Extracting Firm Willingness to Pay Foreign Workers: Allocating H-1B Visas by Auction

Joseph Kuehn* Filippo Rebessi†

March 9, 2026

Abstract

In most years, demand for H-1B visas in the United States exceeds the statutory cap, and visas are allocated by lottery. We develop and estimate a structural model of firm demand for H-1B visas to evaluate alternative allocation mechanisms under a binding quantity constraint. Combining USCIS petition-level wage data with firm-level production data, we recover firm-specific marginal valuations of foreign skilled labor and construct inverse demand curves for visas. Exploiting the randomness of the H-1B lottery, we identify substitution elasticities and marginal products, allowing us to simulate price-based allocation rules. We show that replacing the lottery with a second-price auction substantially reallocates visas from large IT consulting firms to high-productivity technology firms. The implied equilibrium visa price ranges from $12,000 to over $40,000 per visa, increasing cumulative government revenue from approximately $1.5 billion to more than $5 billion over eight years. Aggregate surplus rises as visas are allocated to higher-valuation matches.

JEL Classification: J61, D22, C57

Keywords: H-1B Visa, Immigration Policy, Visa Allocation Mechanism

*California Polytechnic State University, San Luis Obispo (e-mail: jskuehn@calpoly.edu)

†California State University, East Bay (e-mail: filippo.rebessi@csueastbay.edu)

1 Introduction

Under the H-1B visa program, firms in the United States can temporarily employ foreign-born workers in occupations requiring specialized knowledge. Every year the total number of new H-1B visas is capped by Congress. If the number of new H-1B applications exceeds this limit, the visas are allocated via a lottery. While there is ample evidence that H-1B workers increase productivity in the U.S.[1], this program has fallen under scrutiny due to accusations that employers have used foreign workers on H-1B visas to replace their domestic counterparts at lower wages[2]. There are also concerns that the lottery mechanism does not allocate visas to the most productive workers due to the high demand for visas and the random nature of the allocation process.[3]

This paper studies how alternative allocation mechanisms reshape the distribution and economic impact of H-1B visas. We develop and estimate a structural model of firm-level demand for foreign skilled labor that allows firms to differ in production complementarities, wage-setting, and hiring frictions. The model delivers firm-specific willingness to pay for H-1B visas. We use these estimated demand schedules to simulate a counterfactual second-price auction and compare its allocation and surplus outcomes to those generated by the current lottery.

Our framework speaks directly to the design of allocation rules when high-skilled immigration is quantity-constrained. Both the first Trump and Biden administrations have attempted to replace the current lottery with a system that prioritizes allocating visas to workers being paid the highest wages.[4]. The recent Presidential Act setting H-1B fees for selected H-1B visas at $100,000 represents a shift towards a price-based mechanism.

[1]Peri, Shih, and Sparber (2015b) shows that foreign STEM workers increase total factor productivity growth in U.S. cities. Mayda et. al. (2023) finds that the visa cap has led to substantial declines in growth for companies that rely on the H-1B program. Bound, Khanna, and Morales (2017) use a general equilibrium model of the U.S. economy to show that H-1B visa holders increase the overall welfare of U.S. natives.

[2]Wal-Mart was accused by the AFL-CIO in 2015 of using H-1B workers for information technology (IT) jobs that could have been filled by domestic labor (AFL-CIO (2015)). Disney (Preston (2015a)) and Toys 'R' Us (Preston (2015b)) were similarly accused of replacing domestic employees with workers on H-1B visas.

[3]Sharma and Sparber (2020) shows that the randomness of H-1B lotteries creates search externalities and leads to an efficiency loss. Sparber (2018) argues that alternatives to a lottery would increase output, output per worker, and wages. In discussing how critics across the political spectrum have accused Big Tech of exploiting the visa system to reduce labor costs, Cassidy (2017) observes that the lottery mechanism facilitates this by ignoring productivity and merit in its randomness. Neufeld (2025) also notes that the H-1B lottery system's fundamental reliance on randomness undermines its ability to select the most worthwhile candidates.

[4]The Trump administration's *Buy American and Hire American Executive Order* sought to reform the H-1B allocation process by giving preference to the most highly skilled or best-compensated applicants. See Kreighbaum (2025) and Anderson (2025)

By recovering firm-level demand curves for foreign labor, we are able to simulate alternative allocation mechanisms and trace their consequences for government revenue, wage-setting, and surplus distribution across firms. This allows us to evaluate not only whether an auction reallocates visas, but whether it does so in ways consistent with policy objectives.

The analysis uses a new worker-level dataset that combines several sources. Thanks to a Freedom of Information Act request to the United States Citizenship and Immigration Services (USCIS), we have obtained data on wages and other individual characteristics for each H-1B petition received during fiscal years 2014 and 2015. We combine this information with H-1B approved petitions by employer from the USCIS H-1B Employer Data Hub, and data on sales, employment and capital for publicly listed firms from Compustat.

We use this data to estimate a structural model of firm demand for H-1B workers. Firms produce using capital, domestic labor, and foreign labor through a nested CES structure, allowing for imperfect substitution between domestic and foreign skilled labor. Wages are determined through Nash bargaining, and hiring foreign workers entails both a monetary visa cost and firm-specific search costs. The model delivers, for each firm-year, a schedule mapping the number of potential foreign hires into a willingness to pay for additional visas. Estimation exploits the randomness of the H-1B lottery to instrument for foreign employment and recover the marginal product of foreign labor.

The estimated model reveals substantial heterogeneity in firms' valuations of foreign workers. For the median firm-year in our sample, the marginal product of an H-1B worker exceeds that of a comparable domestic skilled worker by approximately $20,000. However, this difference varies widely across firms and industries. In technology-intensive sectors such as information and advanced manufacturing, foreign workers generate particularly high marginal products, while in large IT consulting firms, the relative value of foreign labor is more closely tied to wage differentials. These differences translate into heterogeneous demand schedules. Firms differ both in the level of their marginal valuations and in how quickly those valuations decline as additional visas are obtained. The allocation mechanism therefore determines which high-value margins are realized.

We use the estimated demand curves to simulate a counterfactual second-price auction in which firms bid for visas. Relative to the current lottery, the auction reallocates visas toward firms with higher marginal valuations. Large IT consulting firms lose a substantial share of visas, while firms

in the information sector and advanced manufacturing gain. For example, Cognizant Technology Solutions (one of the largest recipients of H-1B visas in our data) would lose more than 90% of its visas, while Microsoft would increase its H-1B employment by over 100%.

Across years in our sample, the implied auction price ranges from roughly $12,000 to over $40,000 per visa, compared to the current administrative fee of $6,470. Over the eight years in our sample, the auction would increase cumulative government revenue from approximately $1.5 billion under the lottery to more than $5 billion. We also use the model to evaluate alternative price-based reforms, including a high fixed visa fee, and targeted auction credits. These exercises show that while uniform fees can raise revenue, they do not replicate the allocative efficiency gains of an auction, and targeted adjustments generate smaller and more limited reallocations.

More broadly, our results show that when access to foreign labor is restricted, the allocation mechanism plays a central role in shaping economic outcomes. Alternative rules can substantially reallocate employment across firms, alter surplus distribution, and affect aggregate productivity. By estimating a structural model of firm-level demand for foreign labor, we are able to evaluate policy reforms using the underlying production and wage-setting margins that determine firms' valuations.

1.1 Related Literature

High-skilled immigration is widely recognized as a key driver of innovation and entrepreneurship in the U.S. (Kerr (2013)). A large literature documents the effects of the H-1B program across multiple domains. For example, H-1B workers have been shown to increase patenting and firm productivity (Kerr and Lincoln (2010),Doran, Gelber, and Isen (2022), Ghosh, Mayda, and Ortega (2014)), contribute to total factor productivity growth in U.S. cities (Peri, Shih, and Sparber (2015b)), complement native born IT workers in the U.S. (Peri, Shih, and Sparber (2015a)[5]), and improve aggregate welfare in general equilibrium (Bound, Khanna, and Morales (2017)). Other studies highlight effects on the pipeline of international students (Kato and Sparber (2013), Amuedo-Dorantes and Furtado (2017)), the growth of young firms (Mandelman, Mehra, and Shen (2025)), and the access to venture capital funding (Dimmock, Huang, and Weisbenner (2022)). Collectively, this literature supports the claim that increasing the supply of H-1B visas raises U.S. productivity

[5]There is also a broader literature on the degree of complementarity between native and foreign workers using national-level data, including Borjas (2003) and Ottaviano and Peri (2012)

and innovation.[6]

At the same time, concerns persist about the displacement of domestic workers and wage suppression. Kerr, Kerr, and Lincoln (2015), Matloff (2004) and Doran, Gelber, and Isen (2022) examine these crowding-out effects, underscoring that the program's benefits may depend critically on how visas are allocated. Our results contribute to this debate by showing that an auction mechanism could reduce the share of visas issued to firms hiring lower-wage foreign workers, while increasing the share of visas awarded to firms for which foreign-born workers are most productive.

Given the debate over the current allocation system, there is also a large literature on the design of the H-1B allocation mechanism. Sparber (2018) shows that prioritizing highly educated applicants rather than using a lottery would increase output, as well as wages paid to less-educated workers, and estimates a $26.5 billion gain for the economy over a six year period. Mayda et al. (2018) analyze how the quota shapes the employment and selection of foreign-born labor, while Glennon (2024) finds that tighter H-1B constraints spur offshoring by U.S. firms. Pathak, Rees-Jones, and Sonmez (2025) provide a formal market-design perspective, showing that different legally valid implementations of the lottery rules can shift the distribution of advanced-degree visas by as many as 14,000 slots annually. Chassamboulli and Peri (2020) use a structural model to simulate labor-market outcomes under alternative immigration policies, again underscoring how policy design shapes wages, employment, and output.

Our paper contributes to this literature by evaluating a specific mechanism (a second-price auction) and quantifying how it would redistribute visas and better align them with firms' willingness to pay. Both Peri (2012) and Orrenius and Zavodny (2020) analyze the merits of moving from a lottery to an auction, and Kerr and Kerr (2020) discuss an auction as an alternative to the current system. Recent work by Borjas (2026) studies the implications of the H-1B wage gap for optimal visa fees and government revenue, emphasizing pricing rather than allocation under a binding cap. By combining USCIS petition-level data with firm-level production information, we provide a quantitative assessment of how an auction-based mechanism would reallocate visas across heterogeneous firms.

[6] *"I want to emphasize that to address the shortage of scientists and engineers, we must do both - reform our education system and our immigration policies. If we don't, American companies simply will not have the talent they need to innovate and compete"*, Bill Gates, Testimony at the U.S. House of Representatives, Committee on Science and Technology, March 12th, 2008

The remainder of the paper proceeds as follows. Section 2 describes the institutional background and the data. Section 3 presents the structural model. Section 4 estimates the model and recovers firm-level demand for H-1B visas. Section 5 evaluates counterfactual allocation mechanisms and quantifies their effects on visa allocation, revenue, and surplus. Section 6 concludes.

2 The H-1B Visa Program: Background and Data

This section lays out the institutional background and data underlying our analysis. We begin by describing the H-1B visa program and the timing of the application and lottery process that generates the variation we exploit. We then present the administrative and firm-level data sources, and conclude by detailing the construction of the key variables used in the estimation and counterfactual exercises.

Under the H-1B visa program, U.S. firms can temporarily employ foreign workers in specialty occupations that require a bachelor's degree or higher. These visas are sponsored by the firms, but must be tied to a specific individual. The H-1B visas are considered nonimmigrant visas because they only allow the worker to stay in the U.S. for three years. After three years, the firm can apply for a one-time renewal for a maximum of six years, after which the worker must either leave the U.S. or apply for permanent residency.

Since the program began in 1991, the number of new H-1B visas issued by the U.S. Citizenship and Immigration Services (USCIS) has been subject to a congressional cap. The initial cap was set at 65,000 visas, rising to 115,000 in 1999 and 195,000 in 2001, before going back down to its current level of 65,000 visas in 2004. Since 2006, the USCIS has also reserved 20,000 additional visas for foreign worker's with at least a master's degree or higher from a U.S. university. It is important to note that universities, nonprofit, and government research organizations are exempt from the cap. Applications from Chile and Singapore are also not counted towards the cap, thanks to a free trade agreement that allocates 1,400 and 5,400 visas, respectively, to workers from these countries[7]. Renewals are also not subject to the cap.

In order to obtain an H-1B visa, employers submit a labor condition application (LCA) together with their H-1B visa petition, up to 6 months before the starting date of the employment relation-

[7]See `https://www.dol.gov/agencies/whd/immigration/h1b1`

ship (which is set at the beginning of the government's fiscal year, in October). USCIS reviews new petitions starting April 1st of each year, and if the number of petitions received exceeds the congressional cap, H-1B visas are assigned with a lottery. Once an H-1B petition is successfully awarded to a firm, the related processing fee is collected, while petitions not awarded a visa are returned, without any fee collection.

Workers on H-1B visas are allowed to move to other firms. However doing so incurs extra costs compared to a domestic worker. First, H-1B visas are not portable, and the new employer needs to submit a new H-1B visa application and associated fees to USCIS. Second, switching employers may involve the risk of restarting the "green card process" for workers who are the beneficiary of an employment-based petition for permanent residency[8]

Firms petitioning for an H-1B worker need to comply with the following provisions in their application for H1B Nonimmigrants: "H-1B nonimmigrants will be paid at least the actual wage level paid by the employer to all other individuals with similar experience and qualifications for the specific employment in question or the prevailing wage level for the occupation in the area of employment, whichever is higher," and "The employment of H-1B nonimmigrants will not adversely affect the working conditions of workers similarly employed in the area of intended employment" (see the Department of Labor's Labor Condition Application). Petitioners are charged several fees: a $10 registration fee, a $460 filing fee, a $500 antifraud fee, a $1,500 ACWIA Education and Training Fee (for employers with more than 25 employees), and an additional $4,000 "Public Law 114-113 Fee" for employers who have more than 50 employees and at least 50% of their workforce on an H-1B or L-1 visa. The total in mandatory fees ranges between $2,470 and $6,470, with an optional additional $2,500 for premium processing, as well as legal fees.[9] In the analysis below we fix the total cost of acquiring a H-1B visa under the lottery to $6,470.

Once the cap is reached, visa applications are subject to a random lottery. This has happened for every year since 2014 (Table 1). A lottery is first held to allocate the 65,000 standard H-1B visas, and then a second lottery is held to allocate the 20,000 additional visas to master's degree

[8]In order for the green-card application to transfer between employers, the petition has to have been pending for 180 days or more, and the new position has to be considered the same or similar to the current job. The approval decision is at the discretion of USCIS officers. For details, see the American Competitiveness in the Twenty-first Century Act of 2000. Depew, Norlander, and Sorensen (2013) uses data on 6 Indian IT firms for a total of 72,606 employees and finds that, between 2003 and 2011, total voluntary quits were 22%.

[9]See the Instructions by the Department of Homeland Secruity

Table 1: Total Visa Applications Per Year

Year	Cap	Petitions	Probability
2014	85,000	124.000	68.55%
2015	85,000	172,500	49.28%
2016	85,000	233,000	36.48%
2017	85,000	236,000	36.02%
2018	85,000	199,000	42.71%
2019	85,000	190,098	44.71%
2020	85,000	201,011	42.29%
2021	85,000	275,000	30.91%
2022	85,000	308,613	27.54%
2023	85,000	483,927	17.56%
2024	85,000	780,884	10.89%
2025	85,000	479,953	17.71%
2026	85,000	358,737	23.69%

Data on the total number of petitions submitted each year are from redbus2us.com.

holders from U.S. institutions that were not selected in the initial lottery.[10] Winners of the lottery have 90 days to file a complete H-1B petition, while losers must instead hire either a U.S. citizen or permanent resident, hire a foreigner on another visa, or not hire any one for the position.

The majority of awarded H-1B petitions have gone to workers in system analysis and programming (in our dataset, over 50% of the visas are in this occupation), with a strong concentration in a handful of IT consulting firms (e.g. Cognizant, Wipro). The system has been under scrutiny, as it's often seen as a means to replace domestic workers with cheaper foreign workers. Several market solutions to immigration reform have been proposed to allocate visas to drive economic growth and increase government revenue (see for example Orrenius and Zavodny (2020) and Becker and Lazear (2013)).

[10]This order has been in place since 2020. Prior to that year, the lottery for the advanced degrees was held before the general lottery.

2.1 Data Description and Variables Construction

We combine USCIS administrative data on H-1B petitions, approvals, and worker wages with firm-level financial data from Compustat, supplemented by firm payroll data from BrightQuery[11] and occupation- and industry-level wage data from the BLS's Occupational Employment and Wage Statistics (OEWS) and Quarterly Census of Employment and Wages (QCEW). These sources allow us to construct firm-year stocks of H-1B workers and measure wages for foreign and domestic labor. A central element of our analysis is firms' marginal product of foreign labor. Estimating this object requires isolating variation in foreign employment that is not driven by firms' underlying productivity or demand conditions. To obtain this kind of variation and discipline how output responds to foreign labor, we construct the "surprise number of visas". We exploit the fact that, for cap-subject new petitions, the number of H-1B visas ultimately granted to firms differs from the number implied by their applications because of the H-1B lottery. Conditional on the volume of new petitions submitted by a firm, lottery outcomes generate idiosyncratic variation in realized visa allocations that is not chosen by the firm.

Our worker-level administrative data from USCIS reports wages and individual characteristics for each H-1B petition, received during fiscal years 2014 and 2015 (with actual employment starting in fiscal years 2015 and 2016, respectively)[12]. In particular, we observe the annual salary of the beneficiary, their education level, occupation, and birth country. In total, we have 694,092 workers' observations. The mean (median) salary of a worker in this dataset is $85,268 ($76,205). A majority of workers (56.37%) are in the occupation labeled "Systems Analysis and Programming," and more broadly 66.7% of workers are in "Computer Occupations." Most workers have either a bachelor's degree (45.68%) or master's degree (43.38%). The country where the most H-1B workers in our dataset comes from is India with 71.53% of the workers. The second largest is China, with only 8.67%.

The wage data includes the name of the firm employing the H-1B worker. We utilize this information to link workers to publicly traded firms in the Compustat database according to the following procedure. Initially, we match firms with identical names in both datasets. Subsequently, we employ a similarity score algorithm to compare remaining firm names in the USCIS dataset with

[11] https://doi.org/10.82551/Q56V-PX08

[12] This data was obtained from a Freedom of Information Act request to USCIS.

Table 2: Wage Data

	Total	**Matched**
Number of Obs	694,092	208,979
Worker Salary		
Mean ($000)	85.268	92.962
Median ($000)	76.205	88.000
Std Dev ($000)	59.283	27.828
Worker Characteristics		
Computer Occupations	66.72%	82.62%
Bachelor's Degree	45.68%	53.89%
Master's Degree	43.38%	41.94%
China	8.67%	7.12%
India	71.53%	82.11%

Data comes from the USCIS. Each observation is a worker on a H-1B visa. The "Total" column describes the full dataset from USCIS. The "Matched" describes the worker observations that were matched to publicly traded firms in the Compustat database.

those in Compustat. Matches exceeding a specified score undergo manual verification to confirm accuracy. In total, we are able to match 208,979 of the H-1B workers with a firm in Compustat (30.1% of the entire sample). Table 2 summarizes the differences between the matched dataset and the full dataset of wages. Average salaries are higher in the matched dataset (mean of $92,962 versus $85,268), and a larger percentage of workers come from India and are in computer occupations.

We exclude workers with reported salaries that are either below $10,000 (388 observations) or greater than $500,000 (104 observations). The lower cutoff helps to filter out observations reflecting part-time employment or reporting errors, and the upper cutoff mitigates the influence of high-level managerial positions where salary negotiation dynamics differ significantly from those of the typical H-1B worker.

The occupation categories reported in the USCIS H-1B data do not align directly with the Standard Occupational Classification (SOC) system used by the Bureau of Labor Statistics. USCIS instead reports a set of broad administrative occupation labels (e.g., "Occupations in Systems Analysis and Programming," "Physicians and Surgeons," "Other Computer-Related Occupations", etc.). We harmonize these labels to 2-digit SOC occupation groups used in the BLS Occupational

Table 3: Summary of Wage Data by Occupation

		Salary				
Occupation	**N**	**Mean**	**Median**	**Std Dev**	**5%**	**95%**
Executives	126	132.78	115.00	77.05	63.50	300.00
Sales Managers	661	116.03	112.50	44.65	53.30	190.00
Operations Managers	2,470	106.95	100.43	42.57	54.00	182.40
Other Management	3,184	118.52	115.00	40.97	70.00	187.20
Financial Specialist	1,953	94.78	90.00	33.40	53.00	150.00
Computer Occupations	172,657	90.77	85.00	26.13	60.00	139.85
Mathematical Science Occupations	1,047	97.53	96.00	27.50	59.89	141.90
Architects and Surveyors	2,564	106.12	102.79	28.87	65.00	155.00
Engineers	18,922	101.43	97.47	25.07	69.22	145.20
Life Scientists	1,105	77.47	72.00	31.09	43.15	126.83
Physical Scientists	853	105.76	100.00	33.99	63.21	160.00
Social Scientists	1,027	103.69	99.00	37.36	56.74	175.00
Lawyers	23	96.25	95.00	40.87	40.00	165.75
Postsecondary Teachers	10	95.11	101.05	29.69	53.02	130.00
Librarians	14	88.08	91.35	32.34	45.00	140.75
Other Education	17	91.81	90.06	15.79	72.00	140.00
Art and Design Workers	586	104.08	98.05	43.37	50.00	180.49
Entertainers and Sports Workers	39	98.12	95.00	30.36	57.20	146.00
Media and Communication Workers	87	88.34	84.23	36.19	44.85	158.00
Healthcare Practitioners	758	114.11	111.46	49.08	25.84	164.33
Health Technicians	52	80.10	72.82	35.75	47.27	139.93
Other Healthcare	187	113.73	102.71	55.59	60.05	236.00
Other	637	109.37	105.00	36.43	63.00	166.60

Salaries and occupations come from the USCIS. Occupation titles from the USCIS were combined and matched with the SOC minor occupation groupings from the U.S. BLS.

Table 4: Summary of Wage Data by Degree

		Salary				
Education	**N**	**Mean**	**Median**	**Std Dev**	**5%**	**95%**
Bachelor's Degree	112,621	87.89	81.15	25.97	60.00	135.88
Master's Degree	87,647	97.34	94.01	27.63	61.20	145.00
Doctorate Degree	7,838	113.98	111.15	30.62	61.50	161.53
Professional Degree	499	129.77	117.50	60.62	65.00	256.47
Unknown Degree	374	104.56	90.00	51.44	60.00	208.00

Salaries and worker degrees come from the USCIS.

Employment and Wage Statistics (OEWS) using a manually constructed crosswalk that assigns each USCIS occupation category to a single SOC group. This harmonization, which underlies all

Table 5: Summary of Wage Data by Country of Origin

		Salary				
Country	**N**	**Mean**	**Median**	**Std Dev**	**5%**	**95%**
Canada	2,906	126.65	120.00	39.03	78.00	195.00
Mexico	986	103.85	100.68	31.79	62.99	152.00
China	14,888	105.39	105.00	26.23	62.84	148.50
India	171,588	89.58	84.13	25.00	60.00	135.20
America (Low-Middle Income)	100	89.20	84.21	28.60	49.56	142.90
America (Upper-Middle Income)	1,762	106.77	102.54	39.74	58.41	162.16
America (High Income)	221	112.10	108.47	41.94	60.50	171.50
Europe (Low-Middle Income)	357	114.27	111.00	27.72	76.00	160.80
Europe (Upper-Middle Income)	1,753	112.03	110.00	30.70	63.79	160.56
Europe (High Income)	4,102	125.56	120.00	45.89	62.40	202.26
Asia (Low-Middle Income)	3,095	91.04	88.25	27.38	55.51	135.81
Asia (Upper-Middle Income)	395	100.35	98.30	31.24	59.50	150.00
Asia (High Income)	3,330	110.43	109.04	34.83	58.20	165.61
Africa (Low Income)	149	93.49	92.00	30.05	56.73	145.00
Africa (Low-Middle Income)	682	94.73	91.78	31.23	58.00	140.01
Africa (Upper-Middle Income)	143	122.64	110.00	41.56	75.60	200.00
Middle East (Low Income)	33	102.05	110.00	22.10	64.18	126.18
Middle East (Lower-Middle Income)	734	109.78	109.88	26.62	67.65	155.00
Middle East (Upper-Middle Income)	221	107.53	105.18	29.77	66.98	157.80
Middle East (High Income)	724	110.00	105.50	37.29	61.10	170.00
Unknown	810	86.78	76.66	29.62	60.00	139.69

Salaries and country of origin come from the USCIS. Countries were put into categories based on their region and GNI per capita (from the World Bank). Low income is defined as a GNI per capita of less than $1,045, lower-middle income as GNI per capita between $1,046 - $4,095, upper-middle income as GNI per capita between $4,096 - $12,695, and high income as a GNI per capita above $12,696. The five regions are America, Europe, Asia, Africa, and the Middle East. No countries fell into the following categories: Low Income America, Low Income Europe, Low Income Asia, High Income Africa.

occupation-based analyses in the paper, ensures consistency. Table 3 summarizes the salary data by occupation. The occupation with the highest average salary for H-1B workers is executives with a mean salary of $132,782. The lowest is life scientists with a mean salary of $77,467.

Education is classified into five categories: Bachelor's, Master's, Doctorate, Professional, and Unknown. Table 4 summarizes the salary data by education. Average salaries are increasing in the level of the degree with H-1B workers with a bachelor's degree earning the lowest average salaries ($87,893) and those with a doctorate or professional degree earning the highest average salaries

($113,982 and $129,767, respectively).

We observe each worker's country of origin. We categorize them based on their World Bank designation as either low income (GNI per capita less than $1,045), lower middle-income (GNI per capita between $1,046 - $4,095), upper middle-income (GNI per capita between $4,096 - $12,695), and high income (GNI per capita above $12,696), and the country's region.[13] We also create a separate category for Canada, Mexico, China, and India, since they account for roughly 90% of the observations. These classifications are later used to parameterize workers' outside options in the model estimation. Table 5 summarizes the salary data by country of origin. The highest average salary is for workers coming from Canada ($126,653), while the lowest among the individual countries is for workers coming from India ($89,579).

To estimate firm-level production functions and construct the labor inputs used in the model, we combine the worker-level data described above with firm-level financial and employment data from Compustat. We measure output using a standard value-added approach, subtracting material costs from net sales:

$$Y_{jt} = \text{Sales}_{jt} - \text{Materials}_{jt}$$

where material costs are defined as cost of goods sold net of labor expenses. Given the prevalence of missing data on labor expenses in Compustat,we supplement our analysis with data from the Quarterly Census of Employment and Wages (QCEW). From QCEW, we obtain average wages by firm size and industry. We then impute labor expenditure to each firm as this average wage $\widehat{W}_{jt}$ multiplied by the total number of employees from Compustat:

$$\text{Materials}_{jt} \;=\; \text{COGS}_{jt} \;-\; \widehat{W}_{jt} \cdot \text{Employment}_{jt}$$

Both output and material costs are deflated using GDP price indexes by industry sourced from the BEA.

Because changes in H-1B employment may be correlated with unobserved firm-level productivity shocks or demand conditions, using the stock of H-1B workers directly would not, by itself, isolate the marginal contribution of foreign labor to firm output. To overcome this, we use data on firms'

[13] We combine North America, Latin America, and the Caribbean together, and combine South Asia and East Asia and the Pacific.

Table 6: Total Number of H1B Visas and LCAs by Year

	All				**Matched**			
Year	Total H1Bs		Total LCAs		Total H1Bs		Total LCAs	
	New	Continuing	All	Lottery	New	Continuing	All	Lottery
2015	130,616	202,190	906,421	298,202	36,522	76,497	277,661	51,095
2016	111,846	176,656	1,148,959	344,982	33,735	69,173	379,236	80,114
2017	105,092	252,120	1,171,462	343,102	30,306	99,223	376,786	68,545
2018	96,166	277,229	1,063,039	343,802	31,040	109,073	371,144	91,158
2019	87,889	247,072	1,181,113	318,006	23,961	92,389	411,070	85,048
2020	132,967	256,356	973,024	334,421	42,178	96,138	408,302	128,371

Data on total new and continuing H-1B visas comes from the USCIS's H-1B Employer Data Hub. Data on each Labor Condition Application (LCA) submitted to the Department of Labor comes from the Office of Foreign Labor Certification (OFLC). "Lottery" LCAs are applications that were submitted in February or March for work start dates that are either 5 or 6 months in the future. "Matched" visas are visas that are matched to firms in the Compustat dataset.

Labor Condition Applications (LCAs) submitted to the Department of Labor[14] to build firms' "surprise number of visas" for fiscal years ranging from 2015 to 2020[15]. Since the H-1B lottery applies only to a specific subset of visa petitions, it is important to clarify how firms' application timing interacts with the lottery rules. LCAs must be filed for both new and continuing H-1B visas, but continuing visas are not capped and subject to the lottery. In order to focus on new applications only, we use the approach of Mahajan et. al. (2024). In particular, we select only LCAs filed in either March or April, with an employment start date that is in either August or September. Because of the timing of the lottery, and the rules over when a worker on an H-1B visa can start employment, this subset of LCAs have been shown by Mahajan et. al. (2024) to be the most likely applications associated with lottery H-1B visas. Column (5) of Table 6 reports the total number of LCAs filed each year that fall into this "lottery" category, and the last column shows the number of LCAs submitted by firms in our Compustat sample also belonging in this category. We

[14]We match the LCA data with firms listed in Compustat, employing the same matching procedure using firm names previously described.

[15]We limit our estimation to the 2015-2020 period for the following reasons. Prior to 2015, the H-1B cap was not binding in every year, and visa allocations were not consistently determined via lottery. In addition, starting with the 2020 lottery (for H-1B employment in fiscal year 2021), firms are no longer required to submit visa applications in advance, but instead to just register their interest in an H-1B visa. Under this regime, we no longer observe a consistent measure of firm-level applications prior to the lottery and therefore cannot construct our lottery-based instrument after 2020.

construct a firm-year measure of "expected visas" based on the firm's lottery applications' volume and the approval rate for new petitions in that year, and define the surprise number of visas as the residual difference between realized and expected approvals. This residual captures idiosyncratic variation in realized foreign employment that is not driven by firms' intended hiring or demand conditions.

Table 7: Data Summary Statistics

	N	**Mean**	**Median**	**Std Dev**	**5%**	**95%**
Output ($ Mil)	12,326	4,184.27	706.02	12,362.64	10.67	18,933.63
Capital ($ Mil)	12,326	7,843.44	538.73	31,484.11	4.70	36,330.45
Foreign Labor (000s)	12,326	0.15	0.01	1.66	0.00	0.29
Skilled Domestic Labor (000s)	12,326	4.46	0.62	14.88	0.01	19.35
Non-Skilled Domestic Labor (000s)	12,326	18.51	2.46	67.76	0.05	86.70
Domestic Skilled Wage ($000s)	4,029	137.43	118.28	83.98	69.77	249.86
Domestic Non-Skilled Wage ($000s)	4,029	117.52	78.77	106.04	16.53	352.29

All variables are observed at the firm-year level. Output and capital are from Compustat, foreign labor is imputed based on data from the USCIS, and skilled and non-skilled domestic labor are imputed from Compustat and OEWS data. The average wages for domestic workers are computed as specified in the paper.

Finally, to estimate the CES parameters of the production function, we require firm-level variation in domestic wages, which are not directly observed. We therefore construct firm-year measures of skilled and non-skilled domestic wages by combining data from BrightQuery, the BLS Occupational Employment and Wage Statistics (OEWS), and the USCIS H-1B wage records. We use two complementary approaches. Skilled domestic wages are imputed by combining industry-level wage information from OEWS with a public-firm premium and a firm-specific wage component, while non-skilled domestic wages are obtained as a residual by subtracting observed H-1B wage payments and the imputed skilled domestic wage bill from total firm payroll.

In particular, we impute the average skilled domestic wage $\bar{w}_{jt}^{d,\mathrm{s}}$ at firm j in year t as the product of three components. First, we take the average wage of workers in skilled occupations in the firm's industry from OEWS. Second, we scale this wage by the ratio of average payroll per worker at publicly traded firms to the average wage for all firms in the same industry. This information, obtained from BrightQuery, captures systematic wage differences between public and private firms.

Third, we apply a firm-specific skilled wage premium estimated from the USCIS H-1B wage data.[16].

We finally obtain the average non-skilled domestic wage $\bar{w}_{jt}^{d,\text{ns}}$ as follows. First, we need to recover the total wage bill for non-skilled workers $W_{jt}^{d,\text{ns}}$, which is given by:

$$W_{jt}^{d,\text{ns}} \;=\; W_{jt} \;-\; W_{jt}^{H1B} \;-\; \bar{w}_{jt}^{d,\text{s}}\, N_{jt}^{d,\text{s}}.$$

where W_{jt} is the total firm wage bill from Compustat, W_{jt}^{H1B} is the H-1B wage bill from our USCIS data, and $N_{jt}^{d,\text{s}}$ is the number of skilled domestic workers. We don't observe $N_{jt}^{d,\text{s}}$, and we define it as

$$N_{jt}^{d,s} \;=\; \left(\sum_{o\in\mathcal{S}} \text{Share}_{o,\text{ind}(j),t}^{\text{OEWS}}\right) N_{jt} - N_{jt}^{H1B},$$

where $\mathcal{S}$ is set of skilled occupations (i.e. the same that are eligible for H-1B employment), $\text{Share}_{o,\text{ind}(j),t}^{\text{OEWS}}$ is the OEWS employment share of occupation o in firm j's industry at time t, N_{jt} is total firm employment and N_{jt}^{H1B} is H-1B employment. To measure N_{jt}^{H1B} we use the USCIS's H-1B Employer Data Hub, which provides the number of approved H-1B petitions by firm and year. It is important to note that since this dataset does not track separations, we do not directly observe the number of H-1B worker at a given firm. We proceed with the assumption that each approved H-1B petition corresponds to three fiscal years of continuing employment. Consequently, N_{jt}^{H1B} is given by the sum of approved petitions in years t, $t-1$ and $t-2$[17]. We can finally apply the defintion of non-skilled domestic average wage

$$\bar{w}_{jt}^{d,\text{ns}} \;=\; \frac{W_{jt}^{d,\text{ns}}}{N_{jt}^{d,\text{ns}}},$$

with $N_{jt}^{d,\text{ns}}$ being calculated as

$$N_{jt}^{d,\text{ns}} = N_{jt} - N_{jt}^{d,\text{s}} - N_{jt}^{H1B}$$

[16]Using the H-1B wage data, we estimate a log wage equation for skilled workers of the form:

$$ln(w_{ifm}) = \beta X_i + \gamma_f + \delta_m + \epsilon_{ifm}$$

where w_{ifm} is the individual's wage, X_i is the worker's observed characteristics, γ_f is a firm fixed effect, and δ_m is an industry-year fixed effect. The firm fixed effects are interpreted as firm-specific skilled wage premia. To ensure comparability within markets, we normalize these premia by subtracting, in each industry-year, the employment-weighted average firm effect, so that the mean premium within an industry-year equals zero. We then exponentiate the normalized firm effects to obtain multiplicative wage premia that are then applied to the industry-level average skilled wages.

[17]For example, Facebook had 400, 587, and 651 approved petitions in fiscal years 2013, 2014 and 2015, resulting in a "stock" of 1,638 H-1B workers in fiscal year 2015.

The resulting firm-level skilled and non-skilled domestic wages are summarized in the final two rows of Table 7.

3 Model

Time is discrete. Firms produce output using capital and labor with a Cobb-Douglas technology. Labor differs by skill and nationality. Non-skilled positions can only be filled by domestic workers, while skilled positions can be filled by either domestic workers or foreign workers employed under the H-1B program. The firm's central decision is how many of its skilled positions to sponsor for H-1B visas. For each position, the firm compares the surplus from hiring a foreign worker to that from hiring a domestic worker. Positions are ranked by this surplus, and the firm applies for visas for all positions whose surplus exceeds the visa cost. Under the lottery system, which imposes an aggregate cap, realized foreign employment is a random truncation of this demand.

3.1 Technology

Firm j's output in period t is

$$Y_{jt} = A_{jt} K_{jt}^{\alpha_k} L_{jt}^{\alpha_\ell}, \tag{1}$$

where A_{jt} denotes productivity, K_{jt} is capital, and L_{jt} is an aggregate labor input. Labor is a nested CES aggregate. Aggregate labor combines skilled ($L_{s,jt}$) and non-skilled labor (R_{jt}):

$$L_{jt} = \left(\theta_s L_{s,jt}^{\frac{\sigma_s - 1}{\sigma_s}} + (1 - \theta_s) R_{jt}^{\frac{\sigma_s - 1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s - 1}}. \tag{2}$$

Skilled labor is a CES aggregate of foreign (F_{jt}) and domestic workers (D_{jt}):

$$L_{s,jt} = \left(\theta_f F_{jt}^{\frac{\sigma_f - 1}{\sigma_f}} + (1 - \theta_f) D_{jt}^{\frac{\sigma_f - 1}{\sigma_f}} \right)^{\frac{\sigma_f}{\sigma_f - 1}}. \tag{3}$$

This structure allows for imperfect substitution between skill groups and between foreign and domestic skilled labor. Marginal products are defined as the partial derivatives of output with respect to the relevant labor input. In particular, $MPL_{jt}^F \equiv \frac{\partial Y_{jt}}{\partial F_{jt}}$ and $MPL_{jt}^D \equiv \frac{\partial Y_{jt}}{\partial D_{jt}}$, evaluated at the firm's

current input levels.

3.2 Wages

Wages are determined through Nash bargaining between firms and workers. For a foreign worker i employed by firm j in period t, the bargained wage is

$$w_{ijt}^{F} = \eta \, MPL_{jt}^{F} + (1 - \eta) \, \bar{w}_{oec}, \tag{4}$$

where $\eta \in (0, 1)$ denotes the firm's bargaining weight and $\bar{w}_{oec}$ is the worker's outside option. The outside option depends on observable worker characteristics: occupation o, education level e, and country of origin c. These characteristics are observed by the firm at the time of hiring and affect wages through workers' outside options.

We do not model the wage-setting process for domestic skilled workers. Domestic wages enter the model as firm-year wage measures constructed as described in Section 2.1, and serve as empirical benchmarks for the opportunity cost of filling a skilled position domestically. They are treated as measured inputs rather than objects of structural estimation.

3.3 Hiring Costs

Hiring a foreign worker entails two additional costs borne by the firm. First, the firm pays a monetary visa cost c_v per new foreign hire. Second, firms incur a firm-year-specific wedge ψ_{jt}. The model evaluates hiring decisions by comparing, for each skilled position, the marginal surplus from employing a foreign worker to that from employing a domestic worker in that period, net of the visa fee.

We interpret ψ_{jt} as a per-position wedge (common across a firm's openings within year t) that shifts the firm's willingness-to-pay schedule for foreign labor. This term reflects administrative and recruiting frictions, differences in labor market tightness, and other systematic factors not captured by observed wages or marginal products. Because the model abstracts from multi-year match dynamics, ψ_{jt} should be interpreted as a reduced-form firm-year wedge that rationalizes observed foreign hiring under this static surplus comparison. In the estimation, ψ_{jt} is recovered from firms' revealed hiring cutoffs.

3.4 Willingness to Pay for a Visa

Consider a marginal skilled position at firm j in period t. The firm compares hiring a foreign worker to hiring a domestic worker. The net incremental surplus from hiring a foreign worker is

$$WTP_{jtp} = \left(MPL^F_{jt} - MPL^D_{jt}\right) - \left(w^F_{jtp} - w^D_{jt}\right) - \psi_{jt}. \tag{5}$$

This expression defines the firm's willingness to pay for an H-1B visa for position p. As foreign employment rises, diminishing returns imply that WTP_{jtp} is weakly decreasing across positions, generating a downward-sloping firm-level demand curve for visas.

Let N_{jt} denote the number of skilled positions open to new hires. For each position, the firm draws a potential foreign worker from a firm-specific distribution over observable characteristics. The firm applies for visas for all positions satisfying

$$WTP_{jtp} \geq c_v.$$

Visa demand at firm j is therefore

$$H_{jt}(c_v) = \max\{p : WTP_{jtp} \geq c_v\}.$$

Under the current lottery system, realized foreign employment is constrained by the aggregate visa cap. Under alternative allocation mechanisms that price visas, such as auctions, visas are allocated to positions with the highest willingness to pay.

4 Estimation

We estimate the structural parameter vector

$$\Theta = \left(\underbrace{\alpha_0, \alpha_k, \alpha_l}_{\text{Production}}, \underbrace{\sigma_s, \sigma_f, \theta_s, \theta_f}_{\text{Labor substitution}}, \underbrace{\eta, \{w^o_o, w^o_e, w^o_c\}}_{\text{Wage determination}}, \underbrace{\{\psi_{jt}\}}_{\text{Hiring wedge}} \right).$$

which governs production, labor substitution, wage determination, and firm-year hiring wedges.

Identification relies critically on quasi-random variation generated by the H-1B lottery. Conditional on the number of cap-subject applications submitted by a firm, the realized number of visas awarded differs from its expectation due to the lottery outcome. We use these "surprise" visa allocations as exogenous shifts in foreign employment. This variation is central for identifying substitution elasticities and the marginal product of foreign labor in the production function.

The estimation proceeds in four sequential stages. First, we estimate the parameters of the nested CES labor aggregator governing substitution between skilled and non-skilled labor and between foreign and domestic skilled labor. These parameters are identified from firms' optimal labor demand conditions, exploiting exogenous shifts in the composition of skilled labor induced by lottery shocks. Second, given the estimated labor aggregator, we estimate the Cobb–Douglas production function using a proxy-variable approach to address simultaneity in input choices. Foreign employment is instrumented using the firm's surprise visa allocation, while capital and lagged domestic labor serve as additional instruments under standard timing assumptions. Third, we estimate the Nash-bargaining wage equation using worker-level administrative data. Variation in firm-level marginal products identifies the bargaining parameter, while cross-occupation, education, and country-of-origin wage differences identify worker outside options. Finally, we recover firm-year search costs ψ_{jt} using a revealed preference approach. It is set so that for all positions p for which a firm j desires to hire a H-1B visa, $WTP_{jtp} - c_v \geq 0$, and for positions p' for which they prefer a domestic worker, $WTP_{jtp'} - c_v < 0$. For each firm-year, the value of ψ_{jt} is chosen so that the firm is indifferent between hiring the last worker they are observed to submit an H-1B application for (e.g. $WTP_{jtp} = c_v$ for this marginal worker).

Together, these steps allow us to construct each firm-year's inverse demand curve for H-1B visas. For any candidate visa price p, we compute the number of foreign workers the firm would optimally hire,

$$H_{jt}(p) = \max\{m : \text{WTP}_{jtm} \geq p\},$$

where WTP_{jtm} denotes the marginal profit gain from the m-th foreign hire. Aggregating across firms yields market-level visa demand, which we use to evaluate allocation mechanisms.

4.1 Estimating the CES Labor Aggregator

In the first stage, we estimate the parameters governing the nested CES labor aggregator,

$$L_{jt} = \left(\theta_s L_{S,jt}^{\frac{\sigma_s - 1}{\sigma_s}} + (1 - \theta_s) R_{jt}^{\frac{\sigma_s - 1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s - 1}}, \tag{6}$$

$$L_{S,jt} = \left(\theta_f F_{jt}^{\frac{\sigma_f - 1}{\sigma_f}} + (1 - \theta_f) D_{jt}^{\frac{\sigma_f - 1}{\sigma_f}} \right)^{\frac{\sigma_f}{\sigma_f - 1}}, \tag{7}$$

where F_{jt} denotes foreign skilled labor, D_{jt} domestic skilled labor, and R_{jt} domestic non-skilled labor. The parameters (σ_s, σ_f) determine substitution elasticities across labor types, while (θ_s, θ_f) govern their relative weights.

Identification exploits firms' optimal labor demand conditions. For domestic labor types, profit maximization implies

$$p_{jt} \frac{\partial Y_{jt}}{\partial D_{jt}} = ME_{jt}^{D}, \qquad p_{jt} \frac{\partial Y_{jt}}{\partial R_{jt}} = ME_{jt}^{R}, \tag{8}$$

where ME_{jt}^{q} denotes the marginal expense of hiring labor type q. Under Nash bargaining, marginal expenses are proportional to wages, so that $ME_{jt}^{q} = \nu w_{jt}^{q}$. Taking ratios across domestic labor types eliminates the proportionality term ν and yields

$$\frac{\partial Y_{jt} / \partial D_{jt}}{\partial Y_{jt} / \partial R_{jt}} = \frac{w_{jt}^{d}}{w_{jt}^{r}}. \tag{9}$$

Using the nested CES structure, this ratio can be expressed as

$$\ln\left(\frac{w_{jt}^{d}}{w_{jt}^{r}} \right) = \ln\left(\frac{\theta_s (1 - \theta_f)}{1 - \theta_s} \right) + \left(\frac{1}{\sigma_f} - \frac{1}{\sigma_s} \right) \ln(L_{S,jt}) - \frac{1}{\sigma_f} \ln(D_{jt}) + \frac{1}{\sigma_s} \ln(R_{jt}) + \epsilon_{jt}^{CES}. \tag{10}$$

This equation forms the basis of estimation. The residual ϵ_{jt}^{CES} captures measurement error in wages and employment. A central challenge is that F_{jt} and D_{jt} are chosen by firms in response to productivity shocks, rendering the skilled labor composite $L_{S,jt}$ endogenous. To address this, we exploit quasi-random variation generated by the H-1B lottery. Conditional on applications, the realized number of visas differs from its expectation due to the lottery outcome. We construct the

firm-year "surprise" visa allocation,

$$S_{jt} = \text{Realized visas}_{jt} - \text{Expected visas}_{jt},$$

and use S_{jt} as an instrument for foreign employment.

Because S_{jt} is determined by the lottery after applications are submitted and is not a function of contemporaneous productivity shocks, it provides plausibly exogenous variation in F_{jt} and therefore in $L_{S,jt}$. Appendix Table 16 reports first-stage regressions. The surprise-visa measure is strongly correlated with realized foreign employment, with large and precisely estimated coefficients, supporting its relevance as an instrument.[18] We combine S_{jt} with lagged domestic employment $(D_{j,t-1}, R_{j,t-1})$ to form the instrument set

$$Z_{jt}^{CES} = \{S_{jt}, D_{j,t-1}, R_{j,t-1}\}.$$

The parameters $(\sigma_s, \sigma_f, \theta_s, \theta_f)$ are then estimated by GMM using the orthogonality condition

$$E\big[Z_{jt}^{CES}\, \epsilon_{jt}^{CES}(\sigma_s, \sigma_f, \theta_s, \theta_f)\big] = 0 \tag{11}$$

Variation in non-skilled employment relative to skilled employment identifies σ_s, while exogenous shifts in foreign employment induced by S_{jt} identify σ_f. The share parameter θ_s is primarily identified by the average level of the wage ratio, whereas θ_f enters through the construction of the skilled labor composite and is therefore more weakly identified.

Table 8 reports the GMM estimates of the nested CES parameters. The estimated elasticity of substitution between skilled and non-skilled labor is $\hat{\sigma}_s = 1.75$ (s.e. 0.25), indicating that the two labor aggregates are imperfect substitutes. This estimated value is in the range typically found for substitution across broad education or skill groups in U.S. labor market estimates (e.g., Katz and Murphy (1992); Card and Lemieux (2001); Ottaviano and Peri (2012)).

The estimated elasticity of substitution between foreign and domestic skilled labor is substantially larger than the cross-skill elasticity, with $\hat{\sigma}_f \approx 7$ (s.e. 5.55), although it is estimated less precisely. This value is lower than the native-immigrant substitution elasticities obtained in na-

[18]The results in Table 16 also show the importance of using only cap-window LCAs in constructing the instrument, as including all LCAs will lead to weak correlation with won visas and a negative relationship between the subsequently constructed surprise number of visas and foreign labor.

tional education-experience cell frameworks (e.g. Ottaviano and Peri (2012)). Two features of our setting naturally push toward a smaller elasticity. First, we do not further partition skilled labor by experience, and secondly, substitution is identified from within-firm production choices rather than from economy-wide variation.

The share parameter $\hat{\theta}_s = 0.44$ implies that skilled labor accounts for a substantial share of effective labor input. The within-skilled share parameter $\hat{\theta}_f$ is estimated imprecisely, reflecting the limited independent variation available once substitution curvature is accounted for. Due to their weak identification, we assess the robustness of our main results to alternative values of θ_f and σ_f in Appendix B.

The substitution parameters play a central role in shaping firm-level demand for visas. When σ_f is high, foreign and domestic skilled workers are close substitutes, so marginal products converge rapidly as foreign employment rises. In contrast, lower substitution would generate larger productivity premia for foreign workers and steeper firm-level willingness-to-pay schedules. The estimated values imply meaningful but diminishing gains from foreign hiring, generating downward-sloping firm-level visa demand curves that are crucial for the auction counterfactual.

We assess the sensitivity of our counterfactual allocations over a grid of (σ_f, θ_f) values in Appendix B. Specifically, we consider $\sigma_f \in \{3, 7.18, 20\}$ and $\theta_f \in \{0.25, 0.39, 0.50, 0.75\}$. Across these specifications, the qualitative direction of reallocation is stable and the set of major recipient firms in our counterfactual second price auction remains largely unchanged. At the same time, equilibrium outcomes are quantitatively sensitive[19]. This sensitivity highlights that within-skilled substitution is economically consequential for firm-level willingness to pay and motivates estimating these elasticities rather than imposing them a priori.

[19]For illustration, Cognizant receives 17,240 visas under the lottery. Under the baseline parameter estimates, its auction allocation falls to 762 visas. Across the full (σ_f, θ_f) grid in Appendix B, Cognizant's allocation ranges from 237 to 931 visas, corresponding to only 1.4% to 5.4% of its lottery allocation. In contrast, Meta receives 6,523 visas under the lottery and 13,909 under the baseline auction. Across the same parameter grid, its allocation ranges from 11,321 to 23,765 visas, or 174% to 364% of its lottery allocation. Thus, while substitution curvature affects magnitudes, the qualitative ranking of major winners and losers is invariant.

Table 8: Nested CES Labor Aggregator Estimates

	Estimate	Standard Error
σ_s	1.7511	0.2533
σ_f	7.1846	5.5524
θ_s	0.4397	0.3753
θ_f	0.3851	14.2442

GMM estimates of the nested CES labor aggregator parameters. Estimation uses the orthogonality conditions in equation (11) and exploits lottery-induced variation in foreign employment through the surprise visa instrument S_{jt}. Standard errors are clustered at the firm level.

4.2 Estimating the Production Function

In the second stage, we estimate the parameters of the Cobb–Douglas production function,

$$Y_{jt} = A_{jt} K_{jt}^{\alpha_k} L_{jt}^{\alpha_l}, \tag{12}$$

where K_{jt} denotes capital and L_{jt} is the aggregate labor index constructed using the estimated CES parameters from the previous stage. In practice, we rescale L_{jt} within industry groups so that its mean matches mean observed employment. Taking logs yields:

$$y_{jt} = \alpha_0 + \alpha_k k_{jt} + \alpha_l \ell_{jt} + u_{jt}, \tag{13}$$

where $u_{jt} = \log A_{jt}$ represents firm-year productivity.

A central identification challenge is that labor input is chosen by firms in response to productivity shocks. More productive firms optimally hire more workers, implying that ℓ_{jt} is correlated with u_{jt}. Estimating the production function by OLS would therefore bias α_l upward.

We address this simultaneity problem using a proxy-variable approach following Olley and Pakes (1996), Levinsohn and Petrin (2003), and Ackerberg, Caves, and Frazer (2015). We assume that capital is predetermined at time t and therefore uncorrelated with contemporaneous productivity innovations. Domestic labor inputs are assumed to respond to productivity with a lag, while foreign employment is partially determined by the H-1B lottery. Thus, we can again use the surprise number of visas as exogenous variation in foreign employment, which through the CES aggregator varies the total labor input.

In a first stage, we estimate:

$$y_{jt} = \phi_t(k_{jt}, l_{jt}) + \upsilon_{jt} \tag{14}$$

where υ_{jt} are unanticipated shocks to productivity. We assume that firms do not observe υ_{jt} when making input decisions. We obtain nonparametric estimates of ϕ_{jt} by running a third-order polynomial regression of y_{jt} on the inputs and their interactions. We also add year fixed effects to control for time trends. This gives us a measurement-error-free estimate of log output.

We use the first stage estimates to compute an estimate of productivity for any guess at the production function parameters α:

$$u_{jt}(\alpha) = \widehat{\phi_{jt}} - \alpha_0 - \alpha_k k_{jt} - \alpha_l l_{jt} \tag{15}$$

We then assume productivity follows a first-order Markov process,

$$u_{jt} = g(u_{j,t-1}) + \xi_{jt}, \tag{16}$$

where ξ_{jt} is an innovation uncorrelated with predetermined inputs and lottery shocks. We regress $u_{jt}(\alpha)$ on its lag to obtain estimates for the productivity shocks $\xi_{jt}(\alpha)$. The parameters $(\alpha_0, \alpha_k, \alpha_l)$ are then estimated by GMM using the moment conditions

$$E[\xi_{jt}(\alpha)k_{jt}] = 0, \qquad E[\xi_{jt}(\alpha)D_{j,t-1}] = 0, \qquad E[\xi_{jt}(\alpha)R_{j,t-1}] = 0, \qquad E[\xi_{jt}(\alpha)S_{jt}] = 0. \tag{17}$$

where $d_{jt} = log(D_{jt})$, $r_{jt} = log(R_{jt})$, and $s_{jt} = log(S_{jt})$. The first three moments are similar to Ackerberg, Caves, and Frazier (2015). The idea is that capital is decided in the prior period and so is not related to the productivity shock. The two domestic labor decisions are instead expected to be made in response to the productivity shocks, and so lagged skilled and lagged unskilled labor is used for identification. The final moment used to identify the parameter on aggregate labor is new and uses the randomness of the H-1B lottery to generate variation in the number of foreign workers at a firm each year.

Firms in our sample operate in a wide range of industries and rely on fundamentally different production technologies. As a result, it would be implausible to assume a common production

function across all firms. To account for this heterogeneity, we allow the output elasticities of capital and labor to vary across industries. Industries are defined using firms' two-digit NAICS codes, which we aggregate into a set of broad industry groups reflecting similar production processes. We then estimate different values of α_0, α_k, and α_l for each industry. This approach allows factor elasticities to differ flexibly across sectors while maintaining a common functional form and identification strategy within each industry.

Table 9 reports the estimated output elasticities by industry group. Across sectors, the labor elasticity $\hat{\alpha}_\ell$ ranges from 0.46 (real estate) to 0.88 (intermediate goods manufacturing), while the capital elasticity $\hat{\alpha}_k$ ranges between 0.14 and 0.34.These magnitudes fall within the range commonly reported in firm-level production function studies using proxy-variable approaches (see Olley and Pakes (1996), Levinsohn and Petrin (2003), Ackerberg, Caves, and Frazier (2015), De Loecker and Warzynski (2012)).

Combining these elasticities with observed inputs and the estimated CES labor aggregator yields firm-year marginal products of foreign and domestic labor. The distributions of these marginal products of labor can be visualized in Figure 11 in the Appendix. The figure shows that when evaluated at observed input levels, the median (across firm-years) marginal product of labor of H-1B workers is above that for skilled domestic workers, which is greater than the median marginal product of labor for non-skilled domestic workers. Because marginal products scale with $\alpha_\ell Y_{jt}/L_{jt}$, cross-firm heterogeneity in output levels and capital intensity translates directly into dispersion in the value created by an additional skilled worker. This dispersion is central for the counterfactual analysis as differences in marginal products generate heterogeneity in firms' willingness to pay for H-1B visas, which in turn drives reallocation under the auction mechanism.

4.3 Estimating the Nash Bargaining Wage Equation

In the third stage, we estimate the parameters governing wage determination under Nash bargaining. Recall that foreign worker wages are determined by

$$w^f_{jpt} = \eta\, MPL^f_{jt} + (1-\eta) w^o_{oec}, \tag{18}$$

Table 9: Cobb–Douglas Production Function Estimates by Industry Group

Industry Group	$\hat{\alpha}_0$	$\hat{\alpha}_k$	$\hat{\alpha}_\ell$	Obs.
Energy, Utilities, and Construction	-1.2502	0.3355	0.6402	704
	(0.0908)	(0.0100)	(0.0133)	
Consumer Manufacturing	-1.2067	0.3076	0.7043	419
	(0.1178)	(0.0297)	(0.0308)	
Intermediate Goods Manufacturing	-1.6957	0.1842	0.8835	1,713
	(0.0853)	(0.0197)	(0.0251)	
Industrial Manufacturing	-0.9841	0.1475	0.8126	3,532
	(0.0462)	(0.0115)	(0.0130)	
Trade and Transportation	-0.0775	0.2448	0.6174	1,237
	(0.0677)	(0.0109)	(0.0115)	
Information	-0.6587	0.1682	0.8058	2,020
	(0.0500)	(0.0074)	(0.0104)	
Finance and Insurance	0.1059	0.2844	0.6234	543
	(0.1383)	(0.0255)	(0.0308)	
Real Estate	1.0390	0.2949	0.4634	251
	(0.1867)	(0.0237)	(0.0255)	
Professional Services	-0.7199	0.1601	0.7873	526
	(0.0675)	(0.0143)	(0.0147)	
Management	-0.6718	0.2412	0.6801	253
	(0.1379)	(0.0186)	(0.0202)	
Education and Health	-1.3484	0.1356	0.8116	297
	(0.1528)	(0.0313)	(0.0348)	
Other Services	-1.0682	0.3208	0.6094	335
	(0.1531)	(0.0240)	(0.0254)	

Estimates of the Cobb–Douglas production function by industry group. Standard errors in parentheses. Estimation uses the proxy-variable approach described in Section 4.2 and exploits lottery-induced variation in foreign employment through the surprise visa instrument S_{jt}.

where MPL^f_{jt} is the marginal product of foreign labor, $\eta \in [0,1]$ is the firm's bargaining weight, and w^o_{oec} denotes the worker's outside option, which depends on occupation (o), education (e), and country of origin (c). We assume that outside options are additively separable,

$$w^o_{oec} = w^o_o + w^o_e + w^o_c, \tag{19}$$

and that the bargaining parameter η is constant across firms.[20]

We augment the wage equation with an idiosyncratic component ε_{jh} capturing match-specific deviations from predicted outside options. We impose the orthogonality conditions

$$E[\varepsilon_{jh} \mid o_h] = 0, \quad E[\varepsilon_{jh} \mid e_h] = 0, \quad E[\varepsilon_{jh} \mid c_h] = 0, \tag{20}$$

using worker-level wage data from USCIS.

Identification of η relies on variation in marginal products across firms and over time. Given estimates of MPL^f_{jt} from the production stage, cross-firm variation in productivity generates variation in the surplus available from each match. The extent to which wages covary with MPL^f_{jt} identifies the surplus-sharing parameter η. If wages closely track marginal products, η is high; if wages are closer to outside options, η is low.

Identification of the outside option components (w^o_o, w^o_e, w^o_c) relies on systematic wage differences across occupations, education levels, and countries of origin after controlling for firm-level marginal products. Because outside options are only identified up to a normalization, we set the bachelor's degree effect and the Canada country effect to zero.

Table 10: Nash Bargaining Parameter Estimate

	Estimate	Standard Error
η	0.1358	0.0284

Estimate of the surplus-sharing parameter from the Nash bargaining wage equation. Standard errors are clustered at the firm level.

Table 10 reports the estimated surplus-sharing parameter. We obtain $\hat{\eta} = 0.136$ (s.e. 0.028),

[20] We explored allowing η to vary by industry; however, once production parameters are estimated separately by industry, the additional variation was insufficient to precisely identify industry-specific bargaining weights.

implying that wages place substantially more weight on workers' outside options than on firm-specific marginal products. In other words, firms capture the majority of match surplus at the time of hiring. Because $\hat{\eta}$ is well below one-half, variation in marginal products translates primarily into variation in firm surplus rather than wage payments. This is important for the counterfactual analysis as it implies that differences in productivity across firms generate differences in willingness to pay for visas even when wages adjust endogenously. If η were close to one, most productivity gains would accrue to workers, substantially flattening firm-level visa demand curves.

Full estimates for the outside option parameters are reported in Appendix A. Higher outside options are estimated for executives and healthcare practitioners, while lower outside options are estimated for life scientists and media occupations. Education premia are estimated to be modest for a master's degree but much stronger for a doctoral or professional degree compared to a bachelor's degree. The country of origin results show that workers from higher income countries are estimated to have higher outside options than workers coming from lower income countries.

4.4 Recovering Firm-Year Search Costs

In the final stage, we recover firm-year search costs ψ_{jt}, which capture wedges between hiring foreign and domestic workers that are not explained by productivity or wages. These costs reflect administrative burdens, recruitment frictions, and any residual components affecting the relative attractiveness of foreign labor.

In the model, a firm hires a foreign worker in position p if

$$V^f_{jpt} > V^d_{jpt}, \tag{21}$$

where

$$V^f_{jpt} = MPL^f_{jt} - w^f_{jpt} - \psi_{jt} - c_v, \qquad V^d_{jpt} = MPL^d_{jt} - w^d_{jpt}. \tag{22}$$

At the margin, the firm must be indifferent between hiring the last foreign worker and hiring a comparable domestic worker. Denoting the marginal foreign hire by h', the indifference condition implies:

$$\psi_{jt} = \left(MPL^f_{jt} - MPL^d_{jt}\right) - \left(w^f_{h'} - w^d_{ont}\right) - c_v. \tag{23}$$

Because ψ_{jt} enters the model through a discrete hiring threshold, it is recovered using a simulation-based procedure. For each firm-year, we simulate potential foreign hires by drawing worker characteristics from the firm-specific empirical distribution and computing surplus differentials using the estimated production and wage primitives. For a candidate value of ψ_{jt}, the firm submits lottery applications for positions where the estimated surplus differential is positive, yielding implied visa demand $D_{jt}(\psi)$. Given the year-t lottery success rate, the number of applications submitted implies an expected number of visas received for that firm-year. The value of ψ_{jt} is chosen so that the model-implied expected visas received matches the observed number of initial approvals for that firm-year.

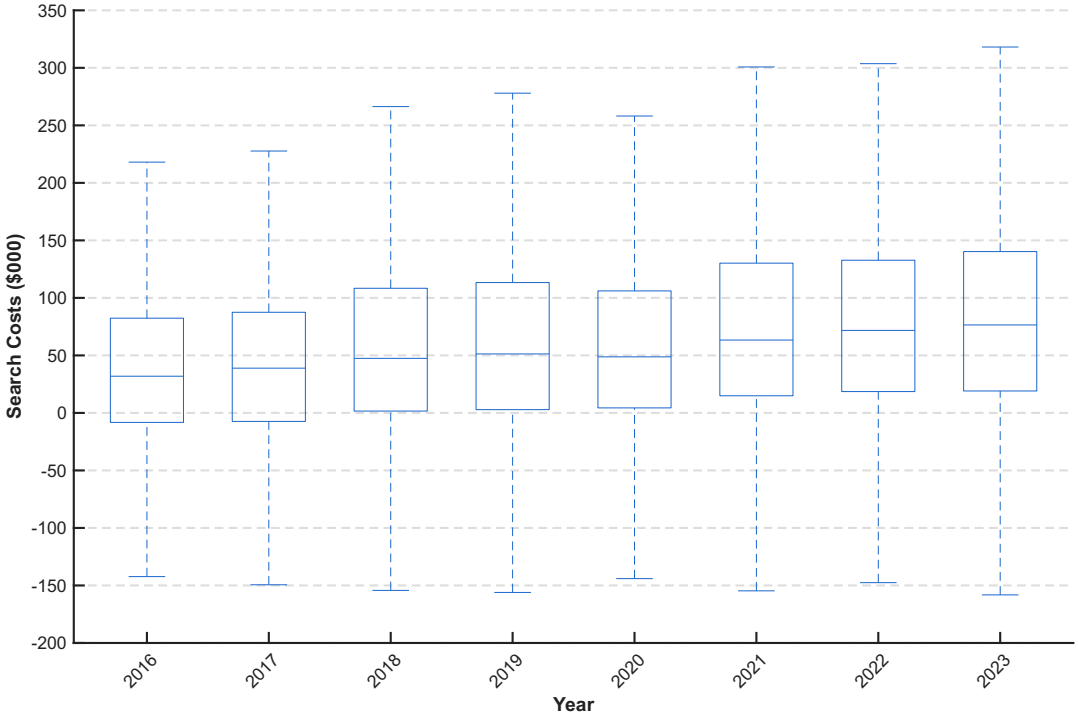


Figure 1: Distribution of Estimated Search Costs by Year

Figure 1 displays the distribution of estimated firm-year search costs by year. The median $\widehat{\psi}_{jt}$ is positive throughout the sample period and increases over time, indicating that, on average, firms face additional costs when hiring foreign workers beyond the monetary visa fee c_v. At the same time, there is substantial cross-firm heterogeneity in the estimated wedge.

Consistent with the interpretation in Section 3.3, ψ_{jt} captures hiring frictions beyond measured productivity and wages. Because ψ_{jt} shifts the threshold at which foreign hiring becomes profitable,

it directly affects the level and slope of firm-level willingness-to-pay schedules. Firms with lower search wedges demand more visas at a given price, while firms with higher wedges require larger productivity premia to justify foreign hiring. This heterogeneity is central to the reallocation patterns observed under the auction counterfactual.[21]

4.5 Implied Firm-Level Visa Demand

Using the estimated parameters, we construct each firm-year's inverse demand curve for H-1B visas. For a given skilled opening, the firm compares the incremental surplus from hiring a foreign worker to that from hiring a comparable domestic worker. The resulting marginal surplus defines the firm's willingness to pay for a visa. Ordering openings by this willingness to pay yields a weakly downward-sloping firm-level demand schedule.

We define the number of skilled openings in year t as N_{jt}, equal to the number of skilled positions filled through new hiring decisions in that year. In our static representation, this corresponds to the number of observed skilled domestic workers (imputed as described in Section 2) plus the number of newly granted H-1B visas. Positions already filled by continuing H-1B workers are excluded. Thus, N_{jt} captures the within-period hiring margin over which the firm chooses between domestic and foreign labor, while continuing H-1B workers are treated as a predetermined state variable.

Because visa allocations are subject to a lottery, the observed number of new H-1B hires reflects realized lottery outcomes rather than unconstrained demand. To recover the firm's latent visa demand, we begin from a counterfactual allocation in which all N_{jt} new skilled positions are initially assigned to domestic workers, holding fixed the stock of continuing H-1B employees. We then allow the firm to convert positions to foreign hires in descending order of willingness to pay. This reconstruction isolates the domestic-versus-foreign composition decision at a fixed level of total skilled employment and recovers the full schedule of willingness to pay that rationalizes the observed allocation under the lottery.

For each opening, a potential H-1B candidate with characteristics (o, e, c) is drawn from the firm-specific empirical distribution $\widehat{F}_j^{oec}$, estimated from the USCIS wage data. For each firm-year,

[21] Negative values of $\widehat{\psi}_{jt}$ are consistent with the model's hiring condition. They arise when, absent search frictions, the net surplus from hiring a domestic worker exceeds that from hiring a foreign worker. Intuitively, they suggest that the search costs to finding a domestic worker are higher than that of a foreign worker on a H-1B visa, implying that the firm is finding it difficult to find a domestic worker to fill the position.

we take N_{jt} random draws from this distribution to generate the set of potential foreign hires.

Let $(\tilde{D}_{jt}, \tilde{F}_{jt})$ denote the current number of skilled domestic and foreign workers. We initialize $\tilde{F}_{jt}$ at the stock of continuing H-1B workers and $\tilde{D}_{jt}$ at N_{jt}, reflecting the all-domestic benchmark for new hires. Using the estimated production function, we compute marginal products at these employment levels:

$$\widehat{MPL^f_{jt}}(\tilde{D}_{jt}, \tilde{F}_{jt}) = \widehat{Y}(K_{jt}, \tilde{D}_{jt} - 1, \tilde{F}_{jt} + 1, R_{jt}, \widehat{u}_{jt}) - \widehat{Y}(K_{jt}, \tilde{D}_{jt} - 1, \tilde{F}_{jt}, R_{jt}, \widehat{u}_{jt}), \tag{24}$$

$$\widehat{MPL^d_{jt}}(\tilde{D}_{jt}, \tilde{F}_{jt}) = \widehat{Y}(K_{jt}, \tilde{D}_{jt}, \tilde{F}_{jt}, R_{jt}, \widehat{u}_{jt}) - \widehat{Y}(K_{jt}, \tilde{D}_{jt} - 1, \tilde{F}_{jt}, R_{jt}, \widehat{u}_{jt}), \tag{25}$$

where $\widehat{u}_{jt}$ is the estimated productivity shock from equation (13).

Foreign wages are determined using the estimated Nash bargaining equation:

$$\widehat{w^f_{jtoecp}} = \widehat{\eta}\,\widehat{MPL^f_{jt}}(\tilde{D}_{jt}, \tilde{F}_{jt}) + (1 - \widehat{\eta})\left(\widehat{w^o_{op}} + \widehat{w^o_{ep}} + \widehat{w^o_{cp}}\right) + \widehat{\epsilon}_{jtp}, \tag{26}$$

where $\widehat{\epsilon}_{jtp}$ is drawn from the firm-specific residual wage distribution. Comparable domestic wages are denoted $\widehat{w^d}_{on_{jt}}$. They are taken directly from observed BLS occupation–industry–year wage data.

The firm's willingness to pay for a visa for position p is then

$$WTP_{jtp} = \left(\widehat{MPL^f_{jt}}(\tilde{D}_{jt}, \tilde{F}_{jt}) - \widehat{MPL^d_{jt}}(\tilde{D}_{jt}, \tilde{F}_{jt})\right) - \left(\widehat{w^f_{jtoecp}} - \widehat{w^d_{on_{jt}}}\right) - \widehat{\psi}_{jt}. \tag{27}$$

We compute WTP_{jtp} for all N_{jt} openings at the initial employment levels $(\tilde{D}_{jt}, \tilde{F}_{jt})$. The opening with the highest willingness to pay is assigned to foreign employment, employment is updated to $(\tilde{D}_{jt} - 1, \tilde{F}_{jt} + 1)$, and the procedure is repeated. Iteration continues until no remaining opening yields positive incremental surplus. Because marginal products decline as foreign employment rises, the resulting sequence $\{WTP_{jtp}\}_{p=1}^{N_{jt}}$ is weakly decreasing and defines a well-behaved inverse demand curve.

It is important to note that this measure of willingness to pay reflects the surplus generated in the initial period of employment. Since H-1B workers who remain with the firm may generate additional future surplus beyond the first year, the constructed WTP_{jtp} should be interpreted as a

lower bound on the firm's total valuation of the visa.[22]

For any visa price c_v, firm-level demand is

$$H_{jt}(c_v) = \max\{p \leq N_{jt} : WTP_{jtp} \geq c_v\}. \tag{28}$$

Aggregating $H_{jt}(c_v)$ across firms yields market-level visa demand.

4.6 Model Fit and Validation

To assess the model's fit, we simulate the H-1B lottery using the estimated firm-level demand schedules and compare the resulting allocation to the observed distribution of visas across firms. Because our sample is restricted to publicly traded firms, we fix the total number of visas in each year to the number observed in our data, V_t, ensuring a consistent basis for comparison.

For each firm-year, all positions p with $WTP_{jtp} \geq c_v = \$6{,}470$ are entered into the lottery. If aggregate demand exceeds V_t, visas are randomly allocated among applicants. We implement this procedure for each year from 2016 to 2023.

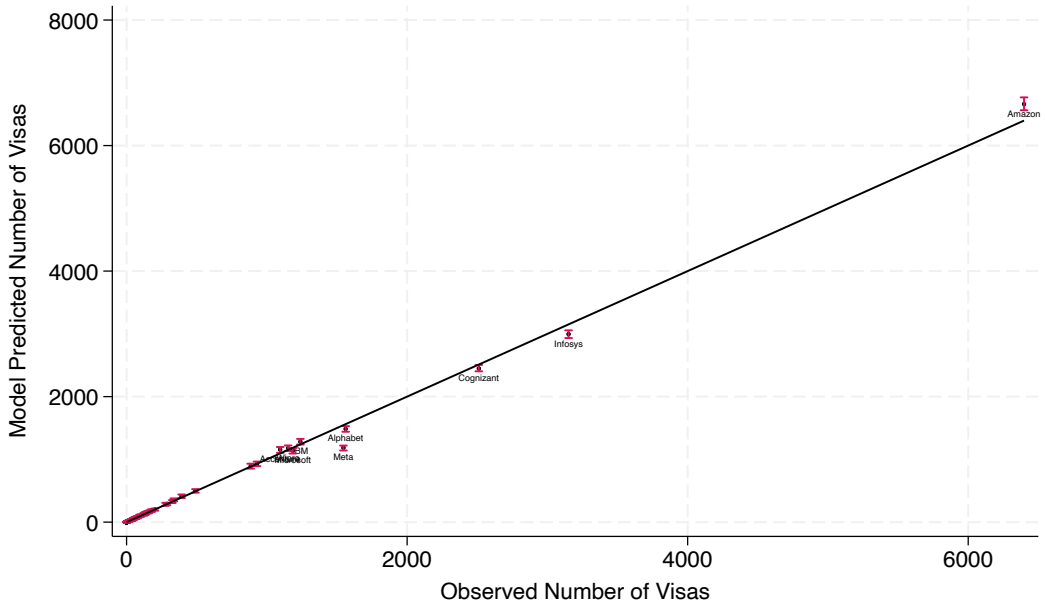
The simulation incorporates two sources of randomness: draws from the estimated wage residual distribution and the lottery itself. We run 100 simulations per year.

Figure 2 illustrates model fit for fiscal year 2023. Each point represents a firm; the horizontal axis reports the observed number of visas received, while the vertical axis reports the mean allocation across simulated lotteries. The 45-degree line corresponds to perfect prediction. The tight clustering of points around this line indicates that the model closely reproduces the cross-firm distribution of visa allocations. Year-by-year plots for 2016-2022, reported in Appendix A.4, show similarly strong fit.

Table 11 summarizes model performance using the firm-level allocation match rate, which is the share of realized visas whose recipient firm matches the simulated lottery outcome. Across years, the model reproduces between 94.6% and 96.6% of firm-level visa allocations, with an average match rate of 95.6%.

[22] After the first year, wages may be renegotiated and workers may switch employers, as a new employer's petition is not subject to the lottery once the worker has already been counted against the cap. Our static surplus calculation therefore captures only the initial-period gain from hiring a foreign worker.

Figure 2: Model Fit in 2023: observed versus simulated H-1B allocations.



Each point represents a firm. The horizontal axis reports the number of visas received in the data, and the vertical axis reports the mean number received across 100 simulated lotteries. The 45-degree line denotes perfect prediction. Vertical bars show 90% simulation intervals reflecting wage and lottery randomness. Year-by-year fit plots for 2016-2022 are reported in Appendix A.4.

Table 11: Summary of Model Fit: Firm-Level Allocation Match Rate

Statistic	Value
Mean match rate (2016–2023)	95.6%
Minimum across years	94.6%
Maximum across years	96.6%

The match rate is computed as the fraction of realized visas in a given year that are assigned to the same firm in the simulated lottery as in the observed data. The table reports the mean match rate across 2016–2023, along with the minimum and maximum across years, based on 100 simulations per year.

5 Counterfactual Second-Price Auction

We first characterize equilibrium under a second-price auction and compare it to the current lottery[23]. We then explain the resulting reallocation through heterogeneity in firm-level demand

[23]We assume firms bid their model-implied willingness to pay. In multi-unit auctions, however, bidders may have incentives to strategically bid below their true evaluation. Our exercise therefore evaluates a truthful-revelation benchmark.

elasticities and decompose the sources of willingness to pay. Finally, we consider policy extensions and a higher fixed visa fee.

5.1 Equilibrium Under a Second-Price Auction

Firms submit bids equal to their willingness to pay for each potential H-1B hire. Visas are awarded to the V_t highest bids, and the equilibrium visa price equals the $(V_t + 1)$-th highest bid. Let $H_{jt}(c_v)$ denote firm j's demand for visas at price c_v. The auction price c_v^A solves:

$$\sum_j H_{jt}(c_v^A) = V_t. \tag{29}$$

Thus, unlike the lottery, the auction clears the market through price adjustment rather than quantity rationing.

Table 12 reports equilibrium outcomes under the lottery and the auction from 2016-2023. In our counterfactual allocations, the equilibrium auction price ranges from approximately $12,372 (2016) to $46,514 (2019), all substantially above the observed fee of $6,470. This increase in prices translates into substantial revenue gains. Summing across years, total government revenue over 2016-2023 more than triples, going from $1.55 billion under the lottery to $5.37 billion under the auction.

Importantly, the auction increases total surplus in every year. We define total surplus as the sum of firm surplus, worker surplus, and government revenue. Firm surplus is the difference between firms' willingness to pay and the visa price, while worker surplus is the difference between wages and outside options. Over 2016-2023, total surplus roughly doubles s from $10.78 billion under the lottery to $21.54 billion under the auction.

These gains reflect improved allocative efficiency as visas are reallocated toward higher-valuation matches. Firm surplus increases despite higher visa prices because high-willingness-to-pay firms are more likely to secure visas, while worker surplus rises due to the higher marginal productivity of the matches that are realized.

The auction changes the marginal allocation rule such that visas are assigned to the highest willingness-to-pay positions rather than randomly among applicants. The resulting reallocation is substantial: pooled across years, several large IT consulting firms lose between 60% and 95% of

Table 12: Auction vs Lottery Comparison (2016–2023)

	2016		2017		2018		2019	
	Auction	Lottery	Auction	Lottery	Auction	Lottery	Auction	Lottery
Number of Visas	30,417	30,417	26,840	26,840	27,294	27,294	19,494	19,494
Price	12.37	6.47	18.53	6.47	31.08	6.47	46.51	6.47
($000)	(12.35,12.39)	(6.47,6.47)	(18.49,18.57)	(6.47,6.47)	(31.04,31.12)	(6.47,6.47)	(46.43,46.59)	(6.47,6.47)
Government	0.38	0.2	0.5	0.17	0.85	0.18	0.91	0.13
Revenue ($Bil)	(0.38,0.38)	(0.2,0.2)	(0.5,0.5)	(0.17,0.17)	(0.85,0.85)	(0.18,0.18)	(0.91,0.91)	(0.13,0.13)
Firm Surplus	0.9	0.42	1.01	0.49	1.67	0.81	0.9	0.63
($Bil)	(0.9,0.9)	(0.42,0.42)	(1.01,1.01)	(0.49,0.49)	(1.67,1.67)	(.81,.81)	(0.90,0.91)	(0.63,0.63)
Worker Surplus	0.67	0.32	0.67	0.35	0.71	0.43	0.52	0.38
($Bil)	(0.67,0.67)	(0.32,0.32)	(0.67,0.67)	(0.35,0.35)	(0.71,0.71)	(0.43,0.43)	(0.52,0.52)	(0.38,0.38)
Total Surplus	1.95	0.94	2.18	1.02	3.23	1.42	2.33	1.13
($Bil)	(1.95,1.95)	(0.94,0.94)	(2.18,2.18)	(1.02,1.02)	(3.23,3.23)	(1.42,1.42)	(2.33,2.33)	(1.13,1.13)

	2020		2021		2022		2023	
	Auction	Lottery	Auction	Lottery	Auction	Lottery	Auction	Lottery
Number of Visas	34,703	34,703	33,217	33,217	36,026	36,026	32,313	32,313
Price	20.35	6.47	15.62	6.47	28.15	6.47	15.28	6.47
($000)	(20.31,20.38)	(6.47,6.47)	(15.6,15.64)	(6.47,6.47)	(28.12,28.18)	(6.47,6.47)	(15.26,15.3)	(6.47,6.47)
Government	0.71	0.22	0.52	0.21	1.01	0.23	0.49	0.21
Revenue ($Bil)	(0.70,.71)	(0.22,0.22)	(0.52,0.52)	(0.21,0.21)	(1.01,1.02)	(0.23,0.23)	(0.49,0.49)	(0.21,0.21)
Firm Surplus	1.50	0.92	1.25	0.72	1.68	0.77	1.12	0.63
($Bil)	(1.50,1.50)	(0.92,0.92)	(1.25,1.25)	(0.71,0.72)	(1.67,1.68)	(0.77,0.77)	(1.12,1.12)	(0.63,0.63)
Worker Surplus	0.94	0.69	0.83	0.55	0.99	0.56	0.82	0.54
($Bil)	(0.94,0.94)	(0.69,0.69)	(0.82,0.83)	(0.55,0.55)	(0.99,0.99)	(0.56,0.56)	(0.82,0.82)	(0.54,0.54)
Total Surplus	3.14	1.84	2.6	1.48	3.68	1.57	2.43	1.38
($Bil)	(3.14,3.15)	(1.84,1.84)	(2.59,2.60)	(1.48,1.48)	(3.68,3.68)	(1.57,1.57)	(2.43,2.43)	(1.38,1.38)

This table compares visa price, government revenue, firm surplus, worker surplus, and total surplus between the simulated lottery and the simulated auction. For both the auction and lottery, the mean across 100 model simulations is reported along with the 5th and 95th percentile in parenthesis below the mean.

their lottery visas, while major technology firms increase their allocations by 90% to over 130%, indicating that the mechanism change fundamentally reshapes the distribution of visas across firms. The 20 largest average gains and losses in visa allocations over 2016-2023 are reported in Table 13.

Table 13: Largest Average Increases in Visas and Decrease in Visas (2016-2023)

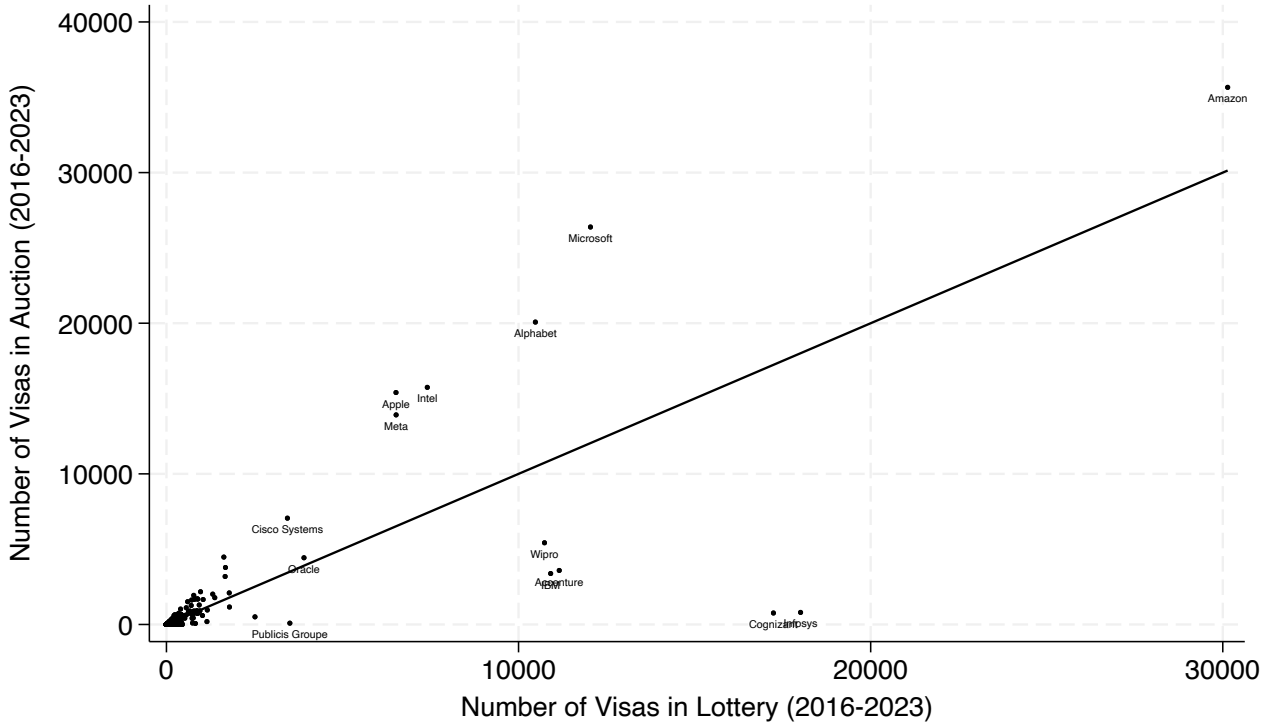
Top Gainers of Visas				Top Losers of Visas			
Firm	Lottery	Auction	Change	Firm	Lottery	Auction	Change
Microsoft	12,043	26,374	119.0%	Infosys	18,011	799	-95.6%
Alphabet	10,475	20,089	91.8%	Cognizant	17,240	762	-95.6%
Apple	6,519	15,395	136.1%	Accenture	11,156	3,584	-67.9%
Intel	7,408	15,741	112.5%	IBM	10,912	3,388	-69.0%
Meta	6,523	13,909	113.2%	Wipro	10,737	5,425	-49.5%
Amazon	30,136	35,664	18.3%	Publicis Groupe	3,504	87	-97.5%
Cisco	3,437	7,060	105.4%	Walmart	2,512	511	-79.7%
Nvidia	1,629	4,476	174.7%	Capital One	1,151	192	-83.3%
Salesforce	1,675	3,790	126.3%	EPAM	822	69	-91.6%
Verizon	1,667	3,181	90.8%	American Express	741	99	-86.6%
Adobe	969	2,171	124.0%	Cummins	1,788	1,164	-34.9%
Juniper Networks	775	1,933	149.5%	CVS Health	450	0	-100.0%
Amgen	715	1,625	127.3%	Fujitsu	439	0	-100.0%
Wayfair	601	1,511	151.2%	Cerner	1,023	595	-41.8%
Expedia	796	1,666	109.5%	T-Mobile	421	2	-99.5%
Applied Materials	872	1,703	95.4%	AECOM	438	92	-79.0%
Intuit	890	1,683	89.0%	Comcast	350	19	-94.6%
Dell	1,315	2,015	53.2%	Anthem	329	0	-100.0%
Marvel Technology	402	1,022	154.4%	Credit Suisse	761	434	-43.0%
VMware	1,040	1,658	59.4%	Synopsis	730	410	-43.8%

The left hand side of the table shows the 20 firms that gain the most average visas in going from a lottery to an auction. The right hand side of the table shows the 20 firms that lose the most average visas in going from a lottery to an auction. For both, displayed is the average number of visas acquired in the simulated lottery, the average number of visas acquired in the simulated auction, and the fraction gained or lost.

Figure 3 illustrates the firm-level reallocation pooling all years[24] by plotting the number of visas allocated under the lottery (horizontal axis) against the number allocated under the auction (vertical axis). The 45-degree line represents no change relative to the lottery. Points above the line correspond to firms that acquire more visas under the auction, while points below the line represent firms that lose visas.

[24]Figures for all individual years are reported in Appendix B.

Figure 3: Simulated Lottery Allocation Versus Auction Allocation of H-1B Visas For All Years



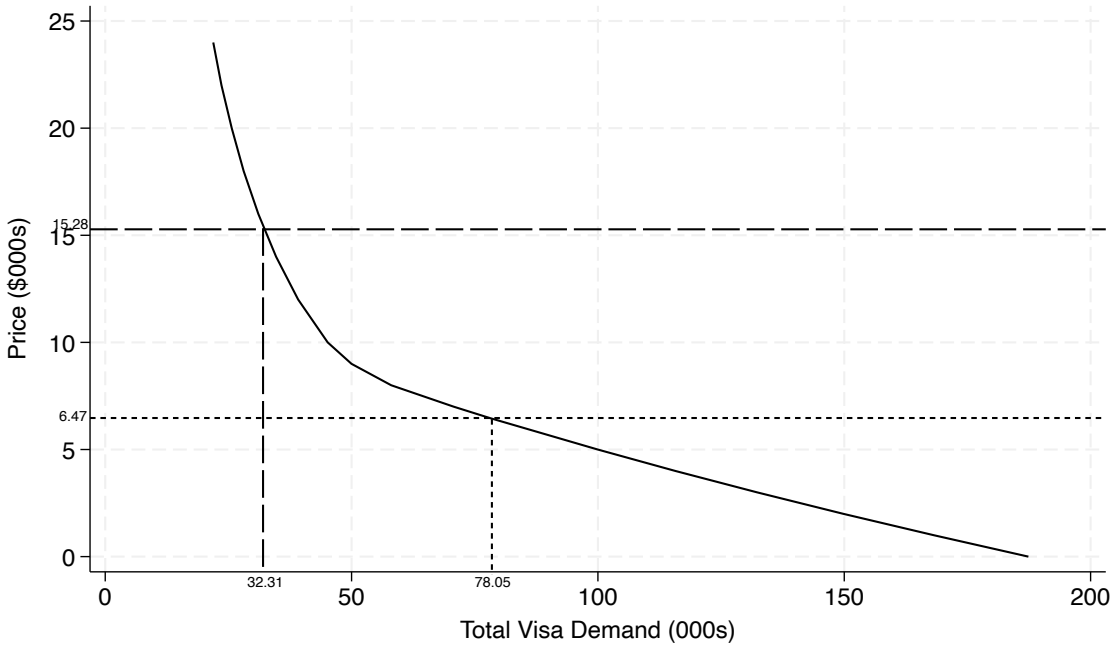
The figure highlights the substantial redistribution among the largest users of the H-1B system. For example, Microsoft, Alphabet, and Amazon expand their allocations under the auction and appear above the 45-degree line, whereas large IT consulting firms such as Cognizant, Infosys and Wipro experience sharp reductions and lie well below the line.

5.2 Why the Allocation Changes: Firm-Level Visa Demand Elasticities

The reallocation documented in Section 5.1 arises from heterogeneity in firm-level visa demand curves. Because firms differ in their marginal willingness to pay for foreign workers, increasing the visa price affects firms asymmetrically.

From the aggregate perspective, Figure 4 plots total visa demand as a function of the visa price for 2023. At the observed fee of \$6,470, aggregate demand (78,051 visas) substantially exceeds the cap V_t (32,313 visas for public firms), implying that the lottery operates in a regime of excess demand. Under the auction, the visa price rises until demand equals the cap. In 2023, the equilibrium

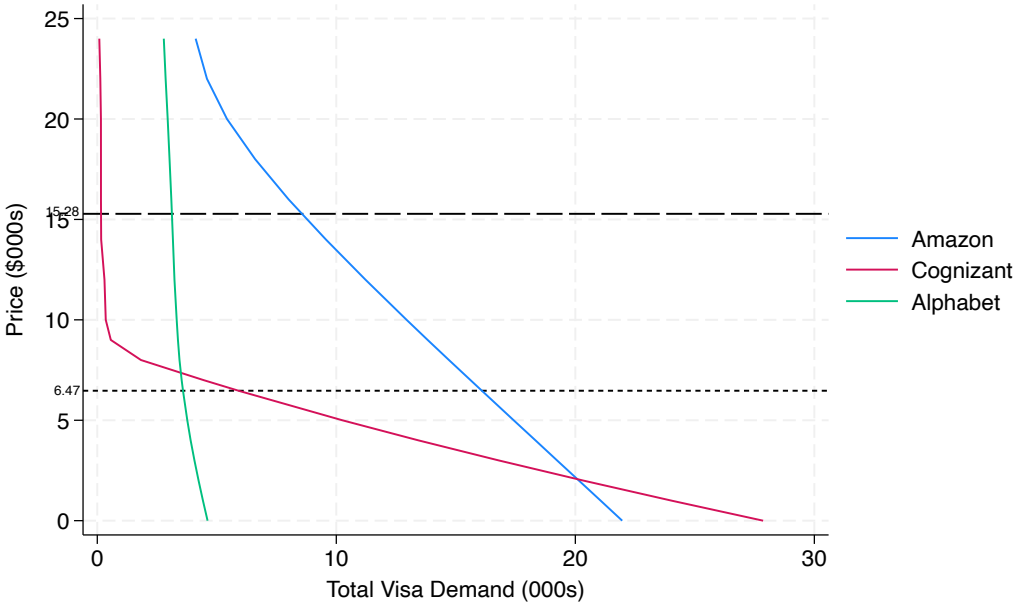
Figure 4: Total Demand for Visas at Different Prices (2023)



price of approximately \$15,277 clears the market by eliminating low-willingness-to-pay positions. The key mechanism is therefore not simply a higher price, but the endogenous shutdown of marginal positions. As the visa price increases, positions with WTP_{jtp} just above \$6,470 are the first to be excluded. Higher-valuation positions remain in equilibrium.

To understand how this affects firms differentially, Figure 5 plots firm-level demand curves for three large users of the H-1B program: Amazon, Cognizant, and Alphabet. These firms differ both in the level and in the slope of their demand curves. Cognizant exhibits high demand at low prices but a steeply downward-sloping demand curve. At the lottery fee, the firm demands 5,911 visas, but demand falls sharply as the price increases, falling to 4,447 visas at a price of \$8,000 and 155 visas at the auction price of \$15,277. By contrast, Alphabet displays a more inelastic demand curve. Although its initial demand is lower than Cognizant's (3,587 visas at the observed price), its demand remains relatively stable as the price rises (for a demand of 3,4440 visas at \$8,000 and 3,118 visas at the auction price). Amazon lies between these two cases, with large initial demand (of 16,091 visas) but substantial sensitivity to price (falling to 8,549 visas at the auction price).

Figure 5: Selected Firm Demand for Visas at Different Prices (2023)



The auction interacts with this heterogeneity in equilibrium. As the visa price rises to clear the market, firms with many low-surplus positions contract sharply, while firms whose marginal positions remain highly productive retain or expand their allocations. Thus, the auction changes the marginal allocation rule so that instead of equalizing allocation probabilities across applicants, it equalizes marginal willingness to pay across firms. The counterfactual reallocation across firms is the equilibrium consequence of the heterogeneity in demand elasticities observed in Figure 5. Approaches that abstract from firm-specific demand heterogeneity would miss precisely this reallocation channel.

5.3 Decomposing Firm Willingness to Pay

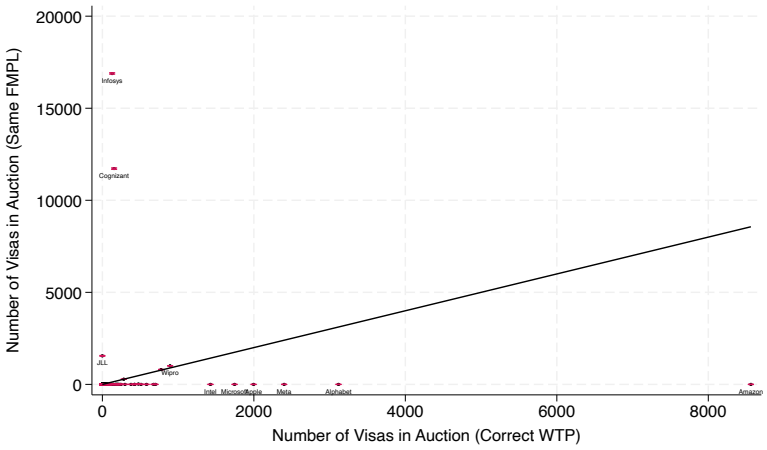
We now examine the three structural components of willingness to pay (in equation 27) determining demand: (i) productivity differences between foreign and domestic workers, (ii) wage differentials arising from bargaining, and (iii) search costs. To assess the relative importance of these channels, we conduct counterfactual exercises in which we selectively shut down cross-firm variation in each component and re-simulate the auction.

Productivity Channel. First, we eliminate cross-firm differences in foreign marginal products by fixing MPL_{jt}^{f} at a common level across firms ($150,000), while leaving all other components unchanged. Under this counterfactual, the allocation shifts markedly. Large technology firms such as Amazon, Alphabet, and Meta lose a substantial share of visas relative to the baseline auction, while several IT consulting firms regain visas. This is illustrated in Panel (a) of Figure 6 for 2023. This result indicates that the high willingness to pay observed among major technology firms in the baseline auction is primarily driven by higher marginal productivity of foreign workers, rather than by wage or search advantages.

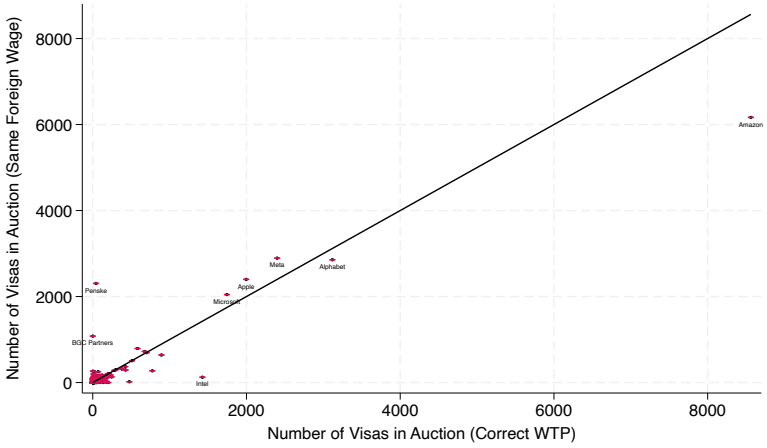
Wage Channel. Second, we eliminate cross-firm variation in foreign wages by fixing w_{jtp}^{f} at a common level ($100,000). In this scenario, firms that rely heavily on relatively low foreign wages experience the largest reductions in visa allocations. Panel (b) of Figure 6 illustrates this counterfactual allocation. While both Amazon and Alphabet receive fewer visas in this scenario, the largest decline is for Intel (who receive an average of 1,427 visas in the main auction, but only 124 visas under wage equalization). This implies that the main driver behind Intel's willingness to pay for visas is the low wages they offer their H-1B workers. Microsoft, Apple, and Meta instead acquire more visas under this alternative, suggesting that they demand H-1B visa workers despite their high salaries.

Search Cost Channel. Third, we analyze the search cost channel by reducing all firm search costs by 10% relative to the mean. Panel (c) of Figure 6 shows the corresponding reallocation for 2023. The biggest losers of visas are firms that have negative search costs, meaning that it is harder for them to find domestic workers for their skilled roles than it is to find a foreign worker. These include a lot of the technology firms such as Amazon (loses 62.6% of their visas), Apple (loses 63.9% of their visas), Intel (loses 86.3% of their visas), and Meta (loses 51.1% of their visas). The largest gainer in this exercise is JLL, a firm that had one of the largest estimated search costs in 2023, indicating that they found it especially difficult to find H-1B workers to fill their skilled openings. Other firms with large search costs included Fresenius Medical Care, Charter Communications, Comcast, Oracle, and IBM.

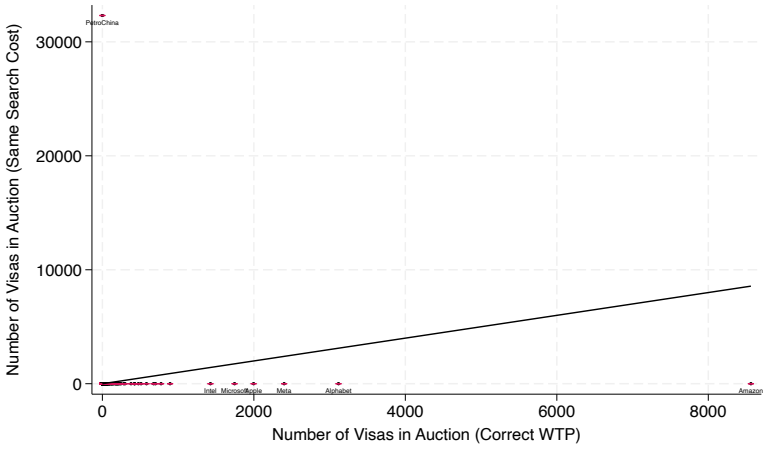
Figure 6: Simulated Auction Allocations of H-1B Visas Adjusting Firm Willingness to Pay



(a) Same Foreign Marginal Product of Labor



(b) Same Foreign Wages

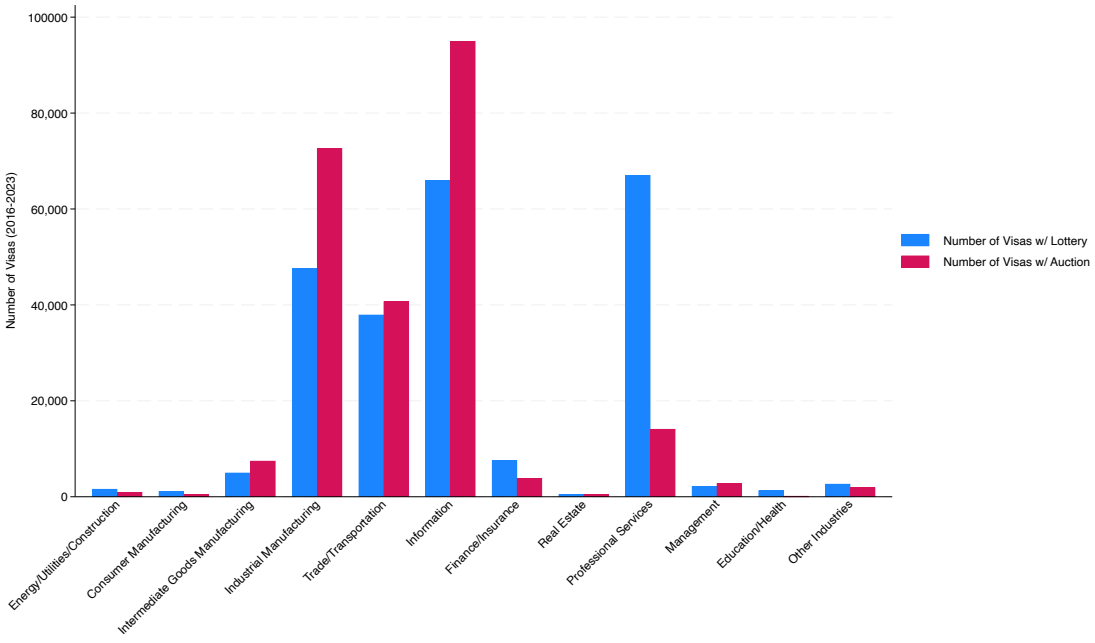


(c) Reduction in Search Costs by 10%

5.4 Industry and Occupational Reallocation

The firm-level reallocation documented above aggregates into systematic shifts across industries and occupations. Figure 7 reports how the distribution of visas across industries changes in going from a lottery to an auction, pooled across 2016–2023. Across industries, visas shift away from

Figure 7: Distribution Across Industries

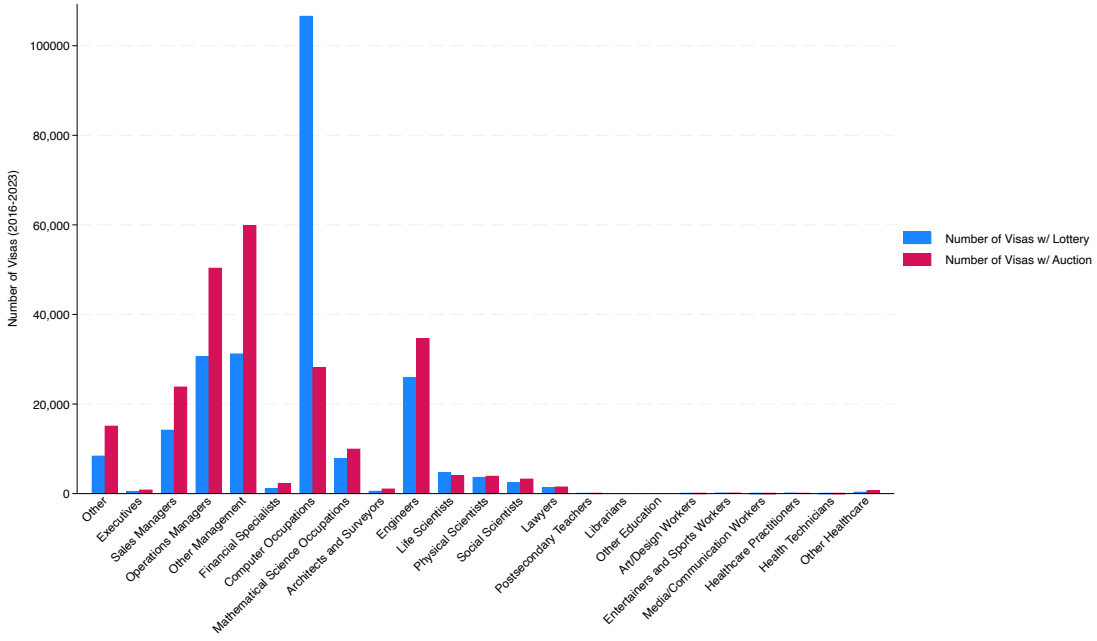


professional services (which includes large IT consulting firms) and toward information and industrial manufacturing. Under the lottery, nearly 70,000 visas over the sample period are allocated to professional services firms. Under the auction, this number falls by approximately 79%. By contrast, allocations to industrial manufacturing and information increase by roughly 53% and 44%, respectively. These shifts are consistent with the firm-level evidence that sectors in which foreign workers generate higher marginal product retain or expand their allocations as the visa price rises.

Occupational composition changes in a similar direction. Computer occupations, which account for the majority of visas under the lottery, experience a 74% decline under the auction. In contrast, management occupations, including sales managers, operations managers, and other management roles, increase their share. These occupations are associated with higher domestic wage baselines

and higher estimated marginal products, which translate into higher willingness to pay in the model. Figure 8 illustrates this shift.

Figure 8: Distribution Across Occupations



Changes in educational attainment and country of origin follow similar patterns. Under the auction, visas shift modestly toward workers with advanced degrees and away from those with only a bachelor's degree (e.g. a 35.8% decrease in visas going to workers with a bachelor's degree and a 48.0% increase in visas going to workers with a master's degree). There is also a reallocation across source countries, with more visas going to workers from high income countries under the auction, though the overwhelming dominance of Indian workers in the H-1B program remains. Detailed results for education and nationality are reported in Appendix C.

5.5 Auction Credits and Targeted Reallocation

One advantage of a price-based allocation mechanism is that it can accommodate distributional objectives through targeted bid credits while preserving market clearing. We illustrate this by introducing several auction credits in 2023 and re-simulating equilibrium.

Table 14: Simulations of Auctions with Credits in 2023

		Visa Price ($000)	Government Revenue ($ Mil)	Visas to Credit Firms (000s)	Percentage of Visas to Credit Firms
Small Firm Credit	No Credit	15.28 (15.26, 15.30)	493.72 (493.09, 494.34)	1.15 (1.14, 1.16)	3.55% (3.51%, 3.59%)
	Credit	16.67 (16.64, 16.69)	506.07 (505.15, 506.77)	3.24 (3.22, 3.27)	10.04% (9.98%, 10.10%)
Industry Credit	No Credit	15.28 (15.26, 15.30)	493.72 (493.09, 494.34)	13.37 (13.35, 13.40)	41.38% (41.33%, 41.46%)
	Credit	22.19 (22.17, 22.21)	545.33 (544.46, 546.01)	17.16 (17.14, 17.19)	53.12% (53.04%, 53.19%)
Education Credit	No Credit	15.28 (15.26, 15.30)	493.72 (493.09, 494.34)	19.96 (19.90, 20.00)	61.78% (61.57%, 61.94%)
	Credit	21.65 (21.60, 21.69)	462.13 (460.25, 463.37)	23.75 (23.71, 23.79)	73.51% (73.37%, 73.62%)
Occupation Credit	No Credit	15.28 (15.26, 15.30)	493.72 (493.09, 494.34)	25.21 (25.16, 25.26)	78.01% (77.85%, 78.17%)
	Credit	21.57 (21.52, 21.61)	406.96 (405.56, 408.19)	29.00 (28.98, 29.02)	89.75% (89.69%, 89.82%)
Country Credit	No Credit	15.28 (15.26, 15.30)	493.72 (493.09, 494.34)	2.67 (2.64, 2.70)	8.27% (8.16%, 8.37%)
	Credit	18.03 (18.00, 18.06)	521.25 (520.20, 522.24)	6.13 (6.09, 6.17)	18.97% (18.86%, 19.09%)

This table summarizes the counterfactual simulations with credits. For each simulation we report the results without an auction credit and with the auction credit. We report the visa price (in thousands of dollars), the resulting government revenue (in millions of dollars), the number of visas going to workers that qualified for the credit (in thousands of visas), and the percentage of visas going to workers that qualified for the credit. The mean result across the 100 simulations is reported as well as the 5% and 95% percentile in parenthesis below the reported mean. The first panel is for a $10,000 credit to firms with less than 10,000 employees. The second is for a $10,000 credit for firms that are not in industrial manufacturing or information. The third is for a $10,000 credit for workers with above a Bachelor's degree. The fourth is for a $10,000 credit for workers that not in the computer occupation. The fifth is for a $10,000 credit for workers flow low or low-middle income countries, excluding India.

Under a credit scheme, eligible firms or workers receive a bid augmentation equal to the credit amount. Operationally, if a firm's willingness to pay for a position is WTP_{jtp} and the position qualifies for a credit of size κ, the effective bid becomes $WTP_{jtp} + \kappa$. The auction price is then determined in the usual way, clearing the market at the $(V_t + 1)$-th highest effective bid.

Table 14. reports the results for several illustrative credits, including: (i) a credit for small firms,

(ii) a credit for firms outside major technology sectors, (iii) a credit for workers with advanced degrees, (iv) a credit for non-computer occupations, and (v) a credit for workers from lower-income countries (excluding India).

The effects are economically meaningful but vary by design. A credit for small firms increases their share of visas from 3.55% to 10.04% and raises the equilibrium visa price modestly. A credit targeting firms outside industrial manufacturing and information increases their share from 41.38% to 53.12%, accompanied by a larger increase in the auction price. Worker-based credits generate compositional shifts across education and occupation categories and can either increase or decrease government revenue depending on how broadly the credit applies.

Overall, these exercises illustrate that an auction framework provides a flexible policy instrument. Legislators seeking to advance particular distributional objectives (such as supporting small firms, encouraging advanced degrees, or diversifying source countries) can do so transparently through bid credits, while preserving market clearing and the underlying price-based allocation rule.

5.6 A Fixed Visa Fee of $33,333

As a final exercise, we consider a higher fixed visa fee of $33,333. This highlights the limits of fixed-price reforms in that even if a high fee extracts revenue, it does not adjust to current demand conditions, and may produce significantly lower total surplus compared to the auction.[25]

Table 15 reports the results for 2023. At a fee of $33,333, aggregate demand falls substantially, and the number of visas allocated declines from 32,313 under both the lottery and the auction to 16,716. Government revenue increases relative to the lottery (from $0.21 billion to $0.56 billion), but remains below the level generated by the auction despite the higher price, reflecting the contraction in quantity. Total surplus under the fixed fee ($1.74 billion) exceeds that under the lottery ($1.38 billion) but remains well below the auction outcome ($2.43 billion).

Figure 9 illustrates how the $33,333 fee reshapes allocations relative to the lottery. We plots firm-level visa shares under the lottery against those under the fixed price. The 45-degree line represents

[25]This specific $33,333 value roughly corresponds to annualizing the proposed $100,000 fee discussed in policy debates in 2026. Because our model evaluates firm willingness to pay as a one-year flow surplus and abstracts from dynamic continuation and separation decisions, mapping a multi-year fee into an equivalent annual value is necessarily approximate.

Table 15: Comparison of Lottery vs. Auction vs. Fixed $33K Price in 2023

	Lottery	Auction	Fixed $33K Price
Visa Price ($000)	6.47	15.28	33.33
No. of Visas Allocated	32,313	32,313	16,716
Government Revenue ($ Bil)	0.21	0.49	0.56
Firm Surplus ($ Bil)	0.63	1.12	0.71
Worker Surplus ($ Bil)	0.54	0.82	0.48
Total Surplus ($ Bil)	1.38	2.43	1.74

This table reports averages over 100 simulations of the lottery, auction, and a fixed $33,333 price for 2023.

no change. Unlike the auction, which redistributes visas across many firms, the fixed price induces a sharp contraction toward the very highest-valuation firms. Several large consulting firms fall close to zero allocation, while a small set of large technology firms increase their share substantially. The compositional effects are correspondingly stronger. As shown in the industry panel (Figure

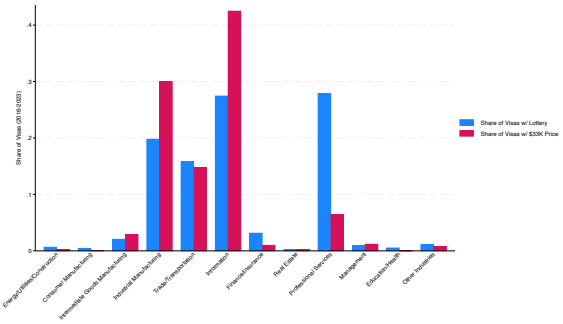
Figure 9: Simulated Lottery Allocation Share Versus Fixed $33.3K Price Allocation Share for All Years



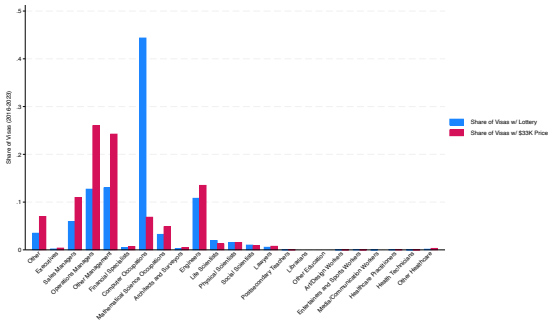
10, Panel (a)), visas shift markedly away from professional services and toward information and industrial manufacturing. At the occupational level (Panel (b)), computer occupations contract sharply while management roles expand. The education distribution tilts further toward master's degrees (Panel (c)), and although the share of Indian workers declines somewhat, the overall country composition remains dominated by India (Panel (d)).

These patterns reflect the mechanical nature of a high fixed price. Because the fee is imposed exogenously rather than determined by market clearing, quantity adjusts discontinuously. If the fee is set above the equilibrium auction price, some high-surplus matches are eliminated along with low-surplus ones.

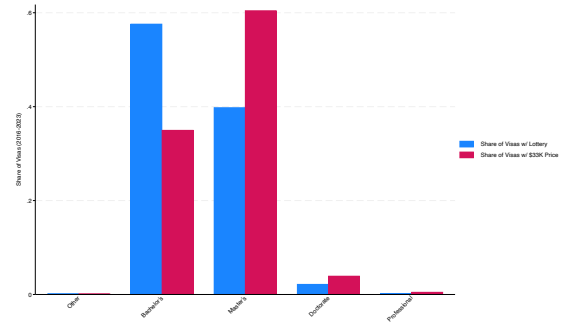
Figure 10: Distribution of Allocations in Lottery Versus Fixed $33.3K Price



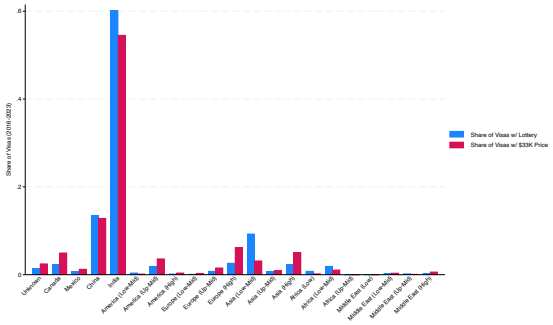
(a) Distribution Across Industries



(b) Distribution Across Occupations



(c) Distribution Across Education



(d) Distribution Across Countries of Origin

6 Conclusion

We study how replacing the H-1B lottery with a second-price auction would reallocate visas across firms and workers and affect government revenue and surplus. We combine newly obtained US-

CIS petition-level wage data with firm-level production information from Compustat to estimate a structural model of firm demand for H-1B workers. In the model, firms fill skilled vacancies with either domestic or H-1B workers. Domestic and foreign workers differ in marginal products, wages (determined via Nash bargaining), and hiring frictions. We exploit the randomness of the lottery to obtain exogenous variation in realized foreign employment, which identifies the marginal product of foreign labor and allows us to recover firm-level willingness to pay for additional visas.

We use the estimated model to simulate a counterfactual second-price auction. Relative to the current lottery, the auction reallocates visas toward large technology firms (e.g., Microsoft, Alphabet, Meta) and away from IT consulting firms (e.g., Cognizant). The reallocation is economically large as some major IT consulting firms lose the vast majority of their visas, while leading technology firms substantially expand H-1B employment. The auction also shifts visas toward more highly educated workers. The implied market-clearing auction price ranges from roughly $12,000 to over $40,000 per visa across years in our sample.

The revenue implications are substantial. Under the current administrative fee, lottery-based allocations generate on the order of $150–$230 million in annual revenue in our sample. Under the auction, annual revenue rises to roughly $380 million to over $1 billion, implying an annual increase of approximately $200 million to $800 million depending on the year. Over the eight years in our sample, cumulative revenue increases from about $1.5 billion under the lottery to more than $5 billion under the auction. At the same time, aggregate surplus rises as visas are sorted toward higher marginal valuations.

These findings are directly relevant for the design of high-skilled immigration policy in environments where access to foreign labor is restricted. In principle, a sufficiently high fixed visa fee could replicate the allocation achieved by an auction if policymakers knew firms' demand conditions. In practice, those demand schedules are unobserved and vary substantially across firms and over time. A second-price auction provides a simple way to allow prices to adjust endogenously to demand while allocating visas to firms with the highest marginal valuations. Our estimates suggest that such a mechanism would generate large reallocations across firms and substantial revenue gains relative to the lottery, without requiring policymakers to calibrate visa prices ex ante.

References

Ackerberg, Daniel, Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica*, 83(6): 2411-2451.

AFL-CIO. "New AFL-CIO Report Exposes Walmart's Abuse of H-1B Visas." AFL-CIO press release, April 2, 2015. https://aflcio.org/press/releases/new-afl-cio-report-exposes-walmarts-abuse-h-1b-visas

Amuedo-Dorantes, Catalina, and Delia Furtado. 2017. "Setting for Academia? H-1B Visas and the Career Choices of International Students in the United States." *Journal of Human Resources*, 54(2): 401-429.

Anderson, Stuart. 2025. "New Trump Immigration Policy: Ending The H-1B Visa Lottery." *Forbes*, July 21, 2025. https://www.forbes.com/sites/stuartanderson/2025/07/21/new-trump-immigration-policy-ending-the-h-1b-visa-lottery

Becker, Gary S., and Edward P. Lazear. 2013. "A Market Solution to Immigration Reform." *The Wall Street Journal*, March 1. https://www.wsj.com/articles/SB10001424127887323375204578271531542362850

Borjas, George J. 2003. "The Labor Demand Curve is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *The Quarterly Journal of Economics*, 118(4): 1335-1374.

Borjas, George J., Jeffrey Grogger, and Gordan H. Hanson. 2011. "Substitution Between Immigrations, Natives, and Skill Groups." *NBER Working Papers 17461*, National Bureau of Economic Research.

Borjas, George J. 2026. "The H-1B Wage Gap, Visa Fees, and Employer Demand." *NBER Working Papers 34793*, National Bureau of Economic Research.

Bound, John, Gaurav Khanna, and Nicolas Morales. 2017. "Understanding the Economic Impact of the H-1B Program on the U.S." NBER Working Paper 23153.

BrightQuery. 2022. Sector Public Companies - Financials/Employment/Benefits. Dewey Data. https://doi.org/10.82551/Q56V-PX08

Caiumi, Alessandro, ad Giovanni Peri. 2025. "Immigration's Effect on US Wages and Employment Redux." Working Paper.

Card, David. 1990. "The Impact of the Mariel Boatlift on the Miami Labor Market." *ILR Review*, 43(2): 245-257.

Card, David, and Thomas Lemieux. 2001. "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." *The Quarterly Journal of Economics*, 116(2): 705-746.

Cassidy, John. 2017. "Why the MAGA Fight over H-1B Visas Is Crossing Party Lines." *The New Yorker*, April 20, 2017. https://www.newyorker.com/news/the-financial-page/why-the-maga-fight-over-h-1b-visas-is-crossing-party-lines

Chassamboulli, Andri, and Giovanni Peri. 2020. "The Economic Effect of Immigration Policies: Analyzing and Simulating the U.S. Case." *Journal of Economic Dynamics and Control*, 114.

Ciccone, Antonio, and Giovanni Peri. 2005. "Long-Run Substitutability Between More and Less Educated Workers: Evidence from U.S. States, 1950-1990." *The Review of Economics and Statistics*, 87(4): 652-663.

De Loecker, Jan, and Frederic Wazynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review*, 102(6): 2437-2471.

Depew, Briggs, Peter Norlander, and Todd Sorensen. 2013. "Flight of the H-1B: Inter-firm Mobility and Return Migration Patterns for Skilled Guest Workers." IZA Discussion Paper No. 7456.

Dimmock, Stephen G., Jiekun Huang, Scott J. Weisbenner. 2022. "Give Me Your Tire, Your Poor, Your High-Skilled Labor: H-1B Lottery Outcomes and Entrepreneurial Success." *Management Science*, 68(9): 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): 2501-2533

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." Working Paper.

Glennon, Britta. 2024. "How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program." *Management Science*, 70(2): 907-930.

Kato, Takao, and Chad Sparber. 2013. "Quotas and Quality: The Effect of H-1B Visa Restrictions on the Pool of Prospective Undergraduate Students from Abroad." *The Review of Economics and Statistics*, 95(1): 109-126.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *The American Economic Review*, 82(3): 371-392.

Kerr, Sari Pekkala, and William R. Kerr. 2020. "Immigration Policy Levers for US Innovation and Startups." NBER Working Paper 27040.

Kerr, Sari Pekkala, William R. Kerr, and William F. Lincoln. 2015. "Skilled Immigration and the Employment Structures of the U.S. Firms." *Journal of Labor Economics*, 33(1): 147-186.

Kerr, William R. 2013. "U.S. High-Skilled Immigration, Innovation, and Entrepreneurship: Empirical Approaches and Evidence." National Bureau of Economic Research.

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(3): 473-508.

Kreighbaum, Andrew. 2025. "DHS Crafting Rule for Weighted Selection of H-1B Petitions." *Bloomberg Law (Daily Labor Report)*, July 18, 2025. https://natlawreview.com/article/dhs-submits-h-1b-weighted-selection-rule-federal-review-implications-employers-and

Levinsohn, James, and Amil Petrin. 2013. "Estimating Production Functions Using Inputs to Control for Unobservables". The Review of Economic Studies, 70(2): 317-341.

Mahajan, Parag, Nicolas Morales, Kevin Shih, Mingyu Chen, and Agostina Brinatti. 2024. "The Impact of Immigration on Firms and Workers: Insights from the H-1B Lottery." Federal Reserve Bank of Richmond Working Paper Series.

Manacorda, Marco, Alan Manning, and Jonathan Wadsworht. 2012. "The Impact of Immigration of the Structure of Wages: Theory and Evidence from Britain." *Journal of the European Economic Association*, 10(1): 120-151.

Mandelman, Federico, Mishita Mehra, and Hewei Shen. 2025. "Skilled Immigration Frictions as a Barrier for Young Firms." *Journal of Monetary Economics*.

Matloff, Norman. 2004. "Globalization and the American IT Worker." *Communications of the ACM*, 47(11): 27-29.

Mayda, Anna Maria, Francesc Ortega, Giovanni Peri, Kevin Shih, and Chad Sparber. 2023. "Coping with H-1B Shortages: Firm Performance and Mitigation Strategies." *ILR Review*

Mayda, Anna Maria, Francesc Ortega, Giovanni Peri, Kevin Shih, Chad Sparber. 2018. "The Effect of the H-1B Quota on the Employment and Selection of Foreign Born Labor." *European Economic Review*, 108: 105-128.

Neufeld, Jeremy. 2025. "Putting America First Requires H-1B Visa Reform." *The Wall Street Journal (Opinion)*, January 7, 2025. https://www.wsj.com/opinion/putting-america-first-requires-h-1b-visa-reform-skilled-labor-policy

Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6): 1263-1297.

Orrenius, Pia M., and Madeline Zavodny. 2020. "An Auctions Approach to Immigration Policy." IZA Policy Paper.

Ottaviano, Gianmarco, and Giovanni Peri. 2008. "Immigration and National Wages: Clarifying the Theory and the Empirics." *Working Paper 14188*, National Bureau of Economic Research.

Ottaviano, Gianmarco, and Giovanni Peri. 2012. "Rethinking the Effect of Immigration on Wages." *Journal of the European Economic Association*, 10(1): 152-197.

Pathak, Parag A., Alex Rees-Jones, and Tayfun Sonmez. 2025. "Immigration Lottery Design: Engineered and Coincidental Consequences of H-1B Reforms." *Review of Economics and Statistics*, 107(1): 1-13.

Peri, Giovanni. 2012. "Rationalizing U.S. Immigration Policy: Reforms for Simplicity, Fairness, and Economic Growth." Brookings Institution Discussion Paper 2012-01.

Peri, Giovanni, Kevin Shih, and Chad Sparber. 2015. "Foreign and Native Skilled Workers: What Can We Learn from H-1B Lotteries?" National Bureau of Economic Research.

Peri, Giovanni, Kevin Shih, and Chad Sparber. 2015. "STEM Workers, H-1B Visas, and Productivity in U.S. Cities." *Journal of Labor Economics*, 33(1): 225-255.

Preston, Julia. 2015. "Pink Slips at Disney. But First, Training Foreign Replacements." *The New York Times*, June 3. https://www.nytimes.com/2015/06/04/us/last-task-after-layoff-at-disney-train-foreign-replacements.html

Preston, Julia. 2015. "Toys R Us Brings Temporary Foreign Workers to U.S. to Move Jobs Overseas." *The New York Times*, September 29. https://www.nytimes.com/2015/09/30/us/toys-r-us-brings-temporary-foreign-workers-to-us-to-move-jobs-overseas.html

Sharma, Rishi, and Chad Sparber. 2020. "Buying Lottery Tickets for Foreign Workers: Search Cost Externalities Indued by H-1B Policy." IZA Discussion Paper No. 13892, IZA, Bonn, Germany.

Sparber, Chad. 2018. "Choosing Skilled Foreign-Born Workers: Evaluating Alternative Methods for Allocating H-1B Work Permits." *Industrial Relations: A Journal of Economy and Society*, 57(1): 3-34.

A Additional Estimation Results

This appendix reports additional estimation results not presented in the main text.

A.1 Instrument relevance

Table 16: Instrument Relevance

Dependent Variable:	Won Visas		Foreign Labor	
Independent Variable:	Lottery LCAs	All LCAs	Surprise Visas (Lottery LCAs)	Surprise Visas (All LCAs)
Coefficient	0.233	0.006	5.535	-6.226
Standard Error	(0.002)	(0.000)	(0.097)	(0.104)
R^2	0.284	0.130	0.220	0.233
Num. Observations	28,105	28,105	14,030	14,030

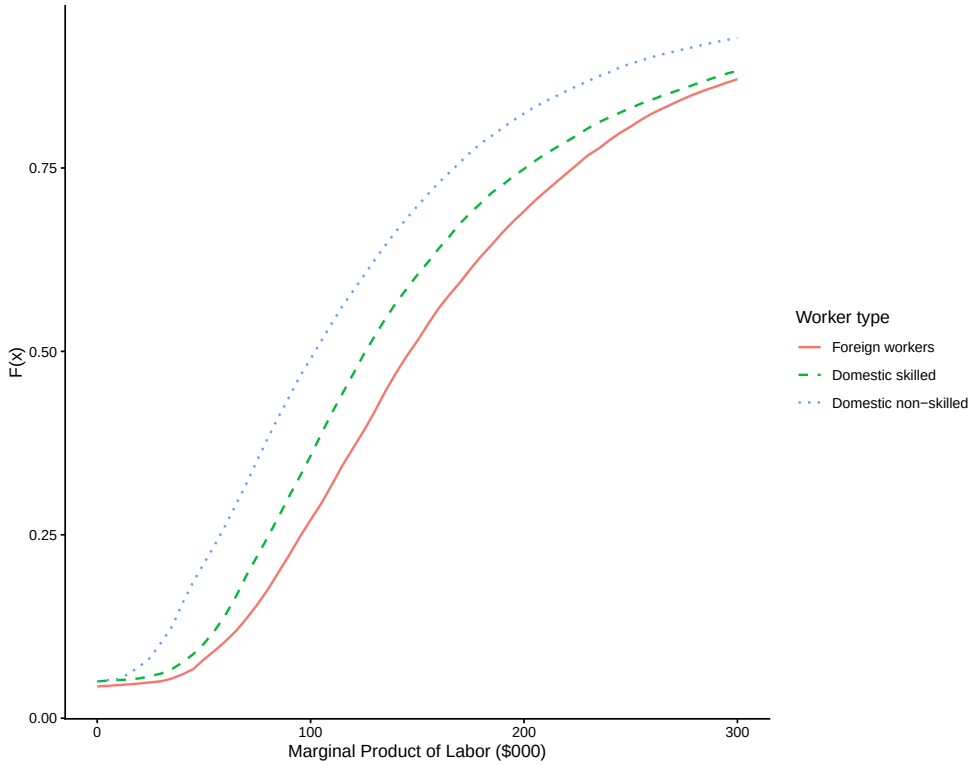
Results from regressing either won visas (columns (2) and (3)) or total foreign labor (columns (4) and (5)) on lottery LCAs (columns (2) and (4)) or all LCAs (columns (3) and (5)). Observations are at the firm-year level. All regressions also control for capital, number of domestic workers, and year indicators.

.

In constructing the instrument of the surprise number of visas, for the number of predicted visas in a given year we use the the number of LCAs that a firm submits in March and April before the lottery for employment start dates in August or September. We count only cap-window LCAs since applications outside this window may be submitted without any intention of actually acquiring a visa. We test the relevance of the instrument and the results are in Table 16. A regression of won visas on lottery LCAs produces a positive and statistically significant coefficient (column (2)), as does a regression of foreign labor on surprise number of visas based on this restricted set of LCAs (column (4)). Columns (3) and (5) of Table 16 show the results of the same regressions if all LCAs had been included. Total LCAs still has a positive and statistically significant effect on visas acquired, but the coefficient and R^2 are significantly lower (a coefficient of 0.006 versus 0.233 and a R^2 of 0.130 versus 0.284). Furthermore, if the surprise number of visas is constructed using all LCAs, it is negatively related to foreign labor. This is likely due to LCAs submitted outside of the "lottery" window not always being submitted with the intention of actually hiring foreign workers.

A.2 Estimated Marginal Products of Labor

Figure 11: Comparison of Empirical CDFs for Marginal Products of Labor for Foreign and Domestic Workers Based on Production Function Estimates



The production function estimates in Table 8 and Table 9 can be used to construct measures of marginal product of labor for each firm at their current number of workers of each type. The distribution of the resulting marginal products of labor are graphed in Figure 11. For the median firm-year in our sample, the marginal product of H-1B labor is $146,619, the marginal product of skilled domestic labor is $125,797, and the marginal product of non-skilled domestic labor is $102,130, all evaluated at the observed values for the inputs. For 64.07% of firm-years in our sample, the marginal product of foreign labor is higher than the marginal product of skilled domestic labor. This is one potential reason why firms may be willing to pay visa costs to hire a H-1B worker over their domestic counterpart.

A.3 Outside-option components

The outside option w^o_{oec} is modeled as additively separable in occupation, education, and country of origin. Tables 17 - 19 report the estimated components used in the wage equation.

Table 17: Worker Outside Option Estimates: Occupation

Occupation	Outside Option ($000)	Standard Error ($000)
Executives	125.962	8.240
Sales Managers	102.428	2.293
Operations Managers	97.141	1.420
Other Management	107.813	1.313
Financial Specialist	85.247	1.350
Computer Occupations	96.319	0.870
Mathematical Science Occupations	64.793	2.841
Architects and Surveyors	87.446	1.141
Engineers	80.630	0.905
Life Scientists	20.885	2.974
Physical Scientists	41.790	2.639
Social Scientists	97.393	1.614
Lawyers	72.825	9.075
Postsecondary Teachers	102.444	9.451
Librarians	71.373	11.143
Other Education	83.367	7.753
Art and Design Workers	82.086	2.051
Entertainers and Sports Workers	93.334	5.502
Media and Communications Workers	63.812	4.837
Healthcare Practitioners	119.151	2.039
Health Technicians	77.500	4.140
Other Healthcare	60.851	5.526

Estimates for the worker's outside option based on their occupation.

The occupations where foreign workers have the highest outside option are executives, healthcare practitioners, other management, and postsecondary teachers. The occupations where H-1B workers have the lowest outside option are life scientists, physical scientists, other healthcare, and media and communications workers.

The estimates for the impact of education in Table 18 show that a master's Degree adds only $2,766 to an H-1B worker's outside option compared to a bachelor's degree. A doctorate degree

Table 18: Worker Outside Option Estimates: Degree

Degree	Outside Option ($000)	Standard Error ($000)
Bachelor's Degree	0	-
Master's Degree	2.766	0.146
Doctorate Degree	12.788	0.604
Professional Degree	26.849	3.465

Estimates for the worker's outside option based on their degree. The parameter for a Bachelor's degree is normalized to zero.

Table 19: Worker Outside Option Estimates: Country of Origin

Degree	Outside Option ($000)	Standard Error ($000)
Canada	0	-
Mexico	-1.111	1.598
China	-9.682	0.937
India	-8.859	0.876
America (Low-Middle Income)	-17.106	4.030
America (Upper-Middle Income)	-4.716	1.504
America (High Income)	1.842	4.018
Europe (Low-Middle Income)	-1.570	2.006
Europe (Upper-Middle Income)	-2.222	1.278
Europe (High Income)	7.718	1.261
Asia (Low-Middle Income)	-12.443	1.043
Asia (Upper-Middle Income)	-11.095	1.911
Asia (High Income)	-4.975	1.193
Africa (Low Income)	-16.721	3.256
Africa (Low-Middle Income)	-14.366	1.693
Africa (Upper-Middle Income)	15.493	4.301
Middle East (Low Income)	-9.039	4.368
Middle East (Low-Miiddle Income)	-6.083	1.590
Middle East (Upper-Middle Income)	-9.370	2.984
Middle East (High Income)	1.685	1.757

Estimates for the worker's outside option based on their country of origin. The parameter for Canada is normalized to zero.

is estimated to add $12,788 and a professional degree is estimated to add $26,849 on top of a

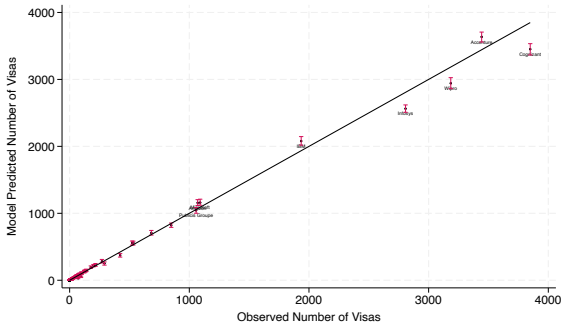
Bachelor's degree.

The parameter estimates for country of origin in Table 19 show that H-1B workers from upper-middle income countries in Africa have the highest outside option, estimated to be $15,493 above that of workers from Canada. Workers from high income countries in Europe, America, and the Middle East also have higher outside options than workers from Canada. Workers from lower income countries in Asia, Africa, and America have the lowest outside options. The overwhelming majority of H-1B workers come from India, where the estimated outside option is $8,859 below that of Canadian workers.

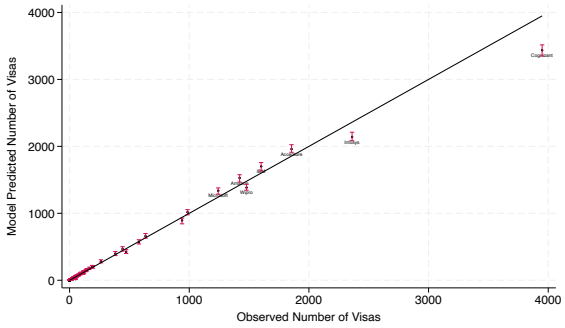
A.4 Model Fit by Year

The graphs in Figure 12 illustrate the model fit. Each figure is a scatterplot for a separate year, comparing the observed number of visas the firm acquired in that year (on the x-axis) with our simulated lottery's prediction for the firm's number of acquired visas that year (on the y-axis). The straight-line represents a perfect model fit. The vertical bands around each point represent 90% confidence intervals for the number of received visas.

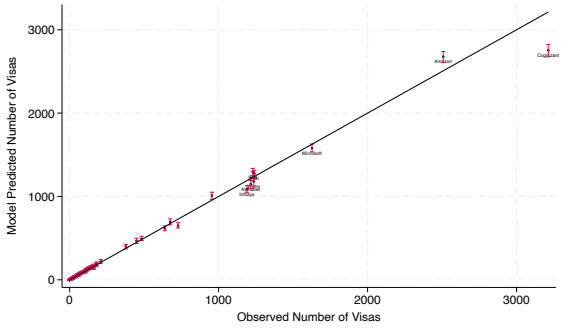
Figure 12: Model Fit by Year (2016–2022)



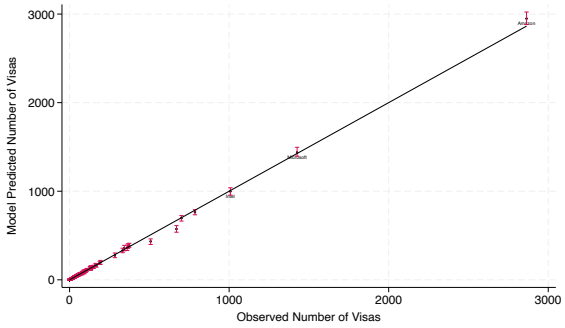
(a) 2016



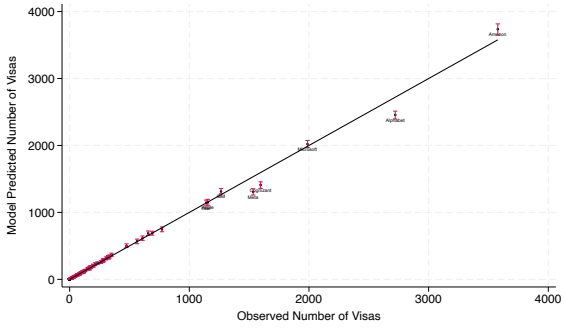
(b) 2017



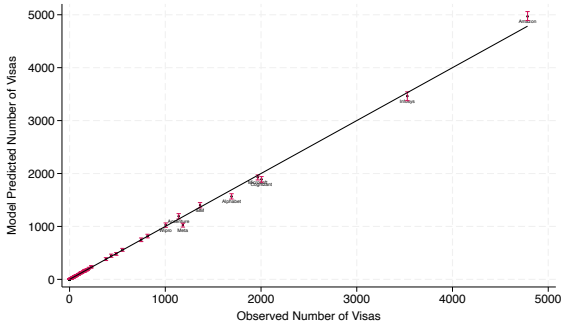
(c) 2018



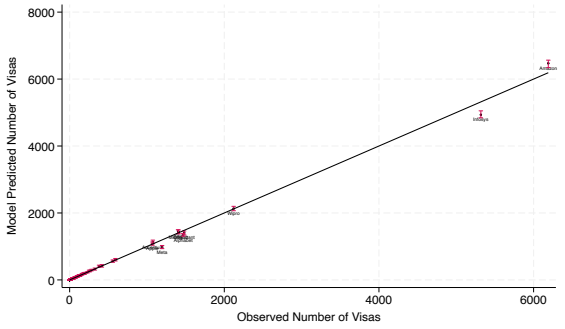
(d) 2019



(e) 2020



(f) 2021



(g) 2022

B Robustness to Alternative Parameter Values for θ_f and σ_f

We re-run the counterfactual analysis for different values for the parameters σ_f and θ_f. In each case, we do not re-estimate the bargained wage parameters, but do re-estimate the search costs. We do this for 12 different combinations of σ_f and θ_f. The different values we use for σ_f are $\{3.00, 7.18, 20.00\}$ (7.18 is the estimated value) and the different values we use for θ_f are $\{0.25, 0.39, 0.50, 0.75\}$ (0.39 is the estimated value).

Table 20: Comparison of Auction Outcomes for Different CES Parameter Values

σ_f	θ_f	Price ($000)	Number of Visas	Government Revenue ($Bil)	Firm Surplus ($Bil)	Worker Surplus ($Bil)	Total Surplus ($Bil)
7.18	0.39	15.25	32,313	0.49	1.12	0.82	2.43
7.18	0.25	15.75	32,313	0.51	1.23	0.24	1.98
7.18	0.50	16.10	32,313	0.52	1.11	1.43	3.06
7.18	0.75	24.41	32,313	0.79	1.38	3.06	5.22
3.00	0.39	25.62	32,313	0.83	1.55	1.13	3.50
3.00	0.25	23.45	32,313	0.76	1.57	0.40	2.72
3.00	0.50	28.02	32,313	0.91	1.60	1.84	4.34
3.00	0.75	37.45	32,313	1.21	1.84	3.57	6.62
20.00	0.39	11.20	32,313	0.36	0.95	0.71	2.03
20.00	0.25	12.89	32,313	0.42	1.08	0.20	1.70
20.00	0.50	11.22	32,313	0.36	0.93	1.26	2.56
20.00	0.75	19.19	32,313	0.62	1.21	2.94	4.77

This table compares visa price, government revenue, firm surplus, worker surplus, and total surplus between simulated auctions with different values for the CES parameters for foreign workers. In all cases the mean values over 100 simulations are reported.

Table 20 compares the auction results across the different possible values of the CES parameters. The auction price varies across parameter values, ranging from roughly $11k to $37k per visa. Holding σ_f fixed, increasing θ_f tends to shift surplus toward workers and raises total surplus, consistent with a larger worker share of match surplus. Holding θ_f fixed, lower values of σ_f are associated with higher equilibrium prices and higher government revenue, while higher values of σ_f reduce prices and revenue. Importantly, however, government revenues remain economically meaningful across all cases, and total surplus remains positive and sizeable throughout.

Table 21: Comparison of Individual Firm Auction Outcomes for Different CES Parameter Values

Firm	$\sigma_f = 7.18$				$\sigma_f = 3.00$				$\sigma_f = 20.00$			
	$\theta_f = 0.39$	$\theta_f = 0.25$	$\theta_f = 0.50$	$\theta_f = 0.75$	$\theta_f = 0.39$	$\theta_f = 0.25$	$\theta_f = 0.50$	$\theta_f = 0.75$	$\theta_f = 0.39$	$\theta_f = 0.25$	$\theta_f = 0.50$	$\theta_f = 0.75$
Accenture	3,586	3,383	3,392	2,213	2,685	2,894	2,546	2,557	4,050	3,811	4,070	2,551
Alphabet	20,076	18,818	21,581	24,740	21,009	19,837	21,961	23,454	19,101	18,274	20,773	25,560
Amazon	35,632	41,201	32,154	31,138	31,328	35,775	29,089	27,358	38,125	43,515	33,833	32,374
Apple	15,399	16,156	15,155	15,689	15,867	16,322	15,561	15,061	14,374	15,403	13,931	16,084
Capital One	192	219	158	74	66	114	60	25	353	276	265	127
Cognizant	761	789	660	400	473	550	388	237	927	877	931	541
Ebay	661	663	715	863	712	735	708	712	484	566	598	1,002
Hewlett Packard	910	901	858	773	696	686	707	755	1,174	998	1,140	724
IBM	3,412	2,680	3,583	3,260	4,150	3,713	4,173	5,026	3,388	2,596	3,544	2,304
Infosys	802	778	769	519	567	638	514	354	934	876	964	646
Intel	15,746	17,377	14,406	13,503	16,163	17,522	14,971	12,526	14,159	16,676	12,138	14,135
Meta	13,906	12,534	16,326	21,630	15,944	14,496	17,251	19,519	11,628	11,321	14,434	23,765
Microsoft	26,428	24,933	27,217	26,631	22,943	22,406	23,066	22,604	28,238	25,598	29,832	28,714
Nvidia	4,477	4,585	4,376	4,026	4,572	4,650	4,478	4,067	4,222	4,457	4,104	4,038
Publicis Groupe	87	77	77	52	34	66	24	11	103	103	115	77
Salesforce	3,759	3,727	3,750	3,359	3,281	3,337	3,195	2,951	3,980	3,912	3,962	3,620
Verizon	3,165	3,139	3,197	2,994	3,032	2,995	3,010	2,980	3,451	3,292	3,425	3,001
Walmart	506	499	494	344	595	457	758	1,154	853	621	830	342
Wipro	5,425	5,301	5,322	4,964	4,606	5,034	4,454	3,556	5,572	5,515	5,567	5,225

This table compares the number of visas received by individual firms across simulated auctions with different values for the CES parameters for foreign workers. In all cases the mean values over 100 simulations are reported.

Table 22: Properties of Visa Receivers in Auctions for Different CES Parameter Values

	$\sigma_f = 7.18$				$\sigma_f = 3.00$				$\sigma_f = 20.00$			
	$\theta_f = 0.39$	$\theta_f = 0.25$	$\theta_f = 0.50$	$\theta_f = 0.75$	$\theta_f = 0.39$	$\theta_f = 0.25$	$\theta_f = 0.50$	$\theta_f = 0.75$	$\theta_f = 0.39$	$\theta_f = 0.25$	$\theta_f = 0.50$	$\theta_f = 0.75$
Industry												
Industrial Manufacturing	72,639	74,975	70,763	69,561	77,118	77,963	76,279	73,821	69,037	73,082	65,692	66,355
Information	95,085	88,821	100,619	107,109	94,716	90,375	97,364	101,097	94,792	88,292	102,046	111,519
Occupation												
Professional Services	14,095	13,578	13,726	10,956	11,784	12,360	11,230	9,991	15,425	14,523	15,777	11,961
Computer Occupations	28,187	29,502	29,037	36,359	32,103	31,961	32,948	35,421	23,840	28,126	23,857	37,897
Operations Managers	50,342	49,372	50,623	49,786	50,456	49,890	50,897	51,033	51,855	50,015	52,193	48,661
Other Management	59,947	58,916	59,636	57,243	55,153	55,861	54,458	52,991	62,310	59,613	62,773	58,644
Sales Managers	23,863	22,682	24,113	23,115	24,168	23,379	24,454	25,014	24,109	22,793	24,496	22,142
Degree												
Bachelors	88,908	90,835	89,780	92,631	90,415	89,811	90,897	92,193	87,646	90,764	88,626	93,138
Masters	141,506	139,915	140,303	137,382	139,869	140,911	139,130	137,430	142,489	139,782	141,203	137,049
Country of Origin												
Canada	9,930	9,707	10,166	10,126	9,704	9,438	9,683	9,350	10,271	9,800	10,774	10,219
India	131,689	132,516	130,272	128,369	131,347	131,845	129,530	126,381	130,804	132,036	128,311	129,522

This table compares the attributes of visa receivers across simulated auctions with different values for the CES parameters for foreign workers. In all cases the mean values over 100 simulations are reported.

The firm-level allocations in Table 21 are also stable across CES parameter values. The identity of the largest visa recipients is unchanged, and the cross-firm ranking of visa holdings is very similar across specifications, indicating that the auction continues to allocate visas toward the same set of high-demand firms even when (σ_f, θ_f) are perturbed.

Finally, Table 22 shows that the observable characteristics of visa receivers are broadly similar across parameter values. The implied composition by industry, occupation, education, and country of origin moves somewhat with (σ_f, θ_f), but the magnitudes are modest relative to the levels, and the qualitative patterns remain intact. Overall, these robustness checks suggest that our counterfactual results can generalize to alternative values of the CES parameters.

C Reallocation across Education and Nationality

Figure 13: Distribution Across Education Levels

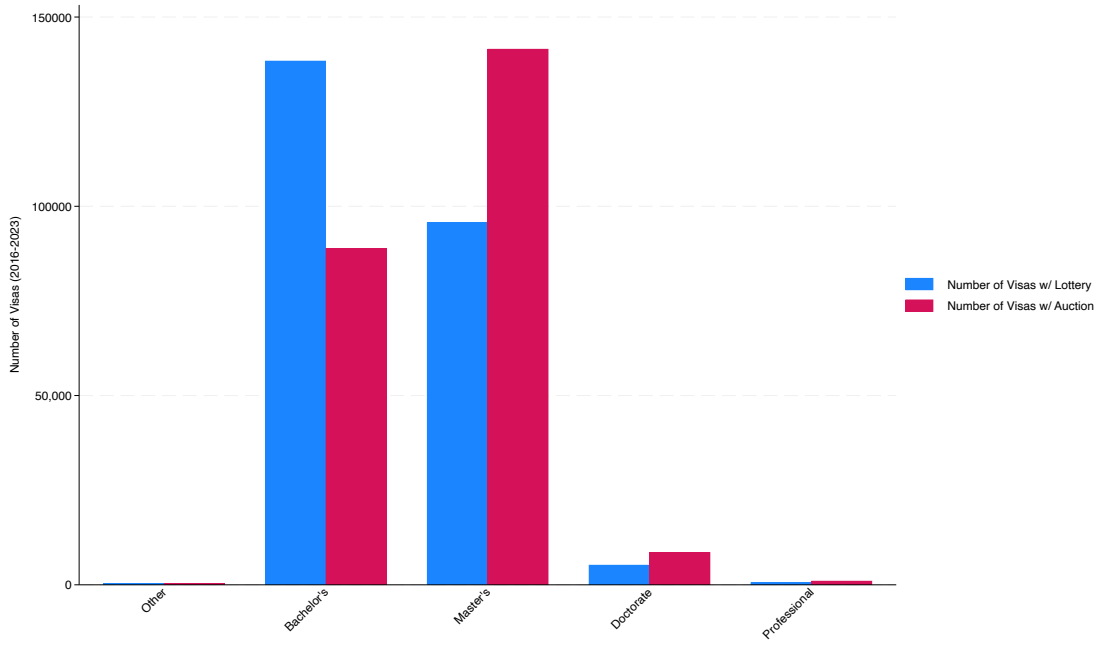


Figure 13 shows that under the auction fewer workers with bachelor's degrees receive H-1B visas, and more are allocated to worker's with master's degrees. Specifically, there is a 35.8% decrease in H-1B visas going to workers with a bachelor's degree and a 48.0% increase in visas going to workers

with a master's degree. This arises because the estimated outside options indicate that master's degree holders face only marginally higher outside options than bachelor's degree holders. At the same time, the relative prevalence of master's degree holders is greater in occupations with higher willingness to pay for visas than in those with lower willingness to pay.

Figure 14: Distribution Across Countries of Origin

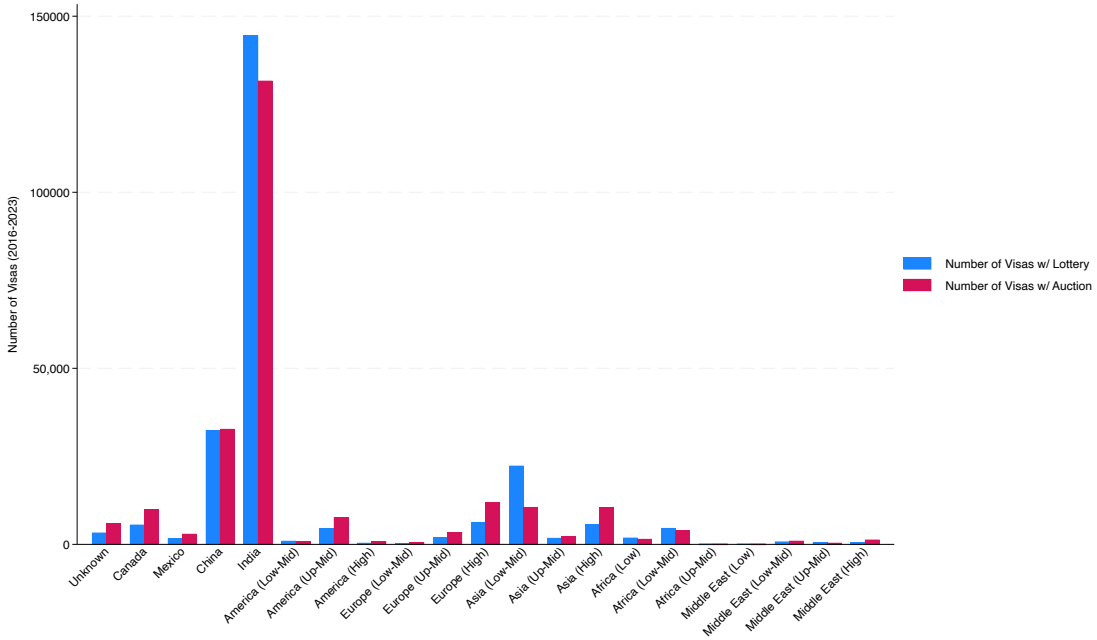


Figure 14 shows how the distribution of visas changes across countries of origin. The overwhelming majority of visas are given to workers coming from India (144,497 under the lottery), and there is a slight decline of 8.9% less visas going to workers from India under the auction. The biggest increase under the auction is in more visas going to H-1B workers from Canada (a 78.4% increase), high-income European countries (a 88.9% increase), and high income Asian countries (a 84.8% increase)