Mexican Migration Flows and Agricultural Labor Markets in the U.S.

Julian Arteaga* Ashish Shenoy†

June, 2024

Abstract

Using information on migratory flows for every Mexican municipality and U.S. county pair throughout the 2006–2019 period, this paper estimates the effect that variations in Mexican migration flows have on U.S. agricultural labor-market outcomes. We instrument for migration-driven changes in local labor supply using a shift-share variable that combines Mexican municipality-level violence levels with preexisting migration network patterns. Our estimates show that, in the short run, decreasing migration rates put upward pressure on wages across all types of agricultural workers, and cause a large increase in the number of H-2A seasonal worker visas requested by employers. Conversely, in the long run, decreasing migration rates lead to *lower* wages in agriculture accompanied by slight reductions in employment levels. Regarding the mechanisms driving this result, we find that an exogenous decrease in the cumulative number of migrants arriving to a county during this period led to reductions in the acreage planted with labor-intensive crops, higher rates of mechanization, and lower average farmland values.

*jgarteaga@ucdavis.edu. University of California, Davis.
†shenoy@ucdavis.edu. University of California, Davis.

1 Introduction

Relative to population size, immigrant workers play an outsized role in several industries of the U.S. economy. In particular, Mexican-born migrants comprise almost 70% of hired agricultural workers in the country (Hernandez and Gabbard, 2019). However, migration from Mexico to the United States has steadily declined over the past twenty years, and the U.S. farm labor supply has been contracting for more than a decade. The prevalence of labor shortages in the agricultural industry is increasing (Zahniser et al., 2018), and the scarcity of available workers has intensified in the last few years (Peri and Zaiour, 2022). A 2019 survey found that 56% of farms in California reported being unable to meet their full labor demand over the previous five years. Common responses to the shortage included raising of wages, adopting labor-saving technologies, and reducing or delaying various cultivation practices.[1] As summarized by Martin (2020), farm employers must face lower migration flows through some combination of satisfying, stretching, substituting, or supplementing the available workforce. While conceptually any of these strategies seems a viable response to dwindling migration inflows, empirical evidence on the way in which agricultural producers have been adapting to these secular changes in labor availability remains scarce. Understanding the type of adaptation process that is taking place, quantifying the relative importance of potential adjustment margins, and—crucially—knowing how short-run responses differ from more structural adjustments in production practices, is necessary for the correct design of agricultural and migratory policy.

This paper estimates the causal effect of migration inflows on local labor market outcomes in the U.S. agricultural industry. Estimating how changes in worker arrivals affect local labor markets in agriculture has so far remained challenging given that a large fraction of the migrant agricultural workforce is undocumented. Reliable data on this type of migratory flows at a sufficiently disaggregated geographical level has usually been unavailable or lacking in representativity. We overcome this obstacle by using a novel, high quality administrative dataset based on consular identity cards issued by the Mexican government to (predominantly undocumented) workers, and in this way are able to measure the strength of migratory flows between every Mexican municipality and U.S. county pair throughout the 2008–2019 period.

We instrument for migration flows of Mexican workers into U.S. counties by using a shift-share design that combines the sudden and spatially heterogeneous increase in violence across Mexican municipalities beginning in 2008 with preexisting migration networks. Violence-driven migration flows in origin municipalities are plausibly orthogonal to labor demand

[1] California Farm Bureau Federation. 2019. "Still Searching for Solutions: Adapting to Farm Worker Scarcity Survey".

conditions in destination counties, while location choices are well predicted by the share of migrants from the same municipality already established in each potential destination before violence levels rose. Based on this identification strategy, and combining migration flows data with county-level information on agricultural wages, employment, and cultivation practices, we are able to quantify how variations in the supply of migrant workers shape employers' hiring and production decisions across different agricultural sub-industries.

Our estimates reveal a stark contrast between the short-run and the long-run effects that migratory flows have on agricultural labor markets. In the short run, we find that yearly reductions in migration rates from Mexico put upward pressure on wages across all types of agricultural workers, with a one percentage point decrease in migration rates leading to wage increases ranging between 0.22% and 0.42% depending on the specific sub-industry analyzed. These magnitudes are in line with other estimates of the response of wages to migration in contexts where incumbent workers and new arrivals appear to be relatively close substitutes (Kleemans and Magruder, 2018; Imbert et al., 2022).

The rise in wages produced by lower migration is accompanied by a reduction in the number of workers directly hired by producers. This reduction, however, is more than offset by a large increase in the number of H-2A seasonal worker visa requests made by producers, with a one percentage point reduction in migration rates simultaneously causing a reduction of 0.6% in directly-hired employment and an *increase* in H-2A visa requests of 3.4%. To the best of our knowledge, our findings provide the first causal estimate of the elasticity of substitution between permanent immigrant worker arrivals and H-2A temporary visa requests, and confirm the hypothesis that increasing the demand for guest workers is the main margin of adjustment through which U.S. producers adapt, in the short run, to reductions in the labor supply.

By contrast, our estimates for the long-run response to migration flows show that U.S. counties experiencing more severe reductions in migrant-labor supply during the 2008–2019 period had *lower* average wage growth by the end of the period, accompanied by slightly lower levels of agricultural employment. Our estimates indicate that a one percentage point reduction in the cumulative–annualized—migration rate between 2008 and 2019 caused a decrease in average weekly wages ranging between 0.15% to 0.27% across different worker types.

Why do higher immigration rates lead to higher wages? The differences between short and long run responses to migration suggest that both labor markets and the broad agricultural production process is adjusting other less flexible factors of production like land and capital as time goes by. Lower expected labor availability might trigger mechanization processes that substitute production from labor to machinery as documented by (Clemens et al., 2018).

Another explanation could rely on the existence of complementarities between domestic and foreign workers (Ottaviano and Peri, 2012), where increases in the supply of foreign workers spurs productivity increases for incumbent workers leading to higher wages. Higher migration rates might lead to higher long-run wages in a sector if, for example, a larger pool of potential employees allows producers to select more skilled workers that lead to productivity increases. It is also possible that increases in the supply of available labor might spur economic activity in other industries of the economy where a large share of workers are also foreign immigrants (Charlton and Castillo, 2022), with this increased competition for migrant labor across sectors leading to faster wage growth. Finally, the tradable nature of agricultural goods might also imply that a higher supply of workers is eventually adjusted for mostly through an expansion in output and revenues rather than through lower wages or higher unemployment (Burstein et al., 2020).

We explore these potential mechanisms and investigate how long-run variations in migration rates change production decisions by comparing county-level changes in agricultural practices as measured in the Census of Agriculture. This analysis reveals that counties with exogenously fewer migrant arrivals between 2007 and 2017 shifted their crop composition away from labor-intensive crops like vegetables, fruits and tree nuts, and increase instead the area devoted to field crops. We further find that lower migration rates led to relative increases in mechanization, measured both as the change in the absolute value of machinery employed and also in proportion to land area cultivated and number of workers employed. Consistent with migration-driven increases in productivity, our estimates show that higher migration rates cause average farmland values in a county to increase, with a one percentage point increase in annual migration causing a 3.4% increase in farmland values measured in dollars per acre. Finally, we find that higher migration rates do not lead to increases in total output or sales, but rather seem to lead to lower average incomes reported by producers.

This paper contributes to the broad literature on the labor market impacts of immigration (Card, 2001; Borjas, 2003), highlighting how the short run responses to migration flows hinge on the degree of substitutability between incumbent and arriving workers, and by showing how in the long run these effects can be attenuated, or even reversed, through adjustments in other margins. The paper also contributes to the literature focused on studying the causes and consequences of migration from Mexico to the U.S. (Hanson and McIntosh, 2010), and in particular to the effects that the sustained slowdown in migration rates can have in the productive potential of the agricultural sector (Charlton and Taylor, 2016; Rutledge and Mérel, 2023).

The rest of the paper is organized as follows. In Section 2 we describe the general trends

Figure 1: U.S. Agricultural Workers and Mexican Migration Inflows - 2007–2019

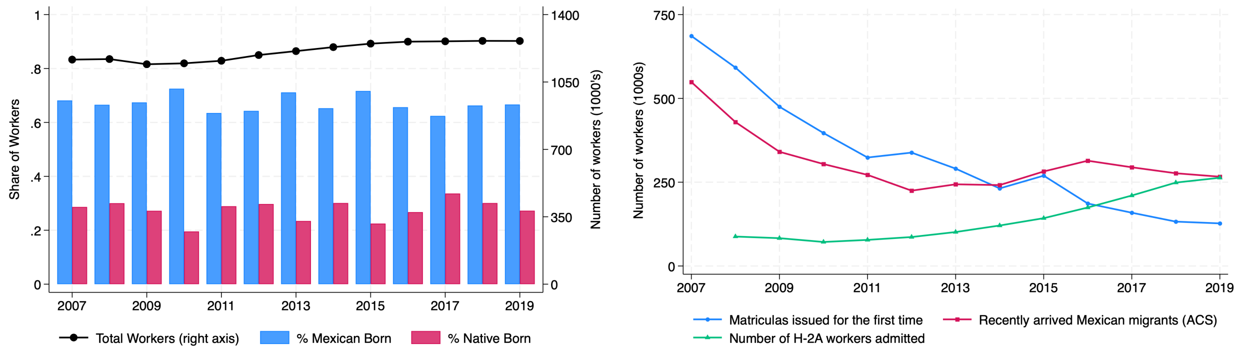

Notes: National-level trends on agricultural sector workers and immigration flows from Mexico to the U.S. Left panel: Total number of hired workers in the agricultural sector according to the Quarterly Census of Employment and Wages (*QCEW*), and share of hired workers who are either Mexican-born or native born according to the National Agricultural Workers' Survey (*NAWS*). Right panel: Number of Matriculas Consulares de Alta Seguridad (*MCAS*) issued for the first time by all Mexican consulates in the U.S. (blue line); number of workers born in Mexico arrived to the U.S. within the previous year according to the American Community Survey (*ACS*) (red line); number of H-2A guest workers admitted by the Department of Labor (green line).

on migratory flows from Mexico and agricultural labor, present the data and describe the samples and variables used for analysis. In Section 3 we motivate the use of our shift-share instrument, and discuss our empirical methodology. Section 4 reports our empirical results, and Section 5 concludes.

2 Data and Background

2.1 Mexican migration and agricultural labor: National-level trends

Our findings coincide with national-level trends showing that, while both the total number of workers and the share of Mexican-born workers hired in agriculture have remained stable during the 2000–2020 period, there has been a substantial drop in immigration flows arriving to the U.S. from Mexico. The sustained decline in worker arrivals from Mexico has been large enough that, at least since 2014, the net flow of migrants from Mexico to the U.S. is roughly zero and might even be slightly negative (Gonzalez-Barrera, 2015). By contrast, as Figure 1 illustrates, the number of temporary H-2A workers being requested by U.S. farmers in order to fill vacant positions increased from under 87,000 yearly workers in 2008, to more than 250,000 in 2019. Taken together, these trends suggest that—against the backdrop of falling permanent immigrant arrivals–U.S. agricultural producers are increasingly relying on H-2A guest workers as a source of labor.

2.2 Migratory flows: the *Matriculas Consulares de Alta Seguridad* data

We measure the yearly inflow of Mexican workers arriving to each U.S. county using information from the *Matrículas Consulares de Alta Seguridad* (MCAS) program maintained by the Mexican government. The MCAS dataset records all *Matrículas Consulares* identification cards issued by Mexican consulates to Mexican-born individuals living in the U.S., and registers both the municipality of origin and the current U.S. county of residence of each cardholder. MCAS are issued to all qualifying Mexican citizens regardless of immigration status or age. While holding a MCAS card does not confer any U.S. immigration status to the person to whom it is issued, many states and local governments allow the document to be used as proof of identity, and so it permits the cardholder to access a services that include opening a bank account, being assigned an Individual Tax Identfication Number (ITIN), or obtaining a driver's license. Since obtaining the MCAS does not entail any additional benefit to authorized migrants, it is generally assumed that MCAS are a measure of unauthorized migration inflows to the U.S. (Massey et al., 2010). Comparing the joint distribution of these inflows both at origin and destination with alternative surveys, Caballero et al. (2018) confirm that MCAS records are in fact a representative and high-quality information source on Mexican migratory flows.

We take yearly MCAS issued for the first time as our main measure of Mexican migration flows to the U.S. from Mexico. Given that MCAS must be renewed every five years it is important, when trying to measure yearly migrant inflows, to separate renewals from first-issuances. The right panel of Figure 1 shows the close correspondence between the observed number of MCAS issued for the first time and the number of newly-arrived Mexican migrants to the U.S. as recorded in the *American Community Survey* (ACS), and is consistent with the documented decline in migration inflows throughout this period (Passel and Cohn, 2018). Figure A1 in Appendix A further shows the correspondence between ACS and MCAS data for the four states in the country with highest immigration rates. A detailed analysis of the validity of MCAS data as a measure of migration flows across time can be found in Tiburcio and Camarena (2023). The spatial distribution of municipality-level outflows, and county-level inflows of migrants is shown in Figure 2.

2.3 Mexican migration networks

The MCAS records allow us to build a measure of the strength of migratory networks between each Mexican municipality of origin and destination U.S. county pair. A large literature (see for example Munshi (2003, 2014)) shows that preexisting migrant community networks

Figure 2: Migration rates in MCAS data – Municipality-level outflows and County-level inflows

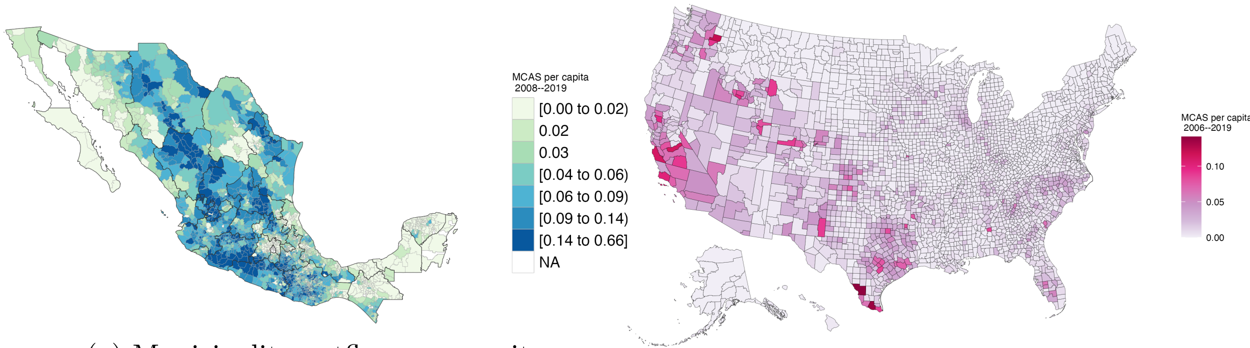

(a) Municipality outflows per capita

(b) County inflows per capita

at destination are a fundamental determinant of the location choice of future migrants. To compute this network measure we use the first two years in our data (i.e. 2006 and 2007) and calculate the share of migrants going to each U.S. county out of the total number of migrants leaving each Mexican municipality during this two-year period. As documented by Tian et al. (2022), we find large differences in the historical destination patterns across different Mexican municipalities, even within the same state. We leverage this spatial variation in the settlement patterns of different origin communities to build the shares of our shift-share instrument for migration inflows. Figure A2 in Appendix A illustrates the differences in migration networks for a pair of nearby municipalities.

2.4 Violence in Mexico

Starting in 2008, Mexico has suffered an unprecedented and drastic increase in violence levels across large parts of its territory. As shown in Figure 3, national homicide rates increased by 300% in 12 years, from 8.6 homicides per 100,000 population in 2007 to 34.4 homicides per 100,000 population in 2019 . While most violence in Mexico is closely related to the illicit drug market supply to the U.S. and the governmental stance on the war on drugs, the specific causes of the sudden surge in homicide rates remain a source of debate (Castillo et al., 2020; Guerrero-Gutiérrez, 2011; Dell, 2015; Williams, 2012). Regardless of their cause, these rapid increases in violence levels across the country have been shown to have negative impacts on labor-market outcomes (Velásquez, 2020), capital accumulation (Brown and Velásquez, 2017), and migration decisions (Orozco-Aleman and Gonzalez-Lozano, 2018).

We use yearly municipal-level data on violent homicide rates from INEGI, Mexico's national statistical agency. As Figure 3 shows, the growth in violence levels has been heterogeneous across space, with some states suffering a more-than-tenfold increase in murder

Figure 3: Homicide Rates in Mexico – 2000–2019

(a) National homicide rate 2000-2019 (b) Homicide rates - selected States

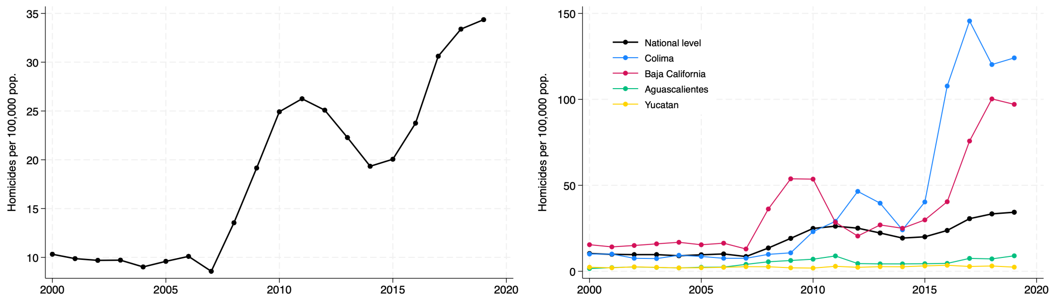

(c) Municipal homicide rates 2000–2007 (d) Municipal homicide rates 2008–2019

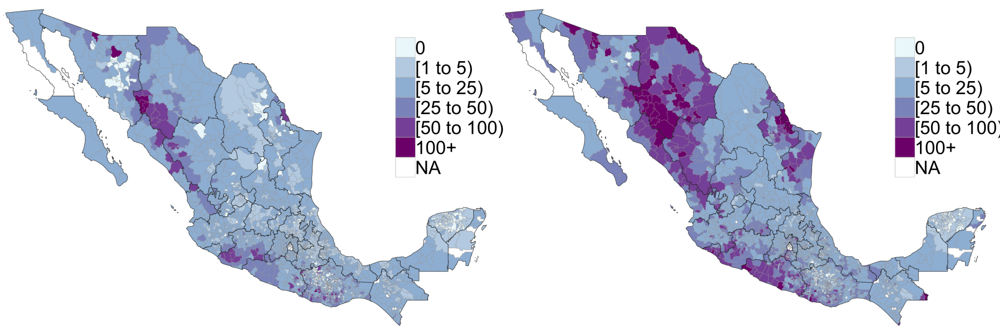

Notes: Yearly homicide rates per 100,000 population. Source: Mexican national statistics institute (INEGI).

rates while others showing no increase at all throughout the period. Combining this data with the MCAS information described above, in Section 3 we document the existence of a strong positive correlation between yearly municipal violence levels and migrant outflows to the U.S.

The spatial variation in the increase in violence across the 2008–2019 period acts as the 'shift' component of our shift-share instrument. By combining this component with municipal-county network shares we are able to aggregate destination origin shocks into a yearly destination-level measure of migrant-supply shocks.

2.5 US agricultural wages and employment

To measure labor market outcomes we use county-industry-year level data on wages and employment from the *Quarterly Census of Employment and Wages* (QCEW) program mantained by the U.S. bureau of labor statistics. The QCEW publishes a quarterly count of establishments, employment level, and total wage bill at the 6-digit NAICS industry level for each county in the United States. It is based on the aggregation of all Quarterly Contribution Reports (QCRs) submitted by employers subject to state and federal unemployment insurance laws to each State's accounting system. The QCEW covers more than 95% OF U.S. jobs. We focus on employment levels and wages recorded in the QCEW under NAICS codes 11 (agriculture, forestry, fishing and hunting), 111 (crop production), and 115 (Agriculture and forestry support activities). Figure A3 in Appendix A show national-level trends of employment and wages for these sub-industries. Agricultural labor tends to be highly seasonal, and about half of all hired workers tend to be hired directly and employed on crop production (NAICS 111). Indirectly-hired workers in agriculture support activities (NAICS 115) make up roughly 20% of the total workforce and, until recently, tended to have lower earnings than their directly-hired counterparts.

We also estimate the impact of migration on the yearly rates of H-2A workers requested for authorization by agricultural employers in each U.S. county. Data on these requests are publicly available at the individual request level from the Department of Labor. [2] To get yearly county-level requests we aggregate the total unique number of workers certified on each intended worksite according to the stated job start date on each application.[3] Figure 1 shows yearly trends in national H-2A request levels, and Figure A6 in Appendix A shows the change in the distribution of H-2A requests across space. H-2A requests are concentrated in the southern states of the U.S.—Louisiana, Florida, North Carolina—, as well as in California, Colorado and the state of Washington. Since 2008, the number of certified H-2A seasonal workers has more than doubled from under 100,000 workers to nearly 250,000 workers in 2018.

We use data from the 2007 and 2017 versions of the Census of Agriculture (*COA*) to measure the effect of migration on agricultural practices. The Census of Agriculture, carried out every 5 years, is a complete count of all farms and ranches in the country and the people who operate them. It collects information on land use and ownership, operator characteristics, production practices, income and expenditures.

3 Methodology

3.1 Instrumental variable motivation

Assessing whether changes in migration rates have an effect on wages or employment rates is challenging due to the fact that observed wages and employment are equilibrium outcomes that are endogenous to labor supply, of which migration is only one component, and to

[2]https://www.dol.gov/agencies/eta/foreign-labor/performance

[3]In some years intended worksites are specified either as cities or zip codes. We harmonize all locations at the U.S. county level using the crosswalks provided by the Missouri Census Data Center https://mcdc.missouri.edu/applications/geocorr2018.html.

labor demand, which is affected by the local economic cycle. To overcome this challenge we use the interaction of changes in violence levels across Mexican municipalities with observed preexisting migration networks as an instrument for the number of migrants arriving to each destination county every year. This instrument is based on the observation that variations in the intensity of violence in origin locations are drivers of the decision to migrate, and that the destination choice is further driven by the strength of social networks created by previous migration waves.

To fix ideas, let $M_{c,t}$ be the number of migrants arriving to U.S. county c on year t. (i.e. the number of first-time issued MCAS assigned to county c). Our goal is to find an instrument for the yearly Mexican immigration rate to each U.S. county c:

$$m_{c,t} = \frac{M_{c,t}}{P_{c,t^0}}$$

while the emigration rate leaving to the U.S. from municipality o on year t is

$$n_{o,t} = \frac{M_{o,t}}{P_{o,t^0}}$$

where $P_{o,t}$ and $P_{c,t}$ are, respectively, municipality and county populations at year t and t^0 indicates a year prior to t.[4]

To test the premise that violence levels at origin municipalities are in fact correlated with migratory outflows we regress yearly municipality-level emigration rates $n_{o,t}$ on yearly homicide rates $V_{o,t} = \frac{\text{Homicides}_{o,t}}{P_{o,t^0}}$

$$n_{o,t} = \alpha + \beta V_{o,t} + \delta_t + \gamma_o + \varepsilon_{o,t} \tag{1}$$

where δ_t and γ_o are respectively year and municipality fixed effects.

Results for regression 1 are shown in Table 1. After accounting for both year and municipality fixed effects, our point estimate indicates that, on average, a 10 percentage point increment in the homicide rate is associated with a 2.8 percentage point increase in municipal emigration rates to the U.S. The magnitude of this correlation is similar to other estimates of the violence-at-origin effect on U.S. migration rates coming from other Central American countries (Clemens, 2021).

While we cannot tell if the observed association of violence and migrant outflows is causal, the existence of this strong correlation is enough motivation to use violence rates

[4] In practice, we normalize all of our per capita variables according to the 2005 county and municipality population estimates calculated respectively by the U.S. Census and INEGI.

Table 1: Homicide rates and yearly emigration rates – Mexican municipality level

	Yearly emigration rate ($n_{o,t}$)			
	(1)	(2)	(3)	(4)
Homicides per capita	0.956***	1.293***	-0.213	0.283**
	(0.176)	(0.188)	(0.178)	(0.138)
Observations	29232	29232	29232	29232
Municipalities	2436	2436	2436	2436
Year FE	No	Yes	No	Yes
Municipality FE	No	No	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the municipality level in parentheses.

as the 'shift' component of our shift-share design. The intuition behind the second part of our instrument is that origin-destination migratory flows can be accurately predicted from aggregate municipality-level outflows multiplied by a measure of the strength of the historical settlement network between each municipality and each U.S. county. More precisely,

$$m_{c,t} = \frac{M_{c,t}}{P_{c,t^0}} = \frac{1}{P_{c,t^0}} \sum_o M_{o,c,t} \approx \frac{1}{P_{c,t^0}} \sum_o \left[M_{o,t} \times \phi_{o,c}^{t^0} \right] = \frac{1}{P_{c,t^0}} \sum_o \left[(n_{o,t} \times P_{mot^0}) \times \phi_{o,c} \right] \quad (2)$$

where $M_{o,c,t}$ is the migration flow from m to c in t, and the share of total migrants from municipality m that arrived to county c during the 2006–2007 two-year period:

$$\phi_{o,c} \equiv \frac{M_{o,c,t^0}}{\sum_c M_{o,c,t^0}}$$

is our measure of migrant-network strength. Whether this measure is indeed a good predictor of subsequent migrant location decisions can be evaluated in the data: Figure A7 shows that this is indeed the case when using the MCAS data, and that network-predicted migration flows (i.e. $[M_{o,t} \times \phi_{o,c}]$ in equation 2 are accurate predictors of observed county-municipality migration flows $M_{o,c,t}$.

Leveraging the fact that changes in municipal homicide rates influence emigration intensity, and that historical migration patterns are good predictors of destination choice, we construct the following shift-share instrumental variable:

$$Z_{c,t} = \frac{1}{P_{c,t^0}} \sum_o \left[\text{Homicides}_{o,t} \times \phi_{o,c}^{t^0} \right] \quad (3)$$

Table 2: County level immigration rates and origin violence shocks – First-stage estimates

	Yearly immigration rate ($m_{c,t}$)			
	(1)	(2)	(3)	(4)
$Z_{c,t}$: Violence shift-share IV	7.617***	7.862***	-5.318***	-4.272***
	(0.463)	(0.486)	(0.869)	(0.774)
Observations	37680	37680	37680	37680
Counties	3140	3140	3140	3140
Year FE	No	Yes	No	Yes
County FE	No	No	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

3.2 First-stage results

To evaluate if the instrument is a strong predictor of county-level migrant inflows we run the following regression:

$$m_{c,t} = \alpha + \delta Z_{c,t} + \delta_t + \gamma_c + \varepsilon_{c,t} \tag{4}$$

where δ_t and γ_c are respectively year and county fixed effects. Results for regression equation 4 are shown in Table 2. While the simple pooled cross-section comparison of county-year observations shows a strong positive relationship between the instrument and migration rates, once county fixed effects are included and unobserved time-invariant county characteristics are accounted for, this relationship reverses and becomes strongly negative. This is a surprising result and could be subject to a number of different explanations. Additional estimations in appendix B show that this negative relationship is not due to *i)* The networks component ($\phi_{o,c}$) of the shift-share variable, *ii)* Noisiness of the yearly data, nor *iii)* The aggregation of violence measures across various municipalities.

We interpret the observed results as suggesting that, when comparing across counties, higher average violence levels in the group of municipalities associated to each county are strongly correlated to higher immigration rates, but that this migration tends to happen in relatively less-violent years. That is, while the yearly comparison across counties shows there is a clear positive relationship between violence and migration rates—i.e. counties with stronger connections to more violent municipalities have higher immigration rates—a within-county comparison yields that yearly deviations from county trend in aggregate municipal violence levels is negatively associated to migratory flows when compared to less violent years.

This interpretation is consistent with a model where violence-induced outmigration is determined by a long-run violence component that follows the migrant network decision rule, and by a short-run violence component that follows some independent short-run decision rule.

Appendix C describes in more detail such a model, and shows that estimations based on the simplest version of this model—carried out on simulated data—are capable of replicating the change in sign of the first-stage regression coefficients observed after the inclusion of county fixed effects.

3.3 Empirical Strategy

To estimate the short-run effect of changes in migration to county c in year t we estimate a regression of the form:

$$y_{c,t} = \alpha + \beta^S m_{c,t} + X'_{c,t}\gamma + \tau_c + \delta_t + e_{c,t}, \tag{5}$$

where the dependent variable yc, t denotes some county-level outcome, $m_{c,t}$ is the immigration rate in county c at year t, and τ_c and δ_t are county and year fixed effects. The vector $X'_{c,t}$ includes the state-level minimum wage at year t and a Bartik-style shock that controls for time-varying changes to local labor demand.[5] Equation (5) is estimated through two-stage least squares using the instrumental variable defined in equation (3).

To estimate the long-run effect of migration inflows on labor markets we compute the county-level change in outcomes between 2008 and 2019 and regress it on the annualized sum of yearly migration flows relative to county baseline population for the same period:

$$\Delta y_c = \alpha + \beta^L \tilde{m}_c + X'_c\rho + v_c, \tag{6}$$

where Δy_c denotes the long difference in some county level outcome between 2008 and 2019, and $\tilde{m}_c \equiv \frac{1}{12}\sum_{t=2008}^{2019} m_{c,t}$, is the cumulative annualized migration rate. The vector X'_c is composed of controls for minimum-wage growth and long-differences in Bartik-style labor demand shocks. Migration rates are instrumented with the sum of the instrument across all periods $Z_c = \sum_{t=2008}^{2019} Z_{c,t}$. All standard errors in both short-run and long-run specifications are clustered at the county level.

Sample choice and regression weights In order to protect respondent's confidentiality, the published QCEW data suppresses information for industry-quarter-county combinations where the number of establishments is deemed small enough as to make individual

[5]The inclusion of these controls is meant to control for time-varying unobservable characteristics that might affect county-level migration inflows. However, excluding them from the regressions does not have any impact in the results for any outcome. The Bartik shock is computed from the Census *County Business Patterns* (CBP) data. It combines the 2-digit NAICS code industry composition of each county in 2007 with yearly national-level industry growth rates.

information identifiable. This implies that the QCEW is, in practice, an unbalanced panel where industry-specific information for a given county in a given year might be undisclosed due to the small number of respondents. For each of the agricultural sub-industries mentioned above, our baseline estimation sample uses all counties that have complete information for all years in the analysis period for the specific sub-industry. Each regression in our baseline results is therefore carried out on a balanced panel of counties, but the set of counties can vary across sub-industry. To test for the sensitivity of our results, we estimate all main results using two alternative estimating samples that respectively consist of *i)* all county-year observations available, and *ii)* the 'fully restricted' sample of counties that have complete information for all industries in all years.

Similarly, our results could be dependent on the choice of weights specified in each estimation. While in our baseline results all counties in the estimating sample are given equal weights, we also test for the sensitivity of results relative to this assumption by estimating regressions weighted by *i)* baseline county population, and *ii)* baseline farm employment levels.[6] Results for these alternative specifications are shown in Figures A4 and A5 in Appendix A. In general terms we find our results are robust to the choice of estimation sample and weighting scheme.[7]

4 Results

4.1 Short-run Results

We estimate the impact of yearly migration flows on wages and employment for four different groups of workers: *i)* all workers hired in the agricultural sector (NAICS 11), *ii)* workers directly hired by employers for crop production (NAICS 111), *iii)* indirectly hired workers for crop support activities (NAICS 115), and *iv)* H-2A seasonal guest workers. Table 3 reports the OLS and IV estimates of equation (5) on the (log) average weekly wage for each of these worker types, while estimates for the effect on (log) employment levels are displayed on Table 4.

Panel A of Table 3 shows that variations in the number of Mexican migrants arriving to a county have in general a weakly negatively correlation with average agricultural wages. However, given that migration is endogenous to labor demand, and migrants are

[6] Baseline defined as population and employment in 2005.

[7] Results only differ qualitatively for a single outcome, long run H-2A employment, for the case of 'fully restricted' sample and for the case of baseline farm employment weights. Both of these specification give particular importance to counties with relatively large agricultural sectors, suggesting that, for counties highly specialized in agriculture, the long-run impact of higher migration rates on the prevalence of the H-2A worker program might be actually positive instead of null.

Table 3: Migration and Agricultural Wages – Short run effects

	(1) Total Agricultural	(2) Directly Hired	(3) Contract Labor	(4) H-2A Workers
Panel A: OLS Estimates				
Migration Rate ($m_{c,t}$)	-0.010	-0.064***	-0.028	-0.030
	(0.021)	(0.023)	(0.050)	(0.022)
N	11508	11808	4572	7848
Num Counties	959	984	381	654
Panel B: IV Estimates				
Migration Rate ($m_{c,t}$)	-0.221**	-0.325***	-0.323***	-0.416**
	(0.107)	(0.069)	(0.118)	(0.170)
N	11508	11808	4572	7848
Num Counties	959	984	381	654
Kleinberg-Paap F	10.49	17.4	7.28	9.955
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

likely attracted to labor markets where wages are higher, the OLS estimates are likely underestimating the causal effect of migration on wages. Indeed, the IV estimates in Panel B of the same table show that increases in yearly migration rates have a much stronger negative effect on wages than the one suggested by the OLS regressions. A one percentage point increase in migration rates causes a reduction of 0.22% in average weekly wages for all workers in the agricultural sector, and this effect intensifies when focusing both on the directly hired or contract labor workers. The largest reduction in average wages due to migration (0.41%) is on the wages paid to H-2A guest workers.

Similarly, the OLS estimates for the correlation between migration and employment shown in panel A of Table 4 underestimate the true effect. The IV estimates show that while a one percentage point decrease in yearly migration rates changes the number of workers hired directly by farmers by -0.63%, this reduction is more than compensated by the increase in the number of H-2A visa requests, which, at 3.4% is about five times larger. For its part, increases in migration also appear to have a negative effect with the number of indirectly-hired contract workers, but this relationship is not statistically distinguishable from zero.

Both sets of results suggest that, in the short run, agricultural labor demand in the U.S. is relatively inelastic, and that fluctuations in labor supply lead to wage responses across all type of farm workers as well as to large increases in the demand for foreign, seasonal

Table 4: Migration and Employment in Agriculture – Short run effects

	(1) Total Agricultural	(2) Directly Hired	(3) Contract Labor	(4) H-2A Workers
Panel A: OLS Estimates				
Migration Rate ($m_{c,t}$)	0.060	0.233***	-0.115	-0.515**
	(0.068)	(0.074)	(0.103)	(0.221)
N	11508	11808	4572	7908
Num Counties	959	984	381	659
Panel B: IV Estimates				
Migration Rate ($m_{c,t}$)	0.246	0.630***	-0.237	-3.418***
	(0.285)	(0.242)	(0.238)	(1.148)
N	11508	11808	4572	7908
Num Counties	959	984	381	659
Kleinberg-Paap F	10.49	17.4	7.28	10
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

guest workers. This is consistent with the seasonal characteristics of agricultural production, where production decisions like crop choice and technology decisions have to be taken well in advance of labor recruitment. However, while yearly labor demand might be relatively fixed, long run adjustments on other margins of production might lead to changes in the elasticity of labor demand as shown below.

4.2 Long-run Results

Tables 5 and 6 present the results from estimating equation (6) on the log-difference of employment and wages between 2008 and 2019 for the same four categories of agricultural workers.

In contrast to short-run responses, these long-difference estimates reveal that, in the long run, increased migration rates induce *higher* average wage growth for all types of agricultural workers. In absolute value, the magnitude of the long-run wage responses to migration is smaller than the short-run response, with a once percentage point increase in annualized migration rates leading to average wage increases of between 0.16% and 0.26%. The weaker magnitude in wage response is consistent with the idea that adjustments in other factors of production across time allow labor demand to become less inelastic.

Regarding employment, the long difference estimates show no effect of higher migration

Table 5: Migration and Agricultural Wages – Long run effects

	(1) Total Agricultural	(2) Directly Hired	(3) Contract Labor	(4) H-2A Workers
Panel A: OLS Estimates				
Migration Rate (m_c)	0.130**	0.268***	0.232**	0.070
	(0.051)	(0.047)	(0.115)	(0.045)
N	959	984	381	654
Num Counties	959	984	381	654
Panel B: IV Estimates				
Migration Rate (m_c)	0.170***	0.235***	0.265**	0.157*
	(0.059)	(0.057)	(0.129)	(0.083)
N	959	984	381	654
Num Counties	959	984	381	654
Kleinberg-Paap F	144.98	130.02	48.21	131.96
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

rates on overall employment rates in agriculture or in the number of seasonal agricultural workers requested. We find increased migration leads to a relatively small increase in the number of indirectly-hired contract workers. The IV estimates stand in contrast with the OLS coefficients, which show a positive correlation between migration rates the log difference in the number of H-2A workers requested by employers in a county.[8]

4.3 Mechanisms

There are several potential explanations for the observed long-run effects of migration on labor markets. One possibility is the existence of potential complementarities between incumbent and newly-arrived workers that might lead to deeper specialization in tasks and higher productivity (Ottaviano and Peri, 2012; Peri and Sparber, 2009). Alternatively, agricultural firms can eventually adjust their capital-labor ratios and technology choices, substituting production away from labor and focusing in less labor-intensive operations (Clemens et al., 2018). Another potential explanation might be that, due to the fact that the agricultural sector produces mainly tradable goods, increases in the relative abundance of

[8]This positive relationship between migration and the growth of the H-2A program could be indicative of seasonal guest worker programs leading to higher migration rates in the long-run. The public DOL data on individual H-2A requests made by employers does not record the place of origin of migrant workers being employed on each county under the program, so it is not possible to correlate the origin location of seasonal guest workers with the origin location of independent migration flows recorded in the MCAS data.

Table 6: Migration and Employment in Agriculture – Long run effects

	(1) Total Agricultural	(2) Directly Hired	(3) Contract Labor	(4) H-2A Workers
Panel A: OLS Estimates				
Migration Rate (m_c)	-0.154	-0.571***	0.257	1.105*
	(0.161)	(0.168)	(0.234)	(0.585)
N	959	984	381	659
Num Counties	959	984	381	659
Panel B: IV Estimates				
Migration Rate (m_c)	0.072	-0.231	0.420*	0.088
	(0.249)	(0.317)	(0.251)	(0.778)
N	959	984	381	659
Num Counties	959	984	381	659
Kleinberg-Paap F	144.98	130.02	48.21	133.843
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

labor might cause agricultural firms to make most of the adjustment through an expansion in output rather than through prices, as suggested by (Burstein et al., 2020).

We investigate which of these explanations are driving the observed long-run effects on wages and employment by analyzing how changes in migration rates differentially change agricultural production outcomes between 2007 and 2017 using data from the Census of Agriculture.

First, Table 7 presents the IV results of estimating equation (6) on outcomes related to the number of farms, farm operators, and crop choice. Columns 1, 2, and 3, show that counties with higher migration rates between 2007 and 2017 had as a result a higher total number farms, more total farm operators, and fewer total acres of cropland planted.[9] Farm operators are defined as individuals involved in the day-to-day decisions of the farm excluding all hired workers unless they were hired specifically as managers. The increase in the total number of operators and in operators per acre indicates a higher proportion of the agricultural workforce shifted towards more managerial roles in counties where migration flows were higher. This is consistent with a task-specialization process where increases in the supply of one type of worker allow other complementary types of workers to become more productive and earn higher wages.

[9]The reduction in total agricultural area is also observed if the measure of total area operated (i.e. including pastures and woodland) is used instead of total cropland.

Table 7: Long run impacts of migration – Task specialization and cultivation practices

			Area Harvested					Area Operated	
	(1) Number of farms	(2) Number of Operators	(3) Total cropland	(4) Field crops	(5) Fruits and tree nuts	(6) Horticulture	(7) Vegetables	(8) Family operated	(9) Non-family operated
Migration Rate (m_c)	1.794***	1.561**	-3.709***	-3.421***	1.650	7.631*	10.382**	-2.259	-18.555***
	(0.586)	(0.634)	(1.015)	(1.069)	(1.972)	(4.318)	(4.966)	(2.289)	(2.570)
Observations	3046	3045	3037	3011	1987	1002	1918	2849	2580
Kleinberg-Paap F	216.388	216.376	216.164	201.812	149.53	80.117	124.088	194.558	173.16

Notes: All outcomes are from the Census of Agriculture and are computed as the county-level log difference between 2007 and 2017. *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

Table 8: Long run impacts of migration – Input substitution

	Machinery value - Owned				Machinery expenses - Rented			
	(1) Total Value	(2) $\frac{\text{Value}}{\text{acre}}$	(3) $\frac{\text{Value}}{\text{worker}}$	(4) $\frac{\text{Value}}{\text{farm}}$	(5) Total Value	(6) $\frac{\text{Value}}{\text{acre}}$	(7) $\frac{\text{Value}}{\text{worker}}$	(8) $\frac{\text{Value}}{\text{farm}}$
Migration Rate (m_c)	-1.812**	-0.682	-0.815	-3.533***	-5.496***	-4.983**	-3.875**	-7.496***
	(0.793)	(0.645)	(0.890)	(0.650)	(2.008)	(1.977)	(1.921)	(2.031)
Observations	3042	3028	2972	3042	2620	2615	2609	2620
Kleinberg-Paap F	260.401	248.476	249.431	260.401	202.052	195.672	201.614	202.052

Notes: All outcomes are from the Census of Agriculture and are computed as the county-level log difference between 2007 and 2017. *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

Regarding crop choice, columns 4 to 7 of Table 7 show that counties where migration rates fell most shifted away from labor-intensive crops (i.e. fruits and tree nuts, horticulture, and vegetables) and experience instead an increase in total acreage devoted to field crops which, in general, demand fewer workers per acre. Consistent with this change in crop composition patterns, results shown in table 8 show that counties where exogenous increases in migration rates took place had as a consequence lower rates of mechanization, measured both as the reported asset value of all machinery owned by producers, and as the cost of machinery rentals. The impact of migration on mechanization remains negative whether measured as total value, or in proportion to total acres harvested, workers hired, or farms in operation. This result is consistent with other evidence that show how reductions in the supply of low-wage labor lead to long-run mechanization (Hornbeck and Naidu, 2014).

Finally, the results presented in Table 9 explore if the positive effect of migration on wages can be explained through increases in land productivity or total ouptut. Columns 1 and 2 reveal that higher migration rates cause average farmland values in a county, measured in dollars per acre, to increase in the long run. This increase in farmland values is indicative of

Table 9: Long run impacts of migration – Productivity

	Farmland Value		Total Income Reported - All commodities			
	(1)	(2)	(3)	(4)	(5)	(6)
	Total Value	$\frac{\text{Value}}{\text{acre}}$	Total Value	$\frac{\text{Value}}{\text{acre}}$	$\frac{\text{Value}}{\text{worker}}$	$\frac{\text{Value}}{\text{farm}}$
Migration Rate (m_c)	2.127**	3.378***	-3.054***	-2.595***	-2.039**	-5.073***
	(0.826)	(0.882)	(1.009)	(0.919)	(0.879)	(1.064)
N	3043	3030	2983	2972	2925	2983
Kleinberg-Paap F	216.096	206.269	205.693	196.106	197.014	205.693

Notes: All outcomes are from the Census of Agriculture and are computed as the county-level log difference between 2007 and 2017. *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the county level in parentheses.

higher land productivity and suggests that increasing migration rates, perhaps by expanding the pool of potential employees allows producers to select more skilled workers that lead to productivity increases. However, columns 3 to 6 of the same table show that higher migration rates caused reductions in the average income reported by farmers, irrespective of whether it is measured in proportion to land cultivated, number of workers hired, or number of agricultural operations.[10] These results suggest that increased migration rates do not seem to translate into expansions of agricultural output.

5 Conclusion

Industry-wide labor shortages and migratory policy are two extensively discussed economic issues in current policy debates. This paper contributes to that debate by quantifying the relative importance of different margins of adjustment through which employers are responding to changes in labor supply, and by showing how these adjustments change over time. Our results quantify the effect that changes in migration flows from Mexico have on the organization of the U.S. agricultural sector. We show the specific way in which agricultural labor markets and farmers' production decisions respond to fluctuations in migrant labor supply, and uncover stark differences in these responses between the short and the long run. While yearly reductions in migration rates push up agricultural wages and lead producers to compensate labor scarcity by increasing their demand for seasonal guest workers, in the long run local economies with fewer migrant arrivals experience

[10]Similar results are obtained if total revenues (i.e. total value of sales of agricultural products) is used instead of income.

broader changes in their agricultural industry. These changes—related to producers crop choices and production practices—have led counties with secular slowdowns in migration rates to experience decreases in agricultural employment and agricultural salaries, as well as reductions in average farmland values. Understanding if these migration-driven changes in production have economically important effects on the price of agricultural products faced by consumers, and whether changes in agricultural productivity might have spillover effects on other sectors of the economy where a large share of workers are also foreign immigrants are two potentially fruitful avenues for future research.

References

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

Brown, R. and Velásquez, A. (2017). The effect of violent crime on the human capital accumulation of young adults. *Journal of development economics*, 127:1–12.

Burstein, A., Hanson, G., Tian, L., and Vogel, J. (2020). Tradability and the labor-market impact of immigration: Theory and evidence from the united states. *Econometrica*, 88(3):1071–1112.

Caballero, M. E., Cadena, B. C., and Kovak, B. K. (2018). Measuring geographic migration patterns using matrìculas consulares. *Demography*, 55(3):1119–1145.

Card, D. (2001). Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics*, 19(1):22–64.

Castillo, J. C., Mejía, D., and Restrepo, P. (2020). Scarcity without leviathan: The violent effects of cocaine supply shortages in the mexican drug war. *Review of Economics and Statistics*, 102(2):269–286.

Charlton, D. and Castillo, M. (2022). Housing booms and h-2a agricultural guest worker employment. *American Journal of Agricultural Economics*, Forthcoming.

Charlton, D. and Taylor, J. E. (2016). A declining farm workforce: Analysis of panel data from rural mexico. *American Journal of Agricultural Economics*, 98(4):1158–1180.

Clemens, M. A. (2021). Violence, development, and migration waves: Evidence from central american child migrant apprehensions. *Journal of Urban Economics*, 124:103355.

Clemens, M. A., Lewis, E. G., and Postel, H. M. (2018). Immigration restrictions as active labor market policy: Evidence from the mexican bracero exclusion. *American Economic Review*, 108(6):1468–87.

Dell, M. (2015). Trafficking networks and the mexican drug war. *American Economic Review*, 105(6):1738–79.

Gonzalez-Barrera, A. (2015). More mexicans leaving than coming to the u.s. *Pew Research Center*, 19.

Guerrero-Gutiérrez, E. (2011). Security, drugs, and violence in mexico: A survey. In *7th North American Forum, Washington DC*. sn.

Hanson, G. H. and McIntosh, C. (2010). The Great Mexican Emigration. *The Review of Economics and Statistics*, 92(4):798–810.

Hernandez, T. and Gabbard, S. (2019). Findings from the national agricultural workers survey (naws) 2015–2016. a demographic and employment profile of united states farmworkers. *Department of Labor, Employment and Training Administration, Washington; 2019. Report No. 13*.

Hornbeck, R. and Naidu, S. (2014). When the levee breaks: Black migration and economic development in the american south. *American Economic Review*, 104(3):963–90.

Imbert, C., Seror, M., Zhang, Y., and Zylberberg, Y. (2022). Migrants and firms: Evidence from china. *American Economic Review*, 112(6):1885–1914.

Kleemans, M. and Magruder, J. (2018). Labour market responses to immigration: Evidence from internal migration driven by weather shocks. *The Economic Journal*, 128(613):2032–2065.

Martin, P. L. (2020). Immigration and farm labor: Challenges and opportunities. In Martin, P., Goodhue, R., and Brian, W., editors, *California agriculture: dimensions and issues*. Giannini Foundation of Agricultural Economics.

Massey, D. S., Rugh, J. S., and Pren, K. A. (2010). The geography of undocumented mexican migration. *Mexican studies/Estudios mexicanos*, 26(1):129–152.

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the us labor market. *The Quarterly Journal of Economics*, 118(2):549–599.

Munshi, K. (2014). Community networks and migration. In *The Oxford handbook of the economics of networks*.

Orozco-Aleman, S. and Gonzalez-Lozano, H. (2018). Drug violence and migration flows lessons from the mexican drug war. *Journal of Human Resources*, 53(3):717–749.

Ottaviano, G. I. and Peri, G. (2012). Rethinking the effect of immigration on wages. *Journal of the European economic association*, 10(1):152–197.

Passel, J. S. and Cohn, D. (2018). Us unauthorized immigrant total dips to lowest level in a decade. *Pew Research Center*, 17.

Peri, G. and Sparber, C. (2009). Task specialization, immigration, and wages. *American Economic Journal: Applied Economics*, 1(3):135–69.

Peri, G. and Zaiour, R. (2022). Labor shortages and the immigration shortfall. *Econofact*.

Rutledge, Z. and Mérel, P. (2023). Farm labor supply and fruit and vegetable production. *American Journal of Agricultural Economics*, 105(2):644–673.

Tian, Y., Caballero, M. E., and Kovak, B. K. (2022). Social learning along international migrant networks. *Journal of Economic Behavior & Organization*, 195:103–121.

Tiburcio, E. and Camarena, K. R. (2023). The local reaction to unauthorized mexican migration to the us. *Working Paper*.

Velásquez, A. (2020). The economic burden of crime evidence from mexico. *Journal of Human Resources*, 55(4):1287–1318.

Williams, P. (2012). The terrorism debate over mexican drug trafficking violence. *Terrorism and Political Violence*, 24(2):259–278.

Zahniser, S., Taylor, J. E., Hertz, T., and Charlton, D. (2018). Farm labor markets in the united states and mexico pose challenges for us agriculture. *United States Department of Agriculture, Economic Research Service*.

Appendix A Additional tables and figures

Figure A1: Migration inflows across data sources – State level trends

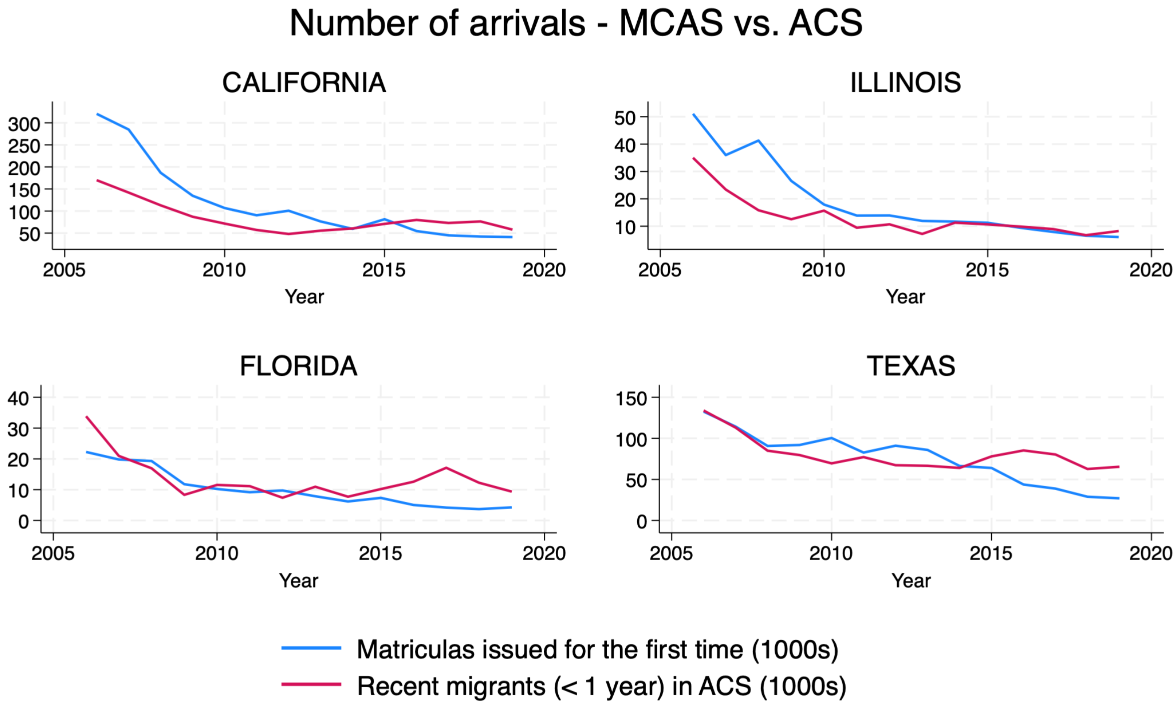

Notes: Number of yearly arrivals from Mexico comparing MCAS and ACS. Blue lines show the number of Matriculas Consulares de Alta Seguridad (*MCAS*) issued for the first time by all Mexican consulates in each state. Red lines show the number of migrants recorded in the ACS as having moved from Mexico to each state within the previous year.

Figure A2: Migration network differences across two Mexican municipalities

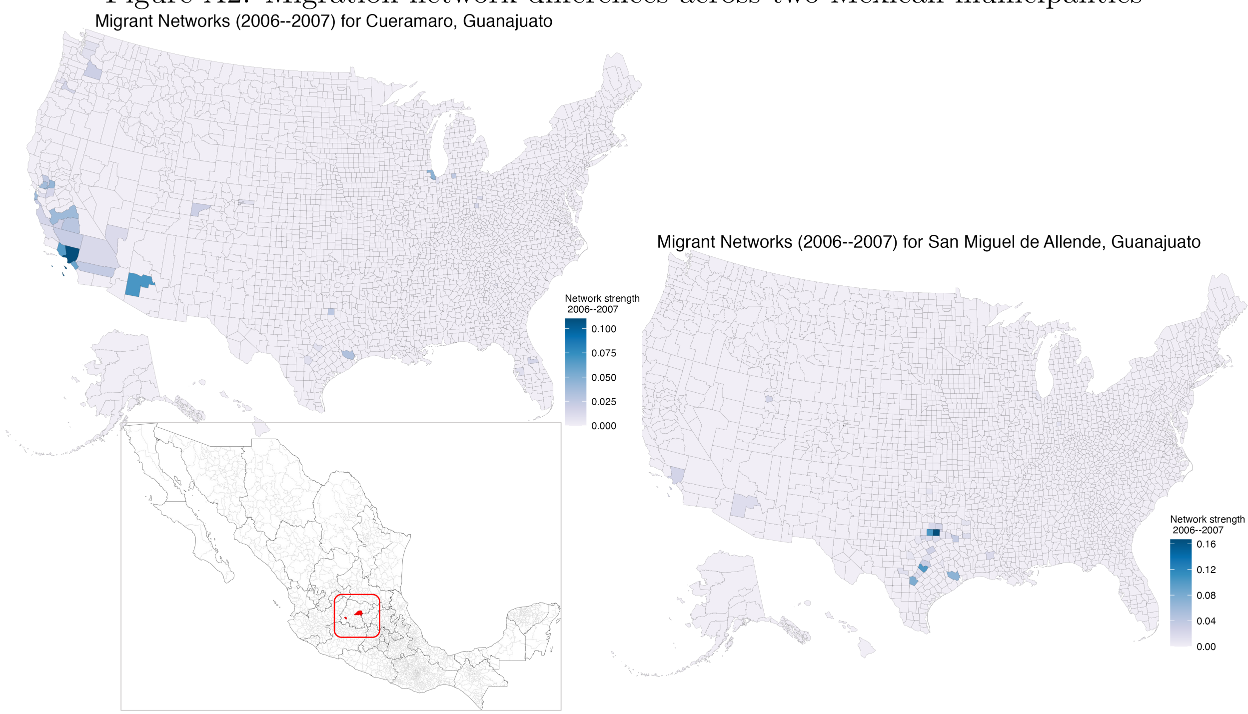

Notes: County-municipality migration network strength measured as the ratio between the number of migrants originated in each Mexican municipality moving to each U.S. county over the total number of migrants in the municipality. Left panel: Migration network for the municipality of Cueramaro, Guanajuato. Right panel: Network for San Miguel de Allende, Guanajuato.

Figure A3: Agricultural labor – National level trends

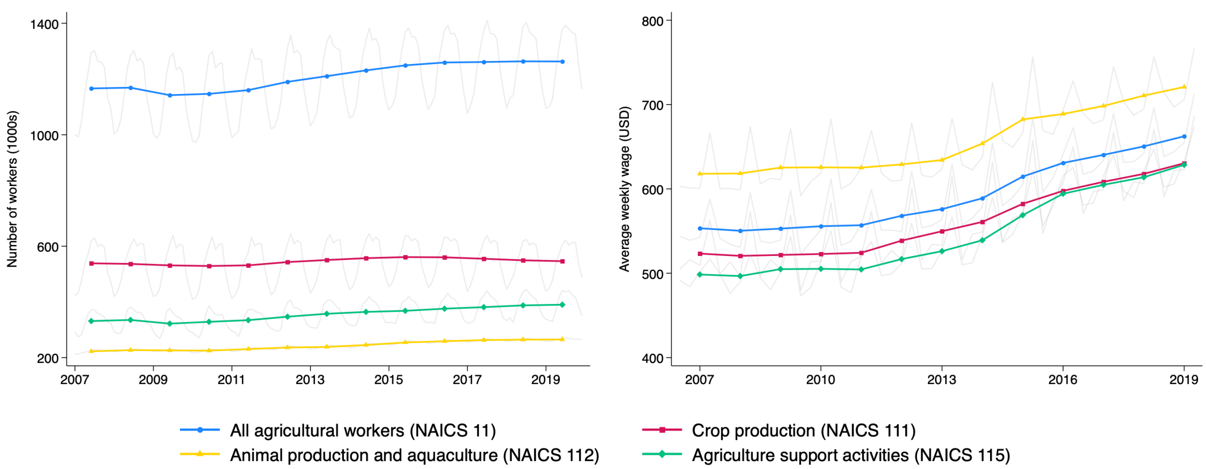

Notes: Data from the QCEW. Gray lines represent quarterly figures; solid color lines yearly averages.

Figure A4: Short run results – Sensitivity to alternative estimation samples and weights

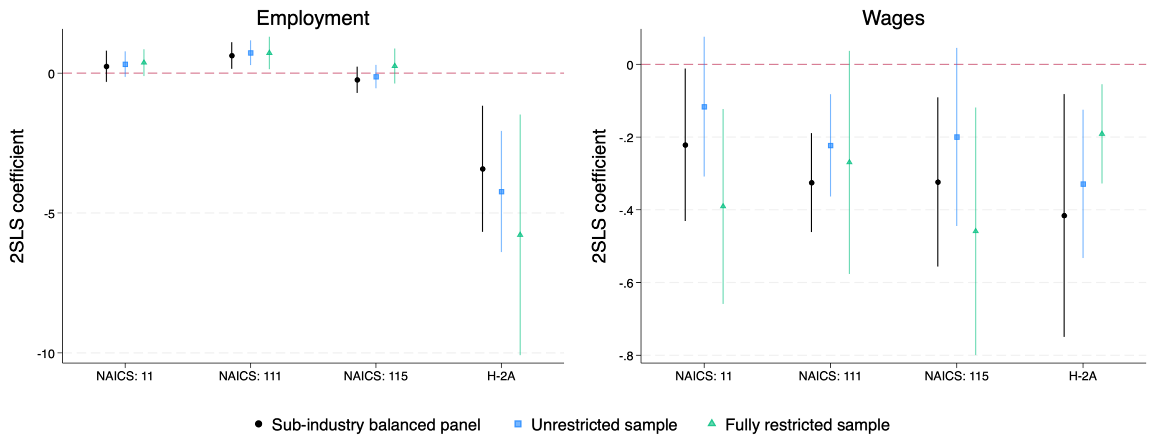

(a) Alternative estimation samples

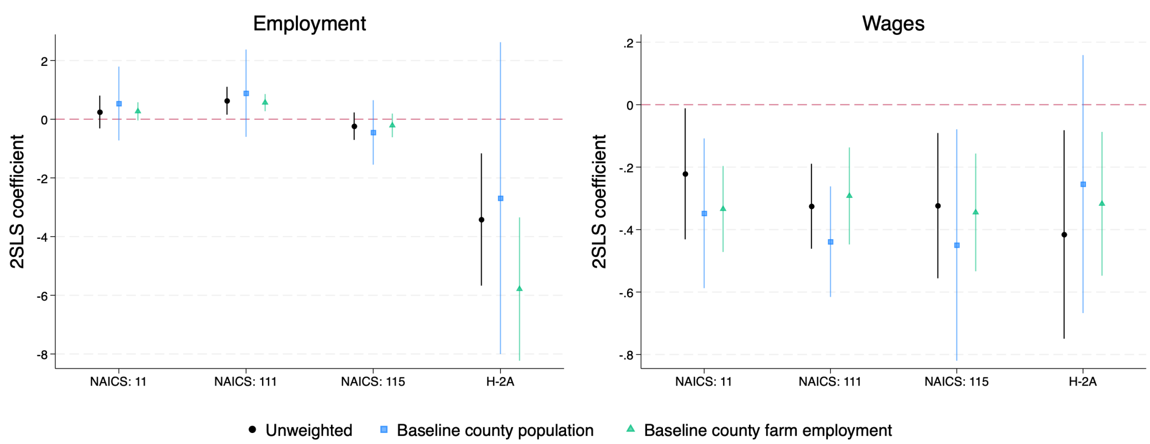

(b) Alternative regression weights

Notes:

Figure A5: Long run results – Sensitivity to alternative estimation samples and weights

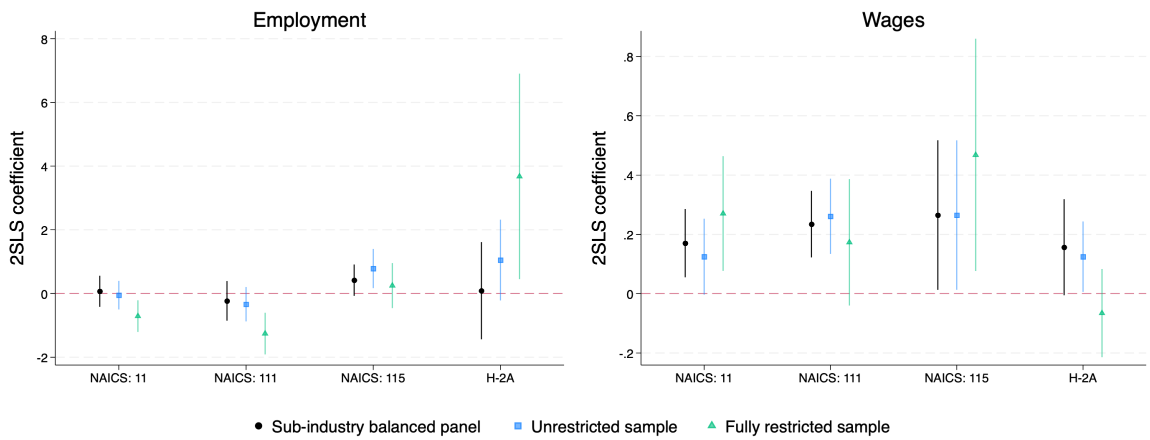

(a) Alternative estimation samples

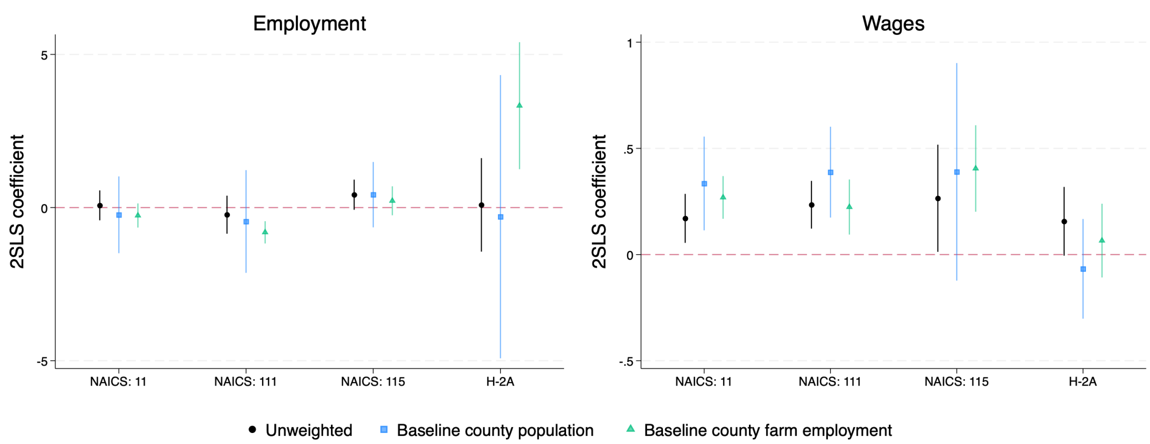

(b) Alternative regression weights

Notes:

Figure A6: H-2A Visa requests by county – 2008-2018

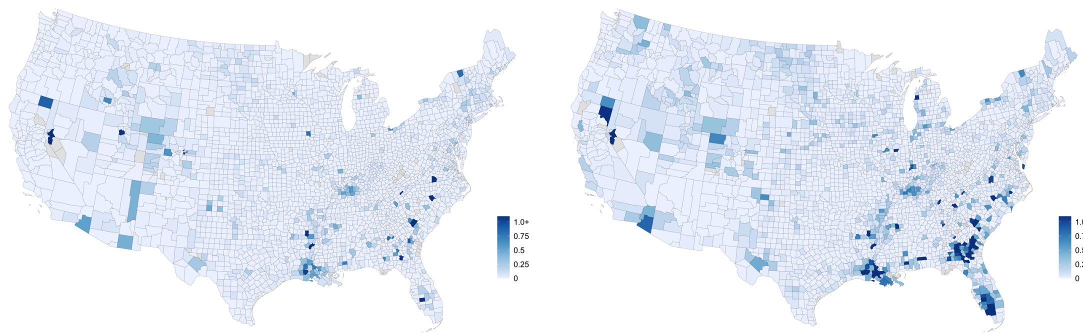

(a) 2008 (b) 2018

Notes: Number of H-2A requests as fraction of total farm employment in county based on the QCEW. Individual H-2A request data from the Department of Labor (*DOL*).

Figure A7: Observed vs. network-predicted migration flows

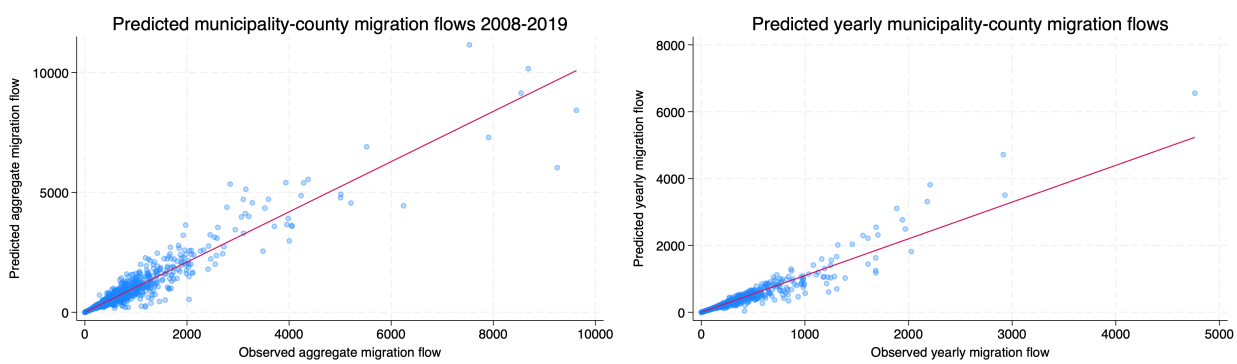

Notes: Left panel: binscatter of aggregate (2008-2019) municipality-county observed migration flows and the predicted flow obtained from multiplying total municipality outflows by the network strength measure $\phi_{o,c}$. Right panel: yearly predicted and observed municipality-county migration flows.

Appendix B Additional First-stage results

To test if the negative relationship between the instrument and migration rates is driven by the 'share' component of the instrument, we define the alternative instrumental variable

$$Z^I_{c,t} = \frac{1}{P_{c,t^0}} \sum_m \left[\text{Homicides}_{m,t} \times \mathbb{1}(\phi^{t^0}_{m,c} > 0) \right]$$

where all (non-zero) origin-destination links are weighted equally. The results of re-estimating regression 4 using $Z^I_{c,t}$ as an instrument are displayed in Table A1 and show that the negative sign is still present once county fixed effects are included.

Table A1: First-stage estimates – Instrumental variable with no network component

	Yearly immigration rate ($m_{c,t}$)			
	(1)	(2)	(3)	(4)
$Z^I_{c,t}$: No migrant network IV	0.004***	0.004***	-0.003***	-0.002***
	(0.000)	(0.000)	(0.001)	(0.000)
Observations	37680	37680	37680	37680
Counties	3140	3140	3140	3140
Year FE	No	Yes	No	Yes
County FE	No	No	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the municipality level in parentheses.

It might also be that the negative relationship observed arises because the year-to-year migration measure is too noisy at such a fine temporal disaggregation. This could happen if for a large number of migrants the arrival date into the U.S. and the decision to get a MCAS card are years apart. While the aggregate trends shown in Figures 1 and A1 do not seem to suggest this, we nonetheless test this by running an alternative version of regression 4, where both migration rates and the instrument are aggregated into three-year bins. Once again, results for this exercise —shown in Table A2— still exhibit the change in sign once county fixed effects are included.

Table A2: First-stage estimates – Three-year migration rate aggregation

	3-Year immigration rate			
	(1)	(2)	(3)	(4)
$Z^{3-year}_{c,t}$: Violence shift-share IV	8.039***	8.252***	-7.871***	-6.489***
	(0.473)	(0.494)	(1.143)	(1.051)
Observations	12568	12568	12568	12568
Counties	3142	3142	3142	3142
Year FE	No	Yes	No	Yes
County FE	No	No	Yes	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the municipality level in parentheses.

We finally show that the inclusion of unit fixed effects renders the relationship between violence and migration inflows negative even when measured at the origin-destination pair level. Given that the MCAS data allows us to observe the magnitude of all migration flows originating in every Mexican municipality headed to each U.S. county, we are able to run the following regression

$$\frac{M_{c,m,t}}{Pop_{m,t^0}} = \beta_0 + \beta_1 \left[\text{Homicides}_{m,t} \times \phi^{t^0}_{m,c}\right] + \delta_t + \gamma_c + \eta_m + \chi_{c,m} + \varepsilon_{c,m,t} \tag{A.1}$$

where $M_{c,m,t}$ is the observed migration from c to m at t, and γ_c, η_m, and $\chi_{c,m}$ are, respectively, county, municipality, and county-by-municipality fixed effects.

Results for regression equation A.1 are shown in Table A3. These results show that while the separate inclusion of either municipality or county fixed effects does not affect the cross-sectional positive relationship between violence and migration, once municipality-by-county fixed effects are included the relationship once again changes sign and becomes negative.

Table A3: First-stage estimates – Origin-destination level regressions

	Yearly Origin-destination immigration rate ($m_{o,m,t}$)					
	(1)	(2)	(3)	(4)	(5)	(6)
Homicides$_{o,t}$ $\times$ $\phi_{o,c}$	0.645***	0.655***	0.731***	0.226**	0.310**	-0.126*
	(0.238)	(0.243)	(0.267)	(0.111)	(0.136)	(0.067)
Observations	5099796	5099796	5099796	5099796	5099796	5099796
Year FE	No	Yes	Yes	Yes	Yes	Yes
Municipality FE	No	No	Yes	No	Yes	No
County FE	No	No	No	Yes	Yes	No
Muni $\times$ County FE	No	No	No	No	No	Yes

Notes: *** p<0.01, ** p<0.05, * p<0.10. Standard errors clustered at the municipality level in parentheses.

Appendix C Violence and Migration - Simulation Exercise

This section describes a simulation exercise that shows how the change in sign of the first-stage regression coefficients shown in Section 3 is consistent with a data generating process where violence-driven origin-destination migration flows are determined by the combination of long-run violence shocks that tend to follow the migrant networks and short-run shocks that follow an independent decision rule.

Setup:

Let the world consist of I origin municipalities indexed by i, J destination counties indexed by j, and T years indexed by t. Each municipality is affected violence that induces outmigration, and migrants select a destination according to some decision rule.

Define:

V_{it}: Violence in municipality i in year t

M_{ijt}: Migration from municipality i to county j in year t.

N_{ij}: Historical migration network. Normalized to sum to 1 within i: $\sum_j N_{ij} = 1, \forall i$.

D_{ij}: Alternate decision rule for destination selection. Also normalized to $\sum_j D_{ij} = 1, \forall i$.

Additionally, for any variable presented without a subscript, let it represent the sum over that subscript. E.g. $V_i \equiv \sum_t V_{it}$.

Yearly violence in a municipality is modeled as a long-run average plus a short-run shock:

$$V_{it} = \bar{V}_i + \tilde{V}_{it}$$

where

$$\bar{V}_i \overset{i.i.d}{\sim} \mathcal{N}(0, Var(\bar{V}_i))$$
$$\tilde{V}_{it} \overset{i.i.d}{\sim} \mathcal{N}(0, Var(\tilde{V}_{it})).$$

Assume outmigration from i to j induced by long-run violence follows the network decision rule N_{ij}, while outmigration caused by short-run violence follows some independent short-run decision rule D_{ij}. That is, migration flows are determined by the following data generating process:

$$M_{ijt} = \beta^L N_{ij} \bar{V}_i + \beta^S D_{ij} \tilde{V}_{it} + \alpha t + \epsilon_{ijt} \tag{A.2}$$

For some normally-distributed mean-zero i.i.d ϵ, and for variables N_{ij}, D_{ij} distributed such that:

$$\begin{pmatrix} log(N_{ij}) \\ log(D_{ij}) \end{pmatrix} \overset{i.i.d}{\sim} \mathcal{N} \begin{pmatrix} 1 & \sigma^2_{dn} \\ \sigma^2_{dn} & 1 \end{pmatrix}$$

where σ^2_{dn} determines the correlation between long-run and short-run decision rules.

Note that, based on (A.2),

$$M_{it} = \sum_j M_{ijt} = \beta^L \bar{V}_i + \beta^S \tilde{V}_{it} + \sum_j \epsilon_{ijt},$$

since $\sum_j N_{ij} = \sum_j D_{ij} = 1$.

The within-i transformation of M_{it} yields:

$$\dot{M}_{it} \equiv M_{it} - \frac{1}{T}\sum_t M_{it} = \beta^S \left[\tilde{V}_{it} - \frac{1}{T}\sum_t \tilde{V}_{it}\right] + \left(\sum_j \epsilon_{ijt} - \sum_t \sum_j \epsilon_{ijt}\right),$$

and so the parameter β^S can thus be recovered through a regression of $\dot{M}_{it}$ on V_{it} or, equivalently, through a municipality fixed effect of the form:

$$\beta^S : \quad M_{it} = \beta V_{it} + \gamma_i + \delta_t + \epsilon_{it} \tag{A.3}$$

Similarly, aggregating across all time periods yields:

$$M_i = \sum_t M_{it} = T\beta^L \bar{V}_i + \beta^S \sum_t \tilde{V}_{it} + \sum_t \sum_j \epsilon_{ijt}$$
$$V_i = \sum_t V_{it} = \sum_t [\bar{V}_i + \tilde{V}_{it}] = T\bar{V}_i + \sum_t \tilde{V}_{it},$$

and thus the parameter β^L can be recovered by a regression of the form:

$$\beta^L : \quad M_i = \beta V_i + \epsilon_i, \tag{A.4}$$

since $E[\tilde{V}_{it}] = 0$, and thus $V_i/T \to \bar{V}_i$ as $T \to \infty$.

Calibration:

Using the expressions above, we calibrate β^S, β^L, and the variance of the violence shocks from the data:

$$\beta^S = 0.283$$
$$\beta^L = 3.274$$
$$Var(\bar{V}_i) = 0.00000007317037$$
$$Var(\tilde{V}_{it}) = 0.00000014502362$$

Simulation:

With the calibrated parameters, we simulate data for an equal number of equally-sized municipalities and counties. Figure A8 shows the resulting distribution of estimating β^S and β^L following equations (A.4) and (A.3) in 100 different data draws and at varying levels of T. The figure confirms that regression (A.4) recovers β_L as T grows.

Figure A8: Distribution of β^L and β^S for 100 regressions - Fixed number of municipalities

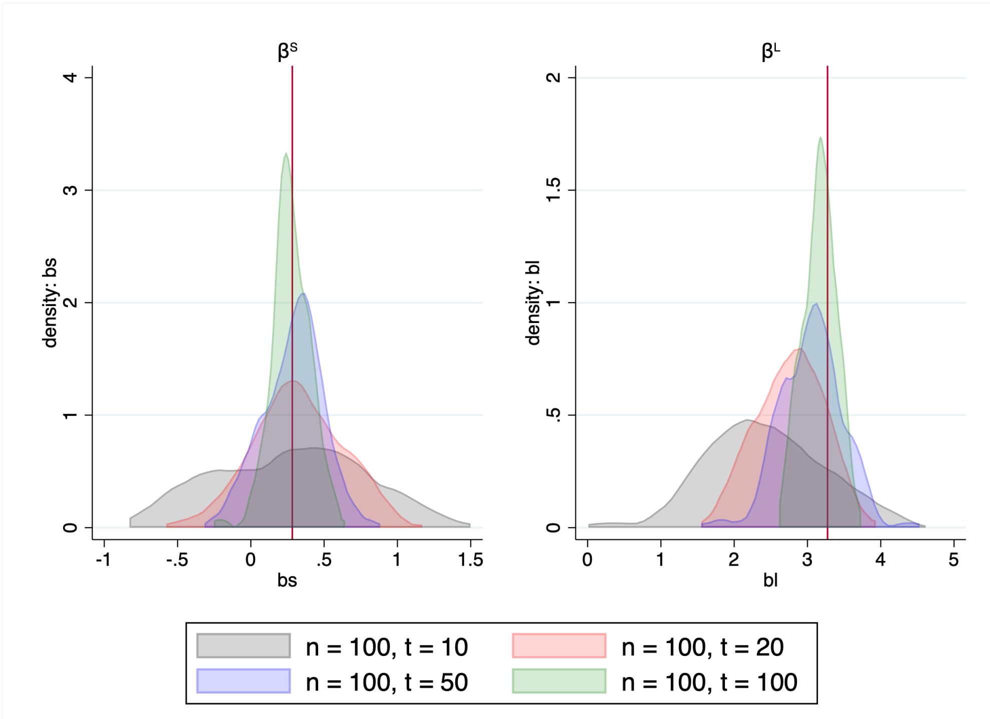

As a check, we estimate the out-migration regressions (A.3) and (A.4) For a single data draw and confirm the estimates are close to those observed in true data. Results for this exercise are shown in Tables A4 and A5.

Table A4: True data: Homicides and migration yearly correlation at the Mexican municipality level

	(1)	(2)	(3)	(4)	(5)
	m_{it}	m_{it}	m_{it}	m_{it}	m_i
Yearly Homicide Rate	0.956***	1.293***	-0.213	0.283**	
	(0.176)	(0.188)	(0.178)	(0.138)	
Aggregate Homicide Rate					3.274***
					(0.457)
Observations	29232	29232	29232	29232	2436
R^2	0.004	0.154	0.000	0.359	0.029
Year FE	No	Yes	No	Yes	No
Municipality FE	No	No	Yes	Yes	No

*** p<0.01, ** p<0.05, * p<0.10

Table A5: Simulated data: Violence and outflows

	(1) M_{it}	(2) M_{it}	(3) M_{it}	(4) M_{it}	(5) M_i
V_{it}	0.693***	0.794***	0.172	0.292***	
	(0.195)	(0.083)	(0.211)	(0.061)	
V_i					3.004***
					(0.147)
Observations	10000	10000	10000	10000	100
R^2	0.001	0.892	0.000	0.895	0.813
Year FE	No	Yes	No	Yes	No
Municipality FE	No	No	Yes	Yes	No

*** p<0.01, ** p<0.05, * p<0.10. $I = 100; J = 100; T = 100$

County-level migration inflows:

Following our empirical strategy, we also aggregate the municipality level simulated outflows to the county level:

$$Z_{jt} \equiv \sum_i (N_{ij} \times V_{it})$$

And, estimate a regression of M_{jt} on Z_{jt}:

$$\begin{aligned} M_{jt} &= \beta \underbrace{\sum_i (N_{ij} \times V_{it})}_{\equiv Z_{jt}} + \gamma_j + \varepsilon_{jt} \qquad \text{(A.5)} \\ &= \beta \sum_i N_{ij} \times (\bar{V}_i + \tilde{V}_{it}) + \gamma_j + \varepsilon_{jt} \\ &= \beta \sum_i N_{ij}\bar{V}_i + \beta \sum_i N_{ij}\tilde{V}_{it} + \gamma_j + \varepsilon_{jt} \end{aligned}$$

Results from estimating regression equation (A.5) on a single draw of simulated data with $\sigma^2_{dn} = 0$ are reported in Table A6. Consistent with the estimations made on real data, the results of this simulation exercise show that even if the correlation between municipality-level violence and outmigration is *defined* to be positive, a negative correlation between the county-aggregated violence variable Z_{jt} and migration inflows can arise. We take this to mean that the underlying relationship between migration and violence driving our results is consistent with a data generating process where violence-driven origin-destination migration flows are determined by the combination of long-run violence shocks that tend to follow the migrant networks and short-run shocks that follow an independent decision rule.

Table A6: Simulated data: Violence and inflows: $Corr(N_{ij}, D_{ij}) = 0$

	(1) M_{jt}	(2) M_{jt}	(3) M_{jt}	(4) M_{jt}	(5) M_j
Z_{jt}	-4.099***	-0.194	-5.122***	-1.011*	
	(0.922)	(0.460)	(1.040)	(0.517)	
Z_j					3.810***
					(1.143)
Observations	10000	10000	10000	10000	100
R^2	0.001	0.895	0.002	0.896	0.099
Year FE	No	Yes	No	Yes	No
Municipality FE	No	No	Yes	Yes	No

*** p<0.01, ** p<0.05, * p<0.10. $I = 100; J = 100; T = 100; \sigma^2_{dn} = 0$.