The Rise of H-2A: Causal Evidence on Guestworker Substitution for Settled Farm Labor

Aakansha Jain*

May 25, 2026

Abstract

U.S. farm labor markets have tightened sharply over the past decade, with rising real wages, persistent harvest-season labor shortages, and a less mobile workforce (Charlton et al., 2019; Fan et al., 2015; Hertz and Zahniser, 2013). Over one-third of U.S. farm workers are unauthorized immigrants, and intensified interior immigration enforcement has reduced the supply of settled non-citizen agricultural workers in affected jurisdictions (Kostandini et al., 2014; Charlton and Kostandini, 2021). Over the same period, certifications under the H-2A agricultural guestworker program, which allows U.S. employers to hire foreign workers for seasonal farm jobs, grew fivefold. This paper estimates how much of the H-2A expansion is causally attributable to the contraction of settled non-citizen farm labor. Using a panel of 740 commuting zones over 2008–2024, I implement a Bartik shift-share instrumental variables design that exploits cross-sectional variation in baseline settled-worker exposure interacted with national enforcement shifts. Our estimates show that for each settled non-citizen agricultural worker lost, farms hire approximately two H-2A guestworkers, accounting for about 48 percent of aggregate H-2A growth over the period.

*PhD candidate, Department of Economics, Iowa State University. Email: aakansha@iastate.edu. I thank my advisors Dr. John Winters and Dr. Peter Orazem for their valuable guidance.

1 Introduction

The U.S. hired farm workforce can be broadly divided into three components: domestic workers (U.S.-born and naturalized citizens), settled non-citizen workers (foreign-born non-citizens with established U.S. residence, predominantly from Mexico), and seasonal guest-workers admitted under the H-2A visa program. The relative sizes of these components have shifted dramatically in recent decades, even as the total has remained nearly constant. Total hired farm employment grew only modestly between 2010 and 2024, from approximately 1.07 million to 1.18 million workers (U.S. Department of Agriculture, Economic Research Service, 2024). Within this stable total, the domestic share has remained stable (TODO- generate graph from ACS), the settled non-citizen workforce peaked in 2015 and has since fallen to roughly 250,000 workers, and H-2A certifications grew fivefold from approximately 75,000 in 2008 to 380,000 by 2024 (Figure 1). Settled non-citizens and H-2A guestworkers crossed around 2020: H-2A workers now outnumber the settled non-citizens whose decline they appear to be replacing. Whether this crossover reflects causal substitution—farms replacing a departing settled workforce with legally sanctioned seasonal labor—or merely coincident national trends is the central question this paper addresses.

The substitution mechanism behind these aggregates is fundamentally between unauthorized labor and legally sanctioned H-2A guestworkers. Over one-third of U.S. farm workers are unauthorized (Charlton, 2025), and intensified immigration enforcement has been shown to reduce the supply of agricultural immigrant labor at the local level (Kostandini et al., 2014; Charlton and Kostandini, 2021). Measuring unauthorized status directly, however, is not possible at the geographic resolution this paper requires. The National Agricultural Workers Survey, the only U.S. source that records self-reported legal status, is not identified below the state level. Residual-method estimates from Pew (Passel and Krogstad, 2024) are national only. The literature on enforcement and agricultural labor has typically used the non-citizen share in the American Community Survey (Kostandini et al., 2014) or broader "likely undocumented" classifications based on demographics (East and Velásquez, 2024), both of which conflate the long-tenured unauthorized workforce with seasonal H-2A arrivals—the very group on the other side of the substitution. I therefore construct a measure I call PANCAL: non-citizens in hired farm labor who did not live abroad in the prior year. The non-citizen filter targets the population most likely to be unauthorized; the settled filter excludes H-2A guestworkers themselves, who would otherwise mechanically appear on both sides of the substitution test.

Understanding why the settled unauthorized workforce exists, and why it is now declining, requires a brief history of U.S. agricultural labor policy. From 1942 to 1964, the Bracero program brought Mexican workers into U.S. fields on short-term seasonal contracts: workers crossed legally, worked for a season, and returned home. After Bracero ended, seasonal labor continued along similar lines but increasingly without legal status: farms relied on a circular flow of unauthorized Mexican workers who crossed informally, worked seasonally, and returned home each year. Massey et al. (2016) show that intensified border enforcement after the 1990s did not reduce the unauthorized population as intended; instead, by raising the cost of crossing, it broke the circular pattern. Workers who had been going home each season began staying year-round, and the settled unauthorized agricultural workforce took its modern shape. Since the mid-2010s, this population has declined as the original cohort ages out, fewer new unauthorized migrants enter agriculture from a demographically changing Mexico (Fan et al., 2015; Charlton et al., 2019), and interior enforcement programs such as Secure Communities and 287(g) partnerships reduce the unauthorized labor supply in affected jurisdictions (Kostandini et al., 2014; Charlton and Kostandini, 2021), causing persistent farm labor shortages (Hertz and Zahniser, 2013; Charlton, 2025).

The H-2A agricultural guestworker program is the primary legal channel through which farms can offset this labor shortfall. Created in 1986 and uncapped, the program allows employers to hire foreign workers for seasonal jobs of up to ten months after demonstrating that domestic workers are unavailable. However, participation is costly: employers must recruit abroad, pay visa and application fees, provide free housing that meets federal standards, supply transportation from the worker's home country, and pay a wage floor set above local market rates. These fixed costs put the program out of reach for smaller farms that need only two or three seasonal workers, concentrating H-2A use among larger operations that can spread costs across many employees (Castillo et al., 2024; Charlton, 2025).

Despite this rapid growth, the question of what is driving it remains open. A growing body of work documents the expansion and geographic concentration of H-2A employment (Castillo et al., 2024; Charlton, 2025) and the labor market consequences of immigration enforcement on agricultural populations (Kostandini et al., 2014; Clemens et al., 2018; East and Velásquez, 2024). Yet no study has established whether the H-2A expansion is *caused* by the contraction of the settled non-citizen workforce, or merely coincides with it. The answer matters for three reasons. First, it speaks to the effectiveness of H-2A as a farm labor policy: the program was designed to provide temporary foreign labor when domestic workers are unavailable, but whether it is in practice serving as the adjustment margin for enforcement-driven labor supply contractions has not been established. Second, the rate of substitution has direct implications for farm production costs, since H-2A workers are substantially more expensive than settled workers. A substitution rate above one-for-one implies that enforcement is not merely reshuffling the same labor through a costlier legal channel but increasing total labor costs per unit of output. Third, the estimate informs the distributional consequences of enforcement across farm sizes. If the substitution is concentrated among large operations that can absorb H-2A's fixed costs, then enforcement may be accelerating consolidation in labor-intensive agriculture, leaving smaller farms without a viable replacement for their settled workforce.

This paper provides estimates bearing on these questions. Using a panel of 740 commuting zones over 2008–2024, I implement a Bartik shift-share instrumental variables design that exploits cross-sectional variation in baseline settled-worker exposure interacted with national enforcement shifts. The main estimates imply that farms hire approximately two H-2A guestworkers for each settled non-citizen agricultural worker lost to enforcement-driven decline.

Applied to the observed 5.2 percentage-point decline in the settled non-citizen workforce between 2008 and 2024, this substitution accounts for roughly 48 percent of aggregate national H-2A growth. Event-study estimates show no differential pre-trends and an H-2A response that lags the workforce decline by two to three years, consistent with the time employers need to navigate H-2A certification and build the housing it requires. The estimates are robust to alternative national shifts—including trending-series benchmarks and the leave-one-out construction—crop-composition controls, and falsification exercises against non-agricultural labor markets and border proximity. This paper contributes to three strands of literature. First, it extends the Bartik shift-share literature (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022) to a setting with two enforcement instruments whose differing temporal profiles permit a formal overidentification test. Second, it adds to the immigration-enforcement literature (Clemens et al., 2018; East and Velásquez, 2024; East et al., 2023; Hanson, 2006; Massey et al., 2016), which has focused on direct effects on deported and displaced workers, by quantifying the induced demand for an alternative legal labor channel. Third, it extends the H-2A program literature. Existing work has documented the program's rapid growth and geographic concentration (Castillo et al., 2024) but has not established what is driving it. This paper estimates that roughly half of aggregate H-2A growth is attributable to enforcement-driven declines in the settled non-citizen workforce.

The remainder of the paper is organized as follows. Section 2 describes the data construction and summary statistics. Section 3 presents the identification strategy and discusses threats to identification. Section 4 reports the main results, event-study estimates, and the aggregate attribution analysis. Section 5 presents the robustness and falsification tests. Section 6 concludes.

2 Data

I construct a commuting-zone-by-year panel covering 740 commuting zones[1] over 2008–2024, excluding 2020 because the Census Bureau did not release a standard 1-year American Community Survey that year due to pandemic-related collection disruptions. The empirical design requires five components: an outcome (H-2A worker certifications), an endogenous variable (the contemporaneous PANCAL share), pre-existing CZ exposure shares for the instrument, a national time-series shift, and crop composition controls. The remainder of this section describes each in turn; full variable definitions and source documentation are in Appendix A (Table 7).

The outcome is the count of H-2A workers certified per commuting zone per year, aggregated from worksite county to commuting zone using DOL OFLC disclosure data. The endogenous variable is the contemporaneous PANCAL share: the number of settled non-citizen agricultural laborers as a share of total agricultural laborers in each commuting zone and year. I define an agricultural laborer as an employed individual aged 16–64 reporting an agricultural industry (IND1990 $\in \{10, 11, 20, 30\}$) and a hired farm-labor occupation (OCC1990 $\in \{479, 484, 488\}$). The occupation restriction excludes farm operators, managers, supervisors, and clerical workers in agricultural firms, isolating the hired labor population that conceptually matches the population surveyed by the NAWS. A worker is classified as a settled non-citizen if they report non-citizen status and did not live abroad in the prior year; the motivation for these filters was given in Section 1. Figure 1 plots the national totals of both series over the panel period.

The shift-share instrument combines a predetermined CZ-level exposure measure with national time-series variation in border enforcement. The exposure measure is the year-2000 PANCAL share, constructed using the same agricultural-laborer and non-citizen filters as the

[1]Commuting zones follow the 1990 Tolbert–Sizer definition (Autor et al., 2013). PUMA- and county-level data are aggregated to CZs using the Dorn (2009) crosswalks for 2000- and 2010-vintage PUMAs, and a custom crosswalk I construct for 2020-vintage PUMAs (used in the 2022–2024 ACS); see Appendix B. CZ 34103 (Valdez–Cordova, AK) is observed in 2008–2021 but absent from 2022 onward following a post-2020 Alaska county reorganization.

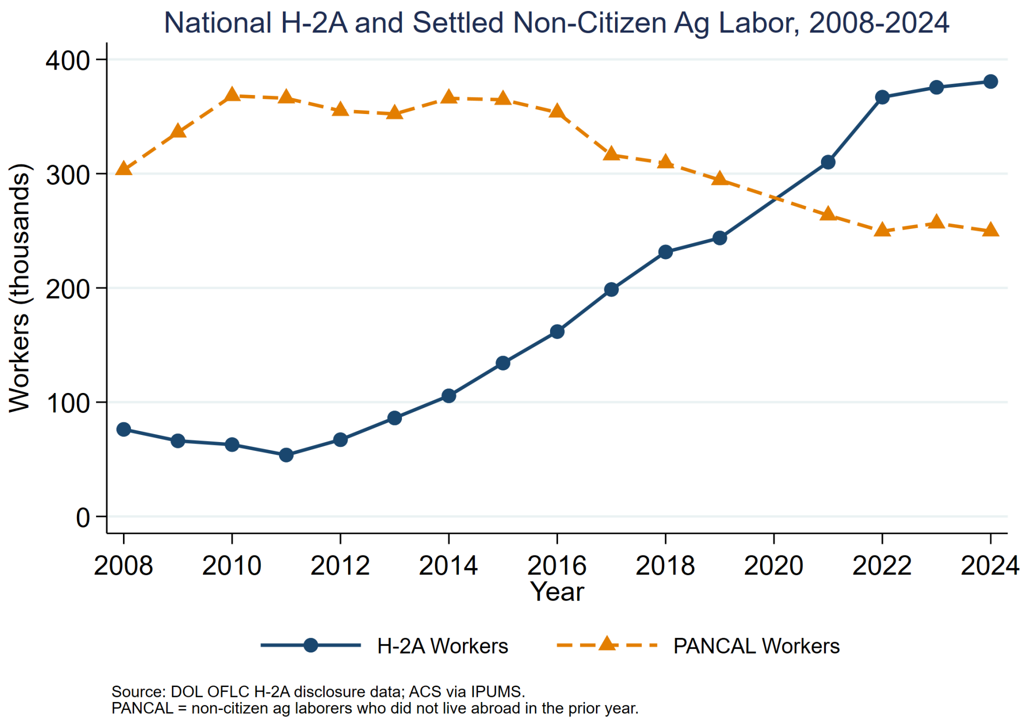

Figure 1: National H-2A certifications and settled non-citizen agricultural labor (PANCAL), 2008–2024. H-2A data from DOL OFLC certification records. PANCAL from the American Community Survey via IPUMS. 2020 is missing because the ACS did not release a standard 1-year file that year.

contemporaneous measure, with a years-in-United-States filter (YRSUSA1 $>$ 1) substituting for the migration filter, which is not consistently available in 2000. Fixing exposure at 2000—eight years before the panel begins—ensures it is predetermined relative to the outcome period. The national shift takes two forms: the log of the real (2024 USD) Border Patrol budget and the log number of US-Mexico border encounters during the agricultural season (March–September). The agricultural-season window is motivated by the seasonal concentration of H-2A demand (Appendix Figure 4). Crop composition controls allow CZs with different baseline crop intensity to follow differential H-2A trajectories; the preferred specification interacts the vegetable share with year dummies.

The final panel contains 11,840 commuting-zone-year observations (740 CZs $\times$ 16 years), of which 11,672 have a non-missing PANCAL share. CZs above the median baseline PANCAL share employ roughly four times as many H-2A workers per year as those below (391 versus 102), have approximately three times the agricultural workforce, a five-fold higher Hispanic population share, a higher vegetable cropland share, and are located closer to the US-Mexico border. These descriptive patterns motivate the vegetable-share controls and the distance-based and non-agricultural falsification tests in Section 5. Full summary statistics are in Appendix A (Table 8).

3 Empirical Strategy

The question this paper asks is whether the decline in settled non-citizen agricultural labor causes farms to hire more H-2A guestworkers. The natural starting point is an OLS regression of H-2A employment on the PANCAL share within commuting zones over time:

$$Y_{cz,t} = \alpha_{cz} + \lambda_t + \delta\, S_{cz,t} + \gamma' X_{cz,t} + \varepsilon_{cz,t}, \tag{1}$$

where $Y_{cz,t}$ is the H-2A worker count in commuting zone cz and year t, $S_{cz,t}$ is the PANCAL share (the number of settled non-citizen agricultural laborers as a share of total agricultural laborers, in percentage points), α_{cz} and λ_t are commuting zone and year fixed effects, and $X_{cz,t}$ is a vector of crop composition controls (vegetable cropland share interacted with year dummies). The parameter of interest is δ: the change in H-2A employment caused by a one percentage-point change in the PANCAL share. If farms substitute toward H-2A workers as their settled workforce contracts, δ should be negative.

OLS estimation of δ is unlikely to recover the causal effect. Simultaneity biases the estimate toward zero: positive local labor demand shocks raise H-2A hiring while also attracting non-citizen workers, compressing the negative substitution coefficient. Measurement error in the PANCAL share, which is constructed from thin ACS cells, attenuates OLS further. Working in the opposite direction, two forces could bias OLS away from zero. First, if high-PANCAL commuting zones also experience independent growth in labor-intensive crop demand that raises H-2A hiring beyond what PANCAL decline alone would generate, OLS would overstate the substitution effect; the vegetable-share $\times$ year controls are designed to absorb this channel. Second, reverse causality: if the availability of H-2A workers causes farms to substitute away from their settled workforce, for instance by shifting to seasonal production models that favor guestworkers over year-round employees, then part of the observed PANCAL decline is a consequence of H-2A expansion rather than a cause of it. The net direction of OLS bias is therefore ambiguous, though the near-zero OLS coefficient reported in Section 4 suggests that attenuating forces dominate in practice. Diagnostics distinguishing the sources of attenuation are also presented there.

To recover the causal effect, I instrument the PANCAL share using a Bartik shift-share design. The instrument interacts each commuting zone's predetermined exposure to settled non-citizen agricultural labor with national time-series variation in border enforcement:

$$Z_{cz,t} = S^0_{cz} \times g_t, \tag{2}$$

where S^0_{cz} is the year-2000 PANCAL share and g_t is a national enforcement measure that varies only over time. The logic is straightforward: commuting zones that had a larger settled non-citizen workforce in 2000 were more exposed when national enforcement subsequently tightened. These high-exposure CZs experienced larger PANCAL declines, generating differential labor shortages that plausibly drove differential H-2A adoption.

In the framework of Goldsmith-Pinkham et al. (2020), identification in a shift-share design comes from the cross-sectional variation in the baseline shares, not from the national shift. Year fixed effects absorb the level of g_t; what remains is the comparison of high-share to low-share commuting zones, reweighted by the temporal pattern of the shift. The identifying assumption is that, conditional on commuting zone and year fixed effects and crop composition controls, commuting zones with higher year-2000 PANCAL shares would have experienced the same changes in H-2A employment as low-share commuting zones in the absence of enforcement-driven PANCAL decline. Because the shares are fixed at 2000, eight years before the panel begins, they are predetermined relative to the outcome period and unaffected by sample-period shocks to either H-2A demand or PANCAL supply.

The national shift g_t nonetheless plays an important role: different shifts reweight the baseline-share comparison along different temporal patterns, and comparing the resulting estimates provides a diagnostic for share-based confounds. I construct two instruments using distinct enforcement measures. The first uses the log of the real (2024 USD) Border Patrol budget:

$$Z^{\text{BP}}_{cz,t} = S^0_{cz} \times \ln(\text{BP budget}^{\text{real}}_t), \tag{3}$$

which captures federal investment in border enforcement capacity. The second uses the log of US-Mexico border encounters during the agricultural season (March–September):

$$Z^{\text{enc}}_{cz,t} = S^0_{cz} \times \ln(\text{Enc}^{\text{ag-season}}_t). \tag{4}$$

The two shifts have very different temporal profiles. The BP budget trends upward over the sample period, producing a monotonic reweighting of the share comparison similar to what other trending national series would generate. Ag-season encounters, by contrast, fall through 2017, surge in 2021–22, and decline by 2024, producing a sharply non-monotonic reweighting that no smoothly trending confound can replicate (Appendix Figure 6).

This distinction matters for the credibility of the design. If a time-varying confound correlated with the baseline shares were driving the results, it would generally produce different second-stage estimates under the BP budget and encounter reweightings, since the two weight years very differently. Finding quantitatively similar estimates from both instruments is consistent with the baseline shares capturing a single underlying channel rather than distinct confounds operating on different time paths. The overidentified specification permits a formal Hansen J test of this logic, reported in Section 4.

The first stage predicts the contemporaneous PANCAL share using the instrument(s):

$$S_{cz,t} = \mu_{cz} + \phi_t + \pi' Z_{cz,t} + \psi' X_{cz,t} + u_{cz,t}, \tag{5}$$

and the second stage estimates the causal effect:

$$Y_{cz,t} = \alpha_{cz} + \lambda_t + \delta\, \hat{S}_{cz,t} + \gamma' X_{cz,t} + \varepsilon_{cz,t}. \tag{6}$$

I report three specifications: just-identified with the BP budget Bartik, just-identified with the ag-season encounter Bartik, and overidentified with both.[2] All specifications are weighted by year-2000 CZ agricultural employment and standard errors are clustered at the CZ level.[3]

Threats to identification

Identification requires that the year-2000 PANCAL shares, conditional on commuting zone and year fixed effects and crop composition controls, are uncorrelated with time-varying determinants of H-2A growth. Several potential violations of this assumption deserve discussion.

Crop demand confounds. The most immediate threat is that high-PANCAL commuting zones are disproportionately concentrated in labor-intensive vegetable and fruit production, and that growing demand for these crops independently drives H-2A adoption. The vegetable-share × year controls in $X_{cz,t}$ absorb this channel directly. Section 5 reports sensitivity to alternative crop controls (orchards, tobacco, and a combined labor-intensive index) and shows that the estimates are stable across specifications.

[2]Both stages are estimated by 2SLS via `ivreghdfe` in Stata

[3]CZ-level clustering allows for arbitrary serial correlation within each CZ and accounts for the fact that the baseline share component of the instrument is constant within CZ. Year fixed effects absorb the common national component of the enforcement shock.

Farm labor contractor networks. H-2A hiring is heavily intermediated by farm labor contractors (FLCs), who account for nearly half of all H-2A certifications (Castillo et al., 2024). If FLC networks are disproportionately established in high-PANCAL commuting zones for historical reasons, these CZs may expand H-2A hiring more readily not because their labor shortages are larger but because the infrastructure to hire guestworkers already exists. Conversely, commuting zones experiencing equivalent PANCAL decline but lacking FLC networks may be unable to access H-2A at all, and the substitution this paper estimates would reflect the joint effect of labor shortages and network availability rather than labor shortages alone. This is a form of heterogeneous treatment effects: the substitution rate may be higher in CZs with established FLC infrastructure. The estimates should therefore be interpreted as the average effect across CZs weighted by agricultural employment, which gives disproportionate weight to exactly the CZs where FLC networks are most developed. Whether the substitution mechanism operates in smaller, less connected labor markets is an open question this paper cannot fully resolve.

Selective migration across commuting zones. If enforcement pushes PANCAL workers from high-share to low-share commuting zones rather than out of the agricultural workforce entirely, the instrument captures geographic reshuffling rather than a true labor supply decline. Commuting zones are relatively large labor markets, which limits the scope for this concern, but some cross-CZ reallocation is likely. To the extent that displaced workers move to other agricultural CZs and take farm jobs there, the estimates capture the spatial redistribution of labor shortages rather than their aggregate magnitude.

Section 5 subjects the identifying assumption to a battery of additional tests: alternative national shifts including trending-series benchmarks (CPI-U, federal health spending, linear time trend) and leave-one-out constructions, crop mix controls, a non-agricultural non-citizen labor falsification, distance-to-border interactions, robustness to instrument timing, and long-difference estimates.

4 Results

Table 1 reports the main estimates. Column (1) gives the OLS coefficient: $\hat{\delta}^{OLS} = -2.1$ (s.e. 1.6, $p = 0.190$), economically small and statistically insignificant. As discussed in Section 3, OLS faces multiple sources of endogeneity, but the near-zero estimate suggests that attenuating forces dominate in practice.

To gauge which source of attenuation binds, I conduct three diagnostics. First, restricting the sample to the quartile of commuting zones with the largest agricultural workforces, where ACS sampling noise is smallest, leaves OLS essentially unchanged at $\hat{\delta} = -1.78$ (SE 2.56). Second, collapsing the panel to three-year averages yields $\hat{\delta} = -5.26$ (SE 4.33). Third, a direct measurement-error diagnostic using lagged PANCAL as an instrument for contemporaneous PANCAL is uninformative ($\hat{\pi} = -0.003$, SE 0.015, $F = 0.05$): within-CZ year-to-year PANCAL variation contains essentially no persistent signal after fixed effects, reflecting the thin-cell ACS structure, so this instrument cannot separate measurement error from endogeneity empirically. The first two diagnostics each move OLS by less than four units, leaving it nearly two orders of magnitude below the IV estimates of −155 to −230 reported below. Classical measurement error in PANCAL alone cannot account for a gap of this size, so endogeneity in the form of simultaneity and reverse causality is the most plausible primary driver, though the diagnostics do not allow a precise decomposition of the two channels.

Columns (2)–(4) of Table 1 report the IV estimates. The just-identified specifications using the BP budget and ag-season encounter Bartiks separately yield $\hat{\delta} = -155.3$ ($p = 0.011$) and $\hat{\delta} = -230.0$ ($p < 0.001$); the overidentified specification combining both instruments yields $\hat{\delta} = -179.1$ ($p = 0.003$), falling between the two just-identified estimates. The Hansen J test fails to reject ($p = 0.128$), consistent with both instruments identifying the same underlying parameter. The weak-instrument-robust Anderson–Rubin test rejects $\hat{\delta} = 0$ at $p < 0.001$ across all IV specifications.

Panel B reports first-stage coefficients. Both instruments enter with the expected negative sign: higher enforcement intensity, interacted with baseline PANCAL exposure, predicts a decline in the contemporaneous PANCAL share. The ag-season encounter Bartik yields a Kleibergen–Paap F-statistic of 32.45 when used alone, comfortably above conventional weak-instrument thresholds; the BP budget Bartik yields $F = 14.82$. In the overidentified specification both instruments are individually significant (BP budget $t = -3.30$, ag-season encounters $t = -4.56$) with a joint $F = 16.31$. Although the raw correlation between the two instruments is near-unity ($\rho = 0.998$) because they share the baseline-share component S^0_{cz}, the within-CZ demeaned correlation is $\rho = 0.28$, confirming that the two national time-series shifts provide distinct identifying variation conditional on fixed effects.

To express these estimates as worker-for-worker substitution rates, I divide $|\hat{\delta}|$ by the regression-weighted mean of total agricultural employment ($\bar{N} = 9{,}047$). The implied substitution rates range from 1.72 (BP budget) to 2.54 (ag-season encounters), with the overidentified estimate at 1.98: for every settled non-citizen agricultural worker lost, farms hire approximately two H-2A guestworkers.

The BP budget Bartik serves as the headline just-identified specification because the budget is set through federal appropriations and is uncorrelated by construction with within-year local agricultural labor demand (Massey et al., 2016). Border encounters reflect both crossing attempts and detection capacity, and so the ag-season encounter Bartik enters as the identification-strengthening complement (Section 5) rather than the headline specification. The overidentified column (4) combines both and yields a substitution rate of 1.98, which is the appropriate single-number summary when both instruments are jointly informative.

Event study

The panel IV specification estimates a single average effect across 2008–2024. To examine the dynamics of the substitution response and test for pre-trends, I estimate reduced-form and first-stage event studies in which the effect of baseline PANCAL exposure is allowed to

Table 1: Main Results: Effect of PANCAL Share on H-2A Employment

	(1) OLS	(2) IV: BP Budget Z^{BP}	(3) IV: Encounters Z^{enc}	(4) IV: Both Overidentified
Panel A: Second Stage				
$\hat{\delta}$ (PANCAL share, pct)	-2.083	-155.346^{**}	-230.023^{***}	-179.127^{***}
	(1.588)	(61.148)	(64.729)	(60.710)
Substitution rate	0.02	1.72	2.54	1.98
Panel B: First Stage				
BP budget Bartik		-0.356^{***}		-0.304^{***}
		(0.093)		(0.092)
Ag-season encounter Bartik			-0.069^{***}	-0.047^{***}
			(0.012)	(0.010)
KP F-statistic	—	14.82	32.45	16.31
Hansen J (p-value)	—	—	—	0.128
Panel C: Specification				
CZ & year FE	Yes	Yes	Yes	Yes
Veg × year controls	Yes	Yes	Yes	Yes
Observations	11,672	11,672	11,672	11,672
CZ clusters	740	740	740	740

Notes: All specifications weighted by year-2000 CZ agricultural employment. Standard errors clustered at the CZ level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The substitution rate equals $|\hat{\delta}|/\bar{N}$, where $\bar{N} = 9{,}047$ is the regression-weighted mean of total agricultural employment.

vary by year:

$$Y_{cz,t} = \alpha_{cz} + \lambda_t + \textstyle\sum_{k \neq 2008} \beta_k^{RF} \big(S_{cz}^0 \times \mathbf{1}[t = k]\big) + \gamma' X_{cz,t} + \varepsilon_{cz,t}, \tag{7}$$

$$S_{cz,t} = \mu_{cz} + \phi_t + \textstyle\sum_{k \neq 2008} \beta_k^{FS} \big(S_{cz}^0 \times \mathbf{1}[t = k]\big) + \psi' X_{cz,t} + u_{cz,t}, \tag{8}$$

where $\hat{\beta}_k^{RF}$ is the additional H-2A workers gained in year k relative to 2008 per unit of baseline PANCAL exposure, and $\hat{\beta}_k^{FS}$ is the corresponding change in the contemporaneous PANCAL share.

Figures 2 and 3 plot the two coefficient sequences. Pre-period coefficients are small and statistically insignificant in both specifications, confirming no differential pre-trends in either H-2A growth or PANCAL decline across high- and low-exposure commuting zones.

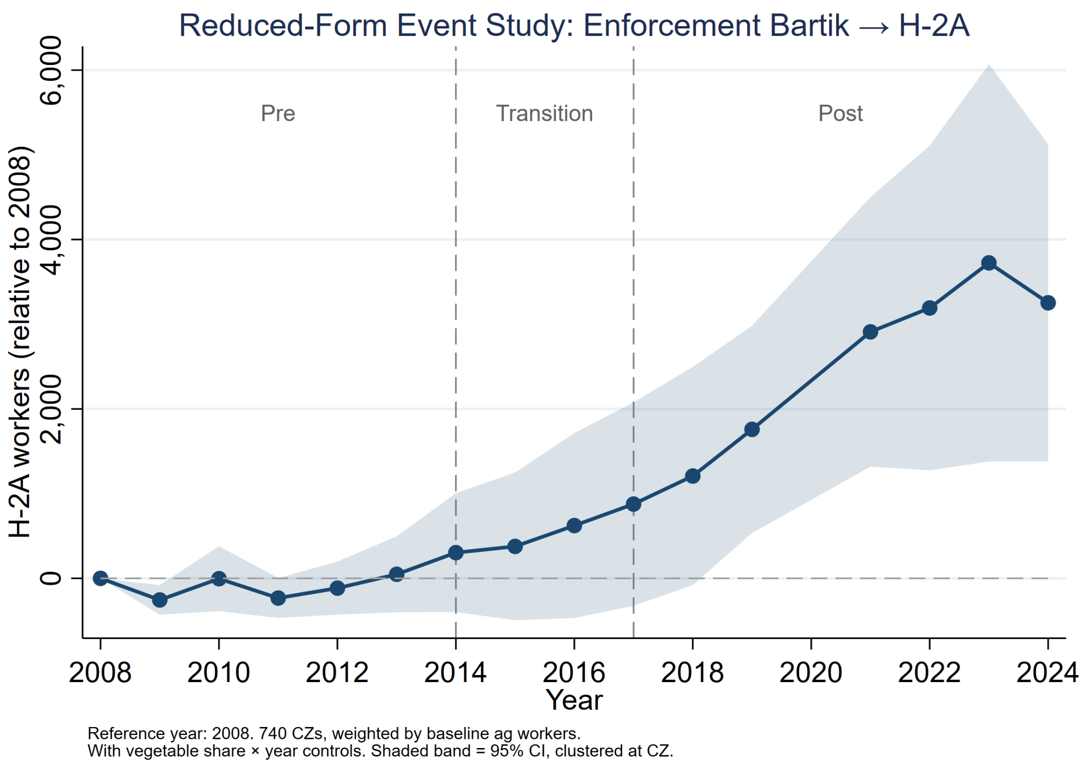

Figure 2: Reduced-form event study: enforcement Bartik → H-2A workers. Each point plots $\hat{\beta}_k^{RF}$ from equation (7). Reference year: 2008. 740 CZs, weighted by baseline agricultural employment. Shaded band: 95% CI, clustered at CZ. With vegetable share × year controls.

The reduced-form H-2A response emerges around 2014–2015, becomes significant at the 5% level from 2017 onward, and grows monotonically through 2023. The first-stage PANCAL response turns negative by 2012 and grows in magnitude through 2023. The H-2A response lags the PANCAL decline by roughly two to three years, consistent with the time employers need to learn the H-2A program, establish contractor relationships, and build the housing required to host guestworkers.

Aggregate attribution

The substitution rate summarizes the CZ-level causal effect, but the policy question is how much of aggregate H-2A growth over 2008–2024 is attributable to PANCAL decline. Using regression-weighted means, H-2A workers per CZ grew from 246 in 2008 to 2,203 in 2024 ($\Delta =$

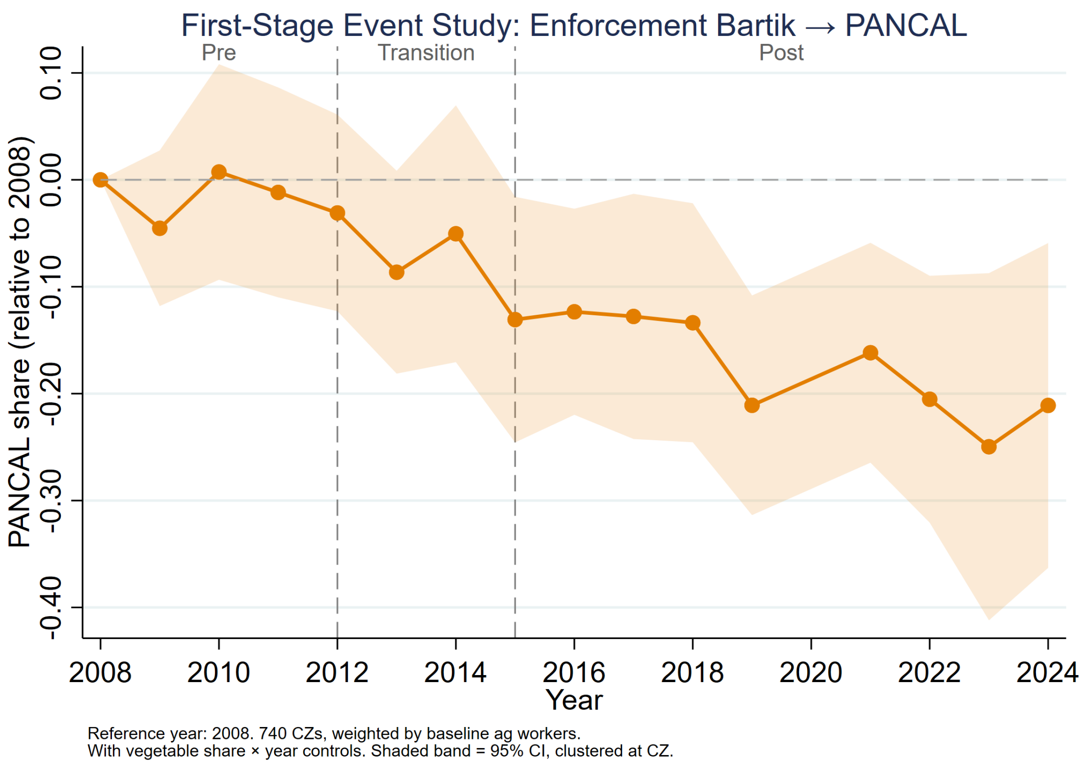

Figure 3: First-stage event study: enforcement Bartik → PANCAL share. Each point plots $\hat{\beta}_k^{FS}$ from equation (8). Reference year: 2008. 740 CZs, weighted by baseline agricultural employment. Shaded band: 95% CI, clustered at CZ. With vegetable share × year controls.

1,957), while the PANCAL share declined by 5.2 percentage points (from 38.7% to 33.5%). Multiplying $|\hat{\delta}|$ by the PANCAL decline yields the predicted H-2A growth attributable to PANCAL loss (Table 2). PANCAL decline explains 41% of national H-2A growth under the BP budget instrument, 61% under ag-season encounters, and 48% in the overidentified specification.

5 Robustness

This section subjects the main estimates to robustness and falsification exercises. The first three assess instrument validity: trending-series benchmarks that test whether the ag-season encounter shift carries identifying content beyond a generic upward trend, leave-one-out and alternative shift constructions that test whether identification depends on the specific en-

Table 2: Share of H-2A Growth Explained by PANCAL Decline

Instrument	$\lvert\hat{\delta}\rvert$	Predicted ΔH-2A	% Explained
BP budget	155.3	807	41.2%
Ag-season encounters	230.0	1,196	61.1%
Overidentified	179.1	931	47.6%

Notes: Predicted ΔH-2A = $\lvert\hat{\delta}\rvert \times 5.2$. % Explained = Predicted / 1,957 (actual regression-weighted H-2A growth per CZ, 2008–2024).

forcement measure, and lag specifications that address reverse causality concerns. The next two address exclusion-restriction threats: crop mix controls for labor-intensive production trends and a non-agricultural non-citizen labor falsification. The final two provide additional identification support: a distance-to-border falsification distinguishing agricultural exposure from geographic proximity, and long-difference estimates confirming the result at lower frequencies.

Trending-series benchmarks and the identifying content of the ag-season encounter shift. U.S. immigration enforcement intensified broadly over 2008–2024 along both interior and border-enforcement margins, and the BP budget series, despite some non-monotonic features Appendix **??**), shares this upward trajectory with many unrelated national series. The first concern is therefore mechanical: when the baseline-share comparison is reweighted by any upward-trending shift, does the design simply read off the same cross-sectional pattern regardless of the shift's economic content? If so, the substitution estimates would be uninformative about enforcement specifically. To test this, I replace the BP budget with three trending series that have no plausible causal link to settled non-citizen agricultural labor—the Consumer Price Index, log real federal health spending, and a simple linear time trend—and re-estimate the IV.

Table 3 reports the results. All three benchmarks yield second-stage estimates similar to the BP budget (substitution rates of 1.86–2.04 versus 1.72), with first-stage F-statistics between 19.00 and 22.36 and reduced-form t-statistics tight at 3.44–3.66 (Appendix Table 9).

Identification in a Bartik design comes from the cross-sectional baseline-share variation; a monotonically upward shift simply reweights years along that comparison, and the underlying enforcement intensity this design is meant to capture has in fact trended upward over the sample period. Multiple trending shifts recovering similar substitution rates is therefore consistent with the baseline shares capturing genuine differential exposure to enforcement.

The more informative result is what happens when the shift is *not* monotonic. The ag-season encounter instrument exhibits a sharply non-monotonic time path—a roughly 9-fold swing from 251,000 encounters in 2017 to 2.3 million in 2022—and is constructed from a window (March–September) timed to the seasonal concentration of H-2A demand. If the baseline shares were proxying for a generic upward trend or other slow-moving regional confound, this non-monotonic, ag-specific reweighting should weaken the first stage relative to monotonic benchmarks. The opposite occurs: the ag-season encounter Bartik delivers a Kleibergen–Paap F-statistic of 32.45, materially above the 19.00–22.36 range produced by the trending benchmarks and above the 14.82 from the BP budget itself. The high-frequency, ag-season-timed component of border enforcement carries identifying power for the contemporaneous PANCAL share that no smoothly trending series can replicate, which is direct evidence that the baseline shares are picking up exposure to enforcement intensity timed to the agricultural season rather than a generic trending confound.

This logic also rationalizes the overidentified specification in Section 4. The BP budget and ag-season encounter Bartiks reweight the baseline-share comparison along very different temporal patterns, and a Hansen J test of whether they identify the same parameter fails to reject ($p = 0.128$). Combining a monotonic enforcement-capacity measure with a non-monotonic, ag-season-timed flow measure—and finding that they overidentify—uses each instrument for the identifying variation it supplies best.

The trending-series benchmarks show that the IV estimate is stable across upward-trending shifts and that the ag-season encounter shift, with its non-monotonic and ag-season-timed time path, supplies identifying content that monotonic series do not. The leave-one-out

Table 3: Trending-Series Benchmarks: IV Estimates Under Alternative Monotonic Shifts

	(1) BP Budget	(2) CPI Spending	(3) Federal Health Spending	(4) Linear Trend
Panel A: Second Stage				
$\hat{\delta}$ (PANCAL share, pct)	-155.3^{**}	-184.3^{***}	-176.6^{***}	-168.5^{***}
	(61.1)	(63.7)	(59.4)	(57.5)
Substitution rate	1.72	2.04	1.95	1.86
KP F-statistic	14.82	19.00	22.35	22.36
Panel B: Specification				
Veg × year controls	Yes	Yes	Yes	Yes
CZ & year FE	Yes	Yes	Yes	Yes
Observations	11,672	11,672	11,672	11,672

Notes: All specifications weighted by year-2000 CZ agricultural employment and use vegetable share × year interactions as controls. Standard errors clustered at the CZ level in parentheses. Each alternative shift is constructed by replacing the BP budget in the standard Bartik construction: column (2) uses CPI-U, column (3) uses log real federal health spending, column (4) uses a simple linear time trend ($year - 2008$). Columns (2)–(4) serve as monotonic-trend benchmarks: stability of the second-stage estimate across these benchmarks shows that the design is robust to the choice of upward-trending shift, while the higher first-stage F of the ag-season encounter Bartik (Table 1) demonstrates that its non-monotonic, ag-season-timed shift adds identifying power beyond what any monotonic series provides. The substitution rate equals $|\hat{\delta}|/\bar{N}$, where $\bar{N} = 9{,}047$ is the regression-weighted mean of total agricultural employment. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

exercise goes further by replacing the national shift entirely with the national PANCAL share itself, testing whether the baseline-share comparison identifies the substitution effect without relying on any external national series.

Leave-one-out and alternative shift construction. Rather than using an external enforcement measure as the shift, I construct two instruments that use the national PANCAL share directly: a simple version using the full national share ($Z^{\text{alt}}_{cz,t} = S^0_{cz} \times \bar{S}_t$) and a leave-one-out version that excludes CZ's own contribution to the national share ($Z^{\text{LOO}}_{cz,t} = S^0_{cz} \times \bar{S}_{-cz,t}$). The LOO construction rules out any mechanical bias from own-CZ variation in the aggregate, following standard shift-share practice (Borusyak et al., 2022).

Both alternatives yield results consistent with the main estimates. The full-national-share instrument produces $\hat{\delta} = -180.7$ ($p = 0.003$, KP $F = 24.4$), essentially identical to the overidentified estimate of -179.1 from the main specification. The LOO instrument produces $\hat{\delta} = -223.8$ ($p = 0.008$, KP $F = 16.1$), implying a substitution rate of 2.47—larger than the benchmark BP budget estimate and close to the ag-season encounter result (-230.0). Because neither alternative uses an enforcement-specific shift, these results provide evidence that the baseline-share comparison captures real differential exposure to PANCAL decline, independent of how national enforcement intensity is measured. Full results are in Appendix Table 10.

The placebo and LOO tests address the instrument; the next two exercises address the exclusion restriction directly.

Crop mix controls. As discussed in Section 3, the most plausible exclusion-restriction threat is that high-PANCAL CZs are disproportionately concentrated in labor-intensive vegetable production, and that vegetable demand growth could independently drive H-2A adoption. To address this, I construct crop-share × year interactions from the 2002 Census of Agriculture and test three specifications: a combined labor-intensive index (vegetables + orchards + berries), separate controls for vegetables, orchards, and tobacco (45 parameters), and vegetable-only controls (15 parameters).

Table 4 reports the main comparison. Without crop controls, the IV implies a substitution rate of 4.55. Adding vegetable share × year controls reduces the second-stage coefficient by approximately 60%, while the first-stage F-statistic rises from 11.64 to 14.82, indicating that the vegetable controls absorb outcome variation correlated with the instrument without weakening the instrument's predictive power for PANCAL. Vegetable-only controls are sufficient: orchard and tobacco share × year interactions are statistically insignificant in nearly every year (Appendix Table 13), and including them reduces the first-stage F-statistic to 6.85 without changing the second-stage estimate substantively (Appendix Table 11). The substantial coefficient reduction when vegetable controls are added confirms that crop demand growth is an empirically meaningful confound—exactly what the controls are designed

Table 4: Crop Mix Controls

	(1) OLS	(2) IV: No controls	(3) IV: Veg × year
$\hat{\delta}$ (PANCAL share, pct)	−2.08	−411.6**	−155.3**
	(1.59)	(177.4)	(61.1)
Substitution rate	0.02	4.55	1.72
KP F-statistic	—	11.64	14.82
Anderson-Rubin p-value	—	< 0.001	< 0.001
CZ & year FE	Yes	Yes	Yes
Crop controls	None	None	Veg × year
Observations	11,672	11,672	11,672

Notes: All specifications weighted by year-2000 CZ agricultural employment. Standard errors clustered at the CZ level in parentheses. The substitution rate equals $|\hat{\delta}|/\bar{N}$, where $\bar{N} = 9{,}047$ is the regression-weighted mean of total agricultural employment. Alternative crop control specifications (labor-intensive index × year; veg + orchard + tobacco × year) and sensitivity to outcome scaling and weighting are reported in Appendix Tables 11 and 12. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

to absorb.

Non-agricultural non-citizen labor. A further concern is that the baseline PANCAL share may be proxying for general immigrant concentration rather than agricultural-specific labor market exposure. High-PANCAL commuting zones also have large non-agricultural non-citizen populations, and if general immigrant concentration correlates with H-2A growth for reasons unrelated to agricultural labor shortages — for instance, because these CZs tend to have established farm labor contractor networks or proximity to recruitment infrastructure — the instrument would capture this broader relationship rather than the substitution mechanism. This concern is mitigated by the fact that agriculture is far more dependent on non-citizen labor than other sectors: over a third of crop workers are unauthorized (Charlton, 2025), compared with roughly 5 percent of the total U.S. workforce (Passel and Krogstad, 2024). The PANCAL share measures sector-specific labor intensity, not general population composition. Nevertheless, I test the distinction directly.

I construct a non-agricultural analog of PANCAL using parallel filters: settled non-citizen workers (CITIZEN = 3, MIGRATE1 ≠ abroad) employed in non-agricultural industries, restricted to working age and low education ($\leq$ high school) to exclude H-1B and other professional visa holders. The cross-sectional correlation between the year-2000 baseline shares of PANCAL and the non-agricultural settled non-citizen share is 0.70; the within-CZ correlation between the time-varying shares is 0.48. The two measures share substantial geographic overlap — both concentrated near the US-Mexico border, in California, and in Florida — but exhibit independent within-CZ variation over time.

Table 5 reports two tests. Column (2) adds the time-varying non-agricultural settled non-citizen share as a control in the second stage: the PANCAL coefficient is -186.4 (implied substitution rate of 2.06), compared with -155.3 in the benchmark, and the coefficient on the non-agricultural share itself is imprecise ($+82.3$, $t = 1.58$). The agricultural channel survives the inclusion of a parallel non-agricultural labor market control. Column (3) reports a reduced-form horse race in which both the agricultural Bartik and a parallel non-agricultural Bartik — constructed by interacting the year-2000 non-agricultural settled non-citizen share with the BP budget series — are entered jointly; the reported coefficients are reduced-form effects on H-2A worker counts, not substitution rates. The agricultural Bartik is statistically significant (75.2, $t = 2.85$); the non-agricultural Bartik flips sign and becomes insignificant (-89.0, $t = -1.09$, $p = 0.27$). Because both Bartiks share the same shift (the BP budget) and the underlying baseline shares have a cross-sectional correlation of 0.70, the two reduced-form coefficients are identified from residualized share variation when entered jointly; the negative point estimate on the non-agricultural Bartik should therefore be read as statistically indistinguishable from zero rather than as an economic claim that non-agricultural immigrant concentration reduces H-2A demand. The relevant content of the horse race is that the agricultural Bartik retains its predictive power after partialling out the non-agricultural component, while the non-agricultural Bartik does not. The baseline PANCAL share captures something specific to agricultural labor markets, not general immigrant concentration.

Table 5: Non-Agricultural Non-Citizen Labor: Control and Falsification

	(1) Benchmark IV	(2) IV + non-ag control	(3) Horse race (RF)
Second-stage / reduced-form coefficient on agricultural exposure			
$\hat{\delta}$ (PANCAL share, pct)	-155.3^{**}	-186.4^{**}	—
	(61.1)	(81.3)	—
Substitution rate	1.72	2.06	—
Ag Bartik (RF coefficient)	—	—	75.2^{***}
	—	—	(26.4)
Non-ag Bartik (RF coefficient)	—	—	−89.0
	—	—	(81.4)
Additional controls			
Non-ag settled non-citizen share (pct)	—	82.3	—
	—	(52.1)	—
KP F-statistic	14.82	12.65	—
Veg × year controls	Yes	Yes	Yes
CZ & year FE	Yes	Yes	Yes
Observations	11,672	11,672	11,840

Notes: Columns (1) and (2) report 2SLS estimates with the BP budget Bartik instrumenting for the PANCAL share. Column (3) reports an OLS reduced form with both Bartiks entered jointly. All specifications are weighted by year-2000 CZ agricultural employment, include vegetable share × year interactions, CZ fixed effects, and year fixed effects. Standard errors clustered at the CZ level in parentheses. The substitution rate equals $|\hat{\delta}|/\bar{N}$, where $\bar{N} = 9{,}047$ is the regression-weighted mean of total agricultural employment. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Distance to the border. CZs with high baseline PANCAL shares tend to be located closer to the US-Mexico border: the cross-sectional correlation between the year-2000 PANCAL share and log distance to the nearest border point is -0.46. Distance itself is time-invariant and therefore absorbed by CZ fixed effects, so it cannot be a direct confound in the panel specification. The indirect threat is that the baseline PANCAL share may be proxying for border proximity rather than agricultural labor market exposure. To test this, I construct a parallel distance-based Bartik instrument that replaces the year-2000 PANCAL share with negative log distance from each CZ centroid to the nearest US-Mexico border point (so that higher values indicate greater proximity, matching the sign convention of the PANCAL share):

$$Z_{cz,t}^{\text{dist}} = -\ln(d_{cz}) \times \ln(\text{BP budget}_t^{\text{real}}).$$

Table 6 reports the key diagnostics; full results are in Appendix Table 14. The distance Bartik fails as an instrument: although its first-stage coefficient is statistically significant (-0.055, $p = 0.031$), its Kleibergen–Paap F-statistic is 4.67, well below conventional weak-instrument thresholds, and its reduced-form effect on H-2A is small and imprecisely estimated ($\hat{\beta} = 1.74$, $t = 0.48$). The Anderson–Rubin test, which is robust to weak instruments, also fails to reject ($p = 0.63$). When both Bartiks are entered jointly in a reduced form, the PANCAL Bartik retains its predictive power (72.9, $t = 3.06$) while the distance Bartik remains insignificant (-7.9, $t = -1.37$). The baseline PANCAL share is not simply proxying for border proximity.

Table 6: Distance Falsification: Key Diagnostics

	Distance Bartik alone	Horse race (RF)
First-stage coefficient	-0.055^{**} ($p = 0.031$)	—
First-stage F	4.67	—
PANCAL Bartik (RF coef.)	—	72.9^{***} ($t = 3.06$)
Distance Bartik (RF coef.)	1.74 ($t = 0.48$)	-7.9 ($t = -1.37$)

Notes: Compact summary of distance-based falsification. Full results in Appendix Table 14. Specifications include CZ and year fixed effects and vegetable share × year controls; standard errors clustered at the CZ level. ** $p < 0.05$, *** $p < 0.01$.

Robustness to instrument timing. The main specification interacts baseline PANCAL shares with contemporaneous enforcement measures. But enforcement may affect the settled non-citizen workforce with a delay: a budget increase or surge in border activity in year t may not fully translate into local PANCAL decline until year $t + 1$. To test sensitivity to this timing assumption, I re-estimate the benchmark specifications using one-year-lagged Bartik instruments (Appendix Table 15). The lagged-BP-budget estimate is $\hat{\delta} = -154.5$ (SE 54.7), and the lagged-encounter estimate is $\hat{\delta} = -233.2$ (SE 96.1, Anderson–Rubin $p < 0.001$); both are statistically indistinguishable from their contemporaneous counterparts

(-155.3 and -230.0, respectively). The first-stage F-statistic for the BP budget instrument rises from 14.82 to 23.97 under the lag, consistent with PANCAL responding to structural enforcement capacity with a delay; the encounter instrument shows the opposite pattern, consistent with encounters capturing contemporaneous border-crossing flows whose effect on local labor supply is more immediate.[4]

Long-difference robustness. A 5-year first-difference IV specification, which removes commuting zone fixed effects mechanically and identifies $\hat{\delta}$ from long-run within-CZ variation rather than year-to-year changes, yields $\hat{\delta} = -132.2$ (SE 64.3, $p = 0.040$) using the BP budget instrument and $\hat{\delta} = -330.9$ (SE 141.8, $p = 0.020$) using ag-season encounters (Appendix Table 16). The BP budget estimate is close to the panel benchmark of -155.3; the encounter estimate is larger but imprecisely estimated, with a confidence interval that comfortably includes the panel benchmark. Both remain statistically significant, confirming that the substitution effect is not an artifact of high-frequency year-to-year variation. Three-year first-difference specifications are not informative because the 2020 ACS gap mechanically eliminates encounter-instrument variation in 3-year windows that span 2020.

6 Conclusion

This paper estimates that farms hire approximately two H-2A guestworkers for each settled non-citizen agricultural worker lost to enforcement-driven decline, accounting for roughly half of aggregate H-2A growth between 2008 and 2024. The result survives alternative national shifts, including trending-series benchmarks and leave-one-out constructions, as well as crop-composition controls and falsification exercises against both non-agricultural labor markets and border proximity. Event-study estimates show no differential pre-trends and an H-2A response that lags the workforce decline by two to three years, consistent with the time

[4] Two minor sample changes accompany the lag: 2008 and 2021 observations are dropped because the lagged instrument requires a non-missing prior year, and two vegetable $\times$ year interactions are absorbed by the year fixed effects in the restricted sample.

employers need to navigate H-2A certification and build the required housing.

The estimates capture substitution toward H-2A specifically and do not reflect other margins through which farms may respond to labor shortages, such as mechanization, crop switching, or reduced production. The total farm-level adjustment to enforcement-driven labor loss is likely larger than the H-2A channel this paper measures. The estimates are also weighted by agricultural employment, giving disproportionate influence to commuting zones where farm labor contractor networks are most developed and H-2A access is most established. Whether the substitution mechanism operates at comparable rates in smaller, less connected labor markets remains an open question.

7 Ongoing and Planned Work

Several extensions are in progress or planned for the next phase of this project. On the identification side, the Bartik design relies on a single cross-sectional exposure measure (the year-2000 PANCAL share); alternative share constructions—such as the share of cropland in labor-intensive categories or historical deportation rates by commuting zone—could provide independent tests of the baseline-share exogeneity assumption. Formal weak-instrument diagnostics beyond the Kleibergen–Paap F-statistic, such as the effective F-statistic of Montiel Olea and Pflueger (2013), would further discipline inference in the just-identified specifications.

The heterogeneity question raised by the FLC network concern could be addressed directly by estimating the substitution rate separately for commuting zones above and below the median agricultural workforce, testing whether smaller CZs that lack H-2A infrastructure respond through different margins—downsizing, crop switching, or reliance on labor contractors based elsewhere—rather than through direct H-2A adoption.

On the outcome side, examining agricultural wage dynamics as a function of enforcement exposure would provide supporting evidence for the labor shortage mechanism: if enforcement creates genuine shortages, wages should rise in high-PANCAL commuting zones during the transition period before H-2A fully substitutes for the settled workforce, and potentially stabilize once H-2A adoption ramps up. The relationship between PANCAL decline and domestic agricultural employment also warrants investigation, both as a control for domestic-worker crowd-out and as an outcome in its own right.

Finally, a sectoral placebo using dairy farming—an industry with a high non-citizen workforce share but limited access to H-2A due to its year-round labor requirements—could sharpen the causal interpretation. If enforcement-driven PANCAL decline affects dairy and crop agriculture equally (a common first stage) but only crop agriculture responds with H-2A growth, this would confirm that the substitution is specific to sectors where H-2A is available as an adjustment margin rather than reflecting a general agricultural trend.

Appendix A: Data Construction Details

This appendix provides construction details for each component of the analysis panel. Variable definitions, source citations, and IPUMS codes are summarized in Table 7.

A.1 H-2A Certifications

The outcome variable comes from the Department of Labor's Office of Foreign Labor Certification (OFLC) disclosure files, which record every employer application for H-2A guestworker certification. Each record corresponds to one application and reports the number of workers requested, the requested employment start and end dates, and worksite location at the state and county level. I aggregate requested worker counts to the commuting-zone level by year using the Dorn (2009) county-to-CZ crosswalk.

The measure captures *requested certifications* rather than actual worker arrivals. Certifications are the standard outcome in the H-2A literature (e.g., Charlton, 2022) because they are universally available, geographically precise, and closely tracked by both employer demand and program participation. While some certified positions go unfilled, the certification measure closely tracks employer demand and is the standard outcome in the H-2A literature.

National H-2A certifications grew from approximately 75,000 in 2008 to 380,000 in 2024 (Figure 1). Demand is strongly seasonal: requested start dates concentrate between March and September, with peaks in March and April (Figure 4). This seasonal concentration motivates the use of agricultural-season border encounters (March–September) rather than full-year encounters as the alternative enforcement shift in the Bartik instrument.

A.2 The PANCAL Share

The endogenous variable is the share of agricultural laborers in each commuting zone and year who are settled non-citizens. I construct it from the American Community Survey

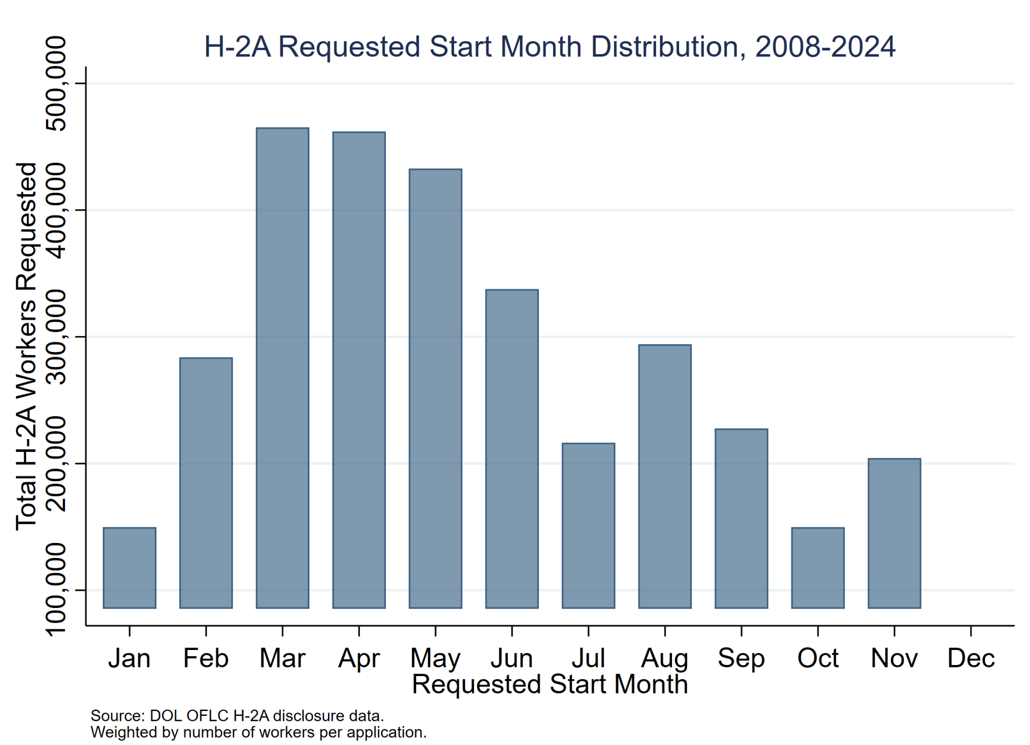

Figure 4: Distribution of H-2A Requested Start Months, 2008–2024. Each bar shows the total number of H-2A workers requested across all employer applications with that requested start month, summed across the sample period. Source: DOL OFLC H-2A disclosure data.

1-year files via IPUMS, applying two filters in sequence.

First, I restrict to *agricultural laborers*: individuals aged 16–64 who are employed (EMPSTAT = 1), report an industry code in agriculture (IND1990 $\in \{10, 11, 20, 30\}$), and report an occupation code in farm labor, nursery work, or agricultural grading and sorting (OCC1990 $\in \{479, 484, 488\}$). The occupation restriction excludes farm operators, managers, supervisors, and clerical workers in agricultural firms, isolating the hired farm labor population. This is the same population that defines the frame of the National Agricultural Workers Survey (NAWS), the standard benchmark dataset for U.S. farm labor.

Second, within this population, the PANCAL subset comprises non-citizens (CITIZEN = 3) who did not live abroad in the preceding year (MIGRATE1 $\neq 4$). The migration filter excludes seasonal migrants and new arrivals, isolating the settled non-citizen workforce. I compute the PANCAL share as the ratio of PANCAL workers to total agricultural laborers within each CZ-year, expressed in percentage points. Across the sample period the share has a weighted mean of approximately 19.6 percent.

Thin cells in the ACS microdata. The ACS sample is thin at the commuting-zone level. The unweighted median CZ-year contains two PANCAL person-records, and 30 percent of CZ-years contain zero PANCAL records. I distinguish two distinct sources of zeros. CZ-years with positive total agricultural employment but zero sampled non-citizens are genuine sampling realizations of a small share and are retained as legitimate zeros. CZ-years with zero total agricultural employment in the ACS—168 of 11,840 observations overall—have an undefined share denominator and are dropped. The regression-weighting scheme (year-2000 total agricultural employment, described in Section 3) downweights CZs with small agricultural workforces precisely because of this measurement issue.

A.3 The 2000 Baseline Share

The share component of the Bartik instrument is the PANCAL share measured in the 2000 Decennial Census 5% sample, accessed through IPUMS. I apply the same agricultural-laborer definition (working age, employed, agricultural industry and occupation codes) and the same non-citizen filter (CITIZEN $= 3$) as in the contemporaneous measure. Because the MIGRATE1 variable is not consistently available in the 2000 Census, I substitute a years-in-United-States filter (YRSUSA1 > 1), which excludes individuals in their first year in the country. The resulting baseline share is fixed at 2000—eight years before the panel begins—to ensure it is predetermined relative to the outcome period.

The 2000 baseline share has a weighted mean of 12.0 percent and a standard deviation of 14.7 percentage points across the 740 commuting zones, with values ranging from 0 to 74 percent. The distribution is right-skewed: most commuting zones have low baseline shares, but a meaningful tail of CZs in California, Florida, Texas, Washington, and Oregon have shares above 40 percent.

A.4 National Enforcement Time Series

The shift component of the Bartik instrument comes from two national time series.

The preferred series is the real (2024 USD) Border Patrol budget. I assemble the nominal annual budget from Department of Homeland Security Budget-in-Brief documents, Congressional Research Service report RL32562, and Congressional appropriations records, and deflate by the CPI-U. The series rises monotonically from approximately $3.3 billion in 2008 to $7.3 billion in 2024, reflecting the bipartisan expansion of border enforcement capacity over this period. Massey et al. (2016) use the same real BP budget series as their preferred enforcement measure.

The second series is the log number of US-Mexico border encounters during the agricultural season (March–September), constructed from Customs and Border Protection encounter records. Unlike the budget, the encounter series is non-monotonic: encounters fell from 681,000 in 2008 to 251,000 in 2017, surged to 2.3 million in 2022, and declined to 1.1 million by 2024. Section 3 discusses why the non-monotonic time path makes the encounter series the more credible identifier and reports the placebo tests that motivate this conclusion.

A.5 Crop Composition Controls

Because high-PANCAL commuting zones are disproportionately concentrated in labor-intensive crop production, I include controls for baseline crop composition. I use the 2002 Census of Agriculture, accessed through USDA Quick Stats, to compute county-level acreage in four categories: total cropland, vegetables, orchards, and tobacco. I aggregate to the commuting-zone level using the Dorn county-to-CZ crosswalk and compute the share of cropland in each category.

The preferred specification interacts the year-2000 vegetable cropland share with year dummies (15 control parameters), allowing CZs with different baseline vegetable intensities to follow differential H-2A trajectories. Section 5 reports robustness to alternative crop control specifications, including a combined labor-intensive index (vegetables + orchards +

berries) and a fuller specification with separate vegetable, orchard, and tobacco interactions; orchard and tobacco interactions are statistically indistinguishable from zero in nearly all years and the additional parameters reduce instrument power without changing the second-stage estimate.

A.6 Final Panel and Summary Statistics

The final panel contains 11,840 commuting-zone-year observations (740 CZs $\times$ 16 years, with 2020 excluded). Of these, 11,672 have a non-missing PANCAL share. Table 8 reports summary statistics for the full sample and for splits by above- and below-median year-2000 PANCAL share. CZs above the median employ substantially more H-2A workers per year (391 versus 102), have a larger total agricultural workforce (1,555 versus 436 workers), a higher Hispanic population share (18.3 percent versus 4.5 percent), a higher vegetable cropland share (1.8 percent versus 0.3 percent), and are located closer to the US-Mexico border (1,281 versus 1,726 kilometers).

Table 7: Variable Definitions and Sources

Variable	Years	Definition	Source
Outcome			
H-2A workers	2008–2024	Count of H-2A workers certified per CZ-year	DOL OFLC disclosure files
Endogenous variable			
PANCAL share	2008–2024	Non-citizen ag laborers who did not live abroad in prior year, as share of total ag laborers (%)	ACS 1-year via IPUMS
Instrument: shares			
PANCAL share, 2000	2000	Same construction as above, with `yrsusa1 > 1` replacing the migration filter	2000 Decennial Census 5% sample via IPUMS
Instrument: shifts			
BP budget (real)	2008–2024	log real Border Patrol budget, 2024 USD	DHS Budget-in-Brief; CRS RL32562; deflated by CPI-U
Ag-season encounters	2008–2024	log Mar–Sep US-Mexico border encounters	CBP encounter records
Controls			
Crop shares	2002	Vegetable, orchard, and tobacco share of total cropland	USDA Census of Agriculture (Quick Stats)

Table 8: Summary Statistics

	(1) Full Sample	(2) Low PANCAL (Below median)	(3) High PANCAL (At/above median)
Panel A: Outcome and endogenous variable			
H-2A workers (count)	246.7	102.4	391.0
	(892.7)	(294.9)	(1,210.5)
PANCAL share (%)	19.6	8.2	30.8
	(22.8)	(13.9)	(24.3)
Panel B: Agricultural employment			
Total ag laborers	995.6	436.4	1,554.8
	(3,317.7)	(538.7)	(4,593.6)
Vegetable cropland share	0.010	0.003	0.018
	(0.034)	(0.009)	(0.046)
Panel C: Demographics and labor market			
Working-age population	280,794	128,718	432,870
	(779,586)	(281,307)	(1,044,134)
Hispanic share	0.114	0.045	0.183
	(0.147)	(0.037)	(0.180)
Employment rate	0.681	0.692	0.669
	(0.070)	(0.076)	(0.062)
Panel D: Geography			
Distance to US-Mexico border (km)	1,503	1,726	1,281
	(768)	(715)	(743)
Panel E: Baseline (year-2000) shares			
PANCAL share, 2000	0.120	—	—
	(0.147)		
Non-ag PANNCL share, 2000	0.033	—	—
	(0.044)		
Total ag laborers, 2000	9,047	—	—
	(weighted mean)		
Observations (CZ-years)	11,840	5,920	5,920
Commuting zones	740	370	370

Notes: Panel-level means with standard deviations in parentheses. Column (1) reports the full sample (740 commuting zones $\times$ 16 years, with 2020 excluded; 168 observations have missing PANCAL share due to zero total agricultural employment). Columns (2) and (3) split the sample by the median year-2000 PANCAL share (0.060). The high-PANCAL group employs nearly four times as many H-2A workers per CZ-year on average, has approximately three times the agricultural workforce, and is located closer to the US-Mexico border. The weighted mean of total agricultural employment ($\bar{N} = 9{,}047$) is the regression-weighted mean used to compute the substitution rate in Section 4; this differs from the simple sample mean of total ag laborers because the regression weights observations by year-2000 CZ agricultural employment.

Appendix B: Construction of the 2020-PUMA-to-Commuting-Zone Crosswalk

The Autor–Dorn–Hanson commuting-zone (CZ) framework relies on PUMA-to-CZ crosswalks that map American Community Survey microdata geographies to the time-consistent 1990 Tolbert–Sizer commuting zones. Dorn has published crosswalks for the 2000-vintage PUMAs (used in the 2006–2011 ACS) and the 2010-vintage PUMAs (used in the 2012–2021 ACS). The 2020-vintage PUMA delineations, in use from the 2022 ACS onward, do not yet have a corresponding published crosswalk. This appendix documents the crosswalk I construct following Dorn's methodology.

B.1 Main Pipeline

The construction chains 2020 PUMAs to 2020 counties to 1990-defined commuting zones using population-weighted allocation factors:

$$\text{2020 PUMA} \xrightarrow[\text{Geocorr 2022}]{\text{pop-weighted}} \text{2020 County} \xrightarrow[\text{Dorn E7}]{\text{exact 1:1}} \text{1990 CZ.}$$

The PUMA-to-county step uses population-weighted allocation factors from the Missouri Census Data Center's Geocorr 2022 application, which assigns each 2020 PUMA to one or more 2020 counties based on the share of the PUMA's 2020 Census population residing in each county. The county-to-CZ step uses Dorn's published 1990 crosswalk, in which each county maps deterministically to a single CZ, so this step introduces no additional approximation. Final allocation factors are computed by summing the PUMA-to-county factors for all counties belonging to the same CZ. The resulting crosswalk contains 3,260 PUMA-CZ pairs across 2,462 unique 2020 PUMAs and 740 commuting zones; 64% of PUMAs map to a single CZ and 36% are split across two to thirteen CZs. Allocation factors sum to 1.0 within each PUMA by construction.

B.2 Connecticut

In 2022, Connecticut replaced its eight historical counties with nine planning regions as the county-equivalent for federal statistical purposes. Geocorr 2022 returns planning region FIPS codes (09110–09190) rather than the legacy county FIPS codes (09001–09015) used in Dorn's 1990 crosswalk. Because the planning regions do not nest into the historical counties (towns were redistributed across boundaries during the redefinition), the standard county-to-CZ chain fails for all 26 Connecticut PUMAs. I resolve this with a separate Geocorr query mapping 2020 PUMAs to 2010 PUMAs for Connecticut only, then chaining the resulting 2010 PUMAs through the published Dorn 2010-PUMA-to-CZ crosswalk:

$$\text{2020 PUMA (CT)} \xrightarrow{\text{Geocorr}} \text{2010 PUMA (CT)} \xrightarrow{\text{Dorn E6}} \text{1990 CZ.}$$

Combined allocation factors are the product of the two stage-specific factors: $a_{2020\to\text{CZ}} = a_{2020\to2010} \times a_{2010\to\text{CZ}}$.

B.3 Alaska

Nine Alaska boroughs and census areas (FIPS codes 02063, 02066, 02068, 02105, 02195, 02198, 02230, 02275, and 02282) were reorganized after 1990 and have no match in Dorn's county-to-CZ crosswalk. These unmatched areas appear in two 2020 PUMAs (200300 and 200400), collectively representing approximately 27,000 residents in remote regions with negligible agricultural employment. Allocation factors for these two PUMAs are rescaled to sum to 1.0 after dropping the unmatched portions. A separate consequence of the post-2020 Alaska reorganization is the disappearance of CZ 34103 (Valdez–Cordova Census Area): the underlying historical county was split into Chugach (FIPS 02063) and Copper River (FIPS 02066) Census Areas, neither of which has a Dorn match, so CZ 34103 is observed in 2008–2021 but not in 2022–2024. The CZ has approximately 10,000 residents and zero agricultural relevance for this analysis.

B.4 Louisiana PUMA 297777

A small number of observations (9,518 person-records, all in the 2006–2011 ACS) are coded to Louisiana PUMA 297777, a special Census Bureau code used for post-Hurricane Katrina displaced residents who could not be reliably geolocated to a specific PUMA of residence. Because PUMA 297777 has no defined geographic footprint, no allocation factors can be computed and these observations cannot be assigned to a commuting zone. They are flagged in the master file and dropped from all CZ-level analyses; the affected records represent approximately 0.05% of the underlying ACS sample over 2006–2011 and have no material effect on the panel.

B.5 Post-1990 County FIPS Recodes

Dorn's crosswalk uses 1990 county FIPS codes; Geocorr 2022 uses 2020 codes. The following recodes are applied to align the two:

Current name	1990 FIPS	2020 FIPS	Change
Miami-Dade, FL	12025	12086	Renamed from Dade County (1997)
Oglala Lakota, SD	46113	46102	Renamed from Shannon County (2015)
Kusilvak, AK	02270	02158	Renamed from Wade Hampton (2015)
Broomfield, CO	—	08014	Created 2001; assigned to CZ 28900

B.6 Probabilistic Assignment and Weight Verification

PUMAs do not nest perfectly within commuting zones. Following the standard Autor–Dorn–Hanson convention, persons residing in PUMAs that span multiple CZs are duplicated into one row per overlapping CZ, with person weights scaled by the corresponding allocation factor. Formally, for an individual i in PUMA p with original ACS person weight perwt_i, the CZ-adjusted weight in commuting zone c is

$$\text{perwt\_cz}_{i,c} = \text{perwt}_i \times a_{p,c},$$

where $a_{p,c}$ is the allocation factor for PUMA p to CZ c, and $\sum_c a_{p,c} = 1$ for every PUMA p. All commuting-zone-level analyses in this paper use perwt_cz.

To verify that the probabilistic expansion preserves total weighted population, I compare pre-expansion and post-expansion weight sums for each ACS year. For 2012, the pre-expansion sum of perwt equals the post-expansion sum of perwt_cz at 313,914,040 (ratio = 1.000); analogous verification holds in every year. The expansion converts the underlying 57.0 million ACS person-records into 96.0 million CZ-assigned records over 2006–2024.

Appendix C: Additional Robustness Results

This appendix contains supplementary figures and tables referenced in the main text. Sections are organized to follow the order of the main text's Robustness section (Section 5).

C.1 Additional Descriptive Figures

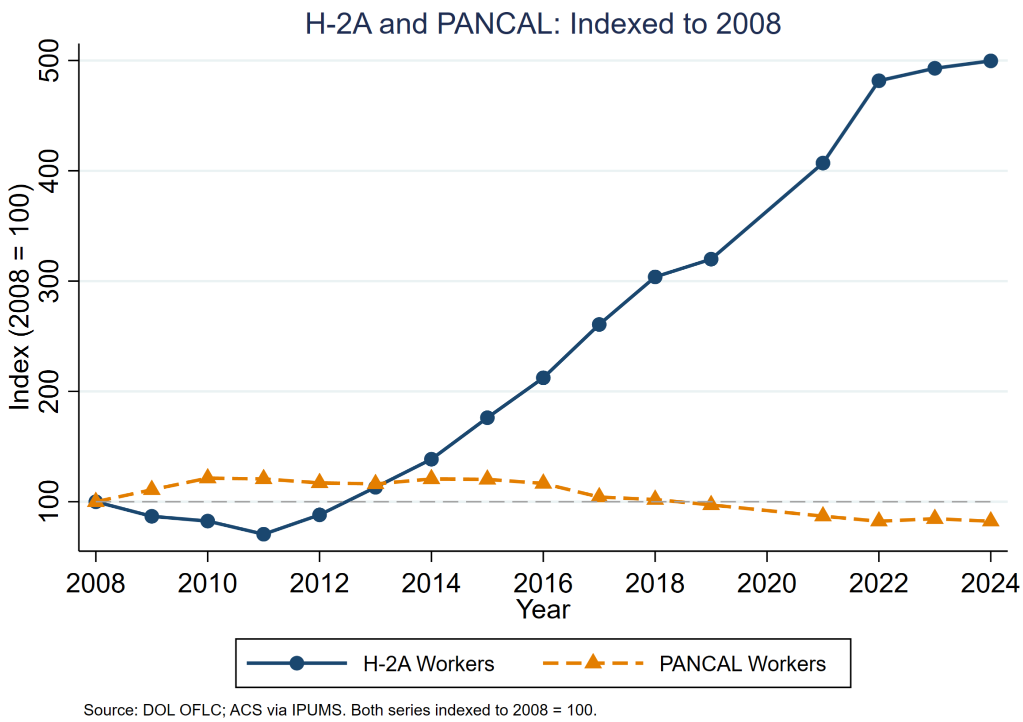

Figure 5: National H-2A and Settled Non-Citizen Agricultural Labor Indexed to 2008, 2008–2024. Both series are indexed to 2008 = 100 to facilitate comparison of growth rates. Levels are presented in main-text Figure 1. Data sources as in Figure 1.

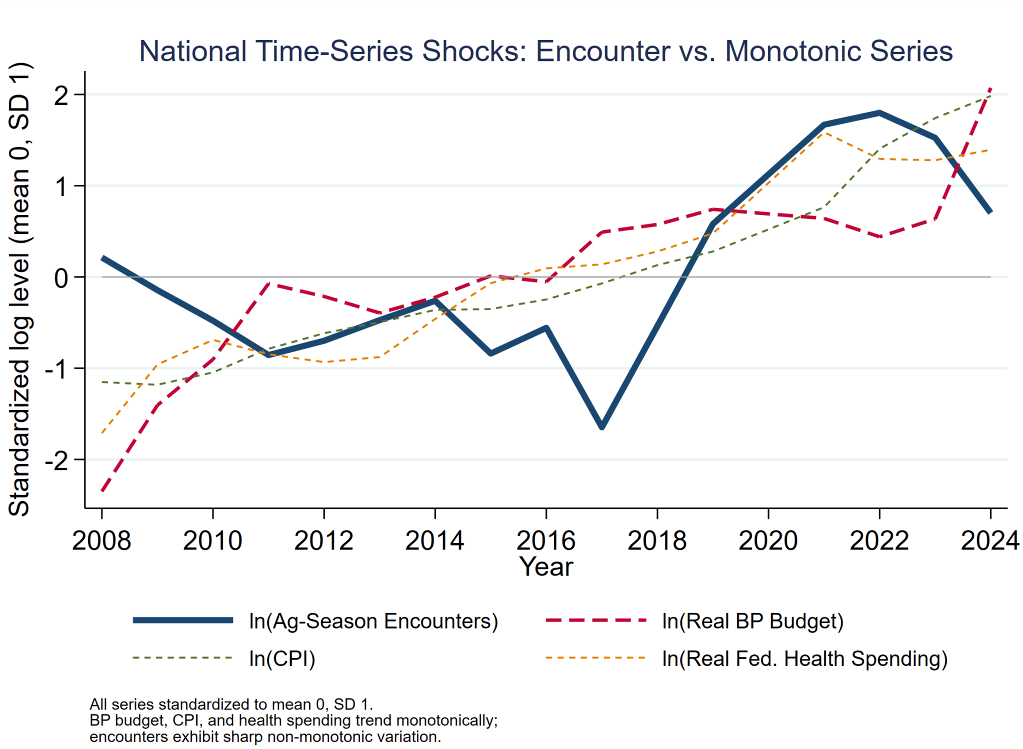

Figure 6: National time-series shocks: enforcement measures and placebo series, 2008–2024. All series are log-transformed and standardized to mean 0, standard deviation 1. The Border Patrol budget, CPI-U, and real federal health spending all trend broadly upward; ag-season border encounters exhibit sharp non-monotonic variation. Sources: DHS Budget-in-Brief and CRS RL32562 (BP budget); CBP encounter records; BLS (CPI-U); OMB Historical Tables (federal health spending).

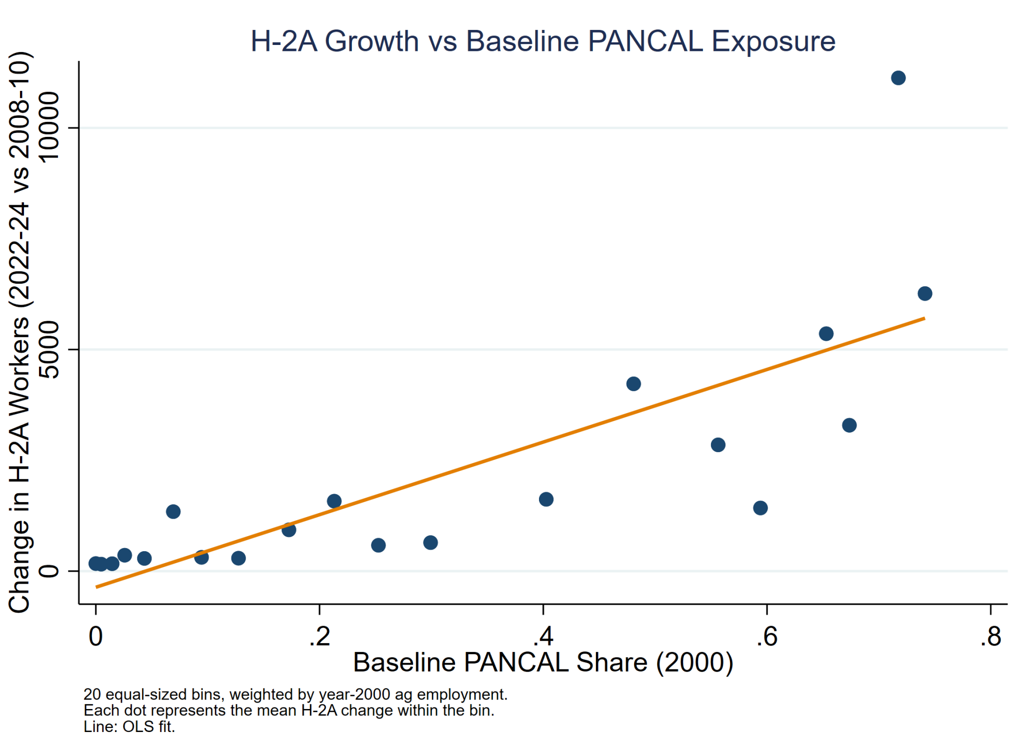

Figure 7: Binned Scatter of H-2A Growth vs. Baseline PANCAL Exposure. Each dot represents the mean change in H-2A workers (2022–2024 average minus 2008–2010 average) within a bin of baseline PANCAL share (2000), using 20 equal-sized quantile bins weighted by year-2000 CZ agricultural employment. The upward-sloping relationship provides the raw cross-sectional visual of the reduced-form effect before any controls or fixed effects.

C.2 Trending-Series Benchmarks: Reduced Form

Table 9: Trending-Series Benchmarks: Reduced-Form Regressions of H-2A on Bartik

	(1) BP Budget	(2) CPI Spending	(3) Federal Health Spending	(4) Linear Trend
Reduced-form coefficient	55.4***	0.435***	62.4***	2.59***
	(16.1)	(0.119)	(17.3)	(0.729)
t-statistic	3.44	3.66	3.60	3.55
Veg × year controls	Yes	Yes	Yes	Yes
CZ & year FE	Yes	Yes	Yes	Yes
Observations	11,840	11,840	11,840	11,840

Notes: Each column reports an OLS regression of H-2A workers on the corresponding Bartik instrument, weighted by year-2000 CZ agricultural employment, with vegetable share × year interactions, CZ fixed effects, and year fixed effects. Standard errors clustered at the CZ level in parentheses. Coefficients are not directly comparable in magnitude across columns because the underlying national time series have different units; the relevant comparison is the t-statistic, which is essentially identical across all four Bartiks. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.3 Alternative Shift Constructions: LOO and Full National Share

Table 10: Alternative Shift Constructions: LOO and Full National Share

	(1) Benchmark BP budget	(2) Alt (full) $\bar{S}_t$	(3) LOO $\bar{S}_{-cz,t}$
$\hat{\delta}$ (PANCAL share, pct)	−155.3**	−180.7***	−223.8***
	(61.1)	(60.6)	(83.8)
Substitution rate	1.72	2.00	2.47
KP F-statistic	14.82	24.43	16.10
AR test p-value	< 0.001	< 0.001	< 0.001
Veg × year controls	Yes	Yes	Yes
CZ & year FE	Yes	Yes	Yes
Observations	11,672	11,672	11,672

Notes: All specifications weighted by year-2000 CZ agricultural employment, with vegetable share × year interactions, CZ fixed effects, and year fixed effects. Column (1) reproduces the BP budget benchmark. Column (2) replaces the BP budget with the full national PANCAL share as the shift: $Z^{\text{alt}}_{cz,t} = S^0_{cz} \times \bar{S}_t$. Column (3) uses the leave-one-out national share excluding CZ cz's contribution: $Z^{\text{LOO}}_{cz,t} = S^0_{cz} \times \bar{S}_{-cz,t}$. Standard errors clustered at the CZ level in parentheses. The substitution rate equals $|\hat{\delta}|/\bar{N}$, where $\bar{N} = 9{,}047$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.4 Crop Control Alternatives

Table 11: Alternative Crop Control Specifications

	(1) Lab-Int × Yr	(2) Veg+Orch+Tob × Yr	(3) Veg × Yr (Preferred)
$\hat{\delta}$ (PANCAL share, pct)	−49.5	−115.8	−155.3**
	(95.5)	(84.6)	(61.1)
KP F-statistic	7.24	6.85	14.82
Anderson-Rubin p-value	0.572	0.091	< 0.001
Control parameters	15	45	15
CZ & year FE	Yes	Yes	Yes
Observations	11,672	11,672	11,672

Notes: All specifications use the BP budget Bartik as the excluded instrument for the PANCAL share, with the outcome in levels and observations weighted by year-2000 CZ agricultural employment. Column (1) uses the labor-intensive index share (vegetables + orchards + berries) interacted with year dummies. Column (2) uses separate vegetable, orchard, and tobacco shares each interacted with year dummies (45 parameters total). Column (3) reproduces the preferred specification with vegetable share × year interactions only (15 parameters). Standard errors clustered at the CZ level in parentheses. The labor-intensive aggregation and the veg+orchard+tobacco specification both reduce first-stage strength below the Stock-Yogo 10% maximal IV size threshold (16.38) while producing larger second-stage standard errors. The vegetable-only specification is preferred because orchards and tobacco have no detectable independent relationship with H-2A growth and their inclusion reduces instrument power without changing the substantive conclusion. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 12: Sensitivity to Outcome Scaling and Weighting

	(1) Levels, Weighted (Preferred)	(2) asinh, Weighted	(3) Levels, Unweighted
$\hat{\delta}$ (PANCAL share, pct)	-155.3^{**}	-0.083^{*}	-63.4^{***}
	(61.1)	(0.043)	(22.0)
Substitution rate	1.72	—	6.36
Semi-elasticity (%)	—	8.3	—
KP F-statistic	14.82	14.82	52.16
Anderson-Rubin p-value	< 0.001	0.019	0.001
Outcome	H-2A workers	asinh(H-2A)	H-2A workers
Weights	total_ag_2000	total_ag_2000	None
Veg × year controls	Yes	Yes	Yes
CZ & year FE	Yes	Yes	Yes
Observations	11,672	11,672	11,672

Notes: All specifications use the BP budget Bartik as the excluded instrument for the PANCAL share and include vegetable share × year interactions (15 parameters), commuting zone fixed effects, and year fixed effects. Column (1) reproduces the preferred specification. Column (2) uses the inverse hyperbolic sine of H-2A workers as the outcome, yielding a semi-elasticity interpretation: a one percentage point decline in the PANCAL share is associated with a 8.3 percent increase in H-2A workers. Column (3) drops the analytic weights. The substitution rate equals $|\hat{\delta}|/\bar{N}$, where $\bar{N}_w = 9{,}047$ for weighted specifications and $\bar{N}_{uw} = 996$ for unweighted specifications. The larger unweighted substitution rate reflects the smaller denominator that arises when commuting zones with small agricultural workforces receive equal weight; the weighted specification is preferred because it gives proportional influence to commuting zones where the substitution mechanism operates at scale. Standard errors clustered at the CZ level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 13: Year-by-Year Crop × Year Coefficients from Spec C

Year	Vegetable × year	Orchard × year	Tobacco × year
2009	1,577	−644	−41,743
	(1,026)	(901)	(44,311)
2010	242	2	−21,639
	(503)	(945)	(24,107)
2011	−1,676	913	−29,934
	(1,088)	(1,225)	(33,179)
2012	−97	4	−40,403
	(1,079)	(1,125)	(35,864)
2013	−1,679	142	−38,218
	(1,118)	(1,531)	(36,431)
2014	−1,076	522	−17,325
	(969)	(2,026)	(30,156)
2015	3,324**	−526	−12,619
	(1,509)	(2,320)	(33,432)
2016	5,412***	−126	25,934*
	(1,881)	(2,980)	(15,283)
2017	6,351***	323	−18,152
	(1,184)	(2,257)	(40,207)
2018	7,962***	75	−35,902
	(1,137)	(2,758)	(52,717)
2019	10,142***	−304	11,853
	(1,412)	(2,955)	(25,403)
2021	12,297***	2,399	−12,926
	(909)	(2,924)	(35,201)
2022	18,174***	2,288	−24,222
	(1,059)	(3,459)	(46,878)
2023	17,803***	1,831	−11,675
	(912)	(3,538)	(36,691)
2024	16,090***	2,058	1,545
	(1,202)	(3,096)	(43,941)
Number of significant coefficients (p < 0.10)			
	9 of 15	0 of 15	1 of 15

Notes: Coefficients on baseline crop-share × year interactions from a single 2SLS specification using the BP budget Bartik as the excluded instrument for the PANCAL share. The specification includes vegetable share × year, orchard share × year, and tobacco share × year interactions (45 control parameters total), commuting zone fixed effects, and year fixed effects, and is weighted by year-2000 CZ agricultural employment. Standard errors clustered at the CZ level in parentheses. The reference year is 2008; 2020 is excluded due to ACS data availability. Vegetable × year coefficients are positive, large, and statistically significant in 9 of 15 years (all post-2014). Orchard × year coefficients are statistically insignificant in every year. Tobacco × year coefficients are statistically insignificant in 14 of 15 years. The Kleibergen-Paap first-stage F-statistic for this specification is 6.85, below conventional weak-instrument thresholds, confirming that adding orchard and tobacco controls reduces instrument power without adding identifying information. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.5 Distance Falsification: Full Results

Table 14: Distance to the Border as Alternative Exposure Measure

	(1) Distance IV	(2) Distance RF	(3) Horse race (RF)
Panel A: Distance Bartik as instrument			
$\hat{\delta}$ (PANCAL share, pct)	−31.4	—	—
	(61.3)	—	—
KP F-statistic	4.67	—	—
Anderson-Rubin p-value	0.63	—	—
Panel B: Reduced form			
PANCAL Bartik	—	—	72.9***
	—	—	(23.8)
Distance Bartik	—	1.74	−7.90
	—	(3.60)	(5.76)
Veg × year controls	Yes	Yes	Yes
CZ & year FE	Yes	Yes	Yes
Observations	11,672	11,840	11,840

Notes: The distance Bartik replaces the year-2000 PANCAL share with negative log distance from the CZ population-weighted centroid to the nearest US-Mexico border point: $Z_{c,t}^{dist} = -\ln(d_c) \times \ln(\text{BP budget}_t)$. Column (1) reports a 2SLS specification using the distance Bartik as the excluded instrument for the PANCAL share. Column (2) reports the reduced-form regression of H-2A on the distance Bartik alone. Column (3) reports a reduced form with both Bartiks entered jointly. All specifications are weighted by year-2000 CZ agricultural employment, with vegetable share × year interactions, CZ fixed effects, and year fixed effects. Standard errors clustered at the CZ level in parentheses. The cross-sectional correlation between the year-2000 PANCAL share and log distance is −0.46. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.6 Lag-Instrument Robustness

C.7 Planned Additional Results (TODO)

The following supplementary analyses are planned for future drafts: Oster (2019) bounding exercises assessing the sensitivity of the main estimate to unobserved confounders, and a sector-specific CBP budget instrument exploiting regional variation in enforcement capacity.

Table 15: Lag-Instrument Robustness

	(1) Contemp BP (Benchmark)	(2) Lag BP	(3) Contemp Enc (Benchmark)	(4) Lag Enc
$\hat{\delta}$ (PANCAL share, pct)	-155.3^{**}	-154.5^{***}	-230.0^{***}	-233.2^{**}
	(61.1)	(54.7)	(64.7)	(96.1)
Substitution rate	1.72	1.71	2.54	2.58
KP F-statistic	14.82	23.97	32.45	9.27
AR test p-value	< 0.001	< 0.001	< 0.001	< 0.001
Instrument	BP budget	L.BP budget	Ag-season encounters	L.Ag-season encounters
Veg $\times$ year controls	15	13	15	13
CZ & year FE	Yes	Yes	Yes	Yes
Observations	11,672	10,223	11,672	10,223
CZ clusters	740	740	740	740

Notes: All specifications weighted by year-2000 CZ agricultural employment, with CZ and year fixed effects and vegetable share $\times$ year interactions. Columns (1) and (3) reproduce the benchmark contemporaneous specifications from Table 1. Columns (2) and (4) replace the contemporaneous Bartik instrument with its one-year lag and restrict the sample to observations with a non-missing lagged instrument, dropping 2008 (no 2007 in panel) and 2021 (no 2020 in panel due to ACS gap). In the restricted sample, VEG_YR2021 is identically zero (no 2021 observations remain) and is omitted by the collinearity check; one additional interaction is absorbed as the reference category since 14 years of data identify at most 13 veg $\times$ year interactions. The substitution rate equals $|\hat{\delta}|/\bar{N}$, where $\bar{N} = 9{,}047$ is the regression-weighted mean of total agricultural employment. Anderson–Rubin (p-values reported) is the weak-instrument-robust test. Standard errors clustered at the CZ level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 16: Long-Difference IV Robustness

	(1) Benchmark (Panel)	(2) 5-yr FD: BP	(3) 5-yr FD: Enc.
$\hat{\delta}$ (PANCAL share, pct)	-155.3^{**}	-132.2^{**}	-330.9^{**}
	(61.1)	(64.3)	(141.8)
Substitution rate	1.72	1.46	3.66
KP F-statistic	14.82	10.13	8.12
AR test p-value	< 0.001	0.002	< 0.001
Specification	Levels + CZ FE	5-yr FD	5-yr FD
Instrument	BP budget	L5.BP budget diff.	L5.Enc. diff.
CZ FE absorbed via	FE	Differencing	Differencing
Year FE	Yes	Yes	Yes
Observations	11,672	7,964	7,964

Notes: Column (1) reproduces the benchmark panel IV from Table 1. Columns (2) and (3) re-estimate the model in 5-year first-differences, which removes commuting zone fixed effects mechanically and identifies $\hat{\delta}$ from long-run within-CZ variation. Three-year first-difference specifications are not reported because the 2020 ACS gap eliminates instrument variation in windows spanning 2020. Standard errors clustered at the CZ level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

References

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Castillo, Marcelo, Philip Martin, and Zachariah Rutledge, "Whither the H-2A Visa Program: Expansion and Concentration," *Choices*, 2024, *39* (1), 1–9.

Charlton, Diane, "Seasonal Farm Labor and COVID-19 Spread," *Applied Economic Perspectives and Policy*, 2022, *44* (3), 1591–1609.

_ **and Genti Kostandini**, "Can Technology Compensate for a Labor Shortage? Effects of 287(g) Immigration Policies on the U.S. Dairy Industry," *American Journal of Agricultural Economics*, 2021, *103* (1), 70–89.

Charlton, Diane E., "Immigration Enforcement and the US Agricultural Sector in 2025," American Enterprise Institute Report April 2025.

Charlton, Diane, Marcelo Castillo, Alexandra E Hill, and J Edward Taylor, "Can Wages Rise Quickly Enough to Keep Workers in the Fields?," *Choices*, 2019, *34* (2).

Clemens, Michael A., Ethan G. Lewis, and Hannah M. Postel, "Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion," *American Economic Review*, 2018, *108* (6), 1468–1487.

East, Chloe N. and Andrea Velásquez, "Unintended Consequences of Immigration Enforcement: Household Services and High-Educated Mothers' Work," *Journal of Human Resources*, 2024, *59* (5), 1458–1502.

__ , **Annie L. Hines, Philip Luck, Hani Mansour, and Andrea Velásquez**, "The Labor Market Effects of Immigration Enforcement," *Journal of Labor Economics*, 2023, *41* (4), 957–996.

Fan, Maoyong, Susan Gabbard, Anita Alves Pena, and Jeffrey M Perloff, "Why Do Fewer Agricultural Workers Migrate Now?," *American Journal of Agricultural Economics*, 2015, *97* (3), 665–679.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Hanson, Gordon H., "Illegal Migration from Mexico to the United States," *Journal of Economic Literature*, 2006, *44* (4), 869–924.

Hertz, Tom and Steven Zahniser, "Is There a Farm Labor Shortage?," *American Journal of Agricultural Economics*, 2013, *95* (2), 476–481.

Kostandini, Genti, Elton Mykerezi, and Cesar Escalante, "The Impact of Immigration Enforcement on the U.S. Farming Sector," *American Journal of Agricultural Economics*, 2014, *96* (1), 172–192.

Massey, Douglas S., Jorge Durand, and Karen A. Pren, "Why Border Enforcement Backfired," *American Journal of Sociology*, 2016, *121* (5), 1557–1600.

Olea, José Luis Montiel and Carolin Pflueger, "A Robust Test for Weak Instruments," *Journal of Business & Economic Statistics*, 2013, *31* (3), 358–369.

Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204.

Passel, Jeffrey S. and Jens Manuel Krogstad, "What We Know About Unauthorized Immigrants Living in the U.S.," Pew Research Center July 2024.

U.S. Department of Agriculture, Economic Research Service, “Farm Labor,” 2024. Accessed April 15, 2026.