

Economic Shocks and Internal Migration: Evidence from the Great Recession

Joan Monras*

May 14, 2014

This paper may change in the near future.

Latest version: http://www.columbia.edu/~jm3364/Spatial_mobility_Joan_Monras_v5.pdf

Abstract

Previous literature shows that internal migration rates are strongly procyclical. This would seem to imply that geographic relocation does not help mitigate negative local economic shocks during recessions. This paper shows that this is not the case. I document that net in-migration rates decreased in areas more affected by the Great Recession. Using various IV strategies that rely on the importance of the construction sector and the indebtedness of households before the crisis, I conclude that internal migration might help to alleviate up to one third of the effects of the crisis on wages in the most affected locations. This would be due to a disproportionate decrease in in-migration into those locations rather than an increase in out-migration.

JEL Classification: J61, J20, J30

*Columbia University. Correspondence: jm3364@columbia.edu. I would like to thank Don Davis, Eric Verhoogen and Bernard Salanié for guidance and encouragement and Jón Steinsson and Brian Kovak for useful comments and discussions. All errors are mine.

1 Introduction

It is a common perception that “Americans have historically been an unusually mobile people, constantly seeking better economic conditions” (Moretti, 2012). We would, then, expect geographic relocation to be an important mechanism for American families to deal with periods of economic crisis. Is this the case?

US internal migration rates are strongly pro-cyclical (Molloy and Wozniak, 2011). This could imply that (most) internal migrants move to take advantage of opportunities created during good economic years. It is less clear, however, whether internal migration also helps in bad times. The fact that aggregate migration rates decline in downturns might suggest that US families do not respond to negative economic shocks by moving to other parts of the country and that instead, families remain in their current location until the economy starts recovering. Is this true? Does this imply that spatial labor reallocation does not help mitigate strong negative local economic shocks in recessions?

In this paper, I use the Great Recession to study how migration decisions are shaped by the effects of the crisis at the local level, i.e. across metropolitan areas. Contrary to previous literature – see Mian and Sufi (2013) and Molloy and Wozniak (2011) – I show that geographic relocation is important in mitigating negative local economic shocks during the period 2006-2011. In particular, I show that the relocation mechanism is decreased in-migration rates into negatively hit locations, rather than out-migration from them.

To understand why reduced in-migration rates are important in mitigating the local effects during recessions it is, first, important to realise that while the average net migration rates across metropolitan areas are close to 0, the gross flows are significantly higher. More than 3.5 percent of households change metropolitan areas in any given year.¹ Second, the decreased migration rates during recessions are, partly, a consequence of fewer people moving towards negatively shocked local labor markets. Third, this implies that the relevant counterfactual, i.e. what would have happened had the Great Recession not occurred, needs to take into account that more people would have moved to the negatively hit locations. This is crucial. In a closely related paper, Yagan (2014) shows how workers who suffered larger local shocks in 2006 still have worse outcomes through 2011, even if they relocated. He takes this as evidence that internal relocation does little to mitigate the negative consequences of local shocks. My results show that fewer workers moved into the locations most affected by the Great Recession. Had they done so, the wage and employment effects would have been even worse.² Adding this new element gives a more positive picture of the ability of

¹More generally, as computed in Molloy et al. (2011), around 1.5 percent of the population moves between Census regions, 1.3 percent move across states within these regions and around 3 percent move across counties within state. These numbers are obviously consistent with the reported migration rates across metropolitan areas.

²In Monras (2013) I show how the local labor demand elasticity is around -1, using the unexpected inflow of Mexicans resulting from the Mexican Peso Crisis of 1995.

internal migration to smooth local shocks.³

In this paper, I show these results using two alternative strategies. My first strategy is to document that locations with larger construction sectors before the crisis (measured by the share of employment in construction in 2000) suffered larger local economic shocks in 2008. These locations, in turn, saw their in-migration rates decline disproportionately.

My second strategy builds on the work by Mian and Sufi (2013). They argue that the 2008 crisis lowered the consumption capacity of highly indebted households. This helps to explain the drop in consumption starting in 2008. Across the territory, the demand for tradables dropped uniformly, while the demand for non-tradables dropped more in highly indebted locations, precisely because non-tradable goods can only be consumed locally. This translates, Mian and Sufi (2013) show, into declines in non-tradable employment in highly indebted locations, while it translates into uniform drops in employment in tradable sectors. Thus, locations that historically had larger employment shares in non-tradable employment *and* were highly indebted at the beginning saw sharper declines in per capita GDP and wages and higher increases in unemployment rates.

I use these strategies to estimate the elasticity of in-migration and out-migration rates to per capita GDP, unemployment rate and wage changes at the local level. This is informative as to how responsive internal migration is to changes in local economic conditions. Both strategies deliver similar results. I find that a 1 percent decrease in per capita GDP leads to around .14 percentage point decrease in the share of in-migrants at the local level. This means that the share of people that were previously living in another metropolitan area (which is usually around 3.00 percent) decreases to around 2.86 percent. In New York, the largest metropolitan area in the US, this means that around 30,000 fewer people move into New York City if GDP per capita drops by 1 percent. I do not find, however, that out-migration rates, i.e. people that leave a certain metropolitan area, significantly increase in areas more severely hit by the crisis. This elasticity is quite similar for high- and low-skilled workers. It is noteworthy, however, that low-skilled workers seem more responsive to unemployment rates, while high-skilled workers are more responsive to local wages.

Similarly, I find that a 1 percent decrease in wages leads to around .30 percentage point decrease in the share of in-migrants at the local level. This helps to mitigate the shock that the crisis has on wages. If the elasticity of the local labor demand is equal to 1 (see Monras (2013) for an estimate of this order of magnitude) then the decrease in in-migration rates reduces the effect of the crisis on wages by around one third. This means that had these in-migrants moved into a city with an actual decrease of 1 percent in wages, the decrease would have been around 1.3 percent.

These results imply that geographic relocation (i.e. net migration) took place as a response to particularly strong negative local economic shocks during the Great Recession. In other words, the populations of more heavily hit locations grew less than before the crisis. This was a consequence

³As emphasized in Topel (1986), the migrants that did not move to the negatively affected locations are themselves labor supply shocks that affect wages in location less affected by the crisis.

of lower in-migration rates rather than out-migration.

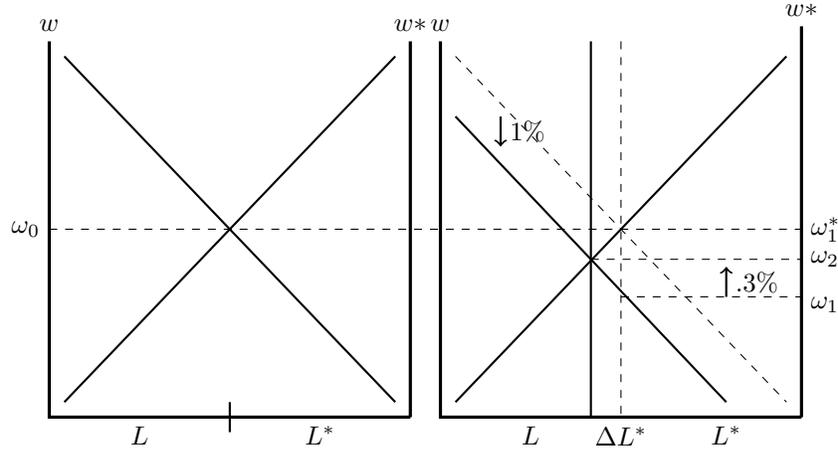
Several papers have looked at the response of local labor markets to negative local labor demand shocks. In a seminal contribution, Blanchard and Katz (1992) argue that locations affected by negative shocks experience permanent losses in employment, temporary increases in unemployment rates and temporary decreases in local wages. Their identification strategy does not distinguish between good and bad times, however. This has become an important issue because several papers have argued that during this last recession internal migration rates have not responded differently across multiple locations. Mian and Sufi (2013) argue that, if anything, populations increased (between 2007 and 2009) in highly indebted counties, rather puzzling for the spatial economic literature (see Glaeser (2008) or Moretti (2011)): with free mobility, people should leave declining locations and move to rising ones. Their findings are explained by the fact that highly indebted counties were attracting more population than other counties before 2008 and stopped doing so as a consequence of the crisis.

In a paper which is closely related to this one, Cadena and Kovak (2013) show that locations more heavily affected by the crisis (using either the construction sector or Mian and Sufi (2013)'s insights as I do) saw declines in Mexican workers but low internal migration responses amongst low-skilled native workers. They do not investigate the possibility that worse affected locations were receiving (perhaps unusually) more people before the crisis, particularly of low skilled workers. This is important. When I account for this, i.e. when I examine the relevant migration rates, I find that natives and immigrants responded similarly to local economic shocks.⁴ In fact, using the same strategies but excluding the non-native population from the computation of migration rates does not change any of the results. Combined with the research presented in this paper, it seems that natives respond by not migrating to affected locations, while Mexicans are more ready to leave from them.

More broadly, this paper is related to the spatial economic literature. Unlike most papers in this literature, I investigate the short-term responses. It is well documented that places where amenities or local labor market conditions improve attract more people (see Glaeser (2008) or more recently Diamond (2013)) while locations with deteriorating local conditions suffer from a shrinking population (see Hornbeck (2012), Hornbeck and Naidu (2012) or Autor et al. (Forthcoming)). It has also been documented that people respond to expectations of future positive prospects (Kennan and Walker, 2011). In line with Glaeser and Gyourko (2005) or Notowidigdo (2013), I show that internal local out-migration rates do not respond swiftly to local economic shocks. However, I also emphasize the importance of the response of in-migration rates into these places.

⁴At the end of the paper I explain carefully why it is possible to obtain Cadena and Kovak (2013) results and still conclude that low skilled natives do respond to economic shocks.

Figure 1: Theoretical framework



2 Theoretical Framework

The theoretical framework needed for my empirical investigation is rather simple. Essentially I need a model with downward-sloping local labor demand in the various local economies of a larger national market. There are a number of ways to accomplish this. For example, if the production technology produces constant returns to scale in the two factors of production - land and labor - and land is assumed to be immobile while labor is mobile, we obtain the desired property of a downward -sloping demand for labor. Alternatively, one can think that each region produces a differentiated product and that technologies are linear in labor in each region (see Blanchard and Katz (1992)). In both cases, we need labor to be mobile across regions.

Labor mobility ensures that the indirect utility equalizes across locations. This could be simply the wage, or it could be the wage net of the amenities and local price indexes. In this set-up, a local demand shock is simply a decrease in the demand for labor in one of the regions. This leads to a decrease in wages in the region receiving the shock, and some relocation of labor from the affected location towards other regions. This restores indirect utility equalization across regions.

This model can be easily represented in Figure 1, where I depict the simpler two-region case. To guide the discussion I denote the two regions as ‘home’ and ‘foreign’. In the first instance the market wage in both regions is ω_0 . This is where ‘home’ demand for labor meets ‘foreign’ demand for labor. From the perspective of ‘home’, the labor demand for ‘foreign’ is its labor supply. This is depicted in the right panel of Figure 1.

When ‘home’ receives a negative labor demand shock, the demand shifts to the left, causing wages to decrease. Assume workers take one period to relocate. Assume that the decrease in wages

is 1%. In this situation, when the shock hits, wages at ‘home’ decrease by 1%, while wages stay at the old level abroad. This causes net in-migration to decrease at home and to increase abroad, bringing the wage to ω_2 . From the perspective of ‘home’, labor relocation mitigates the initial impact, spreading it to the other region.

My estimates imply that the net in-migration rates decrease by around .3% for a 1% decrease in wages. In order to ascertain how much this mitigates the effect on wages one needs to know the (short-run) elasticity of labor demand. Estimates from Monras (2013) suggest that this elasticity is around 1-2. This implies that if .3% fewer people move to the negatively hit location, wages in that location recover by around .3 – .6%. This is around one third of the initial level. Thus, inter-regional mobility spreads the negative effects of the crisis to the broader economy, limiting the impact in hard-hit locations but making it worse for initially mildly hit locations.

In line with the empirical results and the model presented in Monras (2013) the empirical exercise in this paper shows that this is a consequence of decreases in in-migration rates to hard-hit locations. The theoretical intuition behind this result is the following.⁵ When workers are constantly deciding where to live (something that can be modelled as a discrete choice problem where individual workers have idiosyncratic tastes over the different regions) gross flows are larger than net flows.⁶ In the model, the probability that a worker residing in a location s decides to move to s' depends on the wage and amenities in s' relative to the wages and amenities in all other possible locations – which can be summarized as the state of the national economy. The distribution of idiosyncratic tastes ensures that there is always someone who finds attractive to move to s' . This makes gross flows larger than net flows.

When a single location s' (or a small set of locations) suffers a negative shock, this decreases the probability that all the workers not living in s' are attracted to s' , decreasing the in-migration rates to s' . Instead, the probability of people in s' leaving s' is only affected because the change in wages in s' affects the general state of the national economy. If there are many regions in the national economy, the effect of a change in wages in s' on the general state of the economy is small. This makes out-migration rates respond less than in-migration rates. This is what I document in the empirical part of the paper.

3 Empirics

3.1 Data

For this paper I employ to two main data sources. I use the American Community Survey data from Ruggles et al. (2008) to compute migration rates across US metropolitan areas, and in par-

⁵See Monras (2013) the full model and an analysis of the properties of the model.

⁶A result also obtained in (Artuç et al., 2010), among others.

ticular information on the current and past residents' locations to construct in-migration rates, out-migrations and net migration rates. I also use the ACS data to compute unemployment rates and averages wages across metropolitan areas. My second source is the Bureau of Economic Analysis, whose data allow me to obtain a measure of real gross domestic product per capita.

More concretely, I define the in-migration rate to metropolitan area m at time t as follows:

$$in - migration_rate_{m,t} = \sum_{i \in I_{m,t}} mover_{m,t}^i / I_{m,t} \quad (1)$$

Where $I_{m,t}$ denotes the set and the number of individuals that live in m at time t , $mover_{m,t}^i$ is a dummy variable that takes value 1 if individual i lives in city m at t but did not live in m at $t - 1$.⁷

Similarly, I define the out-migration rate from a metropolitan area m as:

$$out - migration_rate_{m,t} = \sum_{k \in M - \{m\}} \sum_{i \in I_{k,t}} mover_{m,k,t}^i / I_{m,t} \quad (2)$$

where, as before, $I_{k,t}$ denotes the set and the number of residents in city k , $mover_{k,m,t}^i$ denotes that individual i was living in m at $t - 1$ and is now living in k and where M is the set of all metropolitan areas in the US. $M - \{m\}$ is the set of all metropolitan areas except for m .

The net migration rate is simply the in-migration rate minus the out-migration rate.

If we limit the count of individuals to people of a certain level of education we obtain the in-migration, out-migration and net migration rates of individuals of education e . In particular, I use a simple distinction between high- and low-skilled workers. High-skilled workers are defined as those who have attended college (SC), graduated from college (CG) or attended graduate school (GS). The low-skilled are high-school drop-outs (HSDO) and high-school graduates (HS) (see Katz and Murphy (1992), Autor and Katz (1999), Acemoglu and Autor (2011) or Card (2009) for papers using similar classifications).

One limitation of the data set is that I only possess information on metropolitan areas of residence from 2005 until 2011. Before that, ACS reports only the state of residence and the state of residence in the previous year. While those could be used to define local labor markets, metropolitan areas are a much better approximation of a local labor market.⁸ An alternative would have been to use CPS data, in which both state and metropolitan areas are reported. The use of CPS data, however, is limited by its small sample size. Furthermore, concerns have been

⁷I use the variable `metarea` and `migmat1` from Ruggles et al. (2008). I do not use the observation where the metropolitan area is not identified.

⁸Autor and Dorn (2009) define local labor markets by Commuting Zones in order to include the entire territory of the US; this is one limitation of using metropolitan areas. In this case, I have limited my analysis to the migration rates of the metropolitan areas.

raised about how the US Census Bureau deals with missing data.⁹ The sample size is particularly important when studying yearly migration rates since these are usually below 4 percent. By using ACS data I can use more than 250 metropolitan areas, whereas it would be hard to work with more than 50 metropolitan areas using CPS. For a detailed discussion of the data sources available to study internal migration, see Molloy et al. (2011). They argue that recent internal migration is best estimated using ACS data.

I have merged two data sets by combining the metropolitan area definitions from the BEA and ACS. The place-names usually coincide perfectly. Sometimes the naming differs slightly but it is always clear on inspection when two different names are referring to the same area. Very occasionally, the aggregation is slightly different across data sets. In those cases I needed to use the more aggregate definition of the two data sets. In total, I obtained 263 metropolitan areas of various sizes and characteristics.

3.2 Summary Statistics

Although life-long migration rates are relatively high in the US, year-on-year migration rates are more modest (Molloy et al. (2011)). In a typical metropolitan area, around 3.5 percent of residents lived in a different location the previous year. In fact, migration rates have declined in the last 20 years or so, as documented in Molloy et al. (2011). This decline in migration rates continued in the 2000s, as can be seen in Table 1 when we compare migration rates before and after 2008.

There is, however, some heterogeneity in how many people in-migrate or out-migrate from various metropolitan areas. There are some extreme examples, usually in college towns like Bloomington, Indiana or Bryan-College Station, Texas, which have in-migration rates consistently above 15 percent. The bulk of metropolitan areas, however, are not far from the average 3.5 percent. The same data can be witnessed in out-migration rates. In fact, it is usually the case that high in-migration metropolitan areas are also high out-migration metropolitan areas, again, for example, many of the college towns in the US.

Unsurprisingly, net migration rates are not always close to 0. Some metropolitan areas are attracting more people than the population they are losing to other metropolitan areas, contributing to their population growth. Figure 2 shows the distribution of net migration rates across metropolitan areas between 2005 and 2010. We see that around half of the metropolitan areas have positive net migration rates, while the other half have negative net migration rates. It is also the case that in every year there are some people leaving and some people moving into every metropolitan area.¹⁰ As can be seen, gross migration flows are larger than net migration flows.¹¹

⁹Molloy et al. (2011) reports lower migration rates in CPS than in ACS, something that is explained in Kaplan and Schulhofer-Wohl (2012) as an undocumented error in the Census Bureau's imputation procedure for dealing with missing data in the Current Population Survey.

¹⁰The zero out-migration rates in Table 1 are due to sample size, and are probably not true zeros.

¹¹This is true both across locations and also industries Molloy and Wozniak (2011), Artuç et al. (2010)

Table 1: Summary statistics

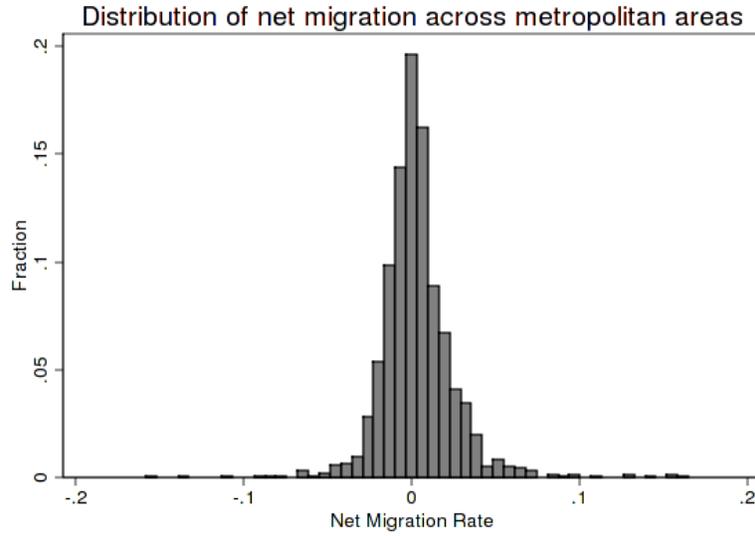
Variable	N.Obs	2006-2010			
		Mean	Std. Dev.	Min	Max
year	1578			2005	2010
population	1578	530500	1069459	58743	1.13e+07
Microdata Obs.	1578	1833.619	1808.536	4 6483	
in-migration rate	1578	.037	.018	.003	.192
in-migration rate, low	1578	.033	.022	0	.24
in-migration rate, high	1578	.041	.017	0	.17
out-migration rate	1578	.037	.014	0	.18
net in-migration rate	1578	0	.015	-.16	.16
(ln) GDP pc	1578	10.71	.27	9.67	11.42
Unemploy. rate	1578	.08	.026	.013	.27
Wage	1578	6.18	.14	5.70	6.59
After 2008					
in-migration rate	789	.035	.018	.003	.19
out-migration rate	789	.035	.013	0	.16
net in-migration rate	789	0	.013	-.11	.16
(ln) GDP pc	789	10.70	.28	9.67	11.42
Unemploy. rate	789	.09	.029	.025	.27
Wage	789	6.16	.14	5.73	6.53

Notes: Those are summary statistics for 263 metropolitan areas between 2006-2010 using ACS data. 'share construction in 2000' uses Census 2000 data. Statistics are computed for working age population.

In terms of per capita GDP, we observe in Table 1 that real GDP per capita remained at around \$44,350 per capita, both for the entire decade of 2000-2010 and for the post-2008 years. This is mainly explained by the crisis. In 2008 and 2009, GDP per capita decreased on average in most metropolitan areas. It increased in all other years. Figure 3 shows the per capita GDP growth rates across metropolitan areas in 2008 and 2009 and all the other years. We can see that in most metropolitan areas GDP per capita declined in 2008 and 2009 whereas it grew in the other years.

GDP per capita is not the only measure of local level economic activity that is relevant for mobility. Unemployment rates and local wages are also very important. Several authors have documented downward nominal wage rigidity during the great recession. Real wages, however, decrease in some locations, even in normal times. Graph 4 shows that in fact more than half of the metropolitan areas experience negative real wage growth rates. In part this is explained by the different evolution of price indexes at the local level, that does not necessarily follow the nationwide price index used to compute real wages (see Handbury (2011)). During the 2008-2009 period the

Figure 2: Net migration rates across metropolitan areas



Notes: Net migration is computed from the people that move into and out of each metropolitan area. I have used data from 263 different metropolitan areas between 2006 and 2010. Source: ACS data from Ruggles et al. (2008).

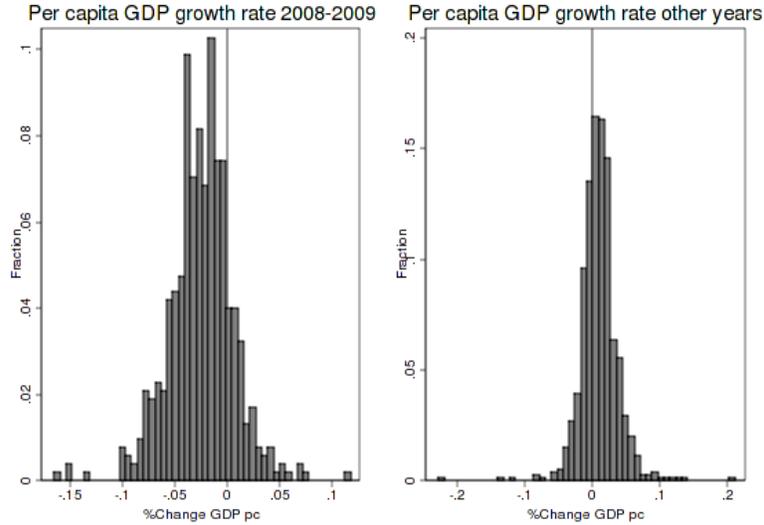
share and the magnitude of wage reductions increased.¹²

A similar picture is obtained when we look at unemployment rates, see Figure 5. Unemployment rates increased in most metropolitan areas. In fact, the downward nominal wage rigidity made unemployment rates to change more than wages.

All three variables move very similarly. In fact if we look at the elasticities between GDP per capita, wages and unemployment rates (controlling for metropolitan area fixed effects and time dummies), we see that the most responsive was the unemployment rate, followed by wages. Table 2 shows these results.

¹²See the Appendix for the breakdown between high and low-skilled wages.

Figure 3: GDP per capita growth rates across metropolitan areas



Notes: This figure shows the per capita real GDP growth rates in all 263 metropolitan areas. Vertical lines distinguish positive from negative growth rates. Source: BEA.

Table 2: The covariance between GDP per capita, wage and unemployment rate

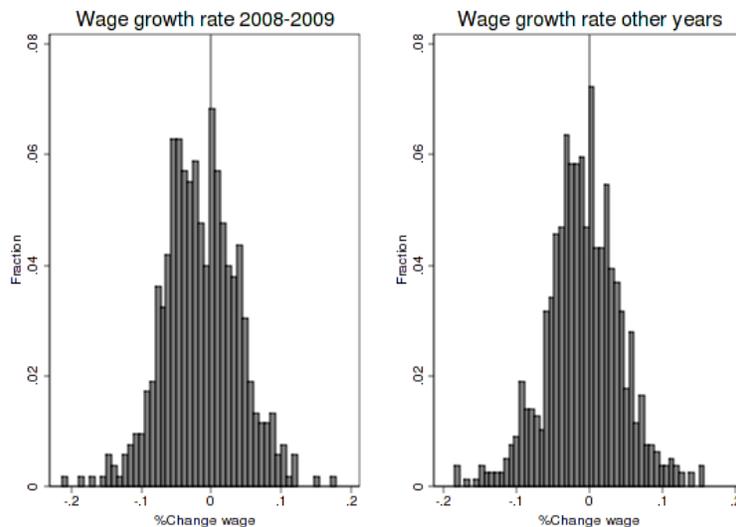
	(ln) GDP pc	(ln) wage	Unemployment rate
(ln) GDP pc	-	0.38***	-0.018***
(ln) wage	1.86***	-	-0.026***
Unemployment Rate	-5.46***	-1.55***	-
msa fe	yes	yes	yes
time fe	yes	yes	yes
N	1315	1315	1315

Notes: This table shows the elasticities between GDP per capita, wages and unemployment rates. Each elasticity is computed from an OLS regression. Only one variable and the fixed effects are included as explanatory variables in the OLS regressions. * $p < .1$, ** $p < .05$ and *** $p < .001$.

3.3 Main Results

The main objective of this paper is to estimate the (short term) response to local economic shocks. This helps to answer the questions of how responsive migration flows are to local economic shocks and how much local economic shocks spread across the economy through the internal migration decisions of workers.

Figure 4: Wage growth rates across metropolitan areas



Notes: This figure shows average real wage growth rates in all 263 metropolitan areas. Vertical lines distinguish positive from negative growth rates. Source: Ruggles et al. (2008).

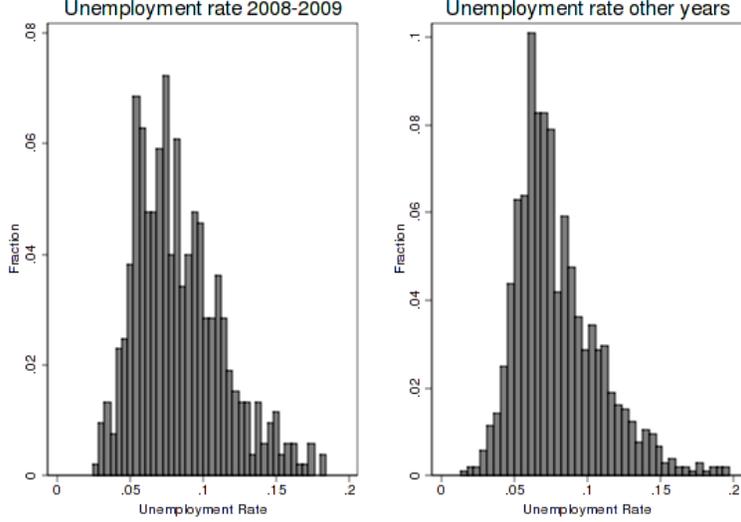
Before using the regression framework to look at the mechanisms of how this reallocation is taking place it is worth taking a look at some of evidence suggesting that internal reallocation does react to economic shocks.

In absence of local shocks it is hard to know how internal migration reacts to changes in local market conditions. An easy way to see this is by plotting the change in population against the change in wages or the lagged change in wages. Figure 6 shows that there is no clear correlation between changes in wages and posterior changes in population levels. This is true both in booming years, i.e. pre 2008, and post 2008. This Figure could be taken as evidence that population does not respond to local wages. It hides, however, two important aspects that I uncover in this paper.

First, there are important differences at the local level between high- and low-skilled workers. In particular, changes in wages of both types of workers do not necessarily coincide and it is important to distinguish between these two type of workers. This can easily be seen in the Figure 7. When comparing pre and post 2008 periods we observe that the relation between population changes and unemployment rates becomes more negative – places with high increases in unemployment loose population, while the relation between population changes and wage changes becomes more positive – places less hit by the crisis gain population.

Figure 8 shows that similar patterns are observed for high skilled workers.

Figure 5: Unemployment rates across metropolitan areas



Notes: This figure shows unemployment rates in all 263 metropolitan areas. Source: Ruggles et al. (2008).

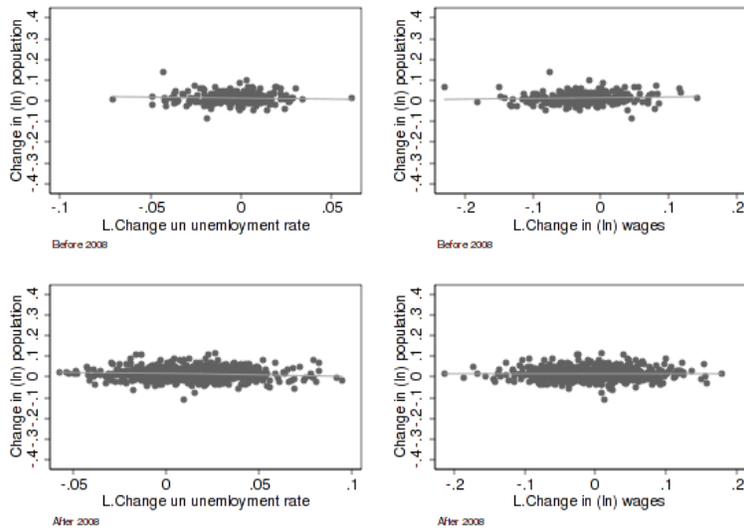
Figures 7 and 8 do not show what is the mechanism through which this labor relocation is taking place. To see this I move to a regression framework where I explain in- and out- migration rates as a function of local economic conditions. To this end I use the following specification:

$$migration_rate_{m,t} = \beta X_{m,t} + \delta_m + \delta_t + \varepsilon_{m,t}^e \quad (3)$$

where $migration_rate_{m,t}$ is either the number of people that move into metropolitan area m (divided by the population in that area), the number of people that move out of metropolitan area m or the net in migration to metropolitan m . When these rates are limited to specific educational groups e it is always specified. δ_m are metropolitan area (MSA) fixed effects, while δ_t are year fixed effects. I use the years 2005-2010, both included. $X_{m,t}$ is a measure of local economic activity. I use the three aforementioned measures: (log) GPD per capita, average (log) wage and unemployment rate.

My main goal is to estimate this elasticity of migration rates to local economic conditions. I could run this regression using OLS, but I might get biased estimates if migration to certain places is also affecting the local market outcomes. There are many reasons why this could be the case. More people in a particular market might put downward pressure in wages, but it may also be that the higher demand in this market is attracting more workers.

Figure 6: Changes in population against lagged changes in wages



Notes This figure plots changes in population against lagged changes in wages across metropolitan areas, strictly before and after 2008. Source: (Ruggles et al., 2008).

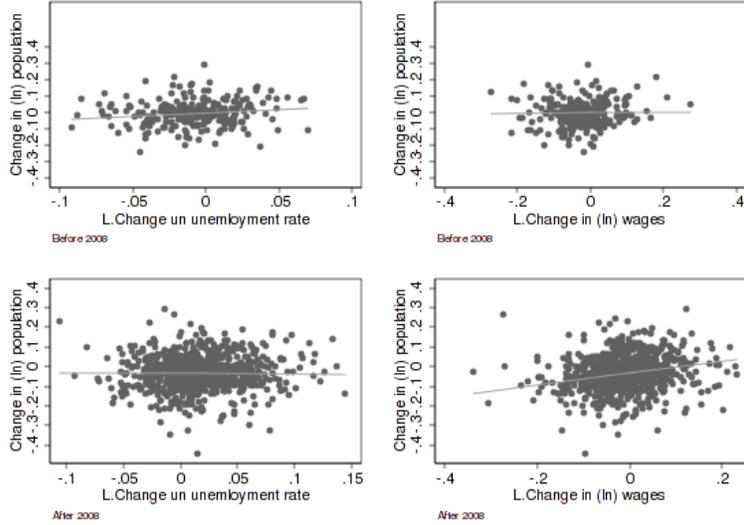
As said in the introduction, I use two alternative strategies to estimate this elasticity. Both rely on the unexpectedness of the current crisis. First, I use the importance of the construction sector and second, the level of indebtedness of the households. This latter measure is directly taken from Mian and Sufi (2013) and is defined as the debt to income ratio of households in a number of US counties. To obtain a measure of indebtedness at the metropolitan area I simply take the weighted (by population) average of the different counties' debt to income ratio whenever the metropolitan area has more than the county.

To make sure that these two measures are indeed good predictors of where the crisis hit hard in 2008 I first present evidence on these measures and the three local economic variables of interest, i.e. GDP per capita, unemployment rates and wages. Once I have shown this I report the elasticity of migration rates on changes in these local economic variables due to the crisis.

3.3.1 Local economic variables and the crisis

This section shows how much the crisis can be linked to either the construction sector or the indebtedness of the households by looking at particular local measures of economic activity. I do

Figure 7: Changes in population against lagged changes in wages, low skilled



Notes This figure plots changes in population against lagged changes in wages across metropolitan areas, strictly before and after 2008. Source: (Ruggles et al., 2008).

so by running the following regressions:

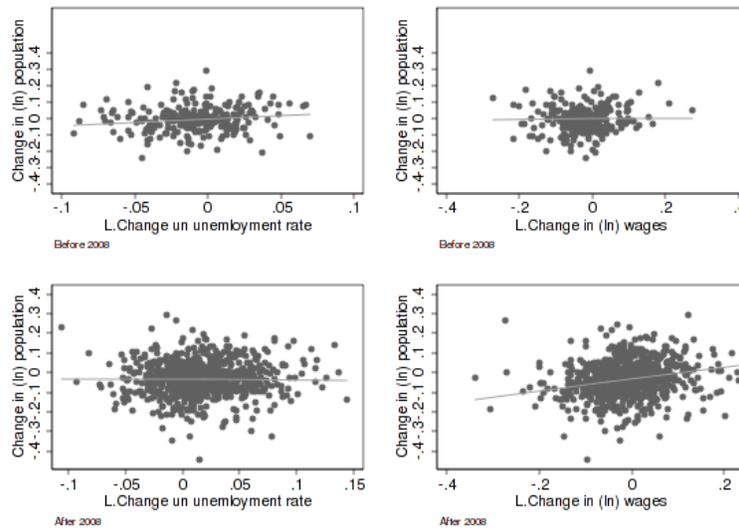
$$X_{m,t} = \beta * shock_t * Z_{m,T} + \delta_m + \delta_t + \eta_{m,t} \tag{4}$$

where $X_{m,t}$ is either per capita GDP, unemployment rate or wage in metropolitan area m at time t , $shock_t$ is a dummy variable that takes value 1 after 2008 and where $Z_{m,T}$ is either the number of workers in construction divided by the total number of workers in metropolitan area m in 2000 (Construction); the debt to income ratio in 2006 (HH debt); or the interaction of the debt to income ratio in 2006 with the share of workers in non-tradable sectors (AD employ.)¹³. δ_m are metropolitan area fixed effects, while δ_t are year fixed effects. Since I also break down unemployment and wages between high and low-skilled workers, this means that I run 21 separate regressions.

Table 3 shows the results of running these regressions. There are at least three remarkable findings. First, all three measures that I use to see what the most affected metropolitan areas were when the crisis hit in 2008 are well correlated with per capita GDP and unemployment rates. Wages, in contrast, seem not to react so much to construction.

¹³I follow Mian and Sufi (2013) to define non tradable sectors. In parenthesis is the variable name in the tables.

Figure 8: Changes in population against lagged changes in wages, high skilled



Notes This figure plots changes in population against lagged changes in wages across metropolitan areas, strictly before and after 2008. Source: (Ruggles et al., 2008).

Second, the measures related to the indebtedness of the households are much more closely correlated to decreases in GDP, increases in unemployment and decreases in wages. If the aggregate demand channel is as important as Mian and Sufi (2013) argue this is exactly what we would expect to find. This means that we can use these measures of how hard the crisis hit to compute the elasticities of migration rates. I do this in the following subsection.

Third, I see how the shock affected the unemployment rates of low-skilled workers relatively more than those of high-skilled workers, while it affected wages for high-skilled workers relatively more than it did for low-skilled workers. These results suggest, thus, that high-skilled workers are more rigid for low-skilled workers than for high-skilled ones and that adjustments take place in quantities rather than in prices for the former group.

3.3.2 Internal migration rates and the crisis

Perhaps the most natural measure of local economic activity is the GDP per capita in different metropolitan areas. Following the overall economy, most metropolitan areas saw decreases in per capita GDP in 2008, as previously documented. Metropolitan areas more dependent on the construction sector and metropolitan areas whose households were more indebted suffered the crisis disproportionately. Previous literature reports that, surprisingly, it seems that *more* people seemed

Table 3: First Stage: Construction sector, Household Debt and Aggregate Demand

Skill level:	ln GDP pc	Unemployment Rate			Wages		
	All	All	Low	High	All	Low	High
Construction Shock	-1.613*** 0.349	0.332** 0.163	0.489** 0.219	0.146 0.115	0.046 0.199	0.196 0.231	-0.158 0.200
HH debt Shock	-0.020*** 0.007	0.011*** 0.002	0.014*** 0.003	0.008*** 0.001	-0.007*** 0.002	-0.007** 0.003	-0.010*** 0.002
AD employ. Shock	-0.097*** 0.022	0.045*** 0.006	0.056*** 0.009	0.031*** 0.005	-0.030*** 0.008	-0.026*** 0.009	-0.042*** 0.008
msa fe	yes	yes	yes	yes	yes	yes	yes
time fe	yes	yes	yes	yes	yes	yes	yes

Notes: The dependent variable is (log) GDP per capita, the unemployment rate and average wages in metropolitan areas. The table shows the results of running 7 x 3 different regressions. 'Construction shock' is the interaction of the share of construction workers in 2000 with a dummy taking value 1 in the years after the beginning of the Great Recession, i.e. 2008-10. 'HH debt shock' is the interaction of Mian and Sufi (2013) measure of Household indebtedness with a dummy taking value 1 in the years after the beginning of the Great Recession. 'HH debt shock' is the interaction of Mian and Sufi (2013) measure of Household indebtedness with a dummy taking value 1 in the years after the beginning of the Great Recession. 'AD employ shock' is the interaction of Mian and Sufi (2013) measure of Household indebtedness with the share of workers in the non tradable sector and a dummy taking value 1 in the years after the beginning of the Great Recession. Regressions are weighted by the number of observations in each metropolitan area. Number of observations: 263 metropolitan areas x 6 years = 1578 or 194 metropolitan areas x 6 years = 1164 when Mian and Sufi (2013) measure used. * p<.1, ** p<.05 and *** p<.001.

to move to these metropolitan areas (Mian and Sufi, 2013).

When thinking about population changes it may be worth taking a wider perspective. Many cities grow over long periods. A crisis in a given city might affect its trend without necessarily implying that the city will necessarily lose population. To see this econometrically one would look at the evolution of the population or look at population growth rates across metropolitan areas, as some papers have done (see Mian and Sufi (2013) and Cadena and Kovak (2013)). There are a few things worth noting in terms of why this exercise might give puzzling results.

First, as can be seen in the Table 4 locations heavily dependent on the construction sector and which were (to some extent, though much less) more leveraged were in fact attracting more people than the average metropolitan area. This may be because the construction sector or the local consumption boom attracted many people looking for jobs or simply because for other unrelated reasons these were metropolitan areas with higher in-migration rates.¹⁴ What the crisis did was to decrease the rate at which these cities were attracting population. It could well be the case, however, that these cities still attracted more people than other cities, despite the fact that they attracted less people than they used to. Thus, when doing the cross-sectional comparison of population growth

¹⁴Note that in all other dimensions metropolitan areas do not differ significantly.

rates we might wrongly conclude that migration did not respond to the crisis when it actually did respond.

Table 4: Comparing different metropolitan areas

	High construction MSAs	Low construction MSAs
Population	523,060	731,542
(ln) GDP pc	10.669	10.818
Unemploy. rate	0.065	0.067
In-migration rate	0.048	0.033
Out-migration rate	0.039	0.037
Net in-migration rate	0.009	-0.004
Share construction in 2000	0.079	0.057
Debt to income ratio	1.793	1.789
AD employment	0.226	0.202
	High leveraged MSAs	Low leveraged MSAs
Population	751,824	502,778
(ln) GDP pc	10.739	10.755
Unemploy. rate	0.065	0.067
In-migration rate	0.042	0.039
Out-migration rate	0.040	0.036
Net in-migration rate	0.002	0.002
Share construction in 2000	0.068	0.067
Debt to income ratio	2.297	1.035
AD employment	0.222	0.201
	High AD employ. MSAs	Low AD employ. MSAs
Population	742,330	512,271
(ln) GDP pc	10.739	10.755
Unemploy. rate	0.065	0.067
In-migration rate	0.042	0.039
Out-migration rate	0.040	0.037
Net in-migration rate	0.002	0.014
Share construction in 2000	0.068	0.015
Debt to income ratio	2.305	0.643
AD employment	0.223	0.045
Number of MSAs	97	97

Notes: This table shows the averages of selected variables splitting MSAs by high/low construction, debt to income ratio and AD employment. AD employment is the interaction of the share of workers in non-tradable sectors in 2000 with the debt to income ratio. The 194 MSAs are always split in two groups of 97. The year is 2006. Source: ACS data and Mian and Sufi (2013).

Second, the population level estimates obtained using ACS or CPS heavily depend on the individual weights assigned to the individual level observations. These weights are meant to make the individuals representative at the local level within a given year. Making the comparison across

years might be slightly more problematic as the way the weights are assigned might change from year to year. Thus, when using population growth rates it is important to be sure that the growth rates are true population changes and not due to any changes in weighting.

A more direct way to look at migration responses is to look at migration rates. This is what I show in Table 5. The results are clear. Net in-migration rates decrease when per capita GDP decreases. A 1 percent decrease in GDP per capita leads to .8-.17 percentage points decrease in the net in-migration rate, as can be seen in Panel C, Table 5. Similarly 1 pp increase in the unemployment rate leads to a .14-.2 pp decrease in net in-migration while a 1 percent decrease in the average wage leads to a .2-.3 pp decrease.

These responses of net migration rates are entirely due to in-migration rates, as shown in Panel A, Table 5. A metropolitan area with a typical in-migration rate of around 3.5 percent would see the in-migration rate drop to around 3.3-3.4 percent as a result of a 1 percent decrease in per capita GDP. In more concrete numbers, this represents .14 percent of the population of any given city. Thus, around 9,000 less people would move into New York City if GDP per capita in New York were to drop by 1 percent. We observe that the estimate from using the IV specification is 3 times larger than the OLS estimate. We also see that the adjustment to the crisis takes place through reductions in in-migration rates, rather than increases in out-migration rates, see Panel B, Table 5.

Cadena and Kovak (2013) suggest that there may be important differences between natives and immigrants. I investigate this possibility using the migration rates of natives only.¹⁵ When doing this, we observe that even when restricting the computation of migration rates to native born individuals, the in-migration rates decrease when the per capita GDP decreases in a location. Thus, the response to the crisis, i.e. fewer people moving to hard hit locations, is the same among natives and immigrants. This can be seen in Table 6.

¹⁵In the next section I also look at this distinction by skill groups.

Table 5: The migration response to the crisis: In-migration rates, total population

Dep. Var. :	Panel A: in-migration rate									
	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	-0.255***	-0.002*	-0.010***	0.158***	0.085***	0.104***	-0.155**	-0.226***	0.238**	0.333***
	0.038	0.001	0.004	0.025	0.029	0.024	0.065	0.077	0.110	0.124
Dep. Var. :	Panel B: out-migration rate									
	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	0.020	-0.000	-0.001	-0.012	0.008	0.013	-0.015	-0.027	0.023	0.040
	0.025	-0.000	0.002	0.014	0.030	0.022	0.054	0.045	0.083	0.067
Dep. Var. :	Panel C: net-in-migration rate									
	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	-0.275***	-0.002	-0.009*	0.171***	0.077*	0.091**	-0.140	-0.199*	0.215	0.293*
	0.052	0.001	0.005	0.030	0.046	0.036	0.097	0.105	0.155	0.163
Exp. Var	Constr. Shock	HH debt Shock	AD employ. Shock	GDP pc			Unemp. rate		Wage	
Instrument				Construction	HH debt	AD emp.	HH debt	AD emp.	HH debt	AD emp.
F-stat				21.372	8.718	19.065	34.333	48.799	11.718	15.361
msa fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
time fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes

Notes: The dependent variable is the in-migration rate, the out-migration rate or the net-in-migration rate, in panels A, B and C respectively. The independent variable is either (log) GDP per capita, unemployment rates or average wages in US metropolitan areas between 2005 and 2010, as indicated. (log) GDP per capita, unemployment rates or wages are instrument with the variable the construction importance before the crisis, and the debt to income ratio introduced in Mian and Sufi (2013), see more details in table 3 or in the text. Regressions are weighted by the number of observations used to compute the shares of construction. F-stats reported are the F-stats of excluded instruments in the first stage regression. Robust standard errors clustered at the metropolitan area level are reported. Number of observations: 263 metropolitan areas x 6 years = 1578 or 194 metropolitan areas x 6 years = 1164 when Mian and Sufi (2013) measure used. * p<.1, ** p<.05 and *** p<.001.

Table 6: The migration response to the crisis: Native In-migration rates, total population

Dep. Var. :	in-migration rate, natives									
	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	-0.233***	-0.001	-0.008**	0.144***	0.066**	0.086***	-0.155**	-0.226***	0.238**	0.333***
	0.034	-0.001	0.003	0.028	0.030	0.024	0.065	0.077	0.110	0.124
Exp. Var	Constr. Shock	HH debt Shock	AD employ. Shock	GDP pc			Unemp. rate		Wage	
Instrument				Construction	HH debt	AD emp.	HH debt	AD emp.	HH debt	AD emp.
F-stat				21.372	8.718	19.065	34.333	48.799	11.718	15.361
msa fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
time fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes

Notes: The dependent variable is the in-migration rate of natives. The independent variable is either (log) GDP per capita, unemployment rates or average wages in US metropolitan areas between 2005 and 2010. (log) GDP per capita, unemployment rates or wages are instrument with the variable the construction importance before the crisis, and the debt to income ratio introduced in Mian and Sufi (2013), see more details in table 3 or in the text. Regressions are weighted by the number of observations used to compute the shares of construction. F-stats reported are the F-stats of excluded instruments in the first stage regression. Robust standard errors clustered at the metropolitan area level are reported. Number of observations: 263 metropolitan areas x 6 years = 1578 or 194 metropolitan areas x 6 years = 1164 when Mian and Sufi (2013) measure used. * p<.1, ** p<.05 and *** p<.001.

3.3.3 Migration and skills

Wozniak (2010) emphasizes that high-skilled workers are 5-15 percent more likely to take advantage of good labor market opportunities.¹⁶ Her analysis, however, does not explain how sensitive the decision is of whether to move to particular places when they have been hit by a negative shock.

An ideal experiment to answer whether in-migration rates respond differently to changes in local labor market conditions would be to have a shock that only affects one type of workers. This is what I do in Monras (2013) with low-skilled workers. There I consider two shocks, one in 1995 when more Mexicans than expected entered high immigration states, and the other in 2005 when (mainly low-skilled) people from Louisiana and Mississippi were displaced by hurricane Katrina. In both cases I observe that fewer low-skilled workers moved into states that either received more Mexicans or more people displaced from Katrina. In this paper, the shock affected both high and low-skilled workers.

¹⁶Literature reviews about internal migration rates include Greenwood (1997).

Table 7: The migration response to the crisis: In-migration rates, low-skilled population

in-migration rate, low-skilled population										
Dep. Var. :	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	-0.275***	-0.002*	-0.011***	0.171***	0.099***	0.110***	-0.146***	-0.191***	0.304*	0.408**
	0.046	0.001	0.004	0.029	0.034	0.024	0.054	0.053	0.171	0.171
Exp. Var	Constr. Shock	HH debt Shock	AD employ. Shock	GDP pc			Unemp. rate		Wage	
Instrument F-stat				Constr.	HH debt	AD emp.	HH debt	AD emp.	HH debt	AD emp.
				21.372	8.718	19.065	24.839	36.008	6.236	9.126
msa fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
time fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
in-migration rate, native low-skilled population										
Dep. Var. :	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	-0.254***	-0.002*	-0.009***	0.158***	0.087**	0.097***	-0.128**	-0.169***	0.267	0.361**
	0.047	0.001	0.003	0.036	0.035	0.025	0.056	0.054	0.164	0.165
Exp. Var	Constr. Shock	HH debt Shock	AD employ. Shock	GDP pc			Unemp. rate		Wage	
Instrument F-stat				Constr.	HH debt	AD emp.	HH debt	AD emp.	HH debt	AD emp.
				21.372	8.718	19.065	24.839	36.008	6.236	9.126
msa fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
time fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes

Notes: The dependent variable is the in-migration rate of low-skilled people. The independent variable is either (log) GDP per capita, unemployment rates or average wages in US metropolitan areas between 2005 and 2010. (log) GDP per capita, unemployment rates or wages are instrument with the variable the construction importance before the crisis, and the debt to income ratio introduced in Mian and Sufi (2013), see more details in table 3 or in the text. Regressions are weighted by the number of observations used to compute the shares of construction. F-stats reported are the F-stats of excluded instruments in the first stage regression. Robust standard errors clustered at the metropolitan area level are reported. Number of observations: 263 metropolitan areas x 6 years = 1578 or 194 metropolitan areas x 6 years = 1164 when Mian and Sufi (2013) measure used. * p<.1, ** p<.05 and *** p<.001.

Table 8: The migration response to the crisis: In-migration rates, high-skilled population

Dep. Var. :				in-migration rate, high-skilled population						
	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	-0.260***	-0.001	-0.010**	0.161***	0.073**	0.105***	-0.178*	-0.328**	0.146*	0.244**
	0.046	0.001	0.005	0.030	0.035	0.035	0.107	0.164	0.086	0.117
Exp. Var	Constr. Shock	HH debt Shock	AD employ. Shock	GDP pc			Unemp. rate		Wage	
Instrument				Constr.	HH debt	AD emp.	HH debt	AD emp.	HH debt	AD emp.
F-stat				21.372	8.718	19.065	52.568	46.145	20.643	25.951
msa fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
time fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Dep. Var. :				in-migration rate, native high-skilled population						
	OLS	OLS	OLS	IV	IV	IV	IV	IV	IV	IV
	-0.234***	-0.001	-0.008*	0.145***	0.047	0.081**	-0.114	-0.252*	0.094	0.187*
	0.043	0.001	0.004	0.030	0.035	0.033	0.101	0.145	0.081	0.104
Exp. Var	Constr. Shock	HH debt Shock	AD employ. Shock	GDP pc			Unemp. rate		Wage	
Instrument				Constr.	HH debt	AD emp.	HH debt	AD emp.	HH debt	AD emp.
F-stat				21.372	8.718	19.065	52.568	46.145	20.643	25.951
msa fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
time fe	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes

Notes: The dependent variable is the in-migration rate of high-skilled people. The independent variable is either (log) GDP per capita, unemployment rates or average wages in US metropolitan areas between 2005 and 2010. (log) GDP per capita, unemployment rates or wages are instrument with the variable the construction importance before the crisis, and the debt to income ratio introduced in Mian and Sufi (2013), see more details in table 3 or in the text. Regressions are weighted by the number of observations used to compute the shares of construction. F-stats reported are the F-stats of excluded instruments in the first stage regression. Robust standard errors clustered at the metropolitan area level are reported. Number of observations: 263 metropolitan areas x 6 years = 1578 or 194 metropolitan areas x 6 years = 1164 when Mian and Sufi (2013) measure used. * p<.1, ** p<.05 and *** p<.001.

However, as argued before, there is variation in how much the crisis affected wages and unemployment rates for workers of different skill levels. In particular, it seems that the crisis affected low-skilled workers especially on employment opportunities. This then translates into higher sensitivities of in-migration rates to unemployment rates for low-skilled than for high-skilled workers. Conversely, the crisis seemed to affect the wages of high-skilled workers relatively more. In-migration rates are, in this case, relatively more sensitive to wages than to unemployment rates. This is shown in Tables 7 and 8.

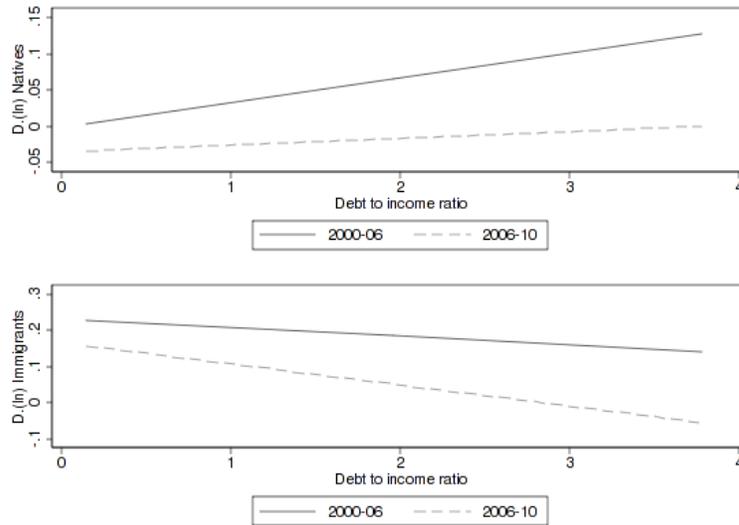
The final point worth emphasizing from these two tables is that this seems to be true both for native and immigrant low skilled workers, although immigrant workers do seem more responsive, as also emphasized in Cadena and Kovak (2013). Contrary to Cadena and Kovak (2013), however, I do find that low skilled natives respond to local conditions. The conflicting results are explained by the pre crisis trends and how these trends reacted to the crisis. Locations that were hit harder during the crisis were increasing relatively more before 2008. Mexicans were more responsive both to positive and negative employment opportunities, thus the pre-crisis trends were more positive and the post-crisis were more negative.

This explains why both the results reported in Cadena and Kovak (2013) and here are possible. It can be the case that when regressing the percentage change in native population between 2006 and 2010 on a measure of how hard the crisis hit across locations we obtain a positive coefficient, while if we do the same exercise with the entire population we find a 0 and a negative for immigrants. This would suggest that overall the population is not responsive to negative shocks – as concluded in Mian and Sufi (2013) –, natives are attracted to hard hit locations, while immigrants, and in particular Mexicans, do respond to negative shocks – as argued in Cadena and Kovak (2013). However, if we do the same regression but using the population change between 2000 and 2006 we realize that those were locations that were growing disproportionately. The change in trend between 2000-2006 and 2006-2010 is evident both for the Immigrants and for natives. This can be seen in Figure 9:

In particular, the first two graphs in Figure 9 show that if we related the growth rate of native population and the debt to income ratio computed in Mian and Sufi (2013) we observe that between 2000 and 2006 there is a strong positive relationship. This relationship becomes less strong between 2006 and 2010, precisely when the crisis hits in this high debt metropolitan areas. If we look at immigrants only, we observe that there was initially a slightly negative relationship, that became even more negative between 2006 and 2010. This change in trend is very similar within natives and immigrants. Understanding these different trends are crucial to interpret whether only low skilled immigrants respond to local shocks or whether also natives do, despite the fact that the relationship between native population growth rates and debt to income ratio is not negative between 2006 and 2010.¹⁷

¹⁷I obtain similar results for the alternative measures used in this paper of how hard the crisis hit across locations.

Figure 9: Differential trends between low-skilled natives and immigrants



Notes: This graph shows the different trends in native and immigrant low-skilled population relative to a measure of how hard the crisis hit at a local level.

4 Conclusion

Contrary to previous literature, such as Mian and Sufi (2013), Molloy et al. (2011) and Yagan (2014), I show in this paper that internal migration rates responded to the crisis. Rather than observing populations leaving the hard hit locations, I have documented that fewer people migrated into the locations that suffered more from the crisis. This is important because it helps decrease the labor supply in those metropolitan areas and it spreads the local shocks spatially.

Furthermore, I show that this is found both when considering native workers alone, or together with immigrants. When distinguishing by skill, low-skilled workers are shown to be more responsive to unemployment rates while high-skilled workers respond more to wage changes. In all, this paper shows that internal mobility may have helped mitigate particularly strong negative shocks in local labor markets.

References

- Acemoglu, D. and D. Autor**, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” *Handbook of Labor Economics Volume 4*, Orley Ashenfelter and David E. Card (eds.), 2011.
- Artuç, E., S. Chaudury, and J. McLaren**, “Trade Shocks and Labor Adjustment: A Structural Empirical Approach,” *American Economic Review*, 2010.
- Autor, D. and D. Dorn**, “This Job Is ‘Getting Old:’ Measuring Changes in Job Opportunities Using Occupational Age Structure.,” *American Economic Review Papers and Proceedings*, 2009.
- **and L. Katz**, “Changes in the Wage Structure and Earnings inequality,” *Handbook of Labor Economics*, 1999.
- **, D. Dorn, and D. Hanson**, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, Forthcoming.
- Blanchard, O. and L. Katz**, “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1992, pp. 1–75.
- Cadena, B. and B. Kovak**, “Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession,” 2013.
- Card, D.**, “Immigration and Inequality,” *American Economic Review Papers and Proceedings*, 2009, *99*(2), 1–21.
- Diamond, R.**, “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000,” *mimeo*, 2013.
- Glaeser, E.**, “Cities, Agglomeration and Spatial Equilibrium,” *Oxford University Press*, 2008.
- **and J. Gyourko**, “Urban Decline and Durable Housing,” *Journal of Political Economy*, 2005, *113*(2), 345–375.
- Greenwood, M.**, “Internal Migration in Developed Countries,” *Handbook of Population and Family Economics* ed. Mark R. Rosenzweig and Oded Stark. New York: Elsevier Science., 1997.
- Handbury, J.**, “Income and Quality: Lessons from Micro Data for Estimating Aggregate Welfare,” *mimeo Columbia*, 2011.
- Hornbeck, R.**, “The Enduring Impact of the American Dust Bowl: Short- and Long-run Adjustments to Environmental Catastrophe,” *American Economic Review*, 2012, *102*(4), 1477–1507.

- **and S. Naidu**, “When the Levee Breaks: Black Migration and Economic Development in the American South,” *NBER WP n. 18296*, 2012.
- Kaplan, G. and S. Schulhofer-Wohl**, “Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey,” *Demography*, 2012.
- Katz, L and K. Murphy**, “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 1992, *107(1)*, 35–78.
- Kennan, J. and J. Walker**, “The Effect of Expected Income on Individual Migration Decisions,” *Econometrica*, 2011, *79(1)*, 211–251.
- Mian, A. and A. Sufi**, “What Explains High Unemployment? The Aggregate Demand Channel,” 2013.
- Molloy, R. and A. Wozniak**, “Labor Reallocation over the Business Cycle: New Evidence from Internal Migration,” *Journal of Labor Economics*, 2011, *29(4)*, 697–739.
- , **C. Smith, and A. Wozniak**, “Internal Migration in the United States,” *Journal of Economic Perspectives*, 2011, *25(3)*, 173–196.
- Monras, J.**, “Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis,” *Columbia University Discussion Paper Series*, 2013.
- Moretti, E.**, “Local Labor Markets,” *Handbook of Labor Economics*, 2011.
- , “What Workers Lose by Staying Put,” *Wall Street Journal*, 2012.
- Notowidigdo, M.**, “The Incidence of Local Labor Demand Shocks,” 2013.
- Ruggles, S., M. Sobek, T. Alexander, C.A. Fitch, R. Goeken, PK Hall, M. King, and C. Ronnander**, “Integrated Public Use Microdata Series: Version 4.0 [Machine-readable database].,” *Minneapolis, MN: Minnesota Population Center [producer and distributor]*, 2008.
- Topel, R.**, “Local Labor Markets,” *Journal of Political Economy*, 1986, *94(3)*, S111–S143.
- Wozniak, A.**, “Are College Graduates More Responsive to Distant Labor Market Opportunities?,” *Journal of Human Resources*, 2010, *45(4)*, 944–970.
- Yagan, D.**, “Moving to Opportunity? Migratory Insurance over the Great Recession,” *Job Market Paper*, 2014.