

LONG-RUN EFFECTS OF H-1B IMMIGRATION ON THE U.S. ECONOMY

By

Ran Abramitzky, Leah Platt Boustan, Ahmet Gulek and Jens Hainmueller

July 2026

COWLES FOUNDATION DISCUSSION PAPER NO. 2551



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

ABSTRACT

We study the effects of H-1B immigration on U.S. industries that employ H-1B workers and their trading partners. Using a novel cross-industry design and the 1999–2003 expansion of the H-1B visa cap for identification, we find that H-1B exposure raised incomes for natives and pre-existing immigrants, with gains concentrated in non-STEM occupations. Income gains propagate forward through supply chains to downstream industries but not backward to upstream industries, consistent with a productivity shock rather than a labor supply shock. We find no direct effect on patenting, suggesting that productivity gains arise from better task execution rather than patentable invention.

Ran Abramitzky
Stanford University
Department of Economics
and NBER
ranabr@stanford.edu

Ahmet Gulek
Stanford University
agulek@stanford.edu

Jens Hainmueller
Stanford University
jhain@stanford.edu

Leah Platt Boustan
Yale University
Department of Economics
and NBER
leah.boustan@yale.edu

We are grateful to Michael Clemens for his invaluable comments. Mi Liu provided excellent research assistance.

JEL No. J60, J68

1 Introduction

How do high-skilled immigrants affect the industries that employ them, and do these effects extend to other industries through production networks? A large literature studies the first question (see Lewis and Peri 2015 for a survey), including work on wages and employment (Borjas, 2003; Ottaviano and Peri, 2012; Dustmann and Glitz, 2015; Beerli et al., 2021), innovation (Akcigit et al., 2017; Kerr and Lincoln, 2010; Hunt and Gauthier-Loiselle, 2010), and task specialization (Peri and Sparber, 2009; Peri et al., 2015). Within this literature, a growing body of work investigates the H-1B program specifically, exploiting either cross-city variation in H-1B concentration (Kerr and Lincoln, 2010; Peri et al., 2015) or the H-1B lottery at the firm level (Doran et al., 2022; Mahajan et al., 2024). Yet these designs each capture only a partial view. Cross-firm designs reflect the short-run impact on the sponsoring firm but miss effects that diffuse across employers within the same labor market. Cross-city designs can over- or understate the effects of immigration, depending on the tradability of the host industries (Burstein et al., 2020) and the structure of the production network (Gulek, 2025). Neither design can trace how H-1B shocks propagate through supply chains, because input–output networks are observed in the United States only at the industry level, not at the city or firm level.

We provide the first industry-level causal estimates of the H-1B program’s long-run effects, exploiting the largest-ever shock to the program for identification: the tripling and reversal of the annual visa cap between 1999 and 2004. We find that H-1B immigration raised incomes for natives and pre-existing immigrants. Income gains concentrate in non-STEM occupations, propagate forward to downstream industries, and seemingly arise from more efficient task execution rather than patentable innovation.

Theory alone cannot sign the effects of H-1B workers. The canonical labor demand framework—the neoclassical model with fixed technology—predicts that immigration increases competition for similar native workers and therefore lowers their wages. But the H-1B program is designed to bring specialized skills,¹ and workers who raise productivity need not simply compete with natives. The program restricts eligibility to “specialty occupations” requiring at least a bachelor’s degree in a specific field.² If H-1B workers raise productivity, whether through innovation or through more efficient execution of specialized tasks, the resulting productivity gains can outweigh the competition effects, benefiting all existing workers. Moreover, whether immigrants primarily compete against native workers or improve productivity has different implications for industry spillovers. Increases in productivity lower production costs and propagate forward through supply chains, benefiting downstream industries (Acemoglu et al., 2012). In contrast, labor supply shocks both lower costs and shift the input mix between labor and intermediate goods, and therefore propagate both forward

¹Historical examples include Huguenot refugees who diffused skills into Prussian manufacturing (Hornung, 2014) and German Jewish émigré scientists who raised U.S. invention (Moser et al., 2014); the displacement channel appears with post-1992 Soviet mathematicians, who crowded out U.S.-born mathematicians from top positions (Borjas and Doran, 2012).

²The majority of H-1B workers are employed in STEM occupations, particularly in computer science and engineering.

and backward (Gulek, 2025). The direction of supply chain propagation therefore distinguishes the two channels empirically.

We first document a new fact: while which exact *firms* sponsor H-1B workers changes dramatically over time, the *industry-level* ranking of H-1B demand is highly persistent. This motivates an industry-level exposure design: any change in H-1B quotas is effectively a shock to particular industries, not particular firms or regions, and lets us study supply chain spillovers, since input-output networks are observed at the industry level in the United States.

The empirical challenge is that more- and less-exposed industries followed different trajectories prior to the policy change, making the parallel trends assumption unlikely to hold. To make progress, we combine synthetic control methods with difference-in-differences in a continuous treatment setting. We compare industries with similar skill intensities and pre-shock income trajectories that happen to be differentially exposed because of pre-existing H-1B demand shares. The most H-1B-exposed industry is Computer Systems Design; its synthetic control is built from other high-education industries that grew similarly through 2000 but received far less H-1B exposure.³ Standard difference-in-differences produces biased estimates in this setting; we show that synthetic control methods remove about half of this bias in an out-of-sample placebo test.

We answer two complementary questions. First, what is the direct effect of H-1B exposure on incomes within an industry? We address this with a continuous treatment design exploiting variation in H-1B exposure across all industries, using census and patent data. Second, how does that effect propagate across industries? We address this with an input-output spillover design that traces H-1B shocks through production networks.⁴ These designs let us progressively narrow the mechanism: we first establish that H-1B immigration raises incomes in exposed industries, then test whether it operates as a productivity shock or a labor supply shock, and finally ask whether the productivity gains come from patentable innovation or from more efficient execution of existing tasks.

First, H-1B immigration raised long-run incomes for native workers in exposed industries, with gains concentrating in non-STEM occupations. A one-percentage-point increase in H-1B exposure raises long-run non-STEM income by 1.2 log points for college-graduate natives. In Computer Systems Design, the most exposed industry, this implies a long-run income gain of approximately 15 percent for native non-STEM workers. That the gains concentrate among non-STEM workers is itself an identification test: an unobserved technology shock driving H-1B demand would raise the wages of STEM workers, whose labor demand is rising, not those of their non-STEM coworkers. The pattern instead points to complementarity: when STEM and non-STEM workers are gross complements, more STEM labor raises the marginal product, and hence the wages, of the non-

³The synthetic control re-weights donor industries to match each outcome’s pre-2000 income trajectory, so the comparison industries differ across outcomes; the leading donors are named in Online Appendix Section B.1. For high-skill natives the largest weight falls on catalog and mail-order houses (e.g., Amazon), reflecting the coincident e-commerce and dot-com boom; for low-skill natives it falls on management and public relations services.

⁴In the Online Appendix, we complement these with a binary treatment design comparing Computer Systems Design, the most exposed industry, to its synthetic control, which transparently illustrates the identification challenge and its resolution.

STEM workers performing complementary tasks. Pre-existing immigrants also gain, which is hard to square with new immigrants simply substituting for incumbents while complementing natives. A structural literature simulates the program’s general-equilibrium effects under the assumption that H-1B and native computer scientists are close substitutes, and predicts lower wages for competing natives (Bound et al., 2015, 2018). We reproduce that loss for native STEM workers in the short run. What their framework does not deliver is the larger, offsetting non-STEM gain, which is where the aggregate effect is made, and which is consistent with the concentration of gains among less-skilled workers that this literature also anticipates (Bound et al., 2018). Consistent with limited substitutability, Mayda et al. (2018) find that the 2004 cap reduction lowered H-1B employment with no offsetting change in native hiring. The firm-level lottery estimates of Mahajan et al. (2024) show the same gradient within sponsoring firms: most incumbents gain—non-college and lower-tenure college workers of all nativities see wage increases of 3–5%—while only the narrow, most-substitutable subgroup of high-tenure native college workers declines. This parallels our STEM/non-STEM split.

Second, if H-1B immigrants improved productivity in exposed industries, production costs and therefore output prices should decrease, benefiting downstream industries that purchase their outputs, while upstream industries should be largely unaffected (Acemoglu et al., 2012). In contrast, a pure labor supply channel would both lower production costs and alter the input mix between labor and intermediate goods, and therefore create both upstream and downstream spillovers (Gulek, 2025).⁵ We find evidence consistent with the former: income gains propagate forward, while backward propagation is not distinguishable from zero. This asymmetry reinforces the interpretation of H-1B immigration as a productivity shock operating through supply chains.

Third, a prominent finding in the literature is that high-skilled immigrants, particularly H-1B workers, increase patenting (Kerr and Lincoln, 2010). We revisit this question using our design and find no evidence that H-1B immigration *directly* increased patenting in exposed industries. As we show in the Online Appendix, patent growth in high-exposure regions comes predominantly from industries that barely hire H-1B workers, making it unlikely to reflect a direct effect of H-1B immigration. The null patent result changes the policy interpretation: the productivity gains documented above arise from better execution of skilled tasks, not from patentable invention.

Together, these results indicate that H-1B immigration operated as a productivity shock: non-STEM income gains, null patent effects, and forward-only supply chain propagation all point to more efficient execution of skilled tasks rather than patentable innovation. We cannot fully rule out unobserved post-treatment shocks differentially affecting exposed industries, but the STEM/non-STEM decomposition is difficult to reconcile with such alternatives. These findings bear on recent restrictions on H-1B visas and the permanent residency pathway, which we discuss in the conclusion.

Our industry-level design complements the firm-level lottery designs (Doran et al., 2022; Maha-

⁵This forward/backward asymmetry is long-established: forward propagation of a cost reduction follows the cost-of-production logic of Hulten (1978) underlying Acemoglu et al. (2012), while the absence of backward propagation under Cobb-Douglas is an instance of the nonsubstitution theorem (Samuelson, 1951). Section 4.2 develops the intuition.

jan et al., 2024) and cross-city designs (Kerr and Lincoln, 2010; Peri et al., 2015) discussed above: it captures long-run within-industry effects, including the cross-firm spillovers that firm-level estimates difference out, and it traces propagation through the input-output networks that city-level designs (Burstein et al., 2020) cannot observe.⁶ The synthetic control methods we employ address the differential pre-trends that bias standard difference-in-differences and shift-share estimates in this setting: industries that attracted H-1B workers were already on steeper growth trajectories during the 1990s technology boom, and failing to account for these trends overstates the causal effect of H-1B immigration.

A cross-regional design identifies only a region’s exposure *relative* to other regions, differencing out the aggregate general-equilibrium response common to all regions (Nakamura and Steinsson, 2014; Chodorow-Reich, 2019; Guren et al., 2021). Because input–output linkages are largely aspatial, the H-1B spillover that propagates through the production network is nearly uniform across regions and falls almost entirely into this differenced-out component—the so-called “missing intercept” (Wolf, 2023). Such designs recover the cross-regional slope almost intact but miss this level component, which we estimate at roughly 40 percent of the total effect (Online Appendix A.7).

We also contribute to an emerging literature on how immigration shocks propagate beyond directly affected markets. Immigration can generate geographic spillovers through native relocation and spatial equilibrium adjustment (Monras, 2020; Albert and Monras, 2022; Burstein et al., 2020), and internationally through trade competition (Brinatti and Guo, 2024) and offshoring (Glennon, 2024). Recent work has begun to trace immigration shocks through production networks: Gulek (2025) shows theoretically and empirically that labor supply shocks propagate both forward and backward through supply chains, and Akgündüz et al. (2024) documents similar network effects.⁷ We provide the first evidence that a high-skilled immigration shock propagates through input-output networks as a productivity shock, and the asymmetry we document—forward propagation only—distinguishes this channel from the labor supply mechanism.

2 Background and Data

2.1 H-1B Institutional Background

The H-1B visa program was established by the Immigration Act of 1990 to allow U.S. employers to temporarily hire foreign workers in “specialty occupations” requiring at least a bachelor’s degree or equivalent. The 1990 Act set the initial annual cap at 65,000 new visas per fiscal year and limited stays to a maximum of six years. Employers were required to file a Labor Condition Application

⁶Peri et al. (2015) also exploit H-1B policy variation, but through a city-level shift-share that interacts 1980 foreign-STEM enclave shares with national H-1B flows; we use the 1999–2003 cap expansion directly, comparing industries with more and less pre-existing H-1B demand. The two designs therefore capture complementary spillover sets: theirs, within-city cross-industry agglomeration (differencing out national general equilibrium and cross-city spillovers); ours, national cross-industry input-output propagation.

⁷Mahajan (2026) documents a similar asymmetry for H-2B workers during the COVID crisis, finding forward propagation and no backward propagation, consistent with temporary migrant workers operating as a productivity input.

attesting that they would pay the prevailing wage. Although the visa is temporary, H-1B holders are eligible to apply for employment-based permanent residency, and many do: the H-1B has historically served as the primary pathway from temporary to permanent high-skilled immigration.

The cap underwent two significant expansions in response to labor market pressure during the technology boom. The American Competitiveness and Workforce Improvement Act of 1998 raised the cap to 115,000 for fiscal years 1999 and 2000. The American Competitiveness in the Twenty-First Century Act of 2000 further increased the cap to 195,000 for fiscal years 2001 through 2003. The cap then reverted to 65,000 beginning in fiscal year 2004. The H-1B Visa Reform Act of 2004 introduced an additional 20,000 visas for applicants holding a master’s degree or higher from a U.S. institution, bringing the combined annual cap to 85,000 beginning in fiscal year 2005. The statutory cap has remained unchanged since.⁸ The variation in the cap between 1999 and 2004 (from 65,000 to 195,000 and back) provides the identifying variation that we exploit in this paper.

2.2 Data

We combine several data sources. *H-1B filings data* on newly issued visas at the six-digit NAICS level starting from 1998 come from data compiled by Mahajan et al. (2024). We complement this with publicly available data from the USCIS Employer Hub, which provides information on all H-1B petitions since 2009.

Census and ACS data consist of the 5 percent samples from the 1980, 1990, and 2000 censuses, and the American Community Survey (ACS) from 2001–2019, obtained from IPUMS (Ruggles et al., 2025). We use the IPUMS *ind1990* consistent industry classification, yielding 206 industries tracked over time. These data allow us to distinguish between natives and immigrants, education groups, and STEM versus non-STEM occupations.

QCEW data from the Quarterly Census of Employment and Wages cover approximately 95 percent of all U.S. jobs and provide industry-level data on establishments, employment, and wages at the four-digit NAICS level, available annually from 1990 through 2024.

Patent data from Kogan et al. (2017) provide the number of patents and forward citations for publicly traded firms, which we match to three-digit SIC industries via Compustat. We date each patent by its *application* (filing) year rather than grant year, because it is far closer to the underlying innovative activity.

Input-output data from the 1997 Bureau of Economic Analysis Benchmark I/O tables provide the Leontief inverse matrix, which we use to construct measures of forward and backward supply chain exposure. We use the 1997 tables because they are pre-determined relative to the H-1B policy variation. Details on data construction and crosswalks are in Section A.1 of the Online Appendix.

⁸Certain employers (higher education, nonprofit research organizations, and government research institutions) are exempt from the cap.

2.3 Exposure Measure

We define an industry’s exposure to the H-1B policy shock as the expected increase in its workforce due to the cap expansion. We first calculate each industry’s average share of H-1B visas in 1998–2000:

$$w_i^{H1B} = \frac{1}{3} \sum_{t=1998}^{2000} \frac{\text{H-1Bs issued for industry } i \text{ in year } t}{\text{Total H-1Bs issued in year } t} \quad (1)$$

Industry-level exposure is then:

$$E_i = \frac{490,000 \times w_i^{H1B}}{L_i} \times 100 \quad (2)$$

where L_i is the total number of workers in industry i in the 2000 Census. E_i measures the number of additional H-1B workers per 100 existing workers, capturing the percentage increase in industry employment that the cap expansion authorized.

The weights w_i^{H1B} are averaged over 1998–2000, the initial years of the cap expansion.⁹ Because the industry-level distribution of H-1B visas is stable over time (Figure 1c), these shares capture predetermined industry demand for H-1B workers rather than endogenous responses to the policy change. Although the cap reverted to 65,000 in fiscal year 2004, the workers admitted during the expansion remained in the labor force, so the treatment we study is the persistent stock of additional H-1B workers in each industry, not the annual flow. The 490,000 figure is the cumulative statutory cap increase above the 65,000 baseline over fiscal years 1999–2003, excluding renewals and extensions. Dividing this total by year-2000 employment, and assuming H-1B workers remain in their sponsoring industry for the duration of their visa, E_i approximates the percentage increase in industry employment due to the temporary cap expansion. Because E_i scales the statutory cap increase rather than realized admissions, it is an *expected*, intent-to-treat measure: the 195,000 cap went unmet in fiscal years 2001–2003 as demand softened in the dot-com downturn, so E_i overstates the realized inflow and our estimates are conservative relative to the effect per realized H-1B worker.¹⁰


Figure 1 introduces the key features of the setting. Panel (a) shows the evolution of the statutory H-1B cap over time. Panels (b) and (c) present a contrast that motivates our research design: while there is dramatic turnover in which *firms* are the top H-1B sponsors (panel b), the *industry-level* ranking of H-1B demand is highly persistent (panel c): Computer Systems Design receives 22–53 percent of all H-1B visas throughout the period, always by far the largest share, despite substantial turnover in the top sponsoring firms. This persistence implies that H-1B demand is fundamentally an industry-level phenomenon driven by the skill requirements of specific sectors, motivating our use of industry-level variation. Panel (d) shows the distribution of H-1B exposure across industries, plotted against skill intensity (college-graduate share). Computer & Data Pro-

⁹Averaging across years lowers noise, but restricting to 1998 shares alone does not change the results (Online Appendix, Section B.5).

¹⁰The 195,000 cap was not reached in fiscal years 2001–2003 (U.S. Government Accountability Office, 2011).

cessing Services (*ind1990: 732*) has the highest exposure. Importantly for identification, there is substantial variation in H-1B exposure *within* skill groups: industries of similar skill intensity have very different exposure levels.

Figure 1: H-1B Policy and Industry Exposure



Notes: Panel (a) shows the statutory H-1B cap. Panel (b) shows the number of new H-1B visas sponsored by the top firms in each year, beginning in 2009 when firm-level public data become available; firm names are visible to illustrate the dramatic turnover. Panel (c) shows the share of H-1B visas by four-digit NAICS industry over time. Panel (d) plots industry-level H-1B exposure (equation 2) against the share of college graduates in the industry workforce (2000 Census).

3 Empirical Design

We exploit the 1999–2004 H-1B cap variation in three analyses: direct effects on income, direct effects on patenting, and indirect supply-chain spillovers. All three rest on the continuous-treatment SC-DiD estimator we explain below, together with the identification strategy; Section 4 presents the results.

3.1 Event-Study Specification

We estimate the following event-study specification:

$$y_{it} = \sum_{j \neq 2000} \beta_j E_i \cdot \mathbb{1}(t = j) + f_i + f_{q(i),t} + \epsilon_{it} \quad (3)$$

where y_{it} is the outcome of interest (e.g., log income of native college graduates in industry i at time t), E_i is the H-1B exposure defined in equation (2), f_i are industry fixed effects, and $f_{q(i),t}$ are skill-quartile-by-time fixed effects, where $q(i)$ assigns each industry to a quartile of the baseline college-graduate share distribution. Identification comes from comparing industries with similar skill intensity but different H-1B exposure: the coefficients $\{\beta_j\}$ trace out the dynamic effects within skill groups. The pre-treatment coefficients test whether differentially exposed industries followed parallel trajectories before the policy change; the post-treatment coefficients estimate the causal effects. In our Census/ACS data, we observe industries from 1980 through 2019, with the policy shock occurring in 2000.

We weight regressions by the size of the relevant subpopulation in each industry in 2000. Standard errors are clustered at the industry level. We focus on income in the main text. Employment effects are mixed and less robust across specifications; we report them in the Online Appendix.

3.2 Short-Run and Long-Run Effects

To succinctly compare average treatment effects across the many subgroups we consider—nativity, education, and occupation—we collapse the dynamic effects into two parameters using a pre/post specification:

$$y_{it} = \beta_{SR} E_i \cdot \mathbb{1}(2001 \leq t \leq 2009) + \beta_{LR} E_i \cdot \mathbb{1}(t \geq 2010) + f_i + f_{q(i),t} + \epsilon_{it} \quad (4)$$

with year 2000 as the omitted base period. The coefficient β_{SR} captures the short-run average treatment effect (2001–2009) and β_{LR} the long-run effect (2010–2019). This specification is used for the STEM/non-STEM decomposition of income effects and, with modified treatment variables, for the supply chain analysis.

3.3 Supply Chain Specification

To test whether H-1B effects propagate through production networks, we construct two measures of indirect exposure using input-output tables. Forward exposure captures how much an industry’s suppliers are H-1B-exposed:

$$F_i = \sum_{j \neq i} \Psi_{ij} \cdot E_j \quad (5)$$

where Ψ_{ij} is the (i, j) element of the Leontief inverse matrix from the 1997 BEA Benchmark I/O tables and E_j is the direct H-1B exposure defined in equation (2). Backward exposure captures

how much an industry’s customers are H-1B-exposed:

$$B_i = \sum_{j \neq i} \Psi_{ji} \cdot E_j \quad (6)$$

Both measures use the off-diagonal 1997 BEA Leontief inverse, pre-determined relative to the policy shock. Industry 732 (Computer & Data Processing Services) is dropped from the regression sample and the SC donor pools. Its exposure still enters other industries’ network sums, which is the propagation we measure. Forward and backward exposure are nearly uncorrelated ($\rho = 0.06$), allowing us to separately identify the two channels.

We estimate equation (4) replacing E_i with F_i and B_i entered jointly, with short-run (2001–2009) and long-run (2010–2019) interactions for both directions. The same SC debiasing procedure applies, with F_i or B_i replacing E_i .

3.4 Synthetic Control Debiasing and Validation

Our design combines predetermined industry shares with a single national policy shock, so identification reduces to a difference-in-differences with parallel trends as the key assumption. This sidesteps the conflation of past and current immigration that troubles conventional shift-share designs (Jaeger et al., 2018).¹¹

Even within skill quartiles, industries with higher H-1B exposure followed steeper pre-treatment trajectories, reflecting the late-1990s IT boom that simultaneously drove H-1B demand and growth in technology-intensive industries. Standard difference-in-differences is therefore biased here; we address this with synthetic control debiasing.

For each industry i , we construct synthetic control weights ω_{ij} ¹² by matching industry i ’s pre-treatment outcome trajectory, demeaned by its own pre-period mean so that the match is on trends rather than levels, to a weighted combination of donor industries in the same skill quartile (Abadie et al., 2010). We then generate counterfactual values for both the outcome and the treatment, $\hat{y}_{it} = \sum_j \omega_{ij} y_{jt}$ and $\hat{E}_i = \sum_j \omega_{ij} E_j$, and estimate equations (3)–(4) on debiased data, defined as the difference between observed and synthetic values: $\tilde{y}_{it} = y_{it} - \hat{y}_{it}$ and $\tilde{E}_i = E_i - \hat{E}_i$. Pre-trends in the raw data imply that unobserved confounders are correlated with H-1B exposure; the matching step proxies these confounders using pre-treatment outcome trajectories, so subtracting the synthetic values removes the confounded component.¹³ This approach is similar in spirit to Synthetic DiD (Arkhangelsky et al., 2021), with the key difference that it generalizes to continuous

¹¹Shift-share identification can rest on exogeneity of the shares (Goldsmith-Pinkham et al., 2020) or of the shocks (Borusyak et al., 2022); with a single aggregate shock, identification runs through the shares. Synthetic-control debiasing relaxes this, requiring only the residual exposure, after subtracting each industry’s synthetic counterfactual, to be orthogonal to post-2000 shocks.

¹²SC weights are computed via constrained quadratic programming in MATLAB. Full algorithmic details, solver parameters, and software versions are in Online Appendix Section A.2.

¹³The SC debiasing procedure is a special case of the Synthetic IV estimator described in Gulek and Vives-i Bastida (2024), where the instrument equals the treatment. For an alternative SC estimator with continuous treatment, see Powell (2022).

treatment exposure.

The training period for the SC weights varies by data source. For Census/ACS outcomes, observed decennially, we train on the full pre-treatment period, 1980–2000; the Online Appendix reports the estimates under alternative training windows (Section A.3). For patents and QCEW outcomes, observed annually, we specify the training windows when presenting those results. The SC-DiD estimates the pre/post specification (equation (4)) on post-training-period data only; this separation between the SC training and estimation samples improves its finite-sample properties.

We validate the SC-DiD procedure by backtesting: we hold out a window in which nothing is treated, where any estimated effect is bias. The Census is decennial, so we backtest in the annual QCEW panel, varying the SC training window from two to nine pre-periods. The SC-DiD beats OLS at every window, but by how much depends on its length: with two pre-periods it removes 12 percent of OLS’s bias; with six or more, 53 to 57 percent. We therefore do not backtest the ACS, which affords only three decennial pre-periods; we show instead that our estimates are stable across training windows (Online Appendix, Section A.3).

All results reported in the main text are robust to alternative SC matching strategies (separate and joint STEM matching), alternative training windows, and multiple data sources (Census/ACS, QCEW, patents). Details on the SC algorithm, the standard-error construction, and complete robustness checks are in the Online Appendix (Sections A.2–A.6).

4 Results


4.1 Direct Effects on Income

We begin with the direct effects of H-1B exposure on worker incomes. Figure 2 presents the evidence. Panels (a)–(d) show event-study estimates for four population groups: native college-graduate, native less-than-college, pre-existing immigrant college-graduate, and pre-existing immigrant less-than-college. Each panel overlays two specifications: OLS with skill-quartile $\times$ year fixed effects and SC with skill-quartile $\times$ year matching. We use pre-existing immigrants to eliminate concerns about compositional changes from new arrivals. The key visual pattern is consistent across panels: OLS exhibits significant pre-trends, confirming that more-exposed industries were on different trajectories before the policy change. The SC-DiD specifications flatten these pre-trends, and from 2003 onward the estimates are positive in every year for all four groups.

Figure 3 summarizes the short-run and long-run income effects from the pre/post specification (equation (4)) with joint STEM matching. For each population group, we decompose the estimates into STEM and non-STEM, and show both the short-run (2001–2009, open circles) and long-run (2010–2019, filled circles) average treatment effects.

The income gains concentrate among non-STEM workers and, for immigrants, grow into the long run. This pattern serves two purposes. First, it provides an *identification test*: if unobserved shocks were driving both H-1B demand and the income gains, the gains should concentrate among STEM workers, whose labor demand is rising; instead, non-STEM workers gain more. Second, it

Figure 2: H-1B Exposure and Industry Income




Notes: Panels (a)–(d) show event-study estimates of H-1B exposure on log income from equation (3). Two specifications: OLS with skill-quartile $\times$ year fixed effects (gray, dashed) and SC with skill-quartile $\times$ year matching (blue, solid). SC weights are constructed at the industry $\times$ year level by matching pre-treatment outcome trajectories (1980–2000). Shaded areas show 95% confidence intervals. Regressions weighted by subpopulation size in 2000. Standard errors clustered at the industry level. “Pre-existing immigrants” arrived before 1998. Reference year: 2000.

points to a *mechanism*: H-1B workers, predominantly in STEM, perform specialized tasks that complement non-STEM employees, so more efficient execution raises the latter’s marginal product. For college-graduate natives in STEM, higher productivity and greater competition roughly offset, yielding a small negative net effect.

Pre-existing immigrants also benefit, by more than natives in most specifications. This is inconsistent with the alternative mechanism that immigrants substitute for other immigrants while complementing natives. That both groups gain is consistent with a *productivity* channel that benefits all existing workers through higher industry output, rather than a compositional shift between groups. Employment effects are sensitive to specifications; for native college graduates, the group most exposed to STEM competition, we find no robust evidence of displacement (Online Appendix, Figures OA.10–OA.11). The binary treatment analysis of Computer Systems Design confirms the same pattern: income gains concentrate in non-STEM occupations across estimators and donor pools (Online Appendix, Section B.1).

Figure 3: Treatment Effects on Income by Population Group and Occupation



Notes: Short-run (2001–2009, open circles) and long-run (2010–2019, filled circles) average treatment effects on log income from the SC-DiD’s second stage: weights are trained on the pre-treatment waves (1980–2000, joint STEM matching, skill-quartile grouping), and equation (4) is then estimated on the post-training sample (2000–2019) with year 2000 as the base period. Each facet shows a population group defined by nativity (native or pre-existing immigrant) and education (college graduate [CG] or less-than-college [LTC]), with two occupation categories: non-STEM and STEM. “Pre-existing immigrants” arrived before 1998. Employment shares in parentheses. Regressions weighted by subpopulation size in 2000. Error bars show 95% confidence intervals, clustered at the industry level. Online Appendix Section A.3 reports the backtest of the SC-DiD procedure and the sensitivity of these estimates to the SC training window.

The composition of STEM employment provides a second identification test. Were STEM supply elastic enough to flatten STEM wages under any shock, exposed industries would clear by drawing in incumbent native STEM workers; our productivity mechanism implies the influx is foreign instead. Native STEM employment is flat, while the native *share* of STEM employment falls with exposure: the marginal STEM worker in exposed industries is an entering immigrant, not a relocating native (Online Appendix A.4).

4.2 Supply Chain Propagation

If H-1B immigrants raise productivity rather than simply add labor, the shock should propagate forward to downstream buyers but not backward to upstream suppliers (Acemoglu et al., 2012), whereas a pure labor supply shock would propagate in both directions (Gulek, 2025).

To build intuition, consider three industries linked through a supply chain: industry 1 supplies industry 2, which supplies industry 3, with industry 2 exposed to H-1B immigration. Both produc-

tivity shocks and labor supply shocks generate forward propagation: when industry 2 becomes more productive or its labor becomes cheaper, its output price falls, and industry 3 benefits regardless of the source. The two channels differ in their backward predictions. Under Cobb-Douglas demand, expenditure shares are fixed, so a more productive industry 2 does not gain market share and its spending on upstream inputs from industry 1 does not change (Acemoglu et al., 2012). A productivity shock therefore generates forward but not backward propagation. A labor supply shock, by contrast, shifts industry 2’s input mix between labor and upstream goods, generating nonzero backward propagation whose sign depends on the substitutability between intermediates and labor (Gulek, 2025). Finding forward-only propagation is therefore consistent with a productivity shock, while nonzero propagation in both directions would suggest a labor supply channel.¹⁴

Figure 4 presents the results. Panels (a) and (b) show forward and backward I/O exposure across industries, plotted against the college-graduate share of each industry’s workforce. Both low- and high-skill industries have substantial variation in I/O exposure, and the quartile boundaries (dashed lines) show which industries are compared within skill groups in the preferred specification.

Panel (c) presents the SC-debiased estimates from equation (4). Forward exposure raises aggregate industry income in the short run (2001–2009) by 6.8 log points for all workers ($p = 0.046$) and 7.6 log points for natives ($p = 0.011$), growing to 10.9 ($p = 0.039$) and 9.7 ($p = 0.053$) log points in the long run (2010–2019). Backward exposure is not distinguishable from zero in both periods. The monotonic increase from short to long run is consistent with gradual diffusion of productivity gains through supply chain networks.


The magnitudes are economically meaningful. Because the most exposed industries are central input suppliers, H-1B shocks transmit widely: the median industry requires 1.6 cents of Computer & Data Processing Services output per dollar of its own final output, 1.2 times that industry’s share of total employment. Comparing workers across the supply chain, our estimates imply that a worker in an industry at the 75th percentile of forward exposure (insurance) gained 1.0 percent in earnings over 2001–2009 relative to a worker at the 25th percentile (child day care services), with the gap growing to 1.6 percent over 2010–2019.¹⁵ This forward/backward asymmetry is consistent with the H-1B immigration shock operating as a productivity shock. H-1B immigrants make exposed industries more productive, and this effect roughly offsets the competition they create for native workers in similar occupations. These findings are robust across SC training windows and most matching strategies (Online Appendix, Section A.6). Industry-level data from the QCEW confirm that forward exposure raises wages downstream (Online Appendix, Section A.6.3).

A corollary of the forward propagation result is that some control industries in the direct effects analysis (Section 4.1) receive indirect treatment through supply chain linkages. We replicate the STEM/non-STEM income analysis after excluding nine of the ten most downstream-exposed

¹⁴This is contingent on Cobb-Douglas final demand across industry products; see Online Appendix A.6.1 for details.

¹⁵This comparison applies the headline forward-exposure coefficients from Figure 4 (synthetic-control, college-share-adjusted)— $\hat{\beta} = 0.068$ for 2001–2009 and $\hat{\beta} = 0.109$ for 2010–2019—to the gap in forward exposure between the industries at the 75th and 25th employment-weighted percentiles, insurance ($F = 0.38$) and child day care services ($F = 0.23$), an interquartile range of 0.15. Thus $0.068 \times 0.15 \approx 0.010$ and $0.109 \times 0.15 \approx 0.016$.

Figure 4: Supply Chain Propagation: I/O Exposure and Income Effects



Notes: Panels (a) and (b): each bubble represents an ind1990 industry; size is proportional to year-2000 employment. I/O exposure constructed from 1997 BEA benchmark input-output tables. Industry 732 (Computer & Data Processing) is excluded; a further 28 industries (constituting 9 percent of employment, chiefly public administration) are dropped because they have no BEA input-output match, leaving 177 industries. Labels in panels (a) and (b) mark the four industries with the highest absolute exposure. Vertical dashed lines mark the quartile boundaries of the college-graduate share distribution, which define the skill-quartile groups used in the preferred specification. Panel (c): coefficients from the joint forward+backward specification on aggregate log income, estimated from equation (4) with F_i and B_i replacing E_i , using SC debiasing (1980-2000 training, skill-quartile matching, industry 732 excluded from donor pool); the second stage is estimated on the post-training sample (2000-2019, base year 2000). Open circles show short-run (2001-2009) effects; filled circles show long-run (2010-2019) effects. Error bars show 95% confidence intervals. Regressions weighted by total or native employment in 2000; standard errors clustered at the industry level.

industries (we retain Radio, TV, and Communication Equipment (341), the second most H-1B-exposed industry) and find similar results (Online Appendix, Section B.6).

4.3 Direct Effects on Innovation

Having established that H-1B immigration raises productivity, we ask through what channel. Kerr and Lincoln (2010) find that H-1B workers increase patenting, suggesting innovation drives the gains; we revisit this with our industry-level design and SC debiasing, which addresses the differential pre-trends that complicate earlier estimates.

Figure 5 presents the evidence. Panel (a) shows event-study estimates for log patents and panel (b) for log future citations, both at the three-digit SIC level. The OLS estimates reveal large pre-trends: H-1B-exposed industries were already on a markedly steeper innovation trajectory throughout the 1990s. The SC-DiD estimator, trained on 1990–1999, matches treated and control industries on these pre-treatment innovation dynamics. A key validation is that SC-DiD also improves fit in the out-of-sample 1980–1990 period, removing most of the OLS pre-gap despite not being directly matched on those years. This suggests the SC-DiD procedure recovers genuine counterfactual trajectories rather than overfitting to the training window.

Figure 5: H-1B Exposure and Innovation



Notes: Event-study estimates of H-1B exposure on patenting outcomes at the three-digit SIC industry level. Panel (a): log number of patents. Panel (b): log future citations. Patent data from Kogan et al. (2017), matched to SIC industries via Compustat. Patents are dated by *application* (filing) year, not grant year; the Kogan et al. (2017) database contains only eventually-granted patents. The sample consists of 79 industries observed in every year 1980–2019 (balanced panel). For citations, we further restrict to industries with positive citations in every year, leaving 72 industries. Regressions weighted by mean outcome value in 1990–2000. Reference year: 1999. Gray: OLS with industry, year, and skill-group $\times$ year fixed effects. Blue: synthetic control (trained on 1990–1999) with skill-group matching. Shaded areas indicate 95% confidence intervals. Standard errors clustered at the industry level. See Online Appendix A.5 for robustness to alternative SC training windows, skill-time fixed effects, and five-year grouped data.

After 2000, the SC-DiD estimates are near zero or slightly negative for both patents and ci-

tations; if anything, they suggest a modest decline. We find no evidence that H-1B immigration *directly* increased patenting in exposed industries. The result is not an artifact of the balanced-panel requirement: in Online Appendix A.5 we replicate it on five-year aggregated data, which relaxes that requirement (135 industries; 126 for citations), and find the same null.

The discrepancy with Kerr and Lincoln (2010) arises from the composition of additional patenting: in high-exposure regions, the extra patents come predominantly from industries that barely hire H-1B workers (Online Appendix, Section A.5). The regional increase is real, but because it occurs largely outside the industries the visas went to, it is unlikely to reflect a direct effect of H-1B immigration.

Together with the income results (Section 4.1) and the supply chain evidence (Section 4.2), the null patent finding completes a consistent picture. Computer Systems Design’s output improves largely through software development, and many software improvements are not patented. Moreover, the H-1B program selects workers for their ability to perform specialty tasks, not for their capacity to generate patentable inventions. A null patent result is therefore unsurprising if the productivity gains in exposed industries arise from better task execution rather than from patentable invention.

This null is the policy-relevant estimand: the cap moves cap-subject, private-sector workers, not the cap-exempt university and nonprofit researchers, so our variation is exactly the margin policy adjusts. It remains local and does not rule out patentable innovation elsewhere; the Kogan et al. (2017) data cover only publicly traded firms, and much innovation is never patented, as Khanna and Lee (2018) find for H-1B workers. Patents are this literature’s standard measure of innovation, yet for workers whose job is to improve a product, better task execution and unpatented innovation are often the same activity. Our result thus separates patentable invention from other improvement, not innovation from its absence: for the policy-relevant H-1B workers, the gains do not surface as patents.

5 Conclusion

This paper studies the long-run, industry-level effects of H-1B immigration on the U.S. economy. Exploiting the tripling and reversal of the H-1B cap between 1999 and 2004, we find that H-1B immigration operated as a productivity shock, with gains arising through task complementarity rather than patentable innovation. Three findings support our conclusion: non-STEM workers benefit more than STEM workers; the effects propagate forward to downstream industries but not backward, consistent with a productivity shock; and the productivity gains do not surface as patents, in contrast to the prior literature.

We cannot fully rule out unobserved post-treatment shocks differentially affecting exposed industries; but the leading alternative—a technology shock simultaneously raising H-1B demand and incomes—predicts STEM-concentrated gains, the opposite of what we find.

These findings have implications for current policy. The second Trump administration imposed

a \$100,000 per-worker fee on new H-1B petitions for entry from outside the United States, and introduced new barriers to the permanent residency pathway.¹⁶ Our estimates suggest that restrictions of this kind could dampen productivity gains and reduce incomes for native workers in both directly exposed industries and in the downstream industries that rely on their output, compounding the distortions that a per-worker fee creates under the program’s mobility restrictions (Clemens, 2026). The costs of restricting high-skilled immigration extend well beyond the industries that sponsor H-1B workers.

¹⁶Presidential Proclamation, “Restriction on Entry of Certain Nonimmigrant Workers,” September 19, 2025.

References

- Abadie, Alberto**, “Using synthetic controls: Feasibility, data requirements, and methodological aspects,” *Journal of Economic Literature*, 2021, 59 (2), 391–425.
- , **Alexis Diamond**, and **Jens Hainmueller**, “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program,” *Journal of the American Statistical Association*, 2010, 105 (490), 493–505.
- , —, and —, “Synth: An R package for synthetic control methods in comparative case studies,” *Journal of Statistical Software*, 2011, 42 (13), 1–17.
- and **Jérémy L’hour**, “A penalized synthetic control estimator for disaggregated data,” *Journal of the American Statistical Association*, 2021, 116 (536), 1817–1834.
- Acemoglu, Daron**, **Vasco M Carvalho**, **Asuman Ozdaglar**, and **Alireza Tahbaz-Salehi**, “The network origins of aggregate fluctuations,” *Econometrica*, 2012, 80 (5), 1977–2016.
- Akcigit, Ufuk**, **John Grigsby**, and **Tom Nicholas**, “Immigration and the rise of American ingenuity,” *American Economic Review*, 2017, 107 (5), 327–331.
- Akgündüz, Yusuf**, **Abdurrahman Aydemir**, **Seyit Cilasun**, and **Murat G Kırdar**, “Propagation of immigration shocks through firm-to-firm trade networks,” IZA Discussion Paper, IZA 2024.
- Albert, Christoph** and **Joan Monras**, “Immigration and spatial equilibrium: The role of expenditures in the country of origin,” *American Economic Review*, 2022, 112 (11), 3763–3802.
- Arkhangelsky, Dmitry**, **Susan Athey**, **David A Hirshberg**, **Guido W Imbens**, and **Stefan Wager**, “Synthetic difference-in-differences,” *American Economic Review*, 2021, 111 (12), 4088–4118.
- Baqaei, David Rezza** and **Emmanuel Farhi**, “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 2019, 87 (4), 1155–1203.
- Beerli, Andreas**, **Jan Ruffner**, **Michael Siegenthaler**, and **Giovanni Peri**, “The abolition of immigration restrictions and the performance of firms and workers: Evidence from Switzerland,” *American Economic Review*, 2021, 111 (3), 976–1012.
- 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.
- and **Kirk B Doran**, “The collapse of the Soviet Union and the productivity of American mathematicians,” *Quarterly Journal of Economics*, 2012, 127 (3), 1143–1203.

- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-experimental shift-share research designs,” *Review of Economic Studies*, 2022, *89* (1), 181–213.
- Bound, John, Breno Braga, Joseph M. Golden, and Gaurav Khanna**, “Recruitment of Foreigners in the Market for Computer Scientists in the United States,” *Journal of Labor Economics*, 2015, *33* (S1), S187–S223.
- , **Gaurav Khanna, and Nicolas Morales**, “Understanding the Economic Impact of the H-1B Program on the United States,” in Gordon Hanson, William Kerr, and Sarah Turner, eds., *High-Skilled Migration to the United States and Its Economic Consequences*, University of Chicago Press, 2018, pp. 109–175.
- Brinatti, Agostina and Xing Guo**, “Third-country effects of US immigration policy,” Technical Report, Bank of Canada 2024.
- Burstein, Ariel, Gordon Hanson, Lin Tian, and Jonathan Vogel**, “Tradability and the labor-market impact of immigration: Theory and evidence from the United States,” *Econometrica*, 2020, *88* (3), 1071–1112.
- Chodorow-Reich, Gabriel**, “Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?,” *American Economic Journal: Economic Policy*, 2019, *11* (2), 1–34.
- Clemens, Michael A.**, “Immigrant-Native Wage Gaps and Immigration Tariffs: Examining the Case for an H-1B Visa Tax,” IZA Discussion Paper 18435, Institute of Labor Economics (IZA) 2026.
- Conley, Timothy G and Christopher R Taber**, “Inference with “difference in differences” with a small number of policy changes,” *Review of Economics and Statistics*, 2011, *93* (1), 113–125.
- Doran, Kirk, Alexander Gelber, and Adam Isen**, “The effects of high-skilled immigration policy on firms: Evidence from visa lotteries,” *Journal of Political Economy*, 2022, *130* (10), 2501–2533.
- Dustmann, Christian and Albrecht Glitz**, “How do industries and firms respond to changes in local labor supply?,” *Journal of Labor Economics*, 2015, *33* (3), 711–750.
- Glennon, Britta**, “How do restrictions on high-skilled immigration affect offshoring? Evidence from the H-1B program,” *Management Science*, 2024, *70* (2), 907–930.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik instruments: What, when, why, and how,” *American Economic Review*, 2020, *110* (8), 2586–2624.
- Gulek, Ahmet**, “Effects of immigrants on non-host regions: Evidence from the Syrian refugees in Turkey,” *Available at SSRN 5094172*, 2025.

- **and Jaume Vives i Bastida**, “Synthetic IV estimation in panels,” 2024. Available at: https://economics.mit.edu/sites/default/files/inline-files/Synthetic_Wald%20%283%29.pdf.
- Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson**, “What Do We Learn From Cross-Regional Empirical Estimates in Macroeconomics?,” *NBER Macroeconomics Annual*, 2021, *35* (1), 175–223. NBER Working Paper No. 26881.
- Hornung, Erik**, “Immigration and the diffusion of technology: The Huguenot diaspora in Prussia,” *American Economic Review*, 2014, *104* (1), 84–122.
- Hulten, Charles R.**, “Growth Accounting with Intermediate Inputs,” *Review of Economic Studies*, 1978, *45* (3), 511–518.
- Hunt, Jennifer and Marjolaine Gauthier-Loiselle**, “How Much Does Immigration Boost Innovation?,” *American Economic Journal: Macroeconomics*, 2010, *2* (2), 31–56.
- Jaeger, David A., Joakim Ruist, and Jan Stuhler**, “Shift-Share Instruments and the Impact of Immigration,” Working Paper 24285, National Bureau of Economic Research 2018.
- Kandel, William A., Sarah A. Donovan, and Jill H. Wilson**, “U.S. Employment-Based Immigration Policy,” CRS Report R47164, Congressional Research Service nov 2024. Version 7, updated November 19, 2024.
- Kerr, William R and William F Lincoln**, “The supply side of innovation: H-1B visa reforms and US ethnic invention,” *Journal of Labor Economics*, 2010, *28* (3), 473–508.
- Khanna, Gaurav and Munseob Lee**, “High-Skill Immigration, Innovation, and Creative Destruction,” Working Paper 24824, National Bureau of Economic Research 2018.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman**, “Technological innovation, resource allocation, and growth,” *Quarterly Journal of Economics*, 2017, *132* (2), 665–712.
- Lewis, Ethan and Giovanni Peri**, “Immigration and the economy of cities and regions,” in “Handbook of Regional and Urban Economics,” Vol. 5, Elsevier, 2015, pp. 625–685.
- Mahajan, Parag**, “Migrant workers and economic resilience: Evidence from the 2020 H-2B visa lottery,” 2026. Upjohn Institute working paper; 25-414.
- , **Nicolas Morales, Kevin Shih, Mingyu Chen, and Agostina Brinatti**, “The impact of immigration on firms and workers: Insights from the H-1B lottery,” *FRB Richmond Working Paper*, 2024.
- Mayda, Anna Maria, Francesc Ortega, Giovanni Peri, Kevin Shih, and Chad Sparber**, “The Effect of the H-1B Quota on the Employment and Selection of Foreign-Born Labor,” *European Economic Review*, 2018, *108*, 105–128.

- Monras, Joan**, “Immigration and wage dynamics: Evidence from the Mexican peso crisis,” *Journal of Political Economy*, 2020, 128 (8), 3017–3089.
- Moser, Petra, Alessandra Voena, and Fabian Waldinger**, “German Jewish émigrés and US invention,” *American Economic Review*, 2014, 104 (10), 3222–3255.
- Nakamura, Emi and Jón Steinsson**, “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *American Economic Review*, 2014, 104 (3), 753–792.
- Ottaviano, Gianmarco IP and Giovanni Peri**, “Rethinking the effect of immigration on wages,” *Journal of the European Economic Association*, 2012, 10 (1), 152–197.
- Peri, Giovanni 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.
- Powell, David**, “Synthetic Control Estimation Beyond Comparative Case Studies: Does the Minimum Wage Reduce Employment?,” *Journal of Business & Economic Statistics*, 2022, 40 (3), 1302–1314.
- Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rodgers, Jonathan Schroeder, and Kari C.W. Williams**, “IPUMS USA: Version 16.0 [dataset],” 2025. Minneapolis, MN: IPUMS. <https://doi.org/10.18128/D010.V16.0>.
- Samuelson, Paul A.**, “Abstract of a Theorem Concerning Substitutability in Open Leontief Models,” in Tjalling C. Koopmans, ed., *Activity Analysis of Production and Allocation*, New York: Wiley, 1951, pp. 142–146.
- U.S. Government Accountability Office**, “H-1B Visa Program: Reforms Are Needed to Minimize the Risks and Costs of Current Program,” Technical Report GAO-11-26, U.S. Government Accountability Office, Washington, DC 2011.
- Wolf, Christian K.**, “The Missing Intercept: A Demand Equivalence Approach,” *American Economic Review*, 2023, 113 (8), 2232–2269.

Online Appendix

A Data, Methods, and Supplementary Results

This appendix collects the data, methods, and analyses that support the results in the main text.

A.1 Data and Measurement

A.1.1 H-1B Filings Data

Our source for the H-1B filings that define treatment exposure is the six-digit NAICS-level data on newly issued H-1Bs starting from 1998, first collected by Mahajan et al. (2024); the exposure weights are built from the 1998–2000 filings. We complement these with publicly available data from the USCIS Employer Hub, which provides information on all H-1B petitions since 2009, including employer name, two-digit NAICS industry, state, and whether the petition is for new employment or a renewal, and which we use for the post-2009 descriptive analysis. We retain only “new employment” petitions, excluding renewals and amendments.

The persistence of the treatment reflects the program’s institutional features. Most employment-based immigrants in recent decades were already living in the United States and adjusted from a nonimmigrant to lawful permanent resident (LPR) status rather than arriving from abroad: between fiscal years 2014 and 2024, 69 to 92 percent of employment-based LPRs adjusted status from within the United States, and adjustments from H-1B and L-1 status and their derivatives accounted for 80 percent of such adjustments under the EB1, EB2 and EB3 categories over fiscal years 2010–2019 (Kandel et al., 2024). The American Competitiveness in the Twenty-First Century Act of 2000 further allows H-1B extensions beyond six years for workers with pending green card applications, so the sponsored cohort can remain in the industry well past the nominal six-year term.

The raw data provide NAICS codes at varying digit lengths (two through six digits). Filings recorded at five or six digits, 84 percent of new-employment petitions with a recorded NAICS code in the 1998–2000 weight window, are truncated to their leading four digits. To construct four-digit weights for the filings recorded at coarser levels, 8 percent of the window’s petitions, we redistribute them. Specifically, filings with only two-digit codes are allocated to four-digit codes using the observed distribution of four-digit codes within each two-digit sector in that year; filings with only three-digit codes are similarly allocated to four-digit codes using the within-three-digit distribution. Redistributed counts are added to the directly observed four-digit counts. Petitions with a missing or unclassified NAICS code, roughly 13 percent of the window’s petitions, are excluded, so the weights are shares of classified petitions. We then compute each four-digit NAICS code’s share of total H-1B filings in each year and average these shares across 1998–2000 to obtain baseline weights w_k^{H1B} for each four-digit NAICS code k . The mapping from these NAICS-level weights to *ind1990*-level exposure is described in the Industry Crosswalks section below.

A.1.2 Census and ACS Data

We use the 5 percent samples of the 1980, 1990, and 2000 U.S. Censuses and the American Community Survey (ACS) from 2001 through 2019, obtained from IPUMS (Ruggles et al., 2025): roughly 0.4 percent samples in 2001–2004 and 1 percent samples from 2005 onward. The sample is restricted to individuals aged 18–65 who are employed (`empstat` = 1). For income variables, we further restrict to observations with positive wage income (`incwage` > 0) and use $\log(\text{incwage})$ as the dependent variable. Log employment outcomes are computed from all employed workers regardless of wage income. We use the IPUMS *ind1990* consistent industry classification. After applying all sample restrictions, we retain only industries with at least one valid observation in each of the 22 survey years, yielding a balanced panel of 206 industries.

We label each wave by its survey year throughout, but the two sources time income differently. Each decennial census reports income for the preceding calendar year, so the 1980, 1990 and 2000 waves report earnings in 1979, 1989 and 1999; the ACS instead records earnings over the twelve months preceding each household’s interview. Our omitted base period, the 2000 wave, therefore reports 1999 earnings, which overlap primarily with fiscal year 1999, the first year of the expanded cap (the last quarter of 1999 falls in fiscal year 2000). The base is thus mildly treated: fiscal year 1999 authorized 50,000 visas above the 65,000 baseline, roughly a tenth of the 490,000 total, and those workers were in place for part of a single year. This attenuates the post-treatment coefficients toward zero, so our estimates are conservative in this respect.

We define the following population groups:

- **Natives:** U.S.-born individuals.
- **Pre-existing immigrants:** Foreign-born individuals who arrived in the United States before 1998. This ensures that the pre-existing immigrant group predates the first H-1B cap increase in fiscal year 1999.
- **Less-than-college (LS):** Workers without a four-year college degree.
- **College graduate (HS):** Workers with a bachelor’s degree or higher.

STEM occupations are defined using the IPUMS *occ2010* classification, comprising the computer and mathematical, architecture and engineering, and life and physical science occupation groups (codes 1000–1240, 1300–1560, and 1600–1760); social science occupations (economists, psychologists, and other social scientists, codes 1800–1860) are not included. The STEM panel is balanced to require both STEM and non-STEM cells present in all 22 survey years within each *ind1990* industry, yielding 190 of the 206 industries, each contributing a STEM and a non-STEM cell (380 industry $\times$ occupation cells).

We apply person weights (`perwt`) throughout. Nativity is defined using IPUMS `bp1`: U.S.-born if `bp1` < 150. Education groups are constructed from IPUMS `educd` (detailed education): college graduate if `educd` $\geq$ 101 (bachelor’s degree or higher), less-than-college otherwise (`ltc` = 1 – `cg`).

For the 1980 Census, where education is recorded as highest grade completed rather than degree attained, we define college as `educd` = 100 (four years of college) and graduate as `educd` $\in$ [110, 113]; for all other years, college is `educd` = 101 and graduate is `educd` $\geq$ 114. We require four completed years in 1980 so that the college-graduate definition is consistent across the baseline years: `educd` = 90 denotes three years of college and no degree, and counting it as college would make the 1980 group broader than the bachelor’s-degree group used from 1990 onward. Pre-existing immigrants are foreign-born individuals who arrived in the United States before 1998. In the 1980 and 1990 Census waves, all foreign-born individuals satisfy this condition by construction. In the 2000 Census and subsequent ACS waves, we restrict to those with `yrimmig` < 1998. Operationally, the condition `yrimmig` < 1998 is applied uniformly across all years; no additional filter on `yrimmig` is needed. Observations with missing `ind1990` are dropped. Skill quartiles for the donor pool restriction are computed using `xtile` on the college-graduate share in the 2000 Census cross-section, where the share is calculated from all workers (native and immigrant) meeting the sample restrictions above. The exposure denominator L_i in equation (2) is total employment (native and immigrant) in industry i from the 2000 Census, subject to the same sample restrictions. Log employment outcomes are the log of the `perwt`-weighted count in each industry-year-subgroup cell; cells with zero count are set to missing. Regression weights are the `perwt`-weighted subpopulation count in industry i in the 2000 Census (not the raw person count).

A.1.3 QCEW Data

The Quarterly Census of Employment and Wages (QCEW) covers approximately 95 percent of all U.S. jobs and provides annual data on the number of establishments, total employment, total wages, weekly wages, and annual pay at the four-digit NAICS level. We use annual QCEW data from 1990 through 2024. The QCEW does not contain demographic information, so we cannot distinguish between natives and immigrants or between education/occupation groups.

A.1.4 Patent Data

Patent data come from Kogan et al. (2017), who provide patent counts, forward citations, and estimated market values for granted patents assigned to publicly traded firms. We date each patent by its *application* (filing) year, constructed from the filing date in the database, rather than by its grant (issue) year. We do so because the application year is much closer to the timing of the underlying innovative activity: the median lag between application and grant is around 2.5 years, and the research work itself precedes the application. Because the database contains only patents that were eventually granted, our sample corresponds to eventually-granted patents dated to their year of application; we apply no additional lag between innovation and application. We match patents to three-digit SIC industries via Compustat firm identifiers. The NAICS-4 to SIC-3 crosswalk is constructed from Compustat records: we extract the SIC-3 and NAICS-4 codes on each CRSP–Compustat link record, tabulate link-record counts in each (SIC-3, NAICS-4) cell, and aggregate H-1B exposure from NAICS-4 to SIC-3 using these counts as weights. Patents are

assigned to industries using the link record in force at the patent’s filing year. Skill groupings for patent regressions (a two-group split in the baseline, quartiles in robustness checks) are constructed from the same count-weighted college-graduate shares at the SIC-3 level. Our analysis covers the period 1980–2019 and includes 79 three-digit SIC industries in the annual balanced panel. Patent regressions are weighted by the mean of the raw outcome level (patent count or citation count) over 1990–2000, computed at the SIC-3 level. The citation panel is further restricted to industries with positive citations in every year, retaining 72 of 79 industries.

A.1.5 Input-Output Data

We use the 1997 Bureau of Economic Analysis (BEA) Benchmark Input-Output tables. The Leontief inverse matrix Ψ is obtained directly from the BEA 1997 Benchmark Industry-by-Industry Total Requirements table and transposed to the Acemoglu et al. (2012) convention (row = buyer, column = seller). We map BEA six-character commodity codes to four-digit NAICS using the BEA’s official concordance, and then to Census *ind1990* codes using the crosswalk described below. BEA sectors that do not map to any *ind1990* code (primarily agriculture and some miscellaneous services) are dropped. The diagonal of the Leontief inverse is zeroed to isolate indirect network propagation from each industry’s own multiplier effect. After aggregation through the crosswalk chain (BEA $\rightarrow$ NAICS-4 $\rightarrow$ *ind1990*), 187 *ind1990* codes have non-zero I/O exposure. In the 206-industry balanced panel, 178 industries carry a non-missing I/O value (one of them, private households, an exact zero); the supply-chain regressions further exclude industry 732, leaving 177. The 28 panel industries without an I/O value have no BEA counterpart (chiefly public administration, household services, and unclassified sectors) and are dropped from the network regressions, so the supply-chain estimates run on 177 industries rather than the 206 of the direct-effects panel. The exposure vector that is propagated through the network is built at the NAICS-4 level with 1999 QCEW employment as the denominator, rather than the Census 2000 employment used in equation (2). We use QCEW here because the BEA tables are NAICS-based, so a NAICS-native denominator avoids a further crosswalk; the Census-based measure is used wherever the analysis is at the *ind1990* level, where 1999 employment is not observed.

A.1.6 Industry Crosswalks

Constructing crosswalks between NAICS (used in H-1B and patent data), Census *ind1990* (used in Census/ACS data), and BEA commodity codes (used in I/O data) is a key data challenge. We describe each crosswalk in turn.

NAICS extraction from Census industry codes. The Census provides the *indnaics* variable, which records each respondent’s industry using a NAICS-based coding scheme. Most *indnaics* codes consist of leading numeric digits followed by an optional alphabetic suffix (e.g., “3399ZM” for miscellaneous manufacturing). We extract the leading numeric characters to obtain two-, three-, and four-digit NAICS codes, assigning each *indnaics* code to the finest NAICS level its digits

support. Two composite codes require special treatment: “52M2” maps to three-digit NAICS 523 (Securities, Commodity Contracts, and Other Financial Investments), and “611M1” maps to four-digit NAICS 6113 (Colleges, Universities, and Professional Schools).

From NAICS weights to *indnaics* weights. Each *indnaics* code is assigned an H-1B weight through hierarchical matching to the four-digit NAICS baseline weights w_k^{H1B} described above. *Indnaics* codes whose leading digits correspond to a four-digit NAICS code with positive H-1B weight receive that weight, split equally when multiple *indnaics* codes share the same four-digit code. Codes that do not match at four digits have their weight aggregated to the three-digit parent and matched to *indnaics* codes that lack a four-digit match; any still-unmatched weight cascades to two-digit. *Indnaics* codes that remain unmatched at all levels receive zero weight. Denote the resulting weight for *indnaics* code k as h_k .

From *indnaics* to *ind1990*: college-graduate weighting. We aggregate *indnaics*-level weights to *ind1990* using college-graduate employment shares from the 2000 Census. Restricting the sample to college graduates ($cg = 1$, employed, ages 18–65), we tabulate person-weighted employment in each (*indnaics*, *ind1990*) cell and compute within-*indnaics* shares ω_{kj} , representing the fraction of code k ’s college-graduate employment belonging to *ind1990* industry j . The industry-level H-1B weight is:

$$s_j = \sum_k \omega_{kj} h_k.$$

We use college-graduate employment rather than total employment because H-1B workers are predominantly college-educated, and the relevant mapping is through the occupational structure of skilled workers. Industry-level exposure is then $E_j = 490,000 \times s_j / L_j \times 100$, where L_j is total employment (all education levels) in *ind1990* industry j from the 2000 Census.

Plurality mapping. For the I/O analysis and QCEW regressions at the NAICS-4 level, we construct a simpler plurality mapping: each four-digit NAICS code is assigned to the single *ind1990* that accounts for the largest share of total employment (person-weighted) in the 2000 Census.

A.1.7 Summary Statistics

Table [OA.1](#) reports the distribution of treatment intensity and outcome variables across the four analysis levels used in the paper.

A.2 The Synthetic Control Estimator

The continuous treatment analysis (Census/ACS, patents, QCEW) uses the standard SC method of Abadie et al. (2010). The binary treatment analysis (Section [B.1](#)) uses the penalized extension of Abadie and L’hour (2021) with cross-validated penalty selection (Section [A.2.5](#)).

Table OA.1: Summary Statistics

	N	Mean	SD	p25	p50	p75	p90	Max
<i>Panel A: Census/ACS (ind1990, 206 industries, year 2000)</i>								
H-1B exposure (%)	206	0.261	0.972	0.022	0.076	0.199	0.511	12.861
Forward I/O exposure	178	0.348	0.199	0.237	0.294	0.398	0.564	1.923
Backward I/O exposure	178	0.867	1.403	0.064	0.482	0.871	2.165	13.184
College-graduate share	206	0.231	0.147	0.122	0.185	0.305	0.472	0.691
Total employment (thousands)	206	591.8	1037.7	116.0	251.7	654.4	1389.1	8576.0
Log income, native college grad.	206	10.603	0.352	10.353	10.703	10.871	10.977	11.108
Log income, native less-than-college	206	9.982	0.360	9.725	10.060	10.244	10.373	10.653
Log income, pre-imm. college grad.	206	10.503	0.343	10.214	10.550	10.745	10.938	11.147
Log income, pre-imm. less-than-college	206	9.878	0.287	9.683	9.877	10.104	10.228	10.651
Log employment, native college grad.	206	10.714	1.477	9.630	10.505	11.843	12.767	15.335
Log employment, native less-than-college	206	12.099	1.161	11.318	11.899	12.884	13.650	15.696
Log employment, pre-imm. college grad.	206	8.685	1.641	7.393	8.613	9.912	10.770	12.708
Log employment, pre-imm. less-than-college	206	9.964	1.353	9.102	9.941	10.832	11.746	13.838
<i>Panel B: Census/ACS, STEM decomposition (ind1990 $\times$ occupation, year 2000)</i>								
Log income, native non-STEM	190	10.149	0.325	9.979	10.216	10.379	10.507	10.767
Log income, pre-imm. non-STEM	190	10.037	0.297	9.811	10.036	10.236	10.441	10.688
Log employment, native non-STEM	190	12.392	1.213	11.540	12.326	13.221	13.966	15.782
Log employment, pre-imm. non-STEM	190	10.266	1.385	9.312	10.403	11.266	12.064	13.907
Total employment, non-STEM (thous.)	190	594.8	1051.6	115.3	261.4	640.6	1311.1	8413.5
Log income, native STEM	190	10.615	0.198	10.506	10.639	10.766	10.844	11.105
Log income, pre-imm. STEM	190	10.709	0.247	10.591	10.731	10.875	10.977	11.295
Log employment, native STEM	190	8.888	1.538	7.777	8.589	9.745	11.252	13.316
Log employment, pre-imm. STEM	190	6.870	1.754	5.666	6.536	8.008	9.480	11.971
Total employment, STEM (thous.)	190	32.9	87.2	2.7	6.2	20.8	88.8	801.9
<i>Panel C: Patents (SIC-3, 81 industries, year 2000)</i>								
H-1B exposure (%)	79	0.560	1.136	0.087	0.176	0.412	1.689	7.290
Log patents	79	5.003	1.869	3.689	4.836	6.001	7.922	9.390
Log forward citations	79	8.179	1.894	6.960	7.999	9.174	11.221	12.765
Number of patents	79	869.9	2109.4	40.0	126.0	404.0	2757.0	11967.0
<i>Panel D: QCEW (NAICS-4, balanced panel, year 2000)</i>								
H-1B exposure (%)	276	0.365	1.420	0.019	0.076	0.212	0.745	20.358
Total employment (thousands)	276	326.9	476.7	76.1	156.1	408.5	759.8	3771.6
Log annual pay	276	10.471	0.446	10.157	10.493	10.724	11.037	11.933
Log employment	276	11.915	1.413	11.240	11.958	12.920	13.541	15.143
Log establishments	276	8.837	1.649	7.702	8.883	10.047	11.018	12.939
College-graduate share	276	0.231	0.151	0.122	0.177	0.302	0.469	0.691

Notes: All moments computed from the year-2000 cross-section, unweighted. H-1B exposure (E_i) measures the expected percentage increase in industry employment due to the 1999–2004 cap expansion (equation 2). Forward and backward I/O exposure are constructed from the 1997 BEA Leontief inverse (equations 5–6), excluding industry 732. “Pre-imm.” denotes pre-existing immigrants (arrived before 1998). STEM occupations defined following the NSF Science and Engineering Indicators classification. Patent and QCEW panels restricted to balanced samples. College-graduate share is the share of workers with a bachelor’s degree or higher. Employment counts are person-weighted totals from the 2000 Census 5% sample.

A.2.1 SC Weight Construction for Continuous Treatment

For each industry i in the sample, we construct SC weights through the following steps.

Step 1: Demeaning. We subtract the pre-period mean of the matching variable (the outcome) for each industry. For Census/ACS outcomes with baseline years $\{1980, 1990, 2000\}$, the demeaned outcome is $\bar{y}_{it}^{dm} = y_{it} - \frac{1}{3} \sum_{s \in \{1980, 1990, 2000\}} y_{is}$. This is equivalent to including industry fixed effects

in the matching step and ensures that SC weights match on within-industry time trends rather than levels.

Step 2: Donor pool restriction. We restrict the donor pool for each industry to other industries in the same college-graduate share quartile. Quartile assignments are computed from the 2000 Census cross-section. This restriction is equivalent to skill-quartile-by-time fixed effects in the regression: identification comes only from comparing industries with similar skill composition. For robustness, we also report specifications with no grouping (all industries as donors) and median-split grouping in Table OA.2; the main results use quartile grouping.


Step 3: Matching. We form the matching matrix X with one row per donor industry and one column per baseline year, containing the demeaned outcome values. Each column is standardized by its cross-sectional standard deviation. We then solve the constrained least squares problem:

$$\min_w \|x_i - X'w\|^2 \quad \text{s.t.} \quad w_j \geq 0, \quad \sum_j w_j = 1 \quad (7)$$

where x_i is the standardized demeaned outcome vector for the target industry and X contains the corresponding vectors for donor industries in the same skill quartile. We solve this problem by quadratic programming, adding a ridge penalty of 4×10^{-2} to the quadratic form. The matching matrix has far fewer baseline years than donor industries, so the unregularized program is degenerate: many weight vectors attain the same pre-treatment fit, and an unregularized solver returns an arbitrary one of them. The ridge selects the unique minimum-norm solution, which makes the estimates reproducible across solvers. They are also insensitive to its size: every coefficient is flat in the ridge parameter from 10^{-6} to 1.6×10^{-1} , and the only discontinuity is at zero. Figure OA.1 shows the Figure 3 estimates at the endpoints of this range and at zero. Specifically, we use MATLAB's `quadprog` with function and step tolerances set to 10^{-14} and a maximum of 2,000 iterations. Weights are rounded to ten decimal places before application; if `quadprog` fails to converge, equal weights are assigned across all donors in the quartile. Standardization in Step 3 uses the cross-sectional standard deviation across all industries, not only those in the donor pool.

Step 4: Missing data. Some industry-subgroup cells have zero population in one or more years (primarily immigrant cells in small industries), producing missing values in the matching variable. Industries with a missing value in *any* year are excluded from the donor pool for all industries in their skill quartile, and their own weight vectors are set to NaN (so that debiased values are also NaN). We screen on the full panel rather than the baseline years alone because the counterfactual $\hat{y}_{it} = \sum_j \omega_{ij} y_{jt}$ is formed in every year, so a donor missing any year cannot support it. These industries are retained in the panel for OLS specifications but are effectively dropped from SC regressions because the debiased outcome is missing.

Figure OA.1: Figure 3 estimates across ridge-penalty values



Notes: Short-run (panel a) and long-run (panel b) average treatment effects on log income for the twelve population-group-by-occupation cells of Figure 3; the two panels share the coefficient axis. Each cell is estimated by four separate SC estimators that differ only in the weight solver's ridge penalty, and every estimate is shown with its own 95% confidence interval clustered at the industry level. Within a cell, from top to bottom: no ridge (open red, 0), a very small ridge (10⁻⁶), the main specification (4 × 10⁻²), and the largest ridge we run (0.16). At zero the matching problem has fewer baseline moments than donors, the weights are not uniquely identified, and the solver returns an arbitrary vertex of the solution set; the three positive ridges are mutually indistinguishable, confirming that the estimates are insensitive to the penalty's size. Everything else is as in Figure 3: joint STEM matching with skill-quartile grouping, second stage estimated on the post-training sample (2000–2019, base year 2000), weighted by subpopulation size in 2000.

A.2.2 Debiasing

Given weights $\hat{w}_{ij}$, we compute counterfactual values for *all* variables in the dataset—not just the matching variable—by taking the weighted average of donor industries:

$$\hat{z}_{it} = \sum_j \hat{w}_{ij} z_{jt}$$

for any variable $z \in \{y_{it}, y_{it}^{dm}, E_i, F_i, B_i, \dots\}$. The debiased variable is $\tilde{z}_{it} = z_{it} - \hat{z}_{it}$. We then estimate equations (3)–(4) on the debiased data with industry and year fixed effects only—no additional skill-quartile-by-time fixed effects, since these are absorbed by the within-quartile matching step. The pooled all-occupation regressions, which stack both occupation groups, additionally in-

teract each fixed effect with the STEM indicator, so that industry and year effects are allowed to differ by occupation group. This procedure is a special case of the Synthetic IV estimator of Gulek and Vives-i Bastida (2024), where the instrument equals the treatment. The cluster-robust standard errors we report are clustered at the industry level and computed from the second-stage regression on the SC-debiased outcomes; they treat the debiased outcomes as observed and do not adjust for the first-stage estimation of the weights. The sample-splitting structure of the procedure justifies this approximation: the SC weights are estimated on the pre-treatment sample, while the treatment-effect coefficients are estimated on the post-treatment sample, so weight-estimation noise and the post-treatment regression scores are largely independent (Gulek and Vives-i Bastida, 2024). This approximation is most credible when the pre-treatment fit is stable across training windows, the synthetic weights are dispersed across donors rather than concentrated on a few industries, and backtesting confirms that good in-sample fit does not reflect overfitting.

Applying the same weights to all variables ensures internal consistency. Because all donor industries in the weighted average belong to the same skill quartile (and the weights sum to one), the synthetic control inherits the skill component of the target industry. Subtracting the synthetic values therefore removes variation common to industries with similar skill intensity, including any skill-quartile-by-time trends.

A.2.3 Training Periods

The training period determines which baseline years are used for matching. We use the following windows:

Census/ACS. SC80 matches on the 1980, 1990, and 2000 Census waves; SC89 matches on the 1980 and 1990 waves only. When both training periods yield similar post-treatment estimates, this provides backtesting support for the identifying assumption.

Patents. For the annual patent panel, we use four training windows: 1980–1990, 1980–1994, 1980–1999, and 1990–1999. The preferred specification (1990–1999) uses the pre-treatment period for the H-1B cap shock.

QCEW. For the annual QCEW panel, we use two training windows: 1990–1999 and 1990–1995.

In all cases, the SC-DiD estimates the pre/post specification (equation (4)) on post-training-period data only, while OLS uses the full panel. This separation between the SC training and estimation samples improves the estimator’s finite-sample properties.

For annual data (patents and QCEW), demeaning subtracts the industry-specific mean of the outcome over all years in the training window. For example, with a 1990–1999 training window, we subtract $\bar{y}_i = \frac{1}{10} \sum_{t=1990}^{1999} y_{it}$ from all years.

Industry exclusions. Industry 732 (Computer & Data Processing Services) is the most H-1B-exposed industry and is excluded from both the regression sample and the SC donor pool in the supply chain specification (Section 4.2), so that the I/O coefficients capture only indirect propagation through input–output linkages. In all other specifications—the main event study (Figure 2), the STEM decomposition (Figure 3), and the patent analysis (Section 4.3)—industry 732 is re-

tained in the sample and donor pool. Retaining it means the most-treated industry can serve as a control for others, so we re-estimate the STEM decomposition with industry 732 removed from the donor pool. The estimates are essentially unchanged: the long-run effect for native college graduates in non-STEM occupations is 0.0112, against 0.0116 in the baseline. This check varies the donor pool, not the estimation sample. It asks whether the synthetic controls for other industries are contaminated by admitting the most exposed industry as a donor, and shows that they are not. It is not a test of what happens when industry 732 leaves the estimation sample, which would remove most of the identifying variation by construction.

A.2.4 Joint STEM Matching

In Figure 3, all rows—including the all-occupation aggregate—use joint STEM SC matching. The algorithm modifies the weight construction as follows:

1. For each industry, demean the outcome separately within each STEM group (STEM and non-STEM), subtracting the within-group, within-industry pre-period mean.
2. Stack the demeaned baseline values for both STEM groups into a single matching vector per industry. With three baseline years and two STEM groups, the matching vector has six elements.
3. Standardize each column of the stacked matching matrix by its cross-sectional standard deviation.
4. Solve the constrained least squares problem (7) for a single weight vector per industry, with the donor pool restricted to the same skill quartile.
5. Apply the resulting weights to debias all variables in both STEM groups.

This yields one weight vector per industry that simultaneously matches both STEM and non-STEM trajectories. The all-occupation estimate is obtained by pooling both STEM groups in the regression (without a STEM filter), using the jointly-constructed weights. This ensures that all three rows in each facet of Figure 3—all, non-STEM, and STEM—use identical counterfactual weights and are directly comparable.

A.2.5 Penalized SC for Binary Treatment

For the binary treatment analysis (Section B.1), where a single treated industry (NAICS 5415) is compared to a synthetic control, we use the penalized SC method of Abadie and L’hour (2021). The penalized weights solve:

$$\min_w \left\| X_{\text{treat}} - X'_{\text{donors}} w \right\|^2 + \lambda \sum_{j=1}^J w_j \|x_j - x_{\text{treat}}\|^2 \quad \text{s.t.} \quad w_j \geq 0, \quad \sum_j w_j = 1 \quad (8)$$

where $\lambda \geq 0$ penalizes placing weight on donors far from the treated unit. The penalty parameter is selected by leave-one-out cross-validation: for each candidate λ on a grid of 51 logarithmically-spaced values $\{0, 10^{-6}, 10^{-5.83}, \dots, 10^2\}$, we sequentially hold out each donor as a pseudo-treated unit, compute SC weights from the remaining donors, and evaluate the post-treatment prediction error. The λ that minimizes the average cross-validated MSE is selected.

A.2.6 Computational Details

SC weights are computed in MATLAB (R2024a or later). Regressions are estimated in Stata using `reghdfe`, which handles multi-way fixed-effect absorption and singleton identification. Figures are produced in R using `ggplot2`, `data.table`, and `cowplot`. I/O matrix construction and aggregation are implemented in Python (NumPy). Standard errors are clustered at the industry level throughout: *ind1990* for Census/ACS regressions, SIC-3 for patents, and NAICS-4 for QCEW. Cluster counts vary by regression because industries with missing cells drop out of SC specifications (Step 4): the all-occupation OLS panels use 206 *ind1990* clusters, the Figure 3 STEM decomposition runs on the 190-industry STEM panel with 56 to 175 effective clusters, the supply-chain regressions use 177, the patent panels use 79 (72 for citations), and the QCEW supply-chain regressions use 275 (NAICS 5415 is excluded by design). Confidence bands, where shown, are pointwise 95% intervals from clustered standard errors. The Census/ACS panel consists of the 1980, 1990, and 2000 Census waves and annual ACS from 2001–2019; years 1981–1989 and 1991–1999 are not observed and no interpolation is applied. The balanced panels contain 206 *ind1990* industries (Census/ACS), 79 SIC-3 industries (patents, annual), and 276 NAICS-4 industries (QCEW, 1990–2024).

A.3 Validating the Estimator

We validate the SC-DiD procedure by backtesting: we hold out a window in which nothing is treated, so that any effect the estimator reports there is bias. The Census is decennial, so this is something the QCEW panel can do and the ACS cannot. We report the QCEW backtest first, and then use what it teaches us to say what the ACS can and cannot support.

The QCEW backtest. The annual QCEW panel has roughly ten pre-2000 periods, so it can vary the length of the SC training window. For each cutoff T from 1991 to 1998 we train the SC on 1990– T , giving between two and nine pre-periods, placebo-treat the held-out years $T + 1$ through 1999, and estimate equation (4) on the debiased panel. The cap first rises in fiscal year 1999, so almost nothing in this window is treated and the coefficients are very nearly all bias; the exception is calendar 1999, which overlaps fiscal year 1999 and may be partly treated. Ending the placebo window at 1998 instead raises the share of bias the SC-DiD removes at every window, so the figures we report are the conservative ones. Figure OA.2 reports the result for all five QCEW outcomes.

The synthetic control removes bias at every training window, and how much it removes depends on how long that window is. Figure OA.2 plots the share of OLS’s placebo bias that the SC-DiD removes at each window; because both estimators are scored on the same held-out years, this share

is comparable across T even though the underlying bias levels are not.¹⁷ With two pre-periods the SC-DiD removes 12 percent of OLS’s bias, and with six or more it removes between 53 and 57 percent. Pooling all eight windows and all five outcomes, mean absolute placebo bias falls from 1.77 log points under OLS to 0.97 under the SC-DiD.

Sensitivity to the training window. For the ACS, observed only decennially, we ask whether the estimated effect depends on which pre-periods the SC is matched on. Figure OA.3 re-estimates the long-run effects of Figure 3 with the SC trained on {1980, 1990}, on {1990, 2000}, and on {1980, 1990, 2000}, which is our main specification. The estimator is otherwise unchanged.

The central pattern does not move. Non-STEM gains exceed STEM gains in all twelve group-by-window comparisons in the long run, and in eleven of the twelve in the short run, and the sign of the estimate is stable across all three windows in ten of twelve cells.¹⁸ For native college graduates in non-STEM occupations, the long-run effect is 0.66 log points training on {1980, 1990}, 1.03 training on {1990, 2000}, and 1.16 in our main specification. The dot-com boom runs through the 1990s, so the {1980, 1990} window selects controls in an environment that changed markedly thereafter; we expect it to fit worst, and it delivers the smallest estimate. The two windows that include 2000, the last pre-period before the cap rises, agree to well within a standard error. Magnitudes are more sensitive for pre-existing immigrants than for natives, and the only signs that move across windows are theirs.

The training window is a researcher’s choice, and this exercise shows that our estimates do not depend on it. The QCEW backtest is the complementary test of whether the debiasing works at all, and it says the debiasing is real but partial. We therefore read the income results as substantially, though not fully, purged of differential pre-trends.


A.4 The Composition of STEM Employment

This section reports the employment evidence behind the second identification test in Section 4.1. If elastic STEM supply flattened STEM wages under any positive shock, exposed industries would absorb the shock by drawing in incumbent native STEM workers; our mechanism instead implies the marginal STEM worker is a new immigrant. Table OA.3 reports synthetic-control-debiased long-run effects, per percentage point of exposure, on STEM-occupation employment. Native STEM employment is statistically flat, while the native *share* of STEM employment falls and post-1998 immigrant STEM arrivals rise sharply with exposure.

¹⁷The bias levels are not comparable across T : the held-out window is $T+1$ through 1999, so as T grows it shrinks toward 1999, when the FY1999 cap increase first binds and when exposed and unexposed industries have diverged the most. Placebo bias therefore rises with T for every estimator, OLS included, and that rise is mechanical rather than a statement about OLS.

¹⁸Both exceptions are pre-existing immigrant college graduates: the estimate pooled across occupations is essentially zero under {1980, 1990} and positive under the other two windows, and the STEM estimate is negative under {1980, 1990} and {1990, 2000} and negligibly positive under our main specification. The pooled estimate is significant only under our main specification, and the STEM estimate only under {1980, 1990}.

Figure OA.2: Backtesting the Synthetic Control in the QCEW Panel




Notes: The share of OLS's mean absolute placebo bias that the SC-DiD removes, by the number of pre-periods in the SC training window. The SC is trained on 1990– T and the held-out years $T + 1$ through 1999 are placebo-treated, for $T = 1991, \dots, 1998$; nothing is treated before 2000, so any coefficient is biased. Because both estimators are scored on the same held-out years, this share is comparable across T even though the underlying bias levels are not (the held-out window slides toward 1999 as T grows, so every estimator faces a larger true divergence at longer windows). The bold line pools the five QCEW outcomes; the thin lines are the individual series (four-digit NAICS level, 276 industries, weighted by 2000 employment). The SC uses the debiased exposure, the ridge penalty of Section A.2, and the strict donor rule, so that it exercises the estimator the paper reports; OLS is scored on the same industries the SC retains. Standard errors clustered at the industry level.

Table OA.3: H-1B exposure and the native/immigrant composition of STEM employment

Outcome (STEM occupations)	Long-run effect	Wild-cluster p
Native log-employment	−0.009 (0.008)	—
Immigrant log-employment	0.041 (0.010)	—
Native share of STEM (percentage points)	−0.33 (0.06)	0.102
Post-1998 immigrant arrivals (first stage)	0.72 (0.06)	< 0.01

Notes: Synthetic-control-debiased long-run effects (2010–2019) per percentage point of H-1B exposure, STEM occupations, *ind1990* panel, base year 2000, joint STEM matching with skill-quartile grouping (SC80, 1980–2000 training; the SC89 1980–1990 training gives a nearly identical native-share estimate of −0.39). Regressions weighted by the relevant group's 2000 STEM employment; standard errors in parentheses clustered by industry. Because H-1B exposure is concentrated in a few industries, cluster-robust standard errors overstate precision, so the native-share p -value is from a wild cluster restricted bootstrap- t (1,999 replications, Rademacher weights, Computer Systems Design retained). The gap between immigrant and native employment growth ($\hat{\beta}_{imm} - \hat{\beta}_{nat} = +5.0$ log points) is itself significant under a stacked wild-cluster test (one-sided $p = 0.005$). The first-stage outcome is post-1998 immigrant STEM employment as a share of contemporaneous STEM employment; because new arrivals are zero before 1998 it cannot be synthetic-control-debiased and is estimated by OLS, and its p -value is cluster-robust rather than wild-cluster (the bootstrap is run for the native-share outcome only).

Figure OA.3: Sensitivity of the Estimated Effects to the SC Training Window



Notes: Long-run (2010–2019) average treatment effects on log income from equation (4), re-estimated with the SC trained on three different sets of pre-periods: {1980, 1990}, {1990, 2000}, and {1980, 1990, 2000}, the last being the specification used throughout the paper. The estimator is otherwise identical to Figure 3, including joint STEM matching, skill-quartile grouping, weights, and clustering; each SC panel carries its own debiased exposure, so both the outcome and the regressor change with the window. This is a sensitivity check, not a backtest: it shows whether the estimate depends on the training window, not whether it is unbiased. Each group is defined by nativity and education (CG = college graduate, LTC = less than college), by occupation group. Error bars show 95% confidence intervals, clustered at the industry level.

A.5 Decomposing Regional Patent Growth

A natural question is whether the industry-level null extends to the regional level. We split states into two groups—above and below the median of mean realized H-1B intensity—and decompose annual patent counts into all patents and a three-way partition of them: NAICS 5415 (Computer Systems Design), other H-1B-exposed industries ($>1\%$ QCEW-based exposure), and non-exposed industries ($<1\%$). Figure OA.4 plots $\ln(\text{patents})$ normalized to zero in 1999. Panel (a) shows that total patenting grew faster in high-exposure states after 1999: by 2019 the above-median group stands 0.56 log points above the below-median group, against a pre-1999 gap that averages zero. The regional increase reported by Kerr and Lincoln (2010) is therefore real in our data. Panel (b) shows that it did not come from the industry that hires H-1B workers: patenting in NAICS 5415 fell sharply in high-exposure states after 2000, and the two groups end only 0.29 log points apart. Panel (c) shows other exposed industries ending 0.40 log points apart. Panel (d) locates the increase in industries with essentially no H-1B presence, which likewise track closely before 1999 and end 0.74 log points apart. A regional design that does not decompose patents by industry attributes all of this growth to the policy shock, when the industries that actually hire H-1B workers account for a minority of it, if any at all.

Table OA.2: Income Effects by Donor-Pool Grouping Rule

Group	Occupation	No grouping (all donors)	Median split (<i>cg-p2</i>)	Quartile (<i>cg-p4</i>)
Native CG	Non-STEM	1.40 (0.48)	0.98 (0.28)	1.16 (0.34)
Native CG	STEM	-0.16 (0.08)	-0.21 (0.08)	-0.12 (0.07)
Native LTC	Non-STEM	1.31 (0.67)	0.98 (0.93)	0.78 (1.10)
Native LTC	STEM	0.07 (0.08)	0.08 (0.08)	0.57 (0.10)
Immigrant CG	Non-STEM	1.88 (0.63)	1.91 (0.59)	1.49 (0.54)
Immigrant CG	STEM	-0.24 (0.17)	-0.03 (0.14)	0.03 (0.25)
Immigrant LTC	Non-STEM	2.61 (0.84)	1.90 (0.34)	1.94 (0.34)
Immigrant LTC	STEM	1.01 (0.30)	1.03 (0.30)	0.96 (0.29)

Notes. Long-run (2010–2019) average treatment effects on log income, in log points, from equation (4). The three columns differ only in the rule restricting which industries may serve as SC donors: no grouping (every other industry is an admissible donor), median split on the college-graduate share (*cg-p2*), and quartile grouping (*cg-p4*). Everything else is held fixed, including joint STEM matching, demeaning, the training window {1980, 1990, 2000}, the ridge penalty, the strict donor rule, weights, and clustering. The quartile column is the specification used throughout the paper and reproduces the estimates behind Figure 3. For every group, the estimated gain is larger in non-STEM than in STEM occupations under all three rules. Standard errors clustered at the industry level in parentheses.

Further patent robustness—across fixed-effects specifications, SC training windows, five-year grouped panels, and an R&D-workforce exposure denominator—is reported in the Additional Robustness Checks (Appendix B.3).


A.6 Supply-Chain Propagation

This subsection interprets the forward/backward asymmetry and reports its robustness and firm-level confirmation.

A.6.1 Interpreting the Forward/Backward Asymmetry

Figure OA.5 gives the intuition behind reading forward-only propagation as a productivity rather than a labor-supply shock, and makes explicit the assumption the reading relies on. Consider three industries linked in a chain—industry 1 supplies industry 2, which supplies industry 3—where each industry also sells to final consumers, and industry 2 receives the H-1B inflow.

Figure OA.4: Regional Patent Decomposition by Industry Group



Notes: Each panel plots $\ln(\text{patents})$ normalized to zero in 1999 for states above and below the median of mean realized H-1B intensity. Panel (a): all patents. Panel (b): patents from NAICS 5415 (Computer Systems Design) firms only. Panel (c): patents from other H-1B-exposed industries (>1% QCEW-based H-1B exposure, excluding 5415). Panel (d): patents from non-exposed industries (<1% QCEW-based H-1B exposure). Panels (b)–(d) partition the patents in panel (a). Patent data from Kogan et al. (2017), matched to four-digit NAICS industries via Compustat. The dashed vertical line marks 1999. We further examine whether the patent growth in non-exposed industries reflects forward (downstream) propagation from H-1B-exposed industries, relating industry patenting to downstream input-output exposure to H-1B; we find no statistically significant relationship.

Forward propagation does not discriminate between the two channels. Whether the H-1B inflow raises industry 2's productivity or lowers its labor costs, industry 2's output price falls, and its downstream buyer (industry 3) and final consumers benefit. Both a productivity shock and a labor-supply shock generate forward propagation, so the forward direction alone is uninformative about the mechanism.


Backward propagation runs through two routes. The first is a demand, or business-stealing, route: if industry 2 becomes cheaper and gains market share among final consumers, it scales up and buys more from its supplier (industry 1). The second is a production route: if H-1B lowers wages in industry 2, industry 2 substitutes between labor and intermediate inputs, changing its demand for industry 1's output (Gulek, 2025). The first route is governed by how consumers

reallocate spending across products; the second by how producers reallocate between labor and intermediates when wages move.

Cobb-Douglas final demand shuts down the first route. With Cobb-Douglas demand across industry products, expenditure shares are fixed, so a more productive industry 2 does not gain market share and does not scale up; the business-stealing route is off. Maintaining this assumption—as we and Gulek (2025) do—the forward/backward contrast becomes a test of a *productivity* channel against a *competition* (labor-supply) channel: a productivity shock leaves input purchases unchanged and propagates forward only, while a competition shock lowers wages, shifts the input mix, and propagates backward. Observed backward propagation of approximately zero therefore favors the productivity channel over the competition channel.

The assumption does real work, so the test is suggestive rather than dispositive. If final demand is not Cobb-Douglas, the business-stealing route reopens, and even a pure productivity shock would propagate backward as industry 2 expands; more generally, propagation is jointly determined by the network and by the elasticities of substitution, of which Cobb-Douglas is a special case (Baqae and Farhi, 2019). A near-zero backward estimate could also reflect a non-Cobb-Douglas demand spillover offsetting a wage-competition spillover. We do not estimate the structural elasticities that would separate these possibilities, leaving that to future work (Gulek, 2025; Baqae and Farhi, 2019). We therefore read the forward-only pattern as consistent with a productivity channel, corroborated by the concentration of income gains among non-STEM workers and the null effect on patents, rather than as decisive on its own.

Figure OA.5: Forward and backward propagation of an H-1B shock



Notes: Industry 1 supplies industry 2, which supplies industry 3; each industry also sells to final consumers. The H-1B inflow hits industry 2. *Forward*: industry 2's price falls, benefiting industry 3 and final consumers (any shock). *Backward* can run two ways: (i) a demand/business-stealing route in which a cheaper industry 2 gains market share among final consumers, scales up, and buys more from industry 1—shut down under Cobb-Douglas final demand; and (ii) a production route in which lower wages in industry 2 shift its mix between labor and intermediate inputs, changing its demand for industry 1's output. Under Cobb-Douglas final demand only route (ii) survives, so backward propagation distinguishes a productivity shock (none) from a competition/labor-supply shock (positive).

This section presents robustness checks for the supply chain propagation analysis in Section 4.2.

We estimate the following specification on the ind1990×year panel:


$$y_{it} = \alpha_i + \gamma_t + \beta_F F_i \times post_t + \beta_B B_i \times post_t + \varepsilon_{it} \quad (9)$$

where F_i and B_i are forward and backward I/O exposure as defined in equations (5)–(6), entered jointly. Industry 732 (Computer & Data Processing) is excluded from both the regression sample and, in SC specifications, from the donor pool.

A.6.2 Aggregate Income Across Specifications

Figure OA.6 shows the aggregate income coefficients across the full specification grid. Forward effects are positive in every specification and significant in most; the joint-STEM variants with skill-time fixed effects are attenuated and not individually significant. Backward effects are uniformly null.

Figure OA.6: I/O Exposure on Aggregate Income: Robustness Across Specifications




Notes: Post-2010 coefficients from the joint forward+backward I/O specification (equation (9)) on aggregate log income across the full specification grid. Panel (a): forward exposure. Panel (b): backward exposure. Blue circles: all-worker estimates. Red triangles: native-only estimates. Rows correspond to specifications varying the estimator (OLS, SC with training on 1980–2000, SC with training on 1990–2000, SC with training on 1980–1990, and joint STEM SC variants) and skill-time fixed effects (none, median split, quartile split). I/O exposure constructed from the 1997 BEA benchmark input-output tables. Industry 732 (Computer & Data Processing) excluded from the sample and SC donor pool. Regressions weighted by total or native employment in 2000. Error bars show 95% confidence intervals, clustered at the industry level.

A.6.3 QCEW Firm-Level Outcomes

Figure OA.7 overlays OLS and SC event study estimates for forward I/O exposure on log annual pay at the NAICS4 level, confirming the ACS income gains using independent firm-level data.

Figure OA.7: OLS vs. SC Overlay Event Study: Log Annual Pay, Forward I/O Exposure (QCEW)



Notes: Annual event-study estimates of forward I/O exposure (F_i , equation (5)) on log annual pay at the four-digit NAICS industry level, using firm-level data from the Quarterly Census of Employment and Wages (QCEW), 1990–2024. Each panel overlays OLS (gray) and SC debiased with 1990–1999 training window (blue, with 95% confidence band). Columns vary the skill-group fixed effects: no grouping, two-group, and four-group. Forward and backward I/O exposure are entered jointly; only forward coefficients are shown. I/O exposure constructed from the 1997 BEA benchmark input–output tables. NAICS 5415 (Computer Systems Design) excluded. Regressions weighted by NAICS4 employment in 2000. Standard errors clustered at the four-digit NAICS level. Reference year: 1999.

A.7 Quantifying the Missing Intercept

Section 1 notes that a cross-regional design differences out the part of the H-1B effect that operates through the production network. Here we quantify it without rebuilding a regional design. We take the industry-level direct exposure E_j and the indirect (network) exposure $D_j = \tilde{\Psi}_{\text{off}} E$ —the Leontief-propagated H-1B shock with the diagonal removed to avoid double-counting—and project both onto regions r using pre-existing industry employment shares ω_{jr} : $E^r = \sum_j \omega_{jr} E_j$ and $D^r = \sum_j \omega_{jr} D_j$. Under full pass-through, with a single coefficient β on direct and indirect exposure alike, a region-and-time fixed-effects design recovers a slope on E^r , while the true total effect is $\beta(E^r + D^r)$. The coefficient β cancels from the share missed, so the result is pure exposure geometry. Two forces matter: how much D^r leaks into the estimated slope (the contamination δ), and how much of the total effect is differenced into the intercept because D^r is nearly uniform across regions. Because input-output linkages are aspatial, D^r varies little across regions ($\text{SD}(D)/\text{SD}(E) \approx 0.08$ – 0.23), so δ is small (the slope is barely biased) but the level is lost. Table OA.4 reports both across four geographies and two data sources.

Table OA.4: Share of the total H-1B effect missed by a cross-regional design

Geography (source)	N	$\bar{E}$	$\bar{D}$	Net share	$SD(D)/SD(E)$	δ	% missed
<i>Panel A: employment-weighted</i>							
State (ACS)	51	0.338	0.284	45.7%	0.14	0.029	44.1
Commuting zone (ACS)	741	0.338	0.284	45.7%	0.15	0.069	41.9
PUMA (ACS)	2,071	0.338	0.284	45.7%	0.14	0.067	42.0
County (QCEW, NAICS4)	3,220	0.490	0.292	37.3%	0.08	0.039	34.9
<i>Panel B: unweighted</i>							
State (ACS)	51	0.302	0.279	48.1%	0.18	0.044	45.8
Commuting zone (ACS)	741	0.198	0.275	58.2%	0.23	0.071	55.2
PUMA (ACS)	2,071	0.323	0.283	46.7%	0.15	0.071	42.9
County (QCEW, NAICS4)	3,220	0.210	0.258	55.1%	0.15	0.014	54.5

Notes. $\bar{E}$ and $\bar{D}$ are mean direct and indirect (network) exposure across regions; “net share” is $\bar{D}/(\bar{E} + \bar{D})$; δ is the contamination of the cross-regional slope by network exposure; “% missed” is the share of the total effect a region-and-time fixed-effects design fails to recover. Weighted means are partition-invariant within a data source. The county rows use QCEW (NAICS4) exposure; all others use ACS (*ind1990*). The production network is one general-equilibrium channel, so these figures are the portion of the missing intercept that flows through it—a lower bound on the full intercept.

B Additional Robustness Checks

This appendix reports additional robustness checks. Each confirms a main-text result under an alternative specification or sample; none is required for the paper’s claims.

B.1 Binary Treatment: Computer Systems Design

The main analysis exploits continuous variation in H-1B exposure across industries. This section complements that design with a binary treatment analysis focusing on the single most-exposed industry: Computer Systems Design (NAICS 5415, *ind1990* = 732), which received about 45 percent of new H-1B filings with a reported industry over 1998–2000, the years from which our exposure weights are constructed, and about 39 percent of all filings; the difference is filings from firms without industry information. The binary setting maps directly into the classical synthetic control framework (Abadie et al., 2010; Abadie, 2021), providing a transparent illustration of the identification challenge and its resolution. We construct a synthetic counterfactual for Computer Systems Design from the remaining industries and trace the treated-minus-synthetic gap over time.


Estimators and inference. We apply four estimators. First, OLS estimates the two-way fixed effects event study with binary treatment, $D_i = \mathbf{1}[i = 5415]$; inference assigns placebo treatment to each donor industry in turn and uses the standard deviation of the resulting placebo gap distribution.¹⁹ Second, synthetic difference-in-differences (SDID) combines unit weights with time weights to reweight the donor pool (Arkhangelsky et al., 2021). Third, classic synthetic control constructs non-negative weights that minimize the distance between the treated industry’s pre-treatment trajectory and a weighted combination of donors (Abadie et al., 2010), implemented via the `Synth` R package (Abadie et al., 2011). Fourth, penalized synthetic control adds a distance-from-treated penalty to the weight objective (Abadie and L’hour, 2021), implemented via the `pensynth` R package with the `clarabel` quadratic programming solver; the penalty parameter λ is selected by leave-one-out cross-validation on the donor pool (see Section A.2.5). For both SC-type estimators, we present results under two training windows to assess sensitivity to the pre-treatment matching period. Confidence intervals for all estimators are constructed from the placebo gap distribution: each donor industry is sequentially treated as the pseudo-treated unit, the full estimation procedure is repeated, and the resulting distribution of placebo gaps provides the basis for 95% pointwise confidence intervals.

First stage: QCEW employment. Figure OA.8 presents the first stage of the binary analysis: the effect on log employment in Computer Systems Design using annual QCEW data from 1990 to 2024. Panel (a) shows OLS and SDID; panels (b) and (c) show classic SC and penalized SC, each under two training windows (1990–1999 and 1990–1996). All estimators show a sharp employment

¹⁹Clustered standard errors are unreliable with a single treated cluster (Conley and Taber, 2011). These placebo bands are the variance estimator of Arkhangelsky et al. (2021); alongside them we report exact permutation p -values in the notes to each exhibit, following Abadie et al. (2010).

increase beginning around 1999, coinciding with the first H-1B cap expansion. The two SC training windows produce similar post-treatment trajectories, confirming that the synthetic control tracks the counterfactual rather than overfitting to a particular pre-treatment period. The employment effect is consistent with Computer Systems Design absorbing a large share of the approximately 490,000 additional H-1B visas that the cap expansion authorized.

Figure OA.8: Binary Treatment First Stage: Log Employment in Computer Systems Design



Notes: Binary treatment event study comparing NAICS 5415 (Computer Systems Design) to its synthetic control, using annual QCEW data from 1990 to 2024. The outcome is log employment. Each panel groups estimators by type: panel (a) shows OLS and SDID (Arkhangelsky et al., 2021); panel (b) shows classic synthetic control (Abadie et al., 2010) under two training windows (1990–1999 and 1990–1996); panel (c) shows penalized synthetic control (Abadie and L’hour, 2021) under the same two training windows. Panels share the same y -axis. Estimates are normalized to zero in the reference year (1999). For OLS, the treated industry’s gap is computed as the difference between its outcome trajectory and the unweighted donor pool average, both demeaned at 1999. For SC estimators, the gap is the difference between the treated industry’s outcome and the weighted synthetic control, demeaned at the reference year. Shaded ribbons show 95% confidence intervals from the placebo gap distribution: each donor industry is sequentially assigned as the pseudo-treated unit, the full estimation is repeated, and the standard deviation of the resulting placebo gaps provides the pointwise confidence band ($\pm 1.96 \times \text{SD}$). The donor pool consists of all other balanced four-digit NAICS industries observed in the QCEW throughout the sample period. The dashed vertical line marks 1999 (reference year).

Income effects: STEM vs. non-STEM decomposition. We next apply the binary design to Census/ACS income outcomes (1980–2019) at the *ind1990* level. The treated unit is industry 732 (Computer & Data Processing Services); the 205 remaining industries serve as controls. Income effects are positive for all four population groups (native and pre-existing immigrant, college graduate and less-than-college), with magnitudes consistent with the continuous treatment estimates.

The synthetic control draws on industries in the highest college-graduate share quartile that were on similar pre-treatment income trajectories. In the preferred penalized SC specification (quartile-restricted donors, 1980–2000 training), the largest weights for native college-graduate income are assigned to Catalog and Mail Order Houses (43 percent) and National Security and International Affairs (4 percent), with the remainder spread across roughly thirty donors at 1–3 percent each. For native less-than-college income, the largest weights are Management and Public Relations Services (15 percent) and Miscellaneous Professional Services (6 percent), with the remainder spread thinly across the pool. The specific donor composition varies across outcomes and estimators because each

configuration matches on its own pre-treatment trajectory, but all draw from the same quartile-restricted set of high-skill industries.

The key identification test is the STEM/non-STEM decomposition. Figure OA.9 presents the long-run average treatment effects (ATT) on native income, decomposed by occupation type (STEM vs. non-STEM) and education group (less-than-college and college graduate). The left panel uses the full donor pool; the right panel restricts donors to the same skill quartile. Income gains are concentrated in non-STEM occupations. For less-than-college natives, non-STEM ATTs range from 9 to 24 log points across estimators, while STEM ATTs are near zero. College-graduate natives show the same ordering, with non-STEM effects of 5 to 14 log points against STEM effects between -4 and 3 . This decomposition holds across all four estimators (OLS, SDID, classic SC, penalized SC) and both donor pool definitions. The consistency of the STEM/non-STEM pattern in the binary setting—where identification comes from a single industry rather than cross-industry variation—reinforces the task complementarity interpretation: H-1B workers, who disproportionately hold STEM positions, raise the productivity of complementary non-STEM workers within the same industry.

B.2 Additional Income Results

This subsection reports the full event-study estimates, their pre/post summary, income-percentile effects, and the STEM decomposition underlying the main-text income results.

B.2.1 Event Studies: Full Results

Figures OA.10 and OA.11 show the full event-study estimates for native and pre-existing immigrant workers across six specifications: OLS, OLS + skill-time FE, SC80, SC80 + skill-time, SC89, SC89 + skill-time.

Figure OA.12 summarizes these dynamics as a single long-run coefficient per specification: the 2010–2019 average treatment effect from equation (4), reported for each group by occupation set.


B.2.2 Income Percentiles

To rule out compositional bias, we estimate the effects of H-1B exposure on the 10th, 25th, 50th, 75th, and 90th percentiles of the native wage distribution within each industry. If the mean income gains were driven by compositional shifts (e.g., low-earning workers exiting exposed industries), we would expect heterogeneous effects across quantiles. Figure OA.13 shows wage increases at every percentile, largest at the bottom of the distribution. Gains at the 75th and 90th percentiles, which exit of low earners cannot raise, indicate the mean effect is not purely compositional.

B.2.3 STEM Decomposition

This subsection presents the full STEM versus non-STEM decomposition of the income and employment event studies. The decomposition is the key identification test discussed in Section 4.1:

Figure OA.9: Binary Treatment ATT: Native Income by Occupation Type



Notes: Long-run average treatment effects (ATT) on native log income from the binary treatment analysis of Computer & Data Processing Services ($ind1990 = 732$). The ATT is computed as the mean of the post-treatment event-study gaps (2001–05, 2006–10, 2011–15, 2016–19) relative to the reference year (2000). Rows decompose by education group (less-than-college, college graduate) and occupation type (STEM, non-STEM). For SC-type estimators, separate rows show results under two training windows (1980–2000 and 1980–1990). Four estimators are shown: OLS (gray circles), classic SC (blue triangles), penalized SC (red squares), and SDID (green crosses). The left panel uses all 205 control industries as donors; the right panel restricts the donor pool to industries in the same skill quartile (college-graduate share quartile 4). Error bars show 95% confidence intervals from the placebo gap distribution: for each control industry, the full estimation procedure is repeated with that industry as the pseudo-treated unit, and the standard deviation of the resulting placebo ATTs provides the confidence band ($\pm 1.96 \times \text{SD}$). We also report exact permutation p -values, $p = (1 + \#\{|\tau_j| \geq |\tau_1|\}) / (1 + J)$ over all J placebo industries, pooling the treated unit into its own reference distribution (Abadie et al., 2010). For less-than-college non-STEM these are 0.016–0.084 with the full donor pool ($J = 189$) and 0.122–0.327 with the quartile-restricted pool ($J = 48$); for college-graduate non-STEM, 0.126–0.571; every STEM p -value exceeds 0.42. Census/ACS data from 1980–2019 at the $ind1990$ level. STEM occupations as defined in Section A.1. “Pre-existing immigrants” arrived before 1998. The vertical line at zero indicates no effect.


if H-1B workers generate productivity gains through task complementarity (rather than through unobserved technology shocks), we expect gains to be concentrated among non-STEM workers.

B.3 Additional Patent Robustness

B.3.1 Robustness Across Specifications and Training Windows

This section presents robustness checks for the patent event studies shown in Figure 5 of the main text.

Figure OA.10: Event-study estimates: Native workers (weighted)




Notes: Event-study estimates of H-1B exposure on native worker outcomes from equation (3). Panels (a)–(b) show log income for less-than-college and college-graduate workers; panels (c)–(d) show log employment for the same groups. Each panel overlays six specifications: OLS (two-way fixed effects only), OLS with skill-quartile $\times$ year fixed effects, SC trained on 1980–2000 without skill-time grouping, SC trained on 1980–2000 with skill-quartile $\times$ year matching, SC trained on 1980–1990 without skill-time grouping, and SC trained on 1980–1990 with skill-quartile $\times$ year matching. SC weights are constructed at the industry level by matching pre-treatment outcome trajectories. Regressions weighted by subpopulation size in 2000. Standard errors clustered at the industry level. Reference year: 2000.

Figure OA.15 varies the skill-time fixed effects specification across three alternatives: no skill-time fixed effects (two-way FE only), two-group skill-time, and four-group skill-time fixed effects. Rows 1–2 show annual data; rows 3–4 show five-year grouped data, which relaxes the balanced panel requirement and increases the sample from 79 to 135 industries for patents (72 to 126 for citations). Across all specifications and both panel definitions, the pattern is the same: OLS exhibits strong pre-trends, and the SC estimator eliminates them within the 1990–1999 training window (the earlier five-year groups, outside the training window, retain a modest gap). Post-2000, the SC estimates are near zero or mildly negative throughout, confirming that patenting does not rise.

Figure OA.16 overlays three SC training windows: SC trained on 1990–1999 (the main specification), SC trained on 1980–1994, and SC trained on 1980–1999. The 1990–1999 and 1980–1999 training windows produce similar, near-zero post-treatment estimates; the 1980–1994 window, which excludes the late 1990s, drifts positive after 2010. The two windows that include the late-1990s run-up agree closely, supporting the null result.

Figure OA.11: Event-study estimates: Pre-existing immigrant workers (weighted)




Notes: Event-study estimates of H-1B exposure on pre-existing immigrant worker outcomes from equation (3). Pre-existing immigrants are defined as foreign-born individuals who arrived in the United States before 1998. Panels (a)–(b) show log income for less-than-college and college-graduate workers; panels (c)–(d) show log employment for the same groups. Six specifications as in Figure OA.10: OLS, OLS with skill-quartile $\times$ year fixed effects, and four SC variants (two training windows $\times$ with/without skill-quartile matching). Regressions weighted by subpopulation size in 2000. Standard errors clustered at the industry level. Reference year: 2000.

B.3.2 H-1B intensity relative to the R&D workforce

Our patent exposure measures new H-1B workers as a share of total industry employment. This denominator may understate H-1B intensity within the workforce that actually produces patents. In industries such as food manufacturing, transportation, and energy, research personnel are a small fraction of employment, so a cohort of H-1B workers that is large relative to the R&D workforce can be negligible relative to all workers. If the patent response scales with H-1B intensity among research workers, an all-worker denominator dilutes it and biases the estimate toward zero. Other studies that normalize by total employment, such as Kerr and Lincoln (2010), are less exposed to this concern because they operate at the city level, where the research share of employment varies far less than it does across industries.

We construct an alternative exposure that normalizes by the R&D workforce, defined as workers in science and engineering STEM occupations (Section A.1) and measured in the ACS: $E_i^{R\&D} = 490,000 w_i / L_i^{S\&E}$. We re-estimate the full synthetic-control patent event study with this exposure. To isolate the role of the denominator, we also report an exposure that retains the same numerator but normalizes by ACS *total* employment, so that the only difference between the two ACS measures

Figure OA.12: Long-run ATT estimates across specifications




Notes: Long-run (2010–2019) average treatment effects of H-1B exposure on log income (panel a) and log employment (panel b), from equation (4) with short-run (2001–2009) and long-run (2010–2019) interactions; the figure reports the long-run coefficient. These are the pre/post counterparts of the event-study estimates in Figures OA.10–OA.11. Six specifications as in Figure OA.10: OLS, OLS with skill-quartile $\times$ year fixed effects, and four SC variants (two training windows $\times$ with/without skill-quartile matching). SC estimates use joint STEM/non-STEM matched weights; the All-occupations estimates pool the occupation panel with occupation-interacted fixed effects. Horizontal bars show 95% confidence intervals; the horizontal scale differs across occupation columns. Regressions weighted by subpopulation size in 2000. Standard errors clustered at the industry level. OLS estimates use the full 1980–2019 panel, the conventional two-way fixed-effects baseline; SC estimates are the SC-DiD’s second stage, estimated on the post-training sample (2000–2019) with base year 2000.

is the denominator.²⁰

²⁰The R&D-workforce denominator is defined for 63 of the 79 patenting industries, because the ACS ind1990

Figure OA.13: H-1B exposure effects on native income percentiles (weighted)



Notes: Event-study estimates of H-1B exposure on percentiles of the native income distribution within each industry, from equation (3). Panels (a)–(e) show the 10th, 25th, 50th, 75th, and 90th percentiles, respectively. Six specifications as in Figure OA.10: OLS, OLS with skill-quartile $\times$ year fixed effects, and four SC variants (two training windows $\times$ with/without skill-quartile matching). If the mean income gains were driven by compositional shifts rather than genuine wage increases, we would expect heterogeneous effects across quantiles. Regressions weighted by native employment in 2000. Standard errors clustered at the industry level. Reference year: 2000.

Table OA.5 reports post-2000 effects on log patents. Because the R&D workforce is a small fraction of employment, the R&D-normalized exposure has a much larger scale: its standard deviation is roughly double that of the all-worker measure, so its coefficients are mechanically smaller and must be read against that scale. The final columns rescale each estimate to a one-standard-deviation increase in exposure. Under the synthetic-control specification, a one-SD increase in R&D-normalized classification is coarser than the NAICS-based exposure used in the main analysis; these 63 industries account for 92% of patents.

exposure changes post-2000 patents by -0.09 log points (s.e. 0.10), statistically indistinguishable from zero and from both the all-worker ACS estimate and our baseline. The OLS estimates are positive for every exposure, reflecting the cross-industry size confounding that the synthetic control is designed to remove. Rescaled to a one-SD increase, they are similar across denominators (0.55 to 0.68 log points), confirming that the three measures capture the same underlying variation. The synthetic control returns all three to zero.

The null is not an artifact of the denominator, because the premise behind the concern does not hold in our setting: the industries that patent are themselves R&D-intensive. STEM scientists and engineers are 11% of employment in the median patenting industry, against 3% in the median industry economy-wide, and the all-worker and R&D-worker exposures rank patenting industries similarly (rank correlation 0.78; correlation in levels 0.93). The industries whose exposure rises most under the R&D normalization are low-research service industries that produce essentially no patents and contribute nothing to identification. We conclude that our finding of no patent response is robust to measuring H-1B intensity relative to the research workforce rather than to all workers.

Table OA.5: Patent response under alternative exposure denominators

Exposure (denominator)	Mean	SD	OLS	SC	OLS/SD	SC/SD
H-1B / all workers, QCEW (baseline)	2.11	2.32	0.238 (.067)	-0.058 (.032)	0.55 (.15)	-0.13 (.07)
H-1B / all workers, ACS	3.41	4.40	0.154 (.038)	-0.015 (.023)	0.68 (.17)	-0.06 (.10)
H-1B / R&D workers, ACS (STEM S&E)	8.00	8.09	0.080 (.020)	-0.011 (.012)	0.65 (.16)	-0.09 (.10)

Notes. Post-2000 (2000–2019) effects on log patents, per unit of exposure (columns OLS, SC) and per one-standard-deviation increase exposure (final two columns). Means and standard deviations are patents-weighted. Standard errors clustered by industry in parentheses. OLS is estimated on the full 1980–2019 panel, the conventional baseline; the SC estimates are the SC-DiD’s second stage, estimated the post-training sample (1999–2019, base year 1999). The SC estimates are never positive: the baseline is marginally negative ($p = 0.0$) and the ACS-denominator estimates are indistinguishable from zero ($p = 0.52$ and 0.35). The R&D-worker denominator is defined for 79 patenting industries (92% of patents).

B.4 Additional Supply-Chain Robustness: Subgroup Decomposition

Figure OA.17 decomposes the I/O exposure coefficients by nativity and education across OLS and SC specifications. The positive forward effects documented in the main text are robust across estimation approaches. Backward exposure coefficients cluster around zero for every specification and subgroup.

B.5 Alternative Exposure: 1998 Shares Only

The baseline treatment variable defines industry-level H-1B exposure using the average share of H-1B visas across 1998–2000. A potential concern is that the 1999 and 2000 shares may partly reflect the policy shock itself: the first cap increase took effect in fiscal year 1999, so averaging over 1998–2000 conflates pre-shock and early-shock filing patterns. To address this concern, we replicate the full analysis using 1998 shares only to construct the exposure variable in equation (2). Because the input-output exposure variables are Leontief-weighted sums of industry-level H-1B

shares, redefining the base shares also changes the forward and backward supply chain exposure measures. All downstream analysis—synthetic control debiasing, regression estimation, and figure construction—is re-run using these alternative exposure variables.

Figure OA.18 reproduces Figure 3 from the main text using 1998-only exposure. The STEM/non-STEM decomposition is qualitatively identical to the baseline: income gains concentrate among non-STEM workers for both natives and pre-existing immigrants, while native college-graduate STEM workers experience near-zero or negative effects. Point estimates are slightly larger than the baseline—for example, the long-run effect on native less-than-college income is 0.8 log points (SE = 0.72) compared to 0.7 log points (SE = 0.73) in the main specification—but the differences are well within confidence intervals.

Figure OA.19 reproduces the patent event studies (Figure 5) using 1998-only exposure. The results are unchanged: OLS exhibits strong pre-trends in both log patents and log citations, the SC estimator eliminates these pre-trends, and the post-treatment SC estimates are near zero. The null effect on innovation is robust to the alternative exposure definition.

Figure OA.20 reproduces the supply chain analysis (Figure 4) using 1998-only exposure. Forward exposure continues to raise aggregate industry income, while backward exposure is not distinguishable from zero. The long-run forward coefficient for aggregate income in the preferred SC specification is nearly identical to the baseline—11.2 log points (SE = 5.1) compared to 10.9 log points (SE = 5.3). The forward/backward asymmetry that distinguishes productivity shocks from labor supply shocks is preserved.

Together, these results confirm that our findings do not depend on the specific years used to construct the exposure variable. The close correspondence between the 1998-only and 1998–2000-average results also reinforces the stability of the industry-level H-1B distribution documented in Figure 1, Panel (c): because industry shares changed little between 1998 and 2000, the two exposure definitions produce nearly identical variation across industries.

B.6 Excluding Downstream-Exposed Industries


The direct effects analysis in Section 4.1 compares differentially exposed industries within skill quartiles, attributing income differences to direct H-1B exposure. A concern is that some industries classified as controls receive indirect treatment through supply chain linkages, violating SUTVA. When downstream industries benefit from H-1B-exposed suppliers (as documented in Section 4.2), including heavily downstream-exposed industries in the control group could attenuate the estimated direct effects or distort the STEM/non-STEM decomposition. To address this, we replicate the STEM/non-STEM income analysis after dropping the nine most downstream-exposed industries from the sample (the top decile of forward I/O exposure, excluding Radio, TV, and Communication Equipment (341), which we retain because it is itself the second most H-1B-exposed industry). The dropped industries are: Petroleum Refining (200), Miscellaneous Petroleum and Coal Products (201), Computers and Related Equipment (322), Electrical Machinery, Equipment, and Supplies (342), Scientific and Controlling Instruments (371), Railroads (400), Gas and Steam Supply Systems

(451), Theaters and Motion Pictures (800), and Libraries (852).


Figure OA.21 reproduces Figure 3 using this restricted sample. The core pattern is preserved: income gains are concentrated among non-STEM workers, and native college-graduate STEM workers experience near-zero or negative effects. For native college graduates, the estimates are essentially unchanged—the long-run non-STEM effect is 1.07 log points ($SE = 0.27$) compared to 1.16 log points ($SE = 0.34$) in the baseline. For native less-than-college workers, point estimates are moderately smaller but remain positive and the confidence intervals overlap substantially with the baseline. For pre-existing immigrants, the pattern is qualitatively identical, with non-STEM gains dominating across both education groups.

The stability of the non-STEM concentration pattern after removing heavily downstream-exposed industries provides additional support for the task complementarity interpretation. If the direct effects were contaminated by supply chain spillovers flowing through downstream-exposed control industries, excluding these industries would alter the estimated treatment effects. The fact that the STEM/non-STEM decomposition is robust to this exclusion indicates that the direct effects and supply chain effects operate through distinct channels.


Figure OA.14: STEM vs. non-STEM event-study estimates: Native workers (weighted)




(a) Native LS income (non-STEM)




(b) Native LS income (STEM)




(c) Native HS income (non-STEM)




(d) Native HS income (STEM)




(e) Native LS employment (non-STEM)



(f) Native LS employment (STEM)




(g) Native HS employment (non-STEM)



(h) Native HS employment (STEM)


Notes: Event-study estimates of H-1B exposure on native worker outcomes, decomposed by occupation type, from equation (3). Left column: non-STEM occupations. Right column: STEM occupations. Panels (a)–(b): less-than-college income. Panels (c)–(d): college-graduate income. Panels (e)–(f): less-than-college employment. Panels (g)–(h): college-graduate employment. Six specifications as in Figure OA.10: OLS, OLS with skill-quartile $\times$ year fixed effects, and four SC variants (two training windows $\times$ with/without skill-quartile matching). Regressions weighted by subpopulation size in 2000. Standard errors clustered at the industry level. Reference year: 2000.

Figure OA.15: Patent event studies: Robustness across specifications




Notes: Event-study estimates of H-1B exposure on patenting outcomes at the three-digit SIC industry level. Each panel overlays OLS (gray) and SC (blue) estimates. Columns vary the skill-time fixed effects specification: no skill-time fixed effects (two-way FE only), two-group, and four-group skill-time fixed effects. Rows 1–2: annual data (79 industries for patents, 72 for citations, balanced panel 1980–2019). Rows 3–4: five-year grouped data, which relaxes the balanced panel requirement and increases the sample to 135 industries for patents and 126 for citations. SC training window: 1990–1999. Patent data from Kogan et al. (2017), matched to SIC industries via Compustat. Regressions weighted by the mean outcome level over the pre-period: 1990–2000 for the annual rows and 1985–1999 for the five-year rows. Shaded areas indicate 95% confidence intervals. Standard errors clustered at the industry level. Reference year: 1999 (annual) or 1995–1999 group (five-year).

Figure OA.16: Patent event studies: Robustness across SC training windows




Notes: Event-study estimates of H-1B exposure on patenting outcomes at the three-digit SIC industry level. Each panel overlays OLS (gray) and three SC specifications trained on different pre-treatment windows: 1990–1999 (blue, solid, with 95% confidence band), 1980–1994 (orange, dashed), and 1980–1999 (green, dashed). All specifications use two-group skill-time fixed effects. Left column: annual data. Right column: five-year grouped data. Top row: log patents. Bottom row: log citations. Patent data from Kogan et al. (2017), matched to SIC industries via Compustat. Regressions weighted by the mean outcome level over the pre-period: 1990–2000 for the annual rows and 1985–1999 for the five-year rows. Standard errors clustered at the industry level.

Figure OA.17: I/O Exposure on Income: Robustness Across Specifications




Notes: Post-2010 coefficients from the joint forward+backward I/O specification (equation (9)) on log income, decomposed by education group. Panel (a): forward exposure. Panel (b): backward exposure. Within each education facet, two population definitions are shown: all workers and native workers only. Rows vary the estimation approach: OLS (with and without skill-quartile $\times$ year fixed effects) and SC Approach 1 (separate matching, training on 1980-2000 and 1980-1990, with and without skill-quartile matching). I/O exposure constructed from the 1997 BEA benchmark input-output tables. Industry 732 (Computer & Data Processing) excluded from both the regression sample and the SC donor pool. Regressions weighted by outcome-specific employment in 2000. Error bars show 95% confidence intervals, clustered at the industry level.

Figure OA.18: Treatment Effects on Income by Population Group and Occupation (1998-Only Exposure)




Notes: Replication of Figure 3 using 1998-only H-1B shares to construct the exposure variable, instead of the 1998–2000 average used in the main specification. Short-run (2001–2009, open circles) and long-run (2010–2019, filled circles) average treatment effects on log income from equation (4), using joint STEM SC matching with skill-quartile grouping. Each facet shows a population group defined by nativity (native or pre-existing immigrant) and education (college graduate [CG] or less-than-college [LTC]), with two occupation categories: non-STEM and STEM. “Pre-existing immigrants” arrived before 1998. Employment shares in parentheses. Regressions weighted by subpopulation size in 2000. Error bars show 95% confidence intervals, clustered at the industry level.

Figure OA.19: H-1B Exposure and Innovation (1998-Only Exposure)




Notes: Replication of Figure 5 using 1998-only H-1B shares to construct the exposure variable. Event-study estimates of H-1B exposure on log patents (panel a) and log forward citations (panel b) at the three-digit SIC industry level, using the KPSS patent database. Patents are dated by application (filing) year, not grant year, as in the main analysis. Six specifications: OLS (with three skill-group fixed effect structures) and SC debiased (with three training window and skill-group combinations). All specifications include industry and year fixed effects. Regressions weighted by SIC-3 employment in 2000. Standard errors clustered at the industry level. Shaded bands show 95% confidence intervals for the preferred SC specification. Reference year: 2000.

Figure OA.20: Supply Chain Propagation: I/O Exposure and Income (1998-Only Exposure)



Notes: Replication of Figure 4 (panel c) using 1998-only H-1B shares to construct both the direct and I/O exposure variables. Short-run (2001–2009, open circles) and long-run (2010–2019, filled circles) average treatment effects from the joint forward/backward I/O specification (equation (9)). Forward (downstream, F_i) and backward (upstream, B_i) exposure constructed from the 1997 BEA benchmark input-output tables. Industry 732 (Computer & Data Processing) excluded. Regressions weighted by outcome-specific employment in 2000. Error bars show 95% confidence intervals, clustered at the industry level.

Figure OA.21: Treatment Effects on Income by Population Group and Occupation (Excluding Downstream-Exposed Industries)



Notes: Replication of Figure 3 after dropping the nine most downstream-exposed industries (top decile of forward I/O exposure from the 1997 BEA Leontief inverse, excluding industry 732). Dropped industries: Petroleum Refining (200), Miscellaneous Petroleum and Coal Products (201), Computers and Related Equipment (322), Electrical Machinery, Equipment, and Supplies (342), Scientific and Controlling Instruments (371), Railroads (400), Gas and Steam Supply Systems (451), Theaters and Motion Pictures (800), and Libraries (852). Radio, TV, and Communication Equipment (341) falls in the top decile but is retained, because it is itself the second most H-1B-exposed industry. These industries are excluded from both the estimation sample and the SC donor pool. Short-run (2001–2009, open circles) and long-run (2010–2019, filled circles) average treatment effects on log income from equation (4), using joint STEM SC matching with skill-quartile grouping. Employment shares in parentheses reflect the restricted sample. Regressions weighted by subpopulation size in 2000. Error bars show 95% confidence intervals, clustered at the industry level.