Online Appendix for
The Migration Accelerator:
Labor Mobility, Housing, and Demand

By Greg Howard

B. Comparison to OLS

To understand the importance of the instrument, I also run the regression simply on the inmigration itself, not the predicted inmigration.[1] The results are shown in Figure B1. The effects are not that different, but the estimates at $t = -2$ and $t = 0$ do suggest some bias. The estimate at $t = -2$ is positive, meaning that when unemployment is on a downward trend, migrants move in more, even conditional on the previous two years of migration. Similarly, the coefficient at $t = 0$ is significantly smaller than the coefficient from the instrument. This suggests that there may be some bias from migrants anticipating that the unemployment rate will decline next year, and deciding to move in because of it.

This is the bias we expect to find from OLS, given that people move to places they are more likely to find jobs. However, it is worth noting that the bias is small. This may be because both regressions include MSA fixed effects, so persistent differences across space, which could drive a lot of migration, do not show up in either.

C. Similarities between high-migration city pairs

In this appendix, I show that migrants move between cities that are similar on two observable dimensions: location and industrial composition. The reason this is important is that it helps to sign any bias that the reader might be concerned about. If there are unobserved shocks causing outmigration in one city and unemployment in the other city to both change, this exercise points out that those shocks are likely similar, and therefore causing more outmigration and more unemployment, biasing my results upward. Since the main result is that unemployment falls, we can interpret my regressions as a lower bound.

To establish this similarity, I regress migration on measures of their similarity.[2] The regression establishes that people do move between places that are similar,

---

[1] The specification can be found in equation (2) in the main text. The migration data used come from Internal Revenue Service (1990-2013) and the unemployment data from Bureau of Labor Statistics (1990-2013).

[2] In addition to the migration data, this also uses United States Census Bureau (2016) and United States Census Bureau (1989-2013).
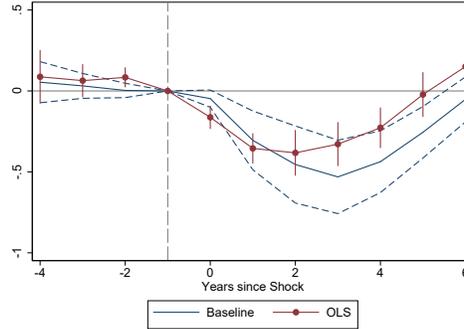
Figure B1. : Unemployment Rate. Comparison to OLS. Standard errors clustered by state.

Table C1—: Migration Network

VARIABLES	(1) Log Migration	(2) Log Migration	(3) Log Migration	(4) Log Migration
Log Distance	-1.691	-1.635		
	(0.032)	(0.037)		
Industry Similarity			2.513	3.281
			(0.264)	(0.431)
Observations	56,630	37,705	56,630	37,705
Origin and Destination Fixed Effects	YES	YES	YES	YES
Flexible Distance Controls	–	–	YES	YES
Unit	CBSA	MSA	CBSA	MSA

Standard errors clustered by from and to MSAs/CBSAs.

Source: Internal Revenue Service (1990-2013), United States Census Bureau (1989-2013), United States Census Bureau (2016), author's calculations

suggesting that places between which people move are likely to experience similar shocks.

In column (1) of Table C1, I estimate a gravity-like relationship between migration on the distance between any two CBSAs using the specification below.[3] In column (2), I show the same result for MSAs.

$$(C1) \qquad \log m_{i \to j, t_0} = \beta \log \text{distance}_{ij} + \alpha_i + \gamma_j + \epsilon_{ij}$$

Another piece of evidence is that migration is higher between MSAs with similar industries. In columns (3) and (4), I control for a quintic in log-distance, and run the regression on an industry similary index, using 2-digit SIC codes from

---

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

1990. I construct the vector of employment in each of those industries, and use the following formula:

$$(C2) \qquad \text{Industry Similarity}_{ij} = \frac{v_i \cdot v_j}{||v_i|| ||v_j||}$$

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

$$(C3) \quad \log m_{i \to j, t_0} = \beta \text{Industry Similarity}_{ij} + P^5(\log \text{distance}_{ij}) + \alpha_i + \gamma_j + \epsilon_{ij}$$

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

D. Robustness

In this appendix, I show that the main results are robust to a variety of controls and other instrumentation strategies.

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

Figure D1 presents the robustness of the result to these concerns. I flexibly control for industry and education. To do this, I interact the shares of 2-digit SIC industries in 1990 with a year fixed effect to control for industry.[4] This specification would nest controlling for Bartik shocks, which would allow less flexibility over the coefficient on the various industry shares. For education, I interact shares of the 11 education codes in the 1990 Census with year dummies.[5]

I also present the results using an alternative shock. To construct this shock, I exclude using migration to or from cities that are similar to the receiving city, based on the correlation of Bartik shocks. In Section III.A, I construct Bartik instruments for each MSA. I compute the correlation of these Bartik shocks over my entire sample, and I reconstruct my inmigration shock, excluding outmigration from the 250 CBSAs that have the most highly correlated Bartik shocks.[6]

---

[4]Industry data comes from United States Census Bureau (1989-2013).
[5]Education data comes from (Ruggles et al., 2019).
[6]With alternative instruments, the estimates of the effect on population change, so I renormalize the coefficients so as still to be interpretable as the effect of a one percent population increase via migration.
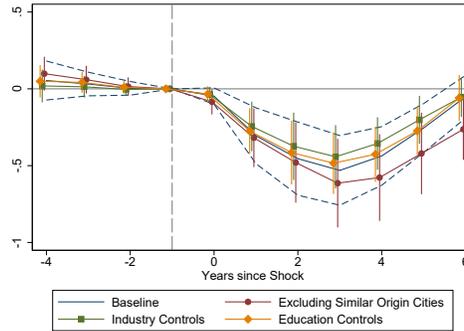
Figure D1. : Unemployment Rate. The effect of an inmigration shock. Robustness to City Charac-
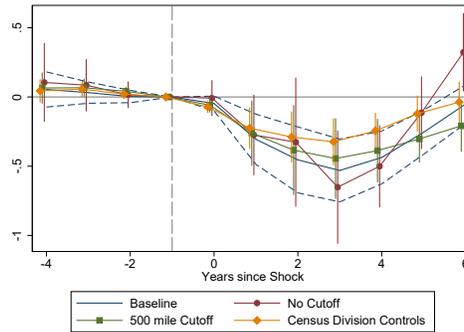teristics. Standard errors clustered by state.

Figure D2. : Unemployment Rate. The effect of an inmigration shock. Robustness to spatial
correlations. Standard errors clustered by state.

ROBUSTNESS TO SPATIAL CORRELATIONS. — In Figure D2, I investigate a variety of
robustness checks aimed at addressing concerns about the spatial structure of my
regression. The concern here is that there may be an omitted variable that affects
areas near the MSA, causing people to move, but which also directly affects the
unemployment rate.

Because my shock relies on the cutoff of 100 miles, I investigate the robustness
to that by using no cutoff, and using a cutoff of 500 miles.[7] The last robustness
check is that I control for the Census Division each MSA is in, interacting each
division with year fixed effects. These controls do not affect the results.

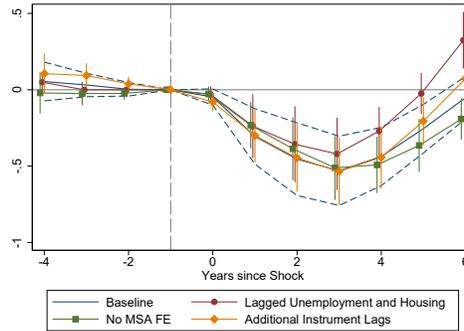[7]I am again renormalizing the effects for these new instruments.

Figure D3. : Unemployment Rate. The effect of an inmigration shock. Robustness to Specification. Standard errors clustered by state.

ROBUSTNESS TO SPECIFICATION. — One set of concerns over the main regression is that the choice of specification is important. In Figure D3, I show that this is not the case. In my baseline specification, I used two lags of the migration instrument to control for the effects of previous migration shocks. Here, I also include two lags of unemployment changes, house price changes, and housing permits changes. Mechanically, this drives the coefficient in $t = -2$ and $t = -3$ to zero.

I also show the results are robust to excluding the MSA fixed effect. I include it in the baseline because migration shocks are persistent over long periods, but it does not seem to have much effect.

Finally, I include additional lags of the migration shock, up to four years.

ROBUSTNESS TO TIME PERIOD. — One concern is that these results may be driven by the Great Recession. In Figure D4, I look at the effect of shocks that occurred in years prior to 2001 (up to 2000), as well as the effect of shocks after 2000 (starting in 2001). In years $t = 2, 3, 4$, the effects are slightly more muted before 2001, but are stronger in year $t = 0$. While they are statistically different than one another, nonetheless they are significantly negative, and within the confidence interval of the baseline regression. I conclude that while the effects may have been amplified by the Great Recession, they are still present and significant prior to it.

ROBUSTNESS TO MIGRATION NETWORK. — Another concern is that by using the 1990-1993 data to construct the migration network, I introduce bias because shocks that occurred during that time period are ongoing in later years, or similar geographic shocks reoccur soon after. In Figure D5, I construct the migration network using the 1940 Census rather than the pre-period of 1990-1993. I use the full 1940 Census, which asked migrants which county they lived in five years previous (see Ruggles et al., 2019). Besides the different pre-period, I construct
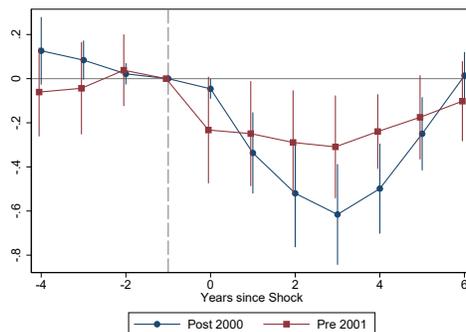
Figure D4. : Unemployment Rate. The effect of an inmigration shock. Robustness to time period. Standard errors clustered by state.

the instrument in the same way and use the same regression. The effect on migration is shown in the top left. The confidence intervals are a bit wider, as expected. Nonetheless, the effect on unemployment is robust to this pre-period. Indeed the magnitude of the effect is larger, though it is also more short-lived. The increase in house prices is also larger.

This robustness check is suggestive that the results are not dependent on any specific shock that affected the pre-period migration patterns and the path of the unemployment rate, as such a shock would have had to delay its effect on unemployment for over 50 years.

HURRICANE KATRINA. — A major source of variation in my data is from Hurricane Katrina, where many people from New Orleans were displaced. This event has been used as a natural experiment to investigate the economic effects of migration on receiving cities, often Houston (see Gagnon and Lopez-Salido, 2014; McIntosh, 2008; De Silva et al., 2010). Here, I show that my results are robust to using only this variation. Figure D6 uses only outflows from the eight counties hit hardest by Katrina: Cameron, Plaquemines, Jefferson, St. Bernard, and Orleans in Louisiana; and Hancock and Harrison in Mississippi. I also only use the outflows from 2005.[8] The instrument for inflows to other cities is based on the 1990-1994 patterns used throughout the rest of the paper.

In the top left of Figure D6, I show that the instrument did a good job of predicting inflows. On the top right, I show this was associated with a decline in the unemployment rate. The bottom shows there was also an increase in house prices. The unemployment results are different in the first period, but then consistent with the rest of the paper. Initially, the unemployment rate rose, but then fell in the next year, and remains low. The initial rise is inconsistent but not

---

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

(a) Inmigration Rate
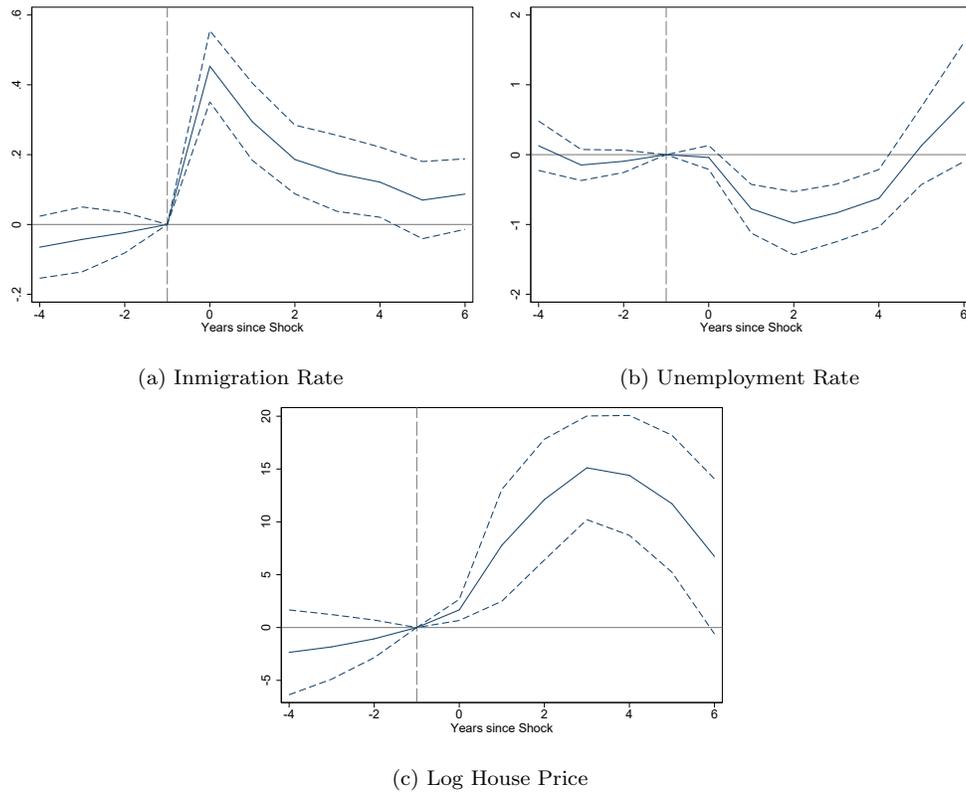

(b) Unemployment Rate

(c) Log House Price

Figure D5. : The effect of inmigration shocks using the 1940 network, with 95 percent confidence intervals. Errors clustered by state.

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

ROBUSTNESS TO MEASUREMENT OF UNEMPLOYMENT. — In the main text, I showed that migration affected several measures of the labor market, including unemployment, employment-population ratios, and unemployment benefits. However, given the centrality of unemployment, it is important to make sure it is measured well. In the main text, I use the Local Area Unemployment Statistics from the BLS, but some of that is imputed. In Figure D7, I use only the raw data from the CPS in order to estimate the effect of migration (Ruggles et al., 2018). The downside of this strategy is that I can construct it for fewer MSAs and in fewer years, due to data confidentiality. The raw data records the CPS for 242 MSAs, only after 2004. I use only these MSAs, using unemployment rates I calculate on

(a) Inmigration Rate
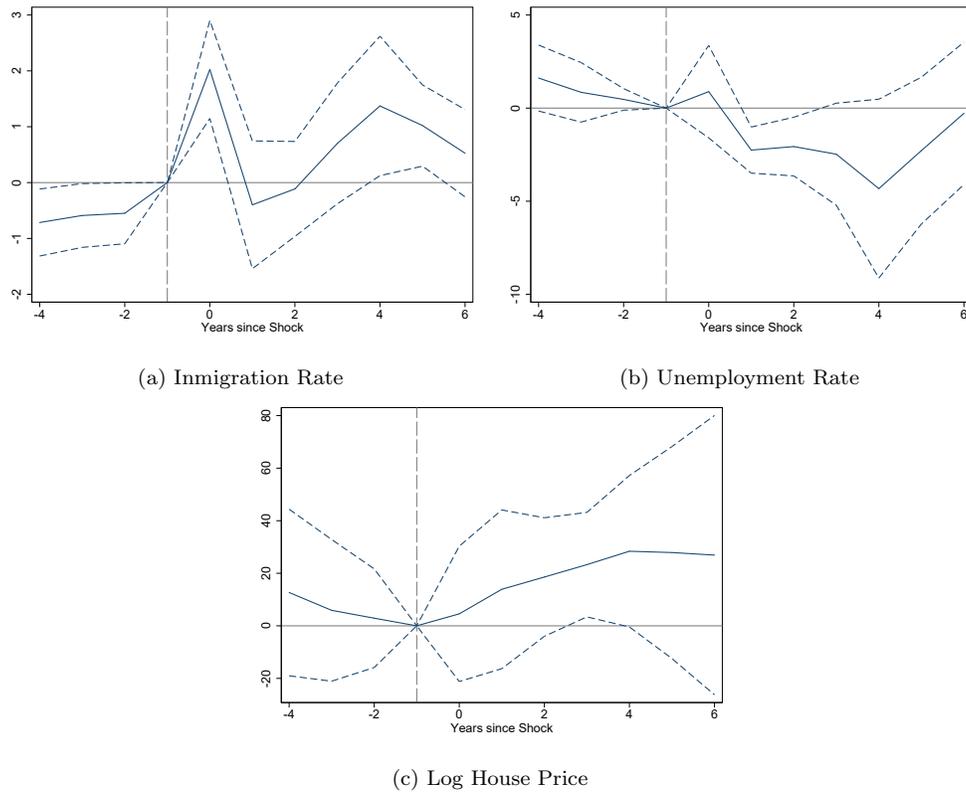

(b) Unemployment Rate

(c) Log House Price

Figure D6. : The effect of outmigration from Katrina hit areas in 2005, with 95 percent confidence intervals. Errors clustered by state.

the raw data, to draw the line "Raw CPS Data" in Figure D7.

Because of the short time period of data availability for the raw CPS data, I drop the CBSA fixed effect. Including it has a large effect on the estimates toward the latter years, which I believe are caused by an over-fitting of the fixed effects.

D1. Robustness of Other Variables

ROBUSTNESS OF EMPLOYMENT COMPOSITION. — In Figure D8, I show the robustness of the employment composition effects, using similar controls as in Section I.D. The same general patterns emerge as in the main body of the paper: sizable increases in construction and non-tradable employment, and no initial effect on tradables.

ROBUSTNESS OF HOUSING PERMITS AND PRICES. — Figure D9 shows the robustness of the increase in house prices and permits. The results are largely the same, with
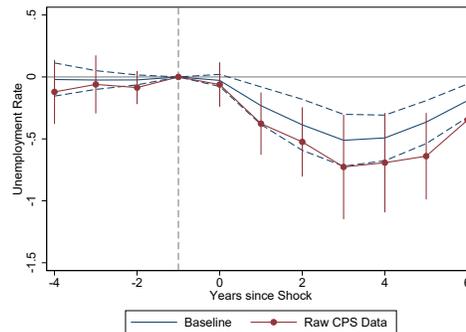
Figure D7. : Raw CPS Unemployment Data. The effect of an inmigration shock, with 95 percent confidence intervals. Errors clustered at MSA level.

house prices and permits increasing.

OTHER MEASURES OF HOUSING MARKETS. — Because of the centrality of the housing market, it may be interesting to better understand the effects inmigration has on other common measures, especially the vacancy and homeownership rates. To measure this, I use the ACS microdata from 2005-2013 (Ruggles et al., 2019). I convert PUMAs to CBSAs using a crosswalk. Because of the crosswalk and the fact the data is only available after 2005, these measures may be noisier than some of my other housing variables.

The results are shown in Figure D10. Vacancies decline in the MSA, as would be expected from an inmigration shock. This is consistent with the increase in housing prices. In addition, the homeownership rate rises. This could be because migrants are more likely to own homes.

THE HOUSING CHANNELS DURING THE MARIEL BOATLIFT. — As a last robustness check, I use a different source of variation to demonstrate the presence of these housing channels. Specifically, I use the famous example of the Mariel boatlift in Miami in 1980, where around 125,000 Cuban immigrants arrived in Miami. The goal of this exercise is not to ask whether the Mariel boatlift was stimulatory, but rather to see if there are similar patterns in employment across industries. Saiz (2003) provides a detailed study of the effect of the immigrants on the housing market, and so my focus is primarily on the labor market consequences. To summarize his results, he finds a large increase in rents due to the immigrants, but not house prices, which he speculates is due to native's preferences to not live near immigrants.[9] In Figure 7 of his paper, he also shows an increase in housing

---

[9]Such a result does not necessarily rule out a house price channel. There could be a composition effect in that higher-quality housing has declining prices and lower-quality housing prices and rental housing
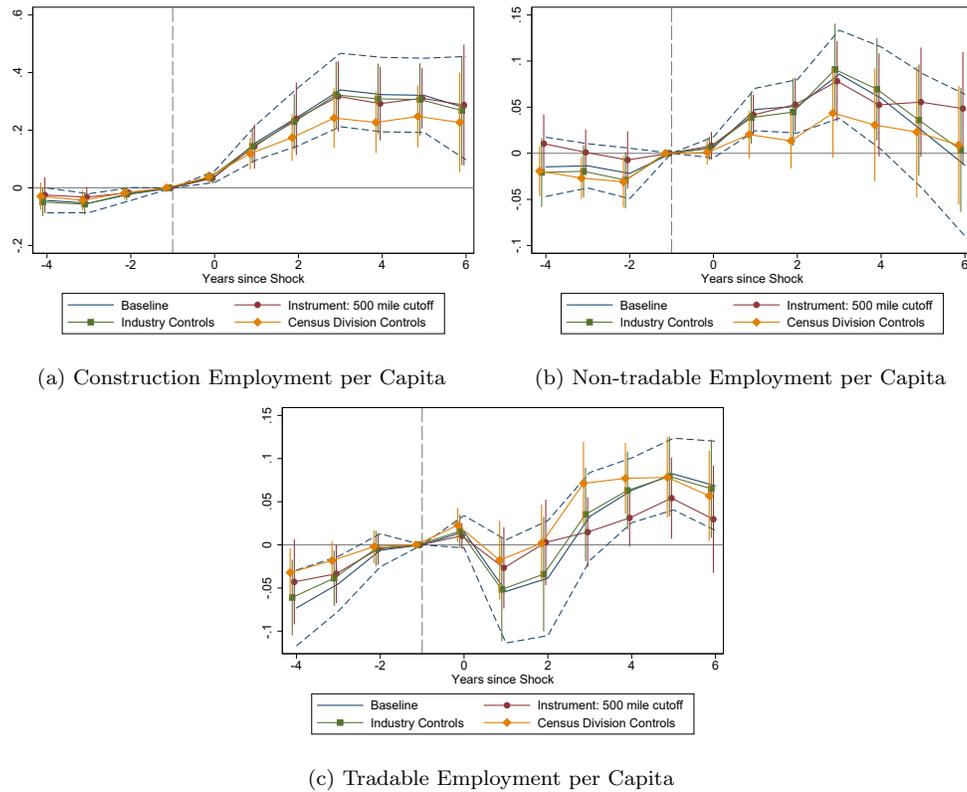
(a) Construction Employment per Capita

(b) Non-tradable Employment per Capita

(c) Tradable Employment per Capita

Figure D8. : The effect of an inmigration shock equal to one percent of the MSA's population, with 95 percent confidence intervals. Errors clustered at the MSA level. Number of CBSAs: 381.

permits per capita in 1980, while each of his control groups is falling. Though this does not last, and is somewhat noisy (in fact, Saiz says "there was no major supply response"), it is consistent with my results below about construction employment.

Because the boatlift was a one-time event, I will compare outcomes in Miami to those of a control group, based on similar cities in the United States at that time. There is debate over the appropriate group, so I will use the baseline group from Card (1990) and the baseline group from Borjas (2017).[10] I can only measure industries using SIC data at coarser levels, so using the decomposition from Mian, Rao, and Sufi (2013) is infeasible. Instead I will use manufacturing as a proxy for tradables, and retail trade as a proxy for non-tradables.[11] Construction

prices are rising. This could lead to a house price channel if the marginal propensity to consume is higher for people who own cheaper housing.

[10]Card (1990) uses Atlanta, Houston, Los Angeles, and Tampa; while Borjas (2017) uses Anaheim, Rochester (New York), Nassau, and San Jose.

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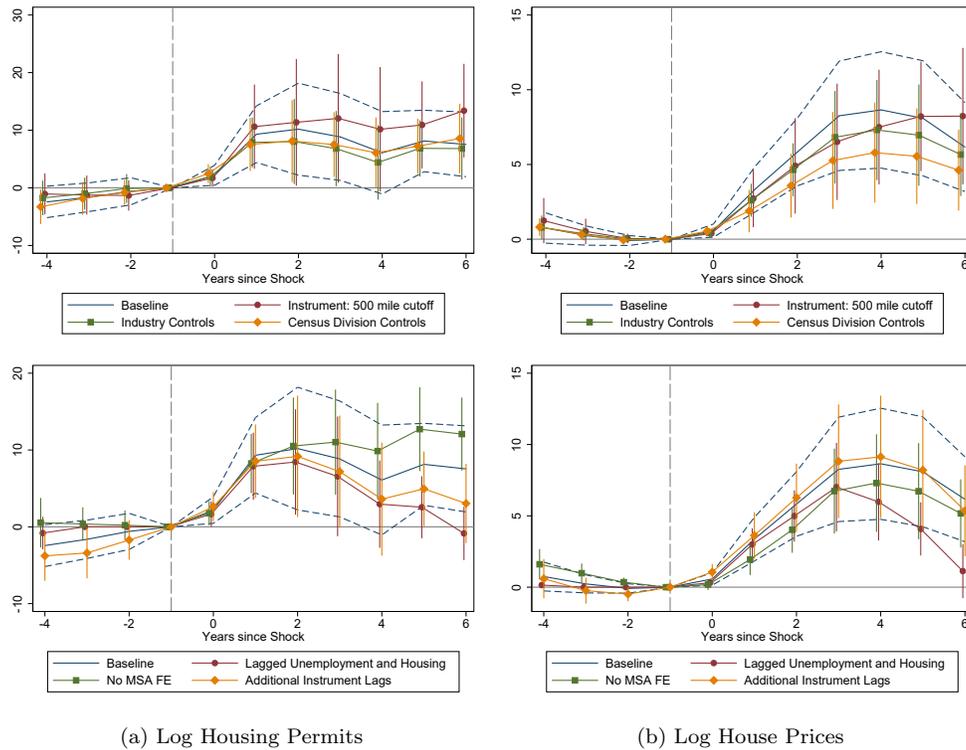
(a) Log Housing Permits

(b) Log House Prices

Figure D9. : The effect of an inmigration shock equal to one percent of the MSA's population, with 95 percent confidence intervals. Errors clustered by state.

is measured using the SIC industry. I plot the employment response of these industries in Figure D11. Data comes from the Bureau of Labor Statistics (1975-1984).

All of these plots show the same qualitative patterns as they do in response to migration shocks in the main body of my paper. In both construction and retail, there seems to be a temporary increase in the employment in that sector. In manufacturing, however, Miami's employment stays quite flat while the comparison cities are growing.

In sum, the data are consistent with the two housing channels. The Mariel boatlift was an approximately 7 percent expansion in the labor force (Card, 1990). This is significantly bigger than any shocks in my data, but the effects of these housing channels are nonetheless modest. Most likely, this is because there is less housing demand from each Mariel immigrant than from a domestic migrant in the United States.
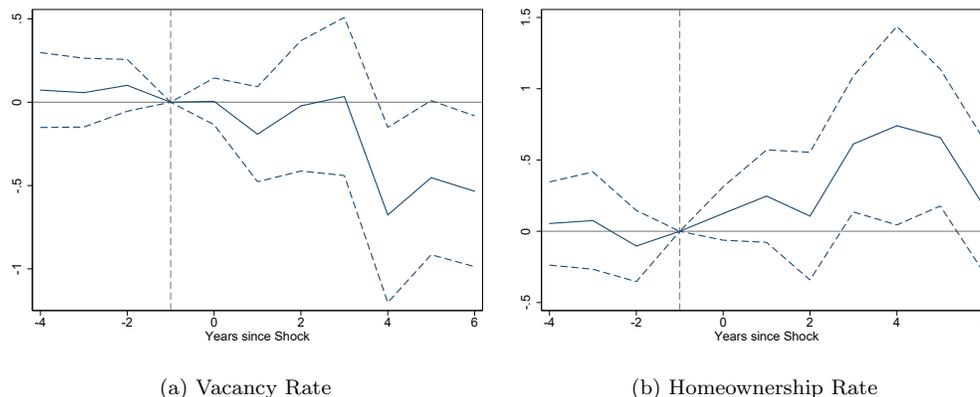
(a) Vacancy Rate   (b) Homeownership Rate

Figure D10. : The effect of an inmigration shock equal to one percent of the MSA's population, with 95 percent confidence intervals. Errors clustered by state.

E. Other Channels

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

Factor Complementarity. — Much of the immigration literature focuses on whether migrants are substitutes or complements with native workers, with implications that substitutes' productivity will decline, while complements' productivity increases. In contrast, in Section IV, I assume everyone is a perfect substitute.

The first step in investigating complementarity is to determine how skilled migrants are. A reasonable proxy might be income, which I plotted in Figure E1. Most people who move over 100 miles make, on average, 700 dollars more per year than those who stay in the same MSA. The 100 miles cutoff is relevant because that is what I use to construct my shock. From Molloy, Smith and Wozniak (2011), we also know they also tend to be younger and more educated. Therefore, a complementarity story would suggest that higher-skilled workers' labor markets would get worse, while the lower-skilled workers would benefit.

One way to investigate this is to use the Occupational Employment Statistics on the wage distribution (Bureau of Labor Statistics, 2001-2015). This data starts in 2001, so does not cover my complete dataset. It also uses a different definition of MSA for the first few years of data. I run the same regressions as I do throughout the rest of the paper, but using the 10th, 25th, 50th, 75th, and 90th percentiles

(a) Construction Employment

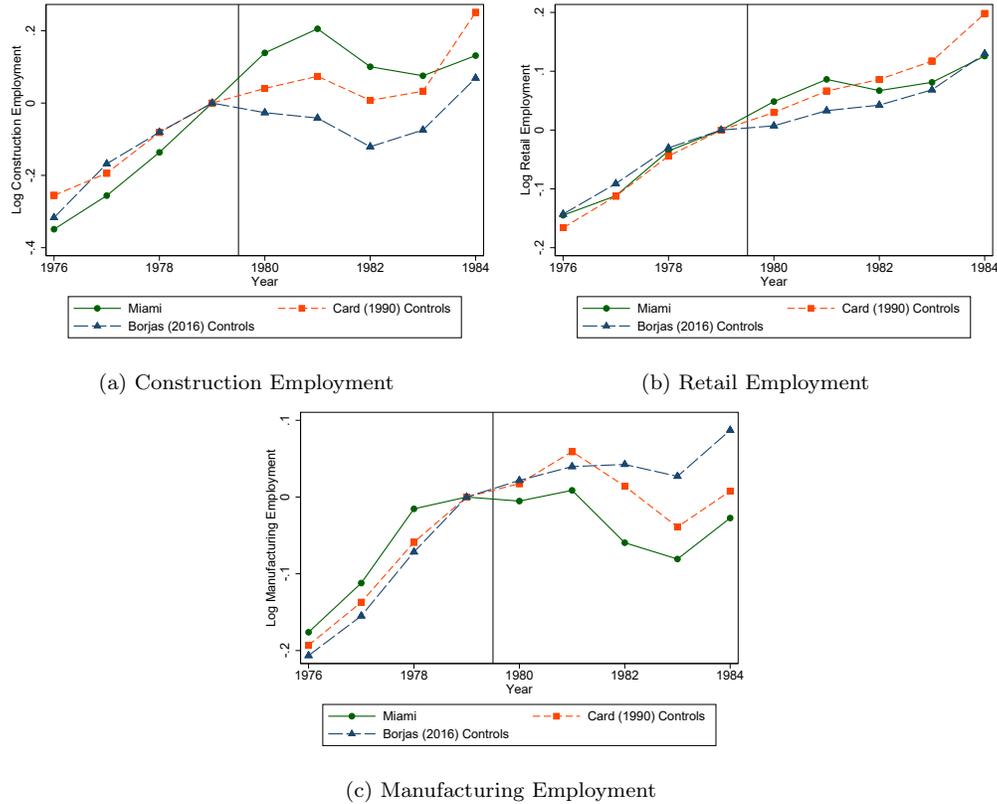(b) Retail Employment

(c) Manufacturing Employment

Figure D11. : The Effect of the Mariel Boatlift. Index, 1979=0.

of the hourly wage distribution as my independent variable. I plot the results in Figure E2.

The migration shock increases the 10th percentile of workers' hourly wages, consistent with a complementarity story. However, there is no negative effect on high wage earners, and even a positive effect in the later years.

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

My results do not depend on the cut-off that I use to construct my instrument. In general, the further the cutoff, the higher-skilled is the migrant. See Figure E1. So under a complementarity story, you might expect the further cut-offs to have larger effects. However, the results shown in Figure D2 are very similar regardless
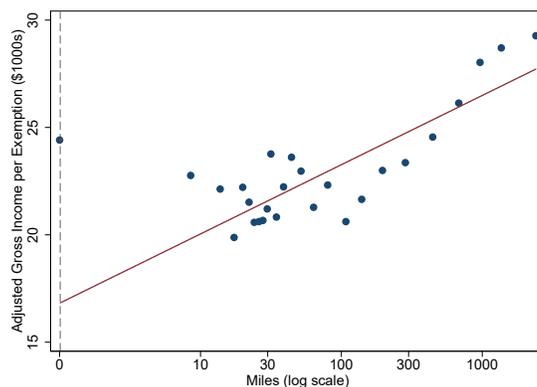
Figure E1. : The correlation between Adjusted Gross Income and distance moved, conditional on moving counties.
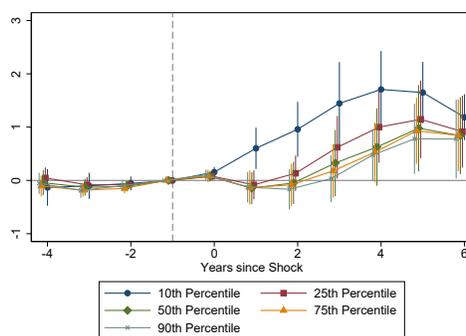

Figure E2. : The effect of a one percent migration shock on the distribution of wages

of cut-off.

Furthermore, a natural implication of this model is that the benefits would accrue more in less highly-skilled communities, assuming migrants are roughly the same skill mix. In Figure E3, I show the opposite is the case; MSAs with a higher percentage of college-educated people (as measured by the ACS in 1990, where I consider anyone with 4 or more years of college to be college-educated) have a larger effect from migration than those without.[12] This result is robust to using any years of college education rather than requiring four. One thing to note is that the college share is negatively correlated to the housing supply elasticity. Controlling for that reduces the difference between the two lines.

---

[12]This exercise assumes the proportion of college educated inmigrants is similar in different cities. Data comes from Ruggles et al. (2019).
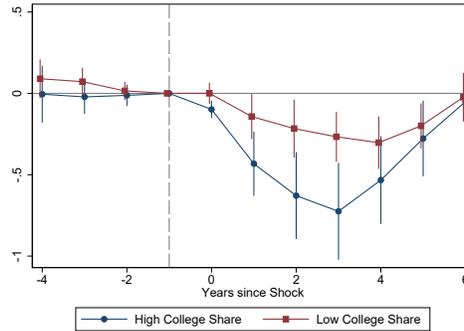
Figure E3. : The effect of a one percent migration shock on the unemployment rate, split by college share.

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

AGGLOMERATION. — A second possibility is that there are agglomeration effects. As migrants move in, knowledge spillovers (Audretsch and Feldman, 2004), the home market effect (Krugman, 1980), thick market externalities (Diamond, 1982), or other increasing returns forces increase employment. However, one might expect tradable goods to be most affected by agglomeration, or at least equally affected. Rather, I find little initial effect on tradable goods, as seen in Figure E4. Second, many agglomeration stories would also imply a larger effect in the long-run than the short-run, and I find the opposite. Finally, the magnitude of my effects is much larger than these forces can explain. Across MSAs in my sample, the biggest cities have, on average, about 0.3 percentage points lower unemployment than the smallest MSAs, despite having populations about 50 times as large. Even if this were all agglomeration effects, I am finding a decline in the unemployment rate of about 0.5 percentage points in response to a one percent increase in population, which is at least two orders of magnitude larger.

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

WEALTH HETEROGENEITY. — A third possibility is that migrants are significantly wealthier than non-migrants. For example, if a retiree moves to Florida, and spends down his savings, he is adding to labor demand but not labor supply. To check whether this is an empirically plausible channel, I use the American Community Survey data, and look at the dividend, interest, and rental income of
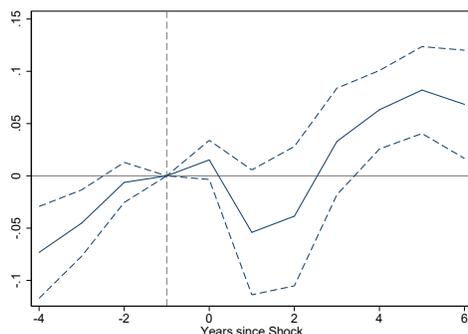
Figure E4. : Tradable goods employment per capita. The effect of an inmigration shock, with 95 percent confidence interval. Errors clustered by state.

interstate migrants versus non-migrants and within-state migrants (Ruggles et al., 2019). The cumulative distribution function of this income is plotted in Figure E5. The first thing to note from this plot is that more than 80 percent of people, migrants and non-migrants, do not report interest income.[13] More crucially, the distribution of interest income for non-migrants first-order stochastically dominates the interest income for migrants. Although the ACS does not measure wealth directly, this is suggestive evidence that non-migrants are wealthier than migrants.

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

In conclusion, none of these other channels—factor complementarity, agglomeration, or wealth heterogeneity—seem likely to be driving the results that I find in the data. But more importantly, none of these channels make the prediction that non-tradables and construction employment would increase while tradable goods employment would fall. Nor would any of these channels have predictions on whether the effect would be stronger in low-housing-supply-elasticity or growing areas. Hence, the evidence that I have shown is strongly supportive of housing causing the decrease in the unemployment rate.

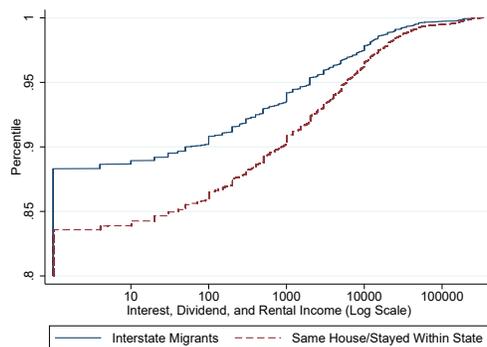[13]A very small fraction report negative interest income.

Figure E5. : The cumulative distribution function of the distribution of interest income, by migration status. American Community Survey, 2000-2014.

*

REFERENCES

Audretsch, David B, and Maryann P Feldman. 2004. "Knowledge spillovers and the geography of innovation." *Handbook of Regional and Urban Economics*, 4: 2713–2739.

Bodvarsson, Örn B, Hendrik F Van den Berg, and Joshua J Lewer. 2008. "Measuring immigration's effects on labor demand: A reexamination of the Mariel Boatlift." *Labour Economics*, 15(4): 560–574.

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

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).

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

De Silva, Dakshina G, Robert P McComb, Young-Kyu Moh, Anita R Schiller, and Andres J Vargas. 2010. "The effect of migration on wages: evidence from a natural experiment." *The American Economic Review: Papers & Proceedings*, 100(2): 321–326.

Diamond, Peter A. 1982. "Wage determination and efficiency in search equilibrium." *The Review of Economic Studies*, 49(2): 217–227.

Gagnon, Etienne, and David Lopez-Salido. 2014. "Small price responses to large demand shocks." Finance and Economics Discussion Series Working Paper.

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

Krugman, Paul. 1980. "Scale economies, product differentiation, and the pattern of trade." *The American Economic Review*, 70(5): 950–959.

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.

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

Ottaviano, Gianmarco, and Giovanni Peri. 2012. "Rethinking the effect of immigration on wages." *Journal of the European Economic Association*, 10(1): 152–197.

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.

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