

ANALYTICS CASE STUDY — FLAGSHIP

Predicting Workforce Stability Among Post-9/11 Veterans

A data-driven analysis for employers. Built on ACS PUMS 2019–2023, BLS LAUS, BEA Regional, VA NCVAS, and the O*NET Military Crosswalk. 178,715 person-year records, PWGTP-weighted, fully reproducible from a single command.

Analyst

Patrick Neil Bradley

Cohort

Post-9/11 veterans, ages 22–64

Survey years

2019 · 2021 · 2022 · 2023

Build

v0.13 · 2026-04-17

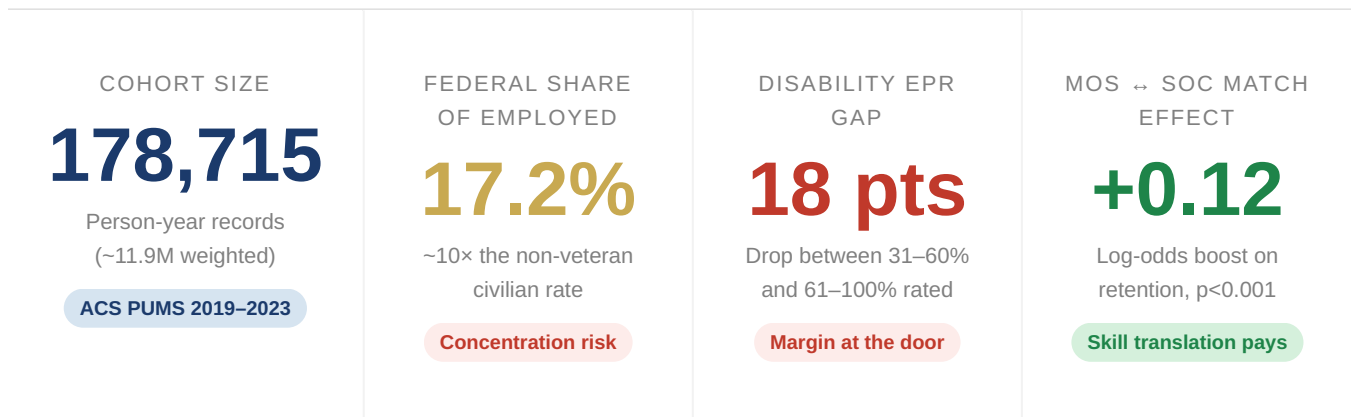
Why this case study exists

Most veteran-employment reporting stops at the headline unemployment rate. That number describes the average veteran. It does not describe the post-9/11 cohort specifically, it does not describe where in the labor market they sit, and it does not describe what happens when a service-connected disability enters the picture.

This case study closes those gaps. It uses public, individually de-identified federal data to build a reproducible picture of where post-9/11 veterans work, how state labor markets shape their outcomes, where the disability-employment margin actually lives, and what happens when you fit a transparent model to the question of who stays on the job.

The goal is to model post-9/11 veteran workforce stability the way a careful practitioner would — grounded in public data, honest about what is observed versus synthetic, and fully reproducible from a single command.

Retention is modeled with a synthetic target because no public longitudinal retention series exists for this cohort. The synthetic label is generated by a fixed-seed logistic rule and is disclosed in the model card. Every descriptive rate on this page is observed, not modeled.



CHAPTER ONE

The cohort in one paragraph

Before any finding, establish who we are looking at and how we weight them.

The analysis universe is post-9/11 veterans aged 22–64 across the 50 states and DC, pooled across ACS PUMS 1-year survey rounds for 2019, 2021, 2022, and 2023. Active-duty rows (ESR 4 / 5) are removed from civilian rates. Every rate on this page is PWGTP-weighted unless marked otherwise; PWGTP is the Census Bureau's person weight, calibrated to population totals. The pooled cohort is 178,715 person-year records (11.9 million weighted person-years) after Build v0.11 categorical-range guards.

2020 is intentionally excluded. ACS published a smaller experimental file that year, and pooling it with the standard 1-year files introduces weighting inconsistencies. The final

cohort distributes across years as 36,714 (2019), 43,458 (2021), 47,928 (2022), and 50,615 (2023).

CHAPTER TWO

Where veterans work is structurally different

Occupation tilt is the first large, clean signal in the data.

Relative to non-veterans in the same ACS rounds, post-9/11 veterans are overrepresented in a handful of occupational families and underrepresented in a different handful. The ratio is the simplest way to see it:

Occupation family	Veteran share / non-vet share
Protective Service	4.15×
Installation / Maintenance / Repair	2.40×
Computer & Mathematical	1.84×
Healthcare Support	0.40×
Food Preparation & Serving	0.42×
Education, Training & Library	0.55×

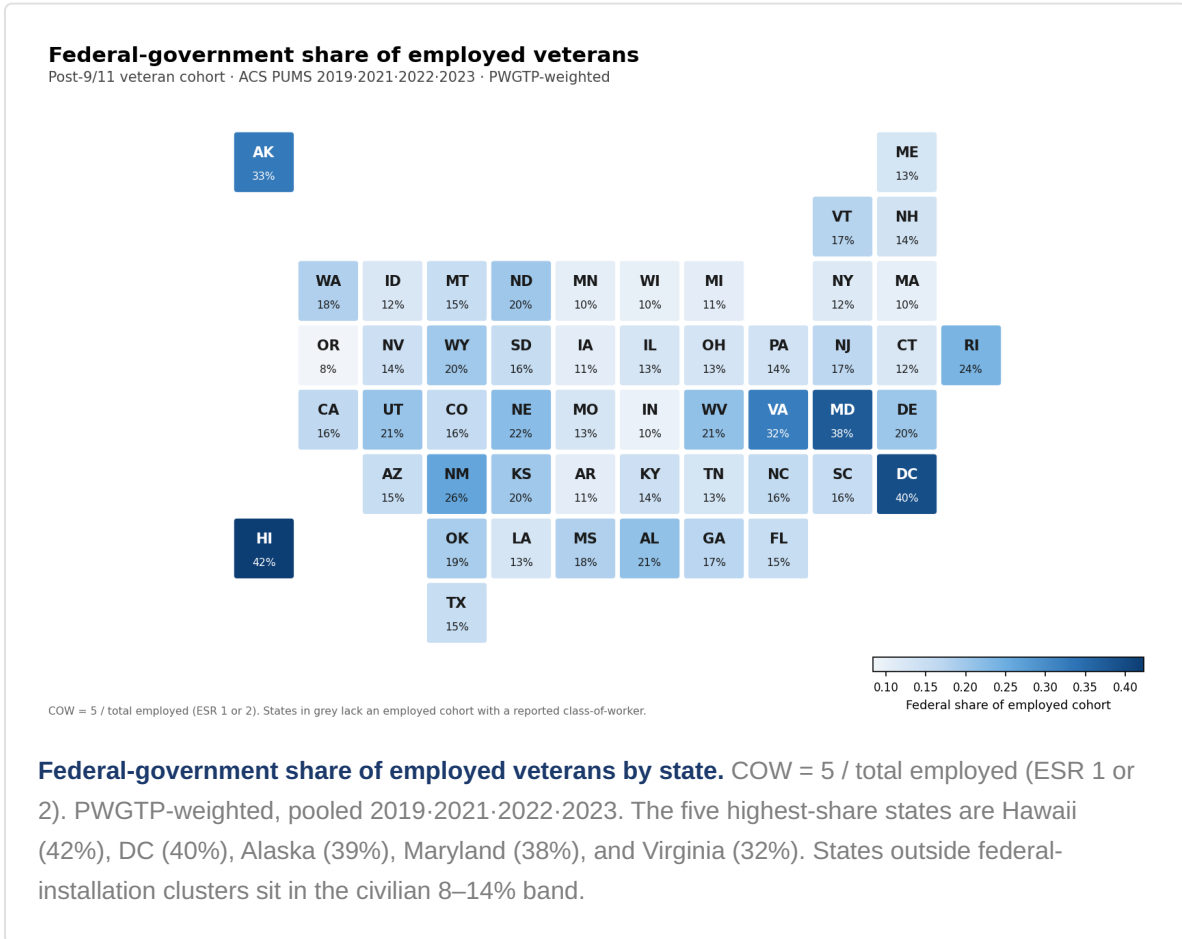
Public Administration leads the industry tilt at 4.03× the non-veteran rate — the federal-civilian cluster that Chapter Three unpacks. Utilities sit at 1.94×. These are not small departures. They translate into labor-market exposure that does not look like the national average in either direction.

CHAPTER THREE

The federal-civilian corridor

Federal employment is where post-9/11 veterans land — concentrated, not dispersed.

17.2% of the employed cohort works directly for the federal civilian government. That is roughly ten times the non-veteran civilian rate. But the share is not spread evenly across the country. It concentrates sharply around federal installations and the National Capital Region.



Employer implication.

Any private employer recruiting in Hawaii, Alaska, or the National Capital Region should model federal employment as a direct competitor for this cohort, not as a background condition.

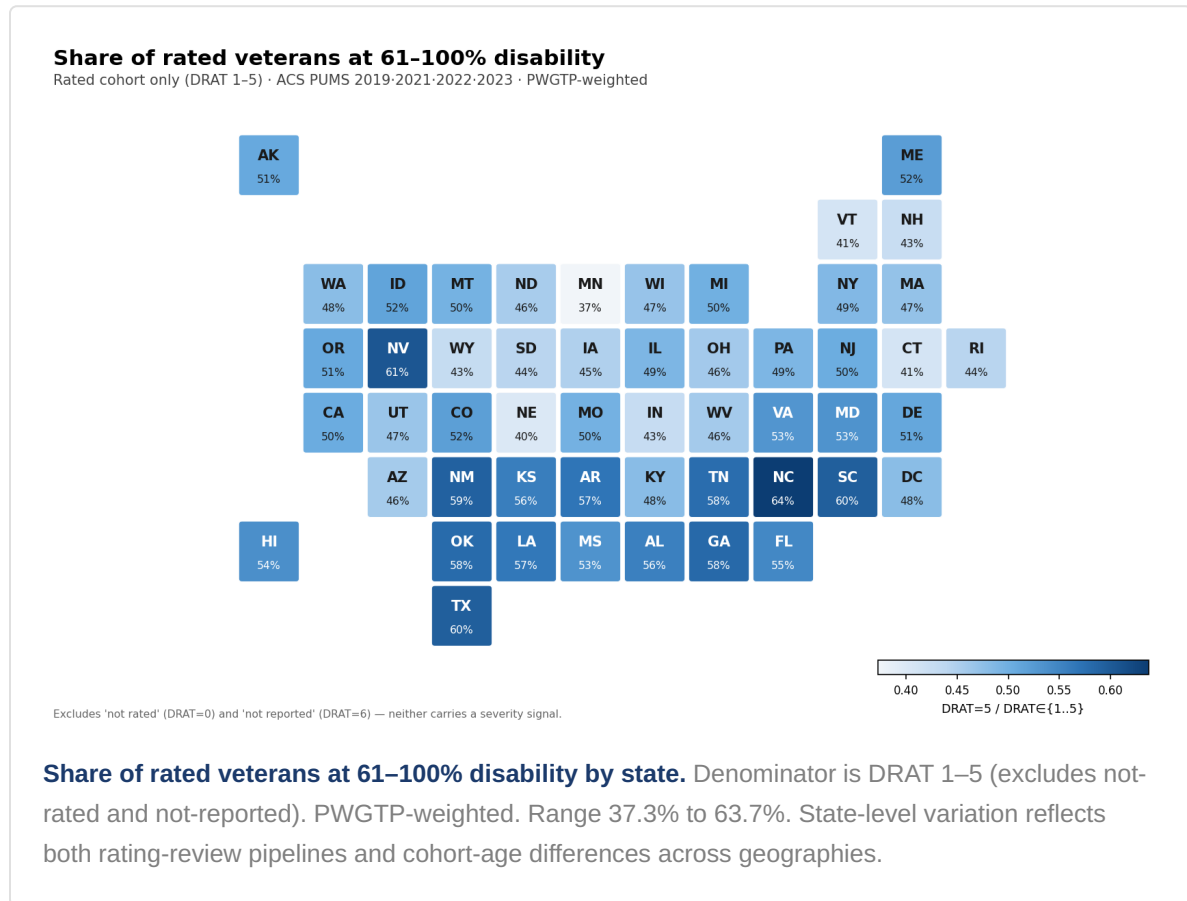
CHAPTER FOUR

Disability rating: the margin is at the door

The employment-to-population gap for rated veterans is large, and it is a hiring margin — not a pay margin.

Between the 31–60% and 61–100% disability rating bands, the employment-to-population ratio falls by 18 percentage points. That is a sizeable cliff. But once inside the employed universe, earnings and hours are nearly flat across disability bands — median wages sit in the \$63,000–\$66,000 range, median weekly hours at 40, regardless of rating severity.

That pattern matters because it tells you where the friction lives. Disability rating predicts *entry*, not conditional earnings. The policy and employer lever is at the hiring door.



CHAPTER FIVE

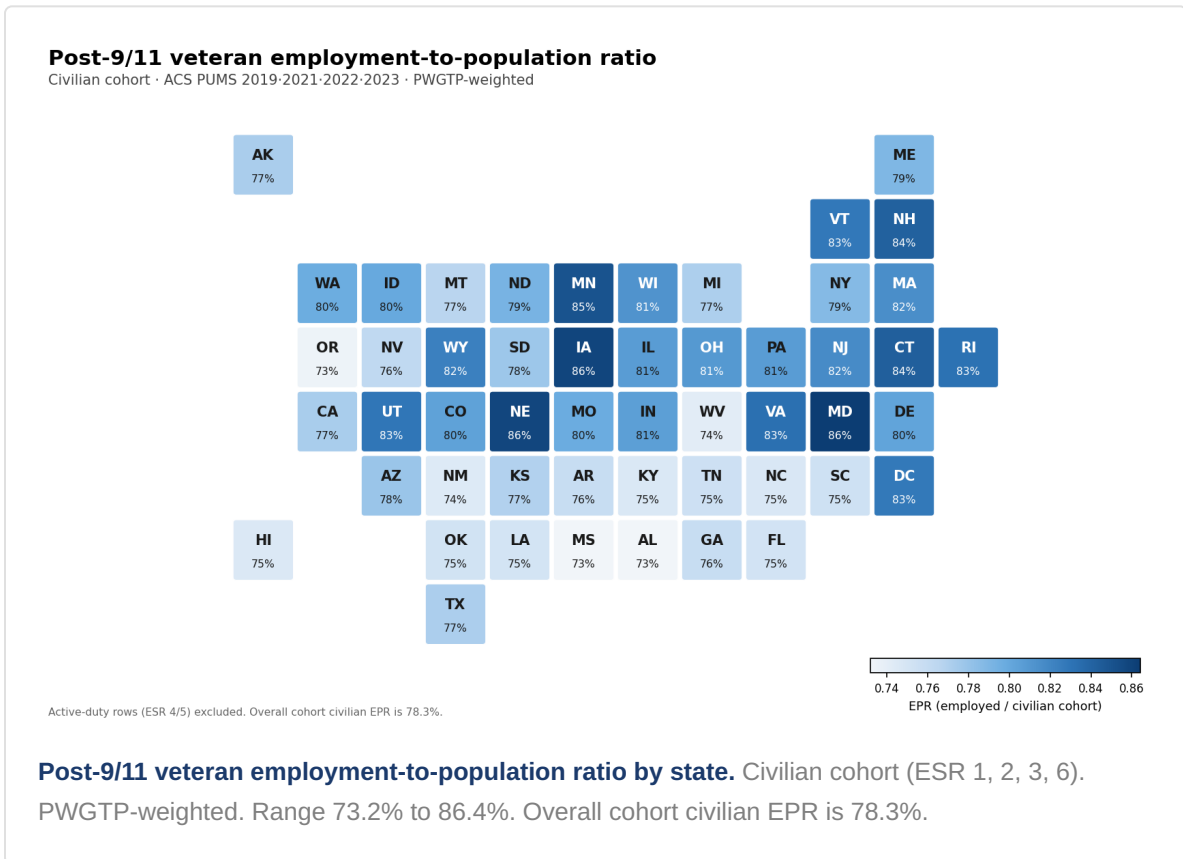
Geography has structure, not noise

State-level labor market conditions correlate cleanly with cohort outcomes.

Three state-level correlations describe most of the geographic structure:

Correlation	Coefficient (r)
State unemployment rate ↔ cohort EPR	-0.30
State per-capita income ↔ cohort median wage	+0.66
State GDP industry concentration (HHI) ↔ cohort EPR	+0.17

Tight state labor markets lift cohort employment; rich state economies lift cohort wages; concentrated state industry mixes correlate weakly but positively with cohort employment, likely because concentration tracks public-administration-heavy states.



CHAPTER SIX

The illustrative retention model

Modeled — Illustrative

A transparent logit on the synthetic target closes the workflow from data to model.

The retention target here is synthetic — generated by a fixed-seed logistic rule over observable features because no public longitudinal retention series exists for this cohort. Fitting a model to it demonstrates the end-to-end workflow (feature engineering, stratified split, weighted fit, calibration) without exposing any real employer record or identifiable work history.


Fit at a glance

Item	Value
Family	Binomial GLM, logit link, IRLS
Training rows	76,145
Features	47
Split	70 / 15 / 15 train / val / test, seed 20260417
Weights	PWGTP passed as <code>freq_weights</code>
Test AUC	0.566
Test Brier score	0.218
Calibration	Within ± 3 pp across all ten weighted deciles

AUC is modest by design. The synthetic target is dominated by a small number of observable signals — disability rating, education floor, and state unemployment — and the residual variance is large relative to those signals. The point of the model is not headline discrimination; it is to demonstrate directional coherence with the descriptive findings.

Illustrative retention model — logit coefficients

Post-9/11 veteran cohort · ACS PUMS 2019-2021-2022-2023 · n=76,167 train, PWGTP-weighted · synthetic retention target



Categorical references: education = HS / GED, disability = not rated, race = white non-Hispanic, occupation = Management. Two small-sample occupation dummies (Military civilian-reported, n=11; Unknown, n=139) are excluded.

Illustrative retention model — logit coefficients with 95% CIs. Grouped by feature family. Categorical references: education = HS / GED, disability = not rated, race = white non-Hispanic, occupation = Management. Two low-sample occupation dummies (Military civilian-reported, n=11; Unknown, n=139) are excluded from the chart — their coefficients were unstable.

What the coefficients say ($p < 0.001$ on every line)

Disability rating 50–60% and 70–100% both pull retention down sharply relative to the non-rated reference. This is the modeled expression of the EPR cliff in Chapter Four. The MOS ↔ SOC skill-match feature — a Build v0.12 indicator that at least one military occupational specialty maps to the veteran's civilian SOC major — enters as a clean positive. Veterans whose service job translates to their civilian job retain at higher rates. State unemployment enters negatively; tight labor markets hold cohort retention up. Occupation-family contrasts against Management are directionally consistent with the descriptive patterns: farming / fishing / forestry strongly positive, personal care & service and food preparation & serving strongly negative. Education at the low end (no HS diploma) is meaningfully negative; the upper end washes out through occupation and wages.

What this model does not claim.

It does not make causal claims — the target is synthetic. It does not produce survey-design-correct standard errors. It does not support individual-level prediction. Full methodology, metrics, calibration deciles, and ethical notes are in the model card below.

CHAPTER SEVEN

What this study does and does not claim

A limitations section that travels with the build log, not one written from memory at the end.

The synthetic retention outcome is the biggest explicit limitation. No real longitudinal retention series for this cohort exists in public data; the analysis demonstrates the workflow that *would* run against a real retention outcome, using a generated one in its place. Every chart or table that uses the synthetic target carries the **Modeled — Illustrative** label.

ACS PUMS carries its own boundaries. It is cross-sectional, survey-based, and subject to Census disclosure-avoidance perturbation. It cannot distinguish GI Bill from VR&E beneficiaries at the individual level. Federal regional datasets (BEA SAGDP2, BLS LAUS) cover 2018 onward only, which determined the pooled-year window.

Two Build v0.12 components are currently blocked on external data-access paths: the O*NET work-context composites (physical demand, schedule variability, autonomy) and the

Tier-2 eight-digit OCCP → SOC crosswalk. Both are scaffolded and unit-tested; neither changes the synthetic-target story. Full limitations live in Section I of the case study source document.

ARTIFACTS

Download and further reading

Executive summary (two-pager)

Thesis, four findings, federal-share tilemap, methodology note. Word document, US Letter.

[DOWNLOAD .DOCX](#)

Model card

Full methodology, metrics table, calibration deciles, feature list, ethical notes.

[READ MODEL CARD](#)

Federal-share tilemap (PNG)

Federal-government share of employed veterans by state.

[VIEW CHART](#)

EPR tilemap (PNG)

Post-9/11 veteran employment-to-population ratio by state.

[VIEW CHART](#)

Severe-DRAT tilemap (PNG)

Share of rated veterans at 61–100% disability by state.

[VIEW CHART](#)

Feature importance chart (PNG)

Illustrative retention model — logit coefficients with 95% CIs.

[VIEW CHART](#)