Lingering Effects of Recessions: Age Differentiated Migration

Nathalie Gonzalez Prieto*, John Coglianese†

September 30, 2025

Abstract

Local economic shocks can change individuals' migration decisions, with this effect potentially varying across the lifecycle. Using data from the New York Fed Equifax Consumer Credit Panel, where we can follow people over time and have information on location and age, joined with sectoral GDP at the county level from the Bureau of Economic Analysis (BEA), we examine the response of migration rates by age group to local economic shocks. Our results indicate that prime-age adults decrease their in-migration rate to areas experiencing a recession. In contrast, the out-migration rate of retirement-age individuals increases after a local labor market experiences a positive shock. This result can explain hysteresis in raw employment and participation rates, not adjusted for demographics, due to the persistent effects of shocks on local demographic composition. Methodologically, we show that our results are robust to including bilateral controls for economic conditions, a specification emphasized in recent work to address spatial confounders.

1 Introduction

Academic literature tends to find little evidence of macro hysteresis in US unemployment rates (Blanchard and Katz, 1992; Dao et al., 2017). That is, local unemployment rates return

*World Bank. I want to thank Katharine Abraham, John Haltiwanger, Judith Hellerstein, Guido Kursteiner, Macarena Kutscher, Edward Olivares, Eugene Ou, John Shea, and Sergio Urzua for their time and suggestions. All remaining errors are our own.

†Federal Reserve Board of Governors. The views expressed in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff.

to their pre-recession levels within a couple of years following the recession's end. However, the evidence is less clear for participation and employment rates in local labor markets. Local areas experiencing a recession see persistent declines in employment and labor force participation. However, recent work emphasizes that this decline stems from changes in local demographic composition, not from a decline in individual labor force attachment. Using state-level variation, Cajner et al. (2020) find that the demographically-adjusted labor force participation rate takes eight years to recover following a recession but eventually goes back to its pre-recession level. As Figure 1 shows, this result does not hold for the raw labor force participation rate, unadjusted for demographic shifts. Using a different methodology, Gonzalez (2020) finds a similar result for unemployment, labor force participation, and employment rates after the Great Recession (GR) using state-level variation, as shown in Figure 2. Cajner et al. (2020) show that states affected by a negative output shock experience a decrease in the population share of prime-age adults even 10 years after the shock. Since prime-age adults have the highest labor force participation rate among all age groups, this means that areas affected by a negative shock are left with a higher share of people with lower labor market attachment. This work builds on the literature on internal migration in the US -see Jia et al. (2022) for a comprehensive review. Our understanding of internal migration has increased a great deal over the last decade, but many open questions remain.

This paper contributes to a growing literature on the persistence of local economic shocks by uncovering an age-specific migration mechanism that reshapes local demographics after recessions. Our results are closely related to Yagan (2019), who documents long-lasting employment hysteresis following the GR, and to Diamond (2016), who highlights heterogeneous location choices across demographic groups. In contrast to these studies, we focus explicitly on the age dimension of migration responses, showing how the relative immobility of retirement-age individuals versus the sensitivity of prime-age adults to labor market conditions can generate persistent demographic shifts. This mechanism has direct implications for fiscal sustainability and local business dynamism, as older residents contribute differently to labor supply, tax bases, and housing markets ((Bloom et al., 2010))

The main goal of this paper is to understand the process through which local labor

Figure 1: Labor Force Participation Recovery: Raw Series and Demographic Adjustments

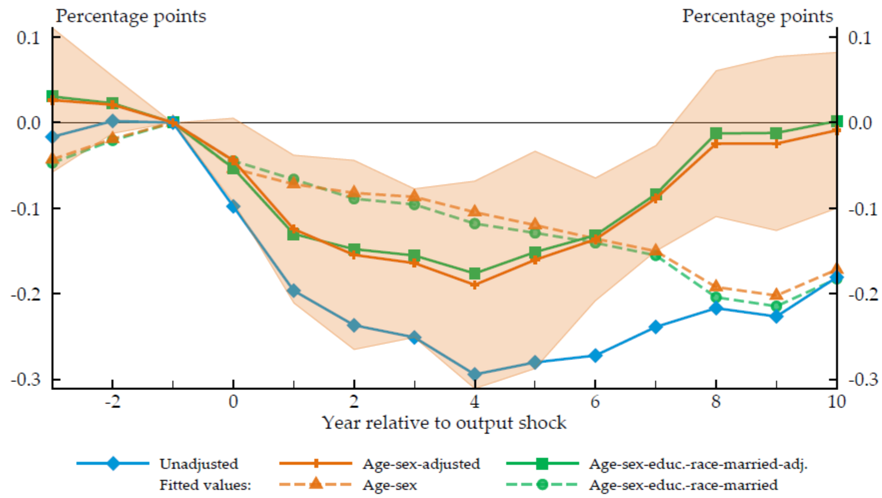

Note: Each line shows the estimated coefficients from Equation 1 using the specified adjusted, unadjusted, and fitted-value LFPR as the outcome. The band around the orange solid line shows a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1% shock to GSP growth in year 0. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population.

Source: Cajner et al. (2020)

Figure 2: Recovery Path of Labor Market Indicators after the Great Recession

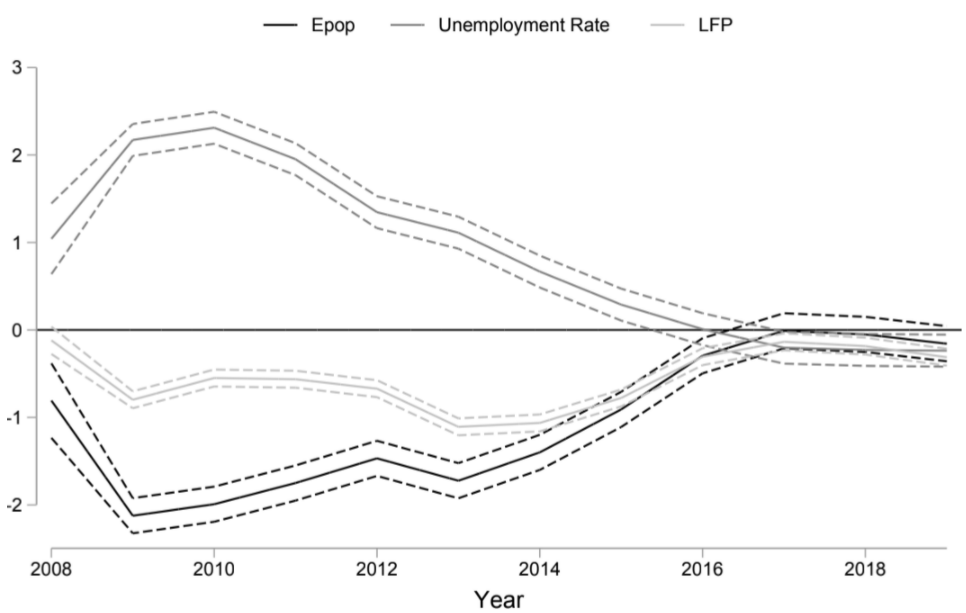

Note: This graph shows the difference in labor market indicators between states severely affected by the Great Recession (GR) and those mildly affected by the GR. States are classified in their corresponding group using the 2009 residuals from Blanchard and Katz (1992) VAR system. These indicators are adjusted for changes in states' demographic composition. Source: Gonzalez (2020) markets are left with a persistently smaller share of prime-age individuals following a negative shock. Using Commuting Zones (CZ) to define labor markets, we study migration patterns by age in response to different growth rates of economic activity induced by changes in national sectoral gross domestic product (GDP) growth as a possible mechanism behind hysteresis. [1]

Reverse causality is the main challenge in identifying the migration response to an increase in GDP. Individuals' migration decisions are potentially affected by the state of the local economy, which in turn is potentially affected by migration decisions via changes in labor supply. In other words, if individuals are basing their migration decisions on the health of the labor market, a lower level of GDP growth could deter them from moving to a specific area, i.e. a reduction in *in-migration*, or lead them to leave, an increase in *out-migration*. If enough people make this decision, the labor supply of the destination could be reduced, which could negatively impact GDP growth.

To overcome this endogeneity challenge, we use an instrumental variable approach using a shift-share instrument based on sectoral GDP. The idea behind this instrument is to identify changes in local GDP growth stemming from national changes in industry-level demand. To the extent that the industrial composition of a CZ reflects structural characteristics of the local labor market, this strategy allows us to identify demand shocks.

Our findings suggest that the net migration—net migration = in-migration − out-migration—rate of prime-age adults increases when a CZ experiences positive GDP growth, while retirement-age individuals decrease their net migration in response to a negative GDP shock. When looking into the source of the net migration response, we find that the response of prime-age individuals comes mainly from a reduction in **in-migration** to areas that experienced a negative shock, whereas the retirement age group experiences a reduction in its **out-migration** rate from areas that experience negative GDP growth.

Importantly, our estimates remain stable after controlling for economic conditions bilaterally. Recent work by Borusyak et al. (2022) highlights that when a local labor market experiences a negative shock, it may be the case that surrounding regions are also affected

[1] As the measure of local economic activity, we use CZ GDP. We built this measure using county-level data from the Bureau of Economic Analysis and the crosswalk of counties to 2000 delineation of CZ provided by the Economic Research Service of the U.S. Department of Agriculture.

by the same negative shock, biasing the response of migration towards zero. We conduct a robustness check in which we estimate the response of bilateral in- and out-migration flows for a shocked CZ, controlling for economic conditions in other CZs, and find nearly identical results as in our baseline regressions. This underscores that the age-differentiated migration elasticities we document are not driven by bilateral correlations but rather by local economic shocks.

Related Literature

This paper is related to literature that studies migration and labor market adjustments summarized by Jia et al. (2022). Blanchard and Katz (1992) is a seminal paper that uses a VAR model on state-level relative employment growth, employment rates, and participation rates to characterize the labor market after a demand shock. They find that the unemployment rate recovered quickly for the period 1972–1990 and conclude that migration was the main equilibrating force after a local recession. Dao et al. (2017) in a follow-up paper, using an alternative identification strategy and additional data sources, including a direct measure of net migration, find that most of those who lose their jobs after a local recession end up in the unemployment pool for the first couple of years. Migration kicks in later and acts as an equilibrating force in the long run, so their results are consistent with the core results in Blanchard and Katz (1992). However, migration has lost power as an equilibrating force in recent decades. This result is in line with other studies that have shown that in recent decades there has been a decrease in the convergence of labor markets across the US (Amior and Manning, 2018; Austin et al., 2018; Ganong and Shoag, 2017) and that there is a negative secular trend in migration, particularly related to work-related moves (Molloy et al., 2011, 2017). Another recent contribution to this literature is Hershbein and Stuart (2021). They use the reduction in the employment rate in the first two years after a national recession to quantify its effect at the MSA level; for example, their measure of the impact of the GR is the reduction in the employment rate in 2007-2009. They find that more affected MSAs experienced permanent declines in their employment rate.

Cadena and Kovak (2016) is one of the few papers that looks at heterogeneity in the migration response to local recessions based on demographic characteristics. They find that, among low-skill workers, Mexican-born immigrants are more responsive to employment changes induced by the GR than native workers. This finding underscores the importance of paying close attention to the heterogeneous migration response of individuals in different demographic groups. We contribute to this literature by analyzing the migration response for different age groups and characterizing the aggregate impact of individuals' decisions on local labor markets.

Cajner et al. (2020), using state-level data, find that there is a long-lived negative effect on local labor force participation rates after a recession. Still, on average, the demographically-adjusted labor force participation rate returns to its pre-recession level after eight years. The demographic adjustment is key to their findings. The authors find that the share of prime-age workers in areas more adversely affected by a recession decreases permanently and assert that this is likely the result of an increase in the out-migration of prime-age workers.

On the structural side, Monras (2020) finds that in-migration is the relevant margin affecting population growth after a local economic shock and develops a general equilibrium model with multiple locations around that stylized fact.

Recent work has also emphasized the importance of heterogeneous migration elasticities and their consequences for local adjustment.Yagan (2019) finds that the Great Recession had permanent effects on employment, providing direct evidence of hysteresis at the local level. Diamond (2016) shows that migration responses vary sharply across education and income groups, underscoring the role of heterogeneity in shaping spatial equilibrium. Notowidigdo (2020) reviews the allocation of workers across regions and stresses the importance of declining migration in amplifying regional disparities.

Our contribution to the literature is threefold. First, using the Consumer Credit Panel, a better-suited data set to understand the migratory response at a granular level, we look directly at the response in the gross migration flows by age group and find that local recessions induce a reduction in the in-migration rate of prime-age adults and a reduction in the out-migration of retirement-age individuals. This result explains how labor markets affected by a recession are left with a lower participation rate via a higher share of retirement-age individuals.

Second, we use Commuting Zone (CZ) level variation. CZs are geographical units closer to a labor market than states. Using CZs as our unit of analysis allows us to include and study within states across CZ migration. This is a relevant difference for two reasons. The gravitational relation between local labor markets in regards to migration documented in the US since Zipf (1946) implies that the distance is a strong predictor of moving flows between areas and labor markets within the same state are likely to be closer to each other than labor markets in different states. Using data from the American Community Survey (ACS), we find that 40% of cross-CZ migration occurs within states. Additionally, a recent paper by Wilson (2022) finds that there is a discontinuity in the migration probability around state borders. This means that by missing the intra-state migration, state-level analyses ignore a relevant proportion of migration decisions.

The third contribution of our paper is methodological. We estimate migration responses both with and without bilateral controls for economic conditions in other CZs. Recent work has emphasized that omitting bilateral controls can lead to biased estimates, for example because persistent dyad-specific characteristics——such as distance, long-run migration networks, or relative amenities——may be correlated with local shocks Olney and Thompson (2024); Chen et al. (2024); Blumenstock et al. (2023); Borusyak et al. (2022). In our context, however, we find that the results are remarkably similar across the two specifications. This suggests that, unlike in settings where migration decisions are heavily shaped by pair-specific frictions, most of the variation we exploit comes from temporal changes in local economic conditions rather than from stable bilateral features of CZ pairs. In other words, while bilateral controls provide an important robustness check, the consistency of our estimates across specifications indicates that our findings are not driven by unobserved spatial confounders. This robustness enhances confidence that the age-differentiated migration elasticities we document reflect genuine responses to local demand shocks.

In the next section, we describe the empirical strategy, including some of its weaknesses and possible avenues for improvement. In Section 4, we describe the data, including how we built the shift-share instrument at the CZ level. Section 5 includes the results using weighted averages of national 2-digit-level GDP growth rates to instrument the CZ GDP growth rate. Finally, in Section 6, we conclude.

2 Labor Force Participation and Migration Over the Life Cycle

In this section, we document average migration and labor participation rates by age using data from the American Community Survey (ACS) between 2005-2019. The labor force participation rate is the share of the population that is either working or looking for a job. The migration rate is the number of people that migrate across CZs as a percentage of the population of that age group in the destination CZ. It is a well-known fact that the likelihood of migration decreases with age. Early papers like Sjaastad (1962) and Schultz (1978) talk about a consensus in the literature on this issue. More time to take advantage of differential economic conditions, increasing opportunity, direct and psychological cost, and cohort selectivity effects have been mentioned as potential explanations for this established fact. More recently, Kaplan and Schulhofer-Wohl (2017) documented that the likelihood of migration continues to decrease with age, with younger workers being more likely to migrate than older workers.

If a relevant reason for moving is the state of the economy, the fact that younger workers migrate more implies that young people are more sensitive to the local labor market conditions than older people. This difference alone could explain why areas affected by a negative economic shock are left with a larger proportion of individuals with a lower labor market attachment.

The left panel of Figure 3 shows that labor market attachment remains high and stable for the prime age population (25–54). After age 55, labor force participation declines steeply. After age 65, the decrease in labor force participation decelerates, but participation continues declining until age 80, the last age reported in the graph. From the right panel, it is evident that migration decreases with age. Additionally, we can see that the slope of the migration profile is not constant. The migration profile has a convex shape. It decreases rapidly between the age of 25 and the mid-40s, and then it levels out above 2 percent up to 65 years of age. After age 65, migration rates decrease and stay constant below 2 percent up to age 80. These patterns indicate that a decrease in the net migration of prime-age individuals would likely be associated with a decrease in the overall labor force participation and employment

Figure 3: Participation and Migration Profiles by Age

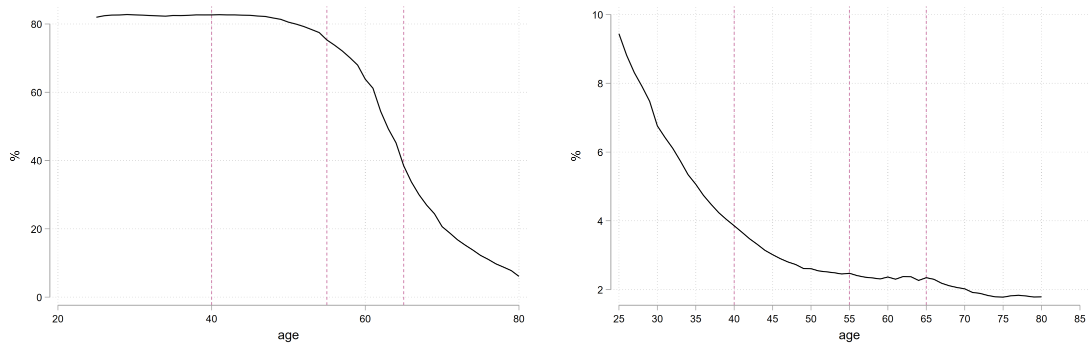

Source: Authors' calculations using ACS microdata from 2005-2019. The left panel plots labor force participation rates by age, and the right panel plots migration rates by age. The migration rate refers to the annual cross-CZ migration rate. A person is classified as a migrant if she moved to a PUMA with no overlapping CZs as the PUMA of origin. CZs are defined using 2000 delineations as described in Appendix A.

rates.

3 Empirical Specification

In this section, we describe the regression specifications we use in our analysis. First, we use a "one-sided" specification in which we examine the response of in- and out-migration for a commuting zone, pooling together all in- and out-migration flows. Next, we introduce an "bilateral" specification in which we separately examine the response of flows between commuting zone pairs. This latter specification allows us to control for economic conditions in other commuting zones, addressing the concerns with one-sided specifications laid out by Borusyak et al. (2022).

3.1 One-sided regressions

Identifying the migration response after a local economic shock is not trivial. The main identification threat is reverse causality. This challenge stems from the fact that people may decide to move to a location based on the economy's health, and at the same time, the economy's health might be affected by migration decisions.

We compare the sensitivity of the migration rates to changes in destination GDP growth across three age groups $A = \{25-39,\ 40-64,\ 65+\}$. The coefficients of interest are β_k^A, the change in the migration rate induced by a change in GDP. More specifically, we allow for the effect to take time by specifying the structural equations of the form

$$MR^A_{c,t+k} - MR^A_{c,t-1} = \beta^A_k \Delta g_{c,t} + \mu^A_c + \mu^A_t + \xi^A_{c,t} \tag{1}$$

where MR is either the net-migration, the in-migration or the out-migration rate, $A = \{25-39,\ 40-64, 65+\}$, $\Delta g_{c,t}$ is GDP of CZ c in time t, and μ_c and μ_t are CZ and time fixed effects respectively. Our goal is to characterize the migration response across age groups to understand the process through which areas affected by a negative labor demand shock are left with a higher share of individuals with lower labor market attachment. Based on this specification, we build an impulse response function with the coefficients of β^A_k over different values of k. β_k is interpreted as the effect of a change in the GDP growth rate over the k periods after the demand shock.

The bias of estimating equation 1 by ordinary least squares (OLS) depends on the right-hand side variable. The sign of the bias for the net and gross migration is the same. If GDP growth and in-migration are positively related in both directions, meaning that as GDP growth increases, the in-migration rate increases, and as the in-migration rate increases, the GDP growth increases, so the sign on the bias of the OLS estimator is positive. The structural relationship between out-migration and GDP growth is likely to be in the opposite direction, meaning that as GDP grows, out-migration is likely to decline. As out-migration increases, GDP is likely to shrink. In this case, the bias of the OLS estimator is negative. [2]

To overcome this challenge, we use the change in local GDP growth induced by changes in local demand. We do this by instrumenting the actual GDP growth rate with the GDP growth rate predicted by the industrial structure of the CZ. More precisely, we define the predicted CZ-level GDP growth as the sum of the shares of local GDP in each sector in 2004 multiplied by the national growth rate of that sector in each subsequent year:

$$\Delta g^*_{c,t} = \sum_s \omega^{2004}_{c,s} \Delta g_{s,t} \tag{2}$$

[2] For details on the sign of the bias in a system with reverse causality, see Basu (2015)

where $\omega_{c,s}^{2004}$ is the share of sector s in nominal GDP in CZ c in 2004, and Δg_t is the national real GDP growth of sector s in period t. The intuition behind this instrument is that it captures the effect on local growth of changes in the national demand for the goods that are produced intensively in the CZ. A decrease in the local GDP growth rate coming from a decrease in demand signals a less favorable environment for workers. This is a shift-share instrument, also known as a Bartik instrument. The literature has traditionally used this type of instrument using employment, not GDP. We use GDP for two reasons: 1) the dynamic behavior of GDP displays less persistence than that of the employment rate. We want to characterize the dynamics of the migration response, so this feature makes GDP growth a better alternative, and 2) the relationship between employment and migration is mechanical, making the reverse causality concern more salient. The first-stage regression is thus:

$$\Delta g_{C,t} = \alpha \Delta g^*_{C,t} + \lambda_t + \lambda_C + \nu_{C,t} \tag{3}$$

where λ_t and λ_C are fixed effects for time and CZs, respectively. For the instrument to be valid, it must be relevant and exogenous. The argument for the instrument's relevance is that a decrease in the national demand for the goods produced intensively by the CZ will translate into a decrease in local demand for labor and, therefore, a decrease in GDP growth. As shown in Table 1, GDP growth is strongly related to predicted GDP growth.

Table 1: First Stage

	(1)
VARIABLES	$\Delta g_{c,t}$
$\Delta g^*_{c,t}$	1.235***
	(0.0274)
Observations	6,900
R-squared	0.410

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note: Includes year-fixed effects and CZ-fixed effects. Annual data from 2005-2019.

As pointed out by Goldsmith-Pinkham et al. (2020), by using a shift-share instrument, we are also implicitly either assuming the exogeneity of the shares or the exogeneity of the shocks. In this case, we rely on the assumption of the shares being exogenous. Concretely the identifying assumption is that differential exposure across CZ to common national-level shocks leads to differential changes in the migration rates. In other words, we take the shares as reflecting long-term structural characteristics of the CZ that make them more or less susceptible to changes in the national level industry labor demand. To make this assumption more sensible, we use the industry shares of a period before the data we use in our estimates.

3.2 Bilateral Specification

One of the advantages of having access to a large panel data set is that we can measure migration flows between CZ pairs. When someone is deciding if they will move and where they will move to, they consider the state of their current location and the state of potential destinations. In other words, the migration decision is likely based on the relative state of the origin and each potential destination. By looking at the one-sided regressions described in the previous section, we are assuming that, in the case of out-migration (in-migration), the state of the potential destinations (origin) is constant. Since people always have the option to stay without paying a moving cost, not considering the two-sided dynamic of the migration decision is a relevant omission. Using data from Brazil, Borusyak et al. (2022) found that results based on a one-sided standard regression provide a misleading picture of the determinants of migration elasticities. We move to a specification and data structure that allows us to characterize the in-migration response of each age group to local economic cycles, controlling for economic conditions in other commuting zones. We implement this idea by including a time-varying control for other CZs, for example in the regression for in-migration we add the term $\theta^A_{o,t}$:

$$IMR^A_{o,d,t+k} - IMR^A_{o,d,t-1} = \beta^A_k \Delta g_{d,t} + \theta^A_{o,t} + \gamma^A_{o,d} + \epsilon^A_{o,d,t} \tag{4}$$

where $IMR^A_{o,d,t}$ is the in-migration rate from origin o (i.e. the number of migrants from o to d divided by population in d) and $\gamma^A_{o,d}$ is an origin-destination fixed effect that captures persistent flows between the CZs. For the time-varying control $\theta^A_{o,t}$, we use either the predicted

GDP growth rate for the origin CZ, $\Delta g^*_{o,t}$, or an origin-CZ-by-time fixed effect, $\gamma^A_{o,t}$. Symmetrically, when we estimate the effects of local economic cycles on out-migration bilaterally, we include time-varying destination controls.

$$OMR^A_{o,d,t+k} - OMR^A_{o,d,t-1} = \beta^A_k \Delta g_{o,t} + \theta^A_{d,t} + \tau^A_{o,d} + \upsilon^A_{o,d,t} \tag{5}$$

This specification allows us to overcome the shortcoming highlighted by Borusyak et al. (2022) by directly controlling for the dynamics of the origin/destination. Since the endogeneity concern operates in the same way in the origin-destination specification, we also instrument the local growth rate with the predicted growth rate using the industrial structure of the CZ in 2004 and national growth rates at the industry level.

4 Data description

We use data from the New York Fed Equifax Consumer Credit Panel (CCP) to measure migration rates by age groups at the CZ level, county-level population estimates from the Census Bureau, and county-by-sector GDP data from the Bureau of Economic Analysis (BEA) to build the GDP growth and the instrument. For the regression estimates, we used data from 2005 to 2019, and for the sectoral shares of GDP, we used data from 2004. Below is an overview of each data source and the variables we use.

4.1 CCP

The CCP is built using individuals' credit reports. It includes anonymized information on age, census block of residence, and financial indicators. The CCP is a nationally representative 5% sample of adults in the US with Equifax credit report data (core sample) and those living in their address. The longitudinal dataset has information starting in 1999 and information up to the second quarter of 2022. As noted by DeWaard et al. (2019), the CCP allows detailed cross-sectional analysis of migration within the US that no public dataset provides, while being consistent with the ACS, the most common data source for migration data in the US. For this project, there are two main advantages of working with the CCP rather than the ACS. First, the CCP contains detailed information on the origin and destination of migrants at the census block level. In contrast, the ACS only reports the geographic data at the PUMA (Public Use Microdata Areas) or MIGPUMA (Migration Public Use Microdata Areas) level, sometimes pooling more than one county together. This means that we can directly measure cross-CZ migration flows. Second, the CCP's sample size is more than three times larger than the ACS. This point is very relevant for the empirical analysis of this paper, given that we require a sample large enough to be representative of age groups within small geographic units.

The main limitation of the CCP is that it is a sample of US residents with a social security number and a credit history. According to Brevoort et al. (2016), 11% of American adults are credit invisible or lack a credit report. The probability of being credit invisible is higher for young, elderly, minority, and low-income consumers. We do not consider this to be a severe drawback since our interest is not focused on financially disadvantaged people. However, it is important to consider this when thinking about the results' external validity. We define migrants as those whose recorded addresses are in a different CZ from the year before. This means we are comparing the reported CZ in period t with the reported CZ in period t-1.[3] We define the in-migration rate for each age group in a particular CZ as the total number of people in that age group who report moving into that CZ in the last year, divided by the number of people in that age group living in the destination CZ in that year. Similarly, the out-migration rate is the number of people in a particular age group who reported moving out of the origin CZ in the last year, divided by the number of people in that age group in the origin CZ that year.

4.2 BEA

Starting in December of 2020, the BEA publishes annual sectoral GDP by county. The data starts in 2001 and is provided in current dollars and in chained dollars of 2012. To build the shares needed for the instrument, we used nominal GDP in 2004. We aggregate the

[3]Equifax has a proprietary algorithm to define the place of residence. We do not have access to the algorithm. Still, from the manuals, we know that address identification relies on several pieces of information, including billing addresses and, if the individual has a mortgage, the address of the mortgage.

sector/county level GDP to sector/CZ level by using the county-to-CZ crosswalk provided by the ERS and described in Appendix 1. The BEA suppresses some cells to avoid disclosing sensitive information, with suppression more prevalent in small counties. To deal with this issue, we use aggregated sectors to ensure the shares add up to one. This process implies that the shares reflect a collection of sectors for some counties. We calculate the equivalent national-level growth rate for all the aggregated sectors necessary to build the CZ predicted GDP growth.

4.3 Descriptive Statistics

Tables 2 and 3 contain the summary statistics of the relevant variables. The first thing to note is that migration rates decrease with age. This is a well-documented fact in line with Figure 3.

Table 2: Summary Statistics: One-Sided Regressions

	mean	sd	p25	p50	p75	min	max
IMR 16+	3.4	1.5	2.3	3.1	4.1	0.0	20.5
IMR 25-39	5.7	2.1	4.2	5.4	6.8	0.0	50.0
IMR 40-64	2.6	1.4	1.6	2.3	3.2	0.0	25.9
IMR 65-80	2.2	1.3	1.3	1.8	2.7	0.0	26.7
OMR 16+	3.4	1.3	2.5	3.1	3.9	0.0	263.0
OMR 25-39	5.8	2.0	4.4	5.4	6.6	0.0	262.5
OMR 40-64	2.6	1.1	1.9	2.4	3.0	0.0	314.0
OMR 65-80	2.2	1.1	1.6	2.0	2.6	0.0	251.9
Observations	10605						

Source: New York Fed Equifax Consumer Credit Panel, BEA, authors' calculations.

Additionally, notice that the average migration rates in the bilateral data structure are two orders of magnitude smaller than in the one-sided data. This comes from the fact most of the O-D pairs have no migration flows while most of the CZ have migration flows. In other words, there are always some people coming or leaving a CZ, but it is not the case that we see people moving to all CZ but to a very reduced set of CZ. This is a fact documented in Sprung-Keyser et al. (2022), where the authors find that the radius of economic opportunity for Americans is very narrow.

Table 3: Summary Statistics: Origin-Destination Regressions

	mean	sd	p25	p50	p75	min	max
IMR 25-39	0.010	0.100	0.000	0.000	0.000	0.000	33.333
IMR 40-64	0.005	0.049	0.000	0.000	0.000	0.000	15.385
IMR 65-80	0.003	0.045	0.000	0.000	0.000	0.000	22.727
IMR 16+	0.006	0.046	0.000	0.000	0.000	0.000	16.190
OMR 25-39	0.010	0.104	0.000	0.000	0.000	0.000	33.333
OMR 40-64	0.004	0.044	0.000	0.000	0.000	0.000	9.091
OMR 65-80	0.003	0.061	0.000	0.000	0.000	0.000	100.000
OMR 16+	0.006	0.047	0.000	0.000	0.000	0.000	5.869
Observations	7151850						

Source: New York Fed Equifax Consumer Credit Panel, BEA, authors' calculations.

5 Results

This section includes the results for the one-sided regressions and the bilateral regressions. The main message is that the results are robust to the specification used. We find that after a local labor market experiences an increase in GDP fueled by national demand, there is a migration response that changes in magnitude and, in some cases, in direction for the different age groups we analyze. The differential migration response by age group is consistent with places that experience GDP shocks having persistent effects via the composition of their demographic composition. The results show that younger people have a larger migration response than older people; this is consistent with the fact that younger people generally have a higher migration probability and that retirement-age people respond less to labor market-related events, consistent with them having a lower reliance on labor market income. Young (26–39) and middle-aged (40–64) people have a positive in-migration response to positive output shocks in both specifications. Even though smaller, the out-migration response of retirement-age people works in the opposite direction. In other words, retirement-age individuals increase their out-migration rate after a positive shock.

5.1 One-Sided Specification

Figure 4 summarizes the results of the local projections on the different migration rates by age group over a ten-year horizon. We include the response of the migration rate to an

Figure 4: One-sided regression results

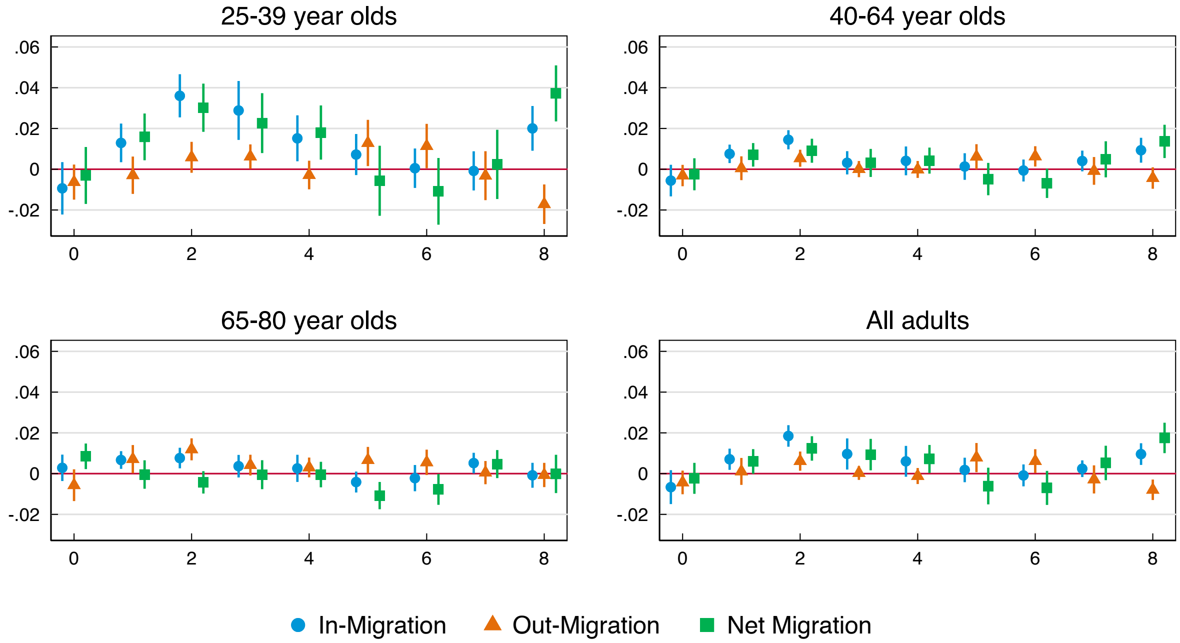

Source: New York Fed Equifax Consumer Credit Panel, BEA, authors' calculations.Each coefficient comes from a different regression of the form: $MR^A_{c,t+k} - MR^A_{c,t-1} = \beta^A_k \Delta g_{c,t} + \mu^A_c + \mu^A_t + \xi^A_{c,t}$ using as an instrument for the GDP growth the predicted GDP growth using a Bartik instrument with industry detail at the two digits level. We use the total population to weigh the results.

increase in GDP—a local economic expansion—induced by an increase in national demand. From top to bottom, the results are ordered by age group, from younger to older, and the final panel of the figure includes the effect for all age groups.

According to the one-sided results, the in-migration rate of all age groups increases with local economic expansion. However, as can be seen by comparing the scale of the Y-axis of the figures, the response is larger for people in the younger age groups. In other words, younger workers are more attracted to areas experiencing an economic expansion, while retirement age individuals respond at a lower scale. On the other hand, we find that out-migration significantly increases after an economic expansion, but only for retirement-age individuals. A muted in-migration and positive out-migration responses indicate a negative net migration response for retirement-age individuals.

For a negative shock, this would symmetrically translate into a decrease in the net migration of young and middle-aged individuals and an increase in the net migration of retirement-age individuals.

5.2 Bilateral Specification

As Figure 5 shows, the results using the bilateral specification are very close to the one-sided regression results.[4] Young and middle age people increase their in-migration to areas that experience an economic expansion, and retirement-age individuals increase their out-migration from these areas. The results are very similar whether we control for time-varying conditions in other CZs either through their predicted GDP growth rate or with CZ-by-time fixed effects.

The overall results after we account for the state of the origin (destination) when analyzing in-migration (out-migration) indicated responses in the migration elasticities that are consistent with the change in the demographic composition of areas that experience an unexpected change in their labor demand. This affirms the findings of our one-sided regressions and indicates that the spatial correlation of shocks is not biasing our results. Most of the migration results in the literature are one-sided and our ability to demonstrate the robustness of one-sided regressions empirically is a contribution of our analysis.

6 Conclusion

Previous literature has shown that local labor markets that experience economic shocks have persistent changes in their demographic composition. Local areas that experience a negative economic shock are left with a higher share of the population above 65. In this paper, we offer novel evidence of the process through which this demographic change takes place. Using detailed data on the universe of US residents with a credit history from the CCP, we show that young adults are more responsive to local economic developments than middle-aged adults and retirement-age adults. Additionally, after a positive GDP shock, we find that the relevant adjustment margin for young adults is in-migration. Young people move more into

[4] We rescale the bilateral coefficients by 710, the number of commuting zones, in order to be comparable to the one-sided regression results.

Figure 5: Bilateral regression results

(a) In-migration

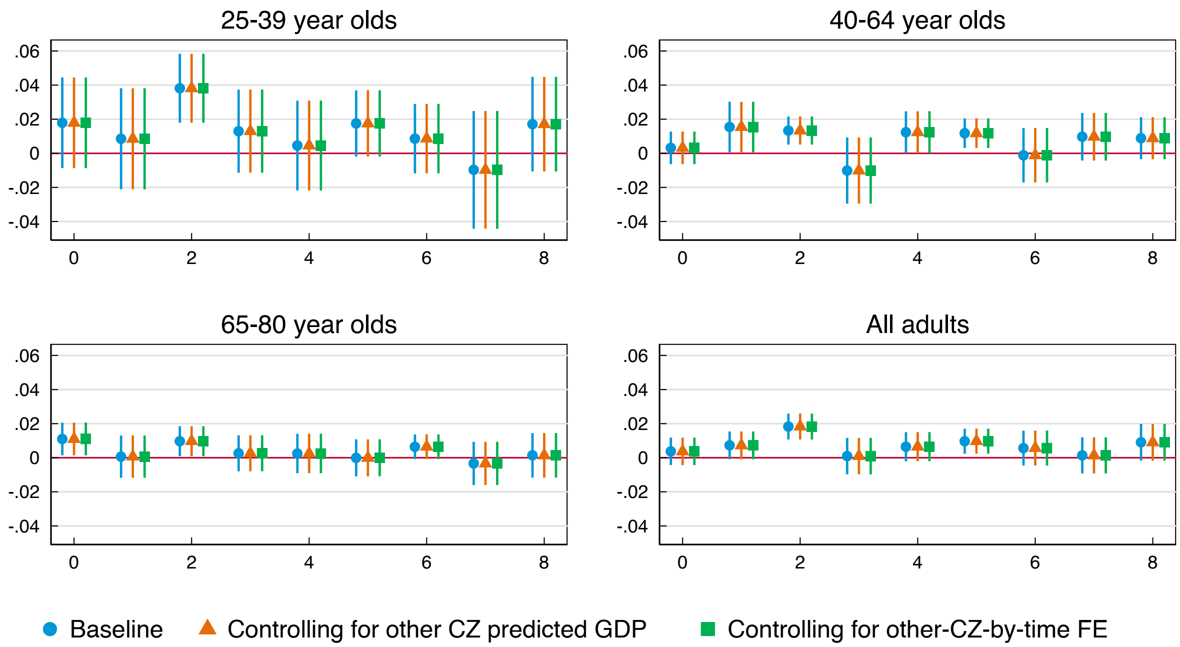

(b) Out-migration

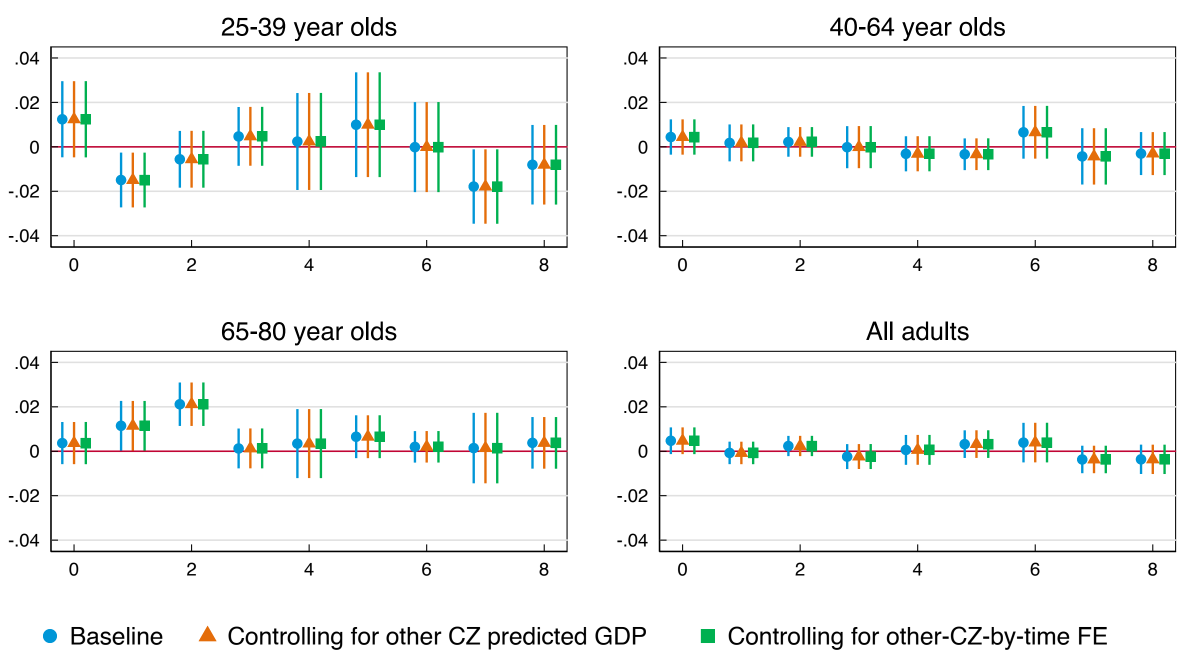

Source: New York Fed Equifax Consumer Credit Panel, BEA, authors' calculations. Bilateral regressions. We instrument destination (origin) GDP with predicted destination (origin) GDP and control for either origin (destination) predicted GDP or origin- (destination)-by-year FE. The predicted GDP growth is a Bartik instrument with industry detail at the two digits level using the industry weights of 2004 and the national growth rates. All specifications include origin-destination-pair FE.

CZs that experience a positive demand shock. In comparison, the most relevant margin of adjustment for retirement-age individuals is out-migration. Retirement age people increase the out-migration from places that experience a positive increase in demand. In addition to implementing one-sided regressions — in- (out)-migration regressions with destination (origin) GDP as the main variable of interest—we implement a novel exercise with a bilateral data structure that allows us to consider the state of the CZ of origin and destination simultaneously. Our results are robust to this identification strategy.

Our results are consistent with younger individuals having a stronger incentive to move where the labor market is performing better. While retirement-age individuals—who are less likely to be active participants in the labor market—place more weight on the price of non-tradable such as housing when making moving decisions.

These results are in line with the broader literature that studies the lack of convergence in labor markets in the US (Amior and Manning (2018); Austin et al. (2018); Ganong and Shoag (2017)) and explains a mechanism behind local hysteresis- the persistent effect of transitory local shocks.

We identify three open questions. First, it would be relevant to identify whether or not the results we find are symmetric to positive and negative shocks. This is particularly relevant because migrating is an expensive process, so it makes sense that the elasticities of migration differ between good and bad times. Second, given that migration in the US occurs mainly between a reduced subset of CZ, as documented by Sprung-Keyser et al. (2022), we would like to quantify the difference in the migration elasticities between highly connected CZs and the rest of the country. Finally, it would be interesting to examine how the elasticity changes across other individual characteristics such as ethnicity, education, or income.

Appendix

A Commuting Zones as Unit of Analysis

Following the recommendations in Molloy et al. (2011) we use CZ as the geographic unit of analysis. CZs are geographic units of analysis intended to reflect more closely the local economy, the area where people live and work. More specifically, CZ are defined using commuting flows from the 2000 census. The strength of the commuting ties between two counties is defined as

$$T_{ij} = \frac{c_{ij} + c_{ji}}{min(w_i, w_j)}$$

where c_{ij} is the number of commuters from i to j and w_i is the total number of workers in i. Then a clustering algorithm is applied for the average linkage that starts by grouping the county pair with largest value of $T_i j$ and subsequently forms clusters of interrelated counties. The final set of CZ is defined such that the average value of T_{ij} is above 0.02. Using this methodology in 2000 the research sector of the department of agriculture defined 709 CZs of which 690 are in the continental US. In contrast to MSAs, another geographic unit defined using commuting patters, CZs do not need to be around a metropolitan area. This means that CZs include rural areas as well and hence cover all the US. Using CZ in contrast to using state-level data to identify migration allows me to capture migration within states across CZ. Within-state migration is a large share of CZ migration, representing, on average, between 2005-2018 40% of CZ migration.

B Commuting Zones in the American Community Survey

This appendix is relevant to the statistics included in the motivation of the paper. The in-migration and out-migration rates used in the regressions come from the CCP so we are able to clearly identify the county of origin/destination and hence do not have to deal with this issue.

Using the information on PUMA (Public Use Microdata Area) delineations and population shares from the CENSUS retrieved from the Missouri Census Data Center, we built the PUMA-CZ crosswalk needed to calculate the cross-CZ migration rates. In cases where the county-group/PUMA crosses CZ boundaries, the crosswalks assign respondents to CZs in proportion to the share of the county-group/PUMA population that falls within each CZ, using county-level population counts from the Census Bureau to generate the shares. In other words, a single PUMA can be assigned to multiple CZs. An additional complication arises from the fact that for migrants, the location of origin is reported at the MIGPUMA (Migration Public Use Microdata Area) level. MIGPUMAs are groups of one or more PUMAs. When Counties are split into various PUMAS, MIGPUMAS group those PUMAS together. To build the CZ-level variables, we use two strategies:

1. A person is classified as a CZ migrant if she moves from a MIGPUMA with no overlapping CZ with the destination PUMA. This is a conservative way of building the migration variable since it is possible that people moving to a PUMA that does not fully overlap with the CZ are moving to a different labor market.

2. To build aggregates at the CZ level we use the product of the individual's weight provided by the Census with the share of the population in the PUMA that belongs to the CZ. This means that every person in the PUMA counts towards all the CZ in the PUMA. To keep the total of number of people consistent with the total population we multiply the weight by the share of the population in each PUMA that belongs to a CZ.

PUMA delineations change approximately every ten years. Between 2005-18 there were two PUMA delineations, the ones based on the 2000 census applied from 2005 to 2011 and the ones from the 2010 census applied from 2012 to 2018. We are using the annual ACS files, and according to the available documentation, these data are representative of areas with populations over 65,000. We exclude CZ that at any point in the sample had a population below 65,000. We perform the current statistical analysis with the 455 CZ above the population threshold.

References

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

Austin, B., Glaeser, E., and Summers, L. (2018). Jobs for the heartland: Place-based policies in 21st-century america. *Brookings Papers on Economic Activity*, pages 151–232.

Basu, D. (2015). Asymptotic Bias of OLS in the Presence of Reverse Causality. UMASS Amherst Economics Working Papers 2015-18, University of Massachusetts Amherst, Department of Economics.

Blanchard, O. J. and Katz, L. F. (1992). Regional Evolutions. *Brookings Papers on Economic Activity*, 23(1):1–76.

Bloom, D., Canning, D., and Fink, G. (2010). Implications of population ageing for economic growth. *Oxford Review of Economic Policy*, 26(4):583–612.

Blumenstock, J. E., Chi, G., and Tan, X. (2023). Migration and the value of social networks. *The Review of Economic Studies*, 92(1):97–128.

Borusyak, K., Dix-Carneiro, R., and Kovak, B. K. (2022). Understanding migration responses to local shocks. Working paper, Working Paper.

Brevoort, K. P., Grimm, P., and Kambara, M. (2016). Credit invisibles and the unscored. *Cityscape*, 18(2):9–34.

Cadena, B. C. and Kovak, B. K. (2016). Immigrants equilibrate local labor markets: Evidence from the great recession. *American Economic Journal: Applied Economics*, 8(1):257–90.

Cajner, T., Coglianese, J., and Montes, J. (2020). The long-lived cyclicality of the labor force participation rate. Working Paper.

Chen, Y., Ma, Q., and Tao, R. (2024). Addressing the fixed effects in gravity model based on higher-order origin-destination pairs. *Journal of Regional Science*. forthcoming.

Dao, M., Furceri, D., and Loungani, P. (2017). Regional Labor Market Adjustment in the United States: Trend and Cycle. *The Review of Economics and Statistics*, 99(2):243–257.

DeWaard, J., Johnson, J., and Whitaker, S. (2019). Internal migration in the United States: A comprehensive comparative assessment of the Consumer Credit Panel. *Demographic Research*, 41(33):953–1006.

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review*, 106(3):479–524.

Ganong, P. and Shoag, D. (2017). Why has regional income convergence in the u.s. declined? *Journal of Urban Economics*, 102:76–90.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Gonzalez, N. (2020). Hysteresis in the us employment rate revised. Working Paper.

Hershbein, B. J. and Stuart, B. A. (2021). Recessions and Local Labor Market Hysteresis. Working Paper 20-325, Upjohn Institute.

Jia, N., Molloy, R., Smith, C., and Wozniak, A. (2022). The economics of internal migration: Advances and policy questions. Finance and Economics Discussion Series 2022-003, Washington: Board of Governors of the Federal Reserve System.

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

Molloy, R., Smith, C. L., and Wozniak, A. (2011). Internal migration in the united states. *Journal of Economic Perspectives*, 25(3):173–96.

Molloy, R., Smith, C. L., and Wozniak, A. (2017). Job changing and the decline in long-distance migration in the united states. *Demography*, 54(2):631–653.

Monras, J. (2020). Economic shocks and internal migration. Working Paper.

Notowidigdo, M. J. (2020). The Incidence of Local Labor Demand Shocks. *Journal of Labor Economics*, 38(3):687–725.

Olney, W. W. and Thompson, O. (2024). The determinants of declining internal migration. NBER Working Paper, latest version.

Schultz, T. P. (1978). Notes on the Estimation of the Micro Economic Determinants of Migration. Center Discussion Paper 283, Econmic Growth Center, Yale University.

Sjaastad, L. A. (1962). The Costs and Returns of Human Migration. In *Investment in Human Beings*, NBER Chapters, pages 80–93. National Bureau of Economic Research, Inc.

Sprung-Keyser, B., Hendren, N., and Porter, S. (2022). The radius of economic opportunity: Evidence from migration and local labor markets. View young adult migration patterns at migrationpatterns.org.

Wilson, R. (2022). Isolated States of America: The Impact of State Borders on Mobility and Regional Labor Market Adjustments. Working paper, Brigham Young University.

Yagan, D. (2019). Employment hysteresis from the great recession. *Journal of Political Economy*, 1(1):1.

Zipf, G. K. (1946). The p1 p2/d hypothesis: On the intercity movement of persons. *American Sociological Review*, 11(6):677–686.