

The Migration Accelerator: Labor Mobility, Housing, and Aggregate Demand

Greg Howard *

March 18, 2017

Abstract

One of the traditional criteria for optimal currency areas is having high labor mobility to smooth differences across regional labor markets (Mundell, 1961). Contrary to this view, I show that labor mobility amplifies the effects of local economic shocks and exacerbates differences across regions, implying that labor mobility is not helpful for having a unified monetary policy. In the first part of this paper, I document that within-U.S. migration causes a large short-term reduction in the unemployment rate of the receiving city. To establish the causal effect of immigration, I construct a plausibly exogenous shock by using the outmigration of other places and predicting its destination based on historical patterns. In the second part of the paper, I document that the increase in the demand for housing explains the boom, through two channels. The construction channel occurs because housing is a durable good: hence there is a surge in the number of new houses and construction jobs. The house price channel occurs because the migrants' housing demand drives up prices, leading to increased borrowing and higher labor demand in non-tradable sectors. Together, these channels account for the size of the labor demand boom. In the final part of the paper, I estimate that migration amplifies these shocks by 20 percent and show that changing patterns of labor mobility increased the geographic variance of the effects of the Great Recession.

*MIT. Email address: goward@mit.edu. Latest version: <http://economics.mit.edu/files/12175>. I would like to thank my advisors, Iván Werning, Arnaud Costinot, and David Autor for their guidance. I would like to thank Andrea Raffo for a comment that inspired the title. I would also like to thank Alp Simsek, David Atkin, Daron Acemoglu, Frank Schilbach, Emi Nakamura, Isaiah Andrews, Ludwig Straub, Jack Liebersohn, John Firth, Arianna Ornaghi, Rachael Meager, Sebastian Fanelli, Yu Shi, Vivek Bhattacharya, Alex Bartik, and seminar and lunch participants at MIT, the Federal Reserve Board of Governors, the Federal Reserve Bank of Chicago, and the Federal Reserve Bank of Minneapolis. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 2014136625. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

For the euro area, too, increased labour mobility across borders is crucial. We should see that the current institutional arrangements of the welfare state do not stand in the way of such mobility.

Mario Draghi, Paris, November 30, 2012

1 Introduction

One of the traditional criteria for optimal currency areas is having high labor mobility across regions (Mundell, 1961). Such movements are traditionally thought to smooth differences in labor markets across space, so the appropriate stance of monetary policy is similar throughout the currency union. Because the United States has a higher rate of labor mobility than the euro area, it is often thought to have an advantage for unified monetary policy (Dao, Furceri, and Loungani, 2014).

In this paper, I present a different view of labor mobility and provide empirical evidence in favor of it. I show that labor mobility amplifies the effects of local economic shocks and exacerbates differences across regions. The results imply that labor mobility is not helpful for unified monetary policy.

This traditional view is that migration smooths regional differences as people move into more prosperous areas. Hence, local shocks are spread across regions, and migration provides insurance even for non-movers. Empirically, this view is most closely associated with Blanchard and Katz (1992), which showed that net migration responded positively to increased labor demand, so that the employment level never returned to trend, even though the unemployment rate did. However, the traditional view also relies on a key implicit assumption, which is present in many standard models: that migration causes slack in the receiving labor market.

In the first part of this paper, I empirically test this assumption by estimating the effect of domestic migrants on the labor market of the receiving MSA. To get a causal effect, I construct an immigration shock using the pre-period migration network and current out-migration from connected MSAs, similar to Altonji and Card (1991). I find that the key

implicit assumption in the traditional view does not hold up empirically: within-U.S. migration from 1997 to 2010 causes a large local labor market boom in the receiving MSA. An immigration shock the size of one percent of the MSA’s population causes a decrease in the unemployment rate by a third of a percent.¹

This result necessitates a new view of migration: instead of mitigating the labor market effects of local shocks, migration amplifies them. When a region experiences an increase in labor demand, it will attract additional migrants, causing an even larger boom. I call this the “migration accelerator,” the additional increase in employment because of the endogenous response of immigration.

The decrease in unemployment is surprising because immigration increases the labor supply, so if labor demand curves are downward sloping, the wage will fall. And if wages are rigid, the unemployment rate will rise instead.² Migrants also move with their labor demand, as stressed in Farhi and Werning (2014), counterbalancing the increase in labor supply. However, as long as agents consume some non-local goods, their labor supply will exceed their labor demand so these effects cannot explain the reversal of sign. In order to make sense of my empirical result, I propose housing as an additional mechanism. My hypothesis is that the housing boom can increase labor demand by more than the corresponding increase in labor supply.

In the second part of the paper, I provide additional empirical evidence of two housing channels. The first channel I label the construction channel. Housing is durable and requires local labor. When migrants move in, the demand for housing rises. In the long-run, the

¹I focus on immigration for two reasons. First, I show the majority of the net migration response to labor demand shocks, as constructed in Bartik (1991), is through immigration, not outmigration. This is consistent with Monras (2015a), which uses different shocks and data, and Coen-Pirani (2010), which notes that immigration is much more volatile than outmigration. In a different setting, Long and Siu (2016) find that during the Dust Bowl era, the fall in net migration was also due to immigration and not outmigration. So to better understand how migration changes the effects of these shocks, it is of primary interest to understand the effects of immigration. Second, because outmigration is less correlated with local economic conditions, its fluctuations are an appealing source of variation for immigration elsewhere.

²Sticky wages are not key for the housing mechanisms I focus on. Without them, a labor market clearing condition would imply that wages would rise instead of unemployment falling. I use wage rigidity because I find large and persistent effects on the unemployment rate but not hourly wages, which would be hard to explain using search models of unemployment, as stressed by Shimer (2005).

steady-state housing stock will increase. But since housing is a durable good, its short-run production must increase in order to reach that steady-state.³ In the data, I see a short-term increase in housing permits and construction employment.

The second channel I label the house price channel. The increased housing demand causes house prices to rise, which has been documented to have a large effect on consumption (Mian, Rao, and Sufi, 2013; Campbell and Cocco, 2007; Attanasio, Leicester, and Wakefield, 2011; Agarwal, Amromin, Chomsisengphet, Piskorski, Seru, and Yao, 2015; Ströbel and Vavra, 2015; Kaplan, Mitman, and Violante, 2016). Berger, Guerrieri, Lorenzoni, and Vavra (2015) summarize the literature and suggest the average estimated elasticity is about 0.2. I document an increase in house prices, second-lien mortgages, and non-tradable goods employment, consistent with this channel.

In areas with inelastic housing supply, one would expect house prices to increase by more in response to the shock, increasing the effect on the unemployment rate through the house price channel.⁴ Indeed, for MSAs with lower elasticities as measured by Saiz (2010), house prices do increase more and the unemployment rate declines more in response to the shock.

Next, I show evidence that these housing channels were in operation in a different context. In 1980, Miami's population increased substantially due to the Mariel boatlift, an unanticipated mass emmigration from Cuba. I document similar patterns in house prices and employment composition during this episode, when compared to comparison cities.⁵

³The logic is similar to how Aiyagari, Christiano, and Eichenbaum (1992) and Baxter and King (1993) model the response of output to government spending. They find large effects on output and employment of permanent changes in government spending because of the response of investment. Similarly, Rognlie, Shleifer, and Simsek (2015) consider a model with too much initial housing stock relative to steady-state, causing a recession due to the zero lower bound.

⁴The theory is ambiguous on the relative magnitude of the construction channel. If housing demand is inelastic and labor shares are higher for multi-family housing, as suggested in Stucke (1949) and Carliner (2003), the construction channel would be larger in inelastic areas as well.

⁵Saiz (2003) investigates the effects of the Mariel boatlift on the housing market, finding a large increase in rental prices. However, he does not find an increase in house prices, in contrast to my findings. A key difference is that he looks at house price changes beginning in 1980Q3, after the boatlift, whereas I use a baseline of 1979. Lewis (2004) looks at changes in industry mix after the Mariel boatlift as well, but focuses exclusively on tradable goods. His hypothesis is that the composition of tradable goods changes according to the Rybczynski (1955) theorem. Hanson and Slaughter (2002) and Lewis (2005) also find evidence of industry composition effects within tradable goods, in other settings. These papers do not focus on the role of non-tradables or construction, the key industries driving my results.

The sign of the effects are similar to what I find for domestic migration, but the magnitudes are smaller. A plausible explanation for the difference is that the domestic migrants which I focus on are likely to demand more housing.⁶

Other theories could explain my main result, that the unemployment rate falls after an immigration shock, including factor complementarity, increasing returns to scale, and love for variety. While they can explain the sign, they cannot explain the timing, magnitude, sectoral composition, and relation to housing supply elasticity. In contrast, the housing channels easily explain my headline results and each of these facts. I expand on this in Section 4.5.

In the final part of the paper, I quantify the economic impact of the previous results through two counterfactual exercises. In the first, I calculate the difference between the effect of a labor demand shock when migration endogenously responds and the counterfactual where migration is held constant. I label this difference the “migration accelerator.” To estimate it, I first calculate the size of the endogenous migration response to a Bartik (1991)-style shock. I then combine that with my main results. I find that migration amplifies the effect on the unemployment rate by 20 percent locally.

Using a similar strategy for the second counterfactual, I calculate the effect that changing migration patterns had on the unemployment rate during the Great Recession. I construct a counterfactual in which migration patterns, but not levels, are held constant from 2004, and consider how the unemployment rate would have moved differently during the Great Recession. I find that the increase in the unemployment rate in Florida and Southern California would have been less severe, and that overall, the standard deviation of changes in the unemployment rate would have been reduced by 13 percent.

My empirical strategy and question are close to a literature on estimating the labor market effects of *international* migration, but with several important differences in setting and

⁶Greulich, Quigley, and Raphael (2004) shows that, compared to natives, immigrants live in houses with fewer rooms and bedrooms, but with more families per unit and more people per room, using data from the 2000 Census.

methodology, and starkly different results. In particular, I use a similar empirical strategy to Altonji and Card (1991), Card (2001), Lewis (2005), Saiz (2007) and Hong and McLaren (2015), which combine the location of immigrant communities and immigrants coming from that country to construct an instrument. The literature, using this and other methodologies, has found a range of effects of international migration on labor market outcomes, typically wages, ranging from a modest positive effect (Ottaviano and Peri, 2006; Card, 2009) to a large negative effect from immigration (Borjas, 2003; Monras, 2015b).⁷ Even different studies of the Mariel boatlift have divergent results. Card (1990) concluded that wages of comparable workers in Miami were largely unaffected, while Borjas (2015) found that they dropped dramatically. Other methodologies have also produced a range of results, as summarized by Dustmann, Schönberg, and Stuhler (2016). My results lie outside of the range of previous findings, documenting a substantial improvement in the local labor market in response to an increase in domestic immigration.

Another strand of the immigration literature looks specifically at housing, but primarily because a rise in housing costs is evidence that immigrants do not completely displace natives, and because changes in house prices are a transfer of wealth from immigrants to natives. Saiz (2003, 2007), Gonzalez and Ortega (2013), Greulich et al. (2004) and Ottaviano and Peri (2006) find positive effects of immigration on housing costs.⁸ I find larger effects on house prices, and I focus on how rising house prices are a part of a housing-led demand boom.

Many papers since Blanchard and Katz (1992) have looked at population adjustments in response to local economic shocks in different settings or time periods (see Decressin and Fatas, 1995; Jimeno and Bentolila, 1998; Bound and Holzer, 2000; Cadena and Kovak,

⁷Hong and McLaren (2015) report a sizable increase in the employment of cities that receive immigrants, finding that each immigrant creates 1.2 jobs in the receiving city. They also find a sizable increase in the number of natives in the labor force. For native workers, the number of jobs increases by 0.86, while the native labor force increases by 0.97, per migrant. If the native unemployment rate were below 10 percent initially, this would imply migration raised the unemployment rate. In 1990 and 2000, the unemployment rate in the U.S. averaged 5.6 and 4.7 percent.

⁸Saiz finds positive effects on rents in both studies. Saiz (2003) finds a negative effect on house prices, while Saiz (2007) finds a positive effect. He suggests these opposite results are because natives prefer not to live near immigrants. The other listed papers find a positive effect.

2016; Monras, 2015a). Charles, Hurst, and Notowidigdo (2016) find that the working-age population in MSAs with housing booms increased while MSAs with manufacturing declines experienced population decreases. All of these papers find evidence that labor mobility responds to the conditions of local labor markets. My results contribute to this literature by providing an explanation for why high rates of labor mobility do not close the gaps in labor market outcomes across space, similar to Amior and Manning (2016).

There is also a large literature on how house prices change migration decisions (see Van Nieuwerburgh and Weill, 2010; Notowidigdo, 2013; Davis, Fisher, and Veracierto, 2013; Head, Lloyd-Ellis, and Sun, 2014; Nenov, 2015). Several papers, including Struyven (2014) and Schulhofer-Wohl (2011), focus on housing lock, the idea that underwater mortgages prevent migration. In contrast to all these papers, I highlight a different role for housing: how increased housing demand can have strong effects on local labor demand. Indeed, because labor demand is a draw for migrants, this might explain why these papers find a small role of house prices in migration decisions.

My work is closely related to the economic geography literature, exploring the role of migration for the propagation of economic shocks. (see Caliendo, Parro, Rossi-Hansberg, and Sarte, 2014; Kline and Moretti, 2014; Allen and Arkolakis, 2014; Diamond, 2016; Redding and Rossi-Hansberg, 2016). I also study how local shocks interact with migration in equilibrium. In fact, many of these papers allow for migration to improve labor markets through increasing returns to scale, or through trade costs similar to Krugman (1980). However, in contrast to these papers, I find evidence for a new short-term local spillover which is typically not present in these models.

The rest of my paper is organized as follows. Section 2 reviews the traditional view and puts forth an alternative new view based on housing, using a simple conceptual framework. Section 3 outlines my empirical strategy and shows that immigration has an expansionary effect on local labor markets in the United States. Section 4 presents the evidence in favor of the construction and house price channels. Section 5 quantifies the size of the migration

accelerator and considers the role migration played during the Great Recession.

2 Conceptual Framework

In this section, I present a simple model of a city, to review the traditional view in a transparent way, and to present the new view with housing. The purpose is to clarify how the two views differ and to present two important channels related to housing. Along the way, I make a number of simplifying assumptions so that the intuitions are more transparent. In Appendix A, I present a microfounded model that extends the insights presented here to a dynamic framework with more general housing demand, relaxing many of those assumptions.

The framework focuses on two equilibrium relationships between immigration and unemployment. In both views, the first relationship is that a lower unemployment rate will draw more migrants. The second relationship is the causal effect migration has on unemployment. Whether it increases or decreases unemployment will separate the traditional and the new views.

2.1 The Traditional View

Normalize the population of the city to be 1, and denote by D the total labor demand in that city. Define m to be the immigration rate, and u to be the unemployment rate. Assume the wage is fixed, so the only endogenous variables are u and m . There are two shocks in the model, a migration shock, ϵ^m , and a labor demand shock, ϵ^d . Consider the following two equations that rationalize the traditional view:

$$m = m^{BK}(u) + \epsilon^m \quad (\text{Blanchard-Katz})$$
$$u = 1 - \frac{D + \epsilon^d}{1 + m} \quad (\text{Traditional})$$

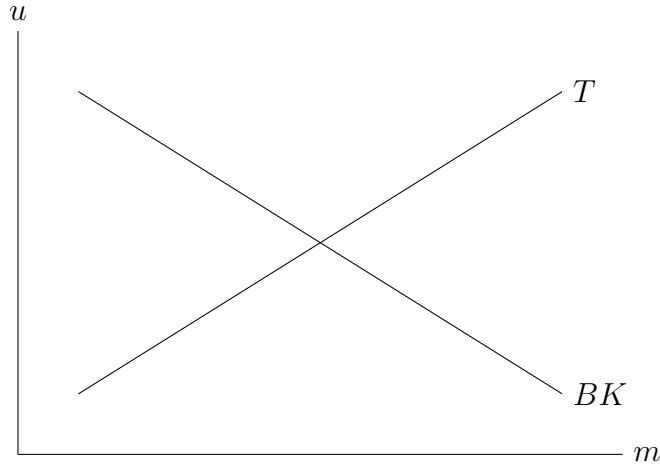


Figure 1: A graphical representation of the traditional-view equilibrium.

where m^{BK} is decreasing in u . The Blanchard-Katz equation simply says that migration will increase when the unemployment rate falls.⁹ Blanchard and Katz (1992) and many other papers empirically establish this relationship. The Traditional equation considers the effect of migration on the unemployment rate. It is an identity: u is equal to one minus the number of jobs divided by the labor supply. If D is held fixed, then in the Traditional equation, u is increasing in m . For a given ϵ^d and ϵ^m , the intersection of these two curves, shown in Figure 1, will determine the equilibrium.

Because the Traditional line is upward sloping, a change in ϵ^d will be dampened by the effect of migration. For example, if ϵ^d is positive, the Traditional line will shift down. This will cause m to increase in equilibrium and the change in u will be less than the magnitude of the shift.

2.2 The New View

Empirically, I can estimate the slope of the second equation by looking at the effect of ϵ^m . In Section 3, I do that, and I find that unemployment falls in response to these shocks, which is inconsistent with an upward sloping curve. This requires a revisiting of the equation

⁹Harris and Todaro (1970) present a model in which people based migration decisions on the unemployment rate, in which workers consider their expected income in the destination city, which depends on both the wage and the probability of finding a job. My microfoundation in Appendix A has similar features.

governing migration's effect on unemployment. A realistic extension is that some labor demand is local, and hence D is not fixed. Migrants will consume non-tradable goods in the city which they move. Furthermore, in equilibrium, the current unemployment rate also affects the labor demand of consumers. This leads us to the New equation:

$$u = 1 - \frac{D(m, u) + \epsilon^d}{1 + m} \quad (\text{New})$$

How should we think about $D(m, u)$? Some of demand is external; for example, the demand for cars produced in Detroit is not largely affected by the number of people living nearby. Denote this by D_x . Some is internal, such as the demand for restaurants. If I assume each person, in expectation, consumes the same amount of non-tradable goods and services, conditional on the unemployment rate, I can write $D(m, u)$ as

$$D(m, u) = D_x + (1 + m)c^{NT}(u)$$

where c^{NT} is normalized to the labor required to produce it. Assume $c^{NT}(u)$ is decreasing in u , but that the slope is less than one so that the equilibrium is unique and stable.

The consumption of non-tradables makes the Traditional line much less steep.¹⁰ However, it can be shown that plugging this new equation for $D(m, u)$ into the New equation cannot lead to a downward sloping line.¹¹ This result is general to a dynamic setting. In Appendix A, I show that a similar setup cannot produce a decline in the unemployment rate. Hence, this model of labor demand requires an extra ingredient in order to explain the empirical findings.

Housing can play that role. I add housing into this model using the following series of simplifying assumptions for tractability, which I relax in Appendix A. Assume that each agent demands exactly one house. Further assume that each non-migrant owns $1 - \delta$ units

¹⁰In the extreme case, where D_x is zero, the New line would be flat.

¹¹The equation simplifies to $u - c^{NT}(u) = 1 - \frac{D_x + \epsilon^d}{1 + m}$. The left hand side is increasing in u , and the right hand side is increasing in m .

of housing, and that their average consumption of non-tradables is increasing in the price of housing.¹² Finally, assume that housing is produced using labor and a fixed factor, land.¹³ The price of housing and the labor required to build housing are both increasing in the amount of new housing demanded. The total labor demand is now given by:

$$D(m, u) = D_x + c_n^{NT}(u, p^h(\delta + m)) + mc_m^{NT}(u, p^h(\delta + m)) + D_h(\delta + m)$$

where $D^h(\cdot)$ is the construction demand from new housing. $c_n^{NT}(u, p^h(\cdot))$, the non-tradable consumption of non-migrants, is increasing in p^h , while $c_m^{NT}(u, p^h(\cdot))$, the non-tradable demand of migrants, is decreasing in p^h . Both are still decreasing in u . In this equation, the increase in labor demand might be proportionally larger than the increase in labor supply. The second term, the consumption of non-migrants, is increasing in the number of migrants because of the effect of house prices, a channel that was not previously present. Furthermore, in the last term, construction demand differs from other non-tradable demand because migrants need an entirely new house, while non-migrants already possess the non-depreciated part of their house from the “previous period.”

These channels become more clear when I plug in the labor demand into the New equation and linearize around $m = 0$.

$$\underbrace{\left(1 + \frac{\partial c_\ell^{NT}}{\partial u}\right)}_{\text{Keynesian multiplier}} \frac{du}{dm} = \underbrace{(1 - u)}_{\text{Labor supply increase}} - \underbrace{c_m^{NT}}_{\text{Migrant non-tradable demand}} - \underbrace{MPCH_n \frac{\partial p^h}{\partial m}}_{\text{House price channel}} - \underbrace{\frac{\partial D_h}{\partial m}}_{\text{Construction channel}} \quad (1)$$

¹²In general, consumption of non-tradables can be influenced by house prices in several ways. Berger et al. (2015) list four channels: the wealth channel, the income channel, the substitution channel, and the collateral channel. In this simplified model, the wealth channel and the income channel cancel out, while the substitution channel is shut down by assumption. So to microfound this, one would need to assume a collateral limit that is influenced by the price of housing. In the general model in Appendix A, I allow for all these channels to be in operation. Empirical estimates of the total effect of house prices on non-tradable consumption are positive.

¹³For simplicity, assume the land is owned by an unmodeled central government, which keeps any profits.

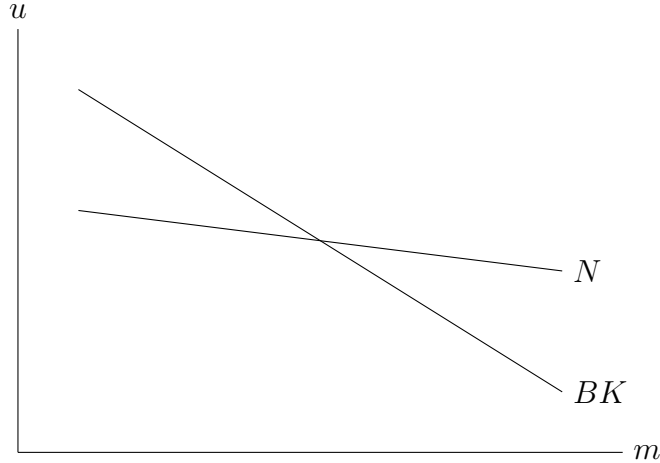


Figure 2: A graphical representation of the new-view equilibrium.

Equation (1) characterizes the four key ways in which migration affects the unemployment rate.¹⁴ First, there is the direct effect of the increase in labor supply, which increases the unemployment rate. It is partially offset by the second effect: the increase in demand for non-tradables. As argued above, these first two effects are always net positive. But there are also two additional channels that come through housing. Third is the house price channel, which is the product of the marginal propensity to consume out of housing wealth, $\frac{\partial c_n^{NT}}{\partial p^h}$, which I denote by $MPCH_n$, and the change in the price of housing. Fourth is the construction channel, which is the additional construction jobs that are needed to build housing for the migrants. If these channels are sufficiently large, the New equation may be downward sloping, a situation illustrated in Figure 2.

Under the new view, if there is an increase in ϵ^d , the New curve still shifts downward, but now the endogenous response of migration lowers the unemployment rate by even more than it did originally.

Equation (1) can be extended dynamically, and with more general housing demand. It frames the empirical investigation in the rest of this paper. Section 3 estimates $\frac{du}{dm}$ in the data. Section 4 estimates the key elements of the construction and house price channels, documenting an effect of migration on housing permits and construction employment, as

¹⁴ On the left-hand side is a Keynesian multiplier between zero and one, which I do not focus on.

well as house prices, mortgages, and non-tradable goods employment.

3 Empirics I: The Expansionary Effect of Migration

Under the traditional view, migration causes an increase in the unemployment rate. In this section, I construct an exogenous shock to test this assumption, using previous migratory patterns and outflows from other MSAs to study the effects of immigration. I find that the unemployment rate falls, the opposite of the traditional view.

3.1 Data

I use two main sources of data. The first is the Internal Revenue Service's Statistics of Income U.S. Population Migration Data. The sample covers the entire United States from 1990-2014, and records migration flows from county to county on a yearly basis. The data records the number of returns filed, as well as any exemptions they claim, proxying for the total number of people in the household. It also includes the adjusted gross income of the migrants. This is a uniquely useful dataset because it allows me to create a network of migration, which I use to construct the shock.¹⁵ Datasets from the ACS only record the state from which someone moved, and matched Census data is not at a high enough frequency to capture the effects I am interested in.

¹⁵There are a few drawbacks to this data. First of all, the address used to determine migration is the address from which the tax return is filed, meaning that the date of migration could be anytime before filing taxes. While much of the migration likely occurred in the previous calendar year, some will have occurred in the first few months of the next year. In 2015, 132 million returns were filed by May 28, out of 148 million filed by November 24, over 85 percent. Marlay and Mateyka (2011) report large seasonality of moves, with summer being the most common season to move during, even more so for people that cross state or county lines. Furthermore, the timing of filing taxes might be endogenous to moving, so the ratio of immigration to non-migrants might be slightly mismeasured. Finally, the sample before 2011 does not include any people who filed after September. These people tend to be richer and have more complex taxes, and they are being missed from the data. So potentially, migrants are undercounted compared to non-migrants, and it might especially be true for rich migrants. Another potential issue is that the data is censored below, and only records data if there are more than 10 returns. Finally, the data covers only people who file taxes and their dependents. The elderly and the jobless are certainly undercounted. Despite all these drawbacks, the data is still very useful in determining patterns of migration, and any of these measurement errors are likely to be small compared to other available datasets.

In general, a county is smaller than a labor or housing market. In estimating the effect of migration, I aggregate to metropolitan statistical areas (MSAs) using the Missouri Census Data Center aggregation tables. I will note when I also use micropolitan statistical areas, which together with MSAs, are referred to as core-based statistical areas (CBSAs). A metropolitan statistical area is a collection of counties with an urban area of at least 50,000 people, while a micropolitan area only requires an urban area of 10,000. I choose MSAs instead of commuting zones because certain housing data is more readily available this way, especially Saiz elasticities. My dataset consists of 381 MSAs and 917 CBSAs.

The second data set comes from the Bureau of Labor Statistics's Local Area Unemployment Statistics (LAUS). I use annual unemployment rates. The LAUS uses a variety of sources to calculate local area unemployment rates, including the Current Population Survey, the Quarterly Census of Employment and Wages, and unemployment insurance claims.¹⁶

For robustness and to explore the housing channel, I also use data from a variety of other sources. Wage and industry employment data comes from the Quarterly Census of Employment and Wages. I sort these into categories based on the decomposition of Mian et al. (2013). House price data comes from the Federal Housing Finance Agency. Housing starts comes from the Census Building Permits Survey. Mortgage data comes from the Home Mortgage Disclosure Act data. Population data comes from the Census. Wage data comes from the BLS Occupational Employment Statistics. Estimates of housing elasticity come from Saiz (2010). The location of counties is taken from the Census Censtats database.

I report means and standard deviations of key variables in Table 1. Each variable is available for 381 MSAs, and $N < 7620$ reflects that variable is not available in all years.

¹⁶Some of the estimate is imputed from demographically-adjusted state-wide estimates, which could imply misleading within-state correlations. However, the primary building blocks are establishment employment counts and unemployment insurance claims, which are area-specific. Adjustments for commuting are made, so I focus on MSAs when using this data because MSAs are constructed to cover popular commuting patterns.

Table 1: Summary statistics

Variable	Mean	Std. Dev.	N
Unemployment Rate (Percent)	6.1	2.8	7620
Employment (1000s)	306.1	712.5	7620
Population (1000s)	643.0	1505.3	7620
Immigration Rate (Percent)	3.3	1.6	7620
Outmigration Rate (Percent)	3.2	1.4	7620
House Permits Issued per 1000 people	6.4	11.2	5714
House Price Growth (Percent)	2.9	5.9	7592
All Mortgage Originations per Capita (\$s)	4.5	3.9	7620
Second Lien Mortgage Originations per Capita (\$s)	0.2	0.4	3810
Non-tradable Employment to Population Ratio (Percent)	7.8	1.7	6096
Construction Employment to Population Ratio (Percent)	3.3	1.5	6096
Tradable Employment to Population Ratio (Percent)	4.9	3.2	6096
Average Weekly Earnings per worker (\$s)	636.4	158.1	7449

3.2 Identifying Immigration Shocks

The goal of this section is to estimate the effect of immigration to an MSA on the MSA's unemployment rate. Isolating the causal relationship requires plausibly exogenous shocks to immigration because immigration and unemployment are likely to be correlated for other reasons.

One concern with using the migration rate itself is reverse causality: people choose to migrate to areas with lower unemployment. This would bias the OLS regression downward, because it induces a negative correlation between immigration and unemployment. My other major concern is omitted variable bias: during my sample, an increase in housing prices lowered unemployment (Mian et al., 2013). If it also affected the immigration rate, there would be omitted variable bias. This would bias the OLS upward, because it induces a positive correlation between immigration and unemployment. These two concerns are not meant to be an exhaustive list, but are likely to be major sources of bias. In this section, I identify shocks to immigration as a strategy to address these concerns.

I use the historical patterns of migration and the outmigration from far-away counties, to construct a shock to immigration to an MSA. I only use the outmigration that goes to places

far from the MSA as well, meaning that the shock is not directly related to the economic conditions of the MSA of interest. This is similar to the strategy used by Altonji and Card (1991), but tailored to suit the domestic migration setting.¹⁷

Specifically, I use the first four years of the IRS data, covering movements from 1990-1994 to map the network of migration around the United States. Then, to construct the predicted immigration for a particular MSA in a particular year, for each county more than 100 miles away, I take the share of people moving into that MSA in the historical network, and multiply by the outmigration of the origin county in that year to places more than 100 miles from the MSA. Then I sum over all counties.

Because the patterns of migration are relatively stable, this measure is strongly correlated to the actual immigration of that MSA. There are many possible explanations over why the patterns are stable, perhaps because of ethnic similarities or family ties (Bartel, 1989). Other determinants, such as distance or the similarity of climate, are quite stable over time as well.

In my baseline construction of predicted migration, I throw out all flows that are to or from a county within 100 miles of the MSA. I check the robustness and exogeneity of this cutoff by using all counties outside the MSA,¹⁸ and by using a cutoff of 500 miles.

As a concrete example, suppose I am constructing a prediction for immigration to the Boston-Cambridge-Newton Metropolitan Statistical Area. To start, I would pick a county, say Montgomery County, Maryland. From 1990-1994, 1.0 percent of the outmigrants from Montgomery County move to the Boston MSA, and 98.6 percent move at least 100 miles away from Boston. In 2007, 42,032 people moved from Montgomery County to other places 100

¹⁷Beaudry, Green, and Sand (2014) use a similar idea to construct instruments based on the migration preferences of specific demographic groups within the United States. Their identification is based on changes in migration across demographic groups rather than counties of origin. They look for longer-term effects using Census data. In contrast, my instrument allows me to use high-frequency variation in order to capture short-term effects.

Shimer (2001) and Foote (2007) also use demographics as an instrument for increases in population. A major difference of my instrument is that demographic trends may be more predictable than changes in outmigration of historically-connected counties, leading to a more gradual change in housing stock that mutes both of the channels I discuss in this paper.

¹⁸For this measure, I also throw out any counties for which more than half of their outmigrants move to the MSA. Including those counties gives noisy and small outmigration shocks large influence over the shock's variation and makes the estimates much less precise.

miles or more away from Boston. To calculate the predicted immigration, I would multiply those 42,032 by 1 percent and divide by 98.6 percent to predict that 426.3 people moved to Boston. I would then sum over all counties in America that are at least 100 miles away from Boston, which would give me a prediction for immigration to Boston in 2007.

In math, the formula for predicted immigration is:

$$\tilde{z}_{n,t} = \sum_{c \in -n} \frac{m_{c \rightarrow n, t_0}}{m_{c \rightarrow -n, t_0}} m_{c \rightarrow -n, t}$$

where $-n$ is the set of all counties that are sufficiently far from n , t_0 is the pre-period, and $m_{c \rightarrow n}$ is the migration from county c to area n . I normalize this measure by the city's population.

The predicted immigration to a county is autocorrelated. So I take a final step to isolate the innovation in the predicted immigration. I assume \tilde{z} follows an AR(2) process.¹⁹

$$\tilde{z}_{n,t} = \beta_1 \tilde{z}_{n,t-1} + \beta_2 \tilde{z}_{n,t-2} + \alpha_n + \alpha_t + z_{n,t}$$

where $z_{n,t}$ is the local innovation to the shock. I include a time fixed effect because there are trends that occur at the national level. I use the Arellano and Bond (1991) estimator to estimate β_1 and β_2 . From this regression, I recover the $z_{n,t}$'s. The final step is to normalize $z_{n,t}$, such that it predicts an increase in immigration equal to one percent of the city's population. This last step is helpful purely for interpretation.

The identifying assumption behind these results is that the outmigration from historically-connected counties is unrelated to other factors that might cause a change in the unemployment rate. In my data, the most pronounced outmigration episodes are because of two hurricanes: Katrina and Irene. In Appendix D, I show similar results to my main regression using only the outmigration from Hurricane Katrina.²⁰ One concern for identification is that

¹⁹This is robust to the choice of lags. Coefficients on lags greater than 2 were insignificant.

²⁰It is easier to study the effects of Katrina because it primarily hit eight counties, whereas the effects of Irene were more widespread. The exercise is quite similar to McIntosh (2008), which finds negative

areas with high mobility between them might experience similar shocks. However, if the shocks go in the same direction, and if positive shocks induce people to stay, then the bias from this story will attenuate my results, suggesting my results might be a lower bound for how expansionary immigration is. I discuss this bias in more detail in Section 3.5.

3.3 Econometric Specification

My specification is

$$\Delta u_{n,t} = \sum_{s=-3}^6 \beta_s \Delta z_{n,t-s} + \alpha_t + \epsilon_{n,t}$$

where $u_{n,t}$ is the unemployment rate in MSA n in period t . The methodology used to construct shocks used up the first six years of migration data, so the sample period is from 1997-2013.²¹

I estimate a moving-average model in order to trace out the impulse response of the migration shock.²² The response of u_{t+s} to the shock in period t is simply β_s . In addition to lags, I include leads of the shock as a placebo test to make sure migration is not “causing” changes in the unemployment rate before it occurs.²³ I chose to use six lags because the effect dissipates after six years, implying that additional years are unlikely to be an important omitted variable. Three lags is appropriate to show a lack of a trend.

I include a year-fixed effect to control for aggregate economic conditions, since it is

effects on wages and employment in Houston in the first-year after Katrina. Indeed, I also find a rise in the unemployment rate in the first year, but a large decline afterward.

²¹One might be concerned this is a special time in U.S. history, especially since the housing boom and bust plays a prominent role throughout most of the time period. However, in Section 4, I argue that the relationship we see between housing construction and house prices match well with previous estimates from before the housing bubble (Poterba, 1984; Topel and Rosen, 1988). Results are robust to splitting the sample to before and after 2007.

²²See Hansen and Sargent (1981) or Plagborg-Møller (2015) for a discussion of moving average models. The following papers have also used a similar econometric set-up, often to answer more aggregate questions: Ramey (2016), Jordà (2005), Angrist, Jordà, and Kuersteiner (2013).

²³A violation of parallel trends is not necessarily a violation of exogeneity, as there may be anticipatory effects from the immigration. Nonetheless, an absence of parallel trends lends credence to the identifying assumption.

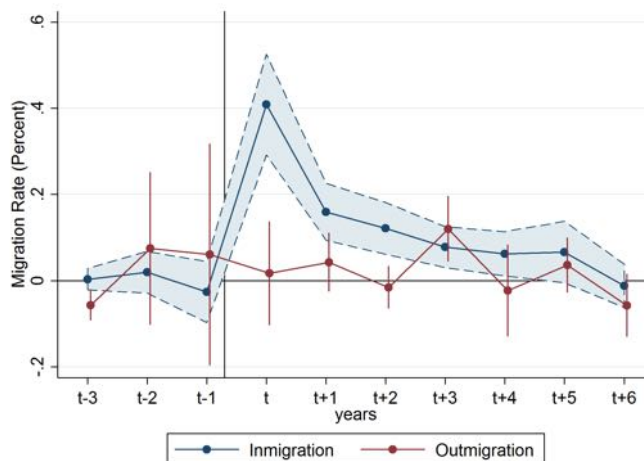


Figure 3: The migration response to the immigration shock, with 95 percent confidence intervals. Standard errors clustered by state.

well-known that gross migration is correlated to economic conditions (See Molloy, Smith, and Wozniak, 2014). I also estimate the equation in first-differences due to concerns that migration or the unemployment rate may be non-stationary. As I show in robustness checks, estimating the equation using MSA fixed effects leads to almost exactly the same results.

Figure 3 shows the response of immigration and outmigration to this shock. For this figure, I run the same specification, but with the migration rate on the left-hand side. I normalize z such that the total immigration response over the six year period is equal to 1 percent of the original city's population. Hence, for future impulse responses, it can be interpreted as the response to a shock which will increase immigration by 1 percent over the course of six years. Note there is not a response from outmigration initially. Whatever the cause of the outmigration from historically-connected counties, it is not common to the receiving city.

3.4 The Effect on Unemployment

Baseline Results

Figure 4 shows the effect of an immigration shock on the unemployment rate. The blue line, with dashed confidence interval bands, is the estimated effect of the one-percent immigration shock. In periods $t - 3$ to $t - 1$, the coefficients are not significantly different from zero, giving no evidence of a pre-trend. In period t , the period of the shock, the unemployment rate falls by 0.1 percentage points. In period $t + 1$, the unemployment rate falls more, to a total effect of 0.3 percentage points, which stays roughly constant through $t + 2$ and $t + 3$, before gradually returning to zero by $t + 6$.

The red line is produced in a similar way, but uses the residuals of actual immigration, rather than the predicted immigration outlined in the previous section. This is a similar exercise to comparing ordinary least squares and instrumental variables regressions. While qualitatively similar, the magnitude of the results is larger using the exogenous variation. As mentioned previously, this could be because the housing bubble during this time period played a major role in both limiting immigration and boosting the economy. During this time period, high house prices were associated with both higher immigration and lower unemployment rates.²⁴

The result is the opposite of the assumption inherent to the traditional view. If, as I have shown here, migration has a negative effect on the unemployment rate, then migration does not dampen local shocks. Rather, as people move to prosperous areas, it will amplify those shocks.

Robustness to Specification

One set of concerns over this regression is that the choice of specification is important. In Figure 5, I show that this is not the case. In my baseline specification, I used first-

²⁴If the sample is restricted to areas with high housing elasticity, the OLS estimate is larger than the IV estimate, suggesting that the housing bubble is biasing the regression in this direction.

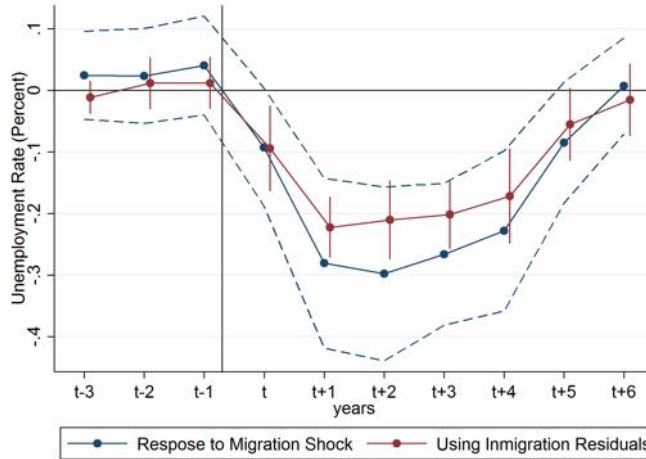


Figure 4: The effect of an immigration shock causing an increase of one percent of the MSA’s population, with 95 percent confidence intervals. Errors clustered by state. Number of MSAs: 381.

differences in case one of the variables was non-stationary. However, using fixed effects does not meaningfully change the results. I also run the regression in first-differences with a fixed effect, effectively allowing for a linear city-specific trend, to show that the result is not driven by differential trends in both migration and unemployment. Lastly, I also include two additional lags in migration, to show that my results are not sensitive to their inclusion.

Robustness to City Characteristics

Another set of concerns over this regression is because the constructed shock might be correlated to other city-time-specific characteristics. For example, one might be concerned that outmigration might be driven by the performance of specific industries, which are also present in the receiving city. Another concern might be that some national shocks could change both migration and unemployment differently in higher-educated cities.

Figure 6 presents the robustness of the result to these concerns. I flexibly control for industry and education. To do this, I interact the shares of 2-digit SIC industries in 1990 with a year dummy to control for industry. For education, I interact shares of the 11 education codes in the 1990 Census with year dummies. I also present the results using an

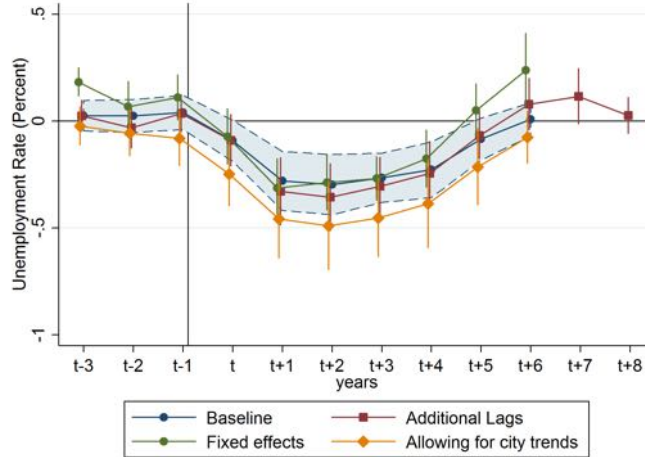


Figure 5: The effect of an immigration shock equal to one percent of the MSA’s population, with 95 percent confidence interval. Errors clustered by state. Number of MSAs: 381.

alternative shock. To construct this shock, I exclude using migration to or from cities that are similar to the receiving city, based on the correlation of Bartik shocks. The result is robust to all of these alternate specifications.

Robustness to Spatial Correlations

In Figure 7, I investigate a variety of robustness checks aimed at addressing concerns about the spatial structure of my regression. The concern here is that there may be an omitted variable that affects areas near the MSA, causing people to move, but which also directly affects the unemployment rate.

Because my shock relies on the cutoff of 100 miles, I investigate the robustness to that by using no cutoff, and using a cutoff of 500 miles. The results are similar, though for no cutoff, the coefficient in period t is much smaller. This could be evidence of bias if a cut-off is not used because the area around the MSA could be doing poorly, causing people to move out. Reassuringly, the estimate for 500 miles is almost right on top of the baseline specification. The last robustness check is that I control for the Census Division each MSA is in, interacting each division with year fixed effects. These controls do not affect the results.

In addition to these robustness checks, in Appendix D, I present a similar regression

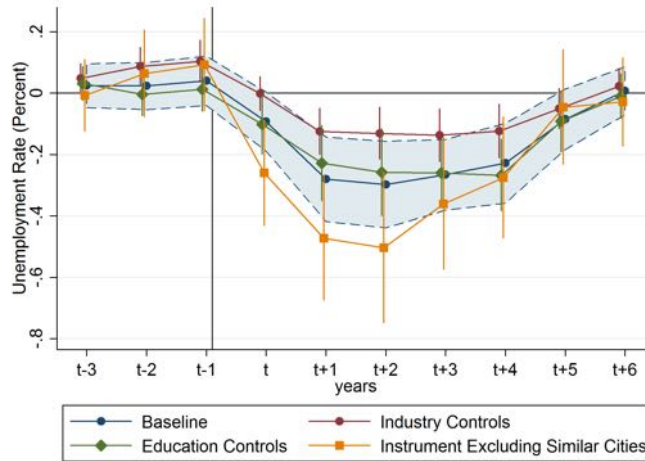


Figure 6: The effect of an immigration shock equal to one percent of the MSA’s population, with 95 percent confidence interval. Errors clustered by state. Number of MSAs: 381.

using a migration network based on the 1940 Census instead of my pre-period, finding similar results, but noisier because the Census does not use the county of origin, only the state.

Robustness to Alternate Measures of the Labor Market

A final robustness check is to make sure we can see this force in other measures of the labor market. In Figure 8, I show the effect on the employment-population ratio, for both MSAs and the broader category of CBSAs. Here I construct the employment-population ratio by dividing employment in the Quarterly Census of Employment and Wages by the population estimates of the U.S. Census. Estimates are consistent with the effects on the unemployment rate, with the employment-population ratio rising.

In Figure 9, I also show that earnings went up and that unemployment benefits fell. Earnings also come from the QCEW, as measured by the annual average weekly wage, while unemployment benefits come from the BEA’s personal current transfer receipts. I allow for a city trend when estimating earnings.

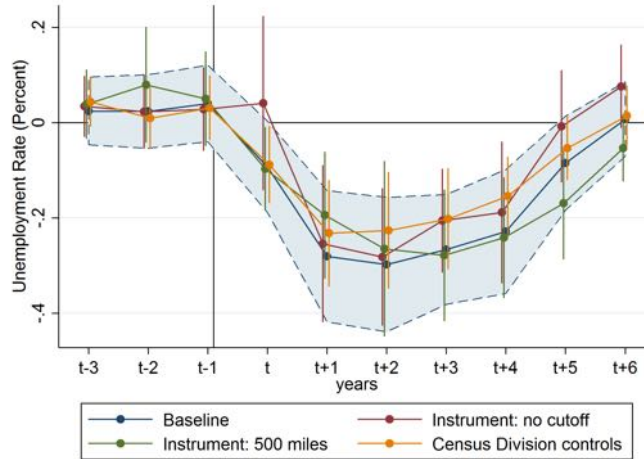


Figure 7: The effect of an immigration shock equal to one percent of the MSA’s population, with 95 percent confidence intervals. Errors clustered by state. Number of MSAs: 381.

Effect on non-migrants

A natural question is whether the effects on the unemployment rate could be driven purely by migrants being more likely to have jobs.²⁵ Ideally, I could distinguish individuals by where they lived previously, but that would require a large panel dataset that tracked both location and employment status. Nonetheless, by focusing on the unemployment rate, there is a natural bound on the direct effect coming from the employment status of migrants.

If I assume each immigrant has a job and each outmigrant is unemployed I can calculate a bound. In $t + 1$, the shock has added about 0.63 percent of the population in immigrants, and about 0.09 percent of the population has left. This implies a bound of 0.63 times the original unemployment rate plus 0.09 times one minus the unemployment rate, assuming the labor force participation rate is roughly similar. The average unemployment rate in my sample is 6 percent, implying that the upper bound on this mechanical effect is 0.11 percent. In that time period, I estimate the unemployment rate has fallen by 0.32 percent, implying most of the effect must have been because of additional jobs for non-migrants, even under these extreme assumptions.

²⁵It appears from the descriptive data that their adjusted gross incomes are not necessarily that much higher. See for example, Figure 24 in Appendix B.

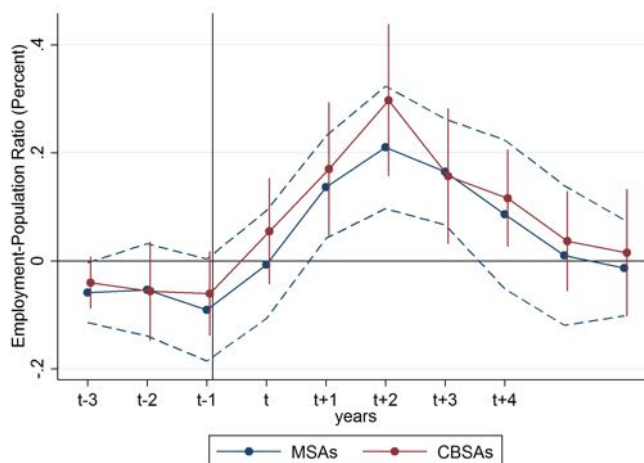


Figure 8: The effect of an immigration shock equal to one percent of the MSA’s population, with 95 percent confidence interval. Errors clustered by state. Number of MSAs: 381. Number of CBSAs: 917.

The fall in the level of unemployment benefits (Figure 9) is also indicative that the effect is not driven purely by immigrants being more likely to have jobs, as that would not change the level of benefits.

3.5 Threats to Identification

In this section, I consider a few threats to the exclusion restriction, and whether they could be driving the results.

Regional Shocks

One threat to identification is economic shocks that affect multiple cities at once. Such a correlation would violate the exclusion restriction if it changed the outmigration rate in one city and the unemployment rate in the other.²⁶ For example, an oil boom would make Dallas more prosperous, and fewer people would migrate out. At the same time, it would lower the unemployment rate in Houston. This would bias my regression.

²⁶It is because of such concerns that I exclude migration within 100 miles of the MSA, and I check for the robustness of the regression with industry and education controls. However, a skeptical reader might not believe these to be sufficient.

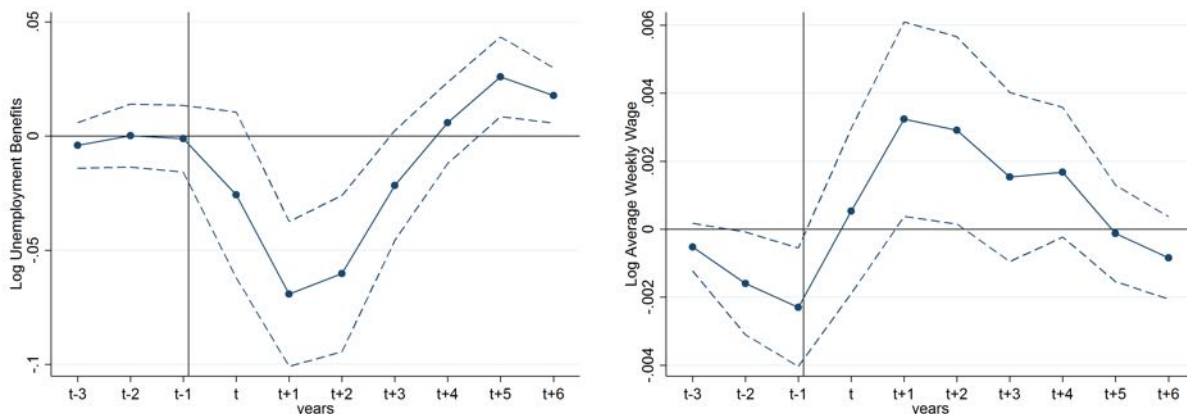


Figure 9: The effect of an immigration shock equal to one percent of the MSA’s population on unemployment benefits (left) and earnings (right), with 95 percent confidence intervals. Errors clustered at MSA level. Number of MSAs: 381.

I expect this bias to be small and positive, suggesting that my results are an upper-bound on immigration’s true effect on the unemployment rate. In Appendix C, I show that city pairs with high migration between them tend to be similar in location and industrial composition. In Section 5, I show that negative local shocks lead to an increase in outmigration. Taking these two facts together, the instrument, the outmigration rate of one city, and the outcome, the unemployment rate of a connected city, are likely to be positively correlated if an economic shock hits both cities. Hence, they would bias my regression upward.

Substitution between Cities

Another threat to identification is that a low unemployment rate in one city would lower migration between other cities. For example, a boom in Boston might cause someone leaving Montgomery County to decide to move there instead of New York. Because I am using migration from Montgomery County to New York to construct my instrument, this would bias my regression.

I again expect this bias to be small and positive. A boom in one city will cause the instrument to be slightly smaller, causing a positive correlation between the unemployment rate and the instrument.

Terms of Trade Effects

One might be concerned that two cities with high migration might compete against each other in the same industries. For example, many people move between Boston and San Francisco, both of which produce pharmaceuticals. If Boston pharmaceuticals were struggling, that might lead to higher outmigration from Boston, and a higher price of pharmaceuticals. The change in price would benefit San Francisco, and could cause the unemployment rate to fall.

However, in Section 4.5, I show that the decline in the unemployment rate does not come from employment in industries that produced tradable goods. Rather, the benefits are concentrated in construction and non-tradable goods and services. So this bias does not seem to be driving the decline in the unemployment rate.

Spatial Equilibrium

More generally, one might be concerned that any economic force that drives migration is also likely to affect unemployment, and that the intuition behind the instrument might not carry into general equilibrium. To address this concern, this section presents a simple static spatial model where all migration is driven by the unemployment rate, and shows that the estimates from the instrumental variables regression are nonetheless close to the true parameter of interest.

In this model, there are 381 cities, and migration from city i satisfies a gravity equation, with the elasticity to the unemployment rate of σ :

$$m_{ij,t} = \frac{(\tau_{ij}u_{j,t})^\sigma}{\sum_k (\tau_{ik}u_{k,t})^\sigma}$$

where τ_{ij} is a shifter in city-to-city migration. The unemployment rate is a function of a local fixed effect, migration, and i.i.d. local labor demand shocks.

$$u_i = \alpha_i + \beta \sum_j m_{ji,t} + \epsilon_{i,t}$$

These two equations describe the equilibrium in any period t . In the following simulation, I assume values for σ and β of -0.4 and -0.3 , respectively. I use the pre-period migration flows to calibrate the τ_{ij} , allowing some of them to be infinity. Next, I linearize around the pre-period equilibrium, and generate migration and unemployment data by drawing ϵ 's from a normal distribution. I use the migration data to construct a similar instrument to the one I use in this paper.²⁷ I then run regressions of the change in unemployment on the change in migration using ordinary least squares, as well as two-stage least squares with the instrument.

I simulate this 10,000 times, and present the results in Figure 10. As you can see, the OLS is biased, but the IV estimates are close to the mean on average.²⁸ In these regressions, the largest bias in OLS comes from reverse causality, which is a negative bias. The simulations confirm this. I conclude that my instrumental variables strategy is valid even in general equilibrium.

4 Empirics II: The Housing Channels

In Section 2, I argued that a standard model cannot explain the decline in the unemployment rate which I found in Section 3. I also showed housing could play that role. In this section, I show the empirical evidence in favor of this hypothesis.

Recall the key equation from Section 2, which outlined two channels. One was the construction channel, which was based on a boom in new houses as the economy adjusted to higher housing demand. Of key importance was new housing, and the increase in construction employment. Second was the house price channel, which was based on an increase in house prices that induce higher consumption. I show evidence in favor of each of these channels,

²⁷Because I have 381 cities, the instrument is based on city-to-city movement, instead of using counties.

²⁸The confidence interval for the IV includes the true value 85 percent of the time, and for OLS, 39 percent of the time. The lack of coverage in the IV may be because of weak instruments in this simulation, a problem I do not have in the actual data. In these simulations, the first-stage F-test for the instrumental variables regression averages 1.4, significantly below most rules-of-thumb for weak instruments. In the data, my first-stage F-statistics are significantly stronger. One difference could be because in the data I can construct the instrument using counties instead of cities.

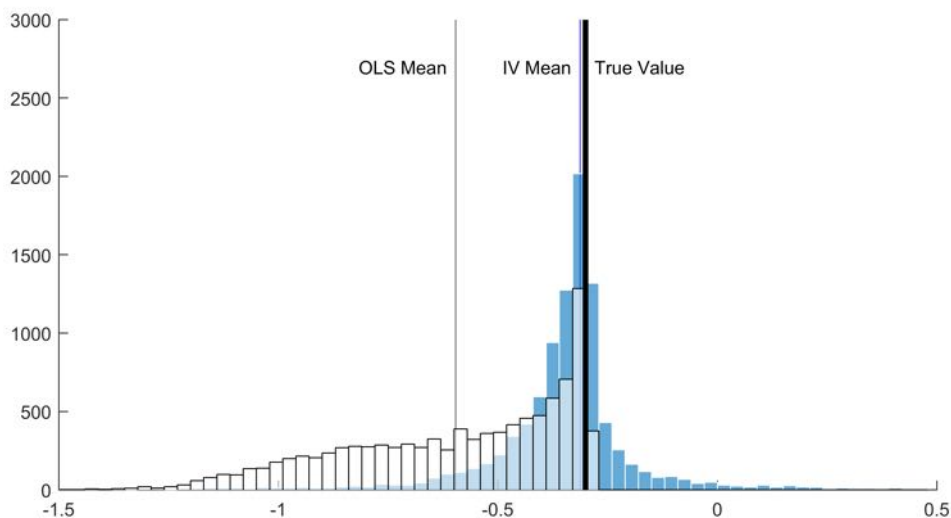


Figure 10: Distribution of regression results in 10,000 simulations of spatial model. OLS regression in white, IV in blue. Not shown: 2.2 percent of IV estimates were greater than 0.5 and 1.4 percent were less than -1.5.

and then show that the unemployment rate's response to migration is dependent on housing price elasticities.

Throughout this section, I use the same specification as in Section 3.

$$\Delta x_{n,t} = \sum_{s=-3}^6 \beta_s \Delta z_{n,t-s} + \alpha_t + \epsilon_{n,t}$$

where x is house permits, construction employment, house prices, mortgages, or non-tradable employment. To study the effect on employment composition, I use the employment categories from Mian et al. (2013). Their decomposition assigns NAICS 4-digit categories to one of four sectors: construction, non-tradable, tradable, and other. They make up respectively 9 percent, 19 percent, 11 percent, and 61 percent of employment in my data. For house prices, which have a strong trend component, I also include a city fixed effect, effectively controlling for a linear trend.

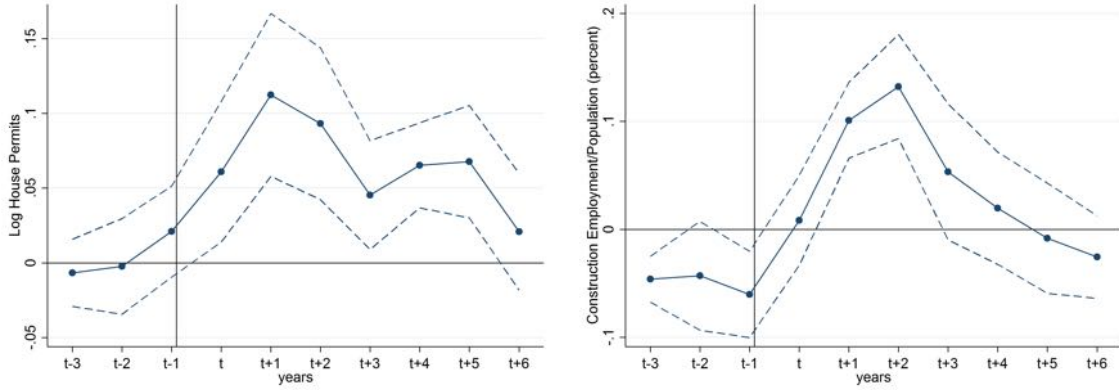


Figure 11: The effect of an immigration shock equal to one percent of an MSA’s population on housing permits issued and construction employment, with 95 percent confidence interval. Errors clustered by state. Number of MSAs: 381.

4.1 Construction Channel

The construction channel requires a build-up of new housing, especially in the short-term. In Figure 11, I show that housing permits, from the Census, increase significantly after a migration shock.²⁹

The effect is quite large, a one percent migration shock causes a 10 percent increase in the number of permits per year. Over the course of six years, the number of houses goes up by approximately 40 percent of a typical year’s permits. In my sample, the number of permits per year averaged 0.65 percent of an MSA’s population, so the cumulative effect corresponds to about one new house for every four immigrants.

In the right half of Figure 11, we see an increase in the construction sector. For one percent immigration, there is a corresponding increase in construction equal to 0.1 percent of the population for about three years. On average, construction employment is about 3.4 percent of the population, so this measure is slightly smaller than the effect found on house permits. One potential reason for this is that Mian et al. (2013) are very broad in their definition of construction, as they aim to show a null result; that construction was not a major cause of the Great Recession, and include industries as diverse as logging and real

²⁹ Ideally, I could measure the stock of housing as well, but unfortunately, the local housing stock only began to be measured by the Census in 2010.

estate agents. Hence, 3.4 percent may be an overestimate of the share of the population working in construction.

Recall from Figure 8 that the total employment-to-population ratio increased by about 0.2 percent in response to the shock, so the construction channel seems to be explaining about half of that. In Appendix D, I show the robustness of these results to many of the same checks I showed in Section 3.

4.2 House Price Channel

I now turn to the house price channel, which posits that house prices go up and cause increased non-tradable demand.

In Figure 12, I show that house prices do increase, responding by roughly five percent in response to one percent immigration. Housing prices come from the Federal Housing Finance Agency, and is based on both sales prices and appraisals. Based on the increase in housing permits, it would suggest a short-run housing supply elasticity of about 2. This is line with Poterba (1984), which estimates a housing supply elasticity of between 0.5 and 2.3; and Topel and Rosen (1988) which estimate a one-quarter-short-run elasticity of 1 and a long-run elasticity of 3 that occurs mostly within a year. Both estimate the elasticity off the time series of aggregate U.S. data.³⁰

In Figure 13, I present some evidence that this housing price increase is leading to additional consumption. On the left is the rise in mortgage lending. Not surprisingly, there is a large increase in the amount of total mortgages. But the percentage increase in second-lien mortgages is much higher. Second-lien mortgages are often taken to finance consumer spending, and as such, are good evidence that people are responding to their increased housing wealth.³¹

³⁰Finding an estimate within this range is important because it suggests the results are being driven by a change in construction, and is less likely to be special to the housing bubble.

³¹The majority of second-lien mortgages are home equity lines of credit (HELOCs). See Lee, Mayer, and Tracy (2012) for a further discussion of second liens in recent years. Of course, second liens are less than 10 percent of the mortgage market, so the increase is smaller in dollar terms.

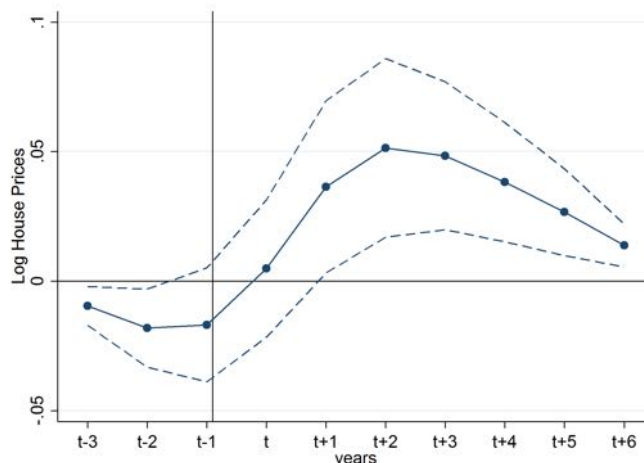


Figure 12: The effect of an immigration shock equal to one percent of an MSA’s population on house prices, with 95 percent confidence interval. Errors clustered by state. Number of MSAs: 381.

On the right is the rise in non-tradable employment, which increases by about .06 percentage points for four years. In the context of the model, non-tradable employment can change because of house prices, but also because of differences in demand between migrants and non-migrants, or a Keynesian multiplier. Given a house price rise of about five percent, and assuming a consumption-to-house-price elasticity of 0.2 (Berger et al., 2015), we would expect non-tradable consumption to rise by 1 percent. The mean non-tradable-employment-to-population ratio is 8 percent in my data, which would predict a 0.08 percentage point increase in non-tradable employment, slightly higher than what we see in Figure 13.

This increase is a bit less than half the total increase in the employment-to-population ratio. Together, the house price channel and the construction channel appear to explain most of the total labor market response.³²

4.3 Cross-Sectional Heterogeneity

Because migration affects unemployment through house prices, one hypothesis is that areas in which house prices are more responsive to demand might experience bigger changes in

³²Given that there are also non-tradable components to many parts of the “other” category, it may require a decline in the tradable-employment-to-population ratio, which I will find in Section 4.5.

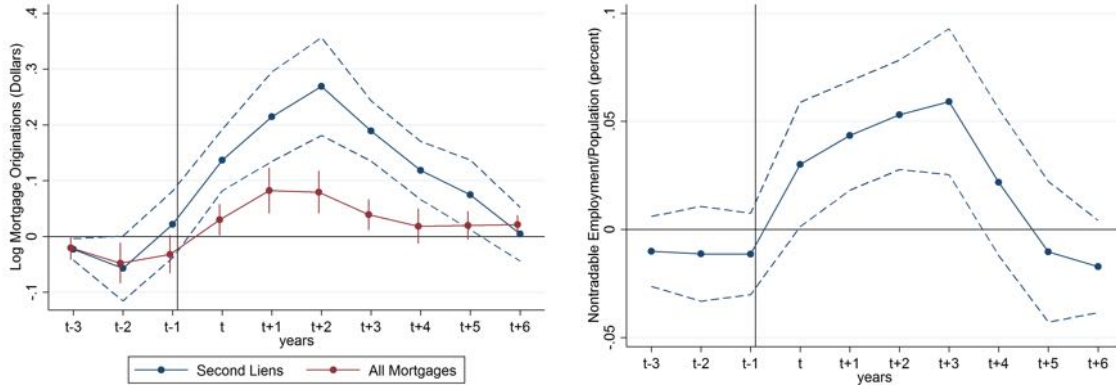


Figure 13: The effect of an immigration shock equal to one percent of an MSA’s population on mortgage originations(left) and non-tradable employment (right), with 95 percent confidence interval. Errors clustered by state. Number of MSAs: 381.

unemployment. In this section, I investigate heterogeneity across cities that differ in housing supply elasticity and population growth.

Heterogeneity by Housing Supply Elasticity

In response to an increase in housing demand, cities with lower housing supply elasticities should experience a larger increase in house prices. Potentially, this provides a channel through which the effects on unemployment vary by housing supply elasticity. I interact the immigration shock with the housing supply elasticity from Saiz (2010).³³ This allows me to see whether the effects of a migration shock are different in areas where we might expect house prices to react more.³⁴

³³Saiz (2010) uses the previous vintage of MSAs. I am able to match 253 of them to current MSAs.

³⁴In my theoretical framework, a lower housing supply elasticity undoubtedly makes the house price channel stronger. Consistent with this, I do find migration’s effect on second lien mortgages is higher in low-elasticity areas. However, for the construction channel, the effect is ambiguous. In Appendix A.4, a model with log-utility over housing would have a smaller construction channel in low-elasticity MSAs. But if I augment the one-period model from Section 2, where the intensive margin of housing demand is inelastic, with Cobb-Douglass production of housing, the construction channel would be larger in low-elasticity MSAs. In the data, I find very little difference in construction employment for high and low elasticity cities (not shown).

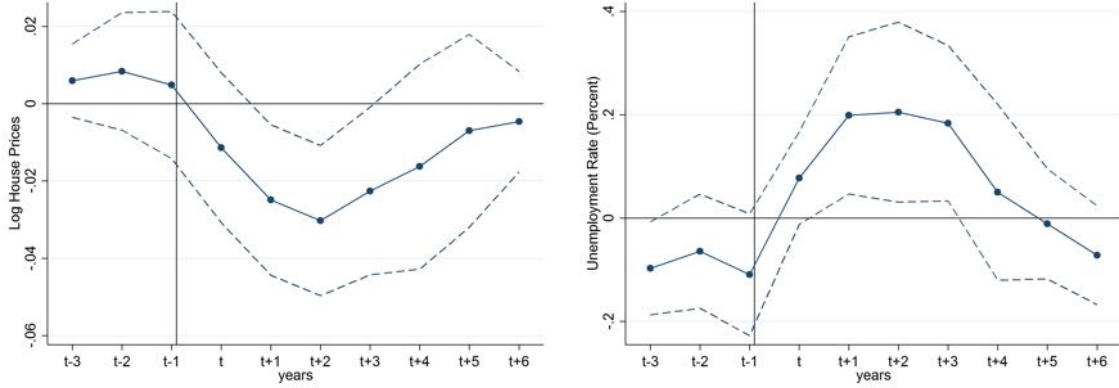


Figure 14: The effect of an immigration shock equal to one percent of the MSA’s population on house prices (left) and unemployment (right) interacted with Saiz (2010) housing supply elasticity, with 95 percent confidence intervals. Errors clustered by state. Number of MSAs: 253.

I run the following regression:

$$\Delta x_{n,t} = \sum_{s=-3}^6 \beta_s^* \Delta z_{n,t-s} \times \text{elasticity}_n + \sum_{s=-3}^6 \beta_s \Delta z_{n,t-s} + \alpha_t \times \text{elasticity}_n + \gamma_t + \epsilon_{n,t}$$

where x is house prices and the unemployment rate. β^* estimates the heterogenous effect of migration by housing supply elasticity.

Figure 14 shows that the effects do differ by housing supply elasticity. On the left, I show that house prices increase by less in more elastic areas, as expected. On the right, I show that the unemployment rate falls by less in those same areas.

Heterogeneity by Population Growth

The marginal effects of migration may differ by a city’s population growth because a declining city has excess housing stock (Glaeser and Gyourko, 2005). With many vacancies, the marginal migrant can move into an empty house with minimal impact on either construction or house prices. Hence, one prediction is that the effect of migration on unemployment ought to be larger in growing cities.

A notable pattern during this time period is that population increased more in areas of

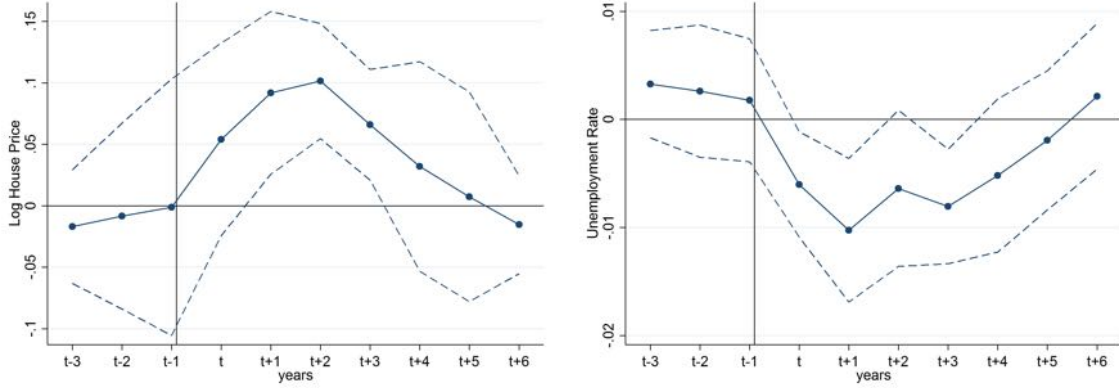


Figure 15: The effect of an immigration shock equal to one percent of the MSA’s population on house prices (left) and unemployment (right) interacted with high January temperatures, with 95 percent confidence intervals. Errors clustered by state. Number of MSAs: 381.

the country with warmer winters (Rappaport, 2007). So to test the prediction, I test if the marginal effects of migration are stronger in warmer climates. The regression is similar as for housing price elasticities:

$$\Delta x_{n,t} = \sum_{s=-3}^6 \beta_s^* \Delta z_{n,t-s} \times \text{Jan. Temp}_n + \sum_{s=-3}^6 \beta_s \Delta z_{n,t-s} + \alpha_t \times \text{Jan. Temp}_n + \gamma_t + \epsilon_{n,t}$$

where I use the average of the high of daily January temperatures from 1990 to 2013, available from the Center for Disease Control and Prevention.

The results are presented in Figure 15. In warmer MSAs, house prices increase more in response to migration and the unemployment rate falls more. These results are consistent with the fact that migration has large effects on unemployment through the housing market and that the housing market is more responsive in growing MSAs.

4.4 The Housing Channels during the Mariel Boatlift

In this section, I use a different source of variation to demonstrate the presence of these housing channels. Specifically, I use the famous example of the Mariel boatlift in Miami in 1980, where around 125,000 Cuban immigrants arrived in Miami. The goal of this exercise is not to ask whether the Mariel boatlift was stimulatory, but rather to see if there are

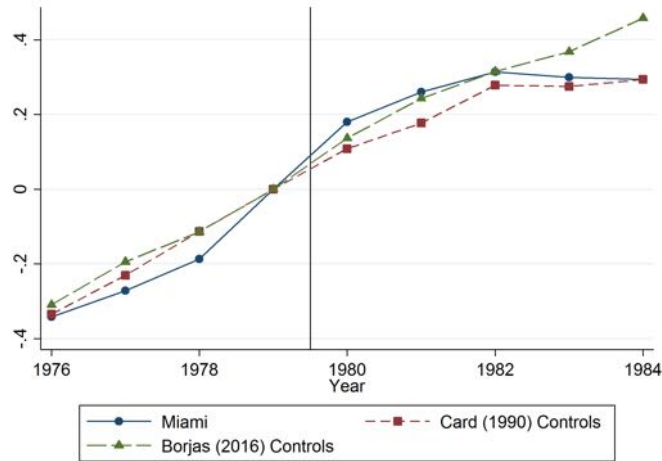


Figure 16: Log House Prices. Index, 1979=0. Miami house prices increased by more than the comparison groups or the rest of Florida in the immediate aftermath of the Mariel boatlift.

similar patterns in house prices and employment across industries. One might expect that housing demand of a Cuban immigrant to Miami would be smaller than an American's, so the magnitudes of this channel might be correspondingly smaller.

Because the boatlift was a one-time event, I will compare outcomes in Miami to those of a control group, based on similar cities in the United States at that time. There is significant debate over the appropriate group, so I will use the baseline group from Card (1990) and the baseline group from Borjas (2015). Card (1990) uses Atlanta, Houston, Los Angeles, and Tampa; while Borjas (2015) uses Anaheim, Rochester (New York), Nassau, and San Jose.

House prices for these cities are only partially in my dataset. For the Borjas control group, I am reduced to only two cities, Rochester and San Jose.³⁵ I plot the average of these two in Figure 16. For the Card control group, I am missing data from Atlanta. Miami, from 1979-1980, does have a bigger increase in house prices than either control group, by 5 percent against Card, and 2 percent against Borjas. It seems to stay higher over the course of three years.

My results are in contrast to Saiz (2003), who, when studying the same event, finds a large increase in rents, but not in house prices. The data we use is similar, though Saiz (2003)

³⁵Anaheim and Nassau have been reclassified since then and are not part of my data

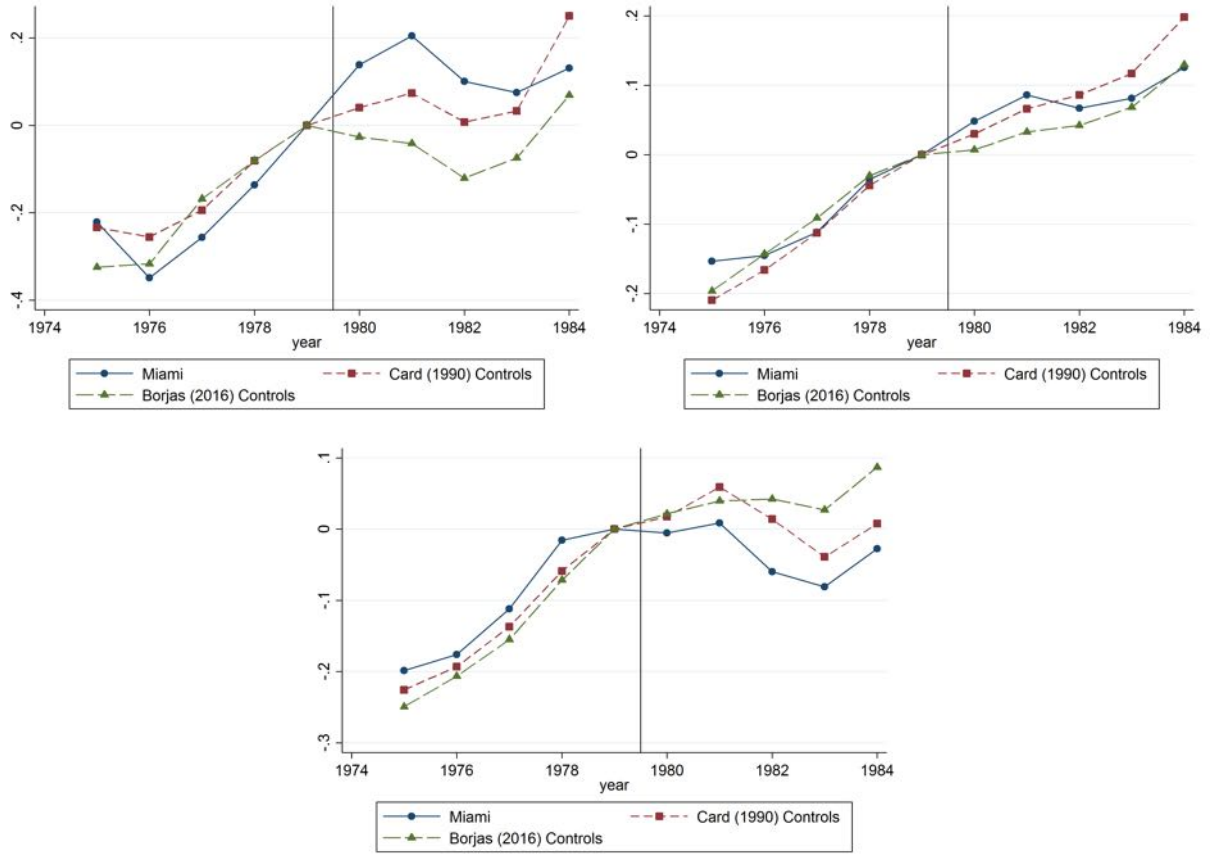


Figure 17: Log Employment. Index, 1979=0. Employment in three sectors: Construction (upper left), retail (upper right), and manufacturing (bottom).

uses yet another control group, with many cities not in my dataset. Critically, though, he indexes house price changes to 1980Q3, when many of the Marielitos had arrived, and any changes in house prices might have already occurred.

For looking at the employment effects, I can only measure industries using SIC data at coarser levels, so using the decomposition from Mian, Rao, and Sufi (2013) is infeasible. Instead I will use manufacturing as a proxy for tradables, and retail trade as a proxy for non-tradables.³⁶ Construction is measured using the SIC industry. For these groupings, all the series exist for all eight cities in the control group. I plot them in Figure 17.

All of these plots show the same qualitative patterns as they do in response to migration

³⁶Bodvarsson, Van den Berg, and Lewer (2008) use a different methodology to consider the effect of the Mariel boatlift on the retail sector, and also find a large role of labor demand.

shocks in the main body of my paper. In both construction and retail, there seems to be a temporary increase in the employment in that sector. In manufacturing, however, Miami's employment stays quite flat while the comparison cities are growing.

Saiz (2003) plots new housing permits for Miami and several comparison groups over the same decade. In 1980, new permits in Miami increased from 1979, while in each of his comparison groups, they decreased. He does not stress this, likely because there are quite significant differences in other time periods as well, but it is consistent with my finding that construction employment expanded.

In sum, the data are consistent with the two housing channels: the house price channel because we see an increase in house prices and retail employment; and the construction channel because we see an increase in permits and construction employment.

I want to conclude this section with a note about magnitudes. The Mariel boatlift was an approximately 7 percent expansion in the labor force (Card, 1990). This is significantly bigger than any shocks in my data, but the effects of these housing channels are nonetheless modest. When compared to the Borjas (2015) controls, the total jobs added in the construction sector and the retail sector are no more than 20,000, less than half of the 45,000 person increase in the labor force. Of course, retail is not an exhaustive list of non-tradable employment, so there could be bigger effects, but it is unlikely that this could drive a large decrease in the unemployment rate. Most likely, this is because there is less housing demand from each Mariel immigrant than from a domestic migrant in the United States.

4.5 Other Channels

Besides housing, there are other theories about why immigration might lead to a decline in the unemployment rate. In this section, I discuss three: factor complementarity, agglomeration, and wealth heterogeneity. Previously, in Section 3.4, I discussed how a selection story, where migrants have a lower unemployment rate than non-migrants, could lead to a small effect, but could not explain the bulk of the finding.

Factor Complementarity

One possibility is that migrants and non-migrants are complements, so that having more migrants lowers the unemployment rate of non-migrants. For example, if most migrants are high-skilled workers and most non-migrants are low-skilled, the incoming migrants could improve the productivity of non-migrants, making it easier for them to find jobs. Indeed, Molloy, Smith, and Wozniak (2011) show that migrants are more educated than non-migrants. Furthermore, I show in Appendix B that the 10th percentile of wages increases, while there is no significant effect on the 25th, 50th, 75th, or 90th percentile. All of this evidence is consistent with a complementarity story. However, a prediction of the complementarity story is that areas with fewer college educated workers would benefit more from immigration, but I find the opposite in the data (see Appendix B).³⁷ Second, Ottaviano and Peri (2012) considers the effects of complementarity in the short and long run, and finds immigration to be more beneficial in the long-run because capital has had time to adjust; the timing of my results is the opposite. And the 10th percentile of wages could also be driven by the fact that many of those workers work in the food preparation and serving or sales, which would be consistent with a housing story. Appendix B includes the figures discussed in this paragraph and a lengthier discussion.

Agglomeration

A second possibility is that there are agglomeration effects. As migrants move in, knowledge spillovers (Audretsch and Feldman, 2004), the home market effect (Krugman, 1980), thick market externalities (Diamond, 1982), or other increasing returns forces increase employment. However, one might expect tradable goods to be most affected by agglomeration, or at least equally affected. Rather, I find that tradable goods decline, as seen in Figure 18. Second, many agglomeration stories would also imply a larger effect in the long-run than the

³⁷This assumes that there are not proportionately more college migrants being pushed to places with high college shares already, which I do not have data on.

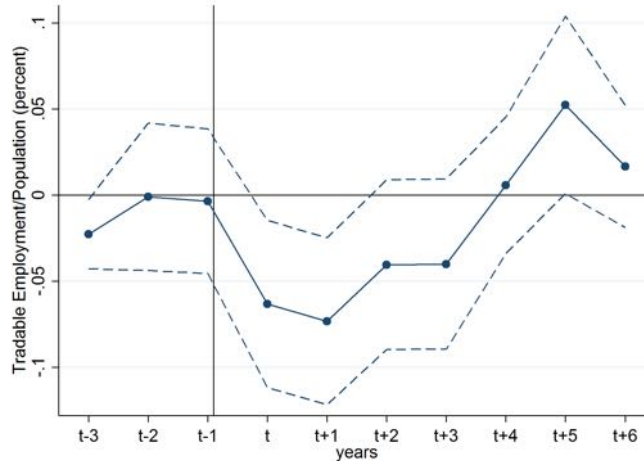


Figure 18: The effect of an immigration shock equal to one percent of an MSA’s population on tradable goods employment (right), with 95 percent confidence interval. Errors clustered by state. Number of MSAs: 381.

short-run, and I find the opposite. Finally, the magnitude of my effects is much larger than these forces can explain. Across MSAs in my sample, the biggest cities have, on average, about 0.3 percentage points lower unemployment than the smallest MSAs, despite having populations about 50 times as large. Even if this were all agglomeration effects, I am finding a decline in the unemployment rate of about 0.3 percentage points in response to a one percent increase in population, which is at least two orders of magnitude larger.

A second possible agglomeration channel is thick-market effects, where an increase in unemployment levels might increase the job finding rates, lowering the unemployment rate. However, the magnitudes of my estimates imply a decrease in the level of unemployment.

Wealth Heterogeneity

A third possibility is that migrants are significantly wealthier than non-migrants. For example, if a retiree moves to Florida, and spends down his savings, he is adding to labor demand but not labor supply. In fact, this force is possible in the model in Appendix A. To check whether this is an empirically plausible channel, I use the American Community Survey data, and look at the dividend, interest, and rental income of interstate migrants

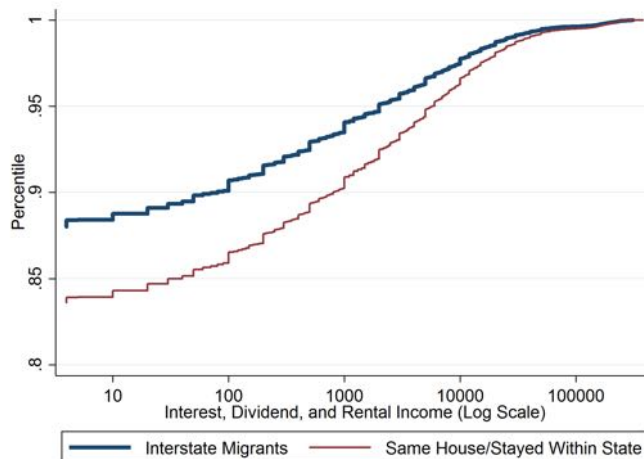


Figure 19: The cumulative distribution function of the distribution of interest income, by migration status. American Community Survey, 2000-2014.

versus non-migrants and within-state migrants. The cumulative distribution function of this income is plotted in Figure 19. The first thing to note from this plot is that more than 80 percent of people, migrants and non-migrants, do not report interest income.³⁸ More crucially, the distribution of interest income for non-migrants first-order stochastically dominates the interest income for migrants. Although the ACS does not measure wealth directly, this is suggestive evidence that non-migrants are wealthier than migrants.

This is consistent with statistics from Molloy et al. (2011), which shows the interstate migration rate is highest for ages 18-24 (4.2 percent), second highest for ages 25-44 (3.0 percent), and lowest for ages 65+ (0.9 percent). Renters also have a higher migration rate (4.7 percent) than homeowners (1.3 percent). These demographics would suggest that migrants are unlikely to be richer than non-migrants.

In conclusion, none of these other channels, factor complementarity, agglomeration, or wealth heterogeneity, seems likely to be driving the results that I find in the data. But more importantly, none of these channels make the prediction that non-tradables and construction employment would increase while tradable goods employment would fall. Nor would any of these channels have predictions on whether the effect would be stronger in high housing-

³⁸A very small fraction report negative interest income.

supply elasticity areas. Hence, the evidence that I have shown is strongly supportive of housing causing the decrease in the unemployment rate.

5 Counterfactual: The Migration Accelerator

One implication of the result from Section 3 is that there exists a “migration accelerator,” an amplification of local labor demand shocks due to migration. When an MSA experiences an increase in labor demand, people move there. Because that migration is expansionary, labor demand increases by even more.

In this section, I quantify how important immigration is in amplifying these shocks. In Figure 20, I return to the framework from Section 2. When the New curve shifts in response to a labor demand shock, both migration and unemployment are affected. If the Blanchard-Katz curve were vertical, the unemployment rate would fall by the amount labeled “demand shock.” But because migration responds in equilibrium, the total effect is larger. This part, I label the “accelerator.”

To estimate this, I first estimate how much migration responds to increases in labor demand, a similar exercise to Blanchard and Katz (1992). I then combine that with my estimates of the expansionary effect of migration in order to calculate the accelerator, but with two important caveats. First, my previous estimates were based on a shock that implied a specific expected path for migration, which is different than the path in response to the labor demand shock. I need to make an assumption about migration’s effect along this path. Second, I assume the effect from migrants who move in response to higher labor demand are similar to the effect of migrants who move in because of a push-factor from other cities. Later, I show that migrants induced by either of these shocks are indeed comparable on two important observable dimensions.

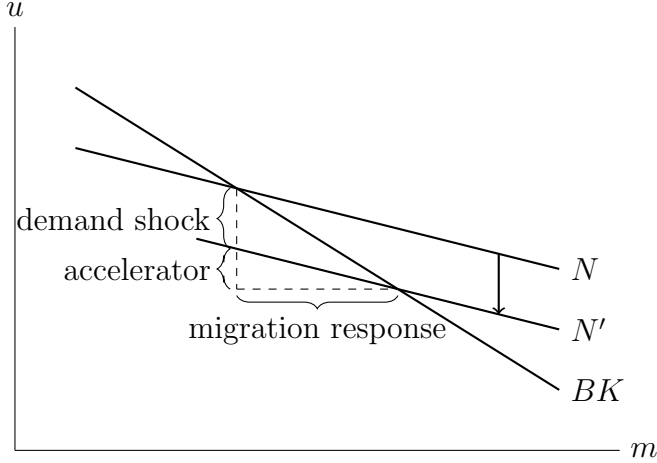


Figure 20: A Graphical Representation of the Accelerator

5.1 Migration's Response to Labor Demand

The first step in calculating the accelerator is to estimate the effect that labor demand has on migration, i.e. the BK curve. Again, there is an endogeneity problem of regressing migration on unemployment because, as I have shown in this paper, reverse causality is a major concern.

To solve this, I use a Bartik (1991)-style instrument, using the share of industries in an MSA and the growth rate of those industries in the rest of America to calculate an instrument for labor demand. I use two-digit SIC, before 1998, and three-digit NAICS codes, after 1998, to construct the instrument in each year. The formula for the instrument is

$$\tilde{z}_{n,t}^b = \sum_j s_{j,n,t-1} g_{j,-n,t}$$

where $s_{j,n,t-1}$ is the employment share of industry j in CBSA n in year $t - 1$, and $g_{j,-n,t}$ is the growth rate of employment in industry j in the rest of the U.S. besides n in year t . As I did for immigration, I then assume $\tilde{z}_{n,t}^b$ is an AR(2) process, and use the residuals, $z_{n,t}^b$, as my labor demand shock.

One endogeneity concern is that nearby CBSAs are likely experiencing similar labor demand shocks. If the economic conditions of those CBSAs are affecting the decisions of

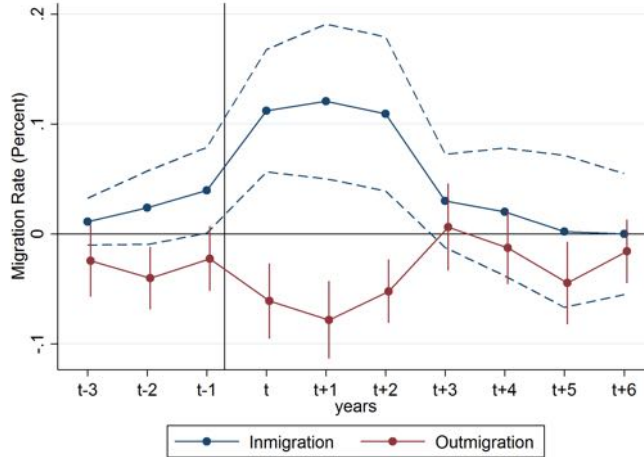


Figure 21: The effect on migration of a labor demand shock, with 95 percent confidence intervals. Standard errors clustered by state. Number of MSAs: 381.

potential migrants, it could bias the regression. To fix this, I control for the Bartik-shock in those other cities. I create this control by weighting cities based on the migration patterns from the pre-period.

I specify the regression as follows

$$\Delta m_{n,t} = \sum_{s=-3}^6 \beta_s^b \Delta z_{n,t-s}^b + \zeta_s \Delta z_{c(n),t-s}^b + \alpha_t + \epsilon_{n,t} \quad (2)$$

where $g_{n,t}$ is the growth rate of employment in CBSA n , z_n^b is the Bartik instrument in CBSA n , and $z_{c(n)}^b$ is the average of the Bartik shocks over all MSAs from which people move to n or to which people move from n . This is a very similar specification to the main regressions run in this paper.

The timing of this regression is a bit different than the timing of the regressions I ran in Section 3. Because migration is measured at the time people file their taxes, while employment is measured quarterly throughout the year, naively running this regression would show a pre-treatment effect. Therefore, I estimate this equation using the previous year's migration. I note this prominently because it matters for calculating the accelerator.

The results for both inmigration and outmigration are shown in Figure 21. The effect

on net migration can also be seen as the difference between the two lines. The effect for outmigration is smaller but significant. In contrast, the effect on immigration is relatively large, about twice the size, and lasting for three years. The cumulative effect of a one percent labor demand shock on net immigration is about 0.25 percent after three years.³⁹

5.2 Accelerator

To calculate the accelerator, I estimate the expected effect of the migration I found in Figure 21 on unemployment. I compare that to the total effect that Bartik shock has on unemployment. Put another way, I am comparing the accelerator line segment of Figure 20 to the total change in the unemployment rate.

Using only my estimates from Section 3, I cannot estimate the effect of any sequence of immigration, I can only do it for the sequence I observed in response to the immigration shock. If Figure 21 looked exactly like the path of immigration induced by this shock (Figure 3), I could directly use those estimates, but because its shape is different, I must make an assumption. The assumption I will use is that the effect of migration in year t on unemployment in $t + s$ is differs only on s , but not when that migration is first anticipated. This may be a reasonable assumption because migration and expectations of the local unemployment rate may not be particularly salient to many people.⁴⁰

With this assumption, I can use my immigration shock as an instrument for immigration's effect on unemployment, and multiply the coefficients from the IV with the estimates from

³⁹This result justifies the focus of this paper on immigration. Immigration is the relevant margin to focus on because it responds more strongly to labor demand. Monras (2015a) also finds that immigration is the more reactive margin. He looks at different shocks more explicitly related to the Great Recession.

⁴⁰Consider how the effects of migration might be different if that migration is anticipated. Regardless of when the new migrants move in, the economy transitions to a new steady-state in terms of the housing stock. If it is known in advance, the construction of new houses will begin before the immigration because non-migrants will anticipate the rise in house prices. So there is no change in the total number of additional construction jobs, only in the time period in which they occur. With rational expectations, the house price channel is driven by the unanticipated response of house prices. Hence, the house price channel would likely be smaller, but it would also begin in the period in which the migration becomes known, not when the migration actually happens. Hence, the total effect of anticipated migration is positive before the migration occurs, decreasing as the migration is further and further out, and is weaker in the periods after the migration than it would have been were it a surprise.

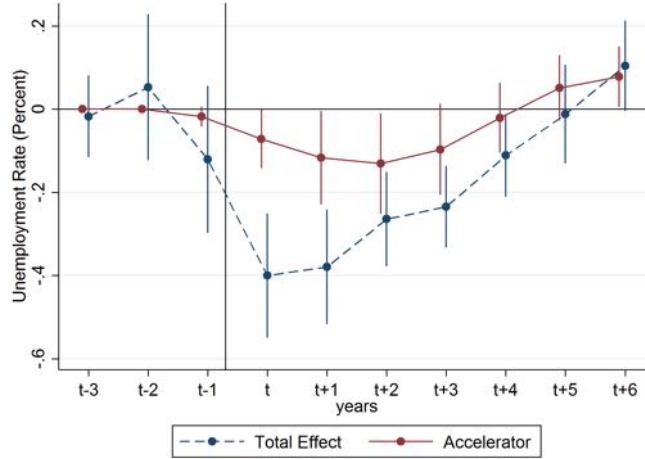


Figure 22: The response of unemployment to a Bartik shock, and the portion that is due to migration, along with 95 percent confidence intervals. Standard errors clustered by state.

migration's response to the Bartik shock, to calculate the part of the unemployment effect that comes from migration.

The instrumental variables specification is:

$$\Delta u_{n,t} = \sum_{s=0}^6 \gamma_s \Delta m_{n,t-s} + \alpha_t + \epsilon_{n,t}$$

$$\Delta m_{n,t-s} = \sum_{r=0}^6 \kappa_{r,s} \Delta z_{n,t-r} + \nu_{t,s} + \eta_{n,t,s}$$

I then combine the estimates of γ_s with the estimates of β_s^b from equation (2), the effect of the Bartik shock on migration. Hence, the accelerator is

$$\text{Accelerator}_s = \sum_{q+r=s} \beta_q^b \gamma_r$$

I show the estimated Accelerator, along with the total effect of the Bartik shock on unemployment in Figure 22. As you can see, the effects from migration explain a small but significant portion of the unemployment rate's response. In the first year, the response is equal to 16 percent of the total effect, implying migration amplifies the effect of a Bartik shock by 20 percent. In subsequent years, it is an even larger fraction of the total effect.

In the counterfactual world where migration did not respond to labor demand, the effect of the Bartik shock would have been smaller, equal to the difference between the two lines. While qualitatively similar, the effect would have been 16 percent more muted.

One implication of this exercise is that migration increases the volatility of the local unemployment rate with respect to Bartik shocks. Hence, for non-movers, migration is amplifying the risk of shocks to local demand, by 20 percent. In addition, much of the persistence in the unemployment rate is driven by this migration as well. We can see this because, at later time periods, the accelerator component is an even larger fraction of the total change. Hence, migration can explain a fraction of the persistent differences in regional outcomes.

The accelerator also has implications for cross-sectional outcomes, with the variance of the unemployment rate likely to be higher due to migration. However, this exercise focuses on shocks affecting only the MSA. It is not straightforward to extrapolate these results to shocks that are felt by large parts of the country, especially when it is regionally concentrated. For example, the migration response to the decline in Midwestern manufacturing might be smaller than these estimates because the same decline is occurring in many of the nearby places. So while shocks that affect many MSAs will be amplified by endogenous migration, it may be smaller than these estimates. I expand on this more in the context of the Great Recession in the next section.

Characteristics of Marginal Migrants

A key assumption for this to be a valid exercise is that the migrants have the same housing and non-tradable demand whether they come because of the migration shock or because of the labor demand shock. I am using the effect of migrants that move because of the migration shock as an approximation for the effect of migrants that move because of the increased labor demand. Within the model, the two important statistics for this exercise were the consumption of non-tradables and housing.

Table 2: Characteristics of Migrants induced by the two different shocks

VARIABLES	(1) AGI (\$1000s)	(2) AGI (\$1000s)	(3) Returns	(4) Returns
Migration	26.00*** (2.123)	26.91*** (3.224)	0.494*** (0.0103)	0.525*** (0.0198)
Kleibergen-Paap F-Statistic	64.8	13.7	64.8	13.7
CBSAs	917	917	917	917
Instruments	Migration Shocks	Labor Demand Shocks	Migration Shocks	Labor Demand Shocks

Standard errors clustered at CBSA level.

*** p<0.01, ** p<0.05, * p<0.1

The IRS Migration Statistics do not measure consumption of non-tradables or housing, but does include the adjusted gross income and the number of returns. These two statistics might be reasonable proxies for housing and non-tradable demand. Certainly, the number of returns per exemption will be related to family size, and the adjusted gross income is probably a good proxy for how rich the migrants are, two important determinants of demand.

I only see the totals for county-to-county flows, similar to exemptions. So I can only estimate the effect on the means of these variables. To find the average income of these migrants, I run the following regression,

$$\Delta \text{AGI migration rate}_{i,t} = \beta \Delta m_{i,t} + \alpha_t + \epsilon_{i,t}$$

where the AGI migration rate is the total income of all migrants into the MSA, normalized by the MSA's population. I then instrument for the migration rate using lags of my migration instrument, or lags of my labor demand instrument. I can do a similar thing for the average returns-to-exemptions ratio.

The results are presented in Table 2. All the first-stage F-statistics are above ten, though not surprisingly, the migration shocks do a better job of predicting migration than the labor demand shocks. The incomes of migrants induced by the two shocks are almost identical.

There does seem to be a small effect on the returns, implying migrants induced by labor demand shocks have slightly smaller families, which might mean they have less housing demand.

5.3 The Effect of Migration on the Great Recession

This paper began by showing a map of how different areas in the country experienced different outcomes during the Great Recession. During this time period, there was a corresponding decline in migration to the areas that were hit hardest. In this section, I consider a counterfactual, where immigration is held at its 2004 patterns, to see how that would have affected the United States during the recession.⁴¹

To construct this counterfactual, I assume that immigration to an MSA is based on its 2004 rate, and adjust that evenly to account for the decline in gross migration over this time period. Hence, in the counterfactual, I would assume that migration in an MSA is equal to its 2004 rate minus an adjustment for national migration.⁴² Once I have the counterfactual migration rate, I compare that to the actual migration rate, creating a spread between the actual immigration and the counterfactual one. I then calculate what the effect of that spread on the unemployment rate. I have to make the same assumption as before in order to apply my estimates, that migration affects unemployment the same way, regardless of when it is anticipated. I then run a similar instrumental variables regression to the last section, except that I allow it to vary based on the housing supply elasticity of Saiz (2010). I use the estimates from that regression, along with the previously calculated spread in order to estimate how the unemployment rate would have changed.

For the data, I consider the change in the unemployment rate between 2006 and 2009, using the average annual unemployment rate. I calculate the spread in the immigration rate for each of these years, each of which has an effect on the unemployment rate in 2009.

⁴¹I chose 2004 because by 2006, migration had slowed considerably, and in 2005, immigration patterns were dominated by people leaving New Orleans after Hurricane Katrina.

⁴²I am still allowing outmigration to evolve as it did, so I have to make the adjustment so that total

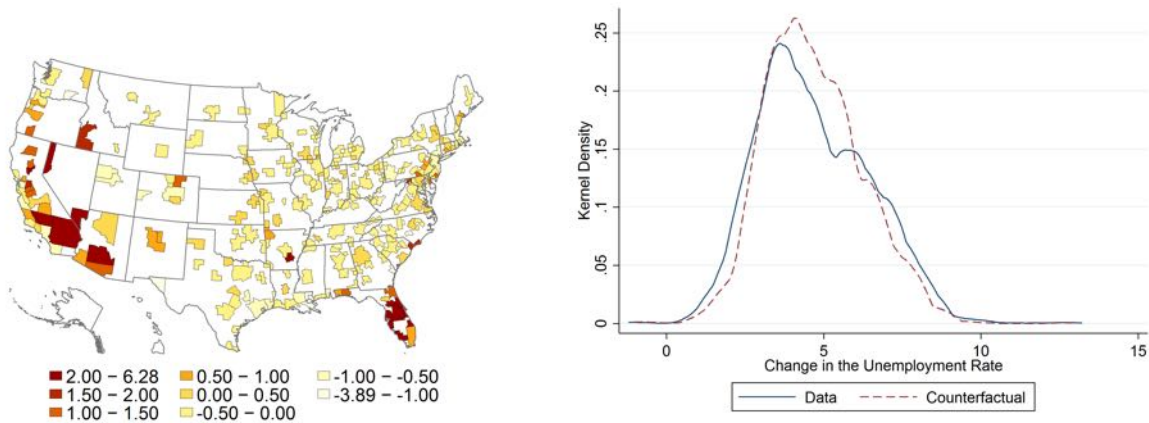


Figure 23: The difference in the change in the unemployment rate during the Great Recession (2006-2009) between data and the fixed-migration counterfactual (left), and the distribution of the change in the unemployment rate during the Great Recession, with and without the endogenous change in the location of immigration (right).

The left-side of Figure 23 maps the difference in the unemployment rate between the data and the counterfactual. In Florida and Southern California, the counterfactual predicts a smaller increase in the unemployment rate. Of course, these two areas were particularly affected by the recession, and had large declines in immigration. Many of these MSAs also have low housing supply elasticities, where the effects of immigration are larger. Taken together, these facts suggest that migration played a role in amplifying the effects of the recession in those regions, so it not a surprise that this is confirmed in the map.

The right-side of Figure 23 plots the distribution of the increase. The counterfactual features a slightly smaller increase in the unemployment rate, by 0.02 percentage points, due to the correlation between changes in migration and the elasticity. But more importantly, the standard deviation of the increase in the unemployment rate decreased by 13 percent.

The main takeaway is that the changing patterns of migration during the Great Recession amplified the differences in regional outcomes, making the recession even worse in the hardest hit areas. Had migratory patterns been held constant at 2004 patterns, there would have been less dispersion in outcomes, with the least affected areas experiencing larger increases

immigration and outmigration are equal.

in unemployment, and the worst-hit areas experiencing smaller increases.

6 Conclusion

In this paper, I document that immigration shocks have a large positive effect on a local labor market. This effect is explained by two housing channels: an increase in construction and an increase in house prices, inducing non-tradable consumption. Because of the positive effect, migration amplifies the effects of other labor demand shocks, counter to the traditional view of migration as an equilibrating force. In fact, I quantify these effects to be large, amplifying Bartik shocks by 20 percent, and increasing the dispersion of unemployment changes during the Great Recession by 13 percent.

There are two implications for policy. Currency unions do not benefit from migration in the way that many suppose from Mundell (1961). As recently as 2012, the head of the European Central Bank said “For the euro area, too, increased labour mobility across borders is crucial” (Draghi, 2012). My results, however, suggest that within a currency union, migration may cause changes in aggregate demand that exacerbates regional differences and hurts non-movers in depressed areas.⁴³ Rather than closing the differences between labor markets, migration amplifies their differences. This means that if the receiving MSA could control its own monetary policy, it would rather tighten by more than it would absent migration. Relative to a world without migration, differences in macro-stabilizing policies are larger.

Second, migration’s amplification of local shocks increases employment risk for people that do not move. However, this risk is smaller in cities with more elastic housing supplies. A large part of the elasticity of housing supply is endogenous to local zoning laws and other housing regulations (see Gyourko, Saiz, and Summers, 2008; Saiz, 2010). My estimates

⁴³Another reason labor mobility is important in currency unions is because it provides insurance for the migrants themselves. Insurance is more important in currency unions because monetary policy cannot play that role. My results have little to say about this role. However, Yagan (2014) shows that migration provided little insurance during the Great Recession.

suggest that increasing the housing supply elasticity would reduce the effects of migration on the unemployment rate, reducing the amplification that I document in this paper.

References

- AGARWAL, S., G. AMROMIN, S. CHOMSISENGPHET, T. PISKORSKI, A. SERU, AND V. YAO (2015): “Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program,” *National Bureau of Economic Research Working Paper*.
- AIYAGARI, S. R., L. J. CHRISTIANO, AND M. EICHENBAUM (1992): “The output, employment, and interest rate effects of government consumption,” *Journal of Monetary Economics*, 30, 73–86.
- ALLEN, T. AND C. ARKOLAKIS (2014): “Trade and the Topography of the Spatial Economy,” *The Quarterly Journal of Economics*, 129, 1085–1140.
- ALTONJI, J. G. AND D. CARD (1991): “The effects of immigration on the labor market outcomes of less-skilled natives,” in *Immigration, trade, and the labor market*, University of Chicago Press, 201–234.
- AMIOR, M. AND A. MANNING (2016): “The persistence of local joblessness,” .
- ANGRIST, J. D., Ò. JORDÀ, AND G. KUERSTEINER (2013): “Semiparametric estimates of monetary policy effects: string theory revisited,” *National Bureau of Economic Research Working Papers*.
- ARELLANO, M. AND S. BOND (1991): “Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations,” *The Review of Economic Studies*, 58, 277–297.
- ATTANASIO, O., A. LEICESTER, AND M. WAKEFIELD (2011): “Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK,” *Journal of the European Economic Association*, 9, 399–435.
- AUDRETSCH, D. B. AND M. P. FELDMAN (2004): “Knowledge spillovers and the geography of innovation,” *Handbook of regional and urban economics*, 4, 2713–2739.
- BARTEL, A. P. (1989): “Where do the new US immigrants live?” *Journal of Labor Economics*, 371–391.
- BARTIK, T. J. (1991): “Boon or Boondoggle? The debate over state and local economic development policies,” .
- BAXTER, M. AND R. G. KING (1993): “Fiscal policy in general equilibrium,” *The American Economic Review*, 315–334.
- BEAUDRY, P., D. A. GREEN, AND B. M. SAND (2014): “In search of labor demand,” National bureau of economic research working paper.
- BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2015): “House prices and consumer spending,” *National Bureau of Economic Research Working Paper*.

- BERMAN, A. AND R. J. PLEMMONS (1979): “Nonnegative matrices,” *The Mathematical Sciences, Classics in Applied Mathematics*, 9.
- BLANCHARD, O. J. AND L. F. KATZ (1992): “Regional evolutions,” *Brookings Papers on Economic Activity*, 1992, 1–75.
- BODVARSSON, Ö. B., H. F. VAN DEN BERG, AND J. J. LEWER (2008): “Measuring immigration’s effects on labor demand: A reexamination of the Mariel Boatlift,” *Labour Economics*, 15, 560–574.
- BORJAS, G. J. (2003): “The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market,” *The Quarterly journal of economics*, 118, 1335–1374.
- (2015): “The Wage Impact of the Marielitos: A Reappraisal,” .
- BOUND, J. AND H. J. HOLZER (2000): “Demand shifts, population adjustments, and labor market outcomes during the 1980s,” *Journal of Labor Economics*, 18, 20–54.
- CADENA, B. C. AND B. K. KOVAK (2016): “Immigrants equilibrate local labor markets: Evidence from the Great Recession,” *American Economic Journal: Applied Economics*, 8, 257–290.
- CALIENDO, L., F. PARRO, E. ROSSI-HANSBERG, AND P.-D. SARTE (2014): “The impact of regional and sectoral productivity changes on the US economy,” *National Bureau of Economic Research Working Paper*.
- CAMPBELL, J. Y. AND J. F. COCCO (2007): “How do house prices affect consumption? Evidence from micro data,” *Journal of Monetary Economics*, 54, 591–621.
- CARD, D. (1990): “The impact of the Mariel boatlift on the Miami labor market,” *Industrial & Labor Relations Review*, 43, 245–257.
- (2001): “Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration,” *Journal of Labor Economics*, 19, 22–64.
- (2009): “How immigration affects U.S. cities,” *Making Cities Work: Prospects and Policies for Urban America*, 158–200.
- CARLINER, M. (2003): “New home cost components,” *Housing Economics*, 51, 7.
- CHARLES, K. K., E. HURST, AND M. J. NOTOWIDIGDO (2016): “Housing Booms, Manufacturing Decline, and Labor Market Outcomes,” Working paper.
- COEN-PIRANI, D. (2010): “Understanding gross worker flows across us states,” *Journal of Monetary Economics*, 57, 769–784.
- DAO, M., D. FURCERI, AND P. LOUNGANI (2014): “Regional labor market adjustment in the United States,” IMF Working Paper.

- DAVIS, M. A., J. D. FISHER, AND M. VERACIERTO (2013): “Gross migration, housing and urban population dynamics,” *FRB of Chicago Working Paper*.
- DE SILVA, D. G., R. P. MCCOMB, Y.-K. MOH, A. R. SCHILLER, AND A. J. VARGAS (2010): “The effect of migration on wages: evidence from a natural experiment,” *The American Economic Review: Papers & Proceedings*, 100, 321–326.
- DECRESSIN, J. AND A. FATAS (1995): “Regional labor market dynamics in Europe,” *European Economic Review*, 39, 1627–1655.
- DIAMOND, P. A. (1982): “Wage determination and efficiency in search equilibrium,” *The Review of Economic Studies*, 49, 217–227.
- DIAMOND, R. (2016): “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000,” *American Economic Review*, 106, 479–524.
- DRAGHI, M. (2012): “Competitiveness: the key to balanced growth in monetary union,” Remarks by Mario Draghi, President of the ECB, Treasury Talks. ‘A European strategy for growth and integration with solidarity,’ A conference organised by the Directorate General of the Treasury, Ministry of Economy and Finance–Ministry for Foreign Trade, Paris [Accessed 7 Sept 2016].
- DUSTMANN, C., U. SCHÖNBERG, AND J. STUHLER (2016): “The Impact of Immigration: Why Do Studies Reach Such Different Results?” *Journal of Economic Perspectives*, 30, 31–56.
- FARHI, E. AND I. WERNING (2014): “Labor Mobility Within Currency Unions,” *National Bureau of Economic Research Working Paper*.
- FOOTE, C. L. (2007): “Space and time in macroeconomic panel data: young workers and state-level unemployment revisited,” .
- GAGNON, E. AND D. LOPEZ-SALIDO (2014): “Small price responses to large demand shocks,” *Feds working paper*.
- GLAESER, E. L. AND J. GYOURKO (2005): “Urban Decline and Durable Housing,” *Journal of Political Economy*, 113, 345–375.
- GONZALEZ, L. AND F. ORTEGA (2013): “Immigration and housing booms: Evidence from Spain,” *Journal of Regional Science*, 53, 37–59.
- GREULICH, E., J. M. QUIGLEY, AND S. RAPHAEL (2004): “The Anatomy of Rent Burdens: Immigration, Growth, and Rental Housing,” *Brookings-Wharton Papers on Urban Affairs*, 149–205.
- GYOURKO, J., A. SAIZ, AND A. SUMMERS (2008): “A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index,” *Urban Studies*, 45, 693–729.

- HANSEN, L. P. AND T. J. SARGENT (1981): “Exact linear rational expectations models: Specification and estimation,” *Staff Report: Federal Reserve Bank of Minneapolis*.
- HANSON, G. H. AND M. J. SLAUGHTER (2002): “Labor-market adjustment in open economies: Evidence from U.S. states,” *Journal of International Economics*, 57, 3–29.
- HARRIS, J. R. AND M. P. TODARO (1970): “Migration, unemployment and development: a two-sector analysis,” *The American Economic Review*, 126–142.
- HEAD, A., H. LLOYD-ELLIS, AND H. SUN (2014): “Search, Liquidity, and the Dynamics of House Prices and Construction,” *American Economic Review*, 104, 1172–1210.
- HONG, G. AND J. MCLAREN (2015): “Are Immigrants a Shot in the Arm for the Local Economy?” *National Bureau of Economic Research Working Paper*.
- JIMENO, J. F. AND S. BENTOLILA (1998): “Regional unemployment persistence (Spain, 1976–1994),” *Labour Economics*, 5, 25–51.
- JORDÀ, Ò. (2005): “Estimation and inference of impulse responses by local projections,” *American Economic Review*, 161–182.
- KAPLAN, G., K. MITMAN, AND G. VIOLANTE (2016): “Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data,” *National Bureau of Economic Research Working Paper*.
- KLINE, P. AND E. MORETTI (2014): “Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority,” *The Quarterly Journal of Economics*, 129, 275–331.
- KRUGMAN, P. (1980): “Scale economies, product differentiation, and the pattern of trade,” *The American Economic Review*, 70, 950–959.
- LEE, D., C. J. MAYER, AND J. TRACY (2012): “A new look at second liens,” *National Bureau of Economic Research Working Paper*.
- LEWIS, E. (2004): “How did the Miami labor market absorb the Mariel immigrants?” Federal reserve bank of philadelphia working paper.
- LEWIS, E. G. (2005): “Immigration, skill mix, and the choice of technique,” *FRB Philadelphia Working Paper*.
- LONG, J. AND H. E. SIU (2016): “Refugees From Dust and Shrinking Land: Tracking the Dust Bowl Migrants,” *National Bureau of Economic Research Working Paper*.
- MARLAY, M. AND P. MATEYKA (2011): “The Seasonality of Moves: 2009,” .
- MCINTOSH, M. F. (2008): “Measuring the labor market impacts of Hurricane Katrina migration: evidence from Houston, Texas,” *The American Economic Review*, 98, 54–57.

- MIAN, A., K. RAO, AND A. SUFI (2013): “Household Balance Sheets, Consumption, and the Economic Slump*.” *Quarterly Journal of Economics*, 128.
- MOLLOY, R., C. L. SMITH, AND A. WOZNIAK (2011): “Internal migration in the United States,” *The Journal of Economic Perspectives*, 25, 173–196.
- (2014): “Declining migration within the US: the role of the labor market,” *National Bureau of Economic Research Working Paper*.
- MONRAS, J. (2015a): “Economic Shocks and Internal Migration,” *Job Market Paper*.
- (2015b): “Immigration and wage dynamics: Evidence from the mexican peso crisis,” Iza discussion paper.
- MUNDELL, R. A. (1961): “A theory of optimum currency areas,” *The American Economic Review*, 51, 657–665.
- NENOV, P. T. (2015): “Regional reallocation and housing markets in a model of frictional migration,” *Review of Economic Dynamics*, 18, 863–880.
- NOTOWIDIGDO, M. J. (2013): “The Incidence of Local Labor Demand Shocks,” Working paper, National Bureau of Economic Research.
- OTTAVIANO, G. AND G. PERI (2006): “Wages, rents and prices: The effects of immigration on US natives,” .
- (2012): “Rethinking the effect of immigration on wages,” *Journal of the European Economic Association*, 10, 152–197.
- PLAGBORG-MØLLER, M. (2015): “Bayesian Inference on Structural Impulse Response Functions,” *Job Market Paper*.
- POTERBA, J. M. (1984): “Tax subsidies to owner-occupied housing: an asset-market approach,” *The Quarterly Journal of Economics*, 729–752.
- RAMEY, V. A. (2016): “Macroeconomic shocks and their propagation,” *Handbook of Macroeconomics*, forthcoming.
- RAPPAPORT, J. (2007): “Moving to nice weather,” *Regional Science and Urban Economics*, 37, 375–398.
- REDDING, S. J. AND E. ROSSI-HANSBERG (2016): “Quantitative Spatial Economics,” *National Bureau of Economic Research Working Paper*.
- ROGNLIE, M., A. SHLEIFER, AND A. SIMSEK (2015): “Investment hangover and the great recession,” *National Bureau of Economic Research Working Paper*.
- RUGGLES, S., K. GENADEK, R. GOEKEN, J. GROVER, AND M. SOBEK (2015): “Integrated Public Use Microdata Series: Version 6.0 [Machine Readable Database],” .

- RYBCZYNSKI, T. M. (1955): “Factor endowment and relative commodity prices,” *Economica*, 22, 336–341.
- SAIZ, A. (2003): “Room in the kitchen for the melting pot: Immigration and rental prices,” *Review of Economics and Statistics*, 85, 502–521.
- (2007): “Immigration and housing rents in American cities,” *Journal of Urban Economics*, 61, 345–371.
- (2010): “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*, 125.
- SCHULHOFER-WOHL, S. (2011): “Negative equity does not reduce homeowners’ mobility,” *NBER Working Paper*.
- SHIMER, R. (2001): “The Impact of Young Workers on the Aggregate Labor Market,” *The Quarterly journal of economics*, 116, 969–1007.
- (2005): “The cyclical behavior of equilibrium unemployment and vacancies,” *The American Economic Review*, 95, 25–49.
- STRÖBEL, J. AND J. VAVRA (2015): “House Prices, Local Demand, and Retail Prices,” *National Bureau of Economic Research Working Paper*.
- STRUYVEN, D. (2014): “Housing Lock: Dutch Evidence on the Impact of Negative Home Equity on Household Mobility,” *Job Market Paper*.
- STUCKE, A. L. (1949): “Labor Share in Construction Cost of New Houses,” *Monthly Labor Review*, 517–520.
- TOPEL, R. AND S. ROSEN (1988): “Housing investment in the United States,” *The Journal of Political Economy*, 718–740.
- VAN NIEUWERBURGH, S. AND P.-O. WEILL (2010): “Why has house price dispersion gone up?” *The Review of Economic Studies*, 77, 1567–1606.
- YAGAN, D. (2014): “Moving to opportunity? Migratory insurance over the Great Recession,” *Job Market Paper*.

A Theoretical Framework

In this section, I present a dynamic microfounded model of a city that includes endogenous migration. I consider the effects of immigration on the equilibrium, and especially its effect on the unemployment rate. I highlight the construction channel and the house price channel, showing that without them, migration is unable to cause a decline in the unemployment rate. In a special case of the model, where I consider log-utility so as to not keep track of the wealth distribution of agents, I show that if the preference parameter on housing consumption is high enough, these two channels will cause a decline in the unemployment rate.

There will be two main shocks to the economy. The first is to the utility of living elsewhere, what I call a “migration shock.” The second will be a labor demand shock. Each of these shocks has an empirical counterpart.

For notation, I will use capital letters as aggregate quantities within the MSA, and lower-case letters as per-capita terms. Prices will also be lower-case. I will use lower subscripts for indexing of people i and time periods t . When there is no time component to an equation, I may omit the lower t subscript.

A.1 Setup

A.1.1 The agent’s problem

Denote Γ as the aggregate state of the economy, which is the previous period’s population of the MSA, N , the wealth distribution of all agents, the housing owned by each agent, and the two shocks: the outside option $G(\cdot)$ for potential migrants, and tradable demand, D_X . The agents use these state variables to forecast future prices. Agents, indexed by i , value their consumption of housing, h ; tradables, c^T ; and non-tradables, c^{NT} . They own one-period bonds a and housing h , and e denotes their employment status. The discount the next period’s utility by β . With probability ρ , they leave the MSA, and receive continuation

utility V^* .⁴⁴

$$V(\Gamma, a_i, h_i, e_i) = \max_{c_i^{NT}, c_i^T, h'_i, a'_i} u(c_i^{NT}, c_i^T, h'_i) + \beta \left((1 - \rho) E[V(\Gamma', a'_i, h'_i, e'_i)] + \rho V^*(a'_i + p^h h'_i) \right)$$

and are subject to a budget constraint:

$$c_i^T + p c_i^{NT} + p^h h'_i + \frac{1}{1 + R} a'_i \leq p^h (1 - \delta) h_i + a_i + w e_i$$

I normalize the price of tradables to one, so p is the price of non-tradable goods, w is the wage, and p^h is the price of housing. δ is the depreciation rate for housing. The agent is also subject to a collateral constraint:

$$(1 - \delta) h'_i p^{h'} \leq \phi a'_i$$

Assume a unit mass of potential migrants considers moving in. If they do move in, their value function is:

$$V(\Gamma, a_i, 0, e) \equiv E_e[V(\Gamma, a_i, 0, e_i)]$$

The expectation is over whether or not they get a job when they first move in. The potential migrants also have an outside option, V_i^* . There is a joint distribution over the outside option and assets, which is stochastic.

$$(V_i^*, a_i) \sim G(\cdot, \cdot)$$

Define G_a as the marginal distribution of a , and $G_{V|a}$ as the partial distribution of V for a given a . Hence $g(\hat{V}, \hat{a}) = g_a(\hat{a}) g_{V|\hat{a}}(\hat{V})$.

Hence the population N is

$$N = (1 - \rho) N_{-1} + m \tag{3}$$

⁴⁴I choose to model outmigration as exogenously determined because it is relatively unresponsive to migration and labor demand shocks, as I show empirically in Section 3.4 and Section 5.1.

where

$$m = \int G_{V|a}^{-1}(E_e[V(\Gamma, a, 0, e)]) dG_a(a) \quad (4)$$

Implicit in this notation is a timing assumption: potential migrants first choose whether to move in, then either receive or do not receive a job, and then make spending decisions.

A.1.2 Land and Housing

The MSA has a fixed amount of land, L , upon which housing is built. Each period, a fraction δ of the housing depreciates, leaving δL available for development. Although it is not key to the model, it is easiest to assume that the land from depreciated houses is lump-sum taxed by the government, who then sells it to house-producers and keeps the profits.⁴⁵ Housing is produced competitively using that land, as well as labor and imported goods. The price of housing is flexible and the production function H is constant returns to scale.

$$H' = (1 - \delta)H + H(D_h, T^H, \delta L)$$

The price of housing is then pinned down by a marginal cost function:

$$p^h = p^h(H' - (1 - \delta)H, w) \quad (5)$$

where p^h is increasing in both arguments. The labor demand from housing is determined by the same things:

$$D_h = D_h(H' - (1 - \delta)H, w) \quad (6)$$

D_h is increasing in the first argument and decreasing in the second argument.

⁴⁵A reasonable alternative assumption is that homeowners keep the land. Because the price of land and the price of houses are monotonically related, Proposition 1 is not affected by this assumption. However, the equations in Appendix A.4 become more complex.

A.1.3 Tradables and Non-tradables

Production of non-tradable goods is linear in labor. I normalize productivity to 1. Hence, labor demand for the production of non-tradables is

$$Y_{NT} = D_c \tag{7}$$

Assume further that the market is competitive and that prices are not sticky beyond any stickiness in wages.

$$w = p \tag{8}$$

Assume $D_x(w)$, the labor demand from the production of tradable goods, is stochastic and exogenously given.

A.1.4 Aggregate demand

Aggregate demand, expressed in labor units, is the sum of labor demand for non-tradable goods, tradable goods, and housing.⁴⁶

$$D = D_c + D_x + D_h \tag{9}$$

A.1.5 Philips Curve

Assume wages, and hence prices, are perfectly sticky.

$$w' = w \tag{10}$$

While this is extreme, it helps to highlight my mechanism.

⁴⁶This equation looks a lot like the $Y = C + I + G + X - M$ equation from introductory macro. In my notation, D_c is equal to consumption minus imports, the only investment is in housing, and, in my baseline model, I am abstracting from government spending.

Labor is allocated randomly between agents.

$$P(e_i = 1) = e = \frac{D}{N} \quad (11)$$

A.2 Equilibrium

Agents solve their maximization problem:

$$c_i^{NT} = c_i^{NT}(\Gamma, a_i, h_i, e_i) \quad (12)$$

$$c_i^T = c_i^T(\Gamma, a_i, h_i, e_i) \quad (13)$$

$$h_i = h_i(\Gamma, a_i, h_i, e_i) \quad (14)$$

Assume that these are all twice continuously-differentiable.

The market clearing conditions are

$$H_t = \int h_i di \quad (15)$$

$$Y_{NT} = \int c_i^{NT} di \quad (16)$$

For any shocks G and D_x and state Γ , an equilibrium is the new population N , the migration m ; prices, w , p , and p^h ; employment, e_i ; consumption, c_i^{NT} , c_i^T , and h_i ; and aggregate variables H , Y_{NT} , D_c , D_h , and D ; such that equations (3)-(16) hold.

A.3 The effect of migration on aggregate demand per capita

Because Γ is so large, it is intractable to solve this model with aggregate shocks. Rather, I will linearize around a deterministic steady-state, where the variance of the outside option and the demand for tradable goods approaches zero. Without these aggregate shocks, agents can perfectly forecast p^h and e , which are the only aggregate variables that matter for their migration and consumption decisions.

The value function can then be written as follows

$$V(\{p_s^h, e_s\}_{s \geq t}, a_i, h_i, e_i)$$

and the consumption functions of agents can similarly replace Γ with $\{p_s^h, e_s\}_{s \geq t}$.

One of the key questions this paper seeks to answer is what is the effect of a migration shock on the employment rate (and also the unemployment rate). Based on our decomposition above, we can break down $\frac{de_t}{dM}$.

$$\begin{aligned} N_t \frac{de_t}{dm} = & \underbrace{-e_t(1-\rho)^t}_{\text{Labor supply increase}} + \underbrace{c_{m,t}^{NT}(1-\rho)^t}_{\text{Migrant demand}} + \underbrace{\sum_{s=0}^{\infty} \int \frac{\partial c_{i,t}^{NT}}{\partial e_s} di \frac{de_s}{dm}}_{\text{Non-tradable Keynesian multiplier}} \\ & + \underbrace{\sum_{s=0}^{\infty} \int \frac{\partial c_{i,t}^{NT}}{\partial p_s^h} di \frac{dp_s^h}{dm}}_{\text{House price channel}} + \underbrace{\frac{dD_{H,t}}{dH_t} \left(\frac{dH_t}{dm} - (1-\delta) \frac{dH_{t-1}}{dm} \right)}_{\text{Construction channel}} \end{aligned}$$

The first two terms are the direct effects of migrants: they increase the labor supply, leading to lower employment rates; and they consume non-tradable goods, requiring labor and increasing employment. The third term is a Keynesian multiplier, amplifying the effects of the other terms because with more employment, agents consume more non-tradables.

The fourth and fifth terms are new, and highlight the two channels. The fourth term is the effect that house prices have on non-tradable consumption. With the increase in housing demand, house prices increase, and non-tradable consumption increases because of that. Berger et al. (2015) break down the effect of house prices into four: a wealth effect, a substitution effect, a collateral effect, and an income effect. They also argue on empirical and theoretical grounds that the total effect is sizable. My model includes all four effects, though not the assumptions to imply that the income, substitution, and collateral effect exactly cancel out.

The fifth term is the construction channel. With the housing demand increase, the number of houses will increase. The change in construction demand is proportional to the

change in the number of houses in period t , but is negatively affected by the number of new houses in period $t - 1$, because there is already a stock of housing that does not require labor to be built. Hence, there is a front-loading effect as the stock of housing is built up.

Note the demand for non-tradable goods does not change because we hold prices fixed.

Of course, housing and house prices are determined in equilibrium. The change is the solution to the following system of equations. But more importantly, we can also estimate these in the data. Hence, the equation above is a useful guide to the empirical exercise even without solving out for $\frac{dp^h}{dm}$ and $\frac{dH}{dm}$.

$$\begin{aligned}\frac{dH_t}{dm} &= dh_m^t(1 - \rho)^t + \sum_{s=0}^{\infty} \int \frac{\partial h_i}{\partial e_s} di \frac{de_s}{dm} + \sum_{s=0}^{\infty} \int \frac{\partial h_{i,t}}{\partial p_s^h} di \frac{dp_s^h}{dm} \\ \frac{dp_t^h}{dm} &= \frac{\partial p_t^h}{\partial H_t} \left(\frac{dH_t}{dm} - (1 - \delta) \frac{dH_{t-1}}{dm} \right)\end{aligned}$$

This decomposition allows me to a proposition about the traditional view, in the spirit of Farhi and Werning (2014).

Proposition 1 (The Traditional View). *Suppose there were no housing, $H_t = 0$. If migrants are less wealthy than non-migrants, i.e. the wealth distribution of migrants is first-order stochastically dominated by the wealth distribution including housing of non-migrants, then the employment rate weakly decreases (the unemployment rate weakly increases) in every subsequent period.*

Proof. Without the housing term,

$$N_t \frac{de_t}{dm} = -e_t(1 - \rho)^t + c_{m,t}^{NT}(1 - \rho)^t + \sum_{s=0}^{\infty} \int \frac{\partial c_{i,t}^{NT}}{\partial e_s} di \frac{de_s}{dm}$$

where $N_t e_t = \int c_{i,t}^{NT} + D_x$. Consumption in any period is increasing in a_0 for a given e_i . Because the probability of employment is the same for all agents, migrants and non-migrants, it must be the case that $c_{m,t}^{NT} < \frac{1}{N_t} \int c_{i,t}^{NT}$ given the assumption that the wealth distribution of non-migrants dominates that of migrants. D_x must be greater than or equal to zero, so

then it must be the case that

$$c_{m,t}^{NT} < e_t$$

Hence, the sum of the first two terms on the right-hand side are negative.

Consider the $T \times T$ matrix A_T where $A_{T,ij} = -\frac{1}{N_t} \int \frac{\partial c_{i,t}^{NT}}{\partial e_s} di$ if $i \neq j$ and $A_{T,ii} = 1 - \frac{1}{N_t} \int \frac{\partial c_{i,t}^{NT}}{\partial e_s} di$. Note that it must be the case that $\sum_{s=0}^T \frac{1}{(1+R)^{t-s}} \frac{\partial c_{i,t}^{NT}}{\partial e_s} \leq 1$ because of the budget constraint.

Hence $B_T^{-1}A_TB_T$ is a M-matrix where B_T is a diagonal matrix and $B_{T,tt}$ is $\frac{1}{(1+R)^t}$. This implies that all the elements of A_T^{-1} are weakly positive, because the inverse of M-matrices have this property (see Berman and Plemmons, 1979).

Now consider $A_T[(e_t - c_{m,t}^{NT})(1-\rho)^t]_{t=1}^T$. The elements of this vector are all strictly negative because of the argument above. Therefore, the product is weakly negative. Now consider the $\liminf_{T \rightarrow \infty} A_T[(e_t - c_{m,t}^{NT})(1-\rho)^t]_{t=1}^T$ which must also be weakly negative.

Furthermore, this must be monotonic in T . To see this, consider $C_T = I_T - B_T^{-1}A_TB_T$ which is weakly positive everywhere. So $(B_T^{-1}A_TB_T)^{-1} = \lim_{r \rightarrow \infty} (C_T)^r$. Because it is weakly positive, for every $i \leq T$ and $j \leq T$, it must be true that $(C_T)_{ij}^r \leq (C_{T+1})_{ij}^r$. Hence $A_{T,ij}$ must also be increasing in T . Hence $A_T[(e_t - c_{m,t}^{NT})(1-\rho)^t]_{t=1}^T$ is also increasing in T . Therefore the limit exists, and is equal to $\frac{de_t}{dm}$. Because the liminf was negative, $\frac{de_t}{dm}$ must be negative. \square

The assumption that migrants are less wealthy is reasonable. From Molloy et al. (2014), we know that most migrants tend to be younger than the general population, and so have had less time to build up wealth. In Section 4.5, I show that the interest, dividend, and rental income of migrants is first-order stochastically dominated by non-migrants in ACS data.

The important part of the proof is that the labor supply goes up by more than the labor demand. Because the migrants are less wealthy, they spend less money on non-tradable goods than non-migrants. Hence the demand for non-tradables goes up by a smaller percentage than the labor supply increase. Combined with the fact that the Keynesian multiplier term

simply amplifies other effects, the employment rate must decrease.

The key insight of this paper is that the logic of Proposition 1 breaks down when housing is incorporated into the model. One way in which it breaks down is that housing is through the construction channel. With non-tradables, the key assumption is that changes in non-tradable demand and changes in labor demand are contemporaneous. But the additional labor demand from housing is $\frac{dD_{H,t}}{dH_t} \left(\frac{dH_t}{dm} - (1 - \delta) \frac{dH_{t-1}}{dm} \right)$, which is not contemporaneous. $D_{h,t}$ depends on both H_t (positively) and H_{t-1} (negatively). For example, if $\frac{dH_t}{dm}$ is constant for all $t \geq 0$, the additional labor demand is much higher in the first period and only slightly larger in all subsequent periods. Hence, for even small increases in the amount of housing demand, the amount of initial labor demand can be quite large.

The other reason housing breaks the logic of Proposition 1 is that the consumption of homeowners increases when house prices appreciate. For one thing, the wealth W_0, i increases when house prices go up. In addition, borrowing limits are relaxed, leading to a collateral channel. These may be mitigated by an income effect because it is more expensive to own housing, and the substitution effect of housing has ambiguous effects on employment. But empirical estimates have shown that increases in house prices generally lead to higher consumption.

In Appendix A.4, I show that for a model with log-utility and Cobb-Douglas housing production, then with a high enough preference for housing, the unemployment rate will fall in response to migration. This is true because agents will own more housing and employ more construction workers.

A.4 Simple Model

In this section, I consider a specific version of this model in which the intuitions are more clear. The linearized version can be solved analytically, it serves as a proof-of-concept that housing can cause a boom, and it illuminates the role of the housing supply elasticity.

Assumption A.1. *Housing production is Cobb-Douglass:*

$$H_t = (1 - \delta)H_{t-1} + A(\delta L)^\alpha D_{h,t}^\gamma T_t^{1-\alpha-\gamma}.$$

Note that α is positively related to the housing supply elasticity because land is available in fixed supply, while labor and tradable goods have a set price.

Assumption A.2. *Suppose utility were log:*

$$U(h, c^{NT}, c^T) = \phi \log h + \psi \log c^{NT} + (1 - \phi - \psi) \log c^T$$

and that there is no collateral constraint except for the natural borrowing limit.

In this model, agents consume proportionally to their total wealth, including their house, any asset holdings, and discounted future earnings. In response to a migration shock, the increase in house prices will have an effect on their wealth, which will increase spending in constant proportions forever. In Berger et al. (2015), the consumption effects of house prices changes are equal to the wealth effects when utility is log. That is the case here as well, although these assumptions imply a low MPC because of the permanent-income nature of these agents.

Assumption A.3. *All agents are identical except for their choice of location. Furthermore, $\rho = 0$ and there is no gross migration in steady-state.*

This assumption simplifies the analysis by assuming that the new migrants consume housing and non-tradable goods in the same amount as the original workers. It also simplifies the expression for the effect on labor supply, because everyone earns the same y .

Assumption A.4. *The economy is at the deterministic steady-state.*

Note that another assumption, required for this to be true, is that $\beta(1 + R) = 1$.

With these four assumptions, the effect of a migration shock, $\frac{dy}{dm}$, can be worked out analytically.

A couple definitions are helpful: Define $\tilde{h}_t = h_t - (1 - \delta)h_{t-1}$ and $\tilde{H}_t = H_t - (1 - \delta)h_{t-1}$. Define $W = p_0^h h_{-1}(1 - \delta) + \sum_{t=0}^{\infty} \beta^t y_t$ and $r_t^h = p_t^h - \frac{1}{1+R}(1 - \delta)p_{t+1}^h$.

Housing demand is given by:

$$h_s r_s^h = \phi \frac{R}{1+R} W$$

which implies $h_0 p_0^h = \frac{1+R}{R+\delta} \phi R W$. The change in housing demand is given by

$$d \log h_s = -d \log r_s^h + d \log W$$

Log-linearizing r and \tilde{h} , we get the following equations, which characterize the path of housing construction and rents.

$$d \log r_t^h = \frac{1+R}{R+\delta} \frac{\alpha}{1-\alpha} \left(d \log \tilde{H}_t - \frac{1-\delta}{1+R} d \log \tilde{H}_{t+1} \right) \quad (17)$$

for all t , and

$$d \log \tilde{H}_t = -\frac{1}{\delta} (d \log r_t^h - (1 - \delta) d \log r_{t-1}^h) + d \log W + dm \quad (18)$$

for all $t > 0$. When $t = 0$,

$$d \log \tilde{H}_0 = -\frac{1}{\delta} d \log r_0^h + \frac{1}{\delta} (d \log W + dm) \quad (19)$$

The solution to these three equations is algebraic, but messy, so I'll introduce some notation:

$$d \log \tilde{H}_t = (AB^{-t} + 1)(1 - \alpha)(dm + d \log W)$$

The solution can be verified to have this form, and A and B are solutions to a messy quadratic equation. Note that there is an increase in both the steady-state and along the transition

path, as the stock of housing rises to the new steady-state. Note that Cobb-Douglas implies that the increase in labor demand for housing is simply $(AB^{-t} + 1)(dm + d \log W)$. Similarly, since $c^{NT} = \psi \frac{R}{1+R} W$, the increase in labor demand for non-tradable employment is $dm + d \log W$ for all t .

The effect on e can be broken down into three simple terms:

$$d \log e_t = \underbrace{\frac{D_h}{we} AB^{-t} (dm + d \log W)}_{\text{construction boom}} + \underbrace{\frac{D_h + D_c}{we} (dm + d \log W)}_{\text{increased steady-state demand}} - \underbrace{dm}_{\text{labor supply increase}}$$

where A and B are functions of α , δ , and R that represent the size and duration of the construction boom, and W is permanent wealth. As you can see, the key determinants of the sign of $d \log e$ are the size of the construction and non-tradable sectors, and the effect of migration on permanent wealth. The construction boom is temporary in this model, as the stock of housing converges to its new steady-state. But there is also a steady-state increase in labor demand which might dominate the increase in labor supply because of the effect on housing wealth.

The effect of migration on $d \log W$ can also be worked out analytically.

$$d \log W = \underbrace{\frac{(1 - \delta)p^h h}{W} \alpha (A + 1) (dm + d \log W)}_{\text{house price appreciation}} + \underbrace{\frac{w D_h}{W} \left(\frac{A(1 + R)}{1 + R - B} \right) (dm + d \log W)}_{\text{extra transition income}} + \underbrace{\frac{1 + R}{R} \frac{w(D_h + D_c)}{W} (dm + d \log W)}_{\text{extra steady-state income}} - \underbrace{\frac{1 + R}{R} \frac{y}{W} dm}_{\text{labor supply increase}}$$

Intuitively, there are four forces for how migration effects permanent wealth. The first is the effect on house prices, the second is the extra construction income from the initial period, the third is the change in steady-state income, which could be positive or negative, and the fourth is the negative effect from extra labor supply.

Proposition 2. *Under assumptions A.1 to A.4, if $\delta > 0$, there exists a ϕ^* such that for all*

$\phi > \phi^*$, a migration shock causes an increase in the employment rate, $\frac{d \log e_0}{dm} > 0$.

Proof. All the terms of the equation above can be rewritten in terms of paramaters.

$$\begin{aligned}
 d \log W = & \underbrace{(1 - \delta) \frac{1 + R}{R + \delta} \phi R \alpha (A + 1) (dm + d \log W)}_{\text{house price appreciation}} + \underbrace{\gamma \frac{R \delta}{R + \delta} \phi \left(\frac{A(1 + R)}{1 + R - B} \right) (dm + d \log W)}_{\text{extra transition income}} \\
 & + \underbrace{\left(\psi + \gamma \frac{(1 + R) \delta}{R + \delta} \phi \right) (dm + d \log W)}_{\text{extra steady-state income}} - \underbrace{\left(1 - \frac{1 + R}{R + \delta} \phi R \right) dm}_{\text{labor supply increase}}
 \end{aligned}$$

As ϕ grows, the coefficients on the first three terms grow, while the last one shrinks. Hence, $d \log W$ grows without bound.

It is also possible to express $d \log e_t$ as a function of parameters which are increasing in ϕ . and $d \log W$. Hence, $d \log e_t$ must also grow with ϕ and also without bound. Therefore, there exists some cutoff such that for $\phi > \phi^*$, $\frac{d \log e_t}{dm}$ is positive. \square

This proposition simply states that, under these assumptions, if you like housing enough, then migration will improve the labor market. It is an important proof of concept, that with housing, it is possible to have an increase in labor market conditions in response to migration. However, this is not proof that housing is enough to have this increase; that remains an empirical question.

Proposition 3. *Under assumptions A.1 to A.4, if $\gamma = 0$, then $\frac{d \log e_t}{dm}$ is decreasing in α .*

Proof. In this case $D_h = 0$, so it suffices to show that $d \log W$ is decreasing in α . Specifically it must be shown that A is decreasing in α . A and B solve the following equations (along with C):

$$\begin{aligned}
 \delta A + C &= 1 \\
 \delta A + C - \frac{C}{B} &= \delta \\
 C &= \frac{1 + R}{R + \delta} \frac{\alpha}{1 - \alpha} \left(A - \frac{1 - \delta}{1 + R} AB \right)
 \end{aligned}$$

Consider $\alpha^* > \alpha$ and let A^*, B^* , and C^* solve that equation. Suppose $A^* \geq A$. By the first equation $C^* \leq C$. Subtracting equation 1 from equation 2 implies $B^* \leq B$. But then the right hand side of the third equation is strictly larger while the left hand side is strictly smaller. This is a contradiction. Therefore, $A^* < A$. \square

This proposition abstracts from the construction channel to focus purely on the house price channel. It states that places with less elastic housing supplies (higher α 's) have a stronger housing price channel.

In this model specifically, the strength of the construction channel is decreasing in α . The reason for this is because of the log-specification. There are two forces at work: a higher α implies that each additional housing unit requires more workers because you have to substitute from land to other inputs. However, a higher α also means that house prices increase by more, lowering demand for housing. In the log-model, the demand force dominates, but this would not be true in a model in which housing demand was less elastic.

So while the effect of α on the construction channel is ambiguous, the effect on the house price channel is unambiguously positive. It is an empirical question how much the elasticity changes.

B Investigating Complementarity

Much of the immigration literature focuses on whether migrants are substitutes or complements with native workers, with implications that substitutes' productivity will decline, while complements' productivity increases. In contrast, in section 2, I assume everyone is a perfect substitute. In this appendix, I investigate whether a complementarity story might be explaining some of the results.

The first step is to determine how skilled migrants are. A reasonable proxy might be income, which I plotted in Figure 24. Most people who move over 100 miles make, on average, 700 dollars more per year than those who stay in the same MSA. The 100 miles cutoff is

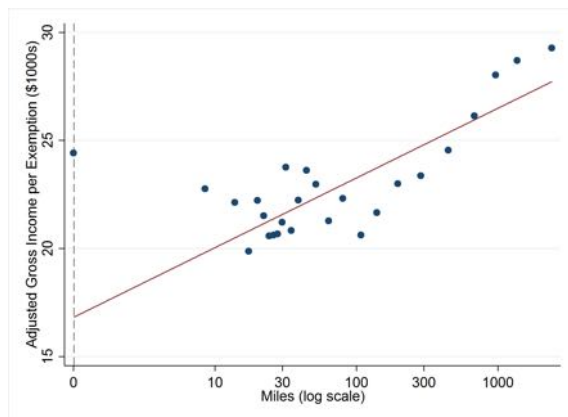


Figure 24: The correlation between Adjusted Gross Income and distance moved, conditional on moving counties.

relevant because that is what I use to construct my shock. From Molloy et al. (2011), we also know they also tend to be younger and more educated. Therefore, a complementarity story would suggest that higher-skilled workers’ labor markets would get worse, while the lower-skilled workers would benefit.

One way to investigate this is to use the occupational employment statistics dataset on the wage distribution. This data starts in 2001, so does not cover my complete dataset. It also uses a different definition of MSA for the first few years of data. I run the same regressions as I do throughout the rest of the paper, but using the 10th, 25th, 50th, 75th, and 90th percentiles of the hourly wage distribution as my independent variable. I plot the results in Figure 25.

The migration shock increases the 10th percentile of workers’ hourly wages, consistent with a complementarity story. However, there is no negative effect on high wage earners. Perhaps it is only because it does not come through in the noisy data.

Even with the result on the 10th percentile, I do not believe that skill-complementarity is driving my results. Although my theory does not speak directly to this, I should note that many of the 10th percentile workers work in non-tradable sectors, specifically “food preparation and serving-related occupations” or “sales and related occupations.” The median wage in these occupations closely tracks the aggregate 10th percentile wage. So perhaps

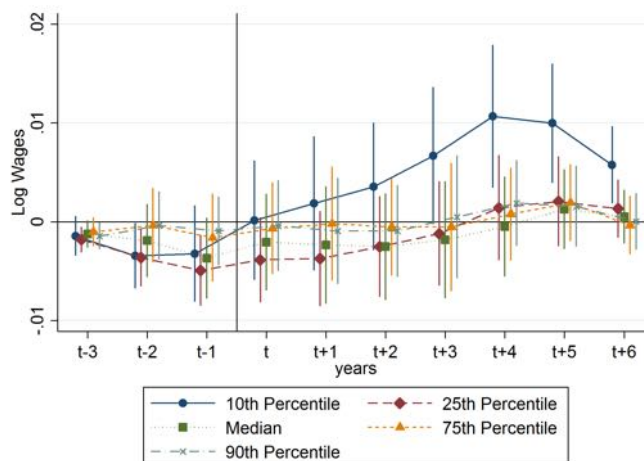


Figure 25: The effect of a one percent migration shock on the distribution of wages

the increase in non-tradable demand, rather than skill-complementarity is driving the wage result.

My results do not depend on the cut-off that I use to construct my instrument. In general, the further the cutoff, the higher-skilled is the migrant. See Figure 24. So under a complementarity story, you might expect the further cut-offs to have larger effects. But there are none.

In addition, the temporary nature of the effects suggests skill-complementarity is not the main force. It would last for a longer amount of time, whereas the channels I highlight, especially the construction channel, are much more temporary in nature.

Furthermore, a natural implication of this model is that the benefits would accrue more in less highly-skilled communities, assuming migrants are roughly the same skill mix. In Figure 26, I show the opposite is the case; MSAs with a higher percentage of college-educated people (as measured by the ACS in 1990, where I consider anyone with 4 or more years of college to be college-educated) have a larger effect from migration than those without. This result is robust to using any years of college education rather than requiring four. One thing to note is that the college share is negatively correlated to the housing supply elasticity. Controlling for that reduces the difference between the two lines.

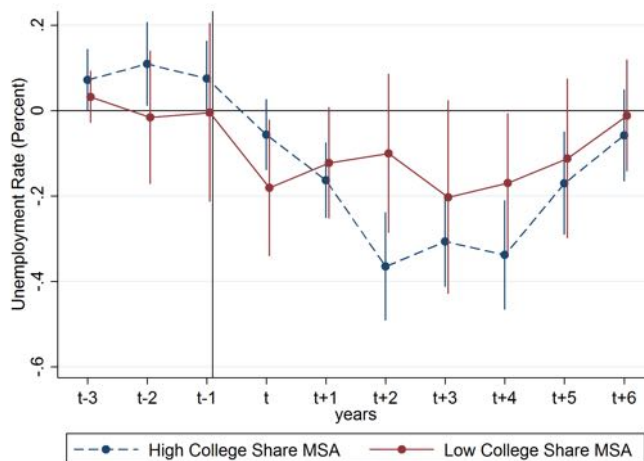


Figure 26: The effect of a one percent migration shock on the unemployment rate, split by college share. Number of MSAs: 158/158

The complementarity story does not have a natural prediction for tradable versus non-tradable goods, high and low elasticity of housing, and would likely make the opposite prediction on timing, such as Ottaviano and Peri (2012).

C Similarities between high-migration city pairs

In this appendix, I show that migrants move between cities that are similar on two observable dimensions: location and industrial composition. To do this, I regress migration on measures of their similarity. The regression establishes that people do move between places that are similar, suggesting that places between which people move are likely to experience similar shocks.

In column (1), I estimate a gravity-like relationship between migration on the distance between any two CBSAs using the specification below.⁴⁷ In column (2), I show the same result for MSAs.

$$\log m_{i \rightarrow j, t_0} = \beta \log \text{distance}_{ij} + \alpha_i + \gamma_j + \epsilon_{ij}$$

⁴⁷There is definitely some misspecification here because the data is censored below by requiring ten tax returns. In fact, CBSAs that are further away are much more likely to be censored, suggesting the true relationship is even stronger than this relationship suggests.

Table 3: Migration Network

VARIABLES	(1)	(2)	(3)	(4)
	Log Migration	Log Migration	Log Migration	Log Migration
Log Distance	-1.684*** (0.032)	-1.627*** (0.037)		
Industry Similarity			2.508*** (0.267)	3.215*** (0.430)
Observations	57,401	38,086	57,401	38,086
Origin and Destination Fixed Effects	YES	YES	YES	YES
Flexible Distance Controls	–	–	YES	YES
Unit	CBSA	MSA	CBSA	MSA

*** p<.01, ** p<.05, * p<.1

Standard errors clustered by from and to MSAs/CBSAs.

Another piece of evidence is that migration is higher between MSAs with similar industries. In columns (3) and (4), I control for a quintic in log-distance, and run the regression on an industry similarity index, using 2-digit SIC codes from 1990. I construct the vector of employment in each of those industries, and use the following formula:

$$\text{Industry Similarity}_{ij} = \frac{v_i \cdot v_j}{\|v_i\| \|v_j\|}$$

where v_i is the vector of employment by sector in MSA i , and $\|\cdot\|$ is the Euclidean norm.

The specification is

$$\log m_{i \rightarrow j, t_0} = \beta \text{Industry Similarity}_{ij} + P^5(\log \text{distance}_{ij}) + \alpha_i + \gamma_j + \epsilon_{ij}$$

There is a strongly positive relationship between industry similarity and migration, even conditional on distance.

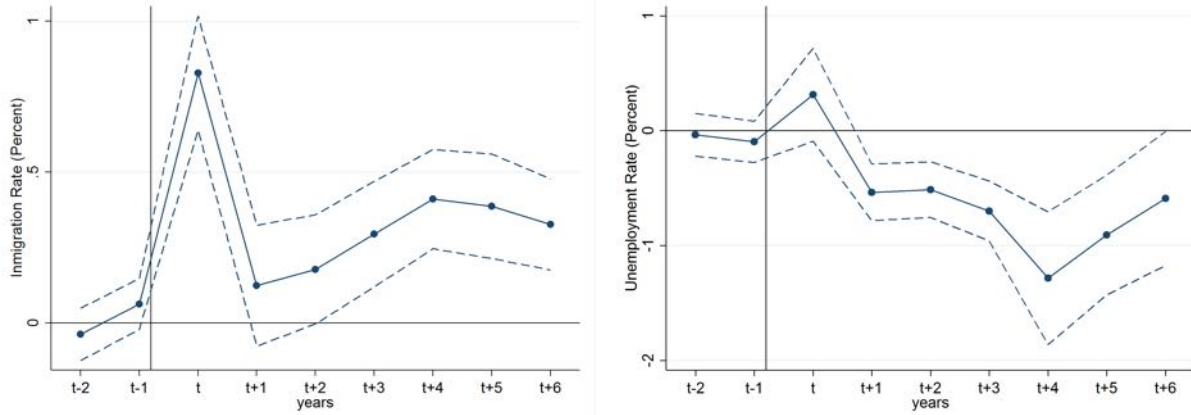


Figure 27: The effect of outmigration from Katrina hit areas in 2005, with 95 percent confidence intervals. Errors clustered by state. Number of MSAs: 381.

D Robustness

Hurricane Katrina

A major source of variation in my data is from Hurricane Katrina, where many people from New Orleans were displaced. This event has been used as a natural experiment to investigate the economic effects of migration on receiving cities, often Houston (see Gagnon and Lopez-Salido, 2014; McIntosh, 2008; De Silva, McComb, Moh, Schiller, and Vargas, 2010). Here, I show that my results are robust to using only this variation. Figure 27 uses only outflows from the eight counties hit hardest by Katrina: Cameron, Plaquemines, Jefferson, St. Bernard, and Orleans in Louisiana; and Hancock and Harrison in Mississippi. I also only use the outflows from 2005.⁴⁸ The instrument for inflows to other cities is based on the 1990-1994 patterns used throughout the rest of the paper. On the left-hand side, I show that the instrument did a good job of predicting inflows. On the right-hand side, I show this was associated with a decline in the unemployment rate. The results are different in the first period, but then consistent with the rest of the paper. Initially, the unemployment rate is predicted to rise, but then falls in the next year, and remains low. The initial rise is

⁴⁸Only using this year necessitates only a two-year lead instead of three, because my data runs only through 2013.

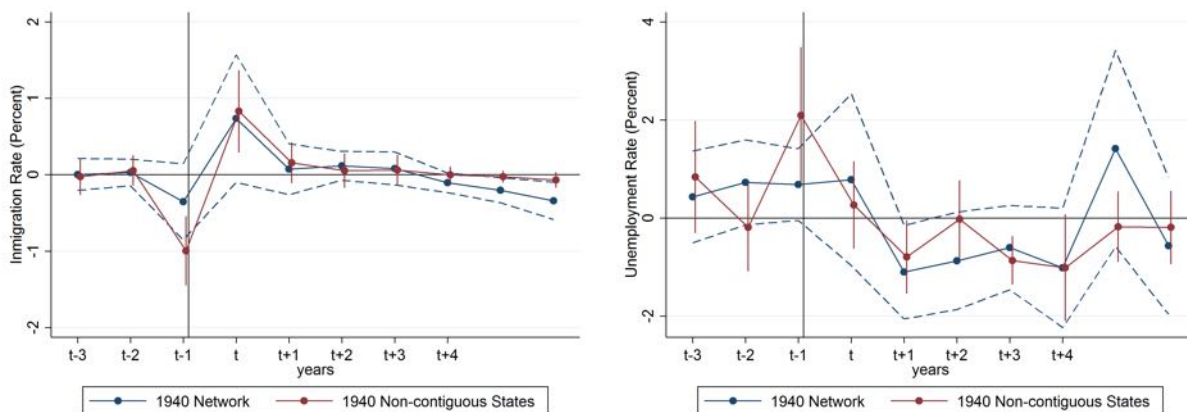


Figure 28: The effect of immigration shocks using the 1940 network, with 95 percent confidence intervals. Errors clustered by state. Number of MSAs: 381.

inconsistent but not surprising: the displaced migrants might be less prepared to find work than an average migrant, and so mechanically raise the unemployment rate; or perhaps 100 miles is not sufficient to rule out direct effects of the hurricane on these other MSAs.

1940 Migration Network

In Figure 28, I construct the migration network using the 1940 Census rather than the pre-period of 1990-1993. I use the 1 percent sample, and construct the migration network based on the reported state of residence five years previously.⁴⁹ The main downside to this approach is that the Census only records the state from which you migrated, and not the county. Hence, the instrument for immigration becomes a significantly worse predictor. Rather than using a cut-off of 100 miles or 500 miles, I only use migration flows between other states, or non-contiguous states. If an MSA spans multiple states, I do not use any of those states (or contiguous ones). The results for the effect on unemployment are largely similar. The response is much noisier, as one would expect from not being able to use county-level data to construct the shocks. The effects in $t - 1$ using the non-contiguous states is also somewhat concerning.

⁴⁹I accessed this data from the Integrated Public Use Microdata Series, (see Ruggles, Genadek, Goeken, Grover, and Sobek, 2015).

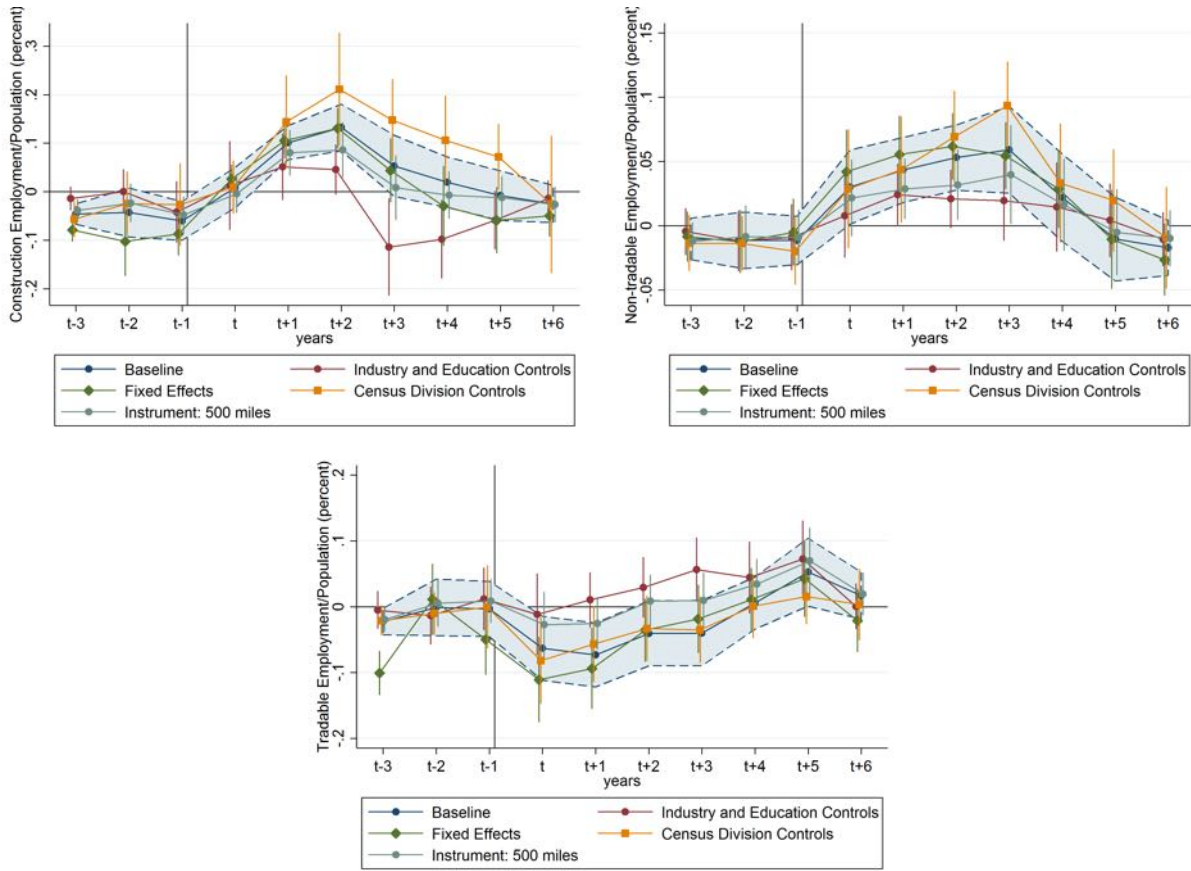


Figure 29: The effect of an immigration shock equal to one percent of the CBSA’s population, with 95 percent confidence intervals. Errors clustered at the CBSA level. Number of CBSAs: 911.

Robustness of Employment Composition

In Figure 29, I show the robustness of the employment composition effects, using similar controls as from section 3.4. The same general patterns emerge as in the main body of the paper: sizable increases in construction and non-tradable employment, with a decrease or no effect on tradable employment. The results are least robust to controlling for the industry and educational controls, but the point-estimates follow the same general patterns.

Robustness of Housing Permits and Prices

Figure 30 shows the robustness of the increase in house prices and permits. The results are largely the same, with house prices and permits increasing, except for using industry

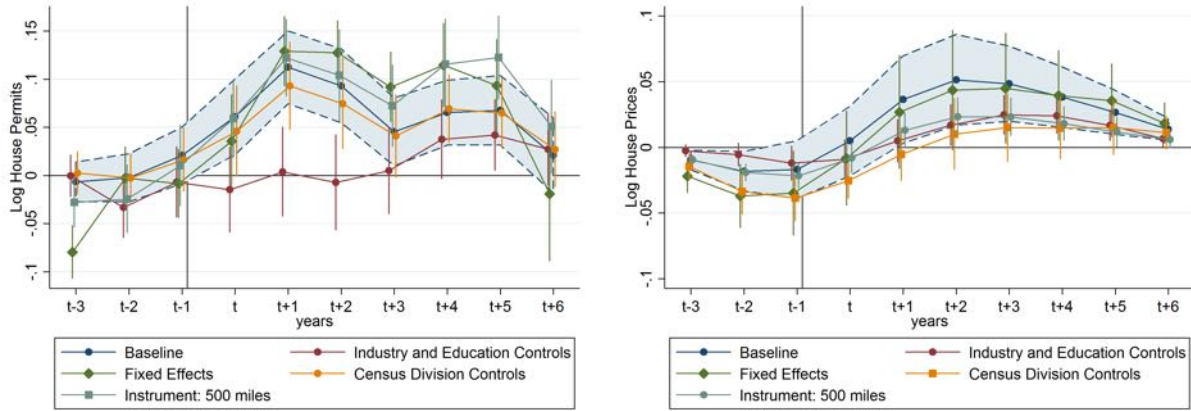


Figure 30: The effect of an immigration shock equal to one percent of the MSA’s population, with 95 percent confidence intervals. Errors clustered by state. Number of MSAs: 381.

and education controls, in which case the permits are not statistically different from zero. Interestingly, using the 500 mile shocks suggests a larger elasticity because permits increase by more, and prices increase by less.