration might have contributed to the dissipation of wage effects. This is the first paper that systematically documents the internal migration response that followed the Mariel Boatlift.

Through the lenses of the model, I estimate that around 50 percent of the indirect utility recovery after the shock is explained by internal migration. This result is robust to a number of alternative estimates of the key parameters of the model, which include the local labor demand elasticity (Borjas, 2017; Card, 1990; Clemens and Hunt, 2018; Peri and Yasenov, 2019), the share of income devoted to housing, the local housing supply elasticity, and long-run internal migration estimates.

This paper is closely related to some of my previous work, most prominently Monras (2020). In Monras (2020), I use the Mexican Peso crisis of 1995 to estimate a dynamic spatial equilibrium model with many locations. Using the estimated model I study the counterfactual evolution of wages under two different assumptions on technology adoption. Relative to Monras (2020), in this paper I extend the conceptual framework and I apply it to the small open economy case of Miami to show how the model can be used to quantify the relative importance of different adjustment mechanisms. On top, the empirical evidence and full exercise also hopefully offers a new perspective on what we learn from the Mariel Boatlift, beyond the discussion on the exact size of the short-run impact.

Related literature

This paper is related to the papers that investigate the link between technology adoption and immigrant shocks (Cascio and Lewis, 2018; Clemens et al., 2018; Lafortune et al., Forthcoming, 2015; Lewis, 2012), since as I argue below, among the factors other than internal migration that help through the recovery, technology adoption is likely to play an important role. Relative to these papers, I offer a model-based measure of the role that technology adoption may play in dissipating the wage effects of immigrant-driven labor supply shocks using data from the Mariel Boatlift episode. The evidence that I present complements this body of prior work. An important difference is that this previous work focuses on how technology or capital adoption can reduce the effect of immigrant shocks on *relative* factor rewards. Instead in this paper, I use the spatial equilibrium assumption to back out how technology adoption and other factors may mitigate the effects of immigration on the *level* of wages.

This paper is also related to the work that has discussed the internal migration responses to immigrant shocks. Borjas et al. (1997) argue that the small estimated effects of immigrant shocks across metropolitan areas may be related to internal migration. Card and DiNardo (2000) show that on average internal migration responses to immigrant shocks are small. Peri and Sparber (2011) corroborate this evidence by defending Card and DiNardo (2000) empirical strategy in contraposition to Borjas (2006). In a recent paper (Albert and Monras, 2020), we argue that the reason why previous literature has found limited evidence for internal migration responses to local shocks is related to two facts. On the one hand, immigrant shocks tend to occur in expensive locations, where, as we show, it is easy for natives to respond by relocating. On the other hand, the immigrant networks instrument tends to put weight on small metropolitan areas close to the Mexican border hence resulting in lower internal mobility estimates than when using other identification strategies. In all, Albert and Monras (2020) show that natives relocate in response to immigrant flows, something that I also find in this paper using variation from the Mariel Boatlift.

Finally, this paper is related to the large literature on the Mariel Boatlift. Card (1990) uses this natural experiment to assess the effect of immigration on the labor market. Using a group of four comparison cities – Tampa, Houston, Atlanta and Los Angeles – Card (1990) reports no differential effect of Cuban immigrants on wages.[1] It is hard to emphasize the importance that this study has had in shaping our thinking about immigration, and more broadly, about using natural experiments in economics. However, Borjas (2017) posed an important challenge to what we had learnt from the Mariel Boatlift episode. Two main things differentiate Borjas' analysis from the original Card (1990). First, he concentrates on studying the wage dynamics of native male workers in Miami in the lowest education group. Second, Borjas (2017) criticizes the control group of cities used in Card (1990) mainly on the grounds that Card chose the control group based on employment trends that included some of the years post-Mariel shock. The conclusion in Borjas (2017) seems to be radically different than in Card (1990). Whereas the initial analysis emphasized that native workers in Miami were not affected by the immigrant shock relative to workers in the control group, Borjas (2017) concludes that there is at least one group of workers that was severely affected. Wage declines for this group are estimated to be as large as 30 percent.

[1] Card distinguishes by racial groups and quartiles in the wage distribution but not by education groups.

Since Borjas' reappraisal, several papers have investigated the episode in detail. The debate has been over two different issues. On the one hand, the micro-level number of observations of male high-school dropouts which are used to compute wage trends is small, in many occasions below 30 individual observations. This means that average wages are not computed with much precision and hence, small changes in the sample of workers used to compute these average wages may have substantial effects on the point estimates. This has been, at least in part, the critique emphasized in Peri and Yasenov (2019) and Clemens and Hunt (2018). On the other hand, there has been some debate over what is the best possible control group of cities (Peri and Yasenov, 2019). The pool of potential control cities is not large, since in the early 1980s there are only 44 metropolitan areas that are covered by the March supplements of the Current Population Survey (CPS) data. Hence, small changes on the metropolitan areas that are used as a control group also lead to large changes in point estimates. None of these previous papers, however, looks at internal migration using the Mariel Boatlift episode. In this paper, I try to take into account the diversity of estimates by showing how the results change when deviating from my baseline estimates rather than taking a stance on what is the best estimate in the literature.

2 Model

In this section I introduce an "open city" spatial equilibrium model of a local labor market, which in the application below will be Miami. It is "open city" because it is a model of just one city that is small relative to the rest of the aggregate economy, hence if workers in Miami leave the city, they are small in numbers relative to workers outside Miami so that they have negligible effects. The model is a spatial equilibrium model in the sense that there is an outside level of utility that workers in Miami can attain if they migrate to another US city.

I assume that there are two sectors in the local economy: a tradable and a non-tradable sector. The tradable sector combines labor and other factors to produce a final good. The non-tradable sector, which can be thought as housing, uses the tradable good and land as inputs to produce homes.

The model focuses on just on type of worker. I assume that these workers are perfect substitutes to immigrants. Other types of labor can be easily introduced as I will explain in what follows and in Appendix A. I also highlight how to analyze the effect of an immigrant shock on workers that are not perfect substitutes to immigrants, although I discuss this point in Appendix B rather than in the main text.

2.1 General setting

Utility

The utility function of a representative worker is given by:

$$U(Y,T) = AY^{1-\alpha}H^{\alpha}$$

Where Y is the tradable good, H is housing, and α is the Cobb-Douglass weight of housing. A denotes the level of amenities in the location. The budget constraint is given by $Y + pH \leq w$.

Utility maximization allows me to compute the indirect utility function. Assuming that the price of the tradable good is the numeraire we have that indirect utility can be represented by:

$$\ln V = \ln A + \ln w - \alpha \ln p \tag{1}$$

Workers can either live in this local labor market and obtain indirect utility $\ln V$ or move elsewhere and obtain $\bar{u}$ instead. Miami is small relative to the rest of the economy in the sense that no matter how many workers leave Miami or move towards Miami, $\bar{u}$ is unaffected. Note also, that since workers do not have disutility from working they supply inelastically their labor endowment.

Tradable sector

The local labor market is defined by the local production function of a representative, perfectly competitive firm that produces a tradable good using the following technology:

$$Y = F(L, O) \tag{2}$$

where Y denotes total output, L denotes labor – which competes with immigrant inflows (in the application below low-skilled workers) – and O is a vector of other factors in the production function – which can include capital, and various other types of labor. $F(L, O)$ is a neoclassical constant returns to scale production function. To keep notation simple I omit specifying explicitly terms like Hicks-neutral technologies, factor-augmenting technologies, or capital. All those could be included explicitly instead of implicitly, see Appendix A.

The representative firm maximizes profits taking factor prices w and w^o as given:

$$\max F(L, O) - wL - w^o O$$

where w is a scalar and denotes the wage of (in the application below, low-skilled) workers and w^o is the vector of prices of the other factors of production.

From the profit maximization problem we obtain the (inverse) demand for labor. This is given by:

$$\ln w = \ln F_L(L, O)$$

Where $F_L(L, O)$ is the marginal product of labor. Note that $F_L(L, O) < 0$.

A first order approximation of this function is given by $\ln F_L(H, L) = \varepsilon - \varepsilon^L \ln L$, where ε is a "residual" that includes high-skill labor, technology, and capital. It is a first order approximation to the extent that I omit interactions between labor and the different factors of production.

Hence (to a first order approximation), we have that:

$$\ln w \approx \varepsilon - \varepsilon^L \ln L \tag{3}$$

This equation relates wages – in the application below, wages of low-skilled workers – to the supply of that factor (since, when in the location, workers supply labor inelastically). It shows that the (inverse) local demand function can be decomposed in two terms. An intercept ε and a slope ε^L multiplied by $\ln L$. The first term captures all the ways in which the aggregate demand for labor can change in a local labor market. These include technological change, changes in industrial composition, or capital adjustment. ε captures the fact that changes in all of these aspects can change the demand for labor. For example, if capital increases and capital and labor are complements, then the intercept ε will be higher. If technology changes, so that labor is favored, the intercept ε will also be at a higher point. I introduce a specific production function in Appendix A with various factors of production and technology parameters to make this point more explicitly.

In contrast ε^L captures the local labor demand elasticity. This is the elasticity of wages with respect to labor holding everything else constant. This measures by how much wages will decline holding technologies and all other factor of production fixed.

Housing sector

The housing sector provides housing for the workers under consideration, which in the application below are low-skilled workers. I assume that housing for these workers is independent of housing for other types of workers.

Construction uses as inputs the final tradable good and land in a Cobb-Douglass production function.[2] I assume that the final tradable good share in production is denoted by η. In this case the supply of housing ($H^S(p)$) is proportional to the housing price p raised to $\frac{1}{1-\eta} = \epsilon$, i.e. $H^S(p) = Hp^{\frac{1}{1-\eta}} = Hp^{\epsilon}$, where H is a positive constant.

The demand for (low-skilled type) housing is given by αwL – as can be derived by maximizing the Cobb-Douglass utility function introduced above subject to the budget constraint. Hence, the equilibrium in the housing market is given by:

$$\alpha wL = pH^S(p) = Hp^{1+\epsilon}$$

or in logs:

$$\ln p = \frac{1}{1+\epsilon}(\ln \alpha - \ln H + \ln w + \ln L) \tag{4}$$

Note that the housing sector is effectively capturing the effect that immigrant shocks may have on the local aggregate demand. When more immigrants enter the economy they expand the demand for housing. Local production of housing reacts to this by increasing the supply of housing.

2.2 Equilibrium

The equilibrium in this model is defined by $\ln V = \bar{u}$. This relationship determines the amount of workers in the local economy. To solve the model we need to determine how $\ln V$ depends on L. For this, we need to use equation 4 and plug it into the indirect utility to obtain:

$$\ln V = \tilde{A} + (1 - \tilde{\alpha}) \ln w - \tilde{\alpha} \ln L$$

where $\tilde{\alpha} = \frac{\alpha}{1+\epsilon}$ and $\tilde{A} = \ln A - \frac{\alpha}{1+\epsilon}(\ln \alpha - \ln H)$. This step shows how we can use the housing sector part of model to get rid of housing prices in the indirect utility function.

Moreover, we can use equation 3 to obtain:

[2] Alternatively, I can assume that it uses labor, but this formulation helps me to avoid to deal with workers' sector location decisions.

$$\ln V = \tilde{\varepsilon} - \tilde{\varepsilon}^L \ln L \tag{5}$$

with $\tilde{\varepsilon}^L = ((1-\tilde{\alpha})\varepsilon^L + \tilde{\alpha})$ and $\tilde{\varepsilon} = \tilde{A} + \varepsilon(1-\tilde{\alpha})$

Equation 5 shows that indirect utility (similar to what happens with wages) can be decomposed in two terms: an intercept and a slope multiplied by $\ln L$. The intercept captures all the ways in which the location is attractive to workers. This includes all aspects that affect local amenities and local demand for labor (net of housing costs) such as local aggregate demand for non-tradable final goods and local technologies. The slope captures all the sources of congestion. More workers add pressure to labor and housing markets.

2.3 Properties

In this subsection, I study what happens to this local economy when there is an inflow of immigrant workers that compete with native workers in the labor and housing markets. In Appendix B, I analyze what happens when native workers are imperfect substitutes to immigrants.

Effect of an immigrant shock

To study the effect of an immigrant shock we need to take the derivative of indirect utility with respect to the size of the immigrant shock, which I measure as $\pi = \frac{I}{L}$, where I is the amount of immigrants that arrive to the local labor market and L is the amount of existing workers in that market. Following from equation 5 we have that:

$$\frac{\partial \ln V}{\partial \pi} = \frac{\partial \tilde{\varepsilon}}{\partial \pi} - \tilde{\varepsilon}^L \frac{\partial \ln L}{\partial \pi} = \frac{\partial \tilde{\varepsilon}}{\partial \pi} - \tilde{\varepsilon}^L \frac{1}{L} \frac{\partial L}{\partial (I/L)} = \nu - \tilde{\varepsilon}^L \lambda$$

where $\lambda = \frac{\partial L}{\partial I}$ measures how many workers stay in the location per immigrant arrival and where $\nu = \frac{\partial \tilde{\varepsilon}}{\partial \pi}$ measures how all other factors that may help to accommodate immigration react to the shock.

At this point it may be worth discussing exactly what ν may be capturing. The simplest interpretation of ν is that it represents an outward shift in the demand for labor. This can be driven by either technological change that increases the productivity of labor or by adjustments in the demand for other factors of production.

Hicks-neutral technological parameters are unlikely to capture the recovery in the demand for labor. For example, when there are various factors of production – like low- and high-skilled labor – then Hicks-neutral technology shifts the demand for both types of labor in exactly the same way. Hence, if one type of labor's indirect utility is affected by the shock and the other is not, changes in Hicks-neutral parameters alone will not be able to return the economy to the pre-shock levels for both types of labor simultaneously. Similar arguments apply to third factors of production. If the elasticity of substitution between capital and low-skilled labor is the same than the elasticity of substitution between capital and high-skilled labor, then adjustment in capital usage alone cannot restore the equilibrium to both types of labor. I illustrate this point in Appendix A.

Other factors like the increased demand for local goods that necessarily comes with immigration (after all immigrants also consume), is already captured in the model through the housing sector. A broader non-tradables sector could be modeled in exactly the same way as housing is introduced. Hence, ν does not capture this channel.

Instead, as I make explicit in Appendix A, factor-biased technological parameters or capital that substitutes low-skilled labor are the most likely candidates behind what ν is capturing.

Short-run

In order to use the model to read the empirical results introduced below, in section 3, I make the following assumption. I define the short-run as a sufficiently short period of time so that there is no time for internal migration and other adjustment mechanisms to help absorb immigrant shocks. Hence, in the short-run any immigrant-induce labor supply shock is absorbed through prices (either wages or rental prices). I.e. a time when $\lambda = 1$ and $\nu = 0$. This assumption captures the idea that internal mobility and other forms of adjustment, like technology or capital adoption are sluggish and often involve large adjustment or fixed costs and occur only after indirect utility has changed. Under this assumption we have that in the short run:

$$\frac{\partial \ln V}{\partial \pi} = -\tilde{\varepsilon}^L$$

Hence, the parameter $\tilde{\varepsilon}^L$ determines by how much indirect utility (of workers directly competing with immigrants in the labor and housing markets) declines with the immigrant shock. Indirect utility is not directly observable. However, under the assumptions of the model we have that $\tilde{\varepsilon}^L = ((1 - \frac{\alpha}{1+\epsilon})\varepsilon^L + \frac{\alpha}{1+\epsilon})$.

We have good estimates of α and ϵ in the literature. Davis and Ortalo-Magne (2011) estimate α at around 0.25. Saiz (2010) provides estimates of the housing supply elasticities for various metropolitan areas. His estimate for Miami is 0.6 – which is among the lowest estimates across metropolitan areas.

Given that we have estimates for α and ϵ we can use the wage regressions presented above to obtain an estimate of ε^L. Hence, given an estimate of ε^L we have that $\tilde{\varepsilon}^L = (1 - \frac{.25}{1+0.6})\varepsilon^L + \frac{0.25}{1+0.6}$. If we were to consider all non-tradables beyond housing we may want to consider what happens to the results with higher levels of α. Note, however, that when $\varepsilon^L = 1$, α and ϵ do not matter.

Long-run

In the long-run, and given the small open city assumption, indirect utility is back to $\bar{u}$. Hence, we have that:

$$0 = \frac{\partial \ln V}{\partial \pi} = \nu - \tilde{\varepsilon}^L \lambda \Rightarrow \nu = \tilde{\varepsilon}^L \lambda \tag{6}$$

This equation says that in the long-run internal migration and other factors respond sufficiently so that indirect utility recovers its pre-shock level. Moreover, if amenities are fixed, a sufficient condition for indirect utility to return to pre-shock levels is that wages and housing prices recover from the shock.

As a result of this equation we have that if we know when the long-run is and we have an estimate of λ and $\tilde{\varepsilon}^L$, we can back out how much all other factors contributed to the absorption of the immigrant-induced labor supply shock. Note, furthermore, that if we had data on some of the other factors, like technology adoption, we would be able to use this framework to quantify its importance explicitly.

Internal migration and other factors

We can use the model to think about the role that internal migration and other factors play in dissipating indirect utility effects of immigration. To do so, it is perhaps useful to illustrate the model with a graph. The y-axis of Figure 1 displays indirect utility levels and the x-axis employment. Initially the market equilibrium is given by the intersection of the initial indirect utility curve (D_L^*) and the initial supply of labor (L^*). The initial market equilibrium is, thus, point A in the figure. At that point indirect utility is at $\bar{u}$. With an unexpected immigrant supply shock, the labor supply curve moves to the right, which in the figure is shown as L^1. Before internal migration and other factors respond, real wages drop and move indirect utility to point B. Using the observed drop in wages and the size of the labor supply shock ($L^1 - L^*$) we can compute the local labor demand elasticity ε^L. Given the assumptions on the relationship between ε^L and $\tilde{\varepsilon}^L$, this wage change allows me to recover the slope of the function D_L^* that moves indirect utility from V^* to V^1.

Figure 1 goes around here

After this initial shock both internal migration and other absorption mechanisms react to bring the economy back to the initial level of indirect utility, in the figure, point D. In the data, we can see how much internal migration responds. This is, we can estimate the difference between L^1 and L^{**}. If only internal migration was contributing to dissipating indirect utility effects we would have that the equilibrium would be at point C_1 and hence at a level of indirect utility which is below the initial one. Hence, it must be that other factors change so that the the indirect utility curve moves from D_L^* to D_L^{**}. This is my proposed estimate of ν. In the Figure we can see directly the importance of all these other factors that contribute to the absorption of immigrants by looking at point C_2 which is the level of indirect utility when internal migration is shut down.

The graph helps to show that we can decompose the indirect utility recovery between the contribution of internal migration and all other factors. This is, we can obtain the indirect utility function D_L^{**} from the estimate of ν. By evaluating indirect utility with the immigrant shock at this level of demand we can compute the level of indirect utility that would prevail if there was no internal migration. This is given by the the level V^2 in the figure. Then, we can compute the difference between V^* and V^1 which is the total short-run indirect utility change, and decompose the recovery as moving from V^2 to V^*, which is the part explained by internal migration, and from V^1 to V^2, which is the part explained by all other factors.

3 Empirical application

In this section, I show how we can read the evolution of local labor markets following immigrant shocks through the lenses of the framework proposed in Section 2. In particular, in this section I document how the large inflow of Cubans that arrived to Miami in 1980 with the Mariel Boatlift likely resulted in a decline in real wages that fully recovered by 1990. I trace internal migration during the decade of the 1980s in response to these local changes. This allows me estimate the key parameters of the model that helps us understand how an unexpected immigrant shock moved the initial (spatial) equilibrium and the forces that brought the economy back into equilibrium.

3.1 Data

I use standard sources of publicly available data. To analyze the short-run effects of the Mariel Boatlift episode I use the March supplements and the outgoing rotation group files of the CPS. The March supplements of the CPS have complete information on wage income during the year prior to the interview and weeks worked, which allows to construct weekly wages. It also contains information on the education level of the individuals in the sample. In particular I can construct 4 education codes: high-school drop-outs, high-school graduates, some college, and college graduates or more. These four groups split the labor market of 1980 in roughly four equally sized groups.

To compute wages I use the exact same sample as Borjas (2017). In particular, I restrict the sample to non-Hispanic prime-age, i.e. 25 to 59 years old, working males. During the 80s women were fast entering the labor market. Hence, when using women to compute wage trends it may be that wage changes are driven by changes in the composition of workers from year to year. This is why I prefer to use only male workers. Including women in the regressions leads to similar results, although substantially more noisy. Arguably we would like to exclude foreign-born individuals if the object of interest is native wages. Birth place is not recorded in the CPS data until 1994, and hence the best approximation is the Hispanic variable, which allows to identify Hispanics of Cuban and of Mexican origin.

An alternative data set to compute wages during this period is the outgoing rotation group files of the CPS. I apply the exact same sample selection when using these data. The number of pre-shock years available in the CPS ORG files is only 1979 and 1980 (which is driven by the coverage of metropolitan areas), whereas the pre-shock years when using the March CPS data include 1975 to 1980.

To study internal migration I trace the share of workers of a certain characteristic that live in Miami. This share could change for reasons other than internal migration. For instance, it could be that mortality rates for say, high-school drop-outs were higher in Miami than in other cities, leading to a decrease in the share of low-skilled workers in Miami. Alternatively, it could be that international migration from places other than Cuba is driving this relative share. From the view point of the model, it does not matter what is driving the change in the composition of workers in Miami. Hence, labeling all worker movements as internal migration is just one way to speak to changes in the relative supply of workers across metropolitan areas. When documenting internal migration I rely on the March CPS data.

To estimate longer-run effects on wages and internal migration, I use the Censuses of 1980 and 1990, provided by Ruggles et al. (2016). From these censuses I can construct weekly wages in 1980 and 1990, following the sample selection applied to the CPS data. I can also obtain a measure of the the size of the Mariel shock. For that I follow Borjas and Monras (2017). In particular, I use data from the 1990 Census on Cuban immigrants arriving in 1980 and 1981 (since these two years are grouped into a single category), which were residing in Miami in 1985, to estimate the number of Cuban migrants that moved to Miami during the Mariel Boatlift. The assumption is that Cubans observed in Miami in 1985 are unlikely to have changed residence during the first five years of the decade and hence represent a good proxy for the size of the shock. If anything we can imagine that the shock was larger than estimated with the 1990 Census data. The Census data allows me to compute the relative size of the shock for each education group, since the Census in 1990 records the educational attainment of the Cuban immigrants. Summary stats tables for these data are provided in Borjas (2017) and Borjas and Monras (2017).

3.2 Identification

In what follows, I run two types of regressions. On the one hand, I use the Mariel Boatlift shock in a standard difference-in-difference setting. The key identification assumption in this case is that Miami would have followed a similar trend than that followed by the control group. Difference-in-difference specifications are quite standard. I use graphical representations of the treatment dummy in each year to analyse the trends in Miami, and in Miami relative to various control groups. I follow Card (1990), Borjas (2017), and Peri and Yasenov (2019) in using three alternative sets of metropolitan areas to construct the control group. I define as the Card control group the metropolitan areas used as control in the initial Card study: Atlanta, Houston, Los Angeles, and Tampa. Borjas proposed an alternative group of metropolitan areas: Anaheim, Rochester, Nassau-Suffolk, and San Jose. In light of this disagreement on the optimal control group, Peri and Yasenov (2019) argue that it is better to construct a synthetic Miami, following Abadie and Gardeazabal (2003). Matching the pre-trends based on weekly wages, the share of low-skilled workers, the share of Hispanics, and the share of manufacturing workers in the labor force, they obtain that a synthetic control for Miami in 1980 consists of New Orleans (43.3%), New York City (30.1%) and Baltimore (24.9%). I define the Peri - Yasenov control group as these three metropolitan areas. I do not directly report results using the synthetic control method due to the difficulties of using this approach in this context.[3] I also report results comparing Miami to all the other identifiable metropolitan areas (43).

On the other hand, I use specifications where I leverage the intensity of the treatment, i.e. where I focus on Cuban-induced increases in the working force of specific factor types of different intensity. More specifically, I estimate equations of the following type, which can be derived directly from the local labor demand equation 3:[4]

$$\Delta \ln y_{ce} = \alpha + \beta \frac{\text{Cub}_{ce}}{\text{Nat}_{ce}} + \delta_c + \delta_e + \varepsilon_{ce} \tag{7}$$

Where c indexes metropolitan areas and e indexes the four education groups: high-school drop-outs, high-school graduates, some college, and college graduates or more. δ_c and δ_e are metropolitan area and education fixed effects, respectively.

As it is well known, this equation identifies the effect of immigrant shocks on outcomes of interest if immigrant location patterns are uncorrelated to the error term. In practice, this is unlikely to be the case. There may be unobserved local labor demand shocks that drive immigrants and improve outcomes of interest like wages. Hence, the need for an instrument.

In this paper I use an instrument inspired in the standard networks instrument used in the literature. The fist stage regression can be expressed as follows:

[3] As argued in Abadie (2020), synthetic control groups work best when the pre-shock period is long and when the pool of donors is large. In this case, the options are a pre-period length which spans 1973 to 1979, with a pool of donors of 33 metropolitan areas, or a pre-shock period of 1976 to 1979 with a pool of donors of 43 metropolitan areas. Moreover, the number of observations in many of these metropolitan areas is small and hence pre-shcok variables are measured with error, which further complicates the use of synthetic control methods in this episode.

[4] Note that in the post-shock period Equation 3 implies that $\ln w \approx \varepsilon - \varepsilon^L \ln L = \varepsilon - \varepsilon^L \ln(N+I) \approx \varepsilon - \varepsilon^L \ln(N) - \varepsilon^L \frac{I}{N}$. Hence, when comparing the pre- to the post-shock periods we obtain: $\Delta \ln w \approx \varepsilon - \varepsilon^L \Delta \ln(N) - \varepsilon^L \frac{I}{N}$. This specification may be problematic when there is substantial skill downgrading, as argued in Dustmann et al. (2013) and Dustmann et al. (2016). Skill-downgrading means that highly educated immigrants are allocated to highly educated natives while instead they are competing in the labor market with low-educated ones. This is not a concern here since a very large share of Cuban immigrants had very low education levels. When skill-downgrading is not a thread, a specification like the one given by equation 7 directly identifies the parameter of interest (from the view point of the model), while other specification like measuring the immigrant shock relative to the overall labor force do not.

$$\frac{\text{Cub}_{ce}}{\text{Nat}_{ce}} = \alpha + \beta \frac{\text{Cub}_{ce,0}}{\text{Nat}_{ce,0}} + \delta_c + \delta_e + \varepsilon_{ce} \tag{8}$$

where Cub_{ce} is the inflow of Cuban workers that arrived in each metropolitan area during the Mariel Boatlift episode with education level e and natives is the size of the local labor force excluding Cuban workers.[5]

The most standard way to use the immigrant networks IV is to assign the flow of immigrants from each country of origin according to the initial distribution of immigrants across metropolitan areas. As argued in Goldsmith-Pinkham et al. (2018), in this setting identification mostly comes from the "shares". A more direct way to use the identifying variation is to directly predict the inflow by the initial share: $\frac{\text{Cub}_{ce,0}}{\text{Nat}_{ce,0}}$. This variable is the size of the Cuban stock relative to the local population at the initial period, in this case 1980, i.e. before the Mariel Boatlift. This variable captures an intensity of treatment, i.e. it measures how important are Cubans (relative to natives) in each metropolitan area-education cell.

If the initial importance of Cubans across cells is uncorrelated with current changes in outcomes of interest then this identification strategy identifies the causal effect of actual Cuban inflows on the variables of interest. Running this regression in the period of the Mariel Boatlift ensures that Cuban inflows are generated by a push, rather than a pull, factor, and hence unlikely to be related to developments in the US economy.

3.3 Short-run estimates

On April 20, 1980, Fidel Castro declared that Cuban nationals could emigrate freely from the port of Mariel. Around 125,000 Cubans took this opportunity and migrated towards the United States during the period that goes from April 23rd to the month of October of 1980. Nearly 70,000 immigrants likely settled in Miami, something that accounts for around 8 percent of the workforce of Miami at the time. Cuban immigrants were very low-skilled. As much as 62 percent lacked a high-school diploma compared to around 23 percent among the natives. Hence, these low-skilled workers experienced a labor supply shock of around 32 percent of the workforce prior to the shock (Borjas and Monras, 2017).

I start the analysis of the Mariel Boatlift episode by analyzing what happened to wages and to the share of low-skilled workers in Miami over the 1980s. This replicates and extends the results reported in Borjas (2017). I refer the reader to Saiz (2003) for an analysis of the short-run effect of Cuban immigrants on Miami's housing market. He shows that rental prices increased on impact in Miami relative to various control groups.

To study how wages of low-skilled workers changed in Miami with the Mariel Boatlift I first use the following difference-in-difference specification:

$$\ln w_{i,c,t} = \delta_c + \delta_t + \beta \text{Post Mariel}_t \text{ x Miami}_c + \gamma X_{i,c,t} + \varepsilon_{i,c,t} \tag{9}$$

where $\ln w_{i,c,t}$ is the wage of worker i in city c at time t, Post Mariel_t is a dummy variable that takes value one after 1980, Miami_c is a dummy variable that takes value one for Miami, and where δ_c and δ_t are city and time fixed effects respectively. I run this regression using only high-school drop-outs. $X_{i,c,t}$

[5] I identify Mariel immigrants as those immigrants arriving to the US in 1981-1983 Census category, as reported in the 1990 Census. I also identify in the 1990 Census the location of each individual in 1985, which I take as a proxy of the location of arrival.

are individual level controls. Note that I can use in equation 9 an interaction of the time fixed effects with the dummy for Miami, instead of Post Mariel$_t$ x Miami$_c$, to plot exactly where the estimate of β comes from.

Equation 9 captures the causal effect of immigration on wages in the short-run as long as the control group is comparable to the treated group. In the particular case of Miami, we have only one treated location and hence inference is complicated from the fact that there may be serial correlation in outcome variables and that we only have one treated location and at most 43 control cities (which is the number of cities available in the CPS data). I report robust standard errors that allow for heteroskedasticity.[6]

The results are reported in Panel A of Figure 2. I report estimates for Miami, in an event-type setting and estimates for Miami relative to four different control groups: the original Card control group – Atlanta, Houston, LA, and Tampa –, the control group proposed in Borjas (2017) – Anaheim, Nasssau, Rochester, and San Jose–, a control group based on Peri and Yasenov (2019) – which is reported only in the regression results in Table 1 and includes New Orleans, New York City, and Baltimore –, and a control group that includes all the metropolitan areas in the US for which we have data in the early 1980s.

Figure 2 goes around here

Irrespective of the control group that I use, Figure 2 shows that there are no systematic trends in the wage evolution in Miami leading to the arrival of the Mariel Boatlift immigrants. Wage declines are small in the first two years after the shock and significantly increase in magnitude thereafter. The largest impact is around 1985 or 1986. After this, wages recover so that by 1990 there is no differential impact in Miami relative to the various control locations. There are many reasons that may explain why wages did not react on impact, but rather after one or two years. It could be that local technologies adapted to the shock, although this alone has a hard time explaining why they later decline. It could also be that there is some wage stickiness, so that wage effects are only observable when new contracts are negotiated. A final explanation could be that it took a couple of years for the Mariel immigrants to enter Miami's labor market, perhaps because they needed to learn English or other specific skills. Whatever the reasons, it seems that there is a decline in wages of the least skilled workers in Miami which may be related to the unexpectedly large flow of immigrants during these years. As explained in Borjas and Monras (2017) the wage decline is only observed for the least skilled native workers. In fact, labor market outcomes of more skilled workers in Miami improved relative to the control groups.

Panels A and B of Table 1 quantify the wage effects using a number of alternative specifications that follow equation 9. In the first column of Panel A, I estimate the wage effects of the Mariel Boatlift using all other 43 metropolitan areas as control group. The second column uses only the Card control, column 3 uses the Borjas control, and column 4 uses Peri - Yasenov's control group. I repeat the estimates in columns 5, 6, 7 and 8 but adding individual level controls (most importantly a dummy for African American workers which is important given Clemens and Hunt (2018) finding that there seems to be a change in the composition in the CPS sample around 1985). All the estimates suggest that wages were

[6] In the regressions where I use all the metropolitan areas I can also control for serial correlation by clustering standard errors at the metropolitan area level. When I do so, standard errors are in general smaller. I obtain similar estimates of the standard errors when I compute bootstrapped standard errors. Stata does not allow to use statistical weights when computing bootstrapped standard errors, which is why I prefer to report robust standard errors.

lower in Miami in the aftermath of the labor supply shock, i.e. between 1981 and 1985, than in the control group. In Panel B, I report the exact same regressions as in Panel A but using CPS ORG data. The results are similar, although smaller, as has been pointed already in the literature. Point estimates vary somewhat across columns, something that I take into account when discussing the meaning of these results in Section 4.

In panel C, I report the estimates using the intensity of treatment as explained in Section 3.2, where the difference in wages is taken between the pre-years 1977-1979, and the post years 1981-1984. The first two columns report the first stage regression. In column 1 without controls, while in column 2 I control for the change of native population which controls for short-run internal migration, see footnote 4 above. It is clear from these columns that the inflow of Cuban migrants was most important in metropolitan-skill cells where Cubans were already a large share. Controlling for native internal migration does not change this result, since, as I document more precisely below, the internal migration response does not start until later in the period. Columns 3, 4, and 5 report the OLS estimates. Column 3 using variation across metropolitan areas for high-school drop-out workers, columns 4 and 5 using variation also across education groups. The point estimate is around -1. This is a direct estimate of the inverse of the local labor demand elasticity which I defined in the model (ε^L). The IV estimates are very similar to the OLS estimates. This is so, because both the initial share of Cuban immigrants and the new inflows concentrate among high-school drop-outs in Miami. This estimate implies that an increase in a metropolitan area-skill cell equivalent to 10 percent of the native workforce in that cell reduces wages by around 10 percent on impact.

Table 1 goes around here

The recovery of wages that starts in Miami around 1985 or 1986 coincides in time with the decrease in the share of low-skilled workers living in Miami relative to the control groups. To investigate this I use the following regression framework:

$$\text{In Miami}_{i,t} = \delta_c + \beta_1 \text{Years 1981 - 1984}_t + \beta_2 \text{Years 1985 - 1990}_t + \varepsilon_{i,t} \tag{10}$$

where In Miami$_{i,t}$ is a variable that takes value one if individual i is in Miami at time t, Years 1981 - 1984$_t$ is a dummy variable that takes value one for the years 1981 - 1984, and Years 1985 - 1990$_t$ is a dummy variable that takes value one for the years 1985 - 1990. I run this regression using all high-school drop-out workers in Miami and in the control group over the period 1977 to 1990. Hence, β_i captures the share of low-skilled workers in Miami relative to the omitted time period (1977-1980), relative to the control group. I can estimate β_i using various types of estimators. I can for example run simple OLS, which would give linear probability model estimates, or I can estimate probit models. The results do not change. I use in what follows probit models. Finally, note that, as before, I can in fact plot an estimate for each of the years in the regression.

To gain intuition on the estimates I first plot the estimate for each of the years in the sample. In Panel B of Figure 2 we see that the share of low-skilled workers living in Miami increases in 1980 coinciding exactly with the arrival of the Mariel Boatlift Cuban immigrants. This is so, both when we compare

Miami to rest of the US, to Card and Borjas' placebos, or when we compare it to all the metropolitan areas in the South Atlantic region.[7]

A second remarkable aspect shown in panel B of Figure 2 is that the relative concentration of low-skilled workers in Miami only seems to last until 1984 or 1985. After that, it seems to decline. Depending on the control group, the decline seems to be complete or it seems that there is a small decline and by the end of the decade there are still more low-skilled workers in Miami than in the control cities.

Table 2 quantifies what we see in Panel B of Figure 2. Panel A of Table 2 shows that there is a sharp increase in the share of low-skilled workers in Miami, which somewhat disappears by the end of the decade. In this table, unlike in the figure, I control for observable characteristics. When comparing Miami to the rest of the US, we see that Miami gained low-skilled workers in the period 1981 to 1984 and then lost some of these workers. In the period 1985 to 1990, however, Miami retained roughly two thirds of the low-skilled workers gained in the early 80s when compared to all of the US. Panel B of Table 2 repeats the exercise but only for high-skilled workers. It is quite clear from this panel that the increased concentration in Miami only affected low-skilled workers.

Table 2 goes around here

3.4 Long-run estimates

To check that indeed wages of low-skilled workers are back to "normal" by 1990 as appreciated in Figure 2, I use the following regression:

$$\Delta \ln w_{ce} = \alpha + \beta \frac{\text{Cub}_{ce}}{\text{Nat}_{ce}} + \delta_c + \delta_e + \varepsilon_{ce} \tag{11}$$

where $\Delta \ln w_{ce}$ is the change in wages of workers of education e between 1980 and 1990 in metropolitan area c, and where $\frac{\text{Cub}_{ce}}{\text{Nat}_{ce}}$ is the Mariel Boatlift induced shock to labor supply in each city and education group, which is measured as the number of Cubans that in the 1990 reported to be living in each city in 1985 and who claim to have arrived in the US in 1980-1981 with education e, divided by the number of non-Cuban workers in each city and education group in 1985. δ_c and δ_e are city and education fixed effects. These allow for city specific and (national) education specific time trends. In some specifications I restrict the regression to low-skilled workers. In this case I cannot include city and education fixed effects.

To control for the possible endogenous location choice of immigrants I instrument $\frac{\text{Cub}_{ce}}{\text{Nat}_{ce}}$ by the share of Cubans in each city before the Mariel Boatlift shock as explained in Section 3.2.

It is worth noting that running this regression between Censuses means that β can be interpreted as the inverse local labor demand elasticity once adjustments have taken place. This is, in the short run, before any adjustment takes place, β is the (inverse) local labor demand elasticity. If there are adjustments, then β contains also these adjustments.

Table 3 goes around here

[7] The same pattern emerges when comparing it to Peri-Yasenov's control group.

Table 3 reports these results. The first column of Panel A shows that the initial share of Cubans (among high-school drop-outs) is a good predictor of the inflow of Cubans during the Mariel Boatlift episode across metropolitan areas. The same is true if I expand the regression to include the four education groups and I include metropolitan area and education fixed effects. In Panel B, columns (1) and (2), I estimate the wage effects over the entire decade using IV regressions. It is clear from these two columns that wages of low skilled workers in high-Cuban locations do not seem to be lower than in lower Cuban migration locations. Similarly, rentals do not seem to have been affected differentially over the decade as a function of the Mariel Boatlift-induced labor supply shock. Point estimates in columns (3) and (4) of Panel B are small and statistically indistinguishable from 0.

To investigate how much internal migration there was during the decade I use the following specification:

$$\Delta \text{Share of low-skilled}_c = \alpha + (1-\lambda)\frac{\text{Cub}_c}{\text{Nat}_c} + \varepsilon_c \tag{12}$$

where Share of low-skilled$_c$ is the number of low-skilled workers as a fraction of the total population, and where the change is taken between 1980 and 1990. In this case, an estimate of $\lambda = 0$ indicates that there is no internal migration. This is, for each Cuban low-skilled immigrant we have that the share of low-skilled workers increases by exactly 1. Instead if $\lambda = 1$ then it means that internal migration completely dissipates the local shock, so that Miami, by 1990 does not have more low-skilled workers despite the sizable unexpected inflow of Cuban low-skilled workers.

Results of regression 12 are shown in columns 3 and 4 of Panel A of Table 3. Both with the OLS and with the IV, I obtain estimates of around .6, i.e $\hat{\lambda} = 0.4$. This means that there was some internal migration but that Miami gained low-skilled workers relative to the other cities in the US.

4 Decomposition

With the estimates provided in Section 3, we can use the model to quantify the relative importance of internal migration and other factors in the absorption of immigration. For this we only need to realize that:

$$\widehat{\nu} = (1 - \frac{.25}{1+0.6})\widehat{\varepsilon^L} + \frac{0.25}{1+0.6}\widehat{\lambda}$$

where again, $\widehat{\varepsilon^L}$ is an estimate of the (inverse) local labor demand elasticity, which under the assumption made in Section 2 can be estimated from the short-run wage response. In Panel C of Table 1 I estimate this parameter to be around -1. This estimate is in line with Table 1. If the labor supply shock was equivalent to .25 percent of the low-skilled labor force and wages are estimated to have declined by between 10 and 30 percent, then it means that the inverse local labor demand elasticity is between .4 and 1.2. $\hat{\lambda}$ is the long-run internal migration response, which we have estimated in Table 3 to be around .4. Finally, as a reminder for the computation of $\widehat{\nu}$ we need estimates of α which I set equal to .25 and $\epsilon = .6$ – which are the estimates available in the literature (Davis and Ortalo-Magne, 2011; Saiz, 2010).

With all these estimates I can decompose the recovery into internal migration and other factors, as explained in Section 2.3. I show this exercise in Table 4, both using the baseline estimates and providing a number of alternative decompositions assuming alternative wage, internal migration, consumption of housing, and housing supply elasticity estimates.

Note that with these estimates I can also report an estimate of the internal migration elasticity – which measures how many low-skilled workers left Miami between 1985 to 1990 given the change in low-skilled wages until 1985. To obtain the change in low-skilled wages I multiply the inverse local labor demand elasticity by the size of the local shock in Miami, which was around 25 to 30 percent. To be conservative I assume that the shock was equivalent to 25 percent of the low-skilled labor market. Having this estimate is useful since it can be compared to the literature, which has estimated this number to be between 1.5 and 3 (Caliendo et al., 2019; Diamond, 2015; Monras, 2020).

The first row shows the baseline estimates. The baseline estimates suggest that around 40 percent of the indirect utility recovery is explained by internal migration. The baseline estimates suggest that the wage and internal migration responses are consistent with an internal migration elasticity of around 1.6, i.e. inside the range of estimates in other literature.

Given the controversy surrounding the wage estimates obtained from the Mariel Boatlift episode, I investigate thoroughly the sensitivity of the decomposition of the recovery between internal migration and other factors. I organize this exercise by showing how the results change if instead of using the baseline estimate of the (inverse) local labor demand elasticity I use an estimate of -0.4, an estimate of -0.7, and an estimate of -1.4. This covers the range of estimates in the literature. I do so for the baseline estimates of the share of income devoted to housing, local housing supply elasticity, and long-run internal migration response. This is shown in panel A of Table 4. Panel A shows that internal migration accounts from 20 to 70 percent of the recovery.

In Panel B of Table 4 I show the same results but assuming that α is 0.3 instead of 0.25. This exercise is justified by the fact that it is sometimes argued that in larger cities the share of income devoted to housing is higher, or by the fact that we can interpret housing as a broader non-tradable sector. The decomposition of the recovery is again similar, with estimates of the importance of internal migration that fluctuate around 50%.

Panel C shows sensitivity of the results to alternative housing supply elasticities. Miami is somewhat special relative to other US cities in that the expansion of its housing stock is relatively constrained. Hence, perhaps this feature of Miami is driving the results, rather than the wage and internal migration estimates. As can be seen in Panel C, assuming a much higher housing supply elasticity of 1.5 does not change significantly the results.

Finally, in Panel D, I return to the baseline estimates of the share of income devoted to housing and housing supply elasticity and I assume instead a larger ten-year horizon internal migration response. Not surprisingly, this exercise increases the relative importance of internal migration, although numbers are still similar to the baseline estimates.

Table 4 goes around here

Taken altogether, Table 4 suggests that (through the lenses of the model introduced in Section 2) internal migration accounts for roughly 50 percent of the recovery of indirect utility. This exercise highlights how the model can be used to understand the full path of adjustment of local economies to immigrant-driven labor supply shocks.

5 Conclusion

In this paper I use the Mariel Boatlift to estimate, through the lenses of a small open economy model of a metropolitan area, the relative importance of internal migration in dissipating wage effects resulting from immigrant-driven labor supply shocks. To do so, I document short-run declines in wages of low-skilled workers in Miami relative to workers in other cities and relative to higher skilled workers. This variation allows to estimate the short-run local labor demand elasticity. Next, I show that internal migration seems to respond to this local shock during the second part of the decade of the 1980s. This seems to coincide with a recovery of wages in Miami.

To interpret this evidence, I develop a model that helps to analyze the relative contribution of internal migration and other factors to the recovery of indirect utility in Miami. I document that in the longer-run, i.e. between 1980 and 1990, wages and rental prices in Miami were not lower than in other local labor markets despite the large inflow of workers following the Mariel Boatlift, suggesting that Miami was back in equilibrium by 1990. At the same time, we see an increase in the supply of low skilled workers in Miami that is smaller than what would have been predicted by the size of the Mariel Boatlift shock. This suggests that internal migration might have helped in the wage recovery, but that other factors were also potentially important. The model helps me quantify the relative importance of these two mechanisms.

Through the lenses of the model and given the estimates that I report in this paper, the evidence suggests that around 50 percent of the recovery is explained by internal migration as emphasized in classical spatial equilibrium models. This result is robust to a number of sensitivity checks.

6 Figures

Figure 1: Graphical representation of the model

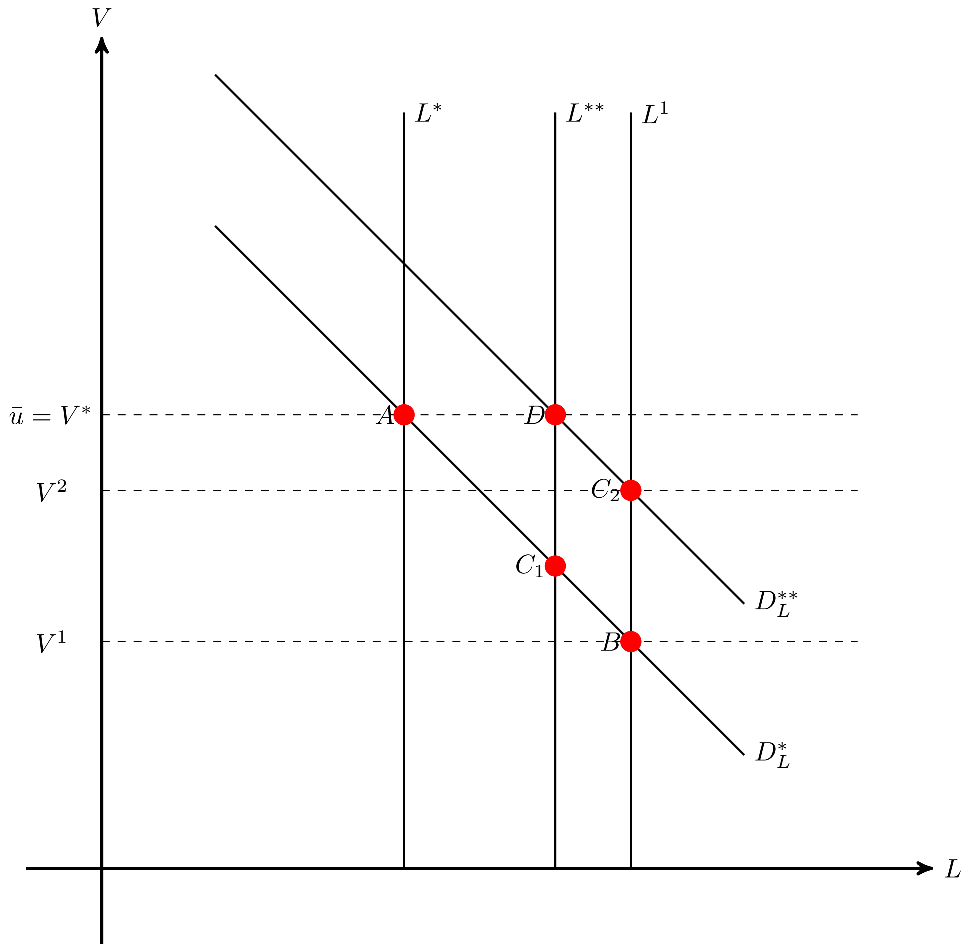

Figure 2: Wage dynamics and internal migration

Panel A: Wages

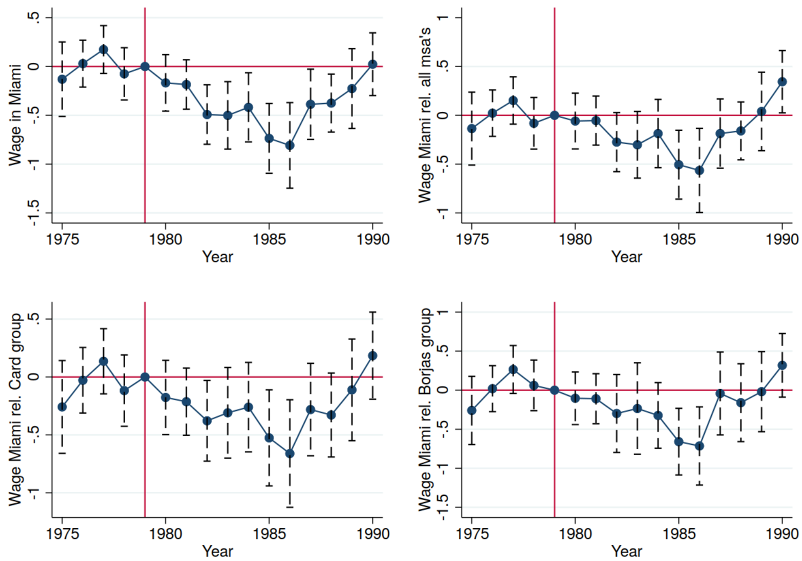

Panel B: Internal migration

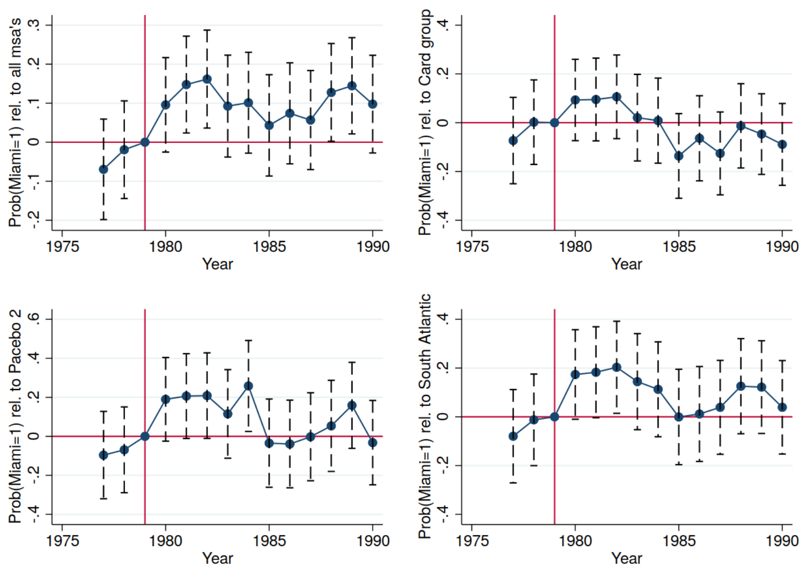

Notes: The graphs in Panel A of this figure show the wage dynamics of low-skilled workers in Miami relative to 1980 (top-left graph), relative to the Rest of the US (RoUS, top-right graph), relative to the Card control group (bottom-left graph) and relative to the Borjas control group (bottom-right graph), see Borjas (2017) for more details on the definitions of the different comparison groups. The graphs in Panel B of this figure show the relative share of low-skilled workers in Miami relative to the rest of the US (top-left graph), relative to the Card control (top-right graph), relative to the Borjas control (bottom-left graph) and relative to the rest of cities in the South Atlantic region (bottom-right graph). Vertical lines display 95 per cent confident intervals.

7 Tables

Table 1: Estimation of the causal effect of Cuban immigration on wages

Panel A: **Wages of Low-Skilled Workers, March supplement**

VARIABLES	(1) (ln) wage OLS	(2) (ln) wage OLS	(3) (ln) wage OLS	(4) (ln) wage OLS	(5) (ln) wage OLS	(6) (ln) wage OLS	(7) (ln) wage OLS	(8) (ln) wage OLS
Post x Miami	-0.239	-0.273	-0.330	-0.222	-0.0992	-0.119	-0.197	-0.140
	(0.0828)	(0.0891)	(0.110)	(0.0893)	(0.0805)	(0.0902)	(0.109)	(0.0951)
Observations	14,105	1,755	855	2,330	14,105	1,755	855	2,330
Year FE	yes	yes	yes	yes	yes	yes	yes	yes
Metarea FE	yes	yes	yes	yes	yes	yes	yes	yes
Controls	no	no	no	no	yes	yes	yes	yes
Comparison to	all metropolitan areas	Card's control group	Borjas' control group	Peri-Yasenov's control group	all metropolitan areas	Card's control group	Borjas' control group	Peri-Yasenov's control group

Panel B: **Wages of Low-Skilled Workers, ORG files**

VARIABLES	(1) (ln) wage OLS	(2) (ln) wage OLS	(3) (ln) wage OLS	(4) (ln) wage OLS	(5) (ln) wage OLS	(6) (ln) wage OLS	(7) (ln) wage OLS	(8) (ln) wage OLS
Post x Miami	-0.0915	-0.0724	-0.145	-0.0991	-0.0670	-0.0271	-0.0969	-0.0646
	(0.0444)	(0.0484)	(0.0510)	(0.0468)	(0.0422)	(0.0468)	(0.0491)	(0.0446)
Observations	19,240	2,388	1,213	3,232	19,240	2,388	1,213	3,232
Year FE	yes	yes	yes	yes	yes	yes	yes	yes
Metarea FE	yes	yes	yes	yes	yes	yes	yes	yes
Controls	no	no	no	no	yes	yes	yes	yes
Comparison to	all metropolitan areas	Card's control group	Borjas' control group	Peri-Yasenov's control group	all metropolitan areas	Card's control group	Borjas' control group	Peri-Yasenov's control group

Panel C: **Short-run inverse local labor demand elasticity**

VARIABLES	(1) Inflows of Cubans First-stage	(2) Inflows of Cubans First-stage	(3) Δ (ln) wage OLS	(4) Δ (ln) wage OLS	(5) Δ (ln) wage OLS	(6) Δ (ln) wage IV	(7) Δ (ln) wage IV	(8) Δ (ln) wage IV
Share of Cubans in 1980	1.260	1.262						
	(0.0529)	(0.0532)						
Inflows of Cubans			-0.854	-1.313	-1.350	-0.857	-1.264	-1.310
			(0.381)	(0.338)	(0.346)	(0.383)	(0.320)	(0.322)
Change in native population		-0.00163			0.0388			0.0385
		(0.00117)			(0.0450)			(0.0382)
Observations	152	152	44	152	152	44	152	152
Education FE	yes	yes	no	yes	yes	no	yes	yes
Metropolitan area FE	yes	yes	no	yes	yes	no	yes	yes
Metropolitan areas	all	all	all	all	all	all	all	all
Sample	all	all	HSDO	all	all	HSDO	all	all

Notes: Panel A and B of this table shows the estimates of the relative wages in Miami relative to various control groups of cities in 1981 to 1985 relative to before 1981. Panel A uses March CPS data, Panel B uses ORG CPS data. All the metropolitan areas refers to the 44 or 45 cities covered by the March CPS and CPS ORG throughout the period. Card's control group includes Atlanta, Houston, Los Angeles, and Tampa and Borjas'control group includes Anaheim, Rochester, Nassau-Suffolk, and San Jose. See details in the text and in Borjas (2017). Panel C replicates and expands the results reported in Borjas and Monras (2017). Controls include age, race and occupation (only in Panel A) dummies. Robust standard errors are reported in parenthesis.

Table 2: Estimation of the causal effect of Cuban immigration on internal migration

Panel A: **Internal Migration of Low-Skilled Workers**

	(1)	(2)	(3)	(4)	(5)
	Prob(Miami=1)	Prob(Miami=1)	Prob(Miami=1)	Prob(Miami=1)	Prob(Miami=1)
VARIABLES	probit	probit	probit	probit	probit
years 1981-1984	0.124	0.0675	0.203	0.136	0.139
	(0.0321)	(0.0446)	(0.0582)	(0.0451)	(0.0494)
years 1985-1990	0.0945	-0.0341	0.0563	0.156	0.0370
	(0.0292)	(0.0403)	(0.0541)	(0.0417)	(0.0453)
Observations	44,845	10,668	3,971	8,158	6,643
Controls	yes	yes	yes	yes	yes
Comparison to	all metropolitan areas	Card's control group	Borjas' control group	Peri-Yasenov's control control group	South Atlantic region

Panel B: **Internal Migration of High-Skilled Workers**

	(1)	(2)	(3)	(4)	(5)
	Prob(Miami=1)	Prob(Miami=1)	Prob(Miami=1)	Prob(Miami=1)	Prob(Miami=1)
VARIABLES	probit	probit	probit	probit	probit
years 1981-1984	0.0249	0.00964	0.0640	0.0474	0.0392
	(0.0205)	(0.0281)	(0.0311)	(0.0297)	(0.0296)
years 1985-1990	0.0485	-0.0128	0.0848	0.0634	-0.0475
	(0.0180)	(0.0247)	(0.0275)	(0.0264)	(0.0260)
Observations	181,054	29,357	17,345	22,587	25,783
Controls	yes	yes	yes	yes	yes
Comparison to	all metropolitan areas	Card's control group	Borjas' control group	Peri-Yasenov's control control group	South Atlantic region

Notes: Panel A and B of this table estimate the probability of being in Miami for low- (Panel A) and high-skilled workers (Panel B) in different periods of time over the 1980s relative to the years before the Mariel Boatlift shock using a probit model. Controls include age and race dummies. All the metropolitan areas refers to the 44 cities covered by the March CPS throughout the period. Card's control group includes Atlanta, Houston, Los Angeles, and Tampa Borjas'control group includes Anaheim, Rochester, Nassau-Suffolk, and San Jose, and Peri-Yasenov control group includes New Orleans, New York City, and Baltimore. See details in the text and in Borjas (2017). Robust standard errors are reported in parenthesis.

Table 3: Estimation of the causal effect of Cuban immigration on long-run wages, rents, and internal migration

Panel A: **First-stage and internal migration**

VARIABLES	(1) Inflow of Cubans First-stage	(2) Inflow of Cubans First-stage	(3) Δ share low-skilled OLS	(4) Δ share low-skilled IV
L.Share of Cubans	0.716	1.231		
	(0.0169)	(0.0845)		
Inflow of Cubans			0.604	0.641
			(0.0903)	(0.113)
Observations	38	152	38	38
Sample	HSDO	all	all	all
Education FE	no	yes	no	no
Metropolitan area FE	no	yes	no	no
widstat				1797

Panel B: **Wages and rents**

VARIABLES	(1) Δ (ln) wage IV	(2) Δ (ln) wage IV	(3) Δ (ln) rent IV	(4) Δ (ln) rent IV
Inflow of Cubans	0.113	-0.0858	0.180	0.0779
	(0.517)	(0.198)	(0.361)	(0.106)
Observations	38	152	38	152
Sample	HSDO	all	HSDO	all
Education FE	no	yes	no	yes
Metropolitan area FE	no	yes	no	yes
widstat	1797	212.3	1797	212.3

Notes: This table estimates the effect of the inflow of Cubans in 1980 (as a fraction of the low-skilled labor force) on the low-skilled wage change, the low-skilled change in rents, and the change in the share of low-skilled workers between 1980 and 1990, using the 1980 importance of Cubans across local labor markets as instrument. HSDO indicates that the regression is restricted to High-school drop-outs. This table uses variation from the 38 metropolitan areas available in the Census and CPS data throughout this period. 'widstat' indicates the F-stat of the excluded instrument in the first stage regression.

Table 4: Contribution of internal migration and local technology adoption to wage recovery

Parameter:	Inv. Local labor demand elasticity $\widehat{\varepsilon^L}$	Internal migration response $\widehat{\lambda}$	Share of income to housing $\widehat{\alpha}$	Housing supply elasticity $\widehat{\epsilon}$	Indirect utility elasticity $\widehat{\tilde{\varepsilon}^L}$	Other factors $\widehat{\nu}$	Contribution to recovery: Other factors	Contribution to recovery: Internal migration	Internal migration elasticity
Panel A: **Baseline**									
Baseline	1.0	0.4	0.25	0.6	1.0	0.6	60%	40%	1.6
Very Elastic LD	0.4	0.4	0.25	0.6	0.5	0.3	30%	70%	4.0
Elastic LD	0.7	0.4	0.25	0.6	0.7	0.4	45%	55%	2.3
Inelastic LD	1.4	0.4	0.25	0.6	1.3	0.8	80%	20%	1.1
Panel B: **Higher share of income devoted to housing**									
Baseline	1.0	0.4	0.30	0.6	1.0	0.6	60%	40%	1.6
Very Elastic LD	0.4	0.4	0.30	0.6	0.5	0.3	31%	69%	4.0
Elastic LD	0.7	0.4	0.30	0.6	0.8	0.5	45%	55%	2.3
Inelastic LD	1.4	0.4	0.30	0.6	1.3	0.8	80%	21%	1.1
Panel C: **Higher housing supply elasticity**									
Baseline	1.0	0.4	0.25	1.5	1.0	0.6	60%	40%	1.6
Very Elastic LD	0.4	0.4	0.25	1.5	0.5	0.3	28%	72%	4.0
Elastic LD	0.7	0.4	0.25	1.5	0.7	0.4	44%	56%	2.3
Inelastic LD	1.4	0.4	0.25	1.5	1.4	0.8	82%	18%	1.1
Panel D: **Higher internal migration response**									
Baseline	1.0	0.6	0.25	0.6	1.0	0.4	40%	60%	2.4
Very Elastic LD	0.4	0.6	0.25	0.6	0.5	0.2	20%	80%	6.0
Elastic LD	0.7	0.6	0.25	0.6	0.7	0.3	30%	70%	3.4
Inelastic LD	1.4	0.6	0.25	0.6	1.3	0.5	54%	47%	1.7

Notes: This table provides estimates on the relative contribution of internal migration and all other factors in dissipating the indirect utility effects of immigrant-driven labor supply shocks. The table provides both the baseline estimates as explained in the main text, and sensitivity analysis of these results to alternative estimates of the key parameters.

References

Abadie, A., "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects," *Journal of Economic Literature*, 2020.

_ **and J. Gardeazabal**, "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Review*, March 2003, *93* (1), 113–132.

Albert, C. and J. Monras, "The Regional Impact of Economic Shocks: Why Immigration is Different than Import Competition," *mimeo*, 2020.

Altonji, J. and D. Card, *The Effects of Immigration on the Labor Market Outcomes of Less-Skilled Natives*, in John Abowd and Richard Freeman (eds.), Immigration, Trade, and the Labor Market, University of Chicago Press, 1991.

Amior, M., "Immigration, Local Crowd-Out and Undercoverage Bias," *mimeo*, 2020.

Borjas, G., "The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market," *Quarterly Journal of Economics*, 2003, pp. 1335–1374.

_ , "Native Internal Migration and the Labor Market Impact of Immigration," *Journal of Human Resources*, 2006, *41(2)*, 221–258.

_ , "The Wage Impact of the Marielitos: A Reappraisal," *Industrial and Labor Relations Review*, 2017.

_ **and J. Monras**, "The Labor Market Consequences of Refugee Supply Shocks," *Economic Policy*, 2017, *32(91)*, 361–413.

_ **, R. Freeman, and L. Katz**, "Searching for the Effect of Immigration on the Labor Market," *American Economic Review Papers and Proceedings*, 1996.

_ **,** _ **, and** _ , "How Much Do Immigration and Trade Affect Labor Market Outcomes?," *Brookings Papers on Economic Activity*, 1997, pp. 1–67.

Caliendo, L., M. Dvorkin, and F. Parro, "Trade and Labor Market Dynamics," *Econometrica*, 2019.

Card, D., "The Impact of the Mariel Boatlift on the Miami Labor Market," *Industrial and Labor Relations Review*, 1990, pp. 245–257.

_ , "Immigration and Inequality," *American Economic Review Papers and Proceedings*, 2009, *99(2)*, 1–21.

_ **and J. DiNardo**, "Do immigrant Inflows lead to Native Outflows," *American Economic Review*, 2000.

Cascio, E. and E. Lewis, "How Much Does Amnesty Strengthen the Safety Net? Evidence from the Immigration Reform and Control Act of 1986," *Journal of Public Economics*, 2018.

Clemens, M. and J. Hunt, "The Labor Market Effects of Refugee Waves: Reconciling Conflicting Results," *Industrial and Labor Relations Review*, 2018.

— , **E. Lewis, and H. Postel**, "Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion," *American Economic Review*, 2018, *108(6)*, 1468–87.

Cortes, P., "The Effect of Low-skilled Immigration on U.S. Prices: Evidence from CPI Data," *Journal of Political Economy*, 2008, pp. 381–422.

Davis, M. and F. Ortalo-Magne, "Household Expenditures, Wages, Rents," *Review of Economic Dynamics*, 2011.

Diamond, R., "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 2015.

Dustmann, C. and A. Glitz, "How Do Industries and Firms Respond to Changes in Local Labour Supply?," *Journal of Labor Economics*, 2015.

— , **T. Frattini, and I. Preston**, "The Effect of Immigration along the Distribution of Wages," *Review of Economic Studies*, 2013.

— , **U. Schonberg, and J. Stuhler**, "The Impact of Immigration: Why Do Studies Reach Such Different Results?," *Journal of Economic Perspectives*, 2016.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift, "Bartik Instruments: What, When, Why, and How," *NBER WP*, 2018.

Gonzalez, L. and F. Ortega, "How Do Very Open Economies Absorb Large Immigration Flows? Evidence from Spanish Regions," *Labour Economics*, 2010.

Hanson, G. and M. Slaughter, "Labor-market adjustment in open economies: Evidence from US states," *Journal of International Economics*, 2002, pp. 3–29.

Lafortune, J., E. Lewis, and J. Tessada, "People and Machines: A Look at The Evolving Relationship Between Capital and Skill In Manufacturing 1860-1930 Using Immigration Shocks," *Review of Economics and Statistics*, Forthcoming.

— , **J. Tessada, and C. Gonzalez-Velosa**, "More hands, more power? Estimating the impact of immigration on output and technology choices using early 20th century US agriculture," *Journal of International Economics*, 2015.

Lewis, E., "Local, Open Economies Within the U.S.: How Do Industries Respond to Immigration?," *FRB of Philadelphia Working Paper No. 04-01*, 2003.

— , "Immigration, Skill Mix, and Capital-Skill Complementarity," *Quarterly Journal of Economics*, 2012, *126(1)*, 1029–1069.

— , "Immigration and production technology," *Annual Review of Economics*, 2013, pp. 165–191.

Monras, J., "Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis," *Journal of Political Economy*, 2020, *128(8)*.

Ottaviano, G. and G. Peri, "Rethinking the Effect of Immigration on Wages," *Journal of the European Economic Association*, 2012, *10* (1), 152–197.

Peri, G. and C. Sparber, "Task Specialization, Immigration and Wages," *American Economic Journal: Applied Economics*, 2009, *1(3)*, 135–169.

_ **and** _ , "Assessing Inherent Model Bias: An Application to Native Displacement in Response to Immigration," *Journal of Urban Economics*, 2011, *69(1)*, 82–91.

_ **and V. Yasenov**, "The Labor Market Effects of a Refugee Wave: Applying the Synthetic Control Method to the Mariel Boatlift," *Journal of Human Resources*, 2019, *54(2)*, 267 − 309.

Ruggles, S., M. Sobek, T. Alexander, C.A. Fitch, R. Goeken, PK Hall, M. King, and C. Ronnander, "Integrated Public Use Microdata Series: Version 4.0 [Machine-readable database].," *Minneapolis, MN: Minnesota Population Center [producer and distributor]*, 2016.

Saiz, A., "Room in the Kitchen for the Melting Pot: Immigration and Rental Prices," *Review of Economics and Statistics*, 2003, *85(3)*, 502–521.

_ , "The Geographic Determinants of Housing Supply," *Quartely Journal of Economics*, 2010, *125(3)*, 1253–1296.

For online publication only

A Production functions

In the main text I assumed a generic production function of the following type: $F(L, O)$ where L refers to labor, with price w, and where O refers to other factors, with price vector w^o. In the main text I assumed that $F_L(L, O) = \varepsilon - \varepsilon^L \ln L$. In tis appendix I show how standard production functions in the literature fall within what I assumed for the model.

Perhaps the most general way to introduce the production function most used in the literature is to assume:

$$Y = F(L, O) = AK^{\beta}((A^H H)^{\frac{\sigma-1}{\sigma}} + (A^L L + K_L)^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}})^{1-\beta}$$

Where A is a Hicks-neutral technology parameter, A^j are factor augmenting technologies, K is capital, H is high-skilled labor, L is low-skilled labor, and K_L is capital that substitutes low-skilled labor.

The marginal product of low-skilled labor in this case is:

$$F_L(L, O) = AK^{\beta}(1-\beta)((A^H H)^{\frac{\sigma-1}{\sigma}} + (A^L L + K_L)^{\frac{\sigma-1}{\sigma}})^{(1-\beta)\frac{\sigma}{\sigma-1}-1}(A^L L + K_L)^{-\frac{1}{\sigma}} A_L$$

So:

$$F_L(L, O) = AK^{\beta}(1-\beta)Y^{\frac{1-\beta\sigma}{\sigma(1-\beta)}}(A^L L + K_L)^{-\frac{1}{\sigma}} A_L$$

Hence

$$\ln w = \ln(1-\beta) + \ln A + \beta \ln K + \frac{1-\beta\sigma}{\sigma(1-\beta)} \ln Y - \frac{1}{\sigma} \ln(A^L L + K_L) + \ln A_L$$

Or:

$$\ln w = \ln(1-\beta) + \ln A + \beta \ln K + \frac{1-\beta\sigma}{\sigma(1-\beta)} \ln Y - \frac{1}{\sigma} \ln(1 + \frac{K_L}{A_L L}) + \frac{\sigma-1}{\sigma} \ln A_L - \frac{1}{\sigma} \ln L$$

Hence, with this production function we would have:

$$\varepsilon = \ln(1-\beta) + \ln A + \beta \ln K + \frac{1-\beta\sigma}{\sigma(1-\beta)} \ln Y - \frac{1}{\sigma} \ln(1 + \frac{K_L}{A_L L}) + \frac{\sigma-1}{\sigma} \ln A_L$$

And

$$\varepsilon^L = \frac{1}{\sigma}$$

where I have ignored the interaction of different factors inside Y.

This expression also allows us to see why factor augmenting technologies or capital that substitutes low-skilled labor are likely candidates to capture the factors other than internal migration that contribute to the recovery of indirect utilities.

To see this, we need to note that in the long-run wages of both high- and low-skilled labor need to be back to equilibrium. The relative factor prices are given by:

$$\ln \frac{w}{w^H} = -\frac{1}{\sigma}\ln(1+\frac{K_L}{A_L L}) + \frac{\sigma-1}{\sigma}\ln\frac{A_L}{A_H} - \frac{1}{\sigma}\ln(\frac{L}{H})$$

This expression shows that hicks-neutral technologies A and capital K cannot help the recovery of factor prices of both factor simultaneously recover their initial level.

B Imperfect substitutability

In a seminal paper, Ottaviano and Peri (2012) argue that the reason why immigration seems to have had a small impact on labor market outcomes of native workers is related to the fact that immigrants and natives sharing the exact same characteristics may be different factors of production. For instance, even if immigrants speak English, their handle of the language may be imperfect something that, as argued in Peri and Sparber (2009) may lead to task specialization.

Knowing to what extend observationally equivalent natives and immigrants are indeed imperfect substitutes is not an easy task. One strategy would be to compare the wage trends of native and immigrant workers following an immigrant labor supply shock. If these wage trends evolve in parallel after the shock, the test would suggest that we cannot rule out that natives and immigrants are perfect substitutes. Instead if there are systematic differences in the wage trends after the shock, then we would conclude that natives and immigrants are imperfect substitutes.

The problem with these type of tests is that most of the data available, particularly in the US, are repeated cross-sections (for example CPS or Census data). Hence, the average wage of immigrants may change either because the price of labor changes, because the new immigrant arrivals are somewhat different than previous ones – for example less productive–, or because there is selected return migration among immigrants already in the host economy upon the arrival of new immigrants. Hence, it is very hard to tell with current data sets whether immigrants and natives are indeed perfect or imperfect substitutes.

In the main text I made the simplifying assumption that natives and immigrants are imperfect substitutes. In this appendix I study how would the model change if instead they were imperfect substitutes (in a model where there are other factors with weight $1-\beta$). In this case, the labor demand equation would be given by:

$$\ln w \approx \ln\beta - (1-\beta)\ln N + \frac{1-\sigma(1-\beta)}{\sigma-1}(\frac{I}{N})^{\frac{\sigma-1}{\sigma}}$$

Note that in this case, an immigrant shock (again measured as immigrants I divided by natives N) has a negative effect on natives wages if and only if $1-\sigma(1-\beta)<0$, or if and only if $\sigma > \frac{1}{1-\beta}$, where σ is the elasticity of substitution between natives and immigrants within that type of labor. This condition means that the effect is negative as long as natives and immigrants are sufficiently good substitutes. Estimates of σ in the literature are typically above 10.

With this derivation it is easy to see how the model predictions change when natives and immigrants are imperfect substitutes (with σ sufficiently large). First, the effect of the immigrant shock on natives is qualitatively similar, but quantitatively less strong. As a result, a smaller than one for one native relocation is enough to return wages to pre-shock levels, even when holding all else fixed. This means that imperfect substitutability can be counted as one of these other factors that “help” to absorb the shock, although in this case it is best represented as a smaller shift in the labor supply curve in Figure 1 rather than an upward shift in the indirect utility curve.