Where Occupations Move:
Interstate Migration by Occupation in the United States, 2005 to 2024

Alessandro Conway

Working Paper AC-WP-2026-02

Version 1.1 (revised September 2026)

Not peer reviewed. Comments welcome.

Abstract

Interstate migration in the United States has been measured for persons and rarely for occupations, although most of the rules that make moving difficult attach to the job. This paper builds bilateral interstate migration flows by occupation from nineteen vintages of American Community Survey microdata, 59.8 million person records across 51 jurisdictions and 475 harmonised occupation codes, and measures the share of each state's occupational workforce arriving from another state within the year, with design-based standard errors. That share varies close to fifteenfold across occupations, from tool and die makers to sailors, and licensure explains almost none of the ordering. Arrivals are younger than the workforce they join in every occupation measured, and arrivals and leavers have the same mean age cell by cell, both about seven years younger than the incumbents. A Poisson pseudo-maximum-likelihood gravity model with origin-year, destination-year, occupation-year, and state-pair effects puts the wage elasticity of flows at 0.114 and the distance elasticity at -0.72, some six times as large. The benchmark omits employment by state and occupation, so its residual recovers occupational concentration from the flows alone. The concentrations hold their position across two decades and agree with the location quotient in an independent establishment survey, while three pre-specified tests on wage growth return tightly estimated nulls.

Keywords: interstate migration, occupational labor markets, gravity model, occupational licensing, American Community Survey
JEL classification: J61, R23, J44, J62, C23

Suggested citation: Conway, Alessandro (2026). "Where Occupations Move: Interstate Migration by Occupation in the United States, 2005 to 2024." Working Paper AC-WP-2026-02, Version 1.1.

Where Occupations Move:
Interstate Migration by Occupation in the United States, 2005 to 2024

Alessandro Conway*

First draft March 2026

This version September 2026

Abstract

Interstate migration in the United States has been measured for persons and rarely for occupations, although most of the rules that make moving difficult attach to the job. This paper builds bilateral interstate migration flows by occupation from nineteen vintages of American Community Survey microdata, 59.8 million person records across 51 jurisdictions and 475 harmonised occupation codes, and measures the share of each state's occupational workforce arriving from another state within the year, with design-based standard errors. That share varies close to fifteenfold across occupations, from tool and die makers to sailors, and licensure explains almost none of the ordering. Arrivals are younger than the workforce they join in every occupation measured, and arrivals and leavers have the same mean age cell by cell, both about seven years younger than the incumbents. A Poisson pseudo-maximum-likelihood gravity model with origin-year, destination-year, occupation-year, and state-pair effects puts the wage elasticity of flows at 0.114 and the distance elasticity at -0.72, some six times as large. The benchmark omits employment by state and occupation, so its residual recovers occupational concentration from the flows alone. The concentrations hold their position across two decades and agree with the location quotient in an independent establishment survey, while three pre-specified tests on wage growth return tightly estimated nulls.

Keywords: interstate migration, occupational labor markets, gravity model, occupational licensing, American Community Survey.
JEL classification: J61, R23, J44, J62, C23.

1 Introduction

Roughly 2.3 per cent of employed Americans crossed a state line in a given year over the period studied here, and the stability of that figure across two decades conceals a distribution that is wide throughout. Each year tool and die makers draw two thirds of one per cent of their number from other states, while sailors draw close to a tenth. Farmers, postal workers, bus drivers, and highway maintenance workers stay within half a percentage point of the floor, and pilots, air traffic controllers, aircraft mechanics, and physical scientists draw between six and nine per cent of theirs. A

*Working paper. A companion paper, *Where workers come from*, builds the national account of arrivals into and departures from each occupation.

policy instrument aimed at labour mobility acts somewhere inside that distribution, which has not previously been measured at occupation detail across a panel of this length.

The absence follows from what the available data could support. Migration has been studied either at a level of aggregation which suppresses occupation entirely, from Molloy et al. (2011) on the national decline to Kennan and Walker (2011) on the individual decision, or one occupation at a time, as in the studies of individual licensure compacts. Neither approach yields what is constructed here, which sits between the two: bilateral flows by occupation, state, and year, over a period long enough to observe change in them.

A workforce authority comes to such data with a narrower question than the academic one. It wants to know whether the people filling a given occupation locally were trained locally or arrived from elsewhere, whether the arrivals are the ones needed, and whether recruitment or training is the more likely to work. All three require bilateral flows resolved by occupation and by age.

The results bear on that choice directly. Licensure, the barrier the mobility literature most often names, explains almost none of which occupations import their workers. Pay is weakly associated with flows and distance strongly, so a premium of any size a state is likely to offer goes with a change in its intake of a few per cent. Arrivals are younger than the workforce they join in every occupation the panel can measure. The claims are descriptive: the contribution is a set of measures constructed from public data, reported with the uncertainty the survey design implies, and tested against a source from which they were not built.

Section 2 describes the data and the construction of the flow panel, Section 3 reports the scale of external dependence across occupations, and Section 4 turns to who moves. Sections 5 and 6 estimate the gravity model and use departures from it to locate occupational agglomeration, and Sections 7 and 8 test the measures against outside data and set out the limitations.

2 Data and construction

Source and coverage

The flows are built from IPUMS USA extracts of the American Community Survey one-year files for 2005 to 2024 (Ruggles et al., 2024), 59.8 million person records in total. There is no 2020 vintage, because the Census Bureau suspended one-year collection in that year, so the panel runs to nineteen vintages with a gap in the middle, and any calculation crossing 2019 to 2021 measures a two-year change. Occupation is coded on the OCC2010 basis, which is the IPUMS harmonisation, and the four military codes are set aside because service members relocate on posting. That leaves 475 civilian occupations, of which 432 appear in every vintage. The remainder enter or leave at the 2010 and 2018 revision boundaries, where the harmonisation absorbs most but not all of the change in the underlying classification, and every occupation named in this paper is among the 432.

A worker is counted as an interstate mover if the state of residence one year previously is a valid state other than the current one. The rule selects a strict subset of the survey's own interstate migration flag, since no record is a mover by the rule alone. It does discard the 0.6 per cent of movers who have the flag without a recorded origin state, since a flow cannot be assigned to a pair of states when one end of the pair is missing.

Flows and zeros

Flows are counted for every origin, destination, occupation and year, and the zeros are retained. The estimation grid is restricted to occupation-years whose wage is measured from at least thirty non-moving workers at both ends of the corridor, which leaves 4.02 million cells across 400 occupations and 2,550 state pairs. Of those cells 91.6 per cent record no movers at all, and a corridor with no movers in a given year is informative about how hard that move is. Dropping the zeros, which taking logarithms of flows requires, would estimate the model on a selected sample of large corridors and would introduce a bias whose direction cannot be signed in advance.

Wages

The wage panel gives the mean log hourly wage by occupation, state, and year, built from wage and salary workers who did not cross a state line in the year, and that restriction is necessary because the destination wage would otherwise absorb the selection the model exists to detect. Workers who move in response to a good offer earn more on arrival, so an elasticity estimated on wages including theirs would be partly a construction of the measurement. Movers are 2.3 per cent of the workforce, so little is lost by excluding them. Hourly pay is computed from annual wage income, usual weekly hours and weeks worked, which removes the arbitrary threshold on hours that an annual measure requires. The variable records what incumbents in the cell were paid across a year, and it stands for the offer facing a prospective mover only to the degree that the two move together.

A cell mean built from thirty or more workers is still partly noise, so the reliability of the wage variable, defined as the share of its observed variation that is not sampling variance, is computed and used in the estimates below.

Distance and standard errors

Distance is the great-circle distance between state centres of population, which places the origin of Nevada near Las Vegas, where its population is concentrated, and away from its sparsely inhabited geographic centre. Contiguity is derived from the Census county adjacency file.

Standard errors on the sourcing measures are computed by Taylor linearisation on the survey's strata and clusters, which the ACS design supports comfortably: there are 2,462 strata in the final year, never fewer than 166 clusters in any stratum of any year, and no singletons. The median relative standard error across occupations is 6.1 per cent, and 345 of 475 occupations fall below 10 per cent.

3 External dependence

External dependence is the share of a state's employed workforce in an occupation that arrived from another state within the year. Pooled across states and vintages, and among the 400 occupations with at least 5,000 person records, it ranges from 0.66 per cent for tool and die makers to 9.70 per cent for sailors and marine oilers, a spread of close to fifteenfold. The two ends stay an order of magnitude apart even at the limits of their confidence intervals, since the most imported occupation

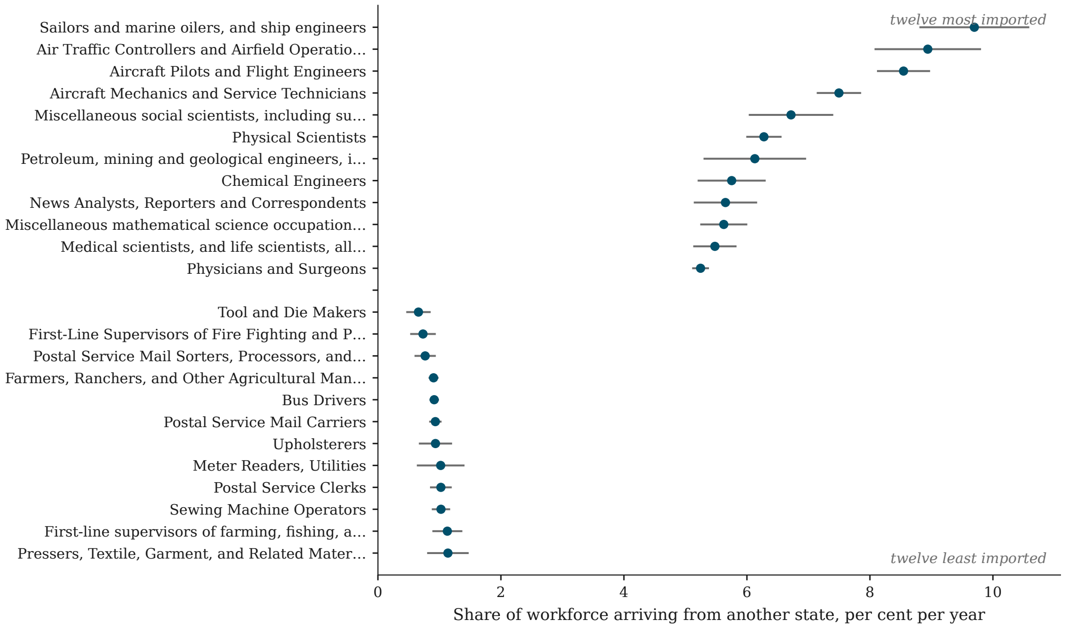

Figure 1: External dependence by occupation, pooled 2005 to 2024. Points are the share of the employed workforce arriving from another state within the year; bars are 95 per cent intervals from Taylor linearisation on the survey design. Occupations with fewer than 5,000 person records are omitted.

is bounded below at 8.80 per cent and the least imported bounded above at 0.86. Figure 1 reports both ends of the distribution.

Occupations at the top of the distribution either serve networks which span states by construction, as sailors, pilots, air traffic controllers, and aircraft mechanics do, or draw on a national market for scarce training, as physical scientists, chemical engineers, and physicians do. Occupations at the bottom are attached to a fixed local asset such as a farm, a postal route, a stretch of highway, or a municipal fire service.

Licensure organises the distribution poorly, although it is the barrier the mobility literature most often names. Physicians and veterinarians are among the most heavily licensed occupations in the country and come near the top of the ranking, while postal workers require no licence at all and fall near the bottom. Through the middle the ordering bears no visible relation to the intensity of licensing, and the credentialling rate measured in the Current Population Survey confirms the impression.[1] Across the 421 occupations for which it is available its rank correlation with external dependence is 0.10, and the 45 occupations in which most workers hold a credential import a median 4.5 per cent of their workforce a year against 4.5 per cent in the rest. A reader who covered the labels and tried to guess which occupations were licensed from their position alone would do no better than chance.

Summed over the country, arrivals and departures are the same people counted twice, so the national gross rate of 4.1 to 5.4 per cent a year is double the arrival rate and contains no separate information. The distinction begins to matter at state level, where a net figure and a gross figure

[1]The share of employed workers in each occupation holding a professional certification or licence, pooled over forty monthly samples from 2015 to 2024. The survey's follow-up on whether a government issued the credential is not harmonised, so the rate is an upper bound on the licensed share.

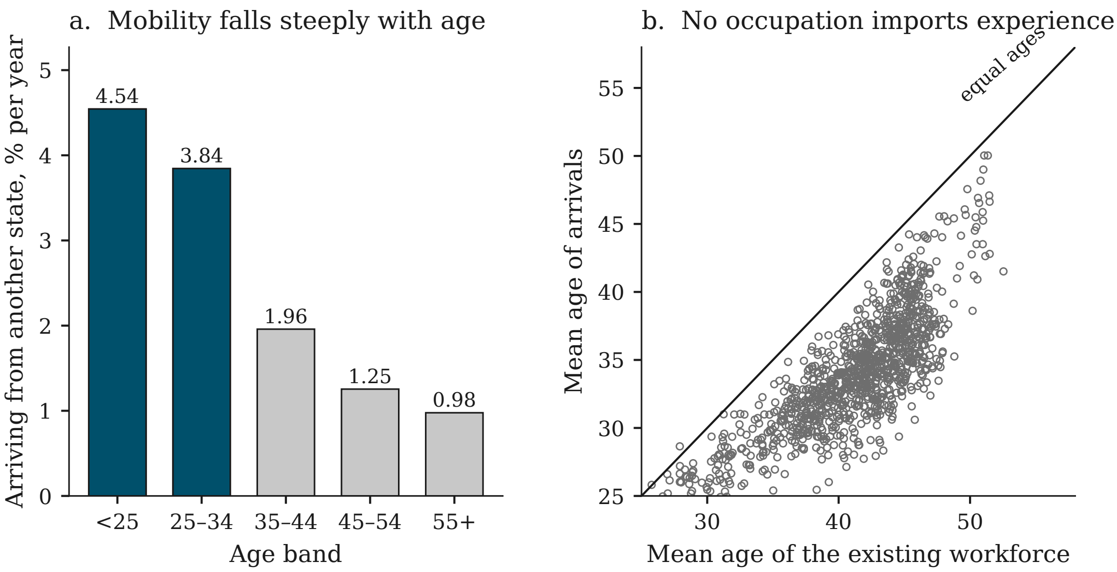

Figure 2: Left: annual external dependence by age band, all occupations pooled. Right: mean age of arrivals against mean age of the existing workforce, one point per state and occupation with at least 100 arrivals and 100 departures. Points below the diagonal indicate arrivals younger than the workforce they join.

can diverge without limit. The District of Columbia replaces close to a fifth of its workforce across state lines each year while its net balance is approximately zero, which a measure built on net flows alone would record as no activity whatever. That is a city in which a fifth of the people at work in any year were living somewhere else twelve months earlier.

4 Age

That migration is concentrated among the young is long established (Rogers and Castro, 1981). Resolving the flows by occupation shows that the pattern holds in every occupation measured, and that the flow is symmetric in age between arrivals and departures.

Annual external dependence falls monotonically with age, from 4.54 per cent among workers under 25 to 0.98 per cent among those aged 55 and over, a gradient of four and a half to one shown in the left panel of Figure 2. That arrivals and leavers have the same mean age nationally, 34.1 years against 34.1, is an identity, since every interstate arrival is somebody else's departure. The equality also holds cell by cell, and that is not an identity.

Across the 1,130 state and occupation cells with at least a hundred arrivals and a hundred departures, the median gap between the mean age of arrivals and the mean age of leavers is -0.2 years. In two thirds of cells the two means fall within two years of each other, against a workforce whose mean age is 41.5. As far as the cells can show, migration exchanges workers for others of much the same age.

All but two of the 1,130 points in the right panel of Figure 2 lie below the diagonal, and the two exceptions are both waiting staff, an occupation whose workforce is itself young. Across the 82 occupations with at least 2,000 observed arrivals the incomers are younger than the workforce they join without exception, from 2.9 years younger for waiters to 11.9 years for aircraft mechanics, 11.0 for physicians, and 10.8 for postsecondary teachers. Interstate migration therefore leaves the age composition of a state's occupational workforce roughly where it found it, exchanging younger members for younger members drawn from elsewhere.

The survey records age and not experience, so what is established here is a statement about age. Arrivals are about seven years younger than the incumbent mean in the typical occupation, and no occupation the panel can measure draws arrivals who are older on average than the workforce they join.

5 The gravity model

Specification

Flows are estimated by Poisson pseudo-maximum likelihood, which accommodates the zeros and the heteroskedasticity that a log-linear specification mishandles (Santos Silva and Tenreyro, 2006), on 4.02 million origin by destination by occupation by year cells. Estimating one occupation at a time is not viable, and the attempt returns a null which is an artefact of the design. Within a single occupation the wage differential is identified only from movement in a state by occupation cell mean over time, and a reliability calculation attributes 54 per cent of that movement to sampling noise. The coefficient is therefore attenuated to roughly half its size, against standard errors too wide to detect it.

Pooling occupations instead identifies the elasticity from variation across occupations within a state pair and year, whose reliability is 0.76 before the fixed effects are applied. It also permits origin-by-year and destination-by-year effects, which absorb the time-varying multilateral resistance that would otherwise remain in the error term (Anderson and van Wincoop, 2003; Bertoli and Fernández-Huertas Moraga, 2013; Head and Mayer, 2014).

Wage and distance elasticities

Table 1 reports the wage elasticity across four specifications, each more demanding than the last, with standard errors in parentheses. The estimate moves very little, and it rises slightly as effects are added, which suggests that omitted multilateral resistance had biased the looser specifications towards zero. The preferred estimate holds with origin-by-year, destination-by-year, occupation-by-year, and dyad effects together, so it is identified from the way wage differentials vary across occupations within one state pair in one year.

An elasticity of 0.114 means that a destination paying ten per cent more than the origin receives about one per cent more migration in that occupation. That is a weak relationship: two corridors have to differ in relative pay by more than a factor of two before their flows differ by a tenth. The coefficient is an equilibrium association, estimated in data where amenities, prices, and labour demand all move alongside the wage, and it does not describe what would follow from a pay rise a state decided to make.

Dividing the coefficient by the reliability of the wage variable, 0.76, gives 0.15. That correction

Table 1: Wage, distance and contiguity elasticities of interstate migration flows, Poisson pseudo-maximum likelihood on 4.02 million cells, standard errors clustered by state pair. Distance and contiguity are absorbed once state pair effects enter.

Fixed effects	Wage	p	Log distance	Contiguity
Origin, destination, year	0.075 (0.022)	0.0006	−0.727 (0.028)	0.671 (0.048)
Origin-year, destination-year, occupation	0.109 (0.026)	<0.0001	−0.722 (0.028)	0.681 (0.048)
+ state pair	0.108 (0.026)	<0.0001		
+ occupation-year	**0.114** (0.026)	<0.0001		

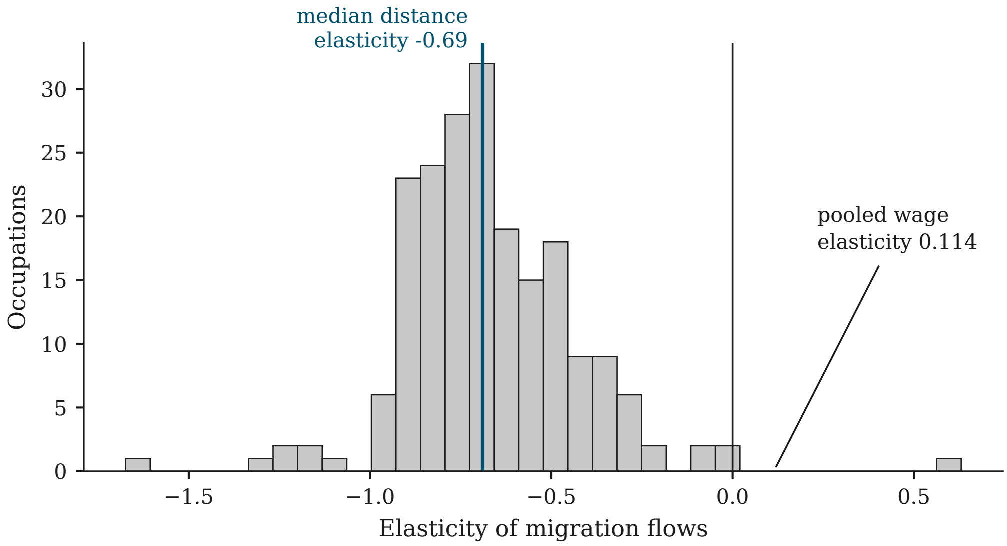

Figure 3: Distribution of the distance elasticity across 203 occupations estimated separately, with the pooled wage elasticity marked for scale.

is the linear errors-in-variables one, and the reliability behind it is measured before the fixed effects, which strip out the wage signal shared within a state and year while leaving the cell-level sampling noise where it was (Griliches and Hausman, 1986). The corrected figure therefore lies above 0.15, and the paper reports 0.114 throughout.

The distance elasticity is −0.72 in the pooled model, some six times the wage elasticity, and unlike the wage elasticity it can be estimated occupation by occupation, because it is identified from variation across 2,550 state pairs and not from movement over time. Estimated separately for the 203 occupations with at least a hundred populated corridors it has a median of −0.69, an interquartile range of −0.81 to −0.53, and a central ninety per cent running from −0.96 to −0.29. The friction therefore varies about threefold across the body of the distribution, and it is distinguishable from zero in 93 per cent of occupations. The tails beyond that range, where editors come out near zero and flight attendants near −1.7, have standard errors wide enough that they are better read as noise than as extremes.

Within the body of the distribution, a coefficient near −0.3 belongs to an occupation recruited nationally, while one near unity belongs to an occupation for which recruitment beyond the neighbouring states returns very little. Figure 3 places the pooled wage elasticity against that distribution.

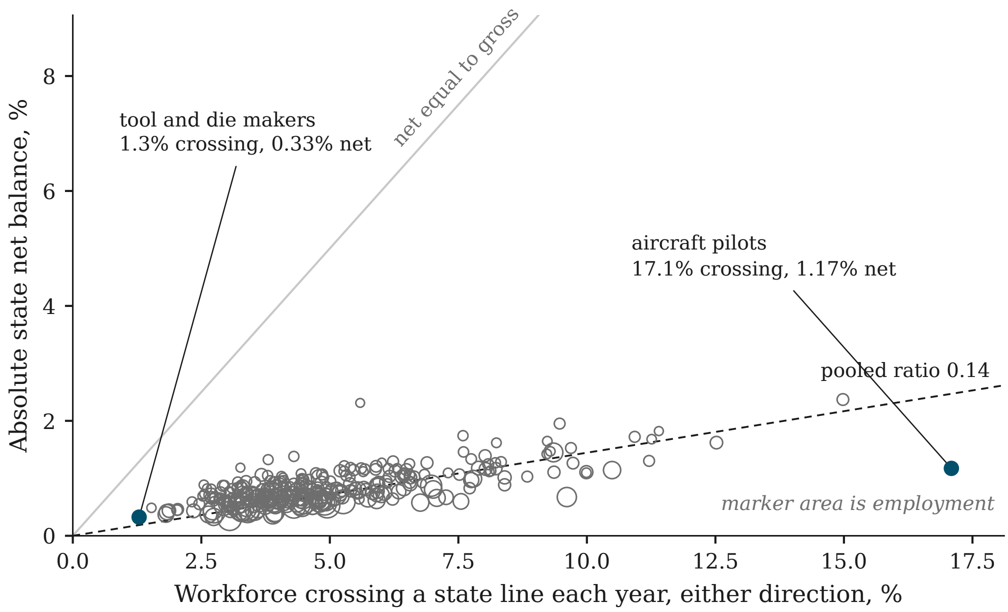

Figure 4: Annual crossing rate against the employment-weighted mean absolute net balance across states, one point per occupation with at least ten state cells, marker area proportional to employment. The grey diagonal is where every arrival would add to the stock; the dashed line is the pooled ratio of net to gross.

Turnover and net balance

Occupations differ twentyfold in the share of their workforce crossing a state line each year in either direction, from 1.29 per cent for tool and die makers to 25.52 per cent for avionics technicians, and very little of that traffic accumulates on either side. Nationally it cannot accumulate, since arrivals and departures are the same people, so the question has to be asked at state level. Across 9,775 state and occupation cells the absolute net balance comes to 14.5 per cent of the gross traffic, and that ratio falls as the traffic rises, from 24 per cent in the least mobile quartile of occupations to 15 per cent in the most mobile.

Aircraft pilots move 17.1 per cent of their workforce across state lines each year while the mean absolute state balance in them is 1.2 per cent, under a tenth of the gross, and aircraft mechanics move 15.0 per cent against a mean absolute balance of 2.4. Figure 4 shows the pattern across all occupations, with the net balance held to a narrow band well below the gross at every level of movement. Sampling noise can only add to the absolute net balance, so the 14.5 per cent is an upper bound on how much of the traffic actually accumulates.

The ratio describes turnover that has occurred and is uninformative about how departures would respond to more arrivals. In the occupations where movement is heaviest, migration changes the stock little, because arrivals and departures run close together in every state. At the bottom of the distribution too few workers move for their balance to register on the stock at all.

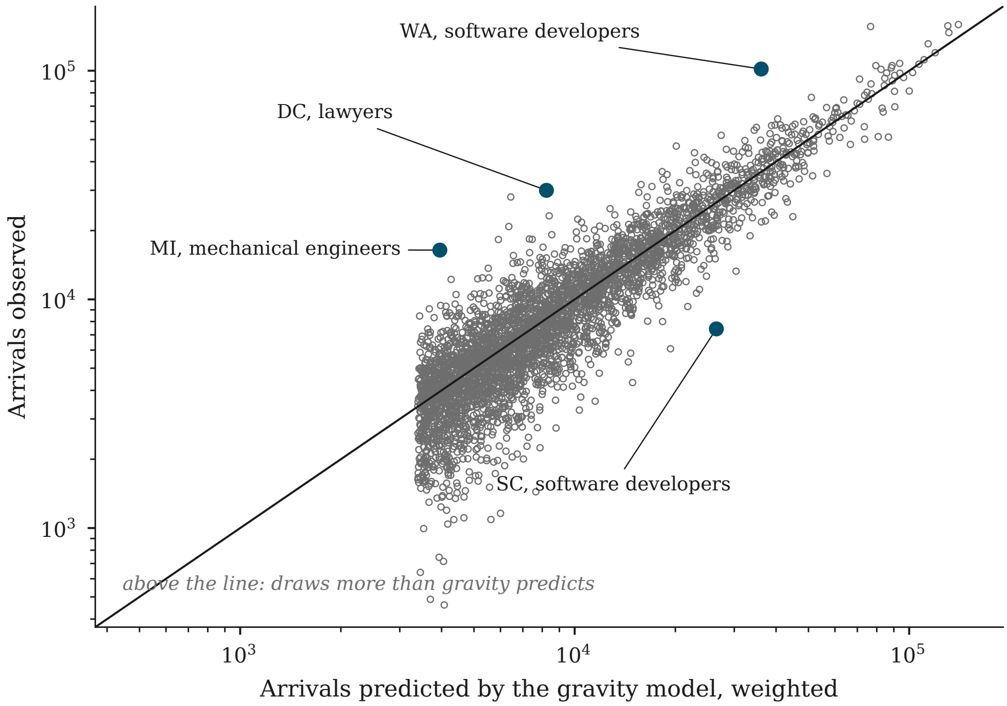

Figure 5: Observed against gravity-predicted arrivals by state and occupation, pooled 2005 to 2024, log scales. Cells below roughly thirty expected sampled movers are suppressed.

6 Departures from the benchmark

The benchmark for this section is the second specification of Table 1: origin-by-year, destination-by-year, and occupation effects, with distance, contiguity, and the wage differential. It absorbs how attractive each state is in each year and how mobile each occupation is nationally, and it deliberately omits how large each occupation is in each state, because that state by occupation cell is what this section sets out to examine. The residual is the intake of an occupation over and above what the state's general attraction and the occupation's general mobility predict, so it measures how concentrated an occupation is in the migration flow, recovered from the flows and from nothing else. Figure 5 plots observed against predicted arrivals.

The concentrations it recovers are the familiar ones. Washington state draws 2.82 times its predicted intake of software developers and California 2.03, the District of Columbia 3.64 times its predicted intake of lawyers, Michigan 4.16 times its predicted intake of mechanical engineers, and New York 2.76 times its predicted intake of securities and financial services agents. Under-supply appears with equal clarity, with South Carolina, Tennessee, and Kentucky drawing between 0.28 and 0.44 of their predicted intake of software developers. California is the instructive case, since it draws twice its predicted software developers and half its predicted home health aides, which is the pattern a high cost of living would produce by sorting arrivals on what they earn.

Deviations are judged on the Poisson scale of the sampled count: the expected flow in each cell is rescaled from weighted to sampled movers, and the ratio of actual to expected has a standard deviation of one over the square root of that count. That scale ignores the clustering in the survey design, whose effect is modest for counts spread across as many clusters as these are. The ranking uses it because a ranking on the raw ratio surfaces only small-cell artefacts: the standard deviation of the ratio falls from 0.39 in cells below 5,000 expected arrivals to 0.20 in cells above 50,000. On the noise-adjusted measure, 43 per cent of cells exceed a z of 2, against the 5 per cent that sampling variation alone would produce.

The concentrations are also stable over the period. The residual was aggregated into four multi-year windows on a balanced panel, and the correlation between the 2005 to 2009 and the 2021 to 2024 positions is 0.543, rising to 0.68 once corrected for the Poisson noise in each window, whose reliability is 0.80. A fifth of the 678 cells move between the first window and the last by more than twice their sampling variation, and the noise-corrected spread of the residual is flat across the four windows, at 0.263, 0.276, 0.257 and 0.275.

Most cells therefore hold their position across twenty years that include a financial crisis, a pandemic, and a rapid expansion of remote work, while a fifth of them move and the dispersion of concentration neither widens nor narrows.

7 Validation

Every measure above is built from household survey records, and the Occupational Employment and Wage Statistics is collected from employers, so agreement between the two is informative in a way that a comparison within a single source is not. Ten hypotheses were fixed in advance and all ten are reported, with p values corrected across the battery by the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). Regressions are weighted by employment and clustered by occupation and by state.

Table 2 reports all ten. Three hold after correction, and the strongest of them is the gravity residual against the location quotient: one standard deviation of the residual is associated with a location quotient 13 per cent higher, with an adjusted p of 4.4×10^{-7}. That is the result shown in Figure 6. The residual measures concentration in the flow of workers and the location quotient measures concentration in the stock, so the comparison sets flows built from household records against a stock measured from employers. The two are far from the same quantity, with a correlation of 0.375 across the 1,706 matched cells, and they agree. External dependence is associated with subsequent employment growth at 0.044 log points per standard deviation, and net inflow at 0.029.

None of the wage tests finds a relationship, whether the predictor is the gravity residual, external dependence, or net inflow. The relationship with five-year wage growth is -0.0004, -0.0014 and $+0.0001$ log points per standard deviation respectively, against standard errors between 0.004 and 0.007, which is enough precision to treat all three as estimated zeros. A state drawing materially fewer workers than the benchmark predicts sees no faster wage growth, so the residual cannot be read as a shortage signal and is not presented as one here.

A tenth result points the same way. States paying above the national rate for an occupation import very slightly fewer of them, wrong-signed against the prediction and not significant after correction, which is the sign a cost of living effect would produce.

Table 2: The validation battery in full. Outcomes are from the Occupational Employment and Wage Statistics: wage and employment growth are log changes from 2019 to 2024, the location quotient is in logs for 2024, and external dependence is the arrival share pooled over 2019 to 2024. Predictors are from the migration panel and standardised, so each coefficient is the change in the outcome per standard deviation of the predictor. Cells are matched through the OCC2010 to SOC crosswalk where the modal SOC code holds at least 80 per cent of the occupation; regressions are weighted by employment and clustered by occupation and by state; adjusted p values are Benjamini-Hochberg across the ten.

Outcome	Predictor	Sign	Coefficient (s.e.)	p	Adj. p	Cells	Held
Wage growth	Gravity residual	$-$	-0.0004 (0.0043)	0.920	0.985	1,706	
Employment growth	Gravity residual	$+$	$+0.0166$ (0.0070)	0.022	0.054	1,706	
Log location quotient	Gravity residual	$+$	$+0.1343$ (0.0208)	<0.0001	<0.0001	1,706	yes
Wage growth	External dependence	$+$	-0.0014 (0.0074)	0.847	0.985	6,793	
Employment growth	External dependence	$+$	$+0.0440$ (0.0143)	0.0035	0.012	6,793	yes
Log location quotient	External dependence	$+$	$+0.0364$ (0.0190)	0.062	0.088	6,793	
Employment growth	Net inflow	$+$	$+0.0293$ (0.0080)	0.0006	0.0030	6,793	yes
Wage growth	Net inflow	$-$	$+0.0001$ (0.0037)	0.985	0.985	6,793	
Employment growth	Gross turnover	$+$	$+0.0361$ (0.0162)	0.031	0.062	6,793	
External dependence	Relative wage, 2019	$+$	-0.0025 (0.0012)	0.051	0.084	6,793	

8 Limitations

Occupation and age are recorded at the time of survey, after any move has taken place. An arrival counted in an occupation either practised it before moving or entered it on arrival, and the data cannot separate the two cases. For occupations with long training requirements the ambiguity is minor, while for easily entered occupations the sourcing measure is an upper bound on imported trained workers, and the claims made here are restricted accordingly. The same timing affects the gravity model, since the origin wage of the destination occupation stands for what a mover gave up, which it does not do for a worker who changed occupation on moving.

Coverage across state and occupation cells is uneven. The results are secure nationally, by occupation, and for large combinations of state and occupation, but only 4,026 of 24,102 state by occupation cells reach thirty arrivals and thirty departures even when all nineteen vintages are pooled. A complete lookup across every state and every occupation is therefore not available from these data and is not offered here, and cells below that threshold are suppressed and not interpolated. The absent 2020 vintage makes any calculation spanning 2019 to 2021 a two-year change.

The sourcing measures come with design-based standard errors. The gravity model has errors clustered by state pair, and its residual is judged against Poisson variation in sampled counts, so neither of these accounts for the full survey design.

Nothing estimated here identifies a causal effect: the wage elasticity is an equilibrium association, the residual locates where a state diverges from its benchmark without explaining why, and the association between migration and employment growth runs in a direction these data cannot determine.

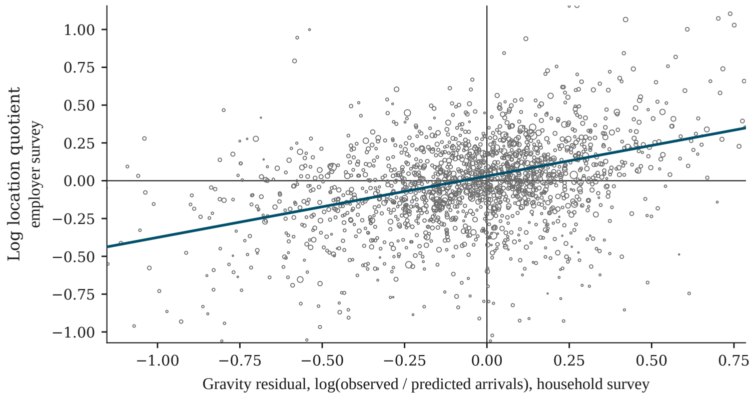

Figure 6: The gravity residual, from household survey microdata, against the location quotient, from the establishment survey. Marker area proportional to employment; the line is the employment-weighted fit.

9 Conclusion

Interstate migration supplies a share of each occupation's workforce that varies close to fifteenfold across occupations, is small in every one of them, and consists of workers younger than those already in place. The measures constructed here put numbers on those facts, report them with the uncertainty the survey design implies, and agree with an independent employer survey wherever agreement is testable.

For a state deciding whether to train an occupation or to recruit it, the results that bear most directly concern licensure, pay, and age. Licensure, the barrier the mobility literature most often names, explains almost none of which occupations import their workers: physicians sit near the top of the ranking and postal workers near the bottom, and across 421 occupations the credentialling rate has a rank correlation of 0.10 with external dependence. Pay is weakly associated with flows and distance strongly. A ten per cent wage advantage goes with about one per cent more arrivals, the distance elasticity is six times the wage elasticity, and the wage tests against outside data find nothing. A pay premium of any size a state is likely to offer therefore goes with a change of a few per cent in its intake of an occupation. Arrivals are younger everywhere, since none of the 82 occupations large enough to measure draws arrivals older than the workforce they join. In the occupations where movement is heaviest, arrivals and departures run close together in every state, so migration changes the stock little.

The deviations from the gravity benchmark are large, stable across two decades, and confirmed by an independent employer survey, and the measures locate them without accounting for them. Occupational licensure compacts, which take effect between specific pairs of states on specific dates, supply variation of the kind that could account for part of them, and that is the direction the work takes next.

References

James E. Anderson and Eric van Wincoop. Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192, 2003.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate. *Journal of the Royal Statistical Society: Series B*, 57(1):289–300, 1995.

Simone Bertoli and Jesús Fernández-Huertas Moraga. Multilateral resistance to migration. *Journal of Development Economics*, 102:79–100, 2013.

Zvi Griliches and Jerry A. Hausman. Errors in variables in panel data. *Journal of Econometrics*, 31 (1):93–118, 1986.

Keith Head and Thierry Mayer. Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of International Economics*, volume 4, pages 131–195. Elsevier, 2014.

John Kennan and James R. Walker. The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251, 2011.

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

Andrei Rogers and Luis J. Castro. *Model Migration Schedules*. International Institute for Applied Systems Analysis, Laxenburg, 1981.

Steven Ruggles, Sarah Flood, Matthew Sobek, et al. IPUMS USA: Version 15.0 [dataset], 2024.

J. M. C. Santos Silva and Silvana Tenreyro. The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658, 2006.