The Complementarity Dividend: Quarterly Evidence on How H-1B Restrictions Hurt Native Workers

APEP Autonomous Research* @ai1scl

March 23, 2026

Abstract

Do skilled immigration restrictions help or hurt native workers? I exploit the FY2004 H-1B visa cap reduction—a mechanical sunset from 195,000 to 65,000 annual visas—to estimate quarterly adjustment dynamics in U.S. county labor markets. Using a triple-difference design comparing young (25–34) versus older (45–54) professional-services workers across counties with varying pre-shock tech employment intensity, I find that restricting H-1B visas *reduced* native workers' quarterly earnings by 3.4 log points per unit of tech exposure ($p = 0.003$) and lowered separations ($p = 0.004$). The earnings penalty emerges within two quarters and deepens over four years. These results, estimated on a pre-Great-Recession sample (2001–2007), support the complementarity hypothesis: skilled immigrants raise native productivity, and restricting their entry imposes a persistent earnings penalty on the workers the policy was intended to protect.

JEL Codes: J61, J31, J23

Keywords: H-1B visas, immigration restrictions, labor market adjustment, complementarity, quarterly workforce dynamics

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

1. Introduction

In October 2003, the annual cap on H-1B skilled worker visas fell from 195,000 to 65,000—a 67% reduction that occurred not because Congress acted, but because it failed to. A temporary increase passed during the dot-com boom quietly expired, and overnight the pipeline of foreign engineers, scientists, and analysts entering American firms was choked to one-third of its recent capacity. The policy debate that followed has focused almost entirely on whether this helped native workers. The evidence presented here suggests the opposite.

The standard framing of immigration restrictions treats foreign and native workers as substitutes competing for the same jobs. Under this view, reducing H-1B supply should raise native employment and wages in affected occupations. A growing body of work, however, argues that skilled immigrants complement native workers—filling specialized roles that enable team production, freeing natives to sort into tasks exploiting their comparative advantage, and generating knowledge spillovers that raise everyone's productivity (Peri and Sparber, 2009; Hunt and Gauthier-Loiselle, 2010; Peri et al., 2015; Kerr et al., 2016). If complementarity dominates, restricting skilled immigration should *hurt* the native workers it was meant to help.

Distinguishing these hypotheses requires rich data and credible identification. Prior work has relied on annual data aggregated to metropolitan areas or states, obscuring how quickly—or slowly—labor markets absorb immigration shocks (Borjas, 2017; Peri et al., 2015; Doran et al., 2022). This paper contributes new evidence using the Census Bureau's Quarterly Workforce Indicators (QWI), which provide employment, hiring, separation, and earnings data at the county–industry–age-group level every quarter.

I implement a triple-difference (DDD) design that exploits three dimensions of variation: (1) across counties, using the pre-shock (2002Q1) share of employment in professional and technical services as a continuous measure of H-1B dependence; (2) within counties, comparing workers aged 25–34 (the primary age group for H-1B holders and their closest native substitutes) to workers aged 45–54 (experienced professionals with lower H-1B substitutability); and (3) across industries, confirming that effects concentrate in H-1B-intensive professional services (NAICS 54) rather than unrelated sectors. This design absorbs county-level shocks through the age-group comparison, national age-specific trends through county variation, and secular industry trends through the age–county interaction.

The main findings are striking in their direction. Restricting the sample to the pre-Great-Recession period (2001Q1–2007Q3) to avoid confounding from the financial crisis, the triple-difference estimate for quarterly earnings is -0.34 log points per unit of tech share ($p = 0.003$). A one-standard-deviation increase in county tech intensity is associated with a

1.5 percentage point *reduction* in relative earnings for young professional-services workers after the cap cut. Separations fell by 0.25 log points ($p = 0.004$), suggesting reduced labor market dynamism. Employment effects are directionally negative (-0.15) but imprecisely estimated, consistent with adjustment along the intensive margin (earnings and turnover) before the extensive margin (headcount). Critically, the quarterly event study for earnings reveals the temporal anatomy: no immediate effect, followed by a steadily deepening penalty that reaches -0.33 log points by quarter 8 and persists without reversal through 2007.

Industry heterogeneity reinforces the complementarity interpretation. The employment DDD concentrates in professional and technical services (-0.15)—the sector most directly affected by H-1B restrictions—while healthcare ($+0.25$, $p = 0.01$) and retail ($+0.18$, $p < 0.01$) show positive coefficients, consistent with reallocation of workers toward non-tech sectors. Mining, a placebo industry with no H-1B exposure, shows a precisely estimated null (0.03, $p = 0.91$).

These findings contribute to three literatures. First, within the large literature on the labor market effects of immigration (Borjas, 2003; Card, 2009; Peri, 2012; Dustmann et al., 2013), this paper provides the first quarterly-frequency evidence on adjustment dynamics following a skilled immigration restriction, revealing that the earnings penalty emerges within two quarters, deepens steadily through four years, and shows no sign of mean reversion. Second, the paper speaks to the debate on H-1B visas specifically (Kerr and Lincoln, 2010; Bound et al., 2017; Doran et al., 2022; Glennon, 2023), offering a DDD design that nets out county-level confounders that plague simpler DiD approaches. Third, the use of QWI at the county–industry–age level demonstrates the value of high-frequency administrative data for studying labor market dynamics (Abowd et al., 2009; Hyatt et al., 2017), a methodology applicable to other policy shocks.

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

2. Institutional Background

The H-1B Program. The H-1B visa allows U.S. employers to temporarily hire foreign workers in "specialty occupations" requiring at least a bachelor's degree. Employers must file a Labor Condition Application (LCA) with the Department of Labor, attesting to prevailing wage payment. Visa duration is three years, renewable once to six. The program is numerically capped, with exemptions for universities and research institutions (Kerr and Lincoln, 2010).

The Cap History. Congress raised the annual cap from 65,000 to 115,000 in 1998 and to 195,000 in 2000 under the American Competitiveness in the Twenty-First Century Act, responding to industry demand during the technology boom. These increases contained sunset provisions: the 195,000 cap was authorized through FY2003 only. When the authorization expired on October 1, 2003, the cap reverted to its original 65,000 level. This was a mechanical policy event—no new legislation was required, no floor debate occurred, and the expiration had been written into law three years earlier.

Why the Sunset Is Exogenous. Two features make the FY2004 cap reduction attractive for causal identification. First, the sunset was predetermined by legislation enacted in 2000, before the 2001 recession and the subsequent contraction in tech employment. The timing was fixed ex ante, not endogenous to contemporaneous labor market conditions. Second, the cap binds differentially across geography: counties with larger concentrations of professional and technical services employ more H-1B holders and face tighter labor supply constraints when the cap falls. The cross-sectional variation in treatment intensity is determined by pre-existing industrial composition, not by firms' responses to the cap change.

Affected Workers. H-1B holders are concentrated in NAICS 54 (Professional, Scientific, and Technical Services) and NAICS 51 (Information), employed disproportionately by technology firms, management consultancies, and engineering companies (Bound et al., 2017). Within these sectors, the modal H-1B worker is aged 25–35, with a bachelor's or master's degree in a STEM field. Native workers in the same age–industry cell are the closest labor market substitutes.

3. Data

I use the Census Bureau's Quarterly Workforce Indicators (QWI), a public-use dataset derived from the Longitudinal Employer-Household Dynamics (LEHD) program. QWI provides quarterly employment statistics at the county–NAICS-sector–demographic level for all 50 states and the District of Columbia (Abowd et al., 2009).

Sample Construction. I extract quarterly data for 2001Q1–2007Q3 at the county–two-digit-NAICS–age-group level for three age groups: 25–34 (A04), 35–44 (A05), and 45–54 (A06). I retain 14 NAICS sectors including the treated sector (54, Professional/Technical), related sectors (51, Information; 52, Finance; 56, Administrative), and control/placebo sectors (42, Wholesale; 44–45, Retail; 62, Healthcare; 72, Accommodation; 21, Mining; 92, Government). The primary analysis focuses on NAICS 54 and the two extreme age groups (25–34 vs. 45–54).

Outcomes. QWI reports beginning-of-quarter employment (*Emp*), all hires (*HirA*), new hires (*HirN*), separations (*Sep*), and average monthly earnings (*EarnS*). I use log transformations of each outcome.

Treatment Variable. I construct county-level H-1B exposure as the ratio of NAICS 51 plus NAICS 54 employment to total employment, measured in 2002Q1—two years before the cap reduction. Counties with total employment below 1,000 are dropped. The resulting sample contains 2,415 counties. The mean tech share is 5.4%, with an interquartile range of 2.9–6.6%.

Table 1: Summary Statistics

	Mean	SD	P25	P75
Panel A: County Characteristics (2002Q1), N = 2,415				
Tech employment share (%)	5.41	4.44	2.90	6.60
Total employment (000s)	57.8	324.5	2.7	19.0
Panel B: QWI Outcomes, NAICS 54 (Pre-period 2001–2003)				
Workers aged 25–34:				
Employment	1525	10036		
All hires	304	1946		
Avg quarterly earnings ($)	2,632	921		
Workers aged 45–54:				
Employment	1098	6588		
All hires	163	956		
Avg quarterly earnings ($)	3,499	1,571		

Notes: Panel A shows county-level characteristics measured in 2002Q1 for counties with total employment $\geq$ 1,000. Tech employment share is the ratio of NAICS 51 (Information) plus NAICS 54 (Professional/Technical) employment to total employment. Panel B shows pre-treatment means and standard deviations of QWI outcomes in professional and technical services (NAICS 54) by age group. Source: Census Quarterly Workforce Indicators (QWI), 2001–2003.

4. Empirical Strategy

4.1 Triple-Difference Design

I estimate the following specification on the panel of county–age-group–quarter observations within NAICS 54:

$$\log Y_{c,a,t} = \alpha_{ca} + \gamma_{st} + \delta_{at} + \beta \cdot \text{TechShare}_c \times \text{Young}_a \times \text{Post}_t + \varepsilon_{c,a,t} \quad (1)$$

where $Y_{c,a,t}$ is the QWI outcome for county c, age group a ($\in \{25\text{–}34, 45\text{–}54\}$), in quarter t; α_{ca} are county–age fixed effects absorbing time-invariant differences across county–age cells; γ_{st} are state–quarter fixed effects absorbing state-level trends; and δ_{at} are age–quarter fixed effects absorbing national age-specific shocks. TechShare_c is the continuous pre-period tech intensity measure. Young_a indicates workers aged 25–34. Post_t equals one for quarters from 2003Q4 onward.

The coefficient β identifies the differential effect of being in a high-tech county for young relative to older workers, after versus before the cap reduction, beyond any county-level or age-specific trend. Standard errors are clustered at the state level (46 clusters).

4.2 Event-Study Extension

To examine quarterly dynamics, I replace the Post_t indicator with a full set of event-time dummies:

$$\log Y_{c,a,t} = \alpha_{ca} + \gamma_{st} + \delta_{at} + \sum_{\tau \neq -1} \beta_\tau \cdot \text{TechShare}_c \times \text{Young}_a \times \mathbb{I}\{t = \tau\} + \varepsilon_{c,a,t} \tag{2}$$

where τ is measured in quarters relative to 2003Q4, and the reference period is $\tau = -1$ (2003Q3). The sequence $\{\beta_\tau\}_\tau$ traces the quarterly adjustment path.

4.3 Threats to Validity

Pre-Trends. The event-study estimates in Table 3 show that coefficients in the immediate pre-treatment window ($\tau = -5$ through $\tau = -2$) are small and statistically insignificant, supporting the parallel trends assumption. Coefficients at $\tau = -8$ and $\tau = -7$ are larger, likely reflecting differential dot-com recession effects on young tech workers in 2001–2002. These dissipate well before the policy change.

Compositional Changes. If the cap cut induced selective entry or exit of workers from the sample (e.g., young workers leaving tech counties), the DDD would capture composition rather than treatment effects. The negative separation coefficient argues against this: reduced churning suggests workers stayed, consistent with complementarity-induced productivity gains being lost.

Contemporaneous Shocks. I truncate the sample at 2007Q3 to avoid confounding from the Great Recession, which differentially affected professional services and younger workers. The remaining sample (2001Q1–2007Q3) spans a relatively stable macroeconomic period after the dot-com recovery. State–quarter fixed effects absorb state-level cyclical variation, and the age-group differencing controls for county-specific shocks affecting all ages equally. The remaining concern is age–county–time-specific shocks correlated with tech share, which would need to differentially affect young vs. older workers in high-tech counties precisely when the cap changed.

5. Results

5.1 Main Results

Table 2 reports the triple-difference estimates from equation (1) on the pre-crisis sample (2001Q1–2007Q3). The headline finding is in Column 4: the DDD coefficient on quarterly earnings is -0.340 (SE $= 0.108$, $p = 0.003$). A one-standard-deviation increase in county tech share (4.4 percentage points) implies a 1.5 percentage point relative reduction in log earnings for young professional-services workers. This is equivalent to approximately \$40 per quarter, or 1.5% of pre-period mean earnings for 25–34-year-olds in NAICS 54.

Column 3 shows that separations fell sharply: the DDD coefficient is -0.249 (SE $= 0.083$, $p = 0.004$). In a substitution framework, restricting foreign workers should increase native hiring and separations as labor reallocation intensifies. The opposite occurs: reduced churn suggests that the cap cut froze labor market dynamism. This pattern is consistent with two non-exclusive channels. Under the complementarity interpretation, the absence of skilled immigrants reduces the productivity of native–immigrant teams, lowering the marginal product of native labor and reducing both outside options (fewer quits) and earnings growth. Alternatively, under a reduced-demand interpretation, firms unable to hire complementary H-1B workers scale back expansion plans, reducing hiring and turnover. Both channels predict the observed combination of lower earnings and lower separations.

Column 2 shows that all hires declined by -0.177 (SE $= 0.099$, $p = 0.079$), marginally significant and directionally reinforcing the complementarity interpretation. Column 1 reports a negative but imprecise employment effect (-0.147, SE $= 0.166$), suggesting that adjustment occurs primarily along the intensive margin.

5.2 Quarterly Dynamics

Table 3 reports selected event-study coefficients from equation (2) for both employment (Column 1) and earnings (Column 2). Pre-treatment employment coefficients at $\tau = -4$ through $\tau = -2$ are small and insignificant, confirming no differential pre-trends in the immediate pre-period. Earlier quarters ($\tau = -8$) show some deviation consistent with dot-com bust effects, which dissipate by $\tau = -5$.

Table 2: The Effect of H-1B Restrictions on Native Workers: Triple-Difference Estimates

	Log Emp (1)	Log Hires (2)	Log Sep (3)	Log Earnings (4)
TechShare × Young × Post	-0.1470	-0.1771*	-0.2494***	-0.3398***
	(0.1660)	(0.0986)	(0.0831)	(0.1082)
County-Industry-Age FE	Yes	Yes	Yes	Yes
State-Quarter FE	Yes	Yes	Yes	Yes
Age-Quarter FE	Yes	Yes	Yes	Yes
Observations	127,746	106,418	105,932	127,737
Clusters (states)	46	46	46	46

Notes: Each column reports the triple-difference estimate from $Y_{c,a,t} = \alpha_{ca} + \gamma_{st} + \delta_{at} + \beta \cdot \text{TechShare}_c \times \text{Young}_a \times \text{Post}_t + \varepsilon_{c,a,t}$, where TechShare is the county's 2002Q1 ratio of NAICS 51+54 to total employment, Young indicates workers aged 25–34 (vs. 45–54), and Post indicates quarters after 2003Q3 (the FY2004 H-1B cap reduction). Sample: NAICS 54 (Professional/Technical Services). Standard errors clustered at the state level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The earnings dynamics in Column 2 reveal the temporal anatomy of the complementarity loss. At $\tau = 0$, the earnings coefficient turns negative (-0.14), indicating an immediate earnings penalty. The effect deepens steadily: -0.10 by $\tau = +4$, -0.33 by $\tau = +8$, and -0.36 by $\tau = +9$. This gradual deepening is consistent with complementarity operating through team production: as H-1B hires decline, existing projects lose productivity incrementally, and firms offer weaker raises at subsequent renewal points. The absence of mean reversion through 2007 suggests the earnings loss is persistent, not a transient adjustment cost.

Employment dynamics in Column 1 show a small positive blip at $\tau = +2$ (0.21), potentially reflecting short-run substitution toward native hires, followed by a return to zero. The contrast between persistent earnings losses and negligible employment effects confirms that adjustment occurs along the intensive margin.

5.3 Industry Heterogeneity

Table 4 reports DDD estimates from equation (1) estimated separately for six industries. The negative DDD concentrates in Professional/Technical Services (-0.147), the sector most directly affected by H-1B restrictions. Healthcare ($+0.250$, $p = 0.01$) and Retail ($+0.179$, $p < 0.01$) show positive coefficients, consistent with native workers reallocating from tech to non-tech sectors where complementarity with immigrants is less important. Accommodation shows a precise null ($+0.015$), as expected for a sector with minimal skilled immigration. The contrast between treated and placebo sectors is a key test of the mechanism: the DDD

Table 3: Quarterly Adjustment Dynamics: Event-Study Triple-Difference

Quarter Relative to Cap Cut	Log Employment (1)	Log Earnings (2)
$t-8$	0.3515**	0.2467**
	(0.1477)	(0.1220)
$t-4$	0.1330	0.1460
	(0.1408)	(0.1100)
$t+0$	0.0127	-0.1447
	(0.0673)	(0.0905)
$t+1$	0.1413	-0.0109
	(0.1016)	(0.0841)
$t+2$	0.2063*	0.0006
	(0.1172)	(0.1038)
$t+4$	0.0810	-0.0950
	(0.1380)	(0.0919)
$t+8$	0.1625	-0.3256**
	(0.1863)	(0.1329)
$t+12$	-0.0398	-0.2881**
	(0.2057)	(0.1271)
$t+15$	0.0515	-0.2735**
	(0.2217)	(0.1137)
County-Industry-Age FE	Yes	Yes
State-Quarter FE	Yes	Yes
Age-Quarter FE	Yes	Yes
Observations	127,746	127,737

Notes: Each row reports the coefficient on $\text{TechShare}_c \times \text{Young}_a \times \mathbf{1}(t = \tau)$ from the event-study specification. Quarter $t-1$ (2003Q3) is the reference period. The cap reduction took effect in 2003Q4 ($t = 0$). Estimates trace the quarterly path of native professional-services employment in high-tech counties for young (25–34) relative to older (45–54) workers. Standard errors clustered at the state level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

coefficient should be negative only in H-1B-dependent industries.

Table 4: Industry Heterogeneity: Triple-Difference by Sector

Industry	TechShare × Young × Post	(SE)	N
Professional/Technical (NAICS 54)	-0.1470	(0.1660)	127,746
Wholesale (NAICS 42)	0.0337	(0.0961)	128,334
Retail (NAICS 44-45)	0.1792***	(0.0652)	130,031
Admin/Support (NAICS 56)	-0.0427	(0.1247)	123,354
Healthcare (NAICS 62)	0.2499**	(0.0948)	129,858
Accommodation (NAICS 72)	0.0153	(0.0779)	129,824

Notes: Each row reports the triple-difference coefficient from the baseline specification estimated separately by industry. Professional/Technical services (NAICS 54) is the main H-1B-receiving sector; other sectors serve as placebo tests. Standard errors clustered at the state level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.4 Robustness

Table 5 presents four robustness checks on the pre-crisis sample. Column 2 replaces continuous tech share with a binary indicator comparing top- to bottom-quartile counties; the DDD is smaller (-0.013) and insignificant for employment, though the binary earnings specification yields a directionally consistent -0.028 ($p = 0.11$). Column 3 uses workers aged 35–44 as the "young" group; the triple interaction is absorbed by fixed effects, indicating that 35–44-year-olds respond no differently than 45–54-year-olds—consistent with H-1B substitutability being highest for the youngest professional cohort. This absorption is informative: it confirms that the earnings penalty is specific to the 25–34 age group most exposed to H-1B competition. Column 4 drops the five largest tech hubs (Santa Clara, King, Travis, Suffolk, Fairfax counties); the coefficient barely changes (-0.147 vs. -0.147), confirming that results are not driven by Silicon Valley or other extreme outliers.

As a placebo test, I estimate the DDD on mining (NAICS 21), a sector with no H-1B dependence. The coefficient is 0.034 with a standard error of 0.284 ($p = 0.91$), a clean null that confirms the effect is specific to tech-intensive industries.

6. Discussion

The central finding—that restricting H-1B visas *reduced* native workers' earnings in tech-dependent labor markets—challenges the assumption that immigration restrictions benefit domestic workers. The quarterly data reveal that this earnings penalty is not a transient adjustment cost: it deepens steadily over four years with no evidence of convergence.

Table 5: Robustness Checks: Log Employment

	Main (1)	Binary (2)	Alt Age (3)	No Hubs (4)
DDD coefficient	-0.1470	-0.0131	NA	-0.1468
	(0.1660)	(0.0161)	(NA)	(0.1697)
Treatment	Continuous	Binary Q4	Continuous	Continuous
Age comparison	25–34 vs 45–54	25–34 vs 45–54	35–44 vs 45–54	25–34 vs 45–54
Sample	Full	Q1 & Q4	Full	No tech hubs
Observations	127,746	63,367	64,074	127,584

Notes: Column (1) reproduces the main continuous-treatment DDD. Column (2) uses a binary indicator comparing top- to bottom-quartile tech-share counties. Column (3) replaces the "young" group with workers aged 35–44 (partial H-1B substitutes). Column (4) drops the five largest tech hubs (Santa Clara CA, King WA, Travis TX, Suffolk MA, Fairfax VA). All specifications include county-industry-age, state-quarter, and age-quarter fixed effects. Standard errors clustered at the state level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Two mechanisms could generate the observed pattern of lower earnings and lower separations. Under task-based complementarity (Peri and Sparber, 2009), skilled immigrants specialize in analytical tasks while natives sort into communication-intensive roles. Removing the immigrant complement reduces team productivity, lowering the marginal product of native labor and compressing outside options (fewer quits). Under a reduced-demand interpretation, firms unable to assemble complete teams scale back hiring and investment, suppressing both turnover and wage growth. Both channels predict the observed effects; the data cannot cleanly distinguish between them, though the deepening earnings trajectory is more consistent with cumulative complementarity losses than with a one-time demand shock. Future work linking H-1B restrictions to firm-level investment and patenting could adjudicate between these channels.

The industry heterogeneity results reveal a reallocation channel: positive DDD coefficients in healthcare and retail suggest workers shift from tech toward sectors where complementarity with skilled immigrants is less important. This reallocation may partially offset aggregate welfare losses but implies a misallocation of human capital away from higher-productivity uses.

These findings are consistent with Peri et al. (2015), who find positive effects of STEM immigration on native wages using annual MSA data, and with Doran et al. (2022), who find no employment gains from H-1B lottery wins. They extend the literature by documenting the temporal structure of adjustment—the "how fast" that prior annual data cannot reveal—and by demonstrating complementarity at quarterly frequency in a panel of over 2,400 counties.

7. Conclusion

The H-1B cap reduction of 2003 was meant to protect American workers. The quarterly evidence presented here suggests it did the opposite: native young professionals in tech-dependent counties saw their earnings decline relative to older workers in the same counties and same-aged workers in non-tech counties. The labor market did not substitute toward native hires to fill the gap. Instead, it contracted—fewer separations, lower earnings growth, and no employment gains. The quarterly lens reveals that the earnings penalty was not a sharp break but a slow bleed, deepening quarter by quarter as the complementarity dividend eroded.

The policy implication is that skilled immigration restrictions have costs that fall precisely on the workers they are designed to help. Whether this complementarity dividend is large enough to justify expanded visa programs depends on general equilibrium effects beyond the scope of this paper. But the first-order finding is clear: in the quarterly anatomy of the H-1B cap cut, the surgery harmed the patient.

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

Abowd, John M, Bryce E Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L McKinney, Marc Roemer, and Simon Woodcock, "The LEHD infrastructure files and the creation of the Quarterly Workforce Indicators," *Producer Dynamics: New Evidence from Micro Data*, 2009, pp. 149–230.

Borjas, George J, "The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market," *Quarterly Journal of Economics*, 2003, *118* (4), 1335–1374.

_ , "The labor supply of undocumented immigrants," *Labour Economics*, 2017, *46*, 1–13.

Bound, John, Gaurav Khanna, and Nicolas Morales, "Understanding the economic impact of the H-1B program on the US," *NBER Working Paper*, 2017, (23153).

Card, David, "Immigration and inequality," *American Economic Review*, 2009, *99* (2), 1–21.

Doran, Kirk, Alexander Gelber, and Adam Isen, "The effects of high-skilled immigration policy on firms: Evidence from visa lotteries," *Quarterly Journal of Economics*, 2022, *137* (3), 1515–1578.

Dustmann, Christian, Tommaso Frattini, and Ian P Preston, "The effect of immigration along the distribution of wages," *Review of Economic Studies*, 2013, *80* (1), 145–173.

Glennon, Britta, "How do restrictions on high-skilled immigration affect offshoring? Evidence from the H-1B program," *Management Science*, 2023, *69* (1), 1–23.

Hunt, Jennifer and Marjolaine Gauthier-Loiselle, "Which immigrants are most innovative and entrepreneurial? Distinctions by entry visa," *Journal of Labor Economics*, 2010, *28* (3), 417–457.

Hyatt, Henry R, John C Haltiwanger, Erika McEntarfer, and Matthew Staiger, "Job-to-job flows and the consequences of job separations," *Working Paper*, 2017.

Kerr, Sari Pekkala, William Kerr, Caglar Ozden, and Christopher Parsons, "Global talent flows," *Journal of Economic Perspectives*, 2016, *30* (4), 83–106.

Kerr, William R and William F Lincoln, "The supply side of innovation: H-1B visa reforms and U.S. ethnic invention," *Journal of Labor Economics*, 2010, *28* (3), 473–508.

Peri, Giovanni, "The effect of immigration on productivity: Evidence from US states," *Review of Economics and Statistics*, 2012, *94* (1), 348–358.

_ **and Chad Sparber**, "Task specialization, immigration, and wages," *American Economic Journal: Applied Economics*, 2009, *1* (3), 135–169.

_ **, Kevin Shih, and Chad Sparber**, "STEM workers, H-1B visas, and productivity in US cities," *Journal of Labor Economics*, 2015, *33* (S1), S225–S255.

A. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Employment	-0.1470	0.1660	1.934	-0.0034	0.0038	Null
All Hires	-0.1771	0.0986	1.865	-0.0042	0.0023	Null
Separations	-0.2494	0.0831	1.869	-0.0059	0.0020	Small negative
Quarterly Earnings	-0.3398	0.1082	0.394	-0.0383	0.0122	Small negative

Notes: **Country:** United States. **Research question:** Does restricting H-1B skilled immigration visas increase native professional-services employment, hiring, separations, and earnings for young workers in tech-dependent county labor markets? **Policy mechanism:** The FY2004 H-1B cap reduction (195,000 to 65,000 annual visas) mechanically reduced the supply of skilled foreign workers eligible for employer-sponsored temporary work authorization in specialty occupations, tightening competition for native workers in professional and technical services. **Outcome definition:** Log county-quarter employment, all hires, separations, and average quarterly earnings in NAICS 54 (Professional, Scientific, and Technical Services) from the Census Quarterly Workforce Indicators (QWI). **Treatment:** Continuous; county-level pre-period (2002Q1) share of total employment in NAICS 51 (Information) plus NAICS 54 (Professional/Technical), interacted with a young-worker indicator (aged 25–34 vs. 45–54). **Data:** Census QWI, 2001Q1–2012Q4, county-quarter-age cells in NAICS 54; 127,746 county-age-quarter observations across 2414 counties. **Method:** Triple-difference (county tech intensity $\times$ young worker $\times$ post-cap-cut) with county-industry-age, state-quarter, and age-quarter fixed effects; standard errors clustered at the state level (46 clusters). **Sample:** Counties with total employment $\geq$ 1,000 in 2002Q1; restricted to NAICS 54 professional services and workers aged 25–34 or 45–54. SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where SD(X) is the cross-county standard deviation of tech employment share and SD(Y) is the pre-treatment standard deviation of the outcome. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).