The Migration Accelerator: Labor Mobility, Housing, and Demand

By Greg Howard *

What is the role of migration in regional evolutions? I document that within-U.S. migration causes a reduction in the unemployment rate of the receiving city, over several years. To establish this causal effect, I construct an instrument using outmigration of other places and predict its destination from historical patterns. The decline in unemployment is due to housing. Housing is durable, so increased demand causes a surge of new houses and construction jobs. Additionally, migrants' housing demand raises prices, increasing borrowing and non-tradable employment. This finding implies the endogenous response of migration amplifies local labor demand shocks by about a third.
JEL: E22, R21, R23

Typical models suggest that domestic migration mitigates demand shocks in local labor markets. For example, this result forms the basis for the criteria of high labor mobility in optimal currency areas (Mundell, 1961). Many studies have noted the high labor mobility rate in the United States, and argued that this makes it a more suitable currency union than, for example, the euro area. However, the literature has overlooked critical features of the housing market which, in theory, could overturn the result.

There are two relevant empirical questions. One is whether labor moves to booming markets. Many studies, most famously Blanchard and Katz (1992), answer this question affirmatively. But also central is the causal effect domestic migration has on the local labor market. I estimate that domestic inmigration to a U.S. city lowers the local unemployment rate substantially and for several years. In total, a one percent increase in the population due to inmigration has a cumulative effect of -2.1 percentage-point-years on the unemployment rate.

* Department of Economics, University of Illinois at Urbana-Champaign, 10 David Kinley Hall, 1407 W. Gregory Drive, Urbana, IL 61801 (email: glhoward@illinois.edu). 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 two anonymous referees, 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, the Federal Reserve Bank of Minneapolis, the University of North Carolina, the University of Illinois, Southern Methodist University, Georgetown University, the International Monetary Fund, Society for Economic Dynamics Conference, and the Housing Urban Labor Macro Conference. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 136625. 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.

The two results together imply migration actually amplifies local shocks in the short-term.

The driving force behind this result, which is largely absent from previous analysis, is that migration causes a boom in housing markets.[1] I show there are two housing channels that play a role in migration's effect on unemployment. The first is a construction channel, which occurs because housing is a durable good and is produced using local labor. Hence, in the short-run, housing investment and construction employment respond more than proportionally to an increase in population.[2] The second is a house price channel, which occurs because consumption responds to house price changes, a well-documented effect (Campbell and Cocco, 2007; Attanasio, Leicester and Wakefield, 2011; Mian, Rao and Sufi, 2013; Agarwal et al., 2015; Ströbel and Vavra, 2019; Kaplan, Mitman and Violante, 2016).

The primary challenge to estimating the effect of domestic migration is reverse causality; people choose to move to prosperous cities, causing a negative correlation between the inmigration rate and unemployment. To address this concern, I construct an instrument similar to Altonji and Card (1991), by estimating inmigration based on the annual outmigration of other places in the U.S. which historically have sent people to the city of interest. This effectively isolates the push-component of migration, excluding the pulls of the receiving city. In constructing the instrument, I do not use any migration to or from places within 100 miles of the city of interest, which largely rules out bias from reverse causality.[3]

The instrument directly addresses reverse causality, but a side benefit is that most omitted-variable biases switch signs and imply that my estimates are, if anything, biased to zero. This is because people tend to move between cities that get similar economic shocks. For example, an industry-level shock tends to affect both the sending and receiving city, raising outmigration in the sending city and unemployment in the receiving city at the same time, a bias that goes in the opposite direction of my result.[4] In addition, I show the robustness of my results to detailed industry, education, and geographic controls, as well as with specifications of my instrument that would address these biases in different ways.

[1]Many studies have examined migration's effects on the housing market (Saiz, 2003; Greulich, Quigley and Raphael, 2004; Ottaviano and Peri, 2006; Saiz, 2007; Gonzalez and Ortega, 2013; Coleman and Karagedikli, 2018), but have not followed through to the effects on the labor market that I highlight in this paper. Kochin (1996) tells a story similar to the construction channel in this paper, with anecdotal evidence from a few historical episodes, including the price of cows in colonial Massachusetts.

[2]This is the prediction of a traditional macro model with capital: if population increases, capital will increase proportionally in the long-run. But because only a fraction of capital is depreciating in each period, the effect on investment is much larger than proportional. A similar prediction about government spending is explored in Baxter and King (1993).

[3]One remaining bias is that a migrant may be dissuaded from moving to the city of interest, and instead choose to move somewhere else, or vice versa. In addition to being quantitatively negligible, this bias would be of the opposite sign, suggesting my results are a lower bound.

[4]There is one exception to this direction of bias. If two cities with similar industries are competing, then a decline in city-specific productivity may be beneficial to the other city. However, as I show, the effects on unemployment are concentrated in non-tradable and construction industries, and are absent in tradable industries, suggesting this bias is of no concern.

Because my results are robust, these biases are likely small, but even if there were unobservables creating bias, they would only suggest that my results understate how much of a boom migration causes.

Another feature of this instrument is that I can decompose my estimates into a weighted average of regressions on instruments based on only one origin area, similar to Goldsmith-Pinkham, Sorkin and Swift (2018). Reassuringly, the two biggest shocks, each accounting for just under 2 percent of the variation, are New Orleans during Hurricane Katrina and Miami during Hurricane Irene. I show the results are consistent using only variation induced by Katrina.[5]

I focus on inmigration because my empirical strategy for predicting inmigration does not work to predict outmigration in the data, an admitted drawback of this paper.[6] However, I still believe my results argue for migration as an accelerator and a destabilizing force across space, for two reasons. First, in Section III, I show the majority of the net migration response to labor demand shocks is through inmigration, not outmigration.[7] So to better understand how migration changes the effects of these shocks, it is of primary interest to understand the effects of inmigration. Second, because of the housing channels I identify, the effects of outmigration and inmigration should be locally symmetric.[8] Outmigration should reduce housing demand, and inmigration should increase it. On the timescales that I consider, it seems unlikely they have large asymmetric effects.[9]

To construct my instrument and to measure migration, I use the IRS migration data. This allows me to measure migration and my instrument at a yearly frequency, and to estimate the impulse responses of migration. One of the key benefits of this data is that I can control for lags of the shift-share instrument, separately identifying short-term and long-term effects. The ability to trace out impulse responses is key to uncovering the housing channels that drive the results. With decadal data such as the Census, this would be quite difficult, a point made in Jaeger, Ruist and Stuhler (2018).[10] In addition, the panel data allows me to use city fixed effects, which soak up many observable characteristics. Finally, the panel data also lets me check for pre-trends.

[5]Hurricane Irene happens quite close to the end of the sample.

[6]A possible explanation for this would be if people made migration decisions sequentially: deciding whether to move and then choosing where to move. Such decision-making would allow me to construct an instrument for inmigration using this strategy, but the reverse would not be relevant for outmigration.

[7]This is consistent with Monras (2015*a*), which uses different shocks and data, and Coen-Pirani (2010), which notes that inmigration is much more volatile than outmigration. In a different setting, Long and Siu (2018) find that during the Dust Bowl era, the fall in net migration was also due to inmigration and not outmigration.

[8]Because of non-linearities in the housing market, there may be non-local asymmetries. I explore this in Section II.

[9]In the model I present in Section IV, there is no difference at all between an inmigrant and an outmigrant. In a search model, the effect on tightness, prices, and quantities would be different for someone leaving a city and destroying a match than for someone moving into a city. However, with the smallest unit of time being one year, this seems unlikely to be important.

[10]Jaeger, Ruist and Stuhler (2018) shows that the persistence of shift-share shocks in the international contexts leads to conflating short and long-term effects when using similar methodologies. My methodology allows me to estimate the impulse responses, which separately show short and long term effects.

Using this strategy, I show that inmigration to a city substantially reduces the unemployment rate in the short-term. An inmigration shock equal to one percent of the city's population causes a decline in the unemployment rate, by about a tenth of a percentage point in the first year, and then by about one half of a percentage point in the subsequent four years.[11] After six years, the unemployment rate returns to baseline. In addition, I show these results are robust to a variety of controls and specifications, and I show consistent results for the employment-to-population ratio and unemployment benefits.

I find that the result on the unemployment rate is driven substantially by the two housing channels. Specifically, I use the same instrument and show that its impact on house prices, housing construction, construction employment, mortgages, and non-tradable employment are all positive and large. I also test two dimensions of cross-city heterogeneity: the effect is stronger in cities with inelastic housing supplies and in growing cities.[12]

In the next part of the paper, I quantify the economic impact of the previous results through a counterfactual exercise. 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 the difference the "migration accelerator." I find that migration amplifies the effect on the unemployment rate by about a third.

Lastly, I provide a stylized framework that rationalizes the empirical results and frames a discussion of welfare and policy. The model provides a contrast to Farhi and Werning (2014) on the role of labor mobility in currency unions. By adding housing, the result that labor mobility is weakly helpful because of aggregate demand externalities is overturned. Second, the model clarifies why migration leading to lower unemployment does not induce unboundedly large inmigration.

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 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, 2015*b*).[13] Even

[11]These effects are too big to be driven by differential unemployment rates of the migrants themselves, as I argue in Section I.D. Since most people in the labor force are employed, adding another employed person has only a small mechanical effect.

[12]Other theories could explain my main result (that the unemployment rate falls after an inmigration 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 each of these facts. I expand on this in Appendix E.

[13]Hong and McLaren (2015) report a sizable increase in the employment of cities that receive immi-

different studies of the Mariel boatlift have divergent results. Card (1990) concluded that wages of comparable workers in Miami were largely unaffected, while Borjas (2017) found that they dropped dramatically. Other methodologies have also produced a range of results, as summarized by Dustmann, Schönberg and Stuhler (2016).[14] 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 inmigration.

The only other paper of which I am aware that uses a similar strategy for domestic migration is Boustan, Fishback and Kantor (2010), who study the effects of internal migration during the Great Depression. They do not have previous migratory patterns and so use distance to negative shocks as an instrument for inmigration, finding negative effects on some measures of employment.

Another strand of the immigration literature looks specifically at housing. Saiz (2003, 2007), Gonzalez and Ortega (2013), Greulich, Quigley and Raphael (2004) and Ottaviano and Peri (2006) find positive effects of immigration on housing costs.[15] 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; Monras, 2015*a*; Cadena and Kovak, 2016). All of these papers find evidence that labor mobility responds to the conditions of local labor markets. However, in contrast to much of this literature, I show that migration is not an mitigating force; rather the migration exacerbates the initial shocks at least in the short-run. To be precise, Blanchard and Katz (1992) and many of the subsequent papers show that migration responds to local labor demand shocks, and that the unemployment rate effects of the local labor demand shocks are short-lived. They interpret this as population movements equilibrating the unemployment rate, but this need not be the case if population movements are also affecting the number of jobs. In recent work, Amior and Manning (2018) show a very persistent effect of local labor demand shocks on employment, despite large migration responses. They frame this as a race between population and employment. My work suggests that the population changes are causing changes in employment, and that those changes are in the short-run larger than the changes in population.

The rest of my paper is organized as follows. Section I outlines my empirical

grants, 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.

[14]In Appendix D, I show that construction and retail employment rose in Miami, compared to control groups, while manufacturing employment did not.

[15]Saiz 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.

strategy and shows that inmigration has an expansionary effect on local labor markets in the United States. Section II presents the evidence in favor of the construction and house price channels. Section III quantifies the size of the migration accelerator. Section IV presents a framework for understanding why a model with housing could allow for inmigration to be expansionary, and discusses policy implications.

I. The Expansionary Effect of Migration

In this section, I test the causal effects of inmigration on unemployment. I use previous migration patterns and concurrent outmigration to construct a plausibly exogenous shock to inmigration. I find that inmigration causes unemployment to fall.

A. 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.[16]

The aggregate patterns of migration in the IRS data are similar to other datasets. Gross migration is procyclical. It has trended downward in recent decades, although the decline in migration in the IRS data is smaller in magnitude than what has been documented using CPS data (Molloy, Smith and Wozniak, 2014). There is significant variation in both inmigration and outmigration rates for different MSAs. Much of the variation is across cities, which makes it important to control for permanent differences across cities using a city fixed-effect. But the within-city standard deviation of migration is still 0.5 percent, meaning there is still economically significant variation.

[16] There 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, 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 inmigration 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 missed in the data. So potentially, migrants are under-counted 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 under-counted. 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.[17]

For robustness and to explore the housing channel, I also use data from a variety of other sources. Industry employment data comes from the County Business Patterns. I sort these into categories based on the decomposition of Mian, Rao and Sufi (2013). House price data comes from the Federal Housing Finance Agency. Housing starts data 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 as well as data coverage. There are more details in Appendix A.

B. Identifying Inmigration Shocks

The goal of this section is to estimate the effect of inmigration to an MSA on the MSA's unemployment rate. Isolating the causal relationship requires plausibly exogenous shocks to inmigration because inmigration 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 inmigration and unemployment. My other major concern is omitted variable bias: during my sample, an increase in housing prices lowered unemployment (Mian, Rao and Sufi, 2013). If it also affected the inmigration rate because housing is less affordable, there would be omitted variable bias. This would bias the OLS upward,

[17]Some 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	MSAs	Mean		Standard Deviation		
		1995-2013	2008-2010	Total	Between	Within
Inmigration Rate (Percent)	381	3.3	3.0	1.5	1.5	0.5
Outmigration Rate (Percent)	381	3.2	2.9	1.4	1.3	0.5
Unemployment Rate (Percent)	381	6.1	8.2	2.8	2.0	2.0
House Price Growth (Log-Difference)	381[a]	2.8	-4.1	5.9	0.8	5.9
House Permits Growth (Log-Difference)	356[a]	-3.2	-25.9	33.6	5.5	33.3

Note: [a]Unbalanced panel. Over 80 percent of MSAs report values in all years.

Source: Internal Revenue Service (1990-2013), Bureau of Labor Statistics (1990-2013), Federal Housing Finance Agency (1975-2015), United States Census Bureau (1995-2014), author's calculations

because it induces a positive correlation between inmigration and unemployment. These two concerns are not meant to be an exhaustive list, but are both likely to be major sources of bias.[18] In this section, I identify shocks to inmigration 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 inmigration 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) and many papers since, but tailored to suit the domestic migration setting.[19]

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 inmigration 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

[18]In Appendix B, I show the same results using OLS. There does seem to be evidence of bias in favor of migrants choosing to move to places where unemployment rates are falling.

[19]Beaudry, Green and Sand (2018) 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 higher-frequency variation in order to capture short-term effects.

Boustan, Fishback and Kantor (2010) also use a similar instrument to look at the labor market effects of internal migration during the Great Depression, interacting bad economic outcomes with distance. Their data also does not allow them to look at dynamic 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.

MSA in the historical network, and multiply by the outmigration of the origin county in the concurrent year to places more than 100 miles from the MSA. Then I sum over all counties.

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,[20] and by using a cutoff of 500 miles.

In math, the formula for predicted inmigration is

$$\tilde{z}_{n,t} = \sum_{c \in -n} \frac{m_{c \to n, t_0}}{m_{c \to -n, t_0}} m_{c \to -n, t} \tag{1}$$

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

As a concrete example, suppose I am constructing a prediction for inmigration 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 98.5 percent move at least 100 miles away from Boston. In 2007, 41,942 people moved from Montgomery County to other places 100 miles or more away from Boston. To calculate the predicted inmigration, I would multiply those 41,942 by 1 percent and divide by 98.5 percent to predict that 434.9 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 inmigration to Boston in 2007.

The measure is unlikely to be directly influenced by the economic conditions of the city of interest because it does not use migration to or from areas near that city. But because the patterns of migration are relatively stable, this measure is strongly correlated to the actual inmigration 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.

The identifying assumption behind these results is that the historical destination of migrants from any county is not correlated to changes in unemployment, except through changes in migration from that county, conditional on year and city fixed effects. For example, the assumption would be that places that were historically sent people from New Orleans were not systematically different in the years after 2005, except for the inmigrants from Hurricane Katrina.

One concern for identification is that areas with high mobility between them

[20]For 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.

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 inmigration is. I discuss this bias in more detail in Section I.E.

In fact, because the way the instrument is constructed, I can decompose by main regression into the weighted average of many regressions. In each year, the instrument is the sum of many instruments, one for each county in the U.S.: $z_{n,t} = \sum_c z^c_{n,t}$. The estimated effect is equal to the weighted sum of regressions on these county-specific instruments. Intuitively, counties with bigger changes in outmigration, and counties whose outmigrants pick some cities over others, not in proportion to their population, will induce more variance and have more weight.

The biggest weights belong to county-year pairs in which there were large hurricanes: Orleans Parish, LA during Hurricane Katrina and Miami-Dade County, FL during Hurricane Irene. In Appendix D, I show similar results to my main regression using only the outmigration from Hurricane Katrina.[21]

Orleans Parish, in 2005, gets 1.9 percent of the weight in the regression, and Miami-Dade in 2011 gets 1.9 percent. However, most other county-year pairs do not get too much weight. The top 100 weights account for 33 percent of the total weight, and the top 1000 account for 69 percent. The top 10,000 (out of more than 54,000) account for 99 percent of the weight.

C. *Econometric Specification*

I use a local projection methodology to estimate the effects of migration on unemployment (Jordà, 2005). The main specification is

$$u_{n,t+s} - u_{n,t-1} = \beta_s z_{n,t} + \gamma X_{n,t} + \alpha_n + \alpha_t + \epsilon_{n,t} \tag{2}$$

where s ranges from -4 to 6, α_n and α_t are MSA and year fixed effects, and X_t are control variables. I run this regression separately for each s, and plot the β_s which traces out the impulse response. The controls always include lags of the instrument because a major endogeneity concern of shift-share migration instruments is the autocorrelation of the shock. As pointed out by Jaeger, Ruist and Stuhler (2018), shift-share instruments are fairly persistent, and this is also true in my data. Later, I will show that they have a causal relationship with unemployment. So if I did not control for lags of the instrument, it would lead to omitted variables bias. In the baseline, I control for two lags, $z_{n,t-1}$ and $z_{n,t-2}$, in every specification. Following Autor, Dorn and Hanson (2013), standard errors are clustered at the state-level in order to allow for spatial correlation across MSAs.

[21] I use all eight counties that Hurricane Katrina hit, which together account for 2.6 percent of the total weight. The exercise is quite similar to McIntosh (2008), which finds negative effects on wages and employment in Houston in the first-year after Katrina. I also find a rise in the unemployment rate in the first year, but a large decline afterward. See Figure D6 in Appendix D.

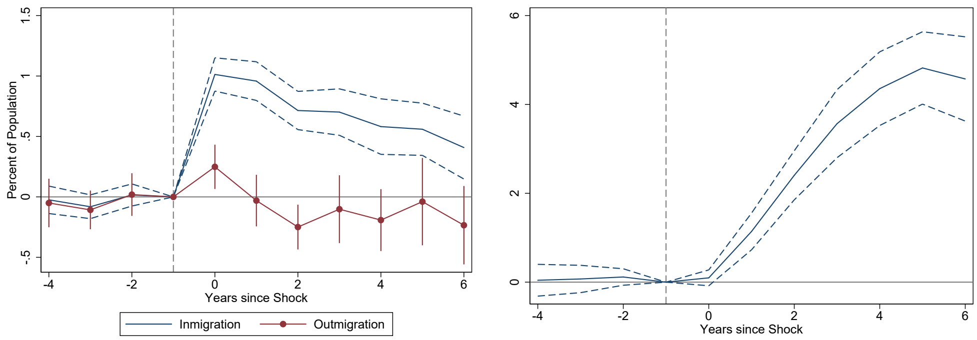

(a) Migration (b) Log Population

Figure 1. : The migration response to the inmigration shock, with 95 percent confidence intervals. Standard errors clustered by state.

The methodology used to construct shocks used up the first four years of migration data and the lags use up the next two, so the sample period is from 1997-2013.[22] I chose to use six lags because the effect on unemployment dissipates after six years. Four leads are used to show a lack of a trend.[23]

Figure 1(a) shows the response of inmigration and outmigration to this shock. For this figure, I run the same specification, but with the migration rate on the left-hand side. Note there is only a small response from outmigration initially. Whatever the cause of the outmigration from historically-connected counties, it is not causing lots of people to move out from the receiving city.

Figure 1(b) shows the response of overall population. After six years, the response has flattened out, with population having increased by 4.57 log-points. In subsequent figures, the magnitude of the coefficients are shrunk by a factor of 4.57, so that the impulse responses can be interpreted as the effect of a migration shock causing a one percent increase in population.

D. *The Effect on Unemployment*

Figure 2 shows the effect of an inmigration shock on the unemployment rate. The blue line, with dashed confidence interval bands, is the estimated effect of the one-percent inmigration shock. In periods $t-3$ to $t-1$, the coefficients are not

[22] One 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 II, I argue that the relationship we see between housing construction and house prices match well with previous estimates from before the housing boom (Poterba, 1984; Topel and Rosen, 1988). Results are robust to splitting the sample to before and after 2001. See appendix D.

[23] The lack of a pre-trend is also helpful in addressing any lingering concerns of reverse causality. If the reader is concerned that good economic conditions are causing an increase in my instrument, it would make sense for unemployment rates to be correlated with future values of the instrument, which it is not.

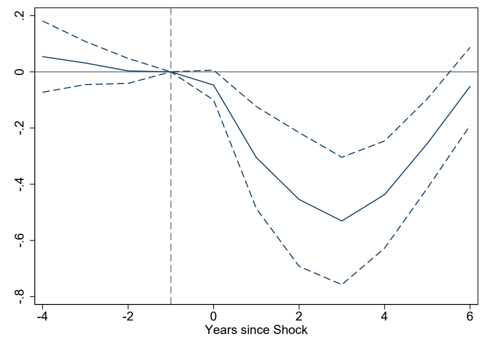

Figure 2. : Unemployment Rate. The response to an inmigration shock, with 95 percent confidence intervals. Errors clustered by state.

significantly different from zero, giving no evidence of a pre-trend. In period t, the period of the shock, the unemployment rate falls by 0.05 percentage points. In period $t+1$, the unemployment rate falls more, to a total effect of 0.3 percentage points, which grows further through $t+2$ and $t+3$, before gradually returning to zero by $t+6$. In total, the inmigration lowers the unemployment rate by 2.1 percentage-point-years.[24]

We can also see the effect in other measures of the labor market. In Figure 3(a), 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 County Business Patterns 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 3(b), I also show that unemployment benefits, as measured by the Bureau of Economic Analysis fell. While not surprising given the results on unemployment, it is constructed from administrative data rather than surveys and so might be more reliable.[25]

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.[26] 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

[24]Appendix B compares these results to ordinary least squares. Without the instrument, there is some evidence of a pre-trend, and the initial effect is stronger.

[25]Given that the LAUS unemployment data is imputed, I also run the regressions on data from the CPS microdata in Appendix D and see similar answers.

[26]It appears from the descriptive data that their adjusted gross incomes are not necessarily that much higher. See for example, Figure E1 in Appendix E.

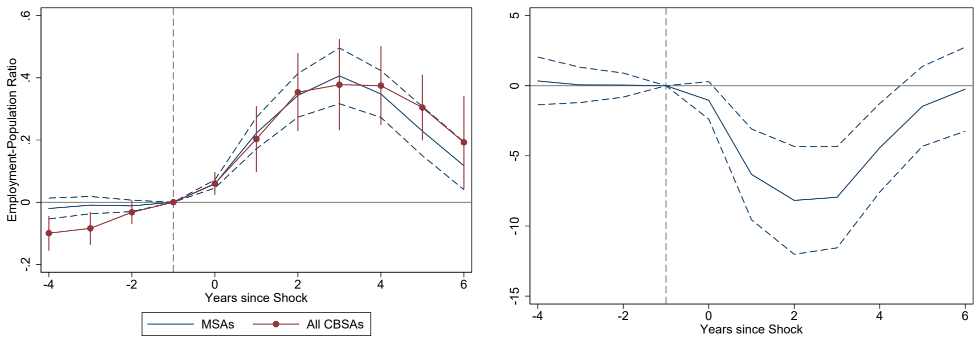

(a) Employment-Population Ratio (b) Log Unemployment Benefits

Figure 3. : The effect of an inmigration shock, with 95 percent confidence intervals. Errors clustered by state.

is a natural bound on the direct effect coming from the employment status of migrants.

If I assume each inmigrant has a job and each outmigrant is unemployed, then I can calculate a bound. In $t+1$, the shock has added about 0.43 percent of the population via inmigration and lost .05 percent of the population due to outmigration. This implies a bound of 0.43 times the original unemployment rate plus 0.05 times one minus the unemployment rate, assuming the labor force participation rate is the same. The average unemployment rate in my sample is 6 percent, implying that the upper bound on this mechanical effect is 0.08 percentage points. In that time period, I estimate the unemployment rate has fallen by 0.31 percentage points, implying most of the effect must have been because of additional jobs for non-migrants, even under these extreme assumptions.

The fall in the level of unemployment benefit, as seen in Figure 3(b), is also indicative that the effect is not driven purely by inmigrants being more likely to have jobs, as that would not change the level of benefits.

E. Threats to Identification

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

Common 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.[27] For example, an oil boom would make Dallas more prosperous, and fewer

[27]It 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. This argument addresses any

people would migrate out. At the same time, it would lower the unemployment rate in Houston. This would bias my regression.

I expect this bias to be small and positive, suggesting that my results are an upper-bound on inmigration'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 III, 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.

In addition, I add controls to address particular concerns, in Appendix D. I show that the results are robust to concerns about industry-specific or skill-specific shocks. Because controlling for the likeliest confounders does not change my estimates meaningfully, the bias from common shocks is probably small. I also show robustness to concerns over spatial characteristics and regression specification by varying the distance cutoff for construction of the instrument, and controlling for Census divisions. Furthermore, I show the results are also not specific to the Great Recession. Nor are they not driven by the choice of 1991-1994 as a pre-period; using the 1940 Census delivers similar results.

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 the receiving 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 Appendix E, 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.

remaining concerns.

II. The Housing Channels

What is the mechanism behind the decline in the unemployment rate? In this section, I show evidence that the inmigration causes a housing boom, and that the labor demand effects from such a boom are large. I also show evidence that the effects differ on two dimensions for which housing should matter: the housing supply elasticity of the city, and whether the city is growing.

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

$$y_{n,t+s} - y_{n,t-1} = \beta_s z_{n,t} + \gamma X_{n,t} + \alpha_n + \alpha_t + \epsilon_{n,t} \tag{3}$$

where y 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, Rao and Sufi (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.

A. *Construction Channel*

The construction channel requires a build-up of new housing, especially in the short-term. In Figure 4, I show that housing permits, from the Census, increase significantly after a migration shock. 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 67 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 2.5 inmigrants.

In the right half of Figure 4, we see an increase in the construction sector. For one percent inmigration, there is a corresponding increase in construction equal to 0.3 percent of the population.

Recall from Figure 3(a) that the total employment-to-population ratio increased by about 0.4 percent in response to the shock, so the construction channel seems to be explaining most of that. In Appendix D, I show the robustness of these results to many of the same checks I discussed in Section I.

B. *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 5, I show that house prices do increase, responding by roughly eight

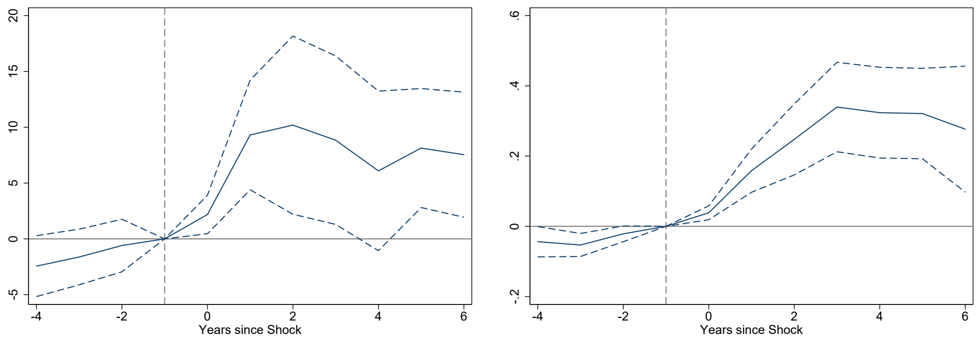

(a) Log Housing Permits (b) Construction Employment per Capita

Figure 4. : The effect of an inmigration shock, with 95 percent confidence interval. Errors clustered by state.

percent in response to one percent inmigration.[28] 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 1.5. This is in line with Poterba (1984), who estimates a housing supply elasticity of between 0.5 and 2.3; and Topel and Rosen (1988) who 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. Finding 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 boom.

In Figure 6, 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 even 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.[29]

On the right is the rise in non-tradable employment, which increases by about 0.08 percentage points. Given a house price rise of about 8 percent, and assuming a consumption-to-house-price elasticity of 0.2 (Berger et al., 2017), we would expect non-tradable consumption to rise by 1 percent. The mean non-tradable-

[28]Given the centrality of the housing market, the reader may be interested in how vacancies or homeownership respond to inmigration. In Appendix D, I show that vacancies decline in response to the shock, which makes sense given the extra demand, and homeownership rises, which also makes sense given that migrants tend to be higher income.

[29]The 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.

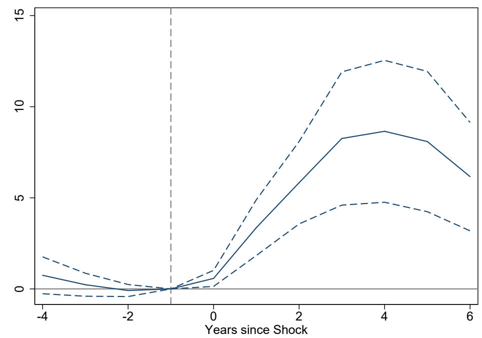

Figure 5. : House prices. The effect of an inmigration shock, with 95 percent confidence interval. Errors clustered by state.

employment-to-population ratio is 8 percent in my data, which would predict a 0.13 percentage point increase in non-tradable employment, slightly higher than what we see in Figure 6.

This increase is a smaller fraction 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. In contrast, I do not find an initial increase in the tradable-employment-to-population ratio, which I show in Appendix E. In fact, the point-estimate declines slightly. After four years, the effect is positive, but is still small compared to the effect on construction employment.[30]

C. 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 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 inmigration shock with the housing supply elasticity from Saiz (2010).[31] This

[30]The reader may be wondering what happens to "other" employment. "Other" employment (not shown) rises. It's contribution to the total increase is comparable to non-tradable goods. This may reflect that some of the "other" category is non-tradable. Because "other" is such a large category, the increase is a much smaller fraction of "other" jobs, and the estimates are fairly noisy, only significantly positive in $t+1$ and $t+2$.

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

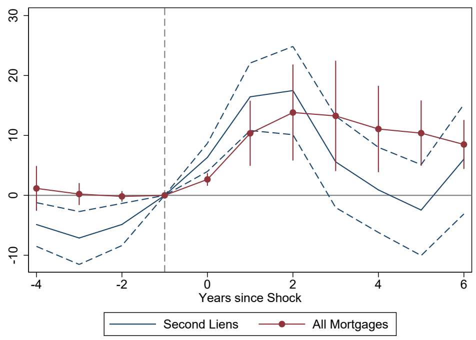

(a) Log Mortgage Originations (Dollars)

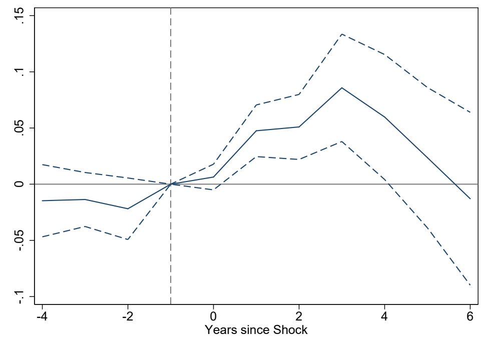

(b) Non-tradable Employment per Capita

Figure 6. : The effect of an inmigration shock, with 95 percent confidence interval. Errors clustered by state.

allows me to see whether the effects of a migration shock are different in areas where we might expect house prices to react more. I run the following regression:

$$y_{n,t+s} - y_{n,t-1} = \beta_s z_{n,t} + \beta_s^* z_{n,t} \times \text{elasticity}_n + \gamma X_{n,t} + \alpha_n + \alpha_t + \epsilon_{n,t} \tag{4}$$

where x is house prices or the unemployment rate. β_s^*, the coefficient of interest, estimates the heterogeneous effect of migration by housing supply elasticity. In this specification, the $X_{n,t}$ include the lags of z, the elasticity, and the interaction of the lags with elasticity.

Figure 7 shows that the effects do differ by housing supply elasticity. On the top left, I show that house prices increase by less in more elastic areas, as expected. On the top right, I show that the unemployment rate falls by less in those same areas. The standard deviation of the Saiz elasticities are 1.44, so these estimates would imply the difference in the size of the effect for cities with one standard deviation difference in elasticities is slightly smaller than the size of the average effect.

The employment difference across elasticities is driven more by the construction employment more than non-tradable employment (Figure 7(d)). A priori, one might have expected bigger differences in non-tradable employment because the house price channel should unambiguosly be stronger in inelastic cities, whereas it is less straightforward whether the construction channel is stronger.[32] In the data, it seems that the construction channel is much stronger in inelastic cities.

[32]For example, if the intensive margin of housing demand was quite elastic, there might be very little construction because people would squeeze in in inelastic cities. Conversely, if the intensive margin of demand is inelastic and the construction sector requires more labor per house in inelastic cities, then the construction effect would be stronger in inelastic cities.

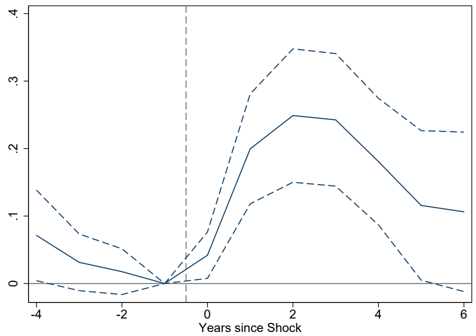

(a) Log House Price

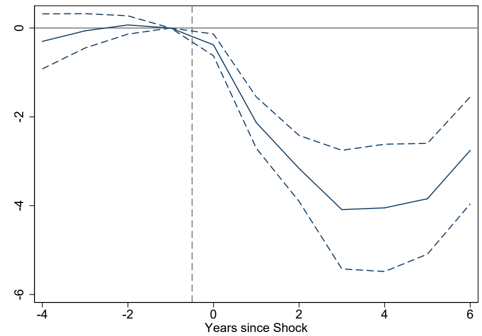

(b) Unemployment Rate

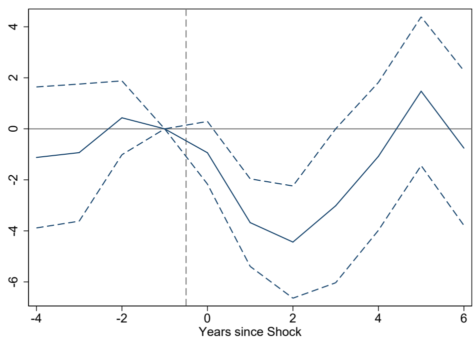

(c) Log Housing Permits

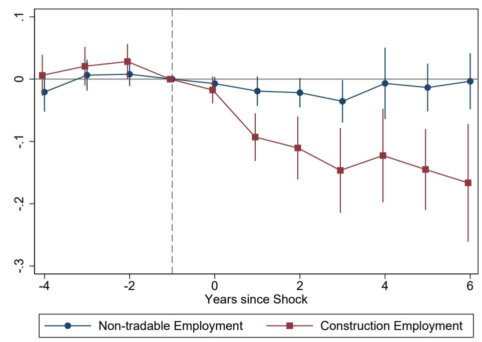

(d) Employment per Capita, by Sector

Figure 7. : The effect of an inmigration shock equal to one percent of the MSA's population interacted with Saiz (2010) housing supply elasticity, with 95 percent confidence intervals. Errors clustered by state.

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 construction. Hence, seeing heterogeneity by the growth rate of an area further supports the hypothesis that housing is central to the story.

To investigate, I use the fact that population increased more in areas of the country with warmer winters (Rappaport, 2007). I test if the marginal effects of migration are stronger in warmer climates. The regression is identical to the one for housing supply elasticity, but uses January temperatures instead.[33] The reason to use January temperatures is that they affect growth rates, but are unlikely to be affected by other shocks that directly affect house prices and unemployment changes.

[33] I use the average of the high of daily January temperatures from 1990 to 2013, available from the

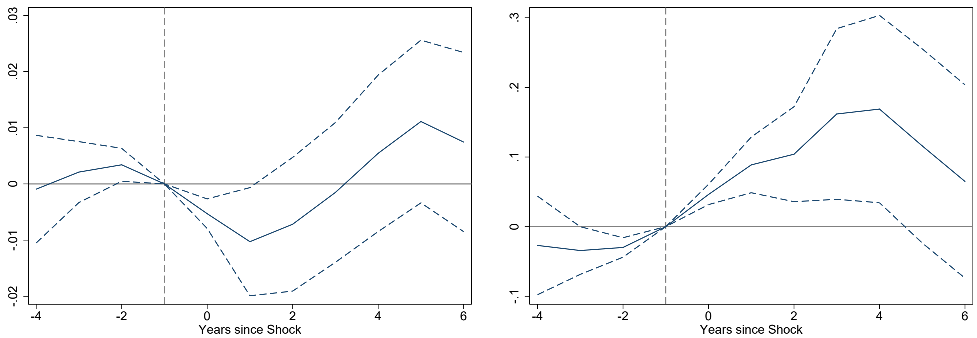

(a) Log House Price (b) Unemployment Rate

Figure 8. : The effect of an inmigration shock interacted with high January temperatures, with 95 percent confidence intervals. Errors clustered by state.

The results are presented in Figure 8. In warmer MSAs, house prices increase more in response to migration and the unemployment rate falls more. The standard deviation of Janurary temperatures is 13 degrees, so a one standard deviation change in the receiving city would add or subtract roughly a third to the baseline effect of inmigration. 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.

III. The Migration Accelerator

One implication of the main result 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.

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.

Center for Disease Control and Prevention.

This section highlights the contrast of this paper to Blanchard and Katz (1992), which argues that migration mitigates the effect of a demand shock on the unemployment rate. Given the results of the previous sections, combined with replicating one of that paper's main empirical findings leads to the opposite result: that migration amplifies the initial shock.

A. *Migration's Response to Labor Demand*

The first step in calculating the accelerator is to estimate the effect that labor demand has on migration. 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 codes before 1998 and three-digit NAICS codes after 1998 to construct the instrument in each year. The formula for the instrument is $\tilde{z}^b_{n,t} = \sum_j s_{j,n,t-1} g_{j,-n,t}$ where $s_{j,n,t-1}$ is the employment share of industry j in MSA 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.

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

$$m_{n,t+s} - m_{n,t-1} = \beta^b_s z^b_{n,t} + \zeta X_{n,t} + \alpha_n + \alpha_t + \epsilon_{n,t} \tag{5}$$

where $m_{n,t}$ is the migration into MSA n at time t, z^b_n is the Bartik instrument in MSA n, and X includes two lags of the Bartik instrument, as well as two lags and the contemporaneous weighted-average of Bartik instruments in cities which send or receive many migrants in that city. This is a very similar specification to the main regressions run in this paper.

The results for both inmigration and outmigration are shown in Figure 9. The effect for outmigration is smaller but significant. The effect on inmigration is about twice the size and lasts for three years.[34]

B. *Accelerator*

To calculate the accelerator, I estimate the expected effect of the migration I found in Figure 9 on unemployment. I compare that to the total effect that

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

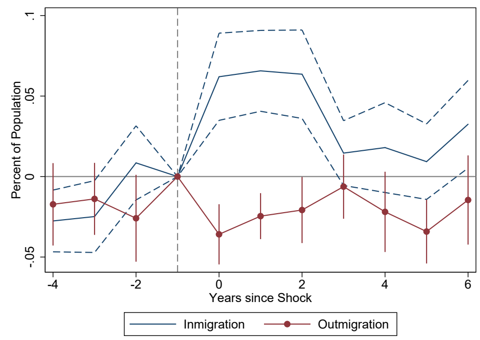

Figure 9. : The effect on migration of a labor demand shock, with 95 percent confidence intervals. Standard errors clustered by state.

Bartik shock has on unemployment.

Using only my estimates from Section I, I cannot estimate the effect of any sequence of inmigration, I can only do it for the sequence I observed in response to the inmigration shock. If Figure 9 looked exactly like the path of inmigration induced by this shock (Figure 1), 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 $t+s$ 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.[35]

With this assumption, I can back out a series of migration shocks that would exactly replicate the path of migration in Figure 9.[36] I then estimate the effect that that series of shocks would have on the unemployment rate, and I call that the "Accelerator."

In math, I back out the shocks by solving the following system of equations

$$\beta_t^b = \sum_{r+s=t;r,s\geq 0} \text{Shock}_r \cdot \beta_s^m \tag{6}$$

[35]Consider 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 inmigration because non-migrants will anticipate the rise in house prices. So there is a similar response in the total number of additional construction jobs, only a change 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.

[36]The biggest shocks are a positive shock in $t = 0$, and a smaller negative shock in $t = 3$ because the effect of the labor demand shock on migration is shorter-lived than the effect of the migration shock.

where β_t^b is the estimated effect of the Bartik shock on inmigration in t periods (Figure 9) and β_s^m is the effect of the migration shock on inmigration (Figure 1). Using those shocks, I then calculate the accelerator as

$$\text{Accelerator}_t = \sum_{r+s=t;r,s\geq 0} \text{Shock}_r \cdot \beta_s \tag{7}$$

where β_s is the effect of the migration shock on unemployment s periods later (Figure 2). To calculate standard errors, I estimate all three regressions jointly and use the delta method.

I show the estimated Accelerator, along with the total effect of the Bartik shock on unemployment in Figure 10. The effects from migration explain a small but significant portion of the unemployment rate's response. At $t = 0$, the response is equal to 10 percent of the total effect, implying the Bartik shock would have had a 10 percent smaller effect sans the migration response. At $t = 1$, the migration effect is about a quarter of the total effect. Cumulatively through $t = 4$, the migration response accounts for approximately fifty percent of the unemployment rate decline, measured in pecent-years.

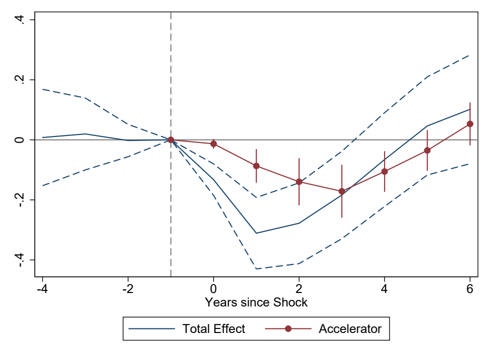

Figure 10. : The response of unemployment to a Bartik shock, and the portion that is due to migration, along with 95 percent confidence intervals. Standard errors clustered by state.

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 about 25 percent smaller at its peak and significantly shorter in duration. This suggest migration amplifies the effects of the shock by about a third, as measured by the largest response.

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. 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.

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.

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:

$$\text{AGI migration rate}_{n,t} = \beta m_{n,t} + \gamma X_{n,t} + \alpha_t + \alpha_n + \epsilon_{n,t} \tag{8}$$

where the AGI migration rate is the total income of all migrants into the CBSA, normalized by the CBSA's population. I then instrument for the migration rate using the migration instrument, or the labor demand instrument. The controls are the same as in the main specifications for the migration and Bartik shocks. I can do a similar thing for the average returns-to-exemptions ratio. I use CBSAs for this exercise in order to have a powerful first-stage when using the Bartik instruments.[37]

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 migration shocks are smaller, though not statistically significantly different. The number of returns are almost identical.

IV. Stylized Model

In this section, I propose a stylized model that is able to rationalize all the facts from the previous empirical sections. There are two main purposes to this model. First, it is a contrast to Farhi and Werning (2014) which argues labor mobility

[37]Using only MSAs has weak instrument issues and wide confidence intervals. The characteristics are not statistically distinguishable from each other, but that is less surprising.

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

VARIABLES	(1) AGI (1000s)	(2) AGI (1000s)	(3) Returns	(4) Returns
Migration	19.71	23.77	0.511	0.507
	(1.727)	(4.722)	(0.0146)	(0.0218)
Instrument	Migration	Labor Demand	Migration	Labor Demand
Kleibergen-Paap F	360.3	61.46	360.3	61.46
CBSAs	917	917	917	917

Standard errors clustered by state

Source: Internal Revenue Service (1990-2013), United States Census Bureau (1989-2013), author's calculations

is weakly beneficial for currency unions due to aggregate demand externalities.[38] By adding housing to the model, that result is overturned. Second, it includes a discussion of the distributional consequences of a migration shock. Because many of the benefits accrue to existing homeowners, this explains why the labor market benefits of the shock do not induce more migration *ad infinitum*.

Setup. — Time is discrete, $t \in \{0, 1, 2, ...\}$. Wages w are rigid, leading to aggregate demand externalities. Interest rates R are fixed because the city is part of a currency union. Housing prices and rents are flexible. The model is populated by two types of consumers: mobile workers and immobile homeowners.[39]

Mobile workers. — Mobile workers earn money in the labor market, pay rent, and consume the remainder hand-to-mouth. Each worker lives in exactly one house. They have the option to live in the city, or elsewhere, where they get utility v_t. Within the city, their expected utility is equal to their expected consumption, so they have an indifference condition where

$$v_t = (1 - u_t)w - r_t^h \tag{9}$$

where u_t is the unemployment rate and r_t^h is the rental rate. Denote the number of workers in the city by N_t, and normalize it to initially be 1.

[38] A more quantitative paper making a similar point is House, Proebsting and Tesar (2018).

[39] Splitting people into these two groups helps to analyze distributional consequences, but obviously, in the real world many people are both workers and homeowners. Having homeowners work would not complicate the model much.

Immobile homeowners. — Immobile homeowners own all the housing and land in a city. Their utility is given by

$$\sum_{t=0}^{\infty} \beta^t v(c_t^h) \tag{10}$$

subject to $p_t^h H_t + c_t^h + a_t = a_{t-1}(1+R) + p_t^h(1-\delta)H_{t-1} + r_t^h N_t + p_t^Z Z$.

a is the liquid assets of these agents, H is the housing stock and p^h is the house price. p^Z is the price of local land and Z is land that becomes available for development each period. They are also potentially subject to borrowing constraints based on the house price. They rent the housing stock out to workers.[40]

I do not try to solve for the homeowner's optimal consumption, but instead point out where such changes affect the equilibrium. Other papers, such as Berger et al. (2017) go into more detail on the way consumers change consumption decisions when house prices change.

Housing. — Housing is produced competitively. Production uses local labor and land:

$$H_t = (1-\delta)H_{t-1} + Z^{\frac{1}{\sigma+1}} {L_t^h}^{\frac{\sigma}{\sigma+1}} \tag{11}$$

where σ is the long-run housing supply elasticity. Each worker lives in a house: $H_t = N_t$.

Labor Demand. — Labor demand comes from local consumption, external demand, and housing. Assume a fixed fraction of consumption is local, and that there is an exogenously determined demand from other cities. Then total labor demand is given by

$$L_t = \alpha N_t v_t + \alpha c_t^h + L_t^h + D_t \tag{12}$$

where α is the fixed share of local consumption and D_t is external labor demand. Unemployment is given by $u_t = 1 - \frac{L_t}{N_t}$.

Equilibrium. — For any path of v and D, an equilibrium is a path of labor demand, unemployment, population, consumption of mobile workers and immobile homeowners, house prices, and rents such that immobile homeowners are maximizing their utility, mobile workers have equal utility to their outside option, housing markets clear, and unemployment is determined by labor demand.

[40]Their marginal utility is used to price housing: $p_t^h = r_t^h + \beta(1-\delta)\frac{v'(c_{t+1})}{v'(c_t)} p_{t+1}^h$. The important part is that a local increase in rents also increases house prices.

THE RESPONSE TO A MIGRATION SHOCK. — Suppose there is a decrease in v. To equalize utility of mobile workers to their outside option, u_t and r_t^h have to respond.

Without housing, there would be only one way to reach equilibrium: an increase in u_t. u_t would move to satisfy $v_t = (1-u_t)w$ and that would be accomplished by increasing N_t until $u_t = 1 - \alpha v_t - \frac{D_t}{N_t}$. This is the insight of Harris and Todaro (1970).

But with housing, there are two margins of adjustment, rents and unemployment, that are linked together by the labor demand equation. If rents rise, that has several effects on labor demand. First, it raises prices of housing leading to more construction and more labor demand for construction. Second, the higher prices change the wealth of homeowners from the existing housing stock, potentially raising labor demand. On the flip side, it also reduces the (non-rent) consumption of workers, lowering labor demand from them. If the net effect of these channels is positive,[41] then the response to the migration shock will be rising rents and lower unemployment.

Another way to think about this is how the economy responds to inflows. First, the price of housing clearly goes up, as defined by the housing supply elasticity:

$$dp_t^h = p_t^h \frac{1}{\sigma\delta} dN_t \tag{13}$$

In future periods, the effect is smaller, given by $dp_{t+s}^h = p_t^h \frac{1}{\sigma} dN_t$ for all $s \geq 1$. So we expect this effect to be larger in the short-run.

So an increase in N certainly increases prices and will also increase rents. But the way migration affects unemployment is more complex:

$$du_t = \underbrace{-\alpha dv_t N_t - \alpha v_m^t dN_t}_{\text{Non-tradable Demand}} + \underbrace{(1-U^t)dN_t}_{\text{Labor Supply}} - \underbrace{\alpha dc_t^h}_{\text{House Price Channel}} - \underbrace{dL_t^h}_{\text{Construction Channel}} \tag{14}$$

Essentially, the change in the unemployment rate can be broken down into four components: the change in labor supply and labor demand from the migrants, and two housing channels. This equation essentially rationalizes the empirical findings by showing how inmigration can lower unemployment because of housing. I provided evidence that the two housing channels were present and large in Section II.

One can show that the construction channel is based only on model parameters: $dL^h = \frac{\sigma+1}{\sigma}\delta^{\frac{1}{\sigma}} Z^{-\frac{1}{\sigma}} dN_t$.[42] I do not try to make a similar characterization of homeowner consumption, but given that the net effect on non-tradable employment per capita is positive in the data, that would suggest the house price channel is

[41] There is a limit: if it's so positive that it decreases unemployment more than it raises rents, the equilibrium is unstable.

[42] Future periods would have a smaller but still positive effect.

positive. The total effect of the four channels is unclear from only the theory, but Section I showed it was negative in the data.

In this sense, the model rationalizes many of the main findings of the paper: that migration causes an increase in house prices, construction, and employment in construction and non-tradable sectors. And it allows for the effects to be a net negative for unemployment because of the housing channels. It rationalizes why the effects are larger in the short-term than the long-term, and why they vary by housing supply elasticity.

The response to a demand shock. — If we instead consider the effects of a small increase in D_t, the model behaves very similarly. Because of the outside option, it must be the case that $du_t w + dr_t = 0$. The direct effect of a change in D_t is to lower the unemployment rate, so mobile workers will inmigrate, causing rents to go up, and unemployment to fall even further. Section III gave empirical estimates of the size of this "migration accelerator."

Even though the main finding of this paper is that inmigration reduces the unemployment rate, the gains of inmigration in the model are captured entirely by homeowners, not workers.[43] In the real world, there is significant overlap between homeowners and workers, but they are gaining because they own a home, not because it is easier to find a job. In fact, the migration shock is making workers worse off because of the declining outside option. Similarly, the welfare changes from a demand shock are borne by the homeowners, not the workers.

Importantly, the accelerator can occur even without increasing the utility of the mobile renters because the rent moves to offset the increase. The demand shock causes only a finite increase in population.

Externalities. — There are two main externalities that occur because of migration which cause a divergence between a social planner's and the agent's decisions. The first is a pecuniary externality, that migrants affect house prices. If homeowners are faced with binding borrowing constraints that are relaxed by the increase in house prices, this may create a desire for the social planner to have more migration than the competitive equilibrium.

The second externality is an aggregate demand externality. Because wages are sticky, increased demand leads to more income for workers. This externality can justify why labor mobility is desirable for currency unions (see Farhi and Werning, 2014). But in the data, the externality goes the wrong way.[44]

[43]I made a stark assumption that the utility of workers was pinned down by the outside option, leading to this result. A more general model, that had mobility costs or idiosynchratic location preferences would change that assumption. But a migration shock would unambiguously raise homeowner utility, while the effect on workers is ambiguous.

[44]In this stylized model, positive demand shocks are always welfare enhancing, but that could be relaxed with vacancy posting costs and a matching function, which would lead to a socially optimal level of unemployment.

Consider what happens after a positive demand shock, when aggregate demand is high. Workers want to move into the city because of low unemployment, not taking into account their effects on demand, further lowering unemployment. Hence the aggregate demand externality goes in the same direction as the initial shock. In the optimal currency area literature, labor mobility is desirable because it stabilizes aggregate demand, but in this setting, labor mobility simply amplifies any aggregate demand changes.

V. Conclusion

In this paper, I document that domestic inmigration causes a large decline in the local unemployment rate. This effect is explained by two housing channels: an increase in construction and an increase in house prices, inducing non-tradable consumption. These two channels are large and overwhelm the effects of increased labor supply. I show these effects are heterogeneous across different types of housing markets. 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 labor demand shocks by a third.

Implications for Regional Evolutions. — These results have important implications for regional evolutions, a topic first analyzed in Blanchard and Katz (1992). Both papers agree that net migration responds positively to a positive labor demand shock. The difference between the two papers is on the effects of that migration.[45]

Blanchard and Katz (1992) interpret the strong response of migration to mean that "most of the adjustment to an adverse shock of employment is through out-migration of labor" (p. 37). Holding the decline in jobs constant, the response of migration is strong enough to account for the unemployment rate returning to average. So even if the employment level never returns to its previous trend, there has been enough migration to bring unemployment back to its previous level.

My contribution is to consider a counterfactual where the number of jobs is not held constant, but instead responds endogenously to the migration. As I show, the labor demand effect of migration is substantial. In fact, the labor demand effect is larger than the labor supply effect, so the initial shock is amplified by the migration accelerator.

My results would suggest that migration is not what brings unemployment back to its previous level. Figure 10 shows that if we took out the effect of migration,

[45]Another major development for regional evolutions is that migration has fallen significantly since Blanchard and Katz (1992) was written. Molloy, Smith and Wozniak (2011) and Kaplan and Schulhofer-Wohl (2017) document and begin to explain this phenomenon. The accelerator that I document is related only loosely in that the magnitude of the accelerator would likely decline with migration.

unemployment would still return to trend before six years and would happen even faster without migration.[46]

Implications for Cross-sectional Estimation. — Another important implication of these results is that they influence how we should think about aggregation for local shocks. Many studies use cross-sectional variation in the intensity of a shock to estimate the effect of the shock nationally (e.g. Autor, Dorn and Hanson, 2013; Nakamura and Steinsson, 2014). One of the reasons aggregation is difficult is because of general equilibrium effects that can occur across cities, of which one of those is the migration response.[47]

My results suggest that the cross-sectional estimates overstate the aggregate impact of a shock. Consider the Bartik (1991)-style shocks I used in Section III. If one wished to estimate the effects of a broad labor demand shock on national unemployment, the researcher would want to purge their cross-sectional estimates of the migration amplification. Given that migration amplifies shocks, it would suggest that estimates based on cross-sectional data are biased upwards compared to the national effects.[48]

[46] If not migration, the force bringing unemployment back to its previous level must be that either job creation declines or labor force participation rises when unemployment is low. In Blanchard and Katz (1992), they find that labor force participation rises briefly after a positive shock, suggesting that job creation is the relevant margin.

[47] Other reasons include the response of monetary policy, network effects, and the price of tradable goods.

[48] Indeed, in work subsequent to this paper, Chen (2018) finds that local fiscal shocks are 30 percent higher than their national counterparts.

A. Data Appendix

Migration data is obtained from the IRS Statistics of Income from 1990-2013 (Internal Revenue Service, 1990-2013). I do not make any adjustment or imputation for data that is censored below at 10 tax returns.[49]

For 1940, migration data comes from the full Census (Ruggles et al., 2019). To convert county codes from ICPSR to FIPS, I take the five leading digits. Each moving respondent was asked which county they lived in five years prior. In several important ways, this is a different measure, and there is no censoring in this data. However, it is still highly predictive of movements in the IRS data 50 years later.

Unemployment comes from the Bureau of Labor Statistics Local Area Unemployment Statistics (Bureau of Labor Statistics, 1990-2013). This is reported at the MSA level. To check robustness, I also use the public use microdata for the CPS, though not all MSAs are identifiable off of this (Ruggles et al., 2018). I also use unemployment insurance payouts from the Bureau of Economic Analysis, which is available at the MSA level (Bureau of Economic Analysis, 1990-2014).

Employment comes from County Business Patterns (United States Census Bureau, 1989-2013), which I aggregate to MSAs (Missouri Census Data Center , MCDC).

Population is from the Census estimates, which I get from the National Bureau of Economic Research webpage (Roth, 2007). Again, I aggregate counties to MSAs (MCDC, 2014).

I use house prices from the Federal Housing Finance Authority, which reports local house price indexes for MSAs, going back to different points in time (Federal Housing Finance Agency, 1975-2015). By 2000, all but one MSA in my sample has a house price index.

Permits come from the Census Building Permits Survey (United States Census Bureau, 1995-2014). Prior to 2003, the Census used a different definition of MSAs, which I converted to the new definitions using weights based on population in 2000 (MCDC, 2014). I then spliced those using the Census series for the new MSA definition, which is available after 2003.

Mortgage data comes from Home Mortgage Disclosure Act, which is available from the National Archives (Federal Reserve Board of Governors Division of Consumer and Community Affairs, 1990-2014). I matched counties to MSAs (MCDC, 2014). Second-lien mortgages were only available after 2004.

I used the Mian and Sufi (2014) definitions of tradable, non-tradable, and construction employment, which I matched with County Business Patterns (United States Census Bureau, 1989-2013). The industry decomposition only matches after 1998.

I use wage data from the Occupational Employment Survey (Bureau of Labor Statistics, 2001-2015), which reports percentiles of the wage distribution by

[49]I do adjust for the fact that Miami-Dade county's FIPS code changes midway through the sample.

MSA after 2005, and for a previous definition of MSAs before that. I create a correspondence using MCDC, 2014.

For the Mariel section of the paper, I use house prices from Federal Housing Finance Agency (1975-2015) and the Quartely Census of Employment and Wages (Bureau of Labor Statistics, 1975-1984).

To get estimates on the college education of the population, I use the Census 5 percent sample in 1990 (Ruggles et al., 2019), and convert the Public Use Microdata Areas to MSAs (MCDC, 2014). I estimate the share of college-education by dividing the number of respondents with 4 or more years of college by the total number of respondents.

For the estimates of interest, dividend, and rental income, I use the American Community Survey from 2001-2014, which records migration status and income from these sources (Ruggles et al., 2019).

I also use the ACS for measures of Vacancy and Homeownership (Ruggles et al., 2019), which is calculable for the latter part of my sample, starting in 2005.

I use the elasticities from Saiz (2010), which are computed for the old definition of MSAs. I convert these by hand.

Average January high daily temperatures are available by county from the Center for Disease Control (Centers for Disease Control and Prevention, 1979-2011). I take the average from 1979 to 2011, and aggregate to MSAs (MCDC, 2014).

The latitude and longitude of counties, which I use to calculate distances is taken from the Gazetteer Files at the Census (United States Census Bureau, 2016). I aggregate to MSAs (MCDC, 2014).

To calculate the Bartik instruments in 1998, I convert NAICS codes to SIC codes using a crosswalk from Autor, Dorn and Hanson (2013).

*

REFERENCES

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao. 2015. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program." National Bureau of Economic Research Working Paper 21512.

Altonji, Joseph G, and David Card. 1991. "The effects of immigration on the labor market outcomes of less-skilled natives." In *Immigration, Trade, and the Labor Market*. 201–234. University of Chicago Press.

Amior, Michael, and Alan Manning. 2018. "The persistence of local joblessness." *American Economic Review*, 108(7): 1942–70.

Attanasio, Orazio, Andrew Leicester, and Matthew 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(3): 399–435.

Autor, David, David Dorn, and Gordon H Hanson. 2013. "The China syndrome: Local labor market effects of import competition in the United States." *The American Economic Review*, 103(6): 2121–2168.

Bartel, Ann P. 1989. "Where do the new US immigrants live?" *Journal of Labor Economics*, 371–391.

Bartik, Timothy J. 1991. *The Debate Over State and Local Economic Development Policies.* WE Upjohn Institute for Employment Research.

Baxter, Marianne, and Robert G King. 1993. "Fiscal policy in general equilibrium." *The American Economic Review*, 315–334.

Beaudry, Paul, David A. Green, and Ben M. Sand. 2018. "In Search of Labor Demand." *American Economic Review*, 108(9): 2714–57.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2017. "House Prices and Consumer Spending." *The Review of Economic Studies*, 85(3): 1502–1542.

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

Borjas, George J. 2003. "The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *The Quarterly Journal of Economics*, 118(4): 1335–1374.

Borjas, George J. 2017. "The Wage Impact of the Marielitos: A Reappraisal." *Industrial and Labor Relations Review*, 70(5): 1077–1110.

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

Boustan, Leah Platt, Price V Fishback, and Shawn Kantor. 2010. "The effect of internal migration on local labor markets: American cities during the Great Depression." *Journal of Labor Economics*, 28(4): 719–746.

Bureau of Economic Analysis. 1990-2014. "CAINC35 Personal Current Transfer Receipts: Unemployment insurance compensation." https://www.bea.gov/data/income-saving/personal-income-county-metro-and-other-areas (accessed August 21, 2019).

Bureau of Labor Statistics. 1975-1984. "Quarterly Census of Employment and Wages." https://www.bls.gov/cew/downloadable-data-files.htm (accessed October 26, 2016).

Bureau of Labor Statistics. 1990-2013. "Local Area Unemployment Statistics." https://www.bls.gov/lau/#data (accessed October 31, 2015).

Bureau of Labor Statistics. 2001-2015. "Occupational Employment Statistics." https://www.bls.gov/oes/tables.htm (accessed July 27, 2016).

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

Campbell, John Y, and Joao F Cocco. 2007. "How do house prices affect consumption? Evidence from micro data." *Journal of Monetary Economics*, 54(3): 591–621.

Card, David. 1990. "The impact of the Mariel boatlift on the Miami labor market." *Industrial and Labor Relations Review*, 43(2): 245–257.

Card, David. 2001. "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration." *Journal of Labor Economics*, 19(1): 22–64.

Card, David. 2009. "How immigration affects U.S. cities." *Making Cities Work: Prospects and Policies for Urban America*, 158–200.

Centers for Disease Control and Prevention. 1979-2011. "North America Land Data Assimilation System (NLDAS) Daily Air Temperatures and Heat Index." https://wonder.cdc.gov/NASA-NLDAS.html (accessed August 15, 2019).

Chen, Vince W. 2018. "Fiscal Multipliers and Regional Reallocation." Boston University. https://drive.google.com/uc?authuser=0&id=1kZCMzJtHQp1m9opyHobDfdIbYi5yQ23w.

Coen-Pirani, Daniele. 2010. "Understanding gross worker flows across U.S. states." *Journal of Monetary Economics*, 57(7): 769–784.

Coleman, Andrew, and Özer Karagedikli. 2018. "Residential construction and population growth in New Zealand: 1996-2016." Reserve Bank of New Zealand Discussion Paper 2018/02.

Decressin, Jörg, and Antonio Fatas. 1995. "Regional labor market dynamics in Europe." *European Economic Review*, 39(9): 1627–1655.

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

Farhi, Emmanuel, and Iván Werning. 2014. "Labor Mobility Within Currency Unions." National Bureau of Economic Research Working Paper 20105.

Federal Housing Finance Agency. 1975-2015. "House Price Index Datasets: All-Transactions Indexes (Estimated using Sales Prices and Appraisal Data): Metropolitan Statistical Areas and Divisions (Not Seasonally Adjusted)." https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx (accessed November 13, 2015).

Federal Reserve Board of Governors Division of Consumer and Community Affairs. 1990-2014. "Home Mortgage Disclosure Data Files." https://catalog.archives.gov/id/2456161 (accessed August 18, 2019).

Foote, Christopher L. 2007. "Space and time in macroeconomic panel data: young workers and state-level unemployment revisited." Federal Reserve Bank of Boston Working Paper 07-10.

Glaeser, Edward L, and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy*, 113(2): 345–375.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2018. "Bartik Instruments: What, When, Why, and How." National Bureau of Economic Research Working Paper 24408.

Gonzalez, Libertad, and Francesc Ortega. 2013. "Immigration and housing booms: Evidence from Spain." *Journal of Regional Science*, 53(1): 37–59.

Greulich, Erica, John M Quigley, and Steven Raphael. 2004. "The Anatomy of Rent Burdens: Immigration, Growth, and Rental Housing." *Brookings-Wharton Papers on Urban Affairs*, 149–205.

Harris, John R, and Michael P Todaro. 1970. "Migration, unemployment and development: a two-sector analysis." *The American Economic Review*, 126–142.

Hong, Gihoon, and John McLaren. 2015. "Are Immigrants a Shot in the Arm for the Local Economy?" National Bureau of Economic Research Working Paper 21123.

House, Christopher L, Christian Proebsting, and Linda L Tesar. 2018. "Quantifying the Benefits of Labor Mobility in a Currency Union." National Bureau of Economic Research Working Paper 25347.

Internal Revenue Service. 1990-2013. "U.S. Population Migration Data." https://www.irs.gov/statistics/soi-tax-stats-migration-data (accessed April 1, 2016).

Jaeger, David A., Joakim Ruist, and Jan Stuhler. 2018. "Shift-Share Instruments and the Impact of Immigration." National Bureau of Economic Research Working Paper 24285.

Jimeno, Juan F, and Samuel Bentolila. 1998. "Regional unemployment persistence (Spain, 1976–1994)." *Labour Economics*, 5(1): 25–51.

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American Economic Review*, 95(1): 161–182.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2017. "Understanding the long-run decline in interstate migration." *International Economic Review*, 58(1): 57–94.

Kaplan, Greg, Kurt Mitman, and Giovanni L 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 22232.

Kochin, Levis A. 1996. "Real (Estate) Booms and Busts on Froniers from the Fall of Cow to the Arrest of Keating." Unpublished.

Lee, Donghoon, Christopher Mayer, and Joseph Tracy. 2012. "A New Look at Second Liens." *Housing and the Financial Crisis*, 205–234. University of Chicago Press.

Lewis, Ethan G. 2005. "Immigration, skill mix, and the choice of technique." Federal Reserve Bank of Philadelphia Working Paper 05-8.

Long, Jason, and Henry Siu. 2018. "Refugees from Dust and Shrinking Land: Tracking the Dust Bowl Migrants." *The Journal of Economic History*, 78(4): 1001–1033.

Marlay, Matthew, and Peter Mateyka. 2011. "The Seasonality of Moves: 2009." Census Bureau Working Paper.

McIntosh, Molly Fifer. 2008. "Measuring the labor market impacts of Hurricane Katrina migration: evidence from Houston, Texas." *The American Economic Review*, 98(2): 54–57.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household balance sheets, consumption, and the economic slump." *The Quarterly Journal of Economics*, 128(4): 1687–1726.

Mian, Atif R, and Amir Sufi. 2014. "What explains the 2007-2009 Drop in Employment?" *Econometrica*, 82: 2197–2223.

Missouri Census Data Center (MCDC). 2014. "Geocorr 2014: Geographic Correspondence Engine." http://mcdc.missouri.edu/applications/geocorr2014.html (accessed August 20, 2019).

Molloy, Raven, Christopher L Smith, and Abigail Wozniak. 2011. "Internal migration in the United States." *The Journal of Economic Perspectives*, 25(3): 173–196.

Molloy, Raven, Christopher L Smith, and Abigail Wozniak. 2014. "Declining migration within the US: the role of the labor market." National Bureau of Economic Research Working Paper 20065.

Monras, Joan. 2015*a*. "Economic Shocks and Internal Migration." IZA Discussion Paper 8840.

Monras, Joan. 2015*b*. "Immigration and wage dynamics: Evidence from the mexican peso crisis." IZA Discussion Paper 8924.

Mundell, Robert A. 1961. "A theory of optimum currency areas." *The American Economic Review*, 51(4): 657–665.

Nakamura, Emi, and Jon Steinsson. 2014. "Fiscal stimulus in a monetary union: Evidence from US regions." *The American Economic Review*, 104(3): 753–792.

Ottaviano, Gianmarco, and Giovanni Peri. 2006. "Wages, rents and prices: The effects of immigration on US natives." Center for Research and Analysis of Migration, Discussion Paper 13-07.

Poterba, James M. 1984. "Tax subsidies to owner-occupied housing: an asset-market approach." *The Quarterly Journal of Economics*, 729–752.

Rappaport, Jordan. 2007. "Moving to nice weather." *Regional Science and Urban Economics*, 37(3): 375–398.

Roth, Jean. 2007. "Census U.S. Intercensal County Population Data, 1970-2014." http://data.nber.org/data/census-intercensal-county-population.html (accessed July 7, 2016).

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. 2018. "Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]." Minneapolis: University of Minnesota Machine Readable Database. https://doi.org/10.18128/D030.V6.0 (accessed August 16, 2019).

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. 2019. "IPUMS USA: Version 9.0 [dataset]." Minneapolis: University of Minnesota Machine Readable Database. https://doi.org/10.18128/D010.V9.0.

Saiz, Albert. 2003. "Room in the kitchen for the melting pot: Immigration and rental prices." *Review of Economics and Statistics*, 85(3): 502–521.

Saiz, Albert. 2007. "Immigration and housing rents in American cities." *Journal of Urban Economics*, 61(2): 345–371.

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *The Quarterly Journal of Economics*, 125(3): 1253–1296.

Shimer, Robert. 2001. "The Impact of Young Workers on the Aggregate Labor Market." *The Quarterly Journal of Economics*, 116(3): 969–1007.

Ströbel, Johannes, and Joseph Vavra. 2019. "House Prices, Local Demand, and Retail Prices." *Journal of Political Economy*, 127(3): 1391–1436.

Topel, Robert, and Sherwin Rosen. 1988. "Housing investment in the United States." *The Journal of Political Economy*, 718–740.

United States Census Bureau. 1989-2013. "County Business Patterns (CBP): CBP Datasets." https://www.census.gov/programs-surveys/cbp/data/datasets.html (accessed February 24, 2016).

United States Census Bureau. 1995-2014. "Building Permits Survey." https://www2.census.gov/econ/bps/Metro/ (accessed August 16, 2019).

United States Census Bureau. 2016. "Gazetteer Files." https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.2016.html (accessed June 28, 2016).