Local Adjustment to Immigrant-Driven Labor Supply Shocks

Joan Monras
Universitat Pompeu Fabra, Centre de Recerca en Economia Internacional, Barcelona Graduate School of Economics, and Centre for Economic Policy Research

When comparing high- to low-immigrant locations, a large literature documents small effects of immigration on labor market outcomes over 10-year horizons. The literature also documents short-run negative effects of immigrant-driven labor supply shocks, at least for some groups of native workers. Taken together, these results suggest 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, with simple reduced-form estimates of the effects of immigrant shocks on the outcomes of interest, the local adjustment to be decomposed through various channels.

I. Introduction

A large part of the literature evaluates the labor market effects of immigration by comparing changes in local labor market outcomes between high- and low-immigrant locations over 10-year horizons using census data. To identify the causal effect of immigration, most of the 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 and when focusing on the "average" native worker (Altonji and Card 1991; Borjas, Freeman, and Katz 1996; Borjas 2003; Cortes 2008; Card 2009). The literature also documents short-run negative effects of immigrant-driven labor supply shocks, at least for some groups of native workers (Borjas 2017; Monras 2020). Taken together, those results suggest that there are mechanisms in place that help local economies recover from the short-run effects of immigrant shocks.

The existing literature has investigated some of the mechanisms that may help local economies absorb immigrant shocks. For example, Lewis

I would like to thank Jinyoung Kim, Isaac Ehrlich, and two anonymous referees for helping me improve this paper. All remaining errors are mine.

Electronically Published April 19, 2021.

[*Journal of Human Capital*, 2021, vol. 15, no. 1]

(2011) 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 share of high school dropouts hired by a plant leads to about a 6% decline in plant-level machinery adoption. Similarly, according to Clemens, Lewis, and Postel (2018), the lack of native employment gains when the Bracero Program was removed can be explained by the patterns of technology adoption in response to immigrant shocks. As explained by Lewis (2011, 2013), the adoption of forms of capital that substitute for particular labor types tends to attenuate the effect that immigrant-driven changes in the skill mix may have on the relative returns to skills.

Other papers have instead focused on how immigrant-induced changes in the 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. Such models require intense cross-sector relocations that are not typically found in the data, in both the United States and other countries. See, for example, the pioneering work by Hanson and Slaughter (2002) and Lewis (2003) and the papers by González and Ortega (2011) and Dustmann and Glitz (2015).

A few recent papers have emphasized that changes in the local supply of labor (mostly through 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 1990s Mexican peso crisis, while Amior (2020) provides systematic evidence that internal migration plays a crucial role. Amior's estimates suggest, in fact, that internal migration may account for the full adjustment. That 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 changes in the local supply of labor—such as internal migration or human capital acquisition—and changes in the local demand for labor—such as technology adoption—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 local labor demand versus labor supply 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 of production. The second sector produces a nontradable 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 the local labor and housing markets.

The model can be used to understand the adjustment dynamics that follow an immigrant-induced labor supply shock in a local labor market. For that, I make two assumptions that can be justified with empirical tests. First, I assume that, in the short run, local technologies and factors of production do not adjust and, hence, that adjustment comes through factor and rental price changes. Second, in the long run, indirect utility must recover to the preshock level for the local (small, open) economy to return to spatial equilibrium. Under these two assumptions, we can use various estimates typically reported in the literature to understand the relative importance of local labor supply and demand factors in dissipating immigrant-induced labor supply shocks.

The first assumption means that short-run regressions relating price changes in the different factors of production to the immigrant-induced shock allow us to recover key parameters of the local production function. The second assumption means that we can use longer-run data on price changes to estimate when the economy is likely to be back to spatial equilibrium. Once we know when prices have returned to equilibrium, we can use data on internal migration and human capital acquisition to estimate how much of the local price recovery can be attributed to local labor supply changes. We can finally use the model to decompose how much of the recovery is due to local labor demand versus supply adjustment.

This model can be applied in a number of settings where we have an exogenous immigrant-induced labor supply shock affecting a local labor market. To illustrate how the model can be used in such contexts, I reanalyze some of the evidence surrounding the well-studied Mariel Boatlift episode. I first document that the wages of a group of (low-educated) workers in Miami declined in the first few years after the shock, relative to those of similar workers in a number of control groups, as has been shown in previous studies. Second, I show that the change between 1980 and 1990 in wages of all factor types and rental prices in Miami was similar to that in the rest of the United States. That point is not always emphasized in previous studies. 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-educated immigrants at the beginning of the decade and the evidence pointing to short-run declines in low-skilled workers' wage and rental price increases, as documented in Saiz (2003). This evidence suggests that Miami was back into equilibrium by 1990 and had fully absorbed the immigrant-driven labor supply shock of the early 1980s.

Next, I document that, although 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 0.6 low-skilled workers for each low-skilled Cuban immigrant. More precisely, 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 had been 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 or endogenous human capital acquisition—the two main factors that can change the share of low-educated workers in a location—might have contributed to the dissipation of wage effects.

The share of low-skilled workers in a location can change for at least two motives. First, workers of particular education levels can move (on net) to other locations. Second, younger workers in the location can change their educational attainment as a response to local shocks. To assess the relative role of internal migration versus human capital acquisition, I compare cohorts of workers born in Florida who by 1980 were just under 18 years old (and, hence, could more easily adjust their educational attainment) to equivalent cohorts born in other states, relative to workers who were just above 18 years old and hence might have had more difficulties in adjusting their educational attainment. Using these comparisons, I show that educational attainment did not change across cohorts differently in Florida than in other locations, which provides suggestive evidence that internal migration is probably more important in explaining the local labor supply adjustment observed in Miami than human capital acquisition. This is the first paper to document systematically the local labor supply response that followed the Mariel Boatlift.

Given these empirical results, I estimate, through the lenses of the model, that around 50% of the indirect-utility recovery after the shock is explained by internal migration, while the rest was likely driven by local labor demand adjustments such as technology adoption. This result is robust to a number of alternative estimates of the model's key parameters, which include the local labor demand elasticity (Card 1990; Borjas 2017; Clemens and Hunt 2019; Peri and Yasenov 2019), the share of income devoted to housing, local housing supply elasticity, and long-run internal migration estimates.

Related literature.—This article is related to papers that investigate the link between technology adoption and immigrant shocks (Lewis 2011; Lafortune, Tessada, and González-Velosa 2015; Cascio and Lewis 2018; Clemens, Lewis, and Postel 2018; Lafortune, Lewis, and Tessada 2019), since, as I argue below, one of the main drivers of local labor demand adjustment is technology adoption. Relative to these papers, I offer a model-based measure of the role that technology adoption and other labor demand factors 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 effects 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 labor demand factors may mitigate the effects of immigration on the level of wages.

This paper is also related to work discussing the internal migration responses to immigrant shocks. Borjas, Freeman, and Katz (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 the Card and DiNardo (2000) empirical strategy, in contraposition to Borjas (2006). In a recent paper (Albert and Monras 2020), we argue that the reason for previous literature finding mixed 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 give weight to small metropolitan areas close to the Mexican border, thereby resulting in lower internal mobility estimates than when other identification strategies are used.

Finally, this paper is related to the large body of 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 enough 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 learned from the Mariel Boatlift episode. Two main points differentiate the Borjas analysis from the original Card (1990) study. 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 on the basis of employment trends that included some of the years following the Mariel shock. The conclusion in Borjas (2017) seems to be radically different from that 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%.

Since the Borjas reappraisal, several papers have investigated the episode in detail. The debate mainly revolves around two different issues. On the one hand, the micro-level number of observations of male high school dropouts used to calculate wage trends is small, often fewer than 30 individual observations. This means that average wages are not calculated with much precision and, hence, that small changes in the sample of workers used to calculate these average wages may have substantial effects on the point estimates. That, at least in part, is the critique emphasized in Peri and Yasenov (2019) and Clemens and Hunt (2019). On the

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

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, as in the early 1980s there were only 44 metropolitan areas covered by the March supplements of the Current Population Survey (CPS) data. Hence, small changes in the metropolitan areas 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.

II. Model

In this section, I introduce an "open-city" spatial equilibrium model of a local labor market, which is Miami in the application below. It is an "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 the number of 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.[2]

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

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

A. *General Setting*

1. Utility

The utility function of a representative worker is given by

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

[2] Ehrlich and Pei (2020, 2021) explore the effects of immigration in models where the unit of analysis is a country rather than a city or a local labor market. Among the differences between a local approach like the one presented in this paper and a more aggregate perspective is the margin of adjustment of the supply of labor in response to immigration, as I discuss in more detail below.

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 calculate the indirect-utility function. Assuming that the price of the tradable good is the numeraire, the 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 or move to Miami, $\bar{u}$ is unaffected. Note also that, since workers do not have disutility from working, they supply inelastically their labor endowment.

In a model with more than one type of labor, for example, with low-educated and highly educated workers, this indirect-utility function would represent the indirect utility of workers most closely resembling immigrants. In that context, leaving the location or acquiring education (so as to become a different factor type) would be (almost) equivalent. Given that in the empirical application that I present below I find no evidence for endogenous human capital acquisition responses, I do not include it in this model explicitly. With human capital acquisition, one would need to track how the returns to other factors of production, not directly affected by the immigrant shock, react.

2. Tradable Sector

The local labor market is defined by the local production function of a representative and 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-educated workers, whom I also refer to, following the literature, as 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. The function $F(L, O)$ is a neoclassical constant-returns-to-scale production function. To keep the notation simple, I omit specifying explicitly terms such as 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:

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

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 $\ln F_{L}(L, O) = \varepsilon - \varepsilon^{L} \ln L$, where ε is a "residual" that includes other labor types, 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

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

This equation relates wages, which, in the application below, means the wages of low-skilled workers, to the supply of that factor. It shows that the (inverse) local demand function can be decomposed into 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. That includes technological change, changes in industrial composition, and capital adjustment. The term ε captures the fact that changes in all those 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; that is, it measures by how much wages decline with labor.

3. Housing Sector

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

Construction uses as inputs the final tradable good and land in a Cobb-Douglass production function.[3] 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 $\eta/(1 - \eta) = \epsilon$, that is, $H^{S}(p) = Hp^{\eta/(1-\eta)} = Hp^{\epsilon}$, where H is a positive constant.

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

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 effectively captures 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 by increasing the supply of housing.

B. *Equilibrium*

The equilibrium in this model is defined by $\ln V = \bar{u}$. This relationship determines the number of workers in the local economy. To solve the model, we need to determine how $\ln V$ depends on $\ln L$. 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 $\tilde{\alpha} = \alpha/(1+\epsilon)$ and $\tilde{A} = \ln A - [\alpha/(1+\epsilon)](\ln \alpha - \ln H)$. This step shows how we can use the housing sector part of the model to get rid of housing prices in the indirect utility function.

Moreover, we can use equation (3) to obtain

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

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 into 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 technologies. The slope captures all the sources of congestion. More workers add pressure to labor and housing markets.

C. *Properties*

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

1. 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 $\pi = I/L$, where I is the number of immigrants who arrive in the local labor market and L is the number of existing workers in that market. Following from equation (5), we have that

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

where $\lambda = \partial L/\partial I$ measures how many workers stay in the location per immigrant arrival and $\nu = \partial\tilde{\varepsilon}/\partial\pi$ 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 either by technological change that increases labor productivity 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—such as 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's is not, changes in Hicks-neutral parameters alone will not be able to return the economy to the preshock 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 as 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, such as the increased demand for local goods that necessarily comes with immigration (after all, immigrants also consume), are already captured in the model through the housing sector. A broader nontradables sector could be modeled in exactly the same way as housing. Hence, ν does not capture this channel.

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

It is also important to stress that, while in the model adjustments to L are associated with internal migration, one can also think that low-skilled immigrant supply shocks may induce low-educated workers to acquire more education and, hence, become another factor of production. I leave this mechanism out of the presentation of the model because it is not empirically relevant in the application that I analyze below. Allowing for endogenous skill upgrading would require the interactions between different factors of production to be modeled in more detail, so I have also left that out of this paper.

2. Short Run

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

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

Hence, the parameter $\tilde{\varepsilon}^{\mathrm{L}}$ 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 $\tilde{\varepsilon}^{\mathrm{L}} = \{1 - [\alpha/(1 + \epsilon)]\}\varepsilon^{\mathrm{L}} + [\alpha/(1 + \epsilon)]$.

We have good estimates of α and ϵ in the literature. Davis and Ortalo-Magné (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 $\tilde{\varepsilon}^{\mathrm{L}}$. Hence, given an estimate of ε^{L}, we have that $\tilde{\varepsilon}^{\mathrm{L}} = \{1 - [0.25/(1 + 0.6)]\}\varepsilon^{\mathrm{L}} + 0.25/(1 + 0.6)$. If we were to consider all nontradables beyond housing, we may want to consider what happens to the results with higher levels of α. Note, however, that when $\varepsilon^{\mathrm{L}} = 1$, α and ϵ do not matter.

3. Long Run

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

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

This equation says that in the long run, labor supply (which in the empirical application is exclusively driven by internal migration) and the adjustment of other (labor demand) factors respond sufficiently that indirect utility recovers its preshock level. Moreover, if amenities are fixed, a sufficient condition for indirect utility to return to preshock levels is that wages and housing prices recover from the shock.

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

4. Labor Supply versus Labor Demand Adjustment

We can use the model to think about the role that labor supply factors, such as internal migration, and other factors play in dissipating the effects of immigration on indirect utility. To do so, it is perhaps useful to illustrate the model with a graph. The figure 1 y-axis displays indirect-utility levels, and the x-axis shows employment. Initially, the market equilibrium is given by the intersection of the initial indirect-utility curve (D_{L}^{*})

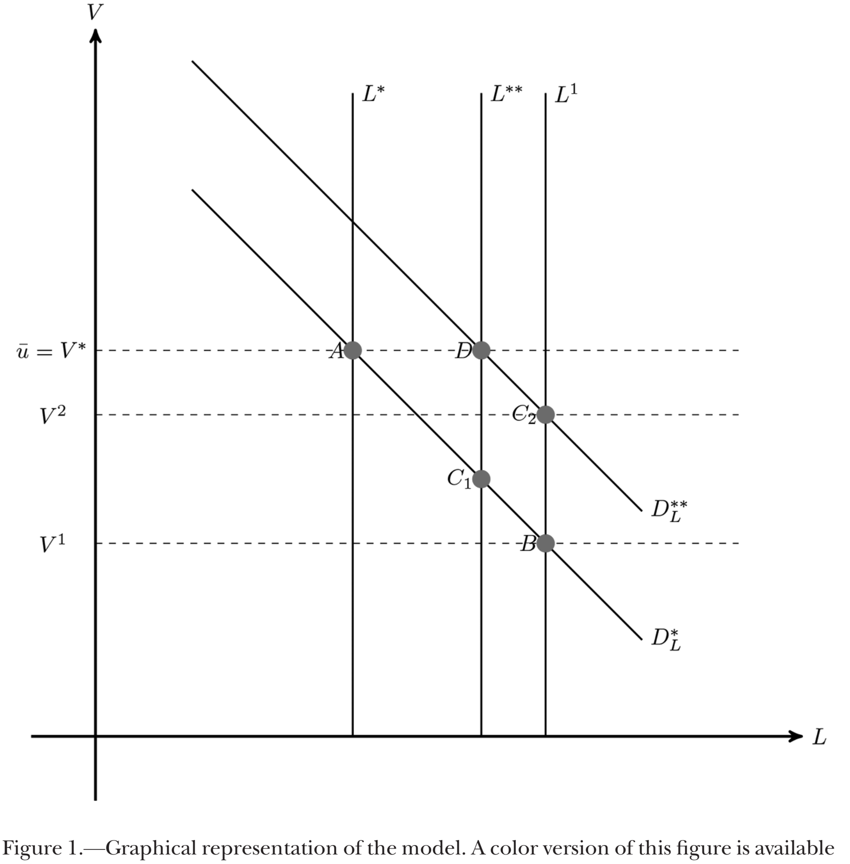

Figure 1.—Graphical representation of the model. A color version of this figure is available online.

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 calculate the local labor demand elasticity ε^{L}. Given the assumptions on the relationship between ε^{L} and $\tilde{\varepsilon}^{\mathrm{L}}$, the wage change allows me to recover the slope of the function D_{L}^* that moves indirect utility from V^* to V^1.

After this initial shock, both internal migration and other absorption mechanisms react to bring the economy back to the initial level of indirect utility at, in the figure, point D. In the data, we can see how much internal migration responds (and potentially how human capital acquisition changes the supply of labor of a particular type). That is, we can estimate the difference between L^1 and L^{**}. If only labor supply factors such as internal migration were contributing to dissipating indirect-utility effects, the equilibrium would be at point C_1 and, hence, at a level of indirect utility below the initial one. Hence, it must be that other factors change so that the indirect-utility curve moves from D_{L}^* to D_{L}^{**}. This is my proposed estimate of ν. We can see in the figure 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 (and other labor supply factors) is shut down.

The graph helps to show that we can decompose the indirect-utility recovery between the contribution of observable factors such as internal migration and all other (potentially unobserved) factors. That 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 calculate the level of indirect utility that would prevail if there was no internal migration. This is given by the level V^2 in the figure. Then, we can calculate 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.

III. Empirical Application

To further illustrate how this model can be used empirically, I reanalyze the evidence around the Mariel Boatlift episode through the lenses of the model in this section. More explicitly, I document how the large inflow of Cubans who arrived in Miami in 1980 with the Mariel Boatlift likely resulted in a decline in real wages, which fully recovered by 1990. I trace internal migration during the 1980s in response to these local changes. That allows me to estimate the key parameters of the model, which helps us understand both how an unexpected immigrant shock moved the initial (spatial) equilibrium and the forces that brought the economy back into that equilibrium.

A. 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 (ORG) files of the CPS. The CPS March supplements have complete information on wage income during the year before the interview and the number of weeks worked, so the weekly wages can be calculated. It also contains information on the education level of the individuals in the sample. In particular, I can construct four education codes: high school dropouts, high school graduates, some college, and college graduates or more. These four groups split the 1980 labor market into roughly equal-sized groups.

To calculate wages, I use the exact same sample as Borjas (2017). In particular, I restrict the sample to non-Hispanic, prime-age—that is, 25–59-year-old—working males. During the 1980s, women were entering the labor market at a much faster rate than previously. Hence, when using women to calculate 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 they are substantially more noisy. Arguably, we would like to exclude foreign-born individuals if the object of interest is native wages. Birthplace is not recorded in the CPS data until 1994, so the best approximation is the Hispanic variable, which allows us to identify Hispanics of Cuban and of Mexican origin.

An alternative data set for calculating wages during this period is provided by the ORG files of the CPS. I apply the exact same sample selection when using these data. The preshock years available in the CPS ORG files cover only 1979 and 1980 (which is driven by the coverage of metropolitan areas), whereas the preshock years when using the March CPS data include 1975–80.

To study internal migration, I trace the share of workers of a certain characteristic who live in Miami, using March CPS data. This share could change for reasons other than internal migration. For instance, it could be that mortality rates for, say, high school dropouts 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. Similarly, it could be that workers in Miami acquire more or less education as a function of immigrant shocks. From the perspective of the model, it does not matter much what is driving the change in the workforce composition 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. To justify further the labeling of local labor supply adjustments as internal migration, in table 4 I use data from the 2000 to show that endogenous educational acquisition does not seem to be the main driver of changes in the local labor composition in Florida over the 1980s.

To estimate longer-run effects on wages and internal migration, I use the 1980 and 1990 censuses, provided by Ruggles et al. (2016). From those, 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 size of the Mariel shock. To do 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) who were residing in Miami in 1985 to estimate the number of Cuban migrants who 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, that they represent a good proxy for the size of the shock. If anything, we can imagine that the shock was larger than can be estimated with the 1990 census data. The census data allow me to calculate the relative size of the shock for each education group, since the census of 1990 records the educational attainment of the Cuban immigrants. Summary statistics tables for those data are provided in Borjas (2017) and Borjas and Monras (2017).

B. *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-differences setting. The key identification assumption in this case is that Miami would have followed a trend similar to that of the control group. Difference-in-differences specifications are quite standard. I use graphical representations of the treatment dummy in each year to analyze the trends in Miami and 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 pretrends 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 because of the difficulties in using this approach in this context.[4] I also report results comparing Miami to all the other identifiable metropolitan areas (43 in total).

On the other hand, I use specifications where I leverage the intensity of the treatment, that is, where I focus on Cuban-induced increases in the workforce 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):[5]

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

where c indexes metropolitan areas and e indexes the four education groups: high school dropouts, high school graduates, some college, and college graduates or more; δ_c and δ_e are metropolitan area and education fixed effects, respectively.

As 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 such as wages. Hence, the need for an instrument.

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

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

[4] As argued in Abadie (2020), synthetic control groups work best when the preshock period is long and the pool of donors is large. In this case, the options are a preperiod length that spans 1973–79, with a pool of donors of 33 metropolitan areas, and a preshock period of 1976–79, with a pool of donors of 43 metropolitan areas. Moreover, the number of observations in many of these metropolitan areas is small, and, hence, preshock variables are measured with error, which further complicates the use of synthetic control methods in this episode.

[5] Note that, in the postshock period, eq. (3) implies that $\ln w \approx \varepsilon - \varepsilon^{\mathrm{L}} \ln L = \varepsilon - \varepsilon^{\mathrm{L}} \ln(N + I) \approx \varepsilon - \varepsilon^{\mathrm{L}} \ln(N) - \varepsilon^{\mathrm{L}}(I/N)$. Hence, when comparing the pre- to the postshock periods we obtain $\Delta \ln w \approx \varepsilon - \varepsilon^{\mathrm{L}} \Delta \ln(N) - \varepsilon^{\mathrm{L}}(I/N)$. This specification may be problematic when there is substantial skill downgrading, as argued in Dustmann, Frattini, and Preston (2013) and Dustmann, Schönberg, and Stuhler (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 that given by eq. (7) directly identifies the parameter of interest (from the perspective of the model), while other specifications measuring the immigrant shock relative to the overall labor force do not.

where Cub_{ce} is the inflow of Cuban workers who arrived in each metropolitan area during the Mariel Boatlift episode with education level e and Nat_{ce} is the size of the local labor force excluding Cuban workers.[6]

The most standard way to use the immigrant-networks instrumental variable (IV) strategy 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, Sorkin, and Swift (2020), in this setting identification mostly comes from the "shares." A more direct way to use the identifying variation is to predict the inflow by the initial share: $\mathrm{Cub}_{ce,0}/\mathrm{Nat}_{ce,0}$. This variable is the size of the Cuban stock relative to the local population at the initial period, in this case 1980, that is, before the Mariel Boatlift. This variable captures an intensity of treatment; that is, it measures how important Cubans are (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 therefore are unlikely to be related to developments in the US economy.

C. *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 the opportunity and migrated toward the United States during the period April 23–October 1980. Nearly 70,000 immigrants likely settled in Miami, something that accounts for around 8% of the Miami workforce at the time. Cuban immigrants were very low educated. As many as 62% lacked a high school diploma, compared to around 23% among the natives. Hence, these low-skilled workers experienced a labor supply shock of around 32% of the workforce before 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-differences specification:

[6] I identify Mariel immigrants as those immigrants arriving in the United States in the 1981–83 census category, as reported in the 1990 census. I also identify in the 1990 census the location of each individual in 1985 and take that as a proxy of the location of arrival.

$$\ln w_{i,c,t} = \delta_c + \delta_t + \beta \text{Post-Mariel}_t \times \text{Miami}_c + \gamma X_{i,c,t} + \varepsilon_{i,c,t}, \quad (9)$$

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 δ_c and δ_t are city and time fixed effects, respectively. I run this regression using only high school dropouts; $X_{i,c,t}$ 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$_t$ × 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, so inference is complicated by the possibility of serial correlation in outcome variables and having only one treated location with, 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.[7]

The results are reported in panel A of figure 2. I report estimates for Miami, in an event-type setting and relative to four different control groups: the original Card control group—Atlanta, Houston, Los Angeles, and Tampa—the control group proposed in Borjas (2017)—Anaheim, Nassau, 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 United States for which we have data for the early 1980s.

Irrespective of the control group that I use, figure 2 shows that there are no systematic trends in the wage evolution in Miami before 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 occurs around 1985 or 1986. After this, wages recover to the extent 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, from that alone, it would be hard to explain why they declined later. It could also be that there is some wage stickiness, so that wage effects are observable only 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

[7] 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 calculate bootstrapped standard errors. Stata does not allow the use of statistical weights when calculating bootstrapped standard errors, so I prefer to report robust standard errors.

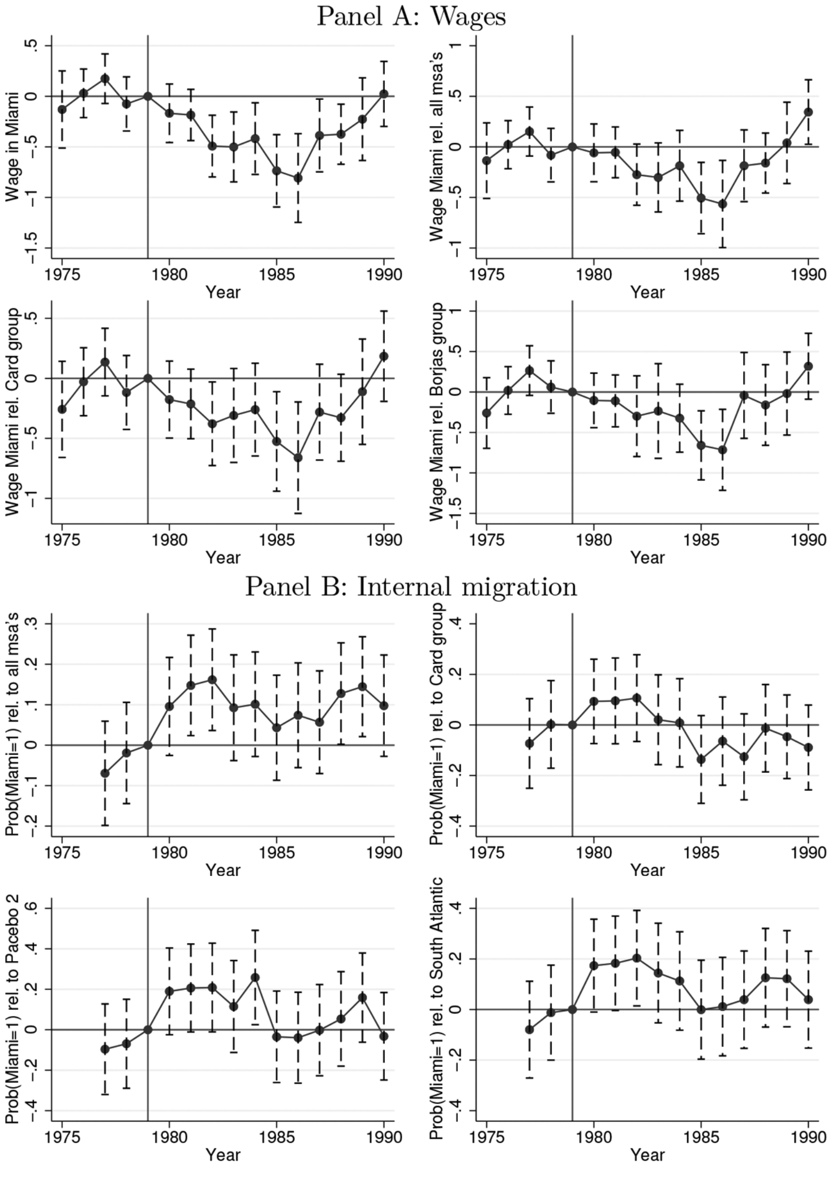

Figure 2.—Wage dynamics and internal migration. The graphs in panel A show the wage dynamics of low-skilled workers in Miami relative to 1980 (*top left*), relative to the rest of the United States (*top right*), relative to the Card control group (*bottom left*), and relative to the Borjas control group (*bottom right*). The graphs in panel B show the relative share of low-skilled workers in Miami relative to the rest of the United States (*top left*), relative to the Card control (*top right*), relative to the Borjas control ("Placebo 2"; *bottom left*), and relative to the rest of cities in the South Atlantic region (*bottom right*). Vertical lines display 95% confidence intervals. msa = metropolitan statistical area. A color version of this figure is available online.

other specific skills. Whatever the reasons, it seems that there is a decline in wages for the least educated workers in Miami that may be related to the unexpectedly large flow of immigrants during these years. As explained in Borjas and Monras (2017), the wage decline is observed only for the least skilled native workers. In fact, labor market outcomes of more skilled workers in Miami actually improved relative to those of 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 column 1 of panel A, I estimate the wage effects of the Mariel Boatlift, using all the other 43 metropolitan areas as a control group. Column 2 uses only the Card control, column 3 uses the Borjas control, and column 4 uses the Peri-Yasenov control group. I repeat the estimates in columns 5–8 but add individual-level controls (most importantly, a dummy for African American workers; see the Clemens and Hunt 2019 finding that there seems to be a change in the composition in the CPS sample around 1985). All the estimates suggest that wages were lower in Miami in the aftermath of the labor supply shock, that is, 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 already been pointed in the literature. Point estimates vary somewhat across columns, so I take that into account when discussing the meaning of these results in section IV.

In panel C, I report the estimates using the intensity of treatment as explained in section III.B, where the difference in wages is taken between the preshock years 1977–79 and the postshock years 1981–84. The first two columns report the first-stage regression. Column 1 show the results without controls, while in column 2 I control for the change of native population that 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 area–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 3–5 report the OLS (ordinary least squares) estimates. Column 3 uses variation across metropolitan areas for high school dropout workers, and columns 4 and 5 use 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 that 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 dropouts in Miami. This estimate implies that an increase in a metropolitan area–skill cell equivalent to 10% of the native workforce in that cell reduces wages by around 10% on impact.

The recovery of wages that starts in Miami around 1985 or 1986 coincides with the decrease in the share of low-skilled workers living in Miami

TABLE 1
Estimation of the Causal Effect of Cuban Immigration on Wages

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ln Wage	ln Wage	ln Wage	ln Wage	ln Wage	ln Wage	ln Wage	ln Wage
	A. Wages of Low-Skilled Workers, March Supplement (OLS)							
Post × Miami	−.239	−.273	−.330	−.222	−.0992	−.119	−.197	−.140
	(.0828)	(.0891)	(.110)	(.0893)	(.0805)	(.0902)	(.109)	(.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
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Comparison to	All MSAs	Card control	Borjas control	Peri-Yasenov control	All MSAs	Card control	Borjas control	Peri-Yasenov control
	B. Wages of Low-Skilled Workers, ORG files							
Post × Miami	−.0915	−.0724	−.145	−.0991	−.0670	−.0271	−.0969	−.0646
	(.0444)	(.0484)	(.0510)	(.0468)	(.0422)	(.0468)	(.0491)	(.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
MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Comparison to	All MSAs	Card control	Borjas control	Peri-Yasenov control	All MSAs	Card control	Borjas control	Peri-Yasenov control

	C. Short-Run Inverse Local Labor Demand Elasticity							
	First-stage Regressions		OLS Estimates			IV Estimates		
	Inflows of Cubans	Inflows of Cubans	Δ(ln Wage)	Δ(ln Wage)	Δ(ln Wage)	Δ(ln Wage)	Δ(ln Wage)	Δ(ln Wage)
Share of Cubans in 1980	1.260	1.262						
	(.0529)	(.0532)						
Inflows of Cubans			−.854	−1.313	−1.350	−.857	−1.264	−1.310
			(.381)	(.338)	(.346)	(.383)	(.320)	(.322)
Change in native population		−.00163			.0388			.0385
		(.00117)			(.0450)			(.0382)
Observations	152	152	44	152	152	44	152	152
Education FE	Yes	Yes	No	Yes	Yes	No	Yes	Yes
MSA FE	Yes	Yes	No	Yes	Yes	No	Yes	Yes
MSAs	All	All	All	All	All	All	All	All
Sample	All	All	HSDO	All	All	HSDO	All	All

Note.—Panels A and B show the estimates of the wages in Miami relative to various control groups of cities in 1981–85, relative to before 1981. Panel A uses March CPS data, and panel B uses CPS ORG data. "All MSAs" refers to the 44 or 45 cities covered by the March CPS and ORG CPS throughout the period. Card's control group includes Atlanta, Houston, Los Angeles, and Tampa; Borjas's control group includes Anaheim, Rochester, Nassau-Suffolk, and San Jose; and Peri-Yasenov's control group includes New Orleans, New York City, and Baltimore. 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 parentheses. MSA = metropolitan statistical area; FE = fixed effects; HSDO = high school dropouts.

relative to the control groups. To investigate that, I use the following regression framework:

$$\mathrm{InMiami}_{i,t} = \delta_c + \beta_1 \mathrm{Years1981–1984}_t + \beta_2 \mathrm{Years1985–1990}_t + \varepsilon_{i,t}, \qquad (10)$$

where $\mathrm{InMiami}_{i,t}$ is a variable that takes value one if individual i is in Miami at time t, $\mathrm{Years1981–1984}_t$ is a dummy variable that takes value one for the years 1981–84, and $\mathrm{Years1985–1990}_t$ is a dummy variable that takes value one for the years 1985–90. I run this regression using all high school dropout workers in Miami and in the control group over the period 1977–90. Hence, β_i captures the share of low-skilled workers in Miami relative to the omitted time period (1977–80), 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 probit models in what follows. Finally, note that, as before, I can in fact plot an estimate for each of the years in the regression.

To gain understanding of 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. That is the case when we compare Miami to rest of the United States, to the Card and Borjas placebos, or to all the metropolitan areas in the South Atlantic region.[8]

A second remarkable aspect shown in panel B of figure 2 is that the relative concentration of low-skilled workers in Miami seems to last only until 1984 or 1985. After that, it seems to decline. Depending on the control group, the decline seems to be complete or 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 but that it disappears somewhat 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 United States, we see that Miami gained low-skilled workers in the period 1981–84, then lost some of these workers. In the period 1985–90, however, Miami retained roughly two-thirds of the low-skilled workers gained in the early 1980s, when compared to the United States overall. 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 affected only low-skilled workers.

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

TABLE 2
Estimation of the Causal Effect of Cuban Immigration on Internal Migration

Variables	Prob (Miami=1) (1)	Prob (Miami=1) (2)	Prob (Miami=1) (3)	Prob (Miami=1) (4)	Prob (Miami=1) (5)
	A. Internal Migration of Low-Skilled Workers				
Years 1981–84	.124	.0675	.203	.136	.139
	(.0321)	(.0446)	(.0582)	(.0451)	(.0494)
Years 1985–90	.0945	−.0341	.0563	.156	.0370
	(.0292)	(.0403)	(.0541)	(.0417)	(.0453)
Observations	44,845	10,668	3,971	8,158	6,643
Controls	Yes	Yes	Yes	Yes	Yes
Comparison to	All MSAs	Card control	Borjas control	Peri-Yasenov control	South Atlantic region
	B. Internal Migration of High-Skilled Workers				
Years 1981–84	.0249	.00964	.0640	.0474	.0392
	(.0205)	(.0281)	(.0311)	(.0297)	(.0296)
Years 1985–90	.0485	−.0128	.0848	.0634	−.0475
	(.0180)	(.0247)	(.0275)	(.0264)	(.0260)
Observations	181,054	29,357	17,345	22,587	25,783
Controls	Yes	Yes	Yes	Yes	Yes
Comparison to	All MSAs	Card control	Borjas control	Peri-Yasenov control	South Atlantic region

Note.—The table presents estimates of the probability of being in Miami for low-skilled (panel A) and high-skilled (panel B) workers 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 MSAs" (metropolitan statistical 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's control group includes Anaheim, Rochester, Nassau-Suffolk, and San Jose; and Peri-Yasenov's control group includes New Orleans, New York City, and Baltimore. Robust standard errors are reported in parentheses.

D. Long-Run Estimates

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

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

where $\Delta \ln w_{ce}$ is the change in wages of workers of education e between 1980 and 1990 in metropolitan area c and $\mathrm{Cub}_{ce}/\mathrm{Nat}_{ce}$ is the Mariel Boatlift–induced shock to labor supply in each city and education group, which is measured as the number of Cubans reported in 1990 to have been living in each city in 1985 who claim to have arrived in the United States in 1980–81 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, respectively. 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 Cub_{ce}/Nat_{ce} by the share of Cubans in each city before the Mariel Boatlift shock, as explained in section III.B.

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. That is, in the short run, before any adjustments, β is the (inverse) local labor demand elasticity. If there are adjustments, then β also contains those adjustments.

Table 3 reports these results. Column 1 of panel A shows that the initial share of Cubans (among high school dropouts) 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 along with the metropolitan area and education fixed effects. In columns 1 and 2 of panel B, I estimate the wage effects over the entire decade, using IV regressions. It is clear from these two columns that the wages of low-skilled workers in high-Cuban locations do not seem to be lower than those in lower-Cuban migration locations. Similarly, rentals do not seem to have been affected differentially over the decade

TABLE 3
Estimation of the Causal Effect of Cuban Immigration on Long-Run Wages, Rents, and Internal Migration

Variables	Inflow of Cubans: First Stage (1)	Inflow of Cubans: First Stage (2)	ΔShare Low Skilled: OLS (3)	ΔShare Low Skilled: IV (4)
	A. First-Stage and Internal Migration			
Lagged share of Cubans	.716	1.231		
	(.0169)	(.0845)		
Inflow of Cubans			.604	.641
			(.0903)	(.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				1,797
	B. Wages and Rents (IV)			
	Δ(ln Wage)	Δ(ln Wage)	Δ(ln Rent)	Δ(ln Rent)
Inflow of Cubans	.113	−.0858	.180	.0779
	(.517)	(.198)	(.361)	(.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	1,797	212.3	1,797	212.3

Note.—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 the instrument. HSDO indicates that the regression is restricted to high school dropouts. This table uses variation from the 38 metropolitan areas available in the census and CPS data throughout this period. "Widstat" indicates the F-statistic of the excluded instrument in the first-stage regression. FE = fixed effects.

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 zero.

To investigate how much the local supply of skills changed 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 the change is taken between 1980 and 1990. In this case, an estimate of $\lambda = 0$ indicates that there is no change in the local supply of skills. That is, for each Cuban low-skilled immigrant, the share of low-skilled workers increases by exactly 1. Instead, if $\lambda = 1$, then it means that internal migration (or other factors, such as human capital acquisition) 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.

The 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 0.6, that is, $\hat{\lambda} = 0.4$. That means that there was some adjustment in the local supply in response to the Mariel Boatlift but that Miami gained low-skilled workers relative to the other cities in the United States.

This local adjustment could come from internal migration or from human capital acquisition. To be more convinced that these estimates capture internal migration and not other forms of local labor supply adjustment, I investigate whether cohorts born in Florida who were around 18 years old around 1980 systematically acquire more education than those born elsewhere. I further compare cohorts under 18 years old in 1980 with those over 18, under the assumption that cohorts younger than 18 years old could more easily adjust their educational attainment. Table 4 shows that the interaction of a Florida dummy with a dummy for cohorts under 18 years old is close to zero and not statistically significant. Panel A shows this exercise by looking at the share of natives who drop out of school, while panel B shows the same results using, instead, natives with at least a high school diploma.

IV. Decomposition

With the estimates provided in section III, 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

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

TABLE 4
Human Capital Acquisition in Florida in Cohorts Just under and over 18 in 1980

Variables	(1)	(2)	(3)	(4)
	A. Share of Natives Who Dropped Out of School (OLS)			
Florida (born in Florida)	.00646	.0471	.0244	.0183
	(.00738)	(.00252)	(.0100)	(.00314)
Treated (cohort 10–18 in 1980)	−.00100	.00864	.00180	−.00245
	(.00868)	(.00278)	(.0123)	(.00348)
Treated × Florida	−.00221	−.0119	−.00501	−.000760
	(.00900)	(.00368)	(.0125)	(.00416)
Observations	52	39	52	663
Comparison to	Card control	Borjas control	Peri-Yasenov control	All MSAs
	B. Share of Natives with a High School Diploma (OLS)			
Florida (born in Florida)	.0201	.0680	.0431	.0121
	(.00787)	(.00785)	(.0148)	(.00500)
Treated (cohort 10–18 in 1980)	−.00401	−.00462	−.0132	−.0144
	(.00898)	(.00907)	(.0181)	(.00421)
Treated × Florida	−.0107	−.0101	−.00146	−.000276
	(.0115)	(.0116)	(.0194)	(.00805)
Observations	52	39	52	663
Comparison to	Card control	Borjas control	Peri-Yasenov control	All MSAs

Note.—This table compares the share of native workers with various levels of education born in Florida and other states for the cohorts who were around 18 years old during the Mariel Boatlift episode. Each observation is an age × state of birth cell, for which I computed the share of workers with a particular education level in the year 2000. The regressions are limited to cohorts who were 10–26 years old in 1980. Cohorts who were 10–18 years old could have adapted their educational attainment to the arrival of Mariel Boatlift immigrants. Robust standard errors are reported. MSA = metropolitan statistical area.

where again, $\widehat{\varepsilon^{L}}$ is an estimate of the (inverse) local labor demand elasticity, which under the assumption made in section II 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 those in the other panels of table 1. If the labor supply shock was equivalent to 25% of the low-skilled labor force and wages are estimated to have declined by between 10% and 30%, it means that the inverse local labor demand elasticity is between 0.4 and 1.2. The term $\hat{\lambda}$ is for the long-run internal migration response, which we have estimated in table 3 to be around 0.4. Finally, as a reminder, for the calculation of $\hat{\nu}$ we need estimates of α, which I set equal to 0.25, and $\epsilon = 0.6$—which are the estimates available in the literature (Davis and Ortalo-Magné 2011; Saiz 2010).

With all these estimates I can decompose the recovery into internal migration and other factors, as explained in section II.C. I show this exercise in table 5, both using the baseline estimates and by providing a number of alternative decompositions assuming alternative wage, internal migration, consumption of housing, and housing supply elasticity estimates.

TABLE 5

CONTRIBUTION OF VARIOUS FACTORS TO WAGE RECOVERY

Parameter	Inverse Local Labor Demand Elasticity $\widehat{\varepsilon^{L}}$	Internal Migration Response $\hat{\lambda}$	Share of Income to Housing $\hat{\alpha}$	Housing Supply Elasticity $\hat{\epsilon}$	Indirect Utility Elasticity $\widehat{\tilde{\varepsilon}^{L}}$	Other Factors $\hat{\nu}$	Contribution to Recovery: Other Factors (%)	Contribution to Recovery: Internal Migration (%)	Internal Migration Elasticity
	A. Baseline								
Baseline	1.0	.4	.25	.6	1.0	.6	60	40	1.6
Very elastic LD	.4	.4	.25	.6	.5	.3	30	70	4.0
Elastic LD	.7	.4	.25	.6	.7	.4	45	55	2.3
Inelastic LD	1.4	.4	.25	.6	1.3	.8	80	20	1.1
	B. Higher Share of Income Devoted to Housing								
Baseline	1.0	.4	.30	.6	1.0	.6	60	40	1.6
Very elastic LD	.4	.4	.30	.6	.5	.3	31	69	4.0
Elastic LD	.7	.4	.30	.6	.8	.5	45	55	2.3
Inelastic LD	1.4	.4	.30	.6	1.3	.8	80	21	1.1
	C. Higher Housing Supply Elasticity								
Baseline	1.0	.4	.25	1.5	1.0	.6	60	40	1.6
Very elastic LD	.4	.4	.25	1.5	.5	.3	28	72	4.0
Elastic LD	.7	.4	.25	1.5	.7	.4	44	56	2.3
Inelastic LD	1.4	.4	.25	1.5	1.4	.8	82	18	1.1
	D. Higher Internal Migration Response								
Baseline	1.0	.6	.25	.6	1.0	.4	40	60	2.4
Very elastic LD	.4	.6	.25	.6	.5	.2	20	80	6.0
Elastic LD	.7	.6	.25	.6	.7	.3	30	70	3.4
Inelastic LD	1.4	.6	.25	.6	1.3	.5	54	47	1.7

Note.—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 a sensitivity analysis of these results to alternative estimates of the key parameters. LD = labor demand.

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 and 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 Miami's local shock, which was around 25%–30%. To be conservative, I assume that the shock was equivalent to 25% 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 (Diamond 2016; Caliendo, Dvorkin, and Parro 2019; Monras 2020).

The first row shows the baseline estimates. The baseline estimates suggest that around 40% 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, that is, within 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, -0.7, or -1.4. This covers the range of estimates in the literature. I do that 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 5. Panel A shows that internal migration accounts for 20%–70% of the recovery.

In panel B of table 5, I show the same results but assume that α is 0.3 instead of 0.25. This exercise is justified, as it is sometimes argued that in larger cities the share of income devoted to housing is higher or because we can interpret housing as a broader nontradable sector. The decomposition of the recovery is again similar, with estimates of the importance of internal migration that fluctuate around 50%.

Panel C shows the 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 the results significantly.

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 10-year horizon internal migration response. Not surprisingly, this exercise increases the relative importance of internal migration, although the numbers are still similar to the baseline estimates.

Taken altogether, table 5 suggests that (through the lenses of the model introduced in sec. II) internal migration accounts for roughly 50% 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.

V. Conclusion

In this paper, I introduce a spatial equilibrium model that allows analysis of the effects of immigrant-induced labor supply shocks in the short and long runs. The model shows how we can use short-run regressions to recover key parameters that govern the local labor demand elasticity. It then shows how we can run longer–time horizon regressions to establish when the economy might return to the initial spatial equilibrium. With these, it is then possible to evaluate how different adjustment factors might have contributed to the local adjustment and recover its relative importance from the structure of the model.

I illustrate how these procedures can be applied, using the well-studied Mariel Boatlift episode. Through the lenses of the model and given the estimates that I report in this paper, the evidence suggests that around 50% of the wage recovery over the 1980s in Miami, relative to a number of potential control locations, is explained by internal migration, with the rest explained by other factors, such as technology adoption. This result is robust to a number of sensitivity checks.

References

Abadie, Alberto. 2020. "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects." *J. Econ. Literature*, forthcoming.

Abadie, Alberto, and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *A.E.R.* 93 (1): 113–32.

Albert, Christoph, and Joan Monras. 2020. "The Regional Impact of Economic Shocks: Why Immigration Is Different from Import Competition." Econ. Working Paper no. 1223, Barcelona Graduate School Econ.

Altonji, Joseph G., and David Card. 1991. "The Effects of Immigration on the Labor Market Outcomes of Less-Skilled Natives." In *Immigration, Trade, and the Labor Market*, edited by John M. Abowd and Richard B. Freeman, 201–34. Chicago: Univ. Chicago Press (for NBER).

Amior, Michael. 2020. "Immigration, Local Crowd-Out and Undercoverage Bias." Res. Online Documents Econ. no. 108490, London School Econ. and Polit. Sci.

Borjas, George J. 2003. "The Labor Demand Curve *Is* Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *Q.J.E.* 118 (4): 1335–74.

———. 2006. "Native Internal Migration and the Labor Market Impact of Immigration." *J. Human Resources* 41 (2): 221–58.

———. 2017. "The Wage Impact of the *Marielitos*: A Reappraisal." *Indust. and Labor Relations Rev.* 70 (5): 1077–1100.

Borjas, George J., Richard B. Freeman, and Lawrence F. Katz. 1996. "Searching for the Effect of Immigration on the Labor Market." *A.E.R.* 86 (2): 246–51.

———. 1997. "How Much Do Immigration and Trade Affect Labor Market Outcomes?" *Brookings Papers Econ. Activity* 1997 (1): 1–67.

Borjas, George J., and Joan Monras. 2017. "The Labour Market Consequences of Refugee Supply Shocks." *Econ. Policy* 32 (91): 361–413.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro. 2019. “Trade and Labor Market Dynamics.” *Econometrica* 87 (3): 741–835.

Card, David. 1990. “The Impact of the Mariel Boatlift on the Miami Labor Market.” *Indust. and Labor Relations Rev.* 43 (2): 245–57.

———. 2009. “Immigration and Inequality.” *A.E.R.* 99 (2): 1–21.

Card, David, and John DiNardo. 2000. “Do Immigrant Inflows Lead to Native Outflows?” *A.E.R.* 90 (2): 360–67.

Cascio, Elizabeth U., and Ethan G. Lewis. 2018. “How Much Does Amnesty Strengthen the Safety Net? Evidence from the Immigration Reform and Control Act of 1986.” Manuscript.

Clemens, Michael A., and Jennifer Hunt. 2019. “The Labor Market Effects of Refugee Waves: Reconciling Conflicting Results.” *Indus. and Labor Relations Rev.* 72 (4): 818–57.

Clemens, Michael A., Ethan G. Lewis, and Hannah M. Postel. 2018. “Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion.” *A.E.R.* 108 (6): 1468–87.

Cortes, Patricia. 2008. “The Effect of Low-Skilled Immigration on U.S. Prices: Evidence from CPI Data.” *J.P.E.* 116 (3): 381–422.

Davis, Morris A., and François Ortalo-Magné. 2011. “Household Expenditures, Wages, Rents.” *Rev. Econ. Dynamics* 14 (2): 248–61.

Diamond, Rebecca. 2016. “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000.” *A.E.R.* 106 (3): 479–524.

Dustmann, Christian, Tommaso Frattini, and Ian P. Preston. 2013. “The Effect of Immigration along the Distribution of Wages.” *Rev. Econ. Studies* 80 (1): 145–73.

Dustmann, Christian, and Albrecht Glitz. 2015. “How Do Industries and Firms Respond to Changes in Local Labor Supply?” *J. Labor Econ.* 33 (3): 711–50.

Dustmann, Christian, Uta Schönberg, and Jan Stuhler. 2016. “The Impact of Immigration: Why Do Studies Reach Such Different Results?” *J. Econ. Perspectives* 30 (4): 31–56.

Ehrlich, Isaac, and Yun Pei. 2020. “Human Capital as Engine of Growth: The Role of Knowledge Transfers in Promoting Balanced Growth within and across Countries.” *Asian Development Rev.* 37 (2): 225–63.

———. 2021. “Endogenous Immigration, Human and Physical Capital Formation, and the Immigration Surplus.” *J. Human Capital* 15 (1): 34–85.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. “Bartik Instruments: What, When, Why, and How.” *A.E.R.* 110 (8): 2586–624.

González, Libertad, and Francesc Ortega. 2011. “How Do Very Open Economies Absorb Large Immigration Flows? Evidence from Spanish Regions.” *Labour Econ.* 18 (1): 57–70.

Hanson, Gordon H., and Matthew J. Slaughter. 2002. “Labor-Market Adjustment in Open Economies: Evidence from US States.” *J. Internat. Econ.* 57 (1): 3–29.

Lafortune, Jeanne, Ethan Lewis, and José Tessada. 2019. “People and Machines: A Look at the Evolving Relationship between Capital and Skill in Manufacturing, 1860–1930, Using Immigration Shocks.” *Rev. Econ. and Statis.* 101 (1): 30–43.

Lafortune, Jeanne, José Tessada, and Carolina González-Velosa. 2015. “More Hands, More Power? Estimating the Impact of Immigration on Output and Technology Choices Using Early 20th Century US Agriculture.” *J. Internat. Econ.* 97 (2): 339–58.

Lewis, Ethan. 2003. “Local, Open Economies within the U.S.: How Do Industries Respond to Immigration?” Working Paper no. 04–01, Fed. Reserve Bank, Philadelphia.

———. 2011. “Immigration, Skill Mix, and Capital-Skill Complementarity.” *Q.J.E.* 126 (2): 1029–69.

———. 2013. “Immigration and Production Technology.” *Ann. Rev. Econ.* 5:165–91.
Monras, Joan. 2020. “Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis.” *J.P.E.* 128 (8): 3017–89.
Peri, Giovanni, and Chad Sparber. 2011. “Assessing Inherent Model Bias: An Application to Native Displacement in Response to Immigration.” *J. Urban Econ.* 69 (1): 82–91.
Peri, Giovanni, and Vasil Yasenov. 2019. “The Labor Market Effects of a Refugee Wave: Synthetic Control Method Meets the Mariel Boatlift.” *J. Human Resources* 54 (2): 267–309.
Ruggles, Steven, Matthew Sobek, Trent Alexander, et al. 2016. “Integrated Public Use Microdata Series [IPUMS]: Version 4.0 [machine-readable database].” Minneapolis: Minnesota Population Center.
Saiz, Albert. 2003. “Room in the Kitchen for the Melting Pot: Immigration and Rental Prices.” *Rev. Econ. and Statis.* 85 (3): 502–21.
———. 2010. “The Geographic Determinants of Housing Supply.” *Q.J.E.* 125 (3): 1253–96.