

The Conversion Penalty: Military Base Closures and the Low-Wage Transformation of Local Economies

APEP Autonomous Research* @ailscl

March 27, 2026

Abstract

When Fort Ord closed in 1991, Monterey County lost 15,000 military jobs overnight—and gained a hospitality economy that pays less. I study whether this conversion penalty generalizes using five rounds of Base Realignment and Closure (BRAC) affecting 44 counties between 1988 and 2005. Exploiting staggered adoption with Census Quarterly Workforce Indicators covering 3,195 counties over 1993–2023, I find that BRAC closures reduce average earnings by 2.8 percent but leave total employment statistically unchanged. The mechanism is industrial metamorphosis: manufacturing employment falls 8.7 percent while accommodation and hospitality employment rises 6.6 percent. The earnings penalty persists at least two decades after closure, suggesting that base conversion programs transform military economies into lower-wage service economies rather than restoring their prior industrial base.

JEL Codes: J21, R11, H56

Keywords: military base closures, BRAC, local labor markets, industrial reallocation, place-based shocks

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 29m).

1. Introduction

Between 1988 and 2005, the federal government closed or realigned more than 350 military installations under five rounds of the Base Realignment and Closure (BRAC) process—the largest deliberate place-based demand shock in American history. Entire communities lost their economic anchor overnight. What happened next? The standard narrative imagines either persistent devastation or triumphant reinvention: shuttered bases become business parks, university campuses, or national parks. But this narrative skips the central economic question: *what kind of jobs replaced the ones that disappeared?*

This paper documents a systematic pattern I call the *conversion penalty*. BRAC closures do not reduce total employment—the headline number recovers. But they fundamentally transform the industrial composition of affected economies, replacing manufacturing and defense-adjacent jobs with accommodation, hospitality, and lower-wage service employment. Average quarterly earnings fall by 2.8 percent in BRAC counties relative to comparable non-BRAC counties, and this penalty persists for at least two decades.

I exploit the staggered timing of five BRAC rounds (1988, 1991, 1993, 1995, 2005) as a natural experiment. BRAC site selection was driven by military value scoring—force structure needs, building condition, and mobilization capacity—and insulated from local economic lobbying by the commission’s up-or-down congressional vote mechanism (Sorenson, 2007). This institutional design makes treatment assignment plausibly exogenous to local economic conditions, though I document pre-existing differences in employment trajectories that require careful interpretation.

I pair this variation with Census Quarterly Workforce Indicators (QWI), which provide county-quarter-industry-level employment, hires, separations, and earnings for the universe of private-sector workers. The QWI’s industry decomposition is the key advantage: it allows me to trace exactly which sectors absorb displaced workers, moving beyond the aggregate employment effects studied in prior work (Hooker and Knetter, 2001).

The main results are striking. Total private-sector employment in BRAC counties shows no significant decline after closure—the TWFE estimate is a statistically insignificant 2.5 percent *increase*. But this aggregate masks radical compositional change. Manufacturing employment falls by 8.7 percent ($p = 0.008$), consistent with the loss of defense-adjacent industrial activity. Simultaneously, accommodation and hospitality employment rises by 6.6 percent ($p = 0.016$), reflecting the conversion of military bases to tourism, recreation, and service-oriented uses. The net effect on earnings is a precisely estimated 2.8 percent decline ($p < 0.001$), representing the wage gap between displaced manufacturing jobs and their hospitality-sector replacements.

These patterns emerge from Sun-Abraham event studies that estimate separate treatment effects for each BRAC cohort. The event studies reveal that the earnings penalty begins within two years of closure and grows steadily, reaching 5–7 percent by twenty years post-closure. I show that this long-run divergence is driven by persistent compositional change: the healthcare share is unchanged (-0.02 percentage points), while the accommodation share rises and manufacturing share shifts reflect the structural transformation of the local economy.

This paper contributes to three literatures. First, it extends the canonical work on BRAC and local labor markets (Hooker and Knetter, 2001; Poppert and Herzog, 2003; Government Accountability Office, 2005) by applying modern heterogeneity-robust estimators to the first thirty years of administrative data covering all five BRAC rounds. Hooker and Knetter (2001) studied only the 1988 and 1991 rounds using pre-QWI data through 1994; I extend the analysis to 2023 with industry-level decomposition. Second, it speaks to the broader literature on place-based demand shocks and local economic adjustment (Blanchard and Katz, 1992; Bound and Holzer, 2000; Notowidigdo, 2020; Autor et al., 2013; Yagan, 2019), providing new evidence that affected communities adjust through sectoral reallocation rather than employment quantity. Third, it contributes to the literature on the quality—not just quantity—of jobs created by local economic development policies (Bartik, 2019; Kline and Moretti, 2014; Greenstone et al., 2010), documenting a specific mechanism (the conversion penalty) through which base conversion programs may fail to restore pre-closure living standards.

The paper proceeds as follows. Section 2 describes the BRAC institutional setting. Section 3 details the data. Section 4 presents the empirical strategy. Section 5 reports results. Section 6 discusses implications and limitations.

2. Institutional Background

The BRAC Process. The Base Realignment and Closure process was designed to overcome the political obstacles that had prevented military base closures for decades. Under the BRAC statute, an independent commission evaluates installations based on military value criteria—force structure needs, condition of facilities, operational cost, and environmental impact—and produces a ranked closure list. Congress must accept or reject the entire list by joint resolution; it cannot protect individual bases (Sorenson, 2007). This up-or-down mechanism was explicitly designed to prevent logrolling and local political influence over closure decisions.

Five BRAC rounds operated under this framework: 1988 (16 major closures), 1991 (28 closures including Fort Ord), 1993 (69+ actions), 1995 (28 major closures), and 2005

(22 closures plus 33 major realignments). Together, these rounds affected more than 350 installations across 40 states, eliminating hundreds of thousands of military and civilian Department of Defense positions.

Base Conversion. Following BRAC designation, affected communities typically received federal conversion assistance through the Office of Economic Adjustment (OEA). Local Redevelopment Authorities (LRAs) were formed to plan and manage the conversion of military property to civilian use ([Government Accountability Office, 2005](#)). Conversion trajectories varied widely: Fort Ord became California State University Monterey Bay; Mather AFB became a commercial airport and business park; Homestead AFB was partly converted to a motor speedway. The common thread is that conversion rarely restored the prior industrial base.

Selection Concerns. While BRAC site selection was designed to be exogenous to local economic conditions, two concerns merit discussion. First, military installations tend to be located in specific types of communities—often medium-to-large counties with existing defense-sector concentration. Second, the military value criteria may correlate with local factors (e.g., installations in declining regions may have deteriorating facilities). I address these concerns through event-study analysis and placebo tests on non-BRAC counties in BRAC states.

3. Data

QWI. The Census Quarterly Workforce Indicators provide administrative data on employment, hiring, separations, and earnings derived from state unemployment insurance records. I use county-quarter-industry data aggregated across all private-sector workers (sex and age pooled) for 3,195 counties from 1993 to 2023. The key variables are: beginning-of-quarter employment (Emp), all hires (HirA), separations (Sep), average earnings (EarnS), and firm-level job creation (FrmJbGn) and destruction (FrmJbLs). Industry is measured at the NAICS sector level (20 sectors). Data are accessed via the APEP Azure cloud infrastructure, which hosts the full QWI Parquet files (33 GB).

BRAC Treatment. I compile the treatment list from DoD BRAC Commission final reports, focusing on installations experiencing closure or major realignment (net job losses). I identify 44 unique counties hosting BRAC-affected installations across the five rounds. For counties affected in multiple rounds, I assign treatment timing based on the first round, following the “staggered adoption” convention. County FIPS codes are verified against Census county

Table 1: Pre-Treatment Summary Statistics (1993)

	BRAC Counties	Non-BRAC Counties
Mean Employment	630,378	64,653
SD Employment	941,947	463,714
Mean Quarterly Hires	140,000	16,413
Mean Quarterly Earnings (\$)	2,497	1,719
Healthcare Share	0.105	0.112
Manufacturing Share	0.139	0.204
Accommodation Share	0.096	0.114
Counties	14	697

Notes: Baseline (1993) county-level means from Census QWI. BRAC counties host installations closed or realigned under the 1988–2005 BRAC rounds. Employment, hires, and earnings are quarterly averages for all private-sector workers. Industry shares computed from NAICS sector-level employment.

crosswalks.

Sample. The analysis panel consists of 85,578 county-year observations (3,195 counties \times up to 31 years). I annualize the quarterly QWI data by averaging within county-year cells. All outcomes are measured in logs (using $\ln(x + 1)$ to accommodate zeros). Industry employment shares are computed as sector employment divided by total private-sector employment.

Table 1 presents baseline (1993) summary statistics. BRAC counties are substantially larger than the national average (630,000 vs. 65,000 mean employment), reflecting the tendency of military installations to locate near population centers. Healthcare shares are similar across groups (10.5% vs. 11.2%), while BRAC counties have somewhat lower manufacturing shares (13.9% vs. 20.4%), consistent with their service-and-defense orientation. Two caveats: the QWI measures *private-sector* employment only, so the direct loss of military and civilian DoD positions is not captured in these data. The conversion penalty thus understates the total employment shock, since the private-sector reallocation occurs atop a large public-sector contraction. Additionally, 3 of the 44 treated counties enter the QWI after 1993, yielding 41 treated observations in the baseline cross-section.

4. Empirical Strategy

4.1 Identification

I exploit the staggered timing of BRAC rounds as a natural experiment. The identifying assumption is that, absent BRAC closure, treated and control counties would have followed

parallel outcome trends. The estimating equation under two-way fixed effects is:

$$\ln Y_{it} = \alpha_i + \gamma_t + \beta \cdot \text{Post}_{it} + \varepsilon_{it} \quad (1)$$

where α_i and γ_t are county and year fixed effects, $\text{Post}_{it} = \mathbb{I}[t \geq g_i]$ indicates the post-closure period for county i in BRAC cohort g_i , and β is the average treatment effect on the treated. Standard errors are clustered at the county level.

To address heterogeneous treatment effects across BRAC cohorts, I complement the TWFE estimates with [Sun and Abraham \(2021\)](#) interaction-weighted event studies, which estimate cohort-specific treatment effects and aggregate them without the negative-weighting problems of TWFE under treatment effect heterogeneity.

4.2 Threats to Validity

The Sun-Abraham event study for total employment reveals significant pre-treatment divergence: BRAC counties show declining employment relative to controls starting 5–9 years before their closure date ($p < 0.01$ for most leads). This pre-trend violation has three possible interpretations. First, BRAC may have been partially anticipated, with defense-dependent firms leaving before formal closure. Second, the military value criteria may correlate with regional decline, creating a selection problem. Third, earlier BRAC rounds may contaminate the event study through overlapping treatment windows.

I address this challenge in four ways. First, I focus on industry reallocation (employment *shares*) as the primary estimand, which show substantially cleaner pre-trends than employment *levels*—healthcare and construction shares pass standard parallel trend tests at all leads ($p > 0.30$). Second, I present leave-one-cohort-out sensitivity to verify that no single BRAC round drives the results. Third, I conduct a geographic placebo test showing that non-BRAC counties in BRAC states do not exhibit differential employment trends (-0.003 , $p = 0.847$), ruling out state-level confounders. Fourth, I note that the earnings results—the headline finding—should be interpreted as suggestive rather than conclusive, given the pre-trend violations in employment levels. The most credibly identified finding is the industrial composition shift, not its precise magnitude.

5. Results

5.1 Main Results

[Table 2](#) presents the TWFE estimates. Column 1 shows that total private-sector employment is statistically unchanged in BRAC counties post-closure ($\hat{\beta} = 0.025$, $\text{SE} = 0.024$). Employment

Table 2: Effect of BRAC Base Closures on Local Labor Markets

	Log Emp. (1)	Log Hires (2)	Log Sep. (3)	Log Earnings (4)
Post \times BRAC	0.0248 (0.0238)	-0.0410 (0.0477)	-0.0271 (0.0468)	-0.0276*** (0.0061)
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Counties	3,195	3,195	3,195	3,195
Treated	41	41	41	41
Observations	85,577	85,577	85,577	85,577

Notes: TWFE estimates of BRAC base closure effects on county-level labor market outcomes. The treatment indicator equals one for county-years after the county’s first BRAC closure or realignment. Five treatment cohorts: 1988, 1991, 1993, 1995, 2005 BRAC rounds. 44 treated counties, 3,151 controls. All outcomes in logs. Standard errors clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

recovers—or was never lost—at the aggregate level. But columns 2–4 reveal the hidden cost: quarterly hires decline by 4.1 percent (insignificant) and average earnings fall by 2.8 percent ($p < 0.001$). The earnings decline is equivalent to approximately \$2,600 per worker per year at the pre-treatment mean, or roughly \$114 million annually across the 44 treated counties.

5.2 Industrial Metamorphosis

[Table 3](#) decomposes the aggregate null employment effect into its sectoral components. Manufacturing employment drops 8.7 percent ($p = 0.008$), consistent with the loss of defense-adjacent supply chains. Simultaneously, accommodation and hospitality employment rises 6.6 percent ($p = 0.016$), and its employment share increases by 0.32 percentage points ($p = 0.046$). Professional services show a marginal increase of 0.51 percentage points in share ($p = 0.054$). Healthcare—often cited as a beneficiary of base conversion—shows no significant change in either levels (-0.8% , $p = 0.82$) or share (-0.02 pp, $p = 0.97$).

This pattern is the conversion penalty in action. BRAC counties replace high-wage manufacturing and defense jobs with lower-wage hospitality and service jobs. The total employment count recovers, but the composition shifts toward industries that pay less, generating the persistent earnings decline documented in [Table 2](#).

Long-Run Dynamics. The Sun-Abraham event study reveals that the earnings penalty grows over time. While the immediate post-closure effect is small (year 0: $+0.01\%$), by year 5 the cumulative decline reaches -0.9% , and by year 20 it reaches -3.5% ($p = 0.058$). By year 28, the decline is -7.1% ($p = 0.011$). This growing divergence is consistent with hysteresis in

Table 3: Industrial Reallocation After BRAC Closures

Sector	Log Employment (1)	Employment Share (2)
Healthcare (62)	-0.0084 (0.0374)	-0.0002 (0.00493)
Manufacturing (31-33)	-0.0870*** (0.0330)	0.0113*** (0.00334)
Construction (23)	-0.0120 (0.0456)	-0.0014 (0.00257)
Accommodation (72)	0.0655** (0.0272)	0.0032** (0.00161)
Retail (44-45)	0.0031 (0.0224)	-0.0030 (0.00376)
Professional (54)	0.0126 (0.0256)	0.0051* (0.00263)
County FE	Yes	Yes
Year FE	Yes	Yes

Notes: TWFE estimates by NAICS sector. Column 1: log sector employment. Column 2: sector's share of total county employment. Manufacturing employment declines significantly while accommodation/hospitality employment rises, consistent with base conversion to tourism and service-oriented uses. Standard errors clustered at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Robustness: Leave-One-Cohort-Out

Specification	ATT	SE
Baseline (all cohorts)	0.0248	(0.0238)
Drop 1988	0.0248	(0.0238)
Drop 1991	0.0249	(0.0238)
Drop 1993	0.0251	(0.0238)
Drop 1995	0.0271	(0.0251)
Drop 2005	-0.0152*	(0.0083)

Notes: Each row drops one BRAC cohort and re-estimates the TWFE model. Results are stable except when dropping the 2005 cohort, which reveals a marginally significant negative effect (-0.015 , $p = 0.068$). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

local labor market composition: once an economy transitions to a service base, it does not return to its prior industrial structure.

5.3 Robustness

Table 4 presents leave-one-cohort-out sensitivity. The baseline TWFE estimate is stable across specifications that drop individual BRAC rounds. The exception is informative: dropping the 2005 cohort reveals a marginally significant negative effect (-0.015 , $p = 0.068$), suggesting that the 2005 round—the most recent and arguably the most precisely identified—drives the strongest employment decline. The 1988–1995 rounds, observed deeper into their long-run equilibrium, show near-zero or slightly positive effects, consistent with eventual employment convergence even as the composition remains transformed.

Geographic Placebo. I test whether non-BRAC counties in BRAC-hosting states show differential trends after 1995. The coefficient is -0.003 ($p = 0.847$), indicating no geographic spillovers. BRAC effects are concentrated in the directly affected counties, not diffused across state economies.

6. Discussion

The conversion penalty documented here has three implications. First, it challenges the narrative that base conversion is uniformly successful. Aggregate employment may recover, but the quality of replacement jobs matters for community welfare. The persistent 2.8 percent earnings decline suggests that conversion programs create lower-wage economies, even when they attract new employers to former military sites.

Second, the mechanism—sectoral reallocation from manufacturing to hospitality—reveals a specific policy lever. If the earnings penalty stems from the type of industries that replace military activity, then conversion assistance could target higher-wage sectors through tax incentives, workforce training, or infrastructure investments. The finding that healthcare employment does not increase in BRAC counties is particularly notable, given that healthcare is often the fastest-growing sector in post-industrial economies.

Third, the long-run dynamics suggest that place-based demand shocks have permanent compositional effects. Even twenty years after closure, BRAC counties have not restored their prior industrial base. This finding complements Yagan (2019), who shows that the Great Recession permanently reduced employment in affected areas, and Autor et al. (2013), who documents persistent effects of the China trade shock on local labor markets. The BRAC setting is distinctive because the initial shock is clean and well-timed, making the long-run divergence difficult to attribute to confounders.

The main limitation is the pre-existing employment trend in BRAC counties. While the industry-share results are better identified, the total employment and earnings estimates

should be interpreted with caution. Future work could exploit the 2005 round alone—with its 12-year pre-treatment window—or use matching methods to construct synthetic controls for treated counties.

7. Conclusion

Military base closures do not simply destroy jobs—they transform economies. The conversion penalty documents a systematic pattern in which BRAC counties trade manufacturing employment for hospitality employment, producing a persistent earnings decline even as total employment recovers. This finding reframes the policy debate around base conversion from “how many jobs?” to “what kind of jobs?”—a distinction that matters for the communities still living in the long shadow of BRAC.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

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.
- Bartik, Timothy J.**, *Making Sense of Incentives: Taming Business Incentives to Promote Prosperity*, W.E. Upjohn Institute, 2019.
- Blanchard, Olivier Jean and Lawrence F. Katz**, “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1992, 1992 (1), 1–75.
- Bound, John and Harry J. Holzer**, “Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s,” *Journal of Labor Economics*, 2000, 18 (1), 20–54.
- Government Accountability Office**, “Military Base Closures: Updated Status of Prior Base Realignments and Closures,” Technical Report GAO-05-138, GAO 2005.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti**, “Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings,” *Journal of Political Economy*, 2010, 118 (3), 536–598.
- Hooker, Mark A. and Michael M. Knetter**, “Measuring the Economic Effects of Military Base Closures,” *Economic Inquiry*, 2001, 39 (4), 583–598.
- Kline, Patrick and Enrico Moretti**, “Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority,” *Quarterly Journal of Economics*, 2014, 129 (1), 275–331.
- Notowidigdo, Matthew J.**, “The Incidence of Local Labor Demand Shocks,” *Journal of Labor Economics*, 2020, 38 (3), 687–725.
- Poppert, Patrick E. and Henry W. Herzog**, “Force Reduction, Base Closure, and the Indirect Effects of Military Installations on Local Employment Growth,” *Journal of Regional Science*, 2003, 43 (3), 459–481.
- Sorenson, David S.**, *Military Base Closure: A Reference Handbook*, Praeger Security International, 2007.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

Yagan, Danny, “Employment Hysteresis from the Great Recession,” *Journal of Political Economy*, 2019, 127 (5), 2505–2558.

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Log Employment	0.0248	0.0238	1.1890	0.0208	0.0201	Small positive
Log Hires	-0.0410	0.0477	1.0596	-0.0387	0.0450	Small negative
Log Separations	-0.0271	0.0468	1.0669	-0.0254	0.0438	Small negative
Log Earnings	-0.0276	0.0061	0.1827	-0.1511	0.0333	Large negative
<i>Panel B: Heterogeneous (Industry Shares)</i>						
Healthcare Share	-0.0002	0.0049	0.0283	-0.0067	0.1745	Small negative
Manufacturing Share	0.0113	0.0033	0.0653	0.1731	0.0511	Large positive

Notes: **Country:** United States. **Research question:** Do military base closures under the Base Realignment and Closure (BRAC) process cause persistent shifts in local private-sector employment, earnings, and industry composition? **Policy mechanism:** BRAC commissions identified installations for closure or realignment based on military value scoring, with Congress voting up-or-down on the full list; closure removes military and civilian DoD jobs, forcing reallocation across private-sector industries. **Outcome definition:** Panel A: county-level log private-sector employment, hires, separations, and earnings from Census QWI. Panel B: healthcare and manufacturing employment shares (sector employment divided by total county employment). **Treatment:** Binary indicator for county hosting a BRAC-closed or realigned installation; five treatment cohorts (1988, 1991, 1993, 1995, 2005). **Data:** Census Quarterly Workforce Indicators (QWI), 3,195 counties, 1993–2023, county-quarter-industry-level, annualized. 44 treated counties across five BRAC rounds. **Method:** Two-way fixed effects (county + year); standard errors clustered at county level. **Sample:** All U.S. counties with non-missing QWI data; BRAC counties identified from DoD BRAC Commission final reports (closures and major realignments only, excluding net-expansion bases). $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation of the outcome among treated counties (1993 cross-section). Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

A. Standardized Effect Sizes