

The Impact of H-1B Immigration on Wages and Unemployment Across U.S. Industries

Divyam Gupta

Econ 103B- Econometric Methods

Prof. Sanchita Mukherjee

Econ 103B- Term Paper

12 May 2026

Abstract

This paper examines the relationship between H-1B employment and labor market outcomes using industry-level panel data from 2015 to 2024. Results from a two-way fixed-effects model suggest that the H-1B share has no statistically significant effect on wages or unemployment, but the results are subject to small-sample-size and endogeneity limitations.

I Introduction

The tech industry has a rich history in the United States, and numerous global employers recruit staff. They do this in part by the H-1B visa program, which lets companies bring in foreigners to work in specific occupations like software developers and engineers. It is a popular program, particularly in the technology and finance sectors.

Over the past few years, there has been a greater awareness on the part of policymakers and economists of the impact the H-1B program has on U.S. workers. Some say that companies can leverage H-1B employees to cut down on labour expenses or to replace domestic employees. Others say that H-1B workers fill skill shortages and contribute to economic development. As a result, the impact of the program on wages and employment within the U.S. labor market is subject to debate.

This paper examines whether the employment of H-1B workers is associated with U.S. wages. In this paper, rather than making general claims about the H-1B program, this paper uses industry-level data to examine whether changes in H-1B share are associated with changes in wages and unemployment over time. The focus is on examining the relationship between H-1B employment and wage changes by industry. The question is relevant because it helps clarify the actual impact of H-1B employees on the labour market.

The results suggest that there is no strong relationship between H-1B share and wages in the industries studied. This could be attributable to the limited number of data points, though. Nevertheless, the H-1B program is still a key policy controversy and will probably continue to be subject to discussion in the future.

II Literature Review

The H-1B visa program has been debated for years because economists and policymakers do not fully agree on whether it helps or harms U.S. workers. Some studies suggest that H-1B workers fill skill shortages and help firms grow, while others argue that they may lower wages or change the composition of workers hired in certain occupations. Because of this disagreement, the research question is not just whether H-1B workers matter, but how they matter and what methods are used to identify their effects.

One of the most important studies in this area is Doran, Gelber, and Isen's paper on H-1B visa lotteries. The authors use the H-1B lottery as a natural experiment and compare firms that won visas with firms that did not. This is a strong research design because lottery outcomes create plausibly exogenous variation in H-1B access, which helps address endogeneity concerns. They find that access to H-1B workers affects firm outcomes such as employment and productivity, showing that the visa program can change firm behavior in meaningful ways rather than simply replacing native workers one-for-one (Doran et al., 2014).

Another important paper is Kerr and Lincoln's study of H-1B visa reforms and U.S. ethnic invention. They use reduced-form analysis to examine how changes in H-1B admissions affect

immigrant science and engineering employment and patenting outcomes across cities and firms. Their results show that high-skilled immigration can increase patenting and immigrant employment in the technology sector, which suggests that the program may also support innovation and firm growth rather than only competing with domestic labor (Kerr & Lincoln, 2010). This is relevant for my paper because it shows that the effects of H-1B workers may be broader than simple job displacement.

A third study by Mayda, Ortega, Peri, Shih, and Sparber examines the effect of the H-1B quota on the employment and selection of foreign-born labor. Their paper uses a quasi-experimental approach based on quota restrictions to study how changes in H-1B availability affect hiring. They find that tighter quota restrictions reduce the employment of new H-1B workers in for-profit firms, which suggests that firms do respond strongly to immigration policy changes (Mayda et al., 2018). This matters for my project because it supports the idea that the H-1B supply is not random and that policy shifts can change labor market outcomes in occupations where H-1B workers are concentrated.

More recently, Borjas studied the H-1B wage gap and employer demand using regression analysis with detailed controls for education, age, gender, occupation, and geography. His results suggest that H-1B workers earn less than comparable U.S.-born workers after controlling for observable differences, which raises questions about how firms use the program and whether wage differences are related to occupation-specific labor demand (Borjas, 2026). This paper is especially useful for my study because it shows why occupation controls and fixed effects are important when comparing labor market outcomes across groups.

Overall, the literature suggests that H-1B workers do not have a single effect across all labor markets. Some studies using natural experiments and reduced-form methods find that H-1B access increases firm growth, immigrant employment, and patenting, while other regression-based studies show that H-1B workers may receive different wages than comparable native workers. The gap in the literature is that fewer studies look at a small number of industries and compare both unemployment and wages in a simple industry-year panel. My paper tries to fill that gap by using fixed effects to study whether industries with higher H-1B shares have different labor market outcomes for U.S.-born workers.

III. Economic Model (with regression equation):

To look at whether the H-1B worker share has any meaningful relationship with wages, this paper uses a two-way fixed effects panel regression. The idea behind this approach is pretty straightforward; instead of comparing different industries to each other, the model looks at how things change within the same industry over time. This helps isolate the effect of the H-1B share from other factors that just happen to differ across industries. The regression equation used is:

$$Y_{it} = \beta_1(H1Bshare_{it}) + \alpha_i + \alpha_t + \varepsilon_{it}$$

Here, i stands for the industry and t stands for the year. The outcome variable in this paper is wages, which allows the model to focus on how the H-1B share is related to earnings in the labor market.

The main variable this paper cares about is $H1Bshare_{it}$. This measures what fraction of workers in a given industry and year are H-1B visa holders. It is calculated by taking the number of

certified H-1B approvals in an industry and dividing it by total employment in that same industry and year. The coefficient β_1 is what the analysis is really focused on, it tells us whether industries that experience an increase in H-1B share also tend to see a change in wages.

If β_1 turns out negative in the wage regression, that would mean a higher H-1B share is associated with lower wages for domestic workers. That would support the substitution argument, the idea that companies are using the H-1B program to bring in cheaper foreign labor instead of hiring and paying domestic workers fairly. On the other hand, if β_1 is positive or not statistically significant, that would support the complement argument, which says H-1B workers are filling roles that are genuinely hard to fill domestically and are not hurting wages in the process.

The α_i term in the equation represents industry fixed effects. These are important because industries are very different from each other in ways that have nothing to do with H-1B workers. Tech has always paid more than Health Care, and Finance has always required certain types of credentials. If those differences are not controlled for, the results would just be picking up on permanent characteristics of each industry rather than anything related to H-1B hiring. The α_t term represents year fixed effects, which account for factors that affect all industries at the same time, such as the 2020 pandemic-era unemployment spike or general wage growth across the economy. Together, these two sets of controls make sure the model is only picking up on within-industry changes over time.

The error term ε_{it} is just a catch-all for anything that affects wages that the model does not explicitly account for.

One important limitation to point out is endogeneity. The $H1Bshare_{it}$ variable is unlikely to be fully exogenous. For example, industries that are growing may both hire more H-1B workers and have different wage trends at the same time. In that case, the model could mistakenly attribute wage changes to the H-1B share when they are actually being driven by underlying industry growth. Fixed effects help reduce this problem by controlling for constant industry characteristics and common time shocks, but they may not fully solve it.

A better approach would be to employ an instrumental variable approach. A potential tool is adjusting of H-1B lottery, or changes to the national visa cap. These are policy driven changes that impact the number of H-1B workers that companies can hire, but do not result directly from wage changes in a specific industry. This fulfills the relevance criterion, since the actual H-1B share in an industry should be strongly connected with lottery outcomes and/or cap changes. The more visas that are available, or the number of lottery slots that are won, the larger the percentage of the H-1B should be.

In addition, the instrument would have to meet the exclusion restriction. If a change in the lottery or visa cap for H-1B were to impact wages, it should be in the context of H-1B jobs, and not through any other means. It is more plausible than just relying on the observed H-1B share, since the outcome of an H-1B lottery depends on the policy rules, not on the trend of wages in individual industries. But the exclusion restriction can still be flawed when there are wider economic conditions that influence visa policy changes. Due to data restrictions, this paper does not apply the IV approach, and the application of a lottery or a cap variation would aid the causal interpretation of future research.

There are a couple of limitations. First, individual-level controls like education level and gender composition are not included because the data do not support them. That said, since those workforce characteristics tend to be pretty stable within an industry over time, the industry fixed effects partially make up for this. Second, the H-1B share here is based on approved applications, not the actual count of H-1B workers physically present in each industry at any given point. This means the measure might slightly undercount true H-1B intensity, and for that reason, the results should be read as showing a correlation rather than proving causation.

The null hypothesis for the regression is:

$$H_0: \beta_1 = 0$$

This means the H-1B share has no statistically significant relationship with wages. The alternative hypothesis is:

$$H_a: \beta_1 \neq 0$$

This would mean there is a meaningful association between H-1B share and wages.

IV Description of Data

This study uses two primary data sources and three industries: Finance, Health Care and Tech between 2015 and 2024. In particular, each observation in the dataset is one industry per year, and there are 30 observations in total since there are three industries and ten years.

The wage and unemployment rate data were obtained from FRED, an online database that collects its information from the Bureau of Labor Statistics (BLS). The total number of approved H-1B visa applications was sourced from the USCIS H-1B Employer Data Hub, which reports on the number of H-1B visa applications approved in each year by industry. This dataset does have a small sample size as an important limitation. The analysis is limited to a number of

industries observed over time, and a relatively small number of observations. This reduces statistical power, and interpretation of the result should be undertaken with caution.

Table 1: Codebook

Variable Name	Description	Construction / Definition	Units	Variable Type	Year	Source
Wage (Annual)	Average annual wage in each industry	Reported wage data converted to annual values	USD	Dependent	2015–2024	FRED
unemp_occ	Unemployment rate in each industry	Directly taken from the FRED dataset	Percentage (%)	Dependent	2015–2024	FRED
H1B_Share	Share of H-1B workers in each industry	H1B_Count / Employment	Ratio (0–1)	Independent	2015–2024	USCIS / FRED
Employment	Total number of employees in each industry	Reported employment	Number of workers	Independent	2015–2024	FRED
H1B_Count	Number of H-1B approvals in each industry-year	Count of approved H-1B applications per industry-year	Number of approvals	Independent	2015–2024	USCIS

Table 2: Descriptive Statistics

Variable	N	Mean	Std Dev	Min	Max
Wage (Annual)	30	71,389	9,397	58,032	88,452

H1B_Count	30	18,515	10,662	4,589	38,884
Employment	30	1,547,500	892,622	423,000	2,895,000
H1B_Share	30	0.0223	0.0194	0.00182	0.0542
unemp_occ	30	3.377	1.232	1.900	7.500

There are 30 total observations, one for each industry-year pair. Wages ranged from about \$58,000 to \$88,000, with an average of around \$71,400. This variation is not surprising. Tech workers consistently paid the highest wages over the period, while Health Care was usually on the lower end.

Unemployment rates went from 1.9% to 7.5%, with the highest level occurring in Tech during 2020, likely due to COVID-19 disruptions. On average, unemployment across all observations was about 3.4%.

The H-1B share varied quite a bit across industries. It ranged from as low as 0.18% to as high as 5.4%. Tech had the highest shares throughout, generally around 4–5%, while Health Care stayed below 0.2% every year. Finance and Insurance fell in the middle, typically between 1% and 2.8%.

To better visualize trends in the data across industries and over time, the following figures are presented.




Figure 1 shows the variation in H-1B share across industries from 2015 to 2024. The technology industry consistently has the highest H-1B share, while health care remains very low throughout the period. Finance lies in between, reflecting moderate reliance on H-1B workers.




Figure 2 presents the trend in average annual wages across industries over time. Wages differ substantially across industries, with technology generally paying the highest wages and health care the lowest. These differences support the use of industry fixed effects in the regression analysis.

V. Model Estimation and Findings

Table 3: Estimated Model Summary Statistics

<i>Variables</i>	<i>Wage Model</i>	<i>Unemployment Model</i>
<i>H1B_Share Estimate</i>	-16,224	17.520
<i>Standard Error</i>	164,192	29.905
<i>t-value</i>	-0.0988	0.5858
<i>p-value</i>	0.9224	0.5657
<i>N</i>	30	30
<i>R²</i>	0.00057	0.01979
<i>Adj. R²</i>	-0.7049	-0.67212

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Assessing the Model: Goodness of Fit

This paper examines the R-squared values from both regressions; R-squared is used for testing the degree of the 'fit' of the models. The values of R^2 for the wage model and the unemployment model are 0.00057 and 0.01979, respectively. These values suggest that the variation in wages and unemployment in this data is attributable to factors other than H-1B share.

The adjusted R^2 values of the two models are negative: the wage model at -0.7049 and the unemployment model at -0.67214. This suggests that the models have limited explanatory power

after accounting for the number of variables included. However, the main concerns are the small sample size and the possibility of endogeneity rather than the R-squared values alone. The number of observations used in each of the three industries is quite small (30 observations over 10 years), so the model might not represent broader trends in the labor market.

Assessing the Model: T-Test

In the wage model, the coefficient on H1B_Share is -16,224. This implies that a one percentage point increase in H-1B share is associated with about a \$162 decrease in average annual wages, holding industry and year fixed effects constant. The negative direction is consistent with the substitution argument, but the result is not statistically significant.

The t value for the wage model is -0.0988, and the p value is 0.9224, which is greater than 0.05, which is a common significance level. The p-value is greater than 0.05, so the model does not reject the null hypothesis: $\beta_1 = 0$; that is, this dataset does not have a statistically significant relationship between H-1B share and wages.

In the unemployment model, H1B_Share's coefficient is 17.520. This suggests that a one percentage point rise in H-1B shares is roughly equivalent to a 0.175 percentage point rise in the unemployment rate, holding industry and year fixed effects constant. This is not statistically significant.

For the unemployment model, $t = 0.5858$ and $p = 0.5657$, which are both greater than 0.05. The model, therefore, does not reject the null hypothesis that $\beta_1 = 0$, that is, that in this data the H-1B share is not statistically significantly related to unemployment.

Overall, the results from both regressions indicate that H-1B share does not have a statistically significant effect on wages or unemployment in the industries studied over the time period considered.

Assessing the Model: F-Test

The F test is used to determine if the entire model is statistically significant. With the wage model, the f-value is 0.00976, and the p-value is 0.922. The F for the unemployment model is 0.343 with a p-value of 0.566.

Since both p-values are greater than 0.05, neither model is statistically significant overall. This suggests that, within this dataset, H-1B share does not provide a strong explanation for changes in wages or unemployment across industries over time.

VI Conclusion

To sum up, both regression results did not produce statistically significant findings. The wage model showed a negative coefficient, while the unemployment model showed a positive coefficient, but neither result was statistically significant. Based on this dataset, there is no strong evidence that H-1B share has a meaningful relationship with wages or unemployment for domestic workers in the industries studied.

These results connect to the larger debate in the literature. Some studies argue that H-1B workers help firms grow, fill skill shortages, and support innovation. Other studies argue that H-1B workers may lower wages or change hiring patterns in ways that hurt domestic workers. The findings in this paper do not clearly support either side. Instead, they suggest that, in this small

industry-level sample, H-1B share is not strongly related to lower wages or higher unemployment.

The biggest limitation of this paper is the small sample size. The dataset only looks at three industries over ten years, which gives a total of 30 observations. After including industry and year fixed effects, there are not many degrees of freedom left. This makes it harder for the model to find a statistically significant effect, even if one exists in the broader labor market. Because of this, the results should be interpreted carefully and should not be seen as proof that H-1B workers have no effect.

Other limitations include endogeneity. The industries that are expanding can simultaneously increase their number of H-1B employees while their average wages or number of job losses change. That makes it difficult to determine whether the H-1Bs are actually hurting labor market results, or whether the industries that are increasingly utilizing H-1B workers are hiring more workers in total, among them H-1B workers. Fixed effects can mitigate some of the industry or year differences, but they do not completely remedy this problem.

Some areas would benefit from future study of this paper, which would be better if more industries were used, a longer period of time, and better identification techniques. A potential way to do this would be to employ an instrumental variable, like the variation in H-1B visa caps or H-1B lottery returns, to more effectively control for the effect of H-1B workers. In general, this paper does not confirm that there is a statistically significant correlation between H-1B share and wages or unemployment in the sample analyzed, but this is an important topic that should be explored further using superior data and methods.

References

- Borjas, G. J. (2026). *The H-1B wage gap, visa fees, and employer demand* (NBER Working Paper No. 34793). National Bureau of Economic Research. <https://doi.org/10.3386/w34793>
- Doran, K., Gelber, A., & Isen, A. (2014). *The effects of high-skilled immigration policy on firms: Evidence from visa lotteries* (NBER Working Paper No. 20668). National Bureau of Economic Research. <https://doi.org/10.3386/w20668>
- Kerr, W. R., & Lincoln, W. F. (2010). The supply side of innovation: H-1B visa reforms and U.S. ethnic invention. *Journal of Labor Economics*, 28(3), 473–508. <https://doi.org/10.1086/651934>
- Mayda, A. M., Ortega, F., Peri, G., Shih, K. Y., & Sparber, C. (2018). The effect of the H-1B quota on the employment and selection of foreign-born labor. *European Economic Review*, 108, 105–128. <https://doi.org/10.1016/j.euroecorev.2018.06.010>
- U.S. Citizenship and Immigration Services. (n.d.). H-1B Employer Data Hub. U.S. Department of Homeland Security. <https://www.uscis.gov/tools/reports-and-studies/h-1b-employer-data-hub>
- Federal Reserve Bank of St. Louis. (n.d.). FRED Economic Data. <https://fred.stlouisfed.org/>