NBER WORKING PAPER SERIES

THE H-1B WAGE GAP, VISA FEES, AND EMPLOYER DEMAND

George J. Borjas

Working Paper 34793
http://www.nber.org/papers/w34793

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026

I am grateful to Nate Breznau, Hugh Cassidy, Daniel Di Martino, Gordon Hanson, Ron Hira, Kirk Doran, Daniel Hamermesh, Joan Monras, Gautam Nair, and Robert VerBruggen for valuable suggestions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The H-1B Wage Gap, Visa Fees, and Employer Demand
George J. Borjas
NBER Working Paper No. 34793
February 2026
JEL No. J08, J18, J69

ABSTRACT

The H-1B program lets firms hire high-skill foreign workers for a six-year term. The annual number of visas allocated to for-profit firms is capped at 85,000 and there is excess demand for those visas. The analysis merges administrative data, including the I-129 petitions that report the wage offer made to specific H-1B beneficiaries, with the American Community Surveys. On average, H-1B workers earn 16 percent less than comparable natives, suggesting that firms may be willing to pay a one-time fee to obtain the visa. The data are examined using a labor demand model to simulate how a fee alters the hiring decision. Depending on the level of excess demand, the unobserved productivity gains or costs from an H-1B hire, and the rate of job separations, the revenue-maximizing fee is between $118,000 and $264,000, has little or no impact on the number of H-1Bs hired, and generates between $6.2 and$22.4 billion in revenues. The demand for visas remains strong even if firms offshore some of the jobs currently held by H-1Bs. The fee also changes the skill composition of the H-1B workforce, making it more skilled.

George J. Borjas
Harvard University
Harvard Kennedy School
and NBER
george_borjas@hks.harvard.edu

The H-1B Wage Gap, Visa Fees, and Employer Demand

George J. Borjas*

I. Introduction

A key insight of the literature on the economics of immigration is that international labor flows produce larger economic gains for the receiving country when the flow is composed of high-skill workers (Blau and Mackie, 2017). The reason is obvious: High-skill immigrants produce more tax revenue and are less reliant on social assistance programs. In addition, high-skill immigrants are more likely to generate human capital externalities that increase the productivity of other workers.

Not surprisingly, many countries have enacted visa programs that target and try to attract high-skill workers, including the Express Entry System in Canada, the National Innovation Visa in Australia, and the Blue Card in the European Union (OECD, 2024; European Migration Network and OECD, 2025). The United States uses the H-1B program to grant temporary admission to immigrants in "specialty occupations."[1] These occupations require a college degree and the H-1B workers typically cluster in science, engineering, or computer-related jobs. The annual number of H-1B visas available to for-profit firms is legislatively capped at 85,000 visas for new workers.[2] The number of H-1B visas granted to non-profit institutions (such as universities, school districts, and medical centers) is not capped, but it is a far smaller program (about 25,000 new visas annually).[3]

The claim that H-1B workers generate beneficial productivity spillovers is used to argue for the continuation and expansion of the program. Testifying before Congress, for example, Bill Gates claimed that: "Microsoft has found that for every H-1B hire we make, we

* Harvard Kennedy School and National Bureau of Economic Research. I am grateful to Nate Breznau, Hugh Cassidy, Anthony Edo, Daniel Di Martino, Gordon Hanson, Ron Hira, Kirk Doran, Daniel Hamermesh, Joan Monras, Gautam Nair, and Robert VerBruggen for valuable suggestions.

[1] The United States has other visa programs aimed at high-skill workers. For example, the O-1 visa program selects persons with "extraordinary ability" in the sciences, business, or athletics. There are also permanent employment-based visas, many of which require advanced education.

[2] The 85,000 visas are allocated as follows: 65,000 to persons with at least a bachelor's degree and 20,000 to persons who have at least a master's degree from an educational institution in the United States.

[3] See https://www.govinfo.gov/content/pkg/FR-2025-12-29/pdf/2025-23853.pdf, p. 60884.

add on average four additional employees to support them in various capacities" (U.S. House of Representatives, 2008, p. 60884).

Beginning with Kerr and Lincoln (2010), many academic studies test the conjecture that H-1B workers produce such spillovers. The Kerr-Lincoln study noted that H-1B visa-holders cluster in a small number of locations (such as San Francisco or New York), and found that increases in the H-1B cap increased patenting in those cities. Other studies have used the Kerr-Lincoln cross-market methodology and find innovation effects of the program (see Peri, Shih, and Sparber, 2020; and the related work of Hunt and Gauthier-Loiselle, 2010). Note, however, that the settlement of H-1B workers in specific markets is not random, making it difficult to find a convincing instrument.

The literature has since shifted to a firm-based experimental approach. Because the number of H-1B visas is capped and there is substantial excess demand, the Department of Homeland Security (DHS) runs a lottery to allocate the 85,000 visas. On average, 450,000 workers participated in the annual lotteries between 2021 and 2026 (U.S. Citizenship and Immigration Services, 2026). The studies then compare the rate of innovation in firms that won the lottery with the rate in firms that lost. The pioneering analysis of Doran, Gelber, and Ilsen (2022) found no evidence of a spillover effect, although more recent studies reach conflicting results (Dimmock, Huang, and Weisbenner, 2021; and Mahajan et al, 2024).[4]

The H-1B program creates a temporary "marriage" between the firm and the worker. A *specific* firm petitions for the temporary employment of a *specific* worker. Although the mobility of this worker to other firms is not prohibited, the worker must find a new job *and* the new firm must again petition for the temporary employment of this person. The most recent data indicates that 63,865 H-1Bs (out of a population of 680,000) switched jobs in FY2024, for an annual separation rate of 9.4 percent.[5] The separation rate

[4] In addition to the debate over the productivity effects of the H-1B program, there is concern over negative wage or employment effects on native workers (Bound, Khanna, and Morales, 2019) and abuse of the program by some employers (Ontiveros, 2017; and Hira and Costa, 2021); see also "Pink Slips at Disney. But First, Training Foreign Replacements," *New York Times*, June 3, 2015.

[5] Although there can be substantial year-to-year variation in the number of approved H-1B petitions for a change of employer, the 2024 number is only slightly above the annual mean of 55,600 between 2011 and 2020 (Bier, 2025). The last DHS estimate of the H-1B population was 583,000 in 2019 (when H-1Bs made up 40 percent of 1.46 million temporary workers). The temporary worker population grew to 1.7 million by 2024. Applying the 40 percent H-1B share yields an estimate of 680,000 H-1Bs in 2024. See U.S. Citizenship and Immigration Services, *H-1B Authorized to Work Population Estimate*, 2019, Table 2; and Office of

for comparable high-educated workers is between 20 and 25 percent.[6] The additional obstacles that H-1B workers encounter in switching jobs may give the firm some market power, which likely reduces their wage. As Johnson, Lavetti, and Lipsitz (2025) show, noncompete agreements (NCAs), a related kind of contractual restrictions on labor mobility, reduce the wage by between 4 and 14 percent.

Very few studies document the wage gap between H-1B workers and comparable Americans because census-type surveys do not provide any information on the type of visa used by the foreign-born to work in the United States. A rare exception is the work of Bourveau et al (2025), which uses payroll data from a Big 4 accounting firm and finds that new H-1B workers earn 10 percent less than comparable natives.[7]

This paper measures the wage gap between H-1B workers and comparable American workers. It also examines the implications of the gap for policy changes (such as a visa fee) that can be adopted to redistribute the payroll savings now enjoyed by lottery-winning firms. The analysis merges (publicly available) data from three sources for the 2021-2024 period: the Labor Condition Application (LCA) filed by the firm to attest that their request to fill temporary employment positions under the H-1B program will not adversely affect conditions for other workers; the I-129 filings where a lottery-winning firm identifies the actual H-1B worker they propose to hire, *reports their salary,* and other characteristics such as gender, education, and age; and the American Community Surveys (ACS), used to construct a comparable sample of native workers. Because of data constraints, the analysis focuses on the capped H-1B visas given to for-profit firms.

Homeland Security Statistics, *Population Estimates for Nonimmigrants Residing in the United States: Fiscal Years 2019 to 2024*, 2025, Table 1. Hunt and Xie (2019) document that the separation rate of temporary workers spikes after they obtain permanent visa status.

[6] The monthly separation rate for workers with at least a college degree in the CPS is 2.1 percent (Kochhar, Parker, and Igielnik. 2022, p. 9). The Job Openings and Labor Turnover (JOLT) data reports a monthly rate of 2.7 percent for workers in the information industry that employs the bulk of H-1Bs (U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey,* November 2025, Table 10).

[7] The other studies that attempt to measure the wage gap include an unpublished working paper by Lofstrom and Hayes (2011) showing that, on average, H-1B workers get paid more; an Economic Policy Institute report by Costa and Hira (2020) showing that H-1B workers are paid 20 to 40 percent less; and an undergraduate thesis by Gustafsson and Lindblad (2025) also showing that H-1B workers get paid less.

The evidence unequivocally documents the existence of a substantial wage gap between H-1B workers and comparable Americans. The average H-1B worker earns about 16 percent less than an American worker with the same education, age, gender, occupation, and who works in the same locality. Since the average salary of these high-skill workers exceeds $100,000, the average payroll savings accruing to a firm that wins an H-1b visa in the lottery are large: nearing $100,000 over the six-year employment term.

These savings suggest that employers might be willing to pay a fee to "buy" a visa for a particular H-1B worker. Such a fee was, in fact, introduced in 2025. Beginning in FY2027, winners of future lotteries must pay a one-time $100,000 visa fee.[8] The paper combines the data on the wage distribution of the H-1B workforce with a model of the hiring decision to determine the impact of such fees on employer demand, and to calculate the size of the fee that would maximize government revenue under alternative assumptions about the unobserved (relative) productivity of an H-1B worker. The model exploits the theoretical insight that the unobserved efficiency gain or cost associated with an H-1B hire must be less than the payroll savings produced by that hire. After all, if the unobserved cost exceeded the payroll savings, the H-1B worker would not have been hired in the first place and would not appear in the sample.

The baseline simulation indicates that the large wage gap between H-1B workers and comparable Americans, combined with the excess demand for the visas, implies that imposing fees in the range of $150,000 to $200,000 may not change the demand for H-1B workers all that much. The fee, however, will generate substantial revenue, between $10 and $20 billion annually, depending on the mean of the unobserved productivity distribution of H-1B workers and on the level of excess demand. The fee will also change the skill composition of the H-1B workforce, making it more skilled.

The analysis also extends the baseline model to allow for job separations in the H-1B workforce (and the implied turnover costs incurred by firms that paid the fee) and for

[8] See https://www.whitehouse.gov/presidential-actions/2025/09/restriction-on-entry-of-certain-nonimmigrant-workers. The fee "applies only to new visas, not renewals, and not current visa holders"; see https://www.shipmangoodwin.com/insights/dhs-clarifies-scope-of-new-dollar100000-h-1b-fee-announces-plans-for-additional-changes-to-h-1b-program.html.

the possibility that some of the jobs currently held by H-1Bs are offshored. These extensions again reveal substantial demand for H-1B workers after the fee is imposed, suggesting that the H-1B wage gap is sufficiently large that it makes those workers competitive relative to both the native workforce and to cheap overseas labor.

II. Data

The employer's choice between hiring an H-1B worker or a comparable U.S.-born worker obviously depends on the wage gap between the two types of workers. Holding other things constant, employer demand for H-1Bs would vanish if comparable native workers were cheaper than the immigrants. As noted earlier, there is excess demand for the capped number of visas available to for-profit firms. This excess demand suggests that the foreign-born workers earn less than comparable natives and/or that they are more productive than natives with the same age, education, occupation, etc.

An important first step in any study of the economic impact of the H-1B program, therefore, is to estimate the wage gap between the two groups. The analysis reported in this paper merges micro data from three distinct sources to directly estimate the wage gap. A description of these data files follows.

A firm that wishes to hire H-1B workers must file a *Labor Condition Application (LCA) for Nonimmigrant Workers* (DOL Form ETA-9035 & 9035E).[9] This form is an attestation by the firm that "the employment of H-1B…nonimmigrant workers in the named occupation will not adversely affect the working conditions of similarly employed U.S. workers."[10] If a firm wishes to hire, say, 15 software programmers through the H-1B program, the firm files a *single* LCA application to cover all these positions. This application is given a unique DOL-ETA Case Number (henceforth "Case Number") that is used to track the firm's application through the various stages.

---

[9] See "Labor Condition Application for Nonimmigrant Workers, Form ETA-9035 & 9035E"; available at https://www.dol.gov/sites/dolgov/files/ETA/oflc/pdfs/eta_form_9035.pdf.

[10] See U.S. Department of Labor, "Labor Condition Application for H-1B, H-1B1 and E-3 Nonimmigrant Workers Form ETA-9035CP –General Instructions for the 9035 & 9035E." https://www.dol.gov/sites/dolgov/files/ETA/oflc/pdfs/ETA_Form_9035CP.pdf.

The information collected in the LCA form provides valuable data used in the empirical analysis. Specifically, it indicates if the positions are for full-time employment. It reports the occupation of the positions to be filled (at the very detailed level of an 8-digit Standard Occupational Classification code). Finally, the firm must classify the position into one of four prevailing wage levels, where the prevailing wage is "the arithmetic mean…of the wages of workers similarly employed in the area of intended employment."[11] The policy discussion often uses these levels to describe the skill composition of the H-1B workforce. As will be shown below, the prevailing wage levels are a *very* poor proxy for skills, and their use can produce a severely distorted picture of the actual skills of H-1B workers.

The second data file used in the analysis contains the information the employer reports in Form I-129, the *Petition for a Nonimmigrant Worker*, which is filed with DHS after the firm has won a slot in the H-1B lottery.[12] This form is "used by an employer to petition…for an alien to come temporarily to the United States as a nonimmigrant to perform services or labor."[13]

Unlike the LCA, the I-129 is an individual-level form: a lottery-winning firm petitions for the temporary employment of a *specific* foreign national. The I-129 names the proposed worker, and provides such identifying information as gender, year and country of birth, educational attainment, and the geographic location of the worksite. The instructions to the I-129 form specifically ask the employer to report the "salary or wages paid to the beneficiary," and those instructions are legally binding.[14] Finally, the I-129 reports the Case Number given to the associated LCA form.

---

[11] See 20 CFR § 656.40: "Determination of prevailing wage for labor certification purposes," available at https://www.ecfr.gov/current/title-20/section-656.40.

[12] The form is available at: https://www.uscis.gov/sites/default/files/document/forms/i-129.pdf. Note that the I-129 is *not* a visa application, as the visa application must be personally filed by the H-1B worker at a foreign consulate or embassy (if the worker is outside the United States). It is instead a petition that, if approved, allows the firm to hire a foreign-born person in a temporary nonimmigrant worker status.

[13] U.S. Citizenship and Immigration Services. "Instructions for Petition for Nonimmigrant Worker." https://www.uscis.gov/sites/default/files/document/forms/i-129instr.pdf.

[14] The United States Citizen and Immigration Service (USCIS) has a regulation stating that its form instructions have the force and effect of law. Section 103, 8 CFR 103.2(a)(1) states: "Every form, benefit request, or other document must be submitted to DHS and executed in accordance with the form instructions."

The study uses data from all I-129s filed between FY2021 and FY2024 by firms subject to the 85,000 annual cap. We merge these data with the information in the LCAs filed between FY2017 and FY2024 (as an LCA is valid for up to three years). Note that the sample consists of *new* H-1B workers hired during the period (as the data only includes lottery winners). We restrict the analysis to the subsample of I-129 petitions that were for full-time positions; were approved by DHS (over 97 percent are approved); and were for persons aged 21-50 at the time of the filing (representing over 99 percent of the sample).

The LCA data is publicly available.[15] The FY2021-FY2024 datafiles for the I-129 applications filed by for-profit firms are also publicly available, making the empirical analysis reported below reproducible. Bloomberg News, through a Freedom of Information Act (FOIA) request, obtained the datafile for all I-129 forms filed by for-profit firms during that period and posted it in a web archive.[16] To summarize, the H-1B worker's education, age, gender, earnings, worksite zip code, and the employer's Federal Employer Identification Number (FEIN) are drawn from the I-129 filings; the full-time employment status, SOC occupation code, and prevailing wage level of the position are drawn from the LCA filings.

Finally, the 2023 ACS provides the sample of "comparable" native workers. To approximate key characteristics of the H-1B workforce, the ACS sample consists of salaried workers aged 21-50, who work full-time year-round (i.e., 50 or more weeks a year, and 35 or more hours a week), have at least a college degree, and were born in the United States.

Table 1 reports summary statistics. Note that H-1B workers are more likely to be male, younger, and better educated. Over 40 percent of the H-1B workforce has at least a master's degree (as compared to a quarter of the native workforce). Both groups of workers have high salaries, slightly over $100,000 (in 2023 dollars).

The raw data suggest that sizable salary gaps will inevitably arise because of differences in location or occupation. Over 40 percent of the H-1B workers work in five high-wage metropolitan areas (New York, Dallas, San Jose, Seattle, and San Francisco), with

---

[15] The LCA data files are available at: https://www.dol.gov/agencies/eta/foreign-labor/performance.

[16] The Bloomberg files contain "Individual-beneficiary-level data from all H-1B registrations and I-129 petitions" and are available at https://github.com/BloombergGraphics/2024-h1b-immigration-data. The analogous data for cap-exempt workers are not publicly available.

over 25 percent living in either New York or the Bay Area. In contrast, only 13.3 percent of the native workers live in those five localities. Similarly, the H-1B workers cluster in a very small number of occupations. The five largest occupations are all in the high-tech sector (including Software Developers, Computer Systems Analysts, and Electrical and Electronics Engineers) and together account for 63.8 percent of total H-1B employment. In contrast, only 5.7 percent of the native workers are employed in those five jobs.

III. Estimates of the Wage Gap

The analysis pools the merged DHS-DOL data file with the 2023 ACS to estimate standard log wage regressions where the wage of worker i ($\log w_i$) is determined by a vector of skill characteristics (X_i) and a vector of fixed effects specifying the worker's location (G_i). The baseline regression model also includes a 0-1 variable indicating if the worker is an H-1B visa recipient (H_i):

$$\log w_i = X_i\beta + G_i\theta + \gamma H_i + \epsilon_n. \tag{1}$$

The variables included in the vector X are education, age, gender, and occupation fixed effects; the vector G consists of fixed effects describing the state-PUMA combination of a worker's worksite (for H-1B workers) or place of residence (for natives).[17] The coefficient γ measures the log wage gap between an H-1B worker and a statistically comparable U.S.-born worker. The observations in the regression are weighted using the ACS sampling weight for native workers and a weight set equal to one for the H-1B workers (since the H-1B sample represents the universe of the H-1B workforce).[18]

[17] The I-129 form filed by the firm reports the worksite zip code. We used the GEOCORR crosswalk to convert these zip codes into state-PUMA combinations, which are available in the ACS. The GEOCORR crosswalk is available at https://mcdc.missouri.edu/applications/geocorr.html. Some zip codes map into more than one state-PUMA combinations. We simplified by allocating the zip code to the PUMA with the largest population.

[18] The analysis excludes observations with outlying wage data in either the I-129 filings or the ACS. Specifically, both natives and H-1B workers who reported full-time annual earnings under about $30,000 are excluded; and H-1B workers with earnings in the top 0.05 percent of the distribution (i.e., earnings exceeding $1.2 million) are also excluded. The annual earnings data in the ACS is capped by the Census Bureau.

Panel A of Table 1 reports the adjusted log wage gap. The various columns of the table present alternative specifications of the regression model that include or exclude covariates.

Consider initially the log wage gaps reported in row 1 of Panel A, which present the generic OLS estimates. The raw gap in column 1 is positive and sizable, indicating that H-1B workers, on average, earn about 13 percent more than natives. It is interesting to note that there is a sizable unadjusted wage gap in log earnings favoring H-1B workers, but that the wage gap in actual earnings (as Table 1 shows) is numerically small and favors natives. This discrepancy may arise because the extreme clustering of H-1Bs in some regions and occupations produces differently shaped wage distributions for the two groups, so that the analysis might be sensitive to measuring earnings differences at the mean or at the median of the distributions. We will show shortly that the results would be very similar if instead of estimating OLS regressions, we estimated quantile regressions to measure the adjusted difference in median log earnings between H-1B workers and natives.

Column 2 shows that the large H-1B log wage advantage vanishes and becomes a disadvantage of almost 5 percent once the regression simply controls for the differences in the geographic location of the two groups. The wage disadvantage increases further after the regression controls for skill differences between the two groups. The preferred estimate of the wage gap in column 4, which adjusts for differences in education, age, gender, and occupation, reveals a 16.1 percent wage disadvantage for H-1B workers.[19] In short, a standard regression analysis that controls for obvious differences in key socioeconomic characteristics (particularly geography and occupation) leads to a 29-point reversal of the wage gap: from a +13 percentage point unadjusted advantage for H-1B workers to a -16 percentage point adjusted disadvantage.

---

[19] The 538 unique occupation codes in the Census classification scheme used in the 2023 ACS (IPUMS variable *occ*) provide far less detail about the worker's job than the 1,835 categories in the LCA data, which are based on 8-digit Standard Occupational Classification (SOC) codes. We created a crosswalk between the two by first collapsing the SOC data in the LCA to 6 digits, and then using crosswalks produced by the Census Bureau (for both the pre- and post-2018 SOC codes) that directly link 6-digit SOC codes to the IPUMS coding. The census crosswalks are available at www.census.gov/topics/employment/industry-occupation/guidance/code-lists.html.

As noted earlier, the average log wage gap might give a distorted impression of the actual wage difference because the wage distributions of the two groups might have very different shapes. The second row of Panel A reproduces the analysis using quantile regressions to estimate the difference in the median log wage between the two groups. The results are very similar: An unadjusted +15.4 percentage point advantage in the median log wage turns into a -15.6 percentage point disadvantage after including the covariates.

It is evident that the sample of U.S.-born salaried college graduates aged 21-50 used as the baseline group differs in important ways from the H-1B workforce, so that the estimated wage gaps might change substantially if the baseline sample more closely resembled the H-1B population. It is then important to determine what the H-1B wage gap would look like if the native workforce had a similar distribution of key socioeconomic characteristics as the H-1B workforce, particularly geography or occupation. This alternative estimation method essentially reweighs the native sample using the geographic or occupational distribution of the H-1B sample so that the geographic or occupational distribution of the two groups would be identical.

To illustrate, let β_j be the share of the H-1B workforce in occupation *j*, and let θ_{ij} be the ACS sampling weight of native person *i* who is in occupation *j*. The sampling weight for natives in the reweighted sample is given by:[20]

$$\omega_{ij} = \beta_j \theta_{ij} \frac{\sum_i \theta_{ij}}{\sum_i \beta_j \theta_{ij}}, \tag{2}$$

where the denominator in equation (2) ensures that the sum of the reweighted native population equals the original sum of the sampling weights in the ACS native sample.

Panel B of Table 2 shows the estimated wage gaps when the native population is reweighted so that the geographic distribution of natives is the same as the geographic distribution of H-1B workers, while Panel C conducts the parallel analysis that equalizes the two occupation distributions. The first column shows that using the baseline group of

[20] Ideally, the reweighting would use both geography and occupation. The sample, however, is not sufficiently large to allow such a detailed reweighting. Over half of the geography-occupation cells required to calculate β have only one observation, and almost 75 percent of the cells have three or fewer observations.

native workers reweighted to more closely resemble the H-1B workforce is sufficient to turn the 13.0 percent wage advantage enjoyed by H-1B workers in the raw data into a negative gap of -8.5 percent when the geographic distributions are equalized and -19.6 percent when the occupation distributions are equalized. After controlling for all the other covariates, column 4 shows that the estimated wage gaps exceed -20 percent, regardless of whether we use OLS or quantile regressions. In short, the construction of "better" native baseline groups reinforces the conclusion that H-1B workers earn far less than comparable native workers, with a wage gap that exceeds 16 percent.

Finally, Panel D of Table 2 reports the estimated wage gaps when the dependent variable in the baseline regression model is annual earnings (rather than its log). The dollar amount of the adjusted wage gaps is sizable. The fully specified model reveals a wage gap of almost $30,900 in mean earnings and $25,400 in median earnings.[21]

A. Wage gaps at the firm level

The I-129 petition that employers file for each individual H-1B worker gives detailed information about the employer, including the name of the firm and the Federal Employer Identification Number (FEIN). This information makes it possible to estimate equation (1) by replacing the H-1B indicator with a vector H_{if} indicating whether H-1B worker *i* is employed by firm *f*. The coefficients of these fixed effects then give the adjusted wage difference between H-1B workers in each firm and the baseline native sample.

Table 3 reports the number of lottery-winning petitions filed in the FY2021-FY2024 period by the 25 firms with the largest number of petitions. The table also reports the average annual earnings of H-1B workers and the adjusted wage gap in each of these firms (using the regression specification that includes all covariates).

There are large differences in the average salary of H-1B workers across firms, ranging from around $150,000 at Meta and Apple to about $80,000 in Infosys and Wipro.

[21] The regressions do not control for job tenure, a variable that will differ between H-1B workers and natives but is not available in the ACS. By construction, the H-1B workforce consists of newly hired workers. One way of approximating the omitted variable bias is to estimate the regression using a sample of younger male workers. Natives would then be likely to also have low job tenure, and the estimated wage gap would not be contaminated by the documented difference in tenure between men and women (Sicherman, 1996). The adjusted log wage gap in the fully specified model using the sample of men aged 21-30 is -0.135 (0.008).

The large American high-tech companies (e.g., Amazon, Google, Microsoft, Meta, Apple, and Tesla) tend to have relatively small and often insignificant wage gaps.[22] For example, the wage gap at Meta and Tesla is about -1 to -3 percent, and both are insignificantly different from zero. In contrast, several companies in the Top 25 (bolded in the table) are often mentioned in discussions of firms that are outsourcing pipelines for Indian H-1B workers, such as Infosys, Tata Consultancy, Wipro, and HCL (Hira, 2010; and Varma and Rogers, 2020). These outsourcing firms have negative wage gaps, and the gaps are sometimes very large. The wage gap at both HCL and Tech Mahindra exceeds -30 percent.

Notably, there is only a modest clustering of H-1B workers into the Top 25 firms: Nearly 75 percent of H-1B hires are outside the Top 25. In fact, the I-129 filings reveal that 46,184 distinct firms (as identified by a unique FEIN) hired at least one H-1B worker. Strikingly, the median firm in the sample hired exactly *one* H-1B worker during the entire 4-year period, and the 75th percentile firm hired *three*.

The large (and significant) wage gap of -18.5 percent reported for "all other firms" in Table 3 then suggests that the low average pay of H-1Bs is not simply due to the large outsourcing firms. Estimating equation (1) in the subsample of firms that hired only one H-1B worker yields a log wage gap of -0.222 (0.004). A similar regression in the subsample of firms that hired at most three yields a gap of -0.212 (0.004).

In short, practically all H-1B workers (except the few that end up in the large American-owned high-tech companies) end up working for firms that pay them far less than the market wage for statistically comparable American workers.[23]

---

[22] Amazon hires H-1B workers in several divisions (Amazon.com Services, Amazon Web Services, and Amazon Development Center). The different divisions, however, have different FEIN numbers.

[23] As noted earlier, Bourveau et al (2025) use payroll data from a Big 4 accounting firm (Deloitte & Touche) and estimate a 10 percent wage disadvantage for H-1B workers relative to other workers in that firm. This gap is not comparable to the 13.2 percent estimate reported for Deloitte Consulting in Table 3. First, although both firms are part of the Deloitte network, it is Deloitte & Touche that is the accounting firm. Second, the wage gap in the Bourveau et al study is a *within-firm* wage gap. Finally, the Deloitte payroll data does not identify which workers are, in fact, H-1B visa recipients (nor does it contain any demographic characteristics). Instead, the Bourveau et al study used characteristics drawn from the LCA files (such as job title, worksite location, and job start date) to construct a "pseudo" sample of Deloitte's H-1B workforce. The data used in this paper indicates that the estimate of the log wage gap (without any covariates) between actual H-1B workers at Deloitte & Touche and the native ACS baseline is -0.116 (with a standard error of 0.008), which closely resembles the Bourveau et al finding.

B. Wage gaps within occupations

The H-1B workforce clusters in a small number of occupations. A single job (Software Developers) employs 38.3 percent of all workers; the five largest occupations employ almost two-thirds; and the ten largest employ three-quarter. Because the relative number of natives in these (typically high-paying) occupations is far smaller, it is not surprising that an earnings regression that adds occupation fixed effects substantially increases the wage disadvantage of H-1B workers relative to comparable Americans.

The comparison of H-1B workers to native workers *within* the same occupation often reveals some very large wage differences. For instance, Table 4 documents an unadjusted wage gap of nearly $40,000 for Software Developers, the largest occupation employing H-1Bs. The last column of the table reports the regression coefficient measuring the adjusted wage gap from earnings functions estimated within each of the ten occupations (as well as the "all other occupations" group). There is substantial variation in the adjusted log wage gap across the occupations. The wage gap is -30 percent for Software Developers and Computer Programmers, -14 percent for Electrical and Electronics Engineers, and +7 percent for Operations Research Analysts.

It would be of interest to determine how these wage gaps depend on supply-demand conditions in specific types of jobs. For example, the positive H-1B wage premium for Operations Research Analysts may reflect the scarcity of that type of skill in the native workforce. Similarly, the 30 percent negative wage gap estimated for Software Developers could be misleading as the wage for native workers in that occupation may have responded to the large supply shock affecting those specific workers. These extensions of the analysis are beyond the scope of this paper, but they should be kept in mind as they affect the interpretation of the results.

IV. The Skill Composition of the H-1B Workforce

The LCA requires that employers classify the position that the H-1B worker will fill into one of four prevailing wage categories. These categories are defined in terms of how the wage associated with that position compares to the wage of other workers in the same occupation *and* region. The prevailing wage levels reported in the LCA are:

Level I: Wage at approximately the 17th percentile of the wage distribution:

Level II: Approximately the 34th percentile.

Level III: Approximately the 50th percentile.

Level IV: Approximately the 67th percentile.

Although these firm-reported prevailing wage levels are often used to describe the skill level of the H-1B workforce, it is easy to see why these levels may not be strongly correlated with tangible skills. Consider the situation of a young high-skill worker living in a high-wage area—for example, a software programmer in San Francisco. This worker will earn a low wage relative to his comparison group (other software programmers in San Francisco), which would lead the firm to place him in a low prevailing wage level. Nevertheless, this worker's actual skills may be relatively high in comparison with many other workers in the national labor market.

The analysis uses an alternative (and more standard) index to describe the skill composition of the H-1B workforce. This index is based on the wage differences that are typically interpreted as productivity differences. Specifically, we use a (log) wage regression to calculate the market value of the set of *observable* skills embodied in each H-1B worker, after netting out extraneous factors that affect wages. One potential extraneous factor in the current context is that H-1Bs disproportionately reside in high-wage areas (such as New York City or the Bay Area). Although geographic wage differences can arise due to both cost-of-living differentials and agglomeration effects that affect the productivity of all workers, a correctly defined skill index that summarizes a worker's *observable* skills would abstract from these geographic wage differences.

Consider the following regression model estimated in the sample of *native* workers:

$$\log w_n = X\beta_n + G\theta_n + \epsilon_n. \tag{3}$$

It is straightforward to use the regression equation in (3) to construct an index based on an H-1B worker's observable skills. Specifically, we use the coefficient vector β_n to predict each H-1B worker's log wage (after "deflating" for geographic wage differences):

$$I_h = X_h\hat{\beta}_n, \tag{4}$$

where the vector X includes educational attainment, age, gender, and occupation fixed effects. The index I_h essentially gives a weighted sum of the H-1B worker's observable skills, where the weights measure how the different skill characteristics are valued in the labor market. For convenience, we classify H-1B workers into four skill categories based on the quartiles of the skill index I_h: workers with an index below the 25th percentile of the native distribution, between the 25th and 50th percentile, between the 50th and 75th percentile, and above the 75th percentile.

Before proceeding to describe the skill composition of the H-1B workforce implied by the index I_h, it is instructive to illustrate with specific examples that the prevailing wage levels reported by employers in the LCA have little to do with the tangible skills embodied in the foreign-born workforce. The first three columns of Table 5 compare the prevailing wage and skill index distributions for three of the top employers of H-1B workers (Amazon.com Services, Microsoft, and Tech Mahindra).

For each of these large users of H-1B visas, the skill index distribution indicates that the H-1B workers are quite skilled relative to natives. Over 40 percent of the workers in each of these firms are in the top quartile of the native distribution, and an additional 40 percent are in the third quartile. These rankings differ dramatically from the prevailing wage levels those firms reported in the LCA: Both Amazon and Microsoft asserted that between 50 and 60 percent of these H-1B workers would be in the bottom prevailing wage level (i.e., around the 17th percentile of the wage distribution). Similarly, the Tech Mahindra LCA filings assert that 100 percent of the H-1B workers would be around the 34th percentile of the distribution.[24]

Column 4 of Table 5 reports the distribution of the entire H-1B workforce across the prevailing wage levels or the quartiles of the skill index. The prevailing wage level

---

[24] Although this paper does not estimate the wage gap for cap-exempt H-1Bs (as the I-129 filings by non-profit firms are not publicly available), the discrepancy between the prevailing wage level classification and the actual skills embodied in cap-exempt H-1B workers is equally stark. Consider the distribution of prevailing wage levels in the LCAs filed by two cap-exempt employers, the Cleveland Clinic and the Savannah Chatham School District. Almost all the positions in the LCAs filed by the school district are for elementary, middle, or high school teaching positions, and the district classified over a quarter of these positions in prevailing wage level IV. At the Cleveland Clinic, where 80 percent of the positions are for medical scientist, physician, or surgeon positions, the Clinic classified half of their workers in prevailing wage level I (and only 10 percent in prevailing wage level IV).

classification obviously gives a misleading picture of the actual skills of H-1B workers. According to the firms' LCA filings, over 80 percent of the workers are in prevailing wage levels I and II, indicating that they are *below* the 50th percentile of the prevailing wage distribution. However, the distribution of the skill index shows that nearly three-quarters of the H-1B workers have observable skills that would put them *above* the median of the native skill distribution.[25]

The analysis in the last section documented that, on average, H-1B workers earn substantially less than statistically comparable native workers. It is of interest to determine how the wage gap varies across skill groups. The last column of Table 5 uses the regression estimated in equation (3) to calculate the mean log wage gap for each of the skill groups:

$$\Delta \log \bar{w}_s = \log \bar{w}_{hs} - \bar{X}_s \, \hat{\beta}_n + \bar{G}_s \hat{\theta}_n, \tag{5}$$

where $\log \bar{w}_{hs}$ is the mean log wage of H-1B workers in skill group *s*; and $\bar{X}_s$ and $\bar{G}_s$ give the mean value of the covariates for that group. The table reveals that the H-1B workers at the *top* of the skill distribution have the *lowest* earnings relative to their native counterparts. There is no wage gap disadvantage in the bottom quartile of the distribution, but this gap increases to -11.8 percent for workers in the third quartile, and to -28.6 percent for workers in the top quartile.[26] As will be shown below, this pattern has important implications: a policy initiative that adds a fixed fee to the cost of hiring an H-1B worker will inevitably change the skill composition of the H-1B workforce, likely increasing its average skills.

V. Visa Fees and Labor Demand

A specific employer must file Form I-129 to hire a specific foreign worker through the H-1B program. The issuance of a visa then tends to create an employment relationship

---

[25] The distinction between the prevailing wage levels and an index that measures the actual skills embodied in the worker has important policy implications. Beginning in FY2027, the H-1B lottery will give greater weight to petitions for workers in the higher prevailing wage levels.

[26] The Bourveau et al (2025) study also suggests that the wage gap at Deloitte & Touche is far greater for workers at the Manager level than at the staff level.

that can last up to six years (with renewal), perhaps giving the employer some market power and likely reducing the relative wage of the H-1B workforce. It is also sometimes argued that H-1B workers introduce productivity gains in the firms that hire them (Chen, Hshieh, and Zhang. 2021; Dimmock, Huang, and Weisbenner, 2022; Ghosh Mayda, and Ortega, 2014; Mahajan et al, 2024), further increasing the firm's profitability.

The lower payroll costs and the potential productivity gains suggest that many employers might be willing to pay a fee to "buy" an H-1B visa for a particular worker. In fact, beginning with the FY2027 lottery, firms will be required to pay a fee of $100,000 at the time the visa is issued. This section of the paper filters the data on the wage distribution of the H-1B workforce through the lens provided by a simple economic model of the hiring decision to estimate the impact of a fee on employer demand, and to calculate the size of the fee that would maximize government revenue.[27]

Let w_h be the annual wage the employer offers H-1B worker h (assumed to be constant over the term of the six-year visa). Let F be a *one-time* fee charged to the employer when first awarded the H-1B visa for that worker. We initially assume that the firm hires the worker for the entirety of the visa term. Let π_h be a random variable, observed to the employer but *unobserved* to the analyst, that captures the productivity gain or additional cost associated with the H-1B hire relative to a native hire. Phrased differently, the variable π_h gives the unobserved "cost differential" of an H-1B hire. The differential may be negative if the H-1B worker brings new or complementary skills into the production process or if they impose lower turnover costs on the firm. It may also be positive because an H-1B hire might introduce additional costs, including legal fees incurred during the lengthy visa process, production inefficiencies created by language barriers, or acculturation issues within the firm's workforce. Finally, let w_n be the annual wage the employer must pay a comparable U.S.-born worker (also assumed to be constant over the six-year period).

[27] One alternative to fees would be to auction off the available H-1B visas. In the U.S. context, however, this (perhaps more efficient) alternative is difficult to implement as it would require new legislation to set up the auction system. The fee has the advantage that it can perhaps be imposed within the context of existing legislation.

For analytical convenience, we scale the cost differential π_h so that it can be interpreted as a fraction of the H-1B worker's wage. The true cost of hiring an H-1B worker is $w_h(1+\pi_h)$. The employer hires an H-1B worker over a comparable American if:

$$w_h(1+\pi_h)R + F \leq w_n R, \tag{6}$$

where R is the discount factor used to add up the payroll outlays over the six-year visa term ($R = \sum_{t=0}^{5} 1/(1+r)^t$). The inequality in equation (6) simply states that the employer hires the H-1B worker if the "true" cost of hiring that worker (i.e., including the unobserved cost differential) is lower than the cost of hiring a comparable American.

Because empirical studies of earnings determination typically use log wage regressions, it is useful to convert the firm's hiring decision rule into log terms:

$$\log w_h + \log(1+\pi_h) + \log\left(1 + \frac{F}{w_h(1+\pi_h)R}\right) \leq \log w_n\,. \tag{4}$$

Consider the hiring decision in the absence of any fee (as was the case in the policy regime during the 2021-2024 period covered by the data). If the employer hired H-1B worker h during that period, it must have been the case that:

$$\pi_h \leq \log w_n - \log w_h = \Delta w_{nh}, \tag{5}$$

using the approximation that $\log(1+\pi_h) \approx \pi_h$. The results reported in Section II indicated that, on average, Δw_{nh} is positive, as hiring a native worker is more costly than hiring an equivalent H-1B worker. To simplify the exposition, the variable Δw_{nh} will be called the "payroll savings" accruing to the firm from the hire of H-1B worker h.

Equation (5) states that for an H-1B worker to have been hired in the absence of a fee it must have been the case that the unobserved cost differential π_h was smaller than the payroll savings. This inequality obviously holds if π_h is negative, or when H-1B workers introduce unobserved productivity gains into the firm. The inequality might also hold if π_h is positive, or when there are unobserved costs associated with the hire. Equation (5) then states that the H-1B worker would still have been hired if the unobserved cost increase is smaller than the payroll savings (in percentage terms).

The empirical exercise uses data from actual H-1B hires in the FY2021-FY2024 period, when the inequality in equation (5) *must* hold, to infer if the employer would still hire the specific foreign worker after the introduction of a one-time fee of F dollars. The simulation assumes that the unobserved cost differential π_h has a normal distribution with mean μ_π. Consider the simplest case where $\mu_\pi = 0$ (although the analysis will show the sensitivity of the results to alternative assumptions). The variance of the normal density is then uniquely determined by the additional assumption that 99 percent of the values of the random variable π_h lie within the (-0.5, +0.5) interval, so that hiring an H-1B worker will typically not reduce the true hiring cost by more than 50 percent of the worker's wage, or increase it by more than 50 percent. The standard deviation of the distribution of π_h then equals 0.194.[28]

Note that these distributional assumptions refer to the *population* of potential H-1B job seekers. Equation (5) implies that the distribution of the unobserved random variable π_h in the subsample of H-1B workers *actually hired* is a normal distribution truncated from above. The truncation point for H-1B worker h is the value of the payroll savings Δw_{nh} associated with that hire.

The simulation requires information on the wage that a U.S.-born worker comparable to H-1B worker h would earn. We predict this wage by using the regression model in equation (3), a log earnings function estimated in the sample of native workers where the covariates include skills (i.e., education, age, gender, and occupation) and geography fixed effects. The wage that H-1B worker h would earn if he were native is:

$$\log \widehat{w}_{nh} = X_h \hat{\beta}_n + G_h \hat{\theta}_n. \tag{6}$$

The employer will hire H-1B worker h even after the introduction of a one-time fee of F dollars if:

[28] The assumption that 99 percent of the area lies within the (-0.5, +0.5) interval implies that the standard deviation is given by: $\sigma_\pi = (\pi_h - \mu_\pi)/2.576 = 0.5/2.576 = .0.194$ if $\mu_\pi = 0$.

$$\pi_h + log\left(1 + \frac{F}{w_h\, e^{\pi_h} R}\right) \leq (\log \widehat{w}_{nh} - \log w_h) = \Delta\widehat{w}_{nh}. \qquad (7)$$

Given the distributional assumptions about the unobserved cost differential π_h and a specific value of the fee F, it is straightforward to determine if the inequality in (7) holds for each of the H-1B workers hired between 2021 and 2024 (when the fee was zero). The hiring rates reported below are averages over 100 replications of the simulation exercise, using a 3 percent rate of discount.

VI. Simulation Results

Figure 1 illustrates the distribution of the estimated payroll savings $\Delta\widehat{w}_{nh}$ for the existing H-1B workforce. The distribution has a mean of 0.148.[29] The exercise produces positive payroll savings for 71.8 percent of the workers. For about 28 percent, however, it seems that it would have been cheaper to hire the comparable native worker. The fact that this subsample of H-1B workers with negative payroll savings was *indeed* hired strikingly illustrates the existence (and importance) of an even more negative unobserved cost differential π_h. After all, in the absence of sizable unobserved productivity gains, these workers would never have been hired in the first place.

To show how visa fees influence the firm's hiring behavior, it is instructive to begin with the simplest (but very unrealistic) scenario where the *distribution* of the random variable collapses to zero (i.e., $\pi_h = 0$)—that is, hiring an H-1B worker *never* introduces a productivity gain and *never* introduces additional costs. The observed data (i.e., the H-1B worker's wage, the wage that a comparable native earns, and the size of the fee) provides all the information required by the employer to make the hiring decision.

As the top panel of Figure 2 shows, only 71.8 percent of existing H-1B workers (i.e., the subsample of workers with positive payroll savings) satisfy the inequality in equation (7). In the absence of an unobserved cost differential, *none* of the 28.2 percent of the

[29] The mean of $\Delta\widehat{w}_{nh}$ is the log wage gap between natives and H-1B workers that would be calculated in a standard Oaxaca-Blinder decomposition, where the coefficients of the earnings function vary between natives and H-1B workers. Note that the gap predicted by this decomposition is very similar to the 16 percent gap estimated in the regression analysis reported in Table 2.

workers who produce a negative payroll savings should have been hired in the first place (when the fee was zero). Despite these initial conditions, it is notable that the simulation shows that even with a fee as high as $89,000, the hiring rate would only fall to 50 percent.

Figure 2 also illustrates the hiring rate in the more realistic scenarios that allow for the presence of an unobserved cost differential.[30] As discussed above, the value of the random variable π_h is drawn from a normal distribution truncated (from above) at the value of the payroll savings, $\Delta\widehat{w}_{nh}$. The random draw from this truncated normal ensures that the unobserved cost differential is always smaller than the payroll savings associated with that worker, providing the economic rationale for the hiring of the 28.2 percent of the H-1B workers with negative payroll savings.

The simulations use three alternative assumptions for the mean of the random variable π_h: -0.1, 0.0, or +0.1.[31] These assumptions allow for the possibilities that the average H-1B worker embodies an unobserved productivity effect that reduces the worker's hiring cost by 10 percent; or that the average H-1B worker neither increases nor decreases the cost of the hire; or that the average H-1B worker introduces an additional cost that raises the cost of hiring by 10 percent.

The simulations obviously predict that all workers are hired when the fee is set to zero and the hiring rate declines as the fee increases. Surprisingly, the hiring rates are not very sensitive to alternative assumptions about the mean of the unobserved cost distribution for positive values of the fee. Consider, for example, the imposition of a $100,000 fee. It turns out that between 47 and 64 percent of the H-1Bs hired between 2021 and 2024 would still be hired. The hiring rate is, of course, lowest at 46.6 percent when the mean of the unobserved cost effect distribution is +0.1 so that there is a substantial additional cost in an H-1B hire, and it is highest at 64.3 percent when the mean is -0.1, so that the average H-1B worker embodies a hidden productivity gain. One implication of the relative constancy of the hiring rate is that the payroll savings associated with hiring H-1B workers is often so

---

[30] The simulations trim the outliers of the distribution of $\Delta\widehat{w}_{nh}$ at the values of the 1st and 99th percentiles. In other words, all values of $\Delta\widehat{w}_{nh}$ below the first percentile are replaced with the 1st percentile value, and all values above the 99th percentile are replaced with the 99th percentile value. This trimming has little effect in the predicted hiring decision because very positive (or very negative) payroll savings dominate the role played by the other variables in the model.

[31] The various scenarios only shift the mean of the distribution of π_h; the variance is held constant.

large that shifting the mean of the distribution of the unobserved cost differential by 20 percentage points (i.e., from -0.1 to +0.1) does little to change the hiring behavior of firms.

The estimated hiring rates illustrated in the top panel of Figure 2 can be used to calculate the total number of H-1B workers who would be hired when firms must pay the one-time fee of F dollars. The government caps the number of H-1B visas issued to for-profit firms at 85,000 visas per year. There has been substantial excess demand for these visas in the past so that USCIS holds an annual lottery to allocate the scarce visas. On average, almost 450,000 distinct registered beneficiaries participated in the lotteries between 2021 and 2026, although there is a huge amount of variation across years, from a minimum of 274,237 in FY2021 to a maximum of 780,884 in FY2024.

The excess demand inevitably plays a crucial role in determining market outcomes once a fee is introduced. Suppose the hiring rate with a fee of, say, \$100,000 is 50 percent. In other words, only half of the H-1B workers hired in the past would have been hired if firms had to pay \$100,000 for the visa. If the lottery winners are a random sample of the population of lottery registrants, firms would still want to hire half of the workers who lost the lottery. After the imposition of the fee, the only firms who will register for the lottery are those who are willing to pay the fee if they win. If the population of lottery registrants was 450,000, there would be a total of 225,000 lottery registrations for the 85,000 slots and total demand for H-1B workers remains unchanged.

More generally, let p_F be the hiring rate when the fee is F dollars, and let N be the baseline number of lottery registrations. The total number of visas V_F that employers will demand after the imposition of a fee is:

$$V_F = \begin{cases} p_F N, & if\ p_F N < 85{,}000, \\ 85{,}000, & if\ p_F N \geq 85{,}000. \end{cases} \tag{8}$$

The simulation uses three alternative values of the baseline number of registrations: twice the number of capped visas or 170,000; three times or 255,000; and four times, or 340,000. The three scenarios reflect “low” excess demand, “moderate” excess demand, and “high” excess demand assumptions. It is instructive to discuss the evidence using the low excess demand assumption, as this scenario provides lower bound estimates of the impact of a fee on employer demand and government revenue.

Panel B of Figure 2 uses equation (8) to calculate the number of visas that will be used for alternative values of the fee, assuming the population of lottery registrants is 170,000. Even if the excess demand for H-1B visas were only 170,000 (below the actual number of registrants in the recent past), firms would still use all or nearly all the 85,000 available visas for fees as large as $100,000. If the average H-1B worker indeed embodied a 10 percent productivity advantage, the number of visas demanded would remain at 85,000 even if the fee were as high as $149,000.

The estimated number of visas demanded provides the data required to calculate the revenue-maximizing fee. The bottom panel of Figure 2 illustrates the revenue curves (i.e., the product of the fee times the number of visas demanded). The revenue-maximizing fee is large: $135,000 when the average H-1B worker increases hiring costs by 10 percent; $156,000 when the average worker neither increases nor reduces costs; and $176,000 when the average worker embodies productivity gains that reduce costs by 10 percent. The total revenue collected at the revenue-maximizing fee in the various scenarios is substantial, ranging from $8.3 billion to $13.0 billion annually.

Table 6 summarizes the results of the various simulations, allowing for both the alternative assumptions about excess demand and the unobserved cost differential. All 85,000 capped visas will be used in either the moderate or high excess demand scenarios, regardless of the value of the cost differential. Even with hiring rates as low as 24 percent, there would still be a sufficiently large number of current lottery losers that would be hired after the imposition of a fee.

Note also that the value of the fee is relatively high in all simulations reported in the table, ranging from $135,000 to over $250,000. Similarly, the revenue collected will also be substantial, ranging from $8 billion to over $20 billion.

Finally, it is important to document that a fee will change the skill composition of the H-1B workforce. As was shown in Table 5, the wage disadvantage of H-1B workers is largest for high-skill workers. Put differently, payroll savings increase when the firm hires more skilled workers. Fees should then shift the demand for H-1B workers towards the higher skill groups.

Figure 3 shows how the hiring rate differs across the quartiles of the skill distribution after the introduction of the revenue maximizing fee for each scenario. Regardless of the assumption about the mean of the unobserved cost differential π_h, the hiring rate is much higher for more skilled H-1B workers. Suppose, for example, that the mean of the unobserved cost differential is zero (i.e., $\mu_\pi = 0$). The hiring rate ranges from 9.3 percent for workers in the bottom quartile of the skill distribution to 63.0 percent for workers in the top quartile. A policy initiative that charges the revenue-maximizing fee at the time an H-1B visa is issued, therefore, not only generates sizable revenues but also markedly increases the skill level of the H-1B workforce.

VII. Extensions

A. Job Separations

The model of the employer's hiring decision explicitly assumed that the hire lasts for the full six-year visa term once a lottery-winning firm pays the required fee. Although H-1B workers are less likely to switch jobs, there is still some mobility. It would be of interest to predict the outcomes after allowing for job turnover (which obviously affects the present value of the cost streams). Suppose that, as noted earlier, the annual separation rate is 9.4 percent. This implies that 9.4 percent of H-1B workers stay on the job 1 year, 8.5 percent stay 2 years, 7.7 percent stay 3 years, 7.0 percent stay 4 years, 6.3 percent stay 5 years, and 61.0 percent stay all six years. The simulation then uses a uniform random variable in the interval [0, 1] to allocate workers to the six possible values of job tenure.

Table 7 summarizes the results when both the cost differential (π_h) and job tenure are randomized. Consider the scenario where there is moderate excess demand and the mean of the unobserved cost differential is zero. The revenue-maximizing fee is $138,000. At this fee, the hiring rate is 33.3 percent, all 85,000 visas are used, and revenue is $11.7 billion. The fee and revenue estimates are about 25 percent lower than those reported in the absence of job turnover in Table 6. Looking across all scenarios, the revenue-maximizing fee is still large, ranging from $100,000 to $200,000; all or nearly all the 85,000 available visas are used; and revenue ranges from about $6 billion to $18 billion.

Although the simulation results reported in Table 7 represent a useful first step for understanding how job turnover in the H-1B workforce affects the size of the fee and employer demand, a large fee might change the nature of the job match in unpredictable ways. The firm will be more hesitant to hire new H-1Bs, and the labor market for "used" H-1Bs (i.e., those who are in the middle of the six-year visa term) would become more competitive. From a worker's point of view, however, the ultimate "prize" of winning the lottery may not be the increased earnings earned during those six years, but the increased probability of obtaining a green card at the end.

In a policy regime where employers pay substantial upfront fees, job separations might interact with the probability of reaching that goal because "over 80 percent of all H–1B visa holders who adjust to lawful permanent resident status do so through an employment-based green card."[32] These employment-based visas, such as EB-2 and EB-3, *require* employer sponsorship. The tradeoff between a higher short-term salary due to a better job opportunity and the possibility that imposing sizable turnover costs on a current employer might have spillover effects on future sponsorship offers could easily become an important determinant of job separations in the H-1B workforce.

B. Offshoring

Although the baseline model emphasized that the hiring decision depends on whether an H-1B worker is cheaper than a comparable native, it ignored the possibility that it might be even cheaper to offshore the job. The burden of sizable visa fees could induce firms to simply offshore many of the jobs currently held by H-1Bs. Suppose the firm can pay wage w_A to a comparable worker abroad. The following two inequalities must hold if the H-1B worker was hired in the absence of a fee:

$$\log w_h + \pi_h \leq \log w_n \,, \qquad (8a)$$

$$\log w_h + \pi_h \leq \log w_A \,, \qquad (8b)$$

[32] U.S. Department of Labor & Department of Homeland Security, "Strengthening Wage Protections for the Temporary and Permanent Employment of Certain Aliens in the United States; Interim Final Rule," *Federal Register* 86 (Jan. 14, 2021), p. 3622. Available at: https://www.govinfo.gov/content/pkg/FR-2021-01-14/pdf/2021-00218.pdf.

Equation (8a) indicates that the H-1B hire occurred only if the "true" hiring cost was less than the wage of a comparable native. Equation (8b) extends the argument and indicates that the H-1B hire occurred only if the true hiring cost was less than what it would have cost to employ a comparable worker abroad.

The evidence suggests that offshoring can generate substantial payroll savings. Daga and Kaga (2006) conclude that "many firms realize labor cost savings in the 20–30 percent range or more when offshoring operational functions." Suppose that the relationship between the offshoring wage and the native wage is:

$$\log w_A = \log w_n\,(1+\theta) \approx \log w_n + \theta, \tag{9}$$

where the random variable θ gives the percent cost savings from offshoring the job. We assume that $\theta \sim N(\mu_\theta, \sigma_\theta^2)$, and that θ is independent of π_h. The independence follows from the fact that the value of the job match between worker h and the firm (which π_h attempts to capture) has little to do with the savings from offshoring a specific job. We set $\mu_\theta = -0.25$, and the variance is determined by the additional assumption that 99 percent of the density lies in the [-0.5, 0] range, so that jobs moved offshore almost always result in payroll savings relative to the native wage.[33] This implies that $\sigma_\theta = 0.097$.

In the absence of a fee, worker h is hired if:

$$\nu = \pi_h - \theta \leq \log w_n - \log w_h = \Delta w_{nh}\,, \tag{10}$$

The comparison of equations (5) and (10) shows that a simulation exercise that accounts for the possibility that H-1B jobs can be offshored is identical to the one conducted earlier, but the random variable has a different mean and variance (i.e., $\nu \sim N(\mu_\pi - \mu_\theta, \sigma_\pi^2 + \sigma_\theta^2)$).

---

[33] As equation (10) implies, the assumption that θ is always negative is conceptually equivalent to an increase in the cost of an H-1B hire (i.e., $\pi_h > 0$). The simulation results then tend to exaggerate the negative impact of a fee on the demand for H-1B workers, giving upper-bound estimates of the drop in demand.

Table 8 summarizes the results. Overall, the impact of offshoring on H-1B demand is modest. Consider the scenario where the mean of the unobserved cost differential is zero and there is moderate excess demand. The baseline simulation reported in Table 6 yields a revenue-maximizing fee of $181,000, which drops to $113,000 if firms have the option to offshore the H-1B jobs once the fee is imposed. Note, however, that all visas would still be used and revenue would equal $9.6 billion dollars.

Regardless of the scenario, Table 8 shows that the revenue-maximizing fee is large (typically exceeding $100,000); that most, if not all, visas are used; and that revenue is substantial, ranging from $5 billion to $15 billion. The relatively inelastic demand for the visas suggests that the H-1B wage disadvantage is sufficiently large that the foreign-born workforce remains competitive with both native workers and with workers abroad even after the introduction of a sizable fee.

VIII. Summary

The economic benefits from immigration are larger when the immigrant flow consists of high-skill workers. As a result, many immigrant-receiving countries set up visa programs to recruit such workers. The United States uses the H-1B program to grant temporary work permits to high-skill immigrants in "specialty occupations." The number of H-1B visas available to for-profit firms is legislatively capped at 85,000 new visas per year, and the H-1B workers typically cluster in science, engineering, or computer-related jobs.

The mechanics of the H-1B program help to create a long-term link between the firm and the worker. A *specific* firm requests permission for the temporary employment of a *specific* worker. In theory, the worker can move to other firms. The new firm, however, needs to go through the process of submitting a petition for the temporary employment of that person. This arrangement may give employers some market power, which likely reduces the wage of the H-1B workforce.

The analysis presented in this paper estimated the wage gap between H-1B workers and comparable U.S.-born workers and examined the implications of the gap for employer demand. The analysis merged data from three distinct sources: the Labor Condition Application (LCA) filed with DOL for temporary positions to be filled by foreign persons, the I-129 forms filed with DHS petitioning the entry of a specific foreign-born person through the H-1B program, and the American Community Surveys (ACS) that provides the "baseline" sample of native workers.

The evidence indicated that H-1B workers are cheaper to hire than comparable Americans. The average H-1B worker earns about 16 percent less than a U.S.-born worker in the same locality and with the same education, age, gender, and occupation. Since these high-skill workers typically earn more than $100,000 annually, the average payroll savings resulting from a single H-1B hire nears $100,000 over the term of the six-year visa term.

The sizable wage disadvantage of H-1B workers suggests that firms might be willing to pay a substantial fee for the "privilege" of hiring such a worker. The simulation of an economic model of the employer's hiring decision, combined with the excess demand for the foreign workers, revealed that imposing a visa fee between $150,000 and $200,000 may not change the number of H-1B workers hired all that much. The fee, however, will generate revenues totaling between $10 billion and $20 billion annually and change the skill composition of the H-1B workforce, making it more skilled.

Immigration policy in the United States partly consists of an alphanumeric soup of visa programs that permit temporary employment of certain classes of persons and that cater to specific sectors or industries (e.g., H-1B, H-2A, H-2B, H-3, O-1A, O-1B, OPT, etc.). The analysis of the H-1B program showed that the specific industry to which this program is targeted (i.e., the high-tech sector) has gained substantially from its access to a foreign-born high-skill workforce through lower labor costs. It would not be surprising if other sectors using other programs (e.g., the agricultural industry that relies on H-2A visas) also benefit substantially from those programs. It would be of interest to document the extent to which each of these programs has increased the wealth of very narrow interests. Such documentation might inform a revamping of guest worker programs so that the economic benefits from immigration are more widely shared.

References

Bier, David J. 2025. "Not Indentured: H-1B Visa Holders Have Changed Jobs 1.1 Million Times," Cato Institute at Liberty blog, October 29, 2025, available at https://www.cato.org/blog/not-indentured-h-1bs-have-changed-jobs-11-million-times

Blau, Francine D. and Christopher Mackie, eds. 2017. *The Economic and Fiscal Consequences of Immigration*. Washington, DC: National Academies Press.

Bound, John, Gaurav Khanna, and Nicolas Morales. 2018. "Understanding the Economic Impact of the H-1B Program on the United States." In *High-Skilled Migration to the United States and Its Economic Consequences*, edited by Gordon H. Hanson, William R. Kerr, and Sarah Turner, University of Chicago Press, 2018, pp. 109–175.

Bourveau, Thomas, Derrald Stice, Han Stice, and Roger White. 2025. "H-1B Visas and Wages in Accounting: Evidence from Big 4 Payroll and the Ethics of H-1B Visas," *Journal of Business Ethics* 199, no. B2 (2025): 309–330.

Chen, Jun, Shenje Hshieh, and Feng Zhang. 2021. "The Role of High-Skilled Foreign Labor in Startup Performance: Evidence from Two Natural Experiments." *Journal of Financial Economics* 142, no. 1: 430-452.

Costa, Daniel and Ron Hira. 2020. "H-1B Visas and Prevailing Wage Levels: A Majority of H-1B Employers—Including Major U.S. Tech Firms—Use the Program to Pay Migrant Workers Well Below Market Wages," Economic Policy Institute, May 4, 2020.

Daga, Vikash, and Noshir Kaka. 2006. "Taking Offshoring Beyond Labor Cost Savings," *The McKinsey Quarterly*, June 19, 2006.

Dimmock, Stephen G., Jiekun Huang, and Scott J. Weisbenner. 2022. "Give Me Your Tired, Your Poor, Your High-Skilled Labor: H-1B Lottery Outcomes and Entrepreneurial Success." *Management Science* 68 (September): 6950-6970.

Doran, Kirk, Alexander Gelber, and Adam Isen. 2022. "The Effects of High-Skilled Immigration Policy on Firms: Evidence from Visa Lotteries." *Journal of Political Economy* 130 (10): 2501–33.

European Migration Network and Organisation for Economic Co-operation and Development. *New and Innovative Ways to Attract Foreign Talents in the EU: EMN-OECD Joint Inform* (February 2025). European Migration Network, Directorate-General for Migration and Home Affairs, European Commission; Organisation for Economic Co-operation and Development. https://www.oecd.org/topics/policy-issues/migration-policies-returns-and-attracting-talent/migration-policy-debates-and-data-briefs.htm#New_and_innovative_ways_to_attract_foreign_talents_in_the_EU.

Ghosh, Anirban, Anna Maria Mayda, and Francesc Ortega. 2014. "The Impact of Skilled Foreign Workers on Firms: An Investigation of Publicly Traded U.S. Firms." IZA Discussion Paper No. 8684.

Gustafsson, Olle and Agnes Lindblad. 2025. Wage Differences in the H-1B Visa Program in the American Tech Sector. Bachelor's Thesis in Economics, University of Gothenburg, June 2025.

Hira, Ron. 2010. "Bridge to Immigration or Cheap Temporary Labor? The H-1B & L-1 Visa Programs are a Source of Both," Economic Policy Institute.

Hira, Ron. 2015. "New Data Show How Firms Like Infosys and Tata Abuse the H-1B Program," Economic Policy Institute, Feb. 19, 2015. https://www.epi.org/blog/new-data-infosys-tata-abuse-h-1b-program.

Hira, Ron, and Daniel Costa. 2021. "New Evidence of Widespread Wage Theft in the H-1B Visa Program: Corporate Document Reveals How Tech Firms Ignore the Law and Systematically Rob Migrant Workers." Washington, DC: Economic Policy Institute, December 9, 2021.

Hunt, Jennifer, and Marjolaine Gauthier-Loiselle. 2010. "How Much Does Immigration Boost Innovation?" *American Economic Journal: Macroeconomics* 2, no. 2 (April): 31–56.

Hunt, Jennifer and Bin Xie. 2019. "How Restricted is the Job Mobility of Skilled Temporary Work Visa Holders?" *Journal of Policy Analysis and Management*, 38 (Winter): 41–64.

Johnson, Matthew S., Kurt J. Lavetti, and Michael Lipsitz. 2025. "The Labor Market Effects of Legal Restrictions on Worker Mobility." *Journal of Political Economy* 133, no. 9 (September 1, 2025): 2735–2793. https://doi.org/10.1086/736217.

Kerr, William R., and William F. Lincoln. 2010. "The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention," *Journal of Labor Economics* 28 (July 2010): 473–508.

Kochhar, Rakesh, Kim Parker, and Ruth Igielnik. 2022. *Majority of U.S. Workers Changing Jobs Are Seeing Real Wage Gains.* Pew Research Center, July 28, 2022. https://www.pewresearch.org/social-trends/2022/07/28/majority-of-u-s-workers-changing-jobs-are-seeing-real-wage-gains/

Lofstrom, Magnus and Joseph Hayes. (2011). "H-1Bs: How Do They Stack Up to U.S. Born Workers?" IZA Discussion Paper No. 6259, Institute for the Study of Labor (IZA).

Mahajan, Parag, Nicolas Morales, Kevin Y. Shih, Mingyu Chen, and Agostina Brinatti. 2024. "The Impact of Immigration on Firms and Workers: Insights from the H-1B Lottery." IZA Discussion Paper No. 16917, Institute for the Study of Labor (IZA).

Ontiveros, Maria L. 2017. "H-1B Visas, Outsourcing and Body Shops: A Continuum of Exploitation for High Tech Workers." *Berkeley Journal of Employment & Labor Law* 38, no. 1 (2017): 1–47.

OECD (Organisation for Economic Co-operation and Development). *International Migration Outlook 2024*. Paris: OECD Publishing, 2024. https://doi.org/10.1787/50b0353e-en.

Peri, Giovanni, Kevin Shih, and Chad Sparber. 2015. "STEM Workers, H-1B Visas, and Productivity in U.S. Cities." *Journal of Labor Economics* 33, no. S1 (April): S225–S255.

Sicherman, Nachum. 1996. *"Gender Differences in Departures from a Large Firm."* Industrial and Labor Relations Review 49, no. 3 (April): 484–505. https://doi.org/10.1177/001979399604900307.

U.S. Citizenship and Immigration Services. 2026. H-1B Electronic Registration Process: Registration and Selection Statistics for FY 2021–FY 2024. Accessed January 2026. https://www.uscis.gov/working-in-the-united-states/temporary-workers/h-1b-specialty-occupations/h-1b-electronic-registration-process.

U.S. House of Representatives, 12 Mar. 2008, 110th Congress. *Testimony before the House Committee on Science and Technology: "Competitiveness and Innovation on the Committee's 50th Anniversary."* www.congress.gov/event/110th-congress/house-event/LC9246/text.

Varma, Roli, and James Rogers. 2020. "To Be or Not to Be on H-1B Visas: Engineers from India in the United States." *Perspectives on Global Development and Technology* 19 (2020): 281–302.

Figure 1. Distribution of payroll savings across H-1B workers

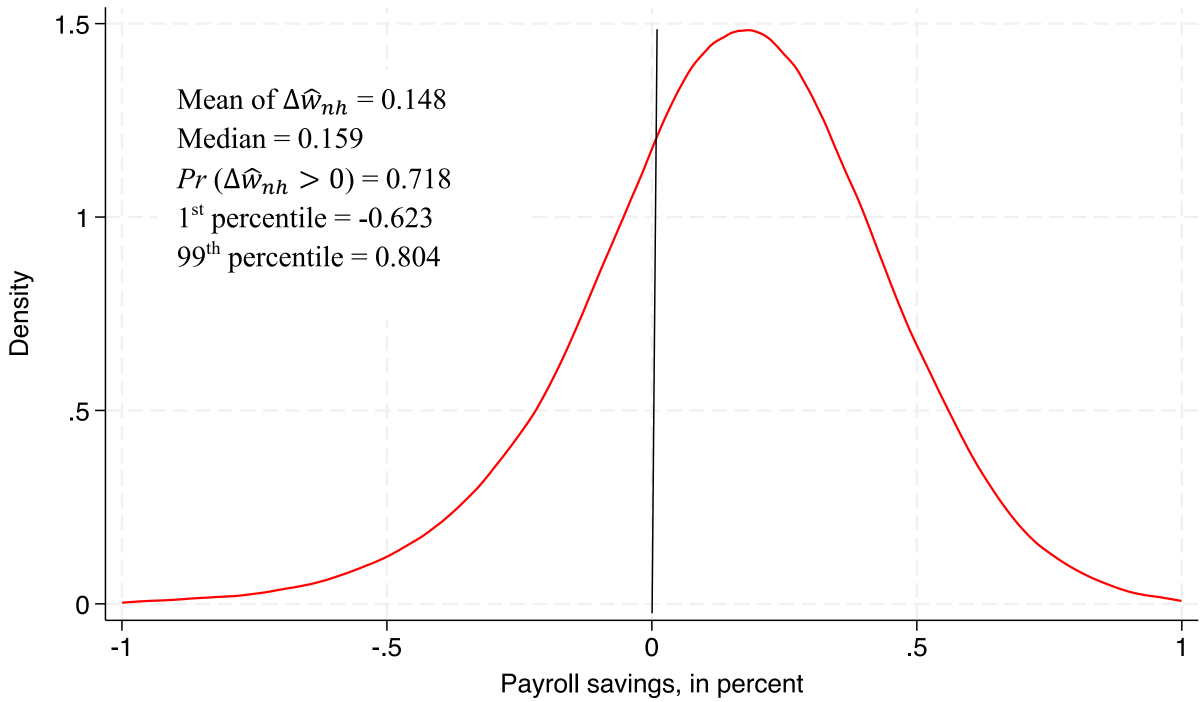

Notes: The "payroll savings" give the difference between the predicted log earnings a particular H-1B worker would earn if he were native and the worker's actual log earnings. The predicted log earnings are calculated from a regression estimated in the sample of native workers that controls for education, gender, age, occupation, and geography.

Figure 2. Impact of a fee on the probability of hire, number of visas, and revenue

A. Hiring rate

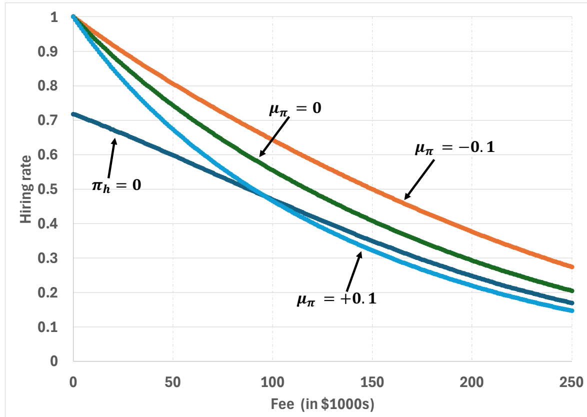

B. Number of visas used, low excess demand scenario

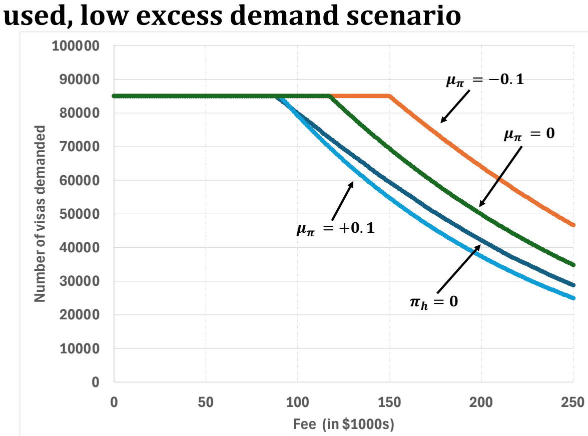

C. Revenue, low excess demand scenario

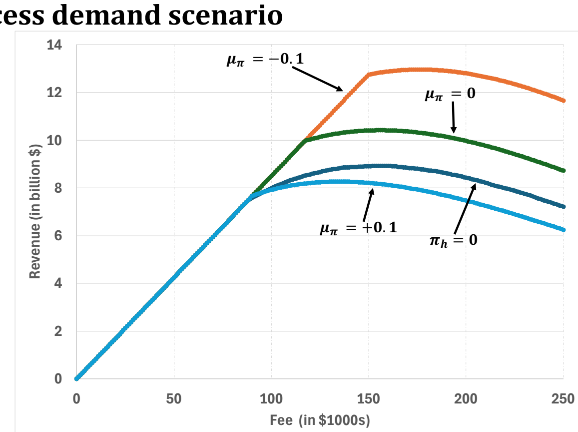

Notes: The hiring rate is calculated across 100 replications of the model simulating the firm's hiring behavior under alternative assumptions about the existence and size of the unobserved cost effect when hiring an H-1B worker. The simulation with $\pi_h = 0$ assumes there is no unobserved cost at all; the simulation with $\mu_\pi = k$ assumes that the mean of the distribution of unobserved cost effects is k (measured as the percent change on the cost of a hire). The "low excess demand" scenarios in panels B and C assume that the number of lottery registrations is twice the capped number of visas, or 170,000. The simulations use a 3 percent rate of discount.

Figure 3. Hiring rate at revenue-maximizing fee, by skill group

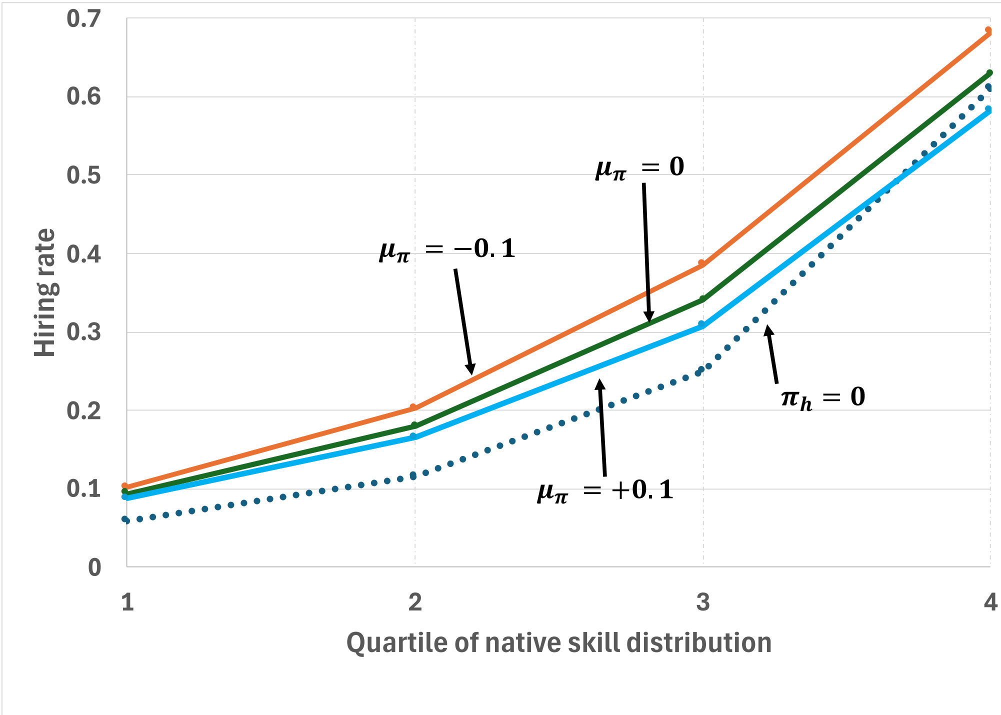

Notes: The hiring rate is the average calculated across 100 replications of the model simulating the firm's hiring behavior at the revenue maximizing fee under alternative assumptions about the existence and size of the unobserved cost effect (see Table 6 for the value of the revenue maximizing fee in each scenario). The average is calculated separately in the subsample of H-1B workers in each quartile of the native skill distribution. The simulations use a 3 percent rate of discount.

Table 1. Summary Statistics

	H-1B workers	Natives
Annual Salary (1000s)	$101.1	$103.1
Log annual salary	11.47	11.34
Demographics:		
Age	31.9	36.2
Male (%)	67.2	48.6
Bachelor's degree (%)	52.3	64.3
Master's degree (%)	41.8	26.8
Professional degree (%)	0.9	5.3
Doctoral degree (%)	5.0	3.6
Top 5 metros:		
1.	New York (13.2%)	New York (6.9%)
2.	Dallas (9.2%)	Chicago (3.5%)
3.	San Jose (8.2%)	Los Angeles (3.5%)
4.	Seattle (7.2%)	Washington (3.1%)
5.	San Francisco (6.0%)	Dallas (2.6%)
% in top 5 H-1B metros	43.9	13.3
Top 5 occupations:		
1.	Software developers (38.3%)	Elem/middle school teachers (5.2%)
2.	Computer occupations, all other (14.0%)	Other managers (4.6%)
3.	Computer systems analysts (5.3%)	Registered nurses (4.2%)
4.	Other math science occupations (3.2%)	Software Developers (2.9%)
5.	Electrical and electronics engineers (3.0%)	Accountants (2.6%)
% in top 5 H-1B occupations	63.8%	5.7
Sample size	343,608	252,664

Notes: The native sample is drawn from the 2023 ACS and consists of U.S.-born salaried workers who have at least a college degree, work full-time year-round, and aged 21-50. The sample of H-1B workers (also aged 21-50) is drawn from the I-129 applications for new visas filed by for-profit employers, selected by the lottery, and approved by DHS. The calculations use the ACS sampling weight for the native sample and a weight set equal to one for H-1B workers.

Table 2. Estimates of the H-1B wage gap

	Regression specification			
	(1)	(2)	(3)	(4)
Dependent variable = Log annual earnings				
A. Baseline regressions				
1. OLS	0.130	-0.048	-0.016	-0.161
	(0.002)	(0.004)	(0.004)	(0.004)
2. Quantile	0.154	-0.020	-0.003	-0.156
	(0.001)	(0.004)	(0.004)	(0.004)
B. Native sample reweighted to have same geographic distribution as H-1Bs				
1. OLS	-0.085	-0.076	-0.011	-0.230
	(0.006)	(0.007)	(0.006)	(0.012)
2. Quantile	-0.057	-0.017	-0.002	-0.227
	(0.006)	(0.006)	(0.006)	(0.012)
C. Native sample reweighted to have same occupation distribution as H-1Bs				
1. OLS	-0.196	-0.166	-0.306	-0.221
	(0.006)	(0.007)	(0.009)	(0.009)
2. Quantile	-0.190	-0.307	-0.308	-0.217
	(0.005)	(0.010)	(0.010)	(0.009)
Dependent variable: Annual earnings ($1000s)				
D. Baseline regressions				
1. OLS	-1.976	-26.177	-20.063	-30.916
	(0.229)	(0.763)	(0.719)	(0.850)
2. Quantile	9.595	-10.388	-11.608	-25.414
	(0.096)	(0.521)	(0.557)	(0.695)
Controls included in regression:				
Geography fixed effects		✓	✓	✓
Education, age, gender fixed effects			✓	✓
Occupation fixed effects				✓

Notes: Robust standard errors reported in parentheses. The regressions have 596,271 observations. The geography fixed effects indicate the state-PUMA combination of the worker's worksite (for H-1B workers) or place of residence (for native workers). The weight used in the regression is the ACS sampling weight for the native sample and is set to one for all H-1B workers.

Table 3. The H-1B wage gap in the 25 largest firms

	Firm	Number of visas	Average salary (1000s)	Adjusted log wage gap
1.	Amazon.com Services	13,301	130.1	-0.018
2.	**Infosys**	**10,198**	**82.7**	**-0.137***
3.	**Tata Consultancy Services**	**7,938**	**83.6**	**-0.115***
4.	**Cognizant Technology Solutions**	**6,308**	**88.0**	**-0.072***
5.	Google	5,179	145.1	-0.078*
6.	Microsoft	4,271	130.3	-0.034
7.	IBM	4,236	101.0	+0.041*
8.	Meta	3,944	149.2	-0.011
9.	**Wipro**	**3,436**	**78.9**	**-0.268***
10.	**Capgemini**	**3,090**	**97.6**	**-0.211***
11.	**HCL**	**2,828**	**86.8**	**-0.322***
12.	Apple	2,609	151.0	-0.009
13.	Intel	2,527	118.3	-0.011
14.	Accenture	2,463	98.3	-0.070*
15.	Ernst & Young	2,127	99.0	-0.039
16.	Amazon Web Services	2,095	117.7	-0.054*
17.	Amazon Development Center	1,889	127.2	-0.072*
18.	**Tech Mahindra**	**1,884**	**86.6**	**-0.348***
19.	Deloitte Consulting	1,836	95.9	-0.132*
20.	Oracle	1,617	133.9	-0.052*
21.	Wal-Mart	1,416	119.9	-0.020
22.	Qualcomm Technologies	1,390	124.5	-0.110*
23.	McKinsey & Company	1,289	152.0	+0.152*
24.	**Larsen & Toubro Infotech**	**1,242**	**102.9**	**-0.134***
25.	Tesla	1,036	118.8	-0.034
	All other firms	253,459	97.9	-0.185*

Notes: * Significant at the 5 percent level. The firm-specific adjusted log wage gaps are estimated from an OLS regression model that pools natives and H-1B workers and that includes education, age, gender, occupation, geography fixed effects, and a vector of indicator variables set to unity if the observation denotes an H-1B worker in a specific firm (and zero otherwise). The regression has 596,271 observations. The weight used in the regression is the ACS sampling weight for the native sample and is set to one for all H-1B workers. The bolded firms are considered to be outsourcing pipelines for H-1B workers.

Table 4. The within-occupation H-1B wage gap for the 10 largest occupations

	Occupation	No. of visas	Average salary ($1,000s)		Adjusted log wage gap
			Natives	H-1Bs	
1.	Software developers	131,435	146.9	107.6	-0.298*
2.	Computer occupations, all other	48,103	95.3	88.7	-0.063*
3.	Computer systems analysts	18,254	98.3	90.8	-0.153*
4.	Other mathematical science occupations	11,086	111.6	109.4	-0.100*
5.	Electrical and electronics engineers	10,444	113.1	110.2	-0.140*
6.	Computer programmers	9,834	109.5	82.2	-0.300*
7	Mechanical engineers	7,528	104.0	89.4	-0.227*
8	Operations research analysts	7,472	96.1	102.2	+0.068*
9.	Financial and investment analysts	6,181	130.7	114.9	-0.400*
10.	Management analysts	6,090	131.0	118.2	-0.160*
	All other occupations	87,181	101.1	98.9	-0.023*

Notes: * Significant at the 5 percent level. The log wage gaps are from a regression estimated separately within each occupation that pools the sample of H-1B and native workers in that occupation and that controls for age, educational attainment, gender, and geography fixed effects. The weight used in the regression is the ACS sampling weight for the native sample and is set to one for all H-1B workers.

Table 5. Distribution of H-1B workers across prevailing wage levels and quartiles of the skill index (% of H-1B workers in each category)

Prevailing wage level:	Amazon.com Services	Microsoft	Tech Mahindra	All firms	Log wage gap
I	61.2	52.8	0.0	26.2	-0.265
II	28.9	33.6	100.0	56.8	-0.156
III	8.8	12.5	0.0	12.1	0.020
IV	1.1	1.1	0.0	4.8	0.177
Skill index:					
1st quartile	2.9	3.0	1.4	6.4	0.042
2nd quartile	12.5	11.4	16.0	22.1	-0.028
3rd quartile	40.0	41.3	42.7	35.8	-0.118
4th quartile	44.8	44.4	40.0	37.7	-0.286

Notes: Firms report the prevailing wage level for each position to be filled by an H-1B worker in the LCA forms filed with DOL. The skill index is based on a log earnings regression estimated in the native sample. The index uses the coefficients from that regression to calculate the weighted sum of the “market value” of a particular H-1B worker’s observable characteristics (i.e., education, age, gender, occupation).

Table 6. Impact of a visa fee on demand for H-1B workers

Baseline number of lottery registrants prior to imposition of fee:	No unobserved costs ($\pi_h = 0$)	Mean of unobserved cost differential		
		μ_π = -0.1	μ_π = 0.0	μ_π = +0.1
A. Low excess demand (*N* = 170,000)				
Revenue-maximizing fee ($1000s)	$156.0	$176.0	$156.0	$135.0
Hiring rate at revenue-maximizing fee	33.7%	43.3%	39.3%	36.0%
Number of visas demanded (1000s)	57.2	73.6	67.2	61.2
Total revenue (billion $)	$8.9	$13.0	$10.4	$8.3
B. Moderate excess demand (*N* = 255,000)				
Revenue-maximizing fee ($1000s)	$159.0	$220.0	$181.0	$146.0
Hiring rate at revenue-maximizing fee	33.1%	33.9%	34.0%	34.2%
Number of visas demanded (1000s)	84.1	85.0	85.0	85.0
Total revenue (billion $)	$13.4	$18.7	$15.4	$12.4
C. High excess demand (*N* = 340,000)				
Revenue-maximizing fee ($1000s)	$199.0	$264.0	$223.0	$184.0
Hiring rate at revenue-maximizing fee	25.0%	25.0%	24.3%	24.3%
Number of visas demanded (1000s)	85.0	85.0	85.0	85.0
Total revenue (billion $)	$16.9	$22.4	$18.9	$15.6

Notes: The hiring rate used to calculate the revenue-maximizing fee and total revenue is the average from 100 replications of the model of the firm's hiring behavior. The level of excess demand gives an estimate of the number of lottery registrations (*N*) prior to the imposition of any fee; it equals twice, three times, or four times the number of capped visas (85,000). The unobserved cost differential in the last three columns is assumed to be normally distributed for the population of H-1B lottery registrants during the 2021-2024 period, and the variance is uniquely determined by the assumption that 99 percent of the values of π_h lie between -0.5 and +0.5 in the case where the mean is zero (μ_π = 0.0). The entire distribution is then shifted to the left or to the right for alternative assumptions about the mean. The simulations use a 3 percent rate of discount.

Table 7. Impact of a visa fee, allowing for job separations

Baseline number of lottery registrants prior to imposition of fee:	Mean of unobserved cost differential $\mu_\pi = -0.1$	$\mu_\pi = 0.0$	$\mu_\pi = +0.1$
A. Low excess demand (*N* = 170,000)			
Revenue-maximizing fee ($1000s)	$157.0	$137.0	$118.0
Hiring rate at revenue-maximizing fee	36.2%	33.5%	31.0%
Number of visas demanded (1000s)	61.5	57.0	52.8
Total revenue (billion $)	$9.7	$7.8	$6.2
B. Moderate excess demand (*N* = 255,000)			
Revenue-maximizing fee ($1000s)	$170.0	$138.0	$118.0
Hiring rate at revenue-maximizing fee	33.3%	33.3%	31.0%
Number of visas demanded (1000s)	85.0	85.0	79.1
Total revenue (billion $)	$14.4	$11.7	$9.3
C. High excess demand (*N* = 340,000)			
Revenue-maximizing fee ($1000s)	$213.0	$177.0	$144.0
Hiring rate at revenue-maximizing fee	25.0%	25.0%	24.9%
Number of visas demanded (1000s)	85.0	85.0	85.0
Total revenue (billion $)	$18.1	$15.0	$12.2

Notes: The hiring rate used to calculate the revenue-maximizing fee and total revenue is the average from 100 replications of the model of the firm's hiring behavior. The level of excess demand gives an estimate of the number of lottery registrations (*N*) prior to the imposition of any fee; it equals twice, three times, or four times the number of capped visas (85,000). The unobserved cost differential is assumed to be normally distributed for the population of H-1B lottery registrants during the 2021-2024 period, and the variance is uniquely determined by the assumption that 99 percent of the values of π_h lie between -0.5 and +0.5 in the case where the mean is zero ($\mu_\pi = 0.0$). The entire distribution is then shifted to the left or to the right for alternative assumptions about the mean. The distribution of job tenure is derived by assuming the annual separation rate of H-1B workers is 9.4 percent. The simulation uses a uniform random variable in [0, 1] to allocate workers to the six possible values of job tenure. The simulations use a 3 percent rate of discount.

Table 8. Impact of a visa fee, allowing for offshoring of H-1B jobs

Baseline number of lottery registrants prior to imposition of fee:	Mean of unobserved cost differential		
	μ_π = -0.1	μ_π = 0.0	μ_π = +0.1
A. Low excess demand (*N* = 170,000)			
Revenue-maximizing fee ($1000s)	$128.0	$108.0	$91.0
Hiring rate at revenue-maximizing fee	36.3%	34.9%	33.1%
Number of visas demanded (1000s)	61.8	59.3	57.6
Total revenue (billion $)	$7.9	$6.4	$5.2
B. Moderate excess demand (*N* = 255,000)			
Revenue-maximizing fee ($1000s)	$139.0	$113.0	$93.0
Hiring rate at revenue-maximizing fee	33.3%	33.3%	33.1%
Number of visas demanded (1000s)	85.0	85.0	84.5
Total revenue (billion $)	$11.8	$9.6	$7.9
C. High excess demand (*N* = 340,000)			
Revenue-maximizing fee ($1000s)	$175.0	$145.0	$119.0
Hiring rate at revenue-maximizing fee	25.0%	24.8%	25.0%
Number of visas demanded (1000s)	85.0	84.4	85.0
Total revenue (billion $)	$14.9	$12.2	$10.1

Notes: The hiring rate used to calculate the revenue-maximizing fee and total revenue is the average from 100 replications of the model of the firm's hiring behavior. The level of excess demand gives an estimate of the number of lottery registrations (N) prior to the imposition of any fee; it equals twice, three times, or four times the number of capped visas (85,000). The unobserved cost differential is assumed to be normally distributed for the population of H-1B lottery registrants during the 2021-2024 period, and the variance is uniquely determined by the assumption that 99 percent of the values of π_h lie between -0.5 and +0.5 in the case where the mean is zero (μ_π = 0.0). The random variable θ giving the (percent) payroll savings from offshoring is assumed to be normally distributed with a mean of -0.25 and a variance uniquely determined by the assumption that 99 percent of the values of θ lie in the [-0.5, 0] interval. The simulations use a 3 percent rate of discount.