Coping with H-1B Shortages: Firm Performance and Mitigation Strategies

January 19, 2023

Abstract

The H-1B visa program allows companies to hire skilled foreign workers. Before 2014, the vast majority of these visas were allocated on a first-come, first-served basis. Since then, the program has been severely oversubscribed and all cap-subject visas have been allocated through lotteries. We merged Compustat data with administrative firm-level data on the universe of approved petitions for H-1B visas. Using DiD and matching estimators, we estimate that the switch in the visa allocation system negatively affected the growth of companies that used the H-1B program. Our analysis indicates that these effects are quantitatively large and their magnitudes grow over time.

JEL Classifications: F22, J61
Keywords: Migration; Skilled Labor; H-1B

1 Introduction

The United States' H-1B visa program provides an important channel through which firms temporarily hire highly-skilled foreign citizens in specialty occupations, and has been surrounded by much debate since its inception. Supporters argue that the U.S. has high and growing demand for skilled labor, and that H-1B restrictions inhibits innovation and growth. Opponents contend that the program creates unfair competition for American workers. Recently in June 2020, President Trump suspended entry of H-1B visa holders, while 324 large US employers and business-related groups had signed a public letter urging him not to do so.[1]

How do restrictions on H-1B hiring affect firm outcomes? This paper examines how changes in visa allocation, which reduced the ability to hire H-1B workers, affected company performance. We merge the universe of H-1B applications between FYs 1999 and 2018 with Compustat data on a wide range of outcomes for all publicly traded companies in the U.S. Though not a random sample of all private-sector employers, publicly traded companies play a crucial role in determining macroeconomic volatility and growth (Gabaix, 2011) and international trade flows (di Giovanni et al., 2018). We leverage a Difference-in-Differences (DiD) approach to evaluate how company outcomes evolved between H-1B program users and non-users, following changes in visa allocation from a first-come first-served procedure to a random lottery.

Though program details have changed over time, there has always been a cap on the number of new H-1B visas issued to employees of private firms each fiscal year (FY).[2]

---

[1] See Shear and Jordan (2020) or CompeteAmerica (2020). Note also that despite a legal distinction between the terms "visa" and "status", we use the terms interchangeably to refer to the right of foreign citizens to work in the United States, as consistent with popular vernacular. See https://internationalaffairs.uchicago.edu/page/visa-vs-status for further discussion.

[2] Universities, government, and non-profit research institutions are exempt. Since FY 2004, the cap restricted new inflows to 65,000, with an additional 20,000 reserved for individuals with a Master's Degree (or higher) from a U.S. university. In tight labor markets, this constraint becomes more binding, as firms competing to hire foreign labor face smaller windows to apply. U.S. Citizenship and Immigration Services (USCIS) closes the applications window when it receives enough petitions to hit the cap. Figure 1 illustrates this constraint by displaying the number of days that applicants could apply for

Demand for new H-1B workers grew so large in FYs 2014-2020 that the cap was reached in the very first week of each year's application period. Instead of the usual first-come first-served method, USCIS decided to randomly select H-1B applications during these "lottery years".

This change in allocation occurred while the overall number of new visas awarded remained constant since the quota did not change. Lottery selection presumably gave rise to a very different allocation of visas – potentially, some high-surplus employer-employee matches that would have succeeded in the first-come first-served allocation, failed in the lottery. From the employer's viewpoint, the lottery-based allocations reduced the probability of forming the desired match with specific workers, and created uncertainty that might have distorted other investments complementary to the hiring of the foreign skilled worker. We use this rationing as a shock to study how skilled immigration affects firm outcomes.

Our analysis focuses on FYs 2010-2018 and examines a broad set of firm-level indicators including employment, sales, profits, market value, capital expenditures and research and development (R&D) expenditures.[3] We consider FYs 2014-2018 as the period of H-1B rationing, and FYs 2010-2013 as the pre-rationing period. Employing DiD and multivariate-distance matching, we draw comparisons between firms that relied upon the H-1B program as a source of labor during the pre-rationing period and those that did not.

We contribute to a growing body of literature on the economic impact of the H-1B visa program in several ways. Most studies focus on the labor market outcomes of native-born workers, relying on changes to the annual H-1B cap (e.g., Kerr and Lincoln, 2010; Peri et al., 2013; Mayda et al., 2018; Kerr et al., 2015b). Recent studies have utilized H-1B lotteries to identify causal impacts on workers and innovation (e.g., Peri

new H-1B status in each fiscal year.

[3] We note that two earlier lotteries occurred for FYs 2008 and 2009. We focus on the more recent lotteries, which helps avoid the economic disruption of the Great Recession.

et al., 2015a; Doran et al., 2014; Dimmock et al., 2019). Our paper contrasts with prior literature by bringing specific attention to firm behavior and outcomes.

Understanding how workers adjust is incomplete without considering firm-level responses. Immigration-induced changes to native-born wages, employment, or invention, for example, are likely the culmination of various firm-level choices to expand or con- tract hiring, or relocate, alter, or expand operations.[4] This is particularly important in the case of the H-1B program, as firms play a disproportionately large role in selecting, sponsoring, and eventually hiring H-1B workers (Kerr et al., 2015a). Complementing existing studies that have looked at firm employment, we examine a wider set of important outcomes, such as sales, profit, and R&D expenditures.

We also contribute to literature by assembling a novel dataset by matching the universe of I-129 petitions to publicly traded companies (in Compustat) that allows us to track firms over time through 2018. While much of the literature has focused on earlier reductions in the cap (circa FY 2004), this allows us to focus on a more recent period of heightened demand for H-1B visas (FYs 2010-2018).

The remainder of the paper proceeds as follows. Section 2 provides a description of important features of the H-1B visa program and presents our empirical strategy. Section 3 describes the data and presents descriptive statistics. Section 4 presents our Difference-in-Difference estimation results. Section 5 presents our multivariate-distance matching analysis. Section 6 concludes.

[4] For example, Glennon (2020) finds that H-1B restrictions leads multinational firms to offshore more jobs.

2 Background on H-1B Rationing

Each year H-1B cap-subject applications begin on the first business day of April, for workers to begin as early as October 1st of that year. Traditionally, USCIS allocates H-1B status to qualified workers on a first-come, first-served basis. However, when the number of applications exceeds the cap within the first few days of the filing period, random lotteries are held to distribute the visas. As we discuss next, in years with robust demand for H-1B visas, the application window is kept open for only a few days and *all* petitions deemed valid are subject to the lottery. This change in the method of allocation has important consequences when labor markets are tight and aggregate demand for new H-1B workers is high.

Figure 1 plots for each fiscal year the number of days it took to reach the final receipt date – the day on which the number of applications received exceeded the cap. Prior to FY 2008, firms were able to secure their desired number of new H-1B workers so long as they submitted their petitions prior to the final receipt date. However, increased demand for H-1B visas moved the final receipt date closer to the first date of the application period, thereby shortening the application window. In FYs 2008-09, USCIS had received more than enough applications for new H-1B visas within a few days of April 1st. As a result, all new-employment H-1B visas were allocated through a lottery. After a pronounced decline in H-1B demand during FYs 2010-2013, USCIS has continued to distribute *all* new H-1Bs by lottery every year since FY 2014.

We refer to these periods (FYs 2008-09, and every year after 2014) as episodes of *visa rationing* or *lottery years*. From the perspective of individual firms, visa rationing periods entail a sharp decline in their ability to hire the desired H-1B worker, relative to years when visas where allocated on a first-come, first-serve basis. Faced with this increased uncertainty, companies may have postponed or canceled the hiring of other workers or investments in equipment deemed complementary to the skilled foreign workers.

Use of the H-1B program varies substantially across firms. One expects that visa rationing would negatively affect the growth and performance of firms that rely on the H-1B program relative to companies that do not rely on the program. To examine this hypothesis, we classify employers into users and non-users of the program. We consider a firm to be a user of the H-1B program in FY 2014 (the onset of the rationing period) if it received at least one approved petition for an H-1B visa in the years prior to the onset of the rationing period (FYs 2010-2013).

The variation over time in the severity of rationing (i.e. the excess demand for new-employment H-1B visas) and firms' heterogeneous participation in the H-1B program affords implementation of a Difference-in-Difference (DiD) estimation strategy. Our main analysis focuses on FYs 2010-2018, where FY 2014-2018 are considered the treatment period. Users of the program are considered treated units and non-users serve as the control group. We also present estimates based on multivariate-distance matching, where we utilize a subset of non-users that most closely resembles treated companies as the control group.

Finally, we note that H-1B status has a 3-year duration, which can be renewed for an additional 3-year period.[5] Importantly, renewals of the H-1B status are not subject to the annual cap, which only applies to new foreign workers at private-sector firms. As we show, this provides firms with an important margin of adjustment in times of tight labor markets, allowing them to retain existing H-1B workers whose status might not have been renewed in normal times.

[5] H-1B visa holders waiting for permanent residency may hold H-1B status beyond six years.

3 Data Sources and Descriptive Statistics

3.1 USCIS Approved H-1B Petitions: FY 1999-2018

Our data covers two decades of the H-1B program (FYs 1999-2018), with the main empirical analysis focusing on the FY 2010-2018 period, which excludes the Great Recession and important changes to immigration policy that might have affected the demand for H-1B visas.[6] To create a consistent firm-level dataset on approved H-1B applications (more precisely called I-129 petitions) over this period, we rely on two sources of data from USCIS. First, we secured individual records of H-1B applicants from 1999-2012 through a Freedom of Information Act (FOIA) request to USCIS. These data contain information about the prospective employee, the employer (firm), the type of request (e.g. new H-1B, continuing H-1B worker, etc.), and the status of the petition (e.g. ap- proved, denied, etc.). We collapse these data to obtain firm-level counts of the number of approved new and continuing H-1B petitions by fiscal year.

Because our FOIA data end in 2012, we also rely upon publicly available data from the USCIS H-1B Employer Data Hub, which provides firm-level data on I-129 petitions by firm and year from FY 2009 onward. The resulting dataset provides a complete firm-level longitudinal dataset tracking I-129 petitions between FYs 1999 and 2018. It is worth noting that our main analysis could be carried out solely on the basis of the Data Hub petitions. However, our exploration of the mitigation strategies developed by users of the H-1B program relies on data going back to year 2003.

Panels (a) and (b) of Figure A.1 show the aggregate U.S. totals for new and continuing H-1B applications, respectively. The graphs report the 1999-2012 data obtained from the FOIA request in the blue solid line and the 2009-2018 data from the Data Hub

[6] In April 2008 the U.S. government extended the period of Optional Practical Training (OPT) from 12 to 29 months for foreign-nationals on F status who had graduated from a US university with a STEM degree.

source in the red dashed line. Clearly, in the overlapping years the totals do not exactly coincide, but the discrepancy is small and the two series move in lockstep. This provides reassurance that the data are consistent across sources. In the overlapping years (FY 2009-2012), we rely on the FOIA data because of the greater detail provided on each case. The Figures clearly illustrate the reduction, and subsequent recovery, in the demand for H-1B visas due to the Great Recession. Interestingly, from 2013 to 2018 we observe a leveling of new-employment H-1B visas, despite the robust economic growth of that period, possibly reflecting the binding annual cap.

3.2 Compustat

We use Compustat data to measure firm-level outcomes. These data include information on all publicly traded firms between FYs 1999 and 2018. Our sample retains only those firms that have positive employment in each year during this period. The available outcome variables we study are: employment, sales, profits (EBITDA), market value, capital expenditures and R&D expenditure.[7]

The only firm identifying information available in both the H-1B and Compustat datasets is firm name. To tally individual visa petitions (I-129s) at the company level and to merge the resulting data with Compustat, we relied on string matching techniques on company names. Our procedure combined automatic string matching with extensive manual checks focused on the top 3,000 petitioners of I-129s and Compustat companies with complete data on outcomes over our period of interest. For a more detailed discussion on this procedure, see Mayda et al. (2020).

As can be seen in Figure 2 (and Table A.1), which shows the number of new (blue solid line) and continuing (red dashed line) H-1B applications by firms in the Compustat

[7] EBITDA stands for earnings before interest, taxes, depreciation and amortization. It is commonly used as a measure of the profitability of a company's operations, before netting out taxes, debt and capital depreciation.

database, our sample of firms accounts for around 20,000 new H-1B issuances per year until 2010 (that is, about 1 in 4 of the 85,000 annual cap) and peaked at 53,000 in FY 2012 (about two thirds of the annual cap). Since then they have gradually declined to 24,000 in FY 2018 and about 10,000 in FY 2019. Approved petitions for continuing employment increased from FYs 1999-2017, peaking at 125,000, and then dropping to 102,000 in FY 2018. Denials for continuing-employment increased in FYs 2016-17 and surged in FY 2018, where there was also an uptick in new-employment denials.

The raw data also highlights an interesting new trend that will play an important role in our analysis. Since FY 2014, there has been a widening gap between the number of H-1B visas for new employment (which declined) and the number of continuing-employment visas (which increased). Incidentally, this tendency is also noticeable during the brief rationing period FYs 2007-08. As we argue later, the widening gap between the two types of visas reflects a mitigation strategy adopted by users of the H-1B program. The change in visa allocation in FY 2014 entailed a reduction in firms' ability to match with the desired new foreign worker, which led them to substitute toward H-1B workers already in the firm because their visa extensions were not subject to the annual cap. Presumably, the H-1B visas for these workers would not have been renewed if the employers had been able to secure new-employment H-1B visas. Panels (a) and (b) of Figure A.2 show trends in the number of new and continuing H-1B petitions separately for the top four H-1B employers over this period (Infosys, Tata Group, Cognizant and Wipro) in Panel (a), and in the remaining Compustat firms in Panel (b). Clearly, the reduction in new H-1B visas from 2013 onward observed in Figure 2 is largely driven by top-receiving firms. Additionally, both top-receivers as well as the other Compustat firms experienced a large rise in approvals for continuing-employment visas over the period.

3.3 Descriptive Statistics

Our final sample consists of a balanced panel of 1,600 firms with positive employment in each fiscal year from 1999 through 2018.[8] We consider a company as a *user* of the H-1B program in 2014 (the onset of the rationing period) if it had at least one approved H-1B petition between FYs 2010-13. We also measure the *intensity* of use, partitioning users based on whether they received a number of approved petitions for H-1B visas above or below the median among all users. We consider two versions of the intensity measures, one based solely on the number of approved petitions (which we refer to as *level dependence*) and another where we divide by average firm employment in the 2010- 2013 period (*ratio dependence*) – hence, a measure of the proportion of H-1B workers hired in total employment.

Table 1 summarizes firm characteristics of users and non-users of the H-1B program (over the 2010-2013 period). According to our data, slightly more than half of Compustat companies (57%) were users of the program in 2014. Users differ from non-users along several dimensions, which is likely driven by unobserved heterogeneity. For instance, the average employment is 33,300 among users of the program, almost 6 times larger than for the average non-user. Disparities of similar magnitudes are also observed along other outcomes, with the exception of R&D expenditures where users of the H-1B program spend 37 times more than the average non-user. Employment and sales growth, measured by the log change between 2010-13 was much larger (1.4 to 2 times) for users relative to non-users. In comparison, low-intensity and high-intensity users of the H-1B program are much more similar on average in terms of both initial levels and growth rates. Only in terms of employment and R&D do we observe significant differences. Employment among high- intensity users is about half the level as for low-intensity users. In contrast, the average high-intensity user spends almost

[8] Company closings, acquisitions, or start-ups that occur during our time period are therefore excluded. Dimmock et al. (2019) and Glennon (2020) analyze closures in connection to the H-1B program.

double on R&D than the average non-user. We also observe important differences in industry composition.[9] About 43% of companies using the H-1B program are manufacturing firms, compared to only 27% of the non-users. Additionally, non-users are much more likely to be in finance & real estate than users, but much less likely to belong to the computer industry. Naturally, there is also a great deal of within-industry heterogeneity in firm characteristics. Hence, our econometric models will include firm fixed-effects and industry-year interaction terms.

Obviously, approved petitions for H-1B visas vary between users and non-users and by intensity of use. Between 2010 and 2013, the average user of the program received 75 new and 108 continuing-employment approvals (for a total of 183). Because H-1B visas typically last for 3 years, the previous figures imply that the employment share of the stock of H-1B workers in the average company using the program (in 2014) was around 1.1%. Naturally, high-intensity users of the program received, on average, a higher number of approved petitions and the employment share of their H-1B workers was 2.2%.

3.4 Visa rationing and hoarding of new-employment visas

Before turning to our econometric analysis, it is instructive to illustrate the effects of the change in visa allocation (from a first-come, first-serve basis to a lottery assignment) on the behavior of firms using the H-1B program. As we noted earlier, users of the program might anticipate the switch in the visa allocation system by monitoring over time the number of days during which the application window for H-1B status remained open. A plausible conjecture is that users of the H-1B program might have wanted to build a buffer of new-employment visas while these were available on a first-come, first-served basis, so as to be able to maintain the stock of H-1B workers and mitigate the effects of

[9] We aggregate SIC industry codes into 8 mutually exclusive groups: Agriculture, Construction, Manufacturing, Trade, Finance/Real Estate, Services, Computer, and Other. The computer industry is an aggregation of several 4-digit SIC industry codes related to the manufacture of computers and related goods and computer-related services.

rationing on the performance of the company.

In order to investigate this prediction, we compare the actual approved petitions received by users of the program in the years immediately preceding the onset of the rationing period (FY 2014) to the 'normal' flow approvals – a measure we construct that estimates how new and continuing approvals would had evolved if the share of H-1B workers in company employment remained constant over time (details provided in the notes to Figure 3).

The results are presented in Figure 3. In regards to new-employment approvals (top figure), we observe that between FYs 2011 and 2014 there is a large, abnormal increase in the flow of approved petitions. In contrast, the pattern for continuing employment is markedly different (bottom figure). Prior to 2014 it either tracks overall employment in the firm or grows at a lower rate. However, after rationing began in 2014, we observe that continuing-employment approvals outpace overall employment in the firm.

These findings illustrate how users of the H-1B program attempted to mitigate the effects of the rationing period that began in FY 2014. Two years prior, these companies 'hoarded' new-employment visas in order to maintain the total stock of H-1B workers in the company, given that continuing visas are not subject to the annual cap. However, because workers on H-1B visas can typically receive a 3-year continuation visa only once, the effectiveness of this mitigation strategy is limited. Accordingly, it is plausible to expect the (negative) effects of rationing on company outcomes to increase with the length of the rationing period. Our analysis in the next section will test this hypothesis.

4 Difference-in-Difference Estimation

We employ DiD estimation to uncover the effects rationing, due to the change in visa allocation in FY 2014, on the firm-level outcomes. We also consider extensions of our basic model to examine if the effects of rationing vary over time or as a function of the intensity of use of the program.

We expect visa rationing to negatively affect the growth of users of the program. In particular, the new visa allocation system reduces the ability of companies to hire new foreign skilled workers. To the extent that these companies are not able to find suitable replacements domestically, these firms' growth may be negatively affected. In addition, this effect could be magnified if companies also decide to postpone investments that are complementary to the hiring of foreign skilled workers, such as hiring other (skilled or unskilled) workers domestically, undertaking capital investments or setting up new R&D projects.

The starting point of our analysis in this section is the model below, where y_{ijt} is the outcome for firm i in industry j and fiscal year $t \geq 2010$:

$$y_{ijt} = \alpha + \beta Rationing_t \times User_i + \gamma_i + \gamma_{jt} + \varepsilon_{ijt} \qquad (1)$$

where the dependent variable is the log of employment (or some other firm outcome), the term γ_i captures a set of firm-level fixed effects accounting for time-invariant heterogeneity across firms, and the term γ_{jt} is a set of sector-by-year effects, capturing industry-specific shocks over time.[10]

The key interaction term in the estimating equation is the product of an indicator for the *Rationing* period (taking a value of one for FYs 2014-2018) and an indicator for *Users* of the H-1B program (taking a value of one for companies that received at least one approved petition between FYs 2010 and 2013). The coefficient of interest is β, capturing the average difference in log employment around the change in the visa allocation system between users and non-users. The inclusion of firm (and industry-by-year) fixed-effects implies that this coefficient is identified by within-company employment changes among users relative to non-users, after netting out the dynamics specific to each firm's industry. We use two-way clustered standard errors on firm and year to address potential within-

[10] We use 8 mutually exclusive industry groups, aggregated from 4-digit SIC codes: Agriculture, Construction, Manufacturing, Trade, Finance/Real Estate, Services, Computer, and Other. We note that our design is *not* a staggered DiD, given recent innovations in staggered DiD designs.

firm autocorrelation, or also within-year cross-firm correlation in errors.

4.1 The Effects of Rationing on Company Growth

Table 2 presents the estimates corresponding to the model in Equation (1). The top panel considers a company as a user of the H-1B program in FY 2014 if it received at least one approved petition between 2010 and 2013. Across all outcomes the point estimates for β are negative and suggest that rationing reduced company growth in terms of employment, as well as for all other outcomes. The estimates imply that program users grew approximately 4 percent less than non-users between FYs 2014 and 2018, or about 1 percent less per year, even though we cannot reject the null hypothesis at the usual 5% significance level.

The other panels in the table consider more demanding thresholds to qualify as a user of the program, ranging from 2 to 11 approved petitions between FYs 2010 and 2013. The point estimates are somewhat larger (in absolute value) than in the top panel and become statistically significant at the usual significance level. In particular, the estimates suggest that users grew by about 6 percent less (or 1.25% less annually) than non-users in terms of overall employment. The largest impact of rationing appears to be for capital expenditures: growth was about 10 percent lower for users of the program (2.5% annually).

We further clarify the magnitude of the effects implied by our estimates on firm employment. As noted earlier, the switch to the allocation of all cap-subject, new-employment visas via lottery reduced employment among users of the H-1B program by about 6% (column 1 in Table 2) between FYs 2014 and 2018, which amounts to a

1.2% annual reduction. According to our data, at the onset of the rationing episode, the average company using the program had a stock of 182 H-1B workers (column 2 in Table 1), which amounts to roughly 0.6% of the 33,000 workers for the average user of the program. Hence, for each H-1B worker lost to the company, employment fell by about one *additional* (non-H-1B) worker.

There are a number of reasons that explain the magnitudes of the effects of rationing on employment. First, H-1B workers can be complementary to other workers in the firm, consistent with the high share of visas going to employers in computer industries and the findings in Peri et al. (2015b). In addition, skilled workers (foreign or domestic) may be highly complementary to capital investments (Krusell et al., 2000). Evidence supporting this channel can be seen in Table 2. Visa rationing led to large reductions in capital investments (column 5) and R&D (column 6) among users of the H-1B program.[11] In particular, these reductions in capital investment may reflect the outsourcing of activities that would have been undertaken by the visa holders to other domestic or foreign-based companies (Glennon, 2020).

Last, we examine whether the effects of rationing on users of the H-1B program are concentrated on a single industry or widespread across all industries. To do so we construct a sub-sample of companies that account for a large share of H-1B visas (Figure A.2). These companies belong to the 3-code SIC industry *737 - Computer Programming, Data Processing, and other Computer Related Services*. As can be seen in Table A.2, excluding the computer-services industry does not affect our main estimates. Thus, visa rationing negatively affected companies using the H-1B program across a wide range of industries.

Importantly, the models estimated above include industry-trends (along with firm fixed-effects) which account for the differential evolution of industries, which surely affects demand for (foreign skilled) workers. Because of sample size considerations, these

[11] Only about 40% of the companies in our sample report positive R&D expenditures.

trends are defined at the level of 1-digit SIC industries. Table A.4 estimates a version of our model with 2-digit industry trends. As a result, the number of interaction terms increases from 72 (8 industries and 9 years) to 567 (63 industries times 9 years) but the results remain qualitatively unchanged (though statistical significance is lower).

In conclusion, our estimates suggest that visa rationing had a negative effect on users of the H-1B program. These companies grew less in terms of employment, sales, profits and market value than non-users, after netting out industry-specific trajectories in these outcomes. We interpret these findings as evidence that the switch in the visa allocation system, and the corresponding reduction in the ability of employers to hire specific foreign skilled individuals, reduced the *growth* of companies that relied on the H-1B visa program. It is worth noting that this finding does not imply that the *profitability* of these companies (i.e. divided by total assets as in Novy-Marx (2013)) also fell. Additional analysis (available upon request) supports the interpretation that visa rationing only affected the *scale* of the company.

4.2 Dynamic effects

As pointed in Section 3.4, users of the program tried to mitigate the effects of visa rationing on the stock of H-1B workers by 'hoarding' new-employment visas in the years immediately prior to the onset of rationing. Because H-1B visas can only be extended for 3-year periods, the effectiveness of this mitigation strategy should decline over time. To investigate this, we extend the model in Equation (1) by subdividing the rationing period into two parts. Accordingly, the coefficient accompanying the dummy variable for FYs 2014-2016 will pick up the short-run effects of rationing, while the coefficient for FYs 2017-2018 will identify the medium-run effects. On the basis of earlier discussions, we expect rationing to have a larger negative impact in the medium run than in the

short run.

Table 3 presents the estimates. The top panel simply reproduces the baseline estimates from Table 2. The bottom panel presents the estimates for the short and medium run effects of visa rationing. As expected, rationing had only a small effect on company growth in FYs 2014-2016 (about 2.2% in terms of employment over the 2-year period). In contrast, users' employment grew by about 7 percent less over FYs 2017-2018 relative to non-users, and a similar pattern is observed for the other outcomes in the Table.

We can provide even more detail on the evolution of the effects of rationing over time by conducting an event study. Specifically, we now consider the following extension of our basic model:

$$y_{ijt} = \alpha + \sum_{r \neq 2013} \beta_r \left(I_i^{User} \times \mathbf{1}(year_r = t) \right) + \gamma_i + \gamma_{jt} + \varepsilon_{ijt} \qquad (2)$$

where β_r is the difference in outcomes between users and non-users in year r. Note also that the model also includes firm fixed-effects and industry-year trends. Figure 4 plots the estimates, which capture employment trajectories of H-1B users relative to non-users.

Two interesting features are worth noticing. First, employment does not exhibit significant differences between the two groups in the pre-rationing period (FYs 2010-2013). This helps bolster confidence in our DiD estimation strategy. Second, during the rationing period the employment trajectory for users of the program gradually fell below the trajectory for non-users, consistent with the declining effectiveness of 'hoarding' in mitigating the effects of the change in the visa allocation system.

4.3 Intensity of treatment

Next, we investigate whether companies using the H-1B program with higher intensity suffered more than lower-intensity users. Presumably, companies relying more heavily on the H-1B program as a way to hire skilled workers should have suffered a larger blow than companies that depended on the program to a lesser extent (or not at all).

To carry out the analysis we need to take a stance on how to measure intensity of use. Unfortunately, our choices are constrained by data availability. Ideally, one would like to know the share of H-1B workers among a company's skilled workforce, but Compustat only reports overall employment in the firm.

We partition the 1,600 companies in our sample, according to two measures of dependence. The first measure is based on the number of approved petitions received by the company between 2010 and 2013, which we refer to as *level dependence*. The second measure divides this value using average total employment in the company (over 2010-2013), and we refer to it as *ratio dependence*.

The estimates corresponding to *level dependence* are presented in Table 4. The top panel partitions the firms in our sample into three groups: non-users and two groups of users based on whether the number of approved petitions they received (between FYs 2010-13) was above or below the median value among users of the program. The estimates indicate that higher-intensity users of the program were more negatively affected during the rationing period than companies with a lower level of dependence. The same pattern can be seen across all outcomes. The analysis in the bottom panel is based on a finer partition (non-users and 5 groups of users) and largely confirms that higher intensity of use is associated with more negative effects of rationing.

Next, we turn now to the partition based on *ratio dependence*. The results are reported in Table 5. The top panel suggests that, according to this measure, low-intensity users were more negatively affected by visa rationing than high-intensity users.

The estimates based on the finer partition suggest a complicated pattern. Users in the second quintile appear to have suffered larger impacts, whereas the most highly dependent users did not experience any adverse effects from rationing.

The disagreement between the two measures is largely related to firm size. Firms with more H-1B approvals in levels are often (but not always) larger firms overall. When measuring dependence as a ratio, firm size appears in the denominator. Hence, all else equal, larger firms will tend to have smaller ratios. To understand this, appendix Table A.6 provides a tabulation matrix of the two measures of dependence for H-1B users (i.e. those with at least 1 H-1B approval in 2010-13). The columns show the count of firms considered having low/high dependence when measuring H-1B approvals in levels, while the rows show the count of firms categorized as having low/high dependence when measuring H-1B approvals as a proportion of average firm employment.

First, the two measures are positively correlated and do have a significant degree of correspondence along the diagonal. Among the 904 firms considered H-1B users, 601 of them (66%) are categorized identically in both measures of dependence – 301 are categorized as low dependence users in both the level and ratio definitions, and 300 are categorized as high users in both definitions. Hence, the disagreement in the findings comes from the firms in the off-diagonal. To understand how these relate to firm size, we examined whether each of these off-diagonal firms was above or below (large/small) the median average firm employment from 2010-13. All 152 firms considered high users under ratio dependence but low users in level dependence are "small" firms with below median firm employment. All 151 firms considered low users under ratio dependence but high users under level dependence are "large" firms with above median employment.

We clarify which of these two groups drives the disagreement in results by simply running a similar difference-in-difference model that tracks each of the four groups of firms in appendix Table A.6. Results in appendix Table A.7 show that much of the

negative impacts are driven by firms considered high users under level dependence and low users under ratio dependence that creates the disagreement. These firms switch from high users when using level dependence, to low users when using ratio dependence. These firms tend to be larger firms on average.

Why are the adverse effects concentrated among this group of firms that hire many H-1B in levels, but whose H-1B hires are a very small fraction of their employment? One possible explanation pertains to the composition and structure of employment within these firms. Such large firms may have a pyramidal like skill composition, where very highly skilled employees make up a very small fraction of employment and very low skilled workers make up a large fraction of workers. Such highly skilled employees often work at headquarters and are crucial to product development/innovation or operation systems designs that are necessary to support the large number of less skilled workers who perform much of the physical labor. An example from our sample is Walmart, which falls into the category of low ratio dependence but high level dependence. Nearly 50% of Walmart employees only possess a high school diploma. The percent of workers with a Master's degree or above is less than 3%.[12] These very skilled workers likely play a crucial role in optimizing the design and functioning of logistics and infrastructure needed throughout the supply chain. Rationing of very highly skilled H-1B workers may harm these core functions that the rest of the employees depend on. We cannot precisely examine these dynamics as they require granular within-firm data, but believe this to be an interesting potential avenue for future work.

5 Matching estimation

This section adopts an alternative approach to the estimation of the effects of visa rationing: the direct estimation of average treatment effects using a matching estimator,

[12] See https://www.zippia.com/walmart-careers-116506/demographics/.

where each treatment unit is matched to one or more control units.[13] For comparability with our earlier analysis, we perform matching estimation on user thresholds of 1 and 6 H-1B approved petitions over the 2010-2013 period.

Table 6 provides a comparison of the average characteristics of users and non-users for each of the thresholds we consider. Our main outcome of interest is the change in the log of employment in the firm between 2014 and 2018. As shown in Column 1, the average for this variable was 8.4 log points (or 2.1% annually) between FYs 2014 and 2018. Columns 2-4 reveal that employment growth was substantially higher among non-users of the program than among treated firms (by 5.6 log points and 6.6 log points when using the thresholds of 1 and 6 approved petitions, respectively).

The Table also makes clear that users and non-users differ in several dimensions. Obviously, the total approved petitions received by users of the program is much higher than for non-users. Depending on the threshold we use, users received between 182 and 273 approved petitions between 2010 and 2013, whereas the average non-user received fewer than 1. More interestingly, users of the program in 2014 are much more likely to have received H-1B visas in the past. While 15 to 28% of the non-users in 2014 had participated in the program between 2003 and 2006, the corresponding rates were 81 to 92% for users in 2014. In addition, users of the H-1B program were much more likely to belong to computer-related industries, grew less in the past (2007-2010), and also differed somewhat in terms of previous employment, profitability and R&D intensity.

Given these differences, we adopt a kernel-based multivariate-distance matching estimator to help better align a control group of non-users with the group of users. The vector of characteristics for each treated unit is compared (using a multivariate distance metric) to the vector of characteristics for all non-treated units in the following manner.

[13] In light of recent criticism to propensity score matching (King and Nielsen (2019)), we adopt a multivariate-distance matching estimator, which uses information more efficiently than the typical uses of propensity score matching. We use the Stata implementation by Jann (2017).

Each treated unit is matched solely to non-treated units within the *same* 2-digit industry. Among these candidate matches, those at a shorter distance to the treated unit are given higher weight (on the basis of the Epanechnikov kernel).[14]

Table 7 presents the estimates. Let us first consider the top panel, where firms are considered as belonging to the treatment group as long as they received at least one approved petition for H-1B visas between 2010 and 2013. Column 1 presents a simple version of our estimator, where we partition all firms into two groups: those in the computer industry and those in other industries. We match treated units to non-treated units within each of these two blocks. The estimated average treatment effect on the treated (ATET) suggests that employment for treated units grew 3.1 log points less than for non-treated units, although we cannot reject the null hypothesis.

We perform a more systematic block matching where each treated unit is matched only within its 2-digit industry in column 2. The estimated ATET is similar, but the standard errors are now much smaller than in column 1. Column 3 restricts the matched observations to the 5 nearest neighbors (rather than applying the kernel as in all other columns in the table), which delivers similar results. Column 4 presents our preferred specification, which again conducts kernel-based matching and bootstraps standard errors (within 2-digit industries). The estimated ATET is -0.0382.

We now turn to the second panel, where the threshold to be considered a user of the program is now 6 approved petitions. Our preferred estimated ATET (column 4) indicates that the change in the allocation system of visas in 2014 reduced the employment growth of users significantly, by approximately 6.3% between FYs 2014 and 2018. Furthermore, columns 5 and 6 show that the effects of rationing were more severe in 2016-2018 than in 2014-2016, in line with our earlier findings. Namely, the effectiveness of the mitigation strategy adopted by users of the program (hoarding new-

[14] Table A.5 reports balancing tests that illustrate the reduction achieved by our matching algorithm in the gaps between the means and variances of the treated and untreated units. Estimation of a probit model for current use in the H-1B program using the same individual characteristics shows that the main predictor of current use of the program is previous use, which has been a commonly used predictor in the literature (see Kerr and Lincoln (2010)).

employment visas) declined over time.

In conclusion, the estimated ATET just discussed are closely aligned with the difference-in-difference estimates reported in Table 2, confirming the negative impact of visa rationing on firms using the H-1B program.[15]

6 Conclusions

The H-1B program provides an important channel to hire skilled foreign labor. Though there has been a cap on the number of new H-1Bs issued to private-sector employers since the inception of the program, that cap has been particularly binding since 2014. *All* cap-subject H-1B visas have been allocated by lottery since that year. The resulting rationing has greatly reduced employers' ability to hire a specific foreign skilled worker and altered the allocation of new visas. In comparison to previous years, when most new-employment visas (at for-profit firms) were allocated on a first-come, first-served basis, the rationing of visas through lotteries is likely to have prevented an important share of high-surplus employer-employee matches from taking place.

Our estimates point to large effects that were probably magnified by strong complementarities between (foreign) skilled workers, domestic workers and capital investments. We estimate that visa rationing lowered employment for users of the H-1B program by an average of 1.2% annually for the duration of the rationing. In the years immediately prior to the onset of the rationing period, users of the program were able to *hoard* new-employment visas and this allowed them to maintain the stock of H-1B workers in the company for a few years. However, our analysis shows that the effectiveness of this mitigation strategy declined rapidly, leading to large negative effects on firm growth in terms of employment, sales and profits.

[15] Additionally, we also performed our main difference-in-differences specification Equation (1) using inverse propensity score re-weighting. The results are reported in Table A.3. Both qualitatively and quantitatively, the estimates are highly consistent with those reported in Table 2.

We demonstrate that the change in allocation process led to rationing of visas which negatively impacted firms that use the program. Furthermore, allocating visas through a lottery system introduces other inefficiencies (Sharma and Sparber, 2021). Other methods of allocation could be more efficient, such as auctioning visas (Peri, 2012), using a points-based system (as is done in Canada), or using a lottery that weights by some scoring system (e.g. based on firm productivity). These methods should retain an important role for employers in identifying the foreign workers possessing the skills needed in their companies (Kerr et al., 2015a). Additionally, flexibility in the annual H-1B cap, for example by considering labor market tightness in specific occupations, would help alleviate negative impacts of rationing during periods of heightened demand.

References

CompeteAmerica, "324 Signatory Business Letter to President and Secretaries-Nonimmigrant Visa Facts," May 2020.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Mejean, "The Micro Origins of International Business-Cycle Comovement," *American Economic Review*, January 2018, *108* (1), 82–108.

Dimmock, Stephen G., Jiekun Huang, and Scott J. Weisbenner, "Give Me Your Tired, Your Poor, Your High-Skilled Labor: H-1B Lottery Outcomes and Entrepreneurial Success," NBER Working Papers 26392, National Bureau of Economic Research, Inc October 2019.

Doran, Kirk, Alexander Gelber, and Adam Isen, "The Effect of High-Skilled Immigration on Patenting and Employment: Evidence from H-1B Visa Lotteries," NBER Working Papers 20668, National Bureau of Economic Research, Inc November 2014.

Gabaix, Xavier, "The Granular Origins of Aggregate Fluctuations," *Econometrica*, May 2011, *79* (3), 733–772.

Glennon, Britta, "How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program," Technical Report 27538, NBER 2020.

Jann, Ben, "KMATCH: Stata module for multivariate-distance and propensity- score matching, including entropy balancing, inverse probability weighting, (coarsened) exact matching, and regression adjustment," Statistical Software Components, Boston College Department of Economics May 2017.

Kerr, Sari Pekkala, William R Kerr, and William F Lincoln, "Firms and the economics of skilled immigration," *Innovation Policy and the Economy*, 2015, *15* (1), 115–152.

_ , _ , and _ , "Skilled immigration and the employment structures of US firms," *Journal of Labor Economics*, 2015, *33* (S1), S147–S186.

Kerr, William R. and William F. Lincoln, "The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention," *Journal of Labor Economics*, 07 2010, *28* (3), 473–508.

King, Gary and Richard Nielsen, "Why Propensity Scores Should Not Be Used for Matching," *Political Analysis*, 2019 2019, *27* (4), 435–454.

Krusell, Per, Lee E. Ohanian, Josè-Victor Rìos-Rull, and Giovanni L. Violante, "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, September 2000, *68* (5), 1029–1054.

Mayda, Anna M., Francesc Ortega, Giovanni Peri, Kevin Shih, and Chad Sparber, "New Data and Facts on H1-B Workers across Firms," in "The Roles of Immigrants and Foreign Students in US Science, Innovation, and Entrepreneurship," in I. Ganguli, S. Kahn and M. MacGarvie, National Bureau of Economic Analysis, 2020.

Mayda, Anna Maria, Francesc Ortega, Giovanni Peri, Kevin Shih, and Chad Sparber, "The effect of the H-1B quota on the employment and selection of foreign-born labor," *European Economic Review*, 2018, *108* (C), 105–128.

Novy-Marx, Robert, "The other side of value: The gross profitability premium," *Journal of Financial Economics*, 2013, *108* (1), 1–28.

Peri, Giovanni, "Rationalizing U.S. Immigration Policy: Reforms for Simplicity, Fairness, and Economic Growth," Discussion Papers 2012-01, The Hamilton Project May 2012.

_ , Kevin Shih, and Chad Sparber, "STEM Workers, H1B Visas and Productivity in US Cities," Norface Discussion Paper Series, Norface Research Programme on Migration, Department of Economics, University College London 2013009, Norface Research Programme on Migration, Department of Economics, University College London March 2013.

_ , _, and _, "Foreign and Native Skilled Workers: What Can We Learn from H-1B Lotteries?," NBER Working Papers 21175, National Bureau of Economic Research, Inc May 2015.

_ , _, and _, "STEM Workers, H-1B Visas, and Productivity in US Cities," *Journal of Labor Economics*, 2015, *33* (S1), 225–255.

Sharma, Rishi and Chad Sparber, "Buying Lottery Tickets for Foreign Workers: Lost Quota Rents Induced by H-1B Policy," Technical Report 2021.

Shear, Michael D and Miriam Jordan, "Trump Suspends Visas Allowing Hundreds of Thousands of Foreigners to Work in the U.S.," *The New York Times*, Jun 2020.

Figure 1: Number of Days Until Final Receipt Day for H-1B Petitions

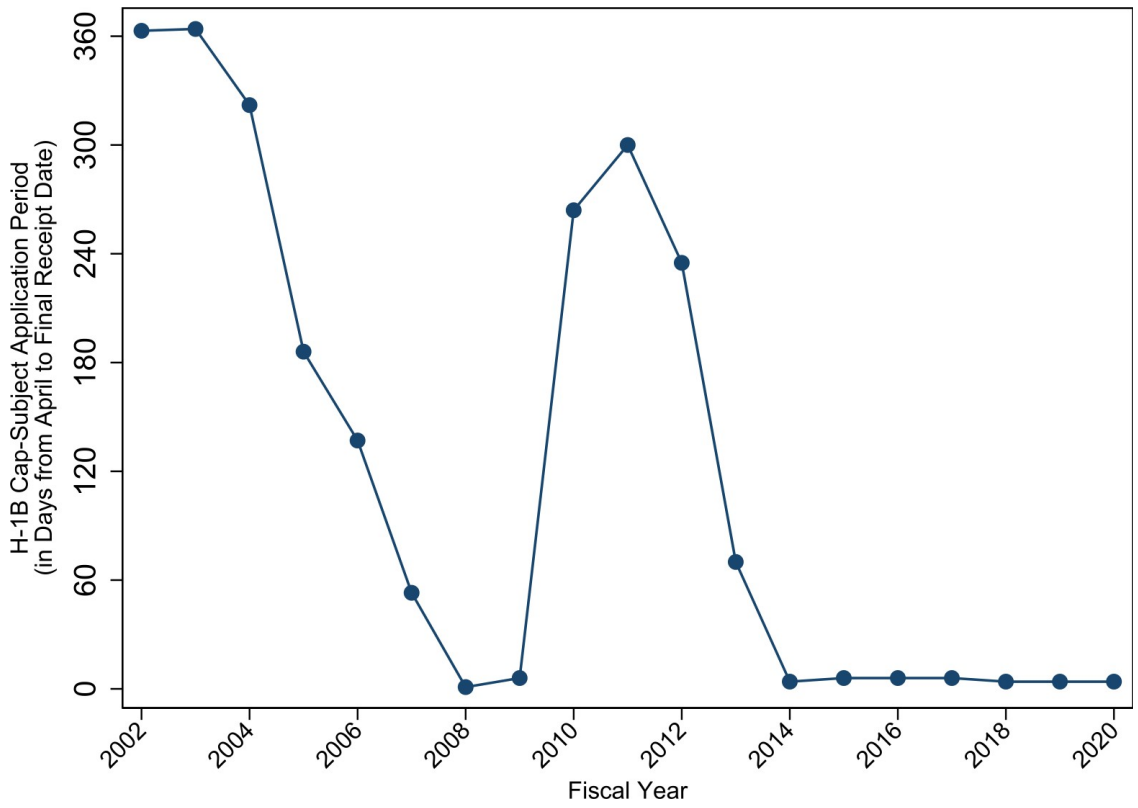

Notes: Number of days between the first week of April (when application window opens) and final receipt day (when USCIS has received enough applications to meet the statutory cap). Zero values correspond to years when USCIS allocated all cap-bound H-1Bs by lottery. In other years, USCIS allocated H-1Bs on a first-come, first-serve basis and used a lottery only for selecting applications received on the last date of receipt.

Figure 2: I129 data on Approved H-1B Petitions for New and Continuing Employment at Compustat Firms.

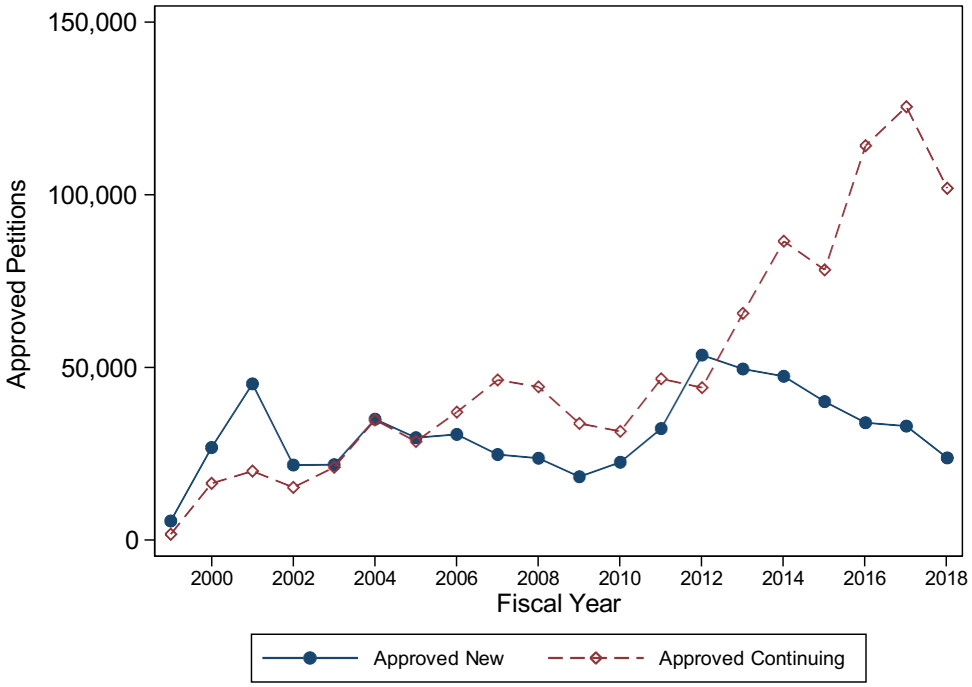

Note: Figures show aggregate totals of approved new and continuing H-1B petitions for the sample of Compustat firms, by fiscal year. Blue series shows approved new H-1B petitions, while the red dashed series shows approved continuing H-1B petitions.

Figure 3: Hoarding

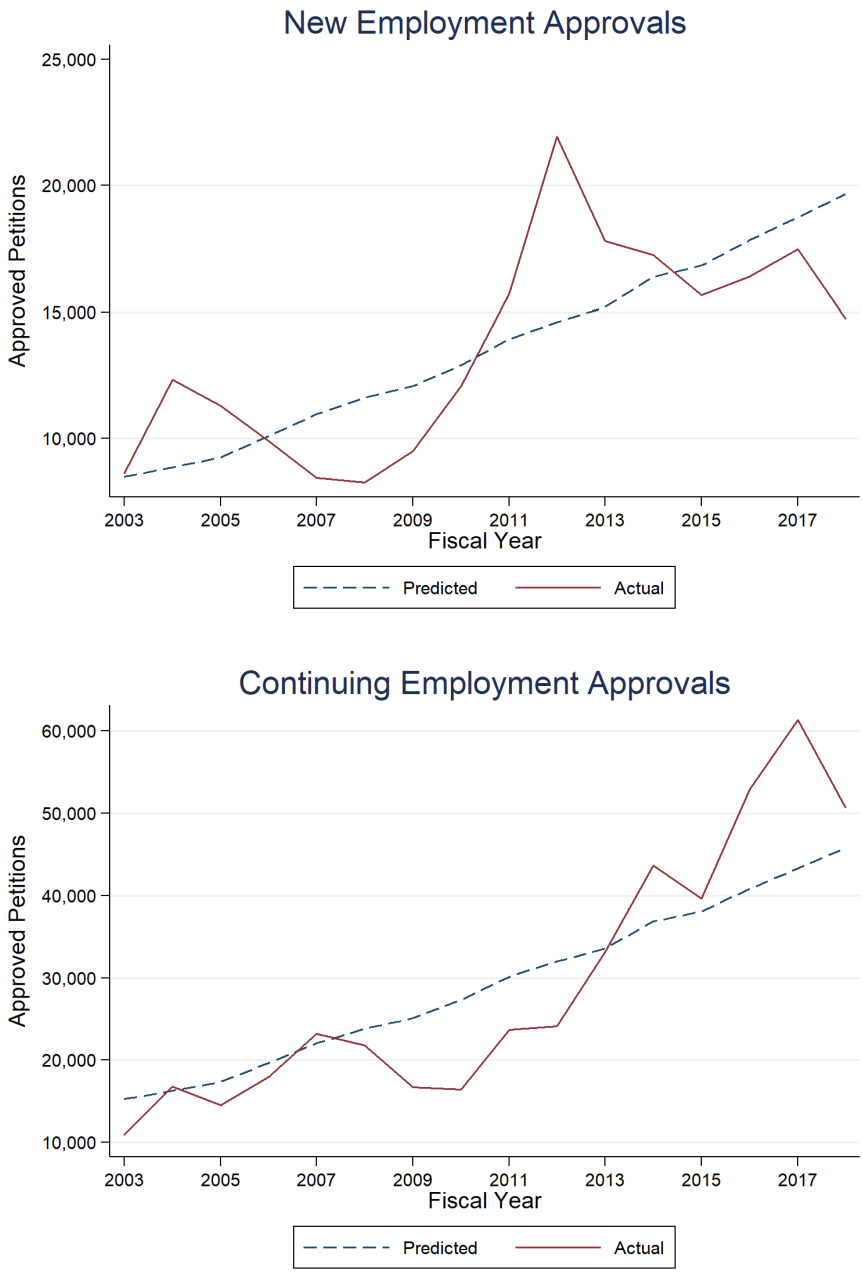

Notes: Figures compare total H-1B approvals (solid line) with *predicted* H-1B approvals had the share of H-1B in company employment remained constant over time (dashed line). We predict annual approved petitions as follows. First, we calculate the share of H-1B workers in company employment by using the sum of new and continuing approved petitions received by the company between 2003 and 2005 and dividing by average employment over the same period. Next, we predict the total stock of H-1B workers in each year by interacting employment in each year with this fixed share of H-1B workers in employment (circa 2005). Last, we create predictors for the annual *flow* of new and continuing-employment visas by regressing, for H-1B users, new (or continuing) company approvals on its predicted total stock. We then sum the predicted values and actual values over all users for each year. The top figure displays these summed values for new approvals, while the bottom figure is for continuing approvals.

Figure 4: Event Study Comparing Total Employment at Firms that Employ H-1B Workers to those that Do not.

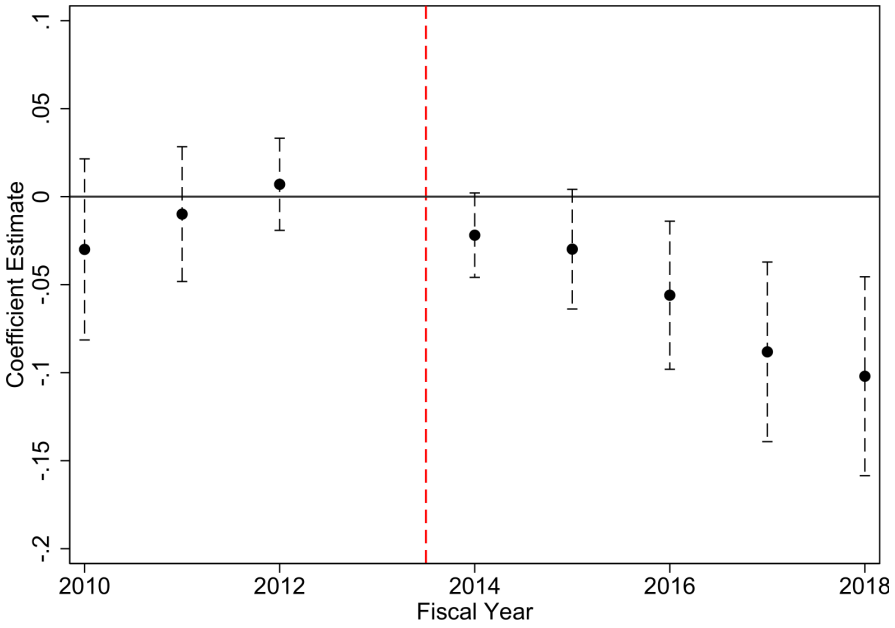

Notes: Dependent variable is the log of employment. Model includes company fixed-effects and industry-year fixed-effects. Rationing period covers fiscal years 2014 through 2018. Point-estimates are relative to FY 2013, the omitted category. 95% Confidence intervals are shown in dashed lines.

Table 1: Firm Characteristics by Intensity of Firms' H-1B Employment. Average 2010-2013 values

	(1) Non-Users	(2) H-1B Users	(3) Low H1B/Emp	(4) High H1B/Emp
Compustat Firm Descriptives:				
Employment (thousands)	5.8	33.3	45.1	21.5
$\Delta Ln(Emp)_{2010-13}$	0.06	0.12	0.11	0.12
Sales ($ million)	1,831.3	12,335.4	13,311.3	11,359.5
$\Delta Ln(Sale)_{2010-13}$	0.12	0.17	0.17	0.17
Profits - EBITDA ($ million)	523.4	4,532.2	4,388.8	4,675.6
Market Value ($ million)	1,432.1	12,202.7	10,581.8	13,823.6
Capital expenditures ($ million)	162.1	794.0	794.5	793.5
R&D expenditures ($ million)	7.4	273.5	184.9	362.0
Agriculture and Mining (%)	4.9	4.9	3.5	6.2
Construction (%)	1.9	1.0	1.6	0.4
Manufacturing (%)	27.3	43.5	43.6	43.4
Trade & Transportation (%)	20.3	20.5	30.5	10.4
Finance & Real Estate (%)	36.8	13.3	10.0	16.6
Business Services (%)	6.5	10.5	8.0	13.1
Other Industries (%)	0.1	0.7	0.7	0.7
Computer Industry (%)	2.3	5.8	2.2	9.3
USA Headquarters (%)	91.4	91.7	89.6	93.8
H-1B Approvals:				
New Emp	0	74.7	7.2	142.2
Continuing Emp	0	107.7	17.8	197.5
Total (New + Cont)	0	182.4	25.0	339.7
Total / Avg Emp (%)	0	1.1	0.08	2.2
Number of Firms	696	904	452	452
Percent of all Firms	43%	57%	28.3%	28.3%

Notes: *H-1B users* are defined as firms that received at least one H-1B approved petition between 2010 and 2013. Columns 3 and 4 report values for firms below (Low) and above (High) ratio of total approved petitions (combining new-employment and continuation) in 2010-2013 over average employment in the same period. The computer industry includes both manufacturing of computers and office equipment (SIC 357, 360 and 5045) and computer-related services (7370, 7373 and 7374).

Table 2: The Effect of Rationing on Company Outcomes, DiD Estimates

Dep. Var. *Ln* of:	(1) Employment	(2) Sales	(3) Profits	(4) Market Value	(5) Capital Expenditures	(6) R&D Spending
User threshold= 1						
Rationing × User	-0.043	-0.053	-0.053	-0.101*	-0.083	-0.083
	[0.028]	[0.031]	[0.031]	[0.045]	[0.046]	[0.070]
Observations	14,400	14,301	12,793	12,856	13,519	5,414
Avg. Approvals User	182	182	182	182	182	182
User threshold= 2						
Rationing × User	-0.069**	-0.077**	-0.079**	-0.100**	-0.127**	-0.096
	[0.029]	[0.030]	[0.030]	[0.038]	[0.045]	[0.061]
Observations	14,400	14,301	12,793	12,856	13,519	5,414
User threshold= 6						
Rationing × User	-0.060*	-0.062*	-0.106**	-0.089**	-0.120**	-0.114*
	[0.026]	[0.029]	[0.033]	[0.038]	[0.041]	[0.053]
Observations	14,400	14,301	12,793	12,856	13,519	5,414
User threshold= 11						
Rationing × User	-0.067**	-0.081**	-0.113**	-0.086**	-0.137***	-0.142**
	[0.024]	[0.027]	[0.035]	[0.037]	[0.036]	[0.049]
Observations	14,400	14,301	12,793	12,856	13,519	5,414

Notes: Results from a regression model that classifies companies as users or non-users of the H-1B program, with varying thresholds to consider a company as being a user of the H-1B program. Users are companies with approved petitions equal to or above the threshold. Non-users are companies with values below the threshold. The thresholds we consider are: 1, 2, 6 and 11 total approved petitions over the period 2010-2013 (combined). The sample period is 2010-2018 and the *Rationing* indicator takes a value of one during years 2014-2018. The dependent variables are the logs of the corresponding outcome. Firm fixed-effects (defined by gvkey, the unique firm identifier in Compustat) and industry-year fixed-effects included in all specifications. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: The Effect of Rationing on Company Outcomes, Short and Medium run effects

Dep. Var. *Ln* of:	(1) Employment	(2) Sales	(3) Profits	(4) Market Value	(5) Capital Expenditures	(6) R&D Spending
Rationing 2014-18						
Rationing × User	-0.043	-0.053	-0.053	-0.101*	-0.083	-0.083
	[0.028]	[0.031]	[0.031]	[0.045]	[0.046]	[0.070]
Sub-periods						
Rationing 2014-16 × User	-0.022	-0.032	-0.042	-0.068	-0.059	-0.056
	(0.024)	(0.026)	(0.030)	(0.042)	(0.044)	(0.059)
Rationing 2017-18 × User	-0.073**	-0.083**	-0.068*	-0.149**	-0.121*	-0.124
	(0.029)	(0.035)	(0.036)	(0.048)	(0.054)	(0.086)
Observations	14,400	14,301	12,809	12,865	13,520	5,418
Avg. Approvals Non-user	0	0	0	0	0	0
Avg. Approvals User	182	182	182	182	182	182

Notes: Results from a regression model that classifies companies as users or non-users of the H-1B program, using 1 approval in 2010-2013 as the threshold to be considered a user of the program. The sample period is 2010-2018 and the *Rationing* indicator takes a value of one during years 2014-2018. The bottom panel considers two sub-periods: 2014-2016 (short run) and 2017-2018 (medium run). The dependent variables are the logs of the corresponding outcome. Firm fixed-effects (defined by gvkey, the unique firm identifier in Compustat) and industry-year fixed-effects included in all specifications. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Standard errors clustered at the firm level. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Intensity of Dependence and the Effect of Rationing, Level Dependence

Dep. Var. *Ln* of:	(1) Employment	(2) Sales	(3) Profits	(4) Market Value	(5) Capital Expenditures	(6) R&D Spending
3 groups: High Users, Low Users, Non-Users						
Rationing × Low User	-0.017	-0.020	-0.003	-0.092*	-0.022	-0.031
	(0.032)	(0.035)	(0.041)	(0.050)	(0.055)	(0.080)
Rationing × High User	-0.069**	-0.086***	-0.098***	-0.111***	-0.146***	-0.129*
	(0.027)	(0.033)	(0.036)	(0.042)	(0.047)	(0.071)
Observations	14,400	14,301	12,809	12,865	13,520	5,418
Avg. Dependence (Low)	4.59	4.59	4.59	4.59	4.59	4.59
Avg. Dependence (High)	361	361	361	361	361	361
6 groups: Dependence Quintiles, & Non-Users						
Rationing × Users Q1	0.024	0.008	0.055	-0.075	0.012	0.001
	[0.040]	[0.035]	[0.055]	[0.071]	[0.063]	[0.098]
Rationing × Users Q2	-0.073	-0.075	-0.048	-0.118	-0.055	-0.014
	[0.049]	[0.052]	[0.049]	[0.069]	[0.072]	[0.087]
Rationing × Users Q3	-0.045	-0.040	-0.071	-0.094	-0.105	-0.131
	[0.038]	[0.040]	[0.045]	[0.058]	[0.059]	[0.084]
Rationing × Users Q4	-0.061	-0.089*	-0.102*	-0.136**	-0.122*	-0.096
	[0.033]	[0.042]	[0.045]	[0.057]	[0.057]	[0.077]
Rationing × Users Q5	-0.091**	-0.097**	-0.121**	-0.098*	-0.174**	-0.155*
	[0.037]	[0.039]	[0.046]	[0.052]	[0.054]	[0.080]
Observations	14,400	14,301	12,809	12,865	13,520	5,418
Avg Dependence (Q1)	1.8	1.8	1.8	1.8	1.8	1.8
Avg Dependence (Q2)	5.64	5.64	5.64	5.64	5.64	5.64
Avg Dependence (Q3)	14.8	14.8	14.8	14.8	14.8	14.8
Avg Dependence (Q4)	42.2	42.2	42.2	42.2	42.2	42.2
Avg Dependence (Q5)	854	854	854	854	854	854

Notes: Results from a regression model that classifies companies as users or non-users of the H-1B program, but we classify users according to their dependence on the program. The *Level Dependence* is based on the total number of approved H-1B petitions for the period 2010-2013. The sample period is 2010-2018 and the *Rationing* indicator takes a value of one during years 2014-2018. The top panel splits the sample between non-users and two groups of users, depending on whether the number of approved petitions received was below (Low User) or above (High User) the median level among users. The bottom panel presents results based on a finer partition of users (quintiles). The omitted category is always non-users. Firm fixed-effects (defined by gvkey, the unique firm identifier in Compustat) and industry-year fixed-effects included in all specifications. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Intensity of Dependence and the Effect of Rationing. Ratio dependence

Dep. Var. *Ln* of:	(1) Employment	(2) Sales	(3) Profits	(4) Market Value	(5) Capital Expenditures	(6) R&D Spending
2 Groups:						
Rationing × Low User	-0.073*	-0.094**	-0.060	-0.143**	-0.087	-0.148*
	[0.032]	[0.037]	[0.033]	[0.055]	[0.054]	[0.075]
Rationing × High User	-0.012	-0.011	-0.045	-0.060	-0.080	-0.035
	[0.030]	[0.032]	[0.038]	[0.047]	[0.054]	[0.074]
Observations	14,400	14,300	12,793	12,856	13,519	5,414
Avg. Dependence (Low)	0.08	0.08	0.08	0.08	0.08	0.08
Avg. Dependence (High)	2.22	2.22	2.22	2.22	2.22	2.22
5 Groups (quintiles)						
Rationing × User Q1	-0.052	-0.076*	-0.043	-0.101	-0.088	-0.228**
	[0.041]	[0.040]	[0.034]	[0.061]	[0.058]	[0.090]
Rationing × User Q2	-0.111**	-0.117**	-0.078	-0.167*	-0.096	-0.102
	[0.041]	[0.047]	[0.048]	[0.074]	[0.072]	[0.080]
Rationing × User Q3	-0.057	-0.081*	-0.088*	-0.177**	-0.118*	-0.091
	[0.033]	[0.038]	[0.042]	[0.054]	[0.052]	[0.085]
Rationing × User Q4	-0.058	-0.086	-0.108*	-0.091	-0.131*	-0.125
	[0.041]	[0.047]	[0.053]	[0.067]	[0.069]	[0.085]
Rationing × User Q5	0.075	0.112*	0.086	0.041	0.028	0.038
	[0.047]	[0.051]	[0.058]	[0.063]	[0.076]	[0.089]
Observations	14,400	14,301	12,809	12,865	13,520	5,418
Avg Dependence (Q1)	0.03	0.03	0.03	0.03	0.03	0.03
Avg Dependence (Q2)	0.09	0.09	0.09	0.09	0.09	0.09
Avg Dependence (Q3)	0.20	0.20	0.20	0.20	0.20	0.20
Avg Dependence (Q4)	0.56	.56	.56	.56	.56	.56
Avg Dependence (Q5)	4.89	4.89	4.89	4.89	4.89	4.89

Notes: Results from a regression model that classifies companies as users or non-users of the H-1B program, but we classify users according to their dependence on the program. The *Ratio Dependence* is based on the total number of approved H-1B petitions for the period 2010-2013 divided by company size (as measured by average employment in 2010-13). The sample period is 2010-2018 and the *Rationing* indicator takes a value of one during years 2014-2018. The top panel splits the sample into non-users and two groups of users, depending on whether their ratio dependence is below (Low User) or above (High User) the median value among users. The bottom panel presents results based on a 6-group partition (non-users and quintiles among users). The omitted category is always non-users. The dependent variables are the logs of the corresponding outcome. Firm fixed-effects (defined by gvkey, the unique firm identifier in Compustat) and industry-year fixed-effects included in all specifications. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Comparison characteristics of users and non-users of the H-1B program in 2014

	None All (1600)	Threshold = 1 Non-Users (696)	Threshold = 1 Users (904)	Threshold = 6 Non-Users (999)	Threshold = 6 Users (601)
$\Delta Ln(Emp)_{2014-18}$	0.084	0.116	0.060	0.109	0.043
User dummies					
User 2003-06	0.520	0.148	0.806	0.277	0.923
User 2010-13 (th1)	0.565	0	1	0.303	1
User 2010-13 (th6)	0.376	0	0.665	0	1
Approvals 2010-13					
New Employment	42.216	0	74.719	0.264	111.950
Continuation	60.839	0	107.680	0.464	161.196
Total	103.056	0	182.399	0.729	273.146
Computer Industry	0.043	0.023	0.058	0.026	0.070
$\Delta Ln(Emp)_{2007-10}$	0.006	0.032	-0.014	0.033	-0.040
$Ln(Emp)_{2010}$	0.963	-0.225	1.878	0.037	2.503
Profits/Sales 2010	-0.256	-0.580	-0.008	-0.417	0.009
R&D/Sales 2010	0.283	0.423	0.177	0.337	0.194

Notes: Total Approvals 2010-13 is the combined approved petitions received between 2010 and 2013. $\Delta Ln(Emp)_{2014-18}$ is the 2014-2018 change in log employment. *User 2003-06* takes a value of 1 if the firm received at least one approved petition between 2003 and 2006. Similarly, *User 2010-13* is an indicator for receiving approvals above the threshold of 1 ($th1$) or 6 ($th6$). $\Delta Ln(Emp)_{2007-10}$ is the 2007-2010 change in log employment. Profits are measured by EBITDA.

Table 7: Matching estimation. Average Treatment Effects on the Treated (ATET)

	(1)	(2)	(3)	(4)	(5)	(6)
Exact Match	Computer ind.	SIC2	SIC2	SIC2	SIC2	SIC2
Matching Algo	Epan	Epan	NN5	Epan	Epan	Epan
Dep. var. $\Delta Ln(Emp)$	2014-18	2014-18	2014-18	2014-18	2014-16	2016-18
User threshold= 1						
ATET	-0.0315	-0.0382	-0.0389	-0.0382	-0.0181	-0.0201
	[0.0556]	[0.0265]	[0.0336]	[0.0251]	[0.0182]	[0.0157]
User threshold= 6						
ATET	-0.0573***	-0.0630**	-0.0730**	-0.0630**	-0.0204	-0.0426***
	[0.0200]	[0.0255]	[0.0347]	[0.0246]	[0.0197]	[0.0159]
Observations	1,591	1,591	1,591	1,591	1,591	1,591
Standard errors	Analytical	Analytical	Bootst.	Bootst.	Bootst.	Bootst.

Notes: Table shows results from matching estimation with the dependent variable being the change in log employment over 2014-18 (cols 1-4), over 2014-16 (col 5), and over 2016-18 (col 6). Multivariate-distance matching estimator based on the following variables: an indicator for having received H-1B visas between 2003 and 2006, 2-digit SIC industry codes, the log of employment in 2010, profitability in 2010, R&D expenditures in 2010 (as a share of sales) and employment growth between 2007 and 2010. The matching algorithm applies the Epanechnikov kernel to the control units (except for column 3 where we use the 5 nearest neighbors). In the selection of the control units we require an exact match for the 2-digit SIC industry code (with the exception of column 1 where we only require exact matching for the computer industry). Balancing statistics are reported in Table A.5. In columns 1-2, we report analytical standard errors (derived assuming independent outcomes across units). In columns 3-6 Standard errors are bootstrapped, using 2-digit SIC industry clusters. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix

A Tables and Figures

Table A.1: Total I-1219 Petitions for Employment at Compustat Companies

Fiscal Year	Approved New	Approved Cont	Denied New	Denied Cont
1999	5,424	1,587	307	77
2000	26,687	16,314	2,575	865
2001	45,118	19,872	4,639	1,676
2002	21,622	15,141	5,170	2,119
2003	21,720	20,977	3,014	1,481
2004	34,891	34,696	2,423	1,599
2005	29,539	28,429	4,305	2,749
2006	30,483	36,935	5,631	4,393
2007	24,658	46,243	4,264	3,578
2008	23,593	44,259	5,203	6,119
2009	18,233	33,658	11,530	9,807
2010	22,425	31,365	6,664	7,085
2011	32,175	46,585	7,848	8,359
2012	53,442	44,023	10,670	9,871
2013	49,428	65,540	2,497	1,127
2014	47,305	86,420	2,462	1,862
2015	39,976	78,107	1,421	1,556
2016	33,868	114,061	1,159	3,686
2017	32,884	125,341	2,056	4,010
2018	23,684	101,728	4,256	12,563

Notes: Data represent our sample of all Compustat firms reporting positive employment in each year from FYs 1999-2018. I-129 data from 1999-2012 come from USCIS FOIA data, while counts from 2013-2018 are from the USCIS Data Hub. The count of denied petitions is severely incomplete because USCIS stops accepting petitions beyond the final receipt date. From that point on, petitions are returned unopened to the sender and are not entered into the selection system.

Table A.2: The Effect of Rationing on Company Outcomes. Computer-services industry vs. Other industries.

Dependent Variable: ln of	(1) Employment	(2) Sales	(3) Profits	(4) MktVal	(5) CapEx	(6) R&D
Computer Services (SIC737)						
Rationing × User	0.076	0.282	-0.257	-0.200	-0.051	-0.173
	(0.110)	(0.262)	(0.244)	(0.146)	(0.187)	(0.151)
Observations	666	666	619	646	665	522
Avg. Approvals Non-user	0	0	0	0	0	0
Avg. Approvals User	1260	1260	1260	1260	1260	1260
Excludes Computer Ind.						
Rationing × User	-0.049	-0.065*	-0.049	-0.100*	-0.084	-0.082
	(0.028)	(0.032)	(0.030)	(0.046)	(0.047)	(0.074)
Observations	13734	13635	12190	12219	12219	12855
Avg. Approvals Non-user	0	0	0	0	0	0
Avg. Approvals User	107	107	107	107	107	107

Notes: Results from a regression model that classifies companies as users or non-users of the H-1B program, using 1 approval in 2010-2013 as the threshold to be considered a user of the program. The sample period is 2010-2018 and the *Rationing* indicator takes a value of one during years 2014-2018. The top panel considers only the computer services (*SIC* = 737) and the bottom panel considers all other industries. Firm fixed-effects (defined by gvkey, the unique firm identifier in Compustat) and industry-year fixed-effects included in all specifications. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.3: DiD Estimates of the Effect of Rationing. Weighted by propensity scores.

Dep. Var. ln of	(1) Emp	(2) Sales	(3) Profits	(4) MktVal	(5) CapEx	(6) R&D
FE ind-year DiD	yes	yes	yes	yes	yes	yes
FE ind-year Probit	no	no	no	no	no	no
Rationing × User	-0.053**	-0.060*	-0.068*	-0.107**	-0.096*	-0.066
	(0.026)	(0.034)	(0.036)	(0.042)	(0.050)	(0.065)
Observations	14319	14265	12800	12784	13460	5374
Avg # H1B 10-13 (Non Users)	0	0	0	0	0	0
Avg # H1B 10-13 (Users)	182	182	182	182	182	182
FE ind-year DiD	yes	yes	yes	yes	yes	yes
FE ind-year Probit	no	no	no	no	no	no
Rationing × User	-0.045*	-0.049	-0.052	-0.103**	-0.079	-0.064
	(0.026)	(0.035)	(0.036)	(0.042)	(0.051)	(0.065)
Observations	13959	13905	12462	12433	13101	5315
Avg # H1B 10-13 (Non Users)	0	0	0	0	0	0
Avg # H1B 10-13 (Users)	182	182	182	182	182	182

Notes: Results from a regression model that classifies companies as users of the H-1B program if they received at least one approved petition between 2010 and 2013. Observations are weighted using the (inverse of the) probability of being a user of the program. These probabilities were estimated on the basis of a Probit model. The estimates for the top panel of Table A.3 correspond to a version of the Probit model that does not include industry-year fixed-effects whereas the bottom panel does include these terms. In all cases the DiD estimation includes firm fixed-effects and industry-year fixed-effects. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.4: DiD Estimates of the Effect of Rationing. 2-digit industry trends

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)
Log of	Emp	Sales	EBITDA	MkVal	CapEx	R&D
Rationing × User	-0.047	-0.043	-0.02	-0.067	-0.06	-0.078
	(0.028)	(0.032)	(0.034)	(0.042)	(0.048)	(0.071)
Rationing1416 × User	-0.025	-0.025	-0.02	-0.044	-0.038	-0.05
	(0.024)	(0.027)	(0.032)	(0.039)	(0.044)	(0.059)
Rationing1718 × User	-0.081**	-0.070*	-0.0184	-0.101*	-0.094	-0.121
	(0.029)	(0.037)	(0.040)	(0.047)	(0.058)	(0.087)
Observations	14400	14300	12793	12856	13519	5414
Avg # H1B 10-13 (Non Users)	0	0	0	0	0	
Avg # H1B 10-13 (Users)	182	182	182	182	182	

Notes: Results from a regression model that classifies companies as users of the H-1B program if they received at least one approved petition between 2010 and 2013. In all cases we include 2-digit industry-year fixed-effects, along with firm fixed-effects. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.5: Balancing statistics Matching estimation.

	(1) Raw	(2) Raw	(3) Raw	(4) Matched	(5) Matched	(6) Matched
Means	Treated	Untreated	Std. Diff.	Treated	Untreated	Std. Diff.
User 2003-2006	0.92	0.28	1.75	0.92	0.47	1.24
Computer ind.	0.07	0.03	0.20	0.07	0.07	0.00
Ln(Emp)2010	2.50	0.08	1.31	2.50	0.65	1.00
Profits/Sales 2010	0.01	-0.42	0.06	0.19	0.14	0.01
R&D/Sales 2010	0.19	0.34	-0.03	0.06	0.05	0.00
$\Delta Ln(Emp)_{2007-2010}$	-0.04	0.03	-0.17	-0.04	0.01	-0.11
Variances	Treated	Untreated	Ratio	Treated	Untreated	Ratio
User 2003-2006	0.07	0.20	0.35	0.07	0.25	0.28
Computer ind.	0.07	0.03	2.54	0.07	0.07	1.00
Ln(Emp)2010	2.94	3.91	0.75	2.79	2.97	0.94
Profits/Sales 2010	19.64	92.58	0.21	0.09	0.15	0.60
R&D/Sales 2010	11.15	28.78	0.39	0.05	0.07	0.76
$\Delta Ln(Emp)_{2007-2010}$	0.14	0.19	0.73	0.09	0.07	1.24

Notes: Balancing statistics corresponding to the multivariate-distance matching estimator reported in 7 when the threshold to be considered a user of the program is having received at least 6 approved petitions for H-1B visas. Recall that each treated unit has been matched only to non-treated units in the same 2-digit industry.

Figure A.1: Aggregate I-129s. FOIA and Data Hub comparison

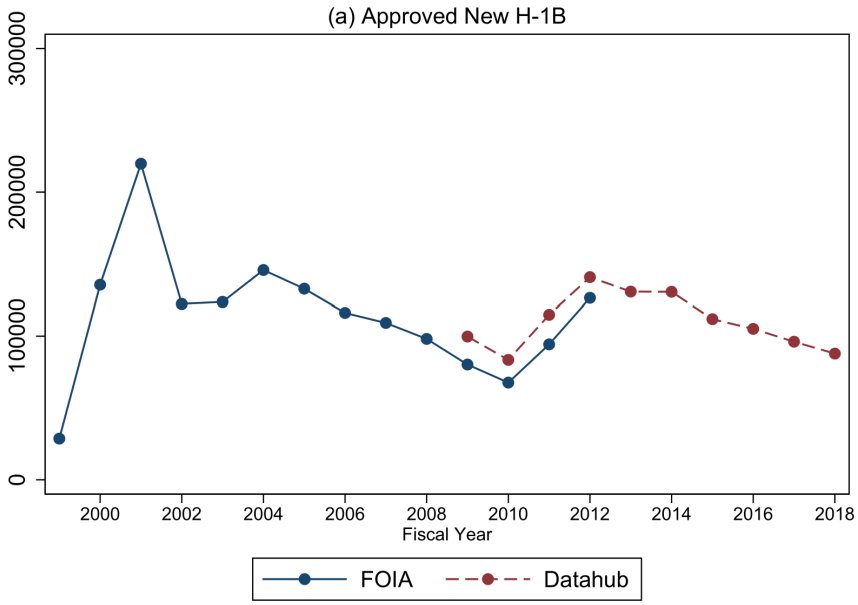

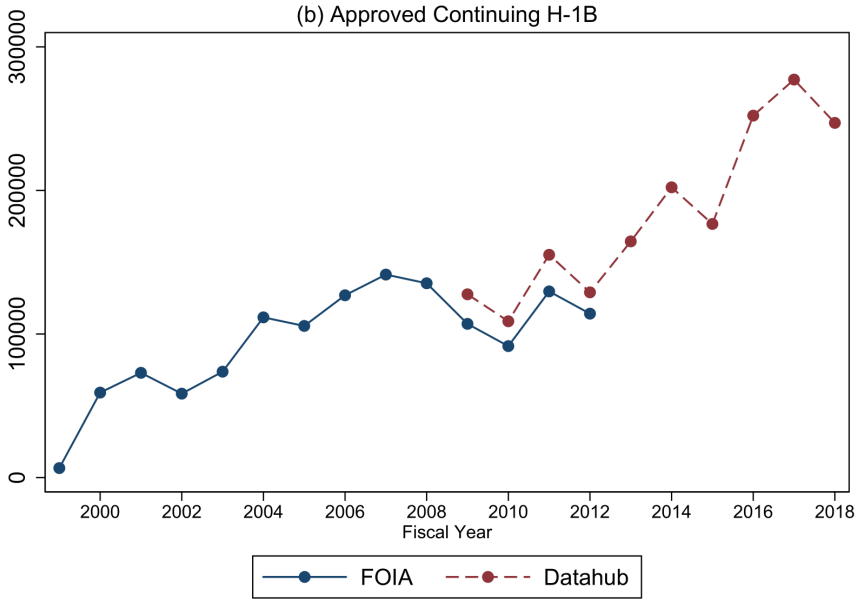

Note: Figures show aggregate totals of approved new H-1B petitions in panel (a), and approved continuing H-1B petitions in panel (b), by fiscal year. Blue series shows data from USCIS I-129 FOIA data. Red series displays data from the USCIS Data Hub.

Figure A.2: Approved H-1B Petitions at Four Largest H-1B Employers (Top Panel) and Other Compustat Firms (Bottom)

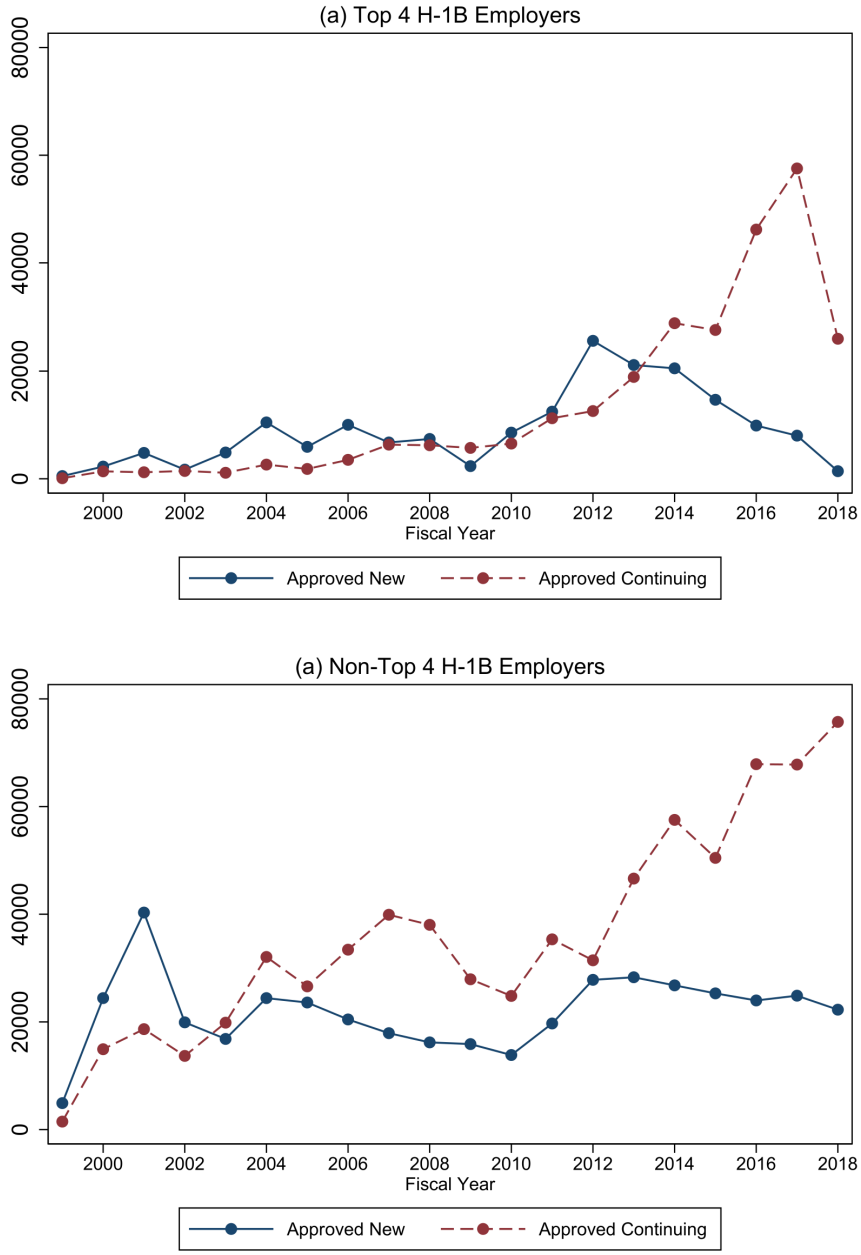

Notes: Figures show aggregate totals of approved new and continuing H-1B petitions for the sample of Compustat firms, by fiscal year. Blue series shows approved new H-1B petitions, while the red dashed series shows approved continuing H-1B petitions. Totals for the top four H-1B employers, in panel (a), are based on the total number of approved petitions for new H-1B employment over the whole period (1999-2018). The top 4 companies are: Infosys Ltd., Tata Group, Cognizant Tech Solutions, and Wipro Ltd. Totals for all other Compustat firms are shown in panel (b).

Table A.6: Firm Count Matrix by User Definition

	Low Level User	High Level User	Total
Low Ratio User	**301**	151	452
High Ratio User	152	**300**	452
Total	453	451	**904**

Notes: Table shows a tabulation matrix of firm counts of H-1B users (defined as having 1 or more approved H-1B petitions in the 2010-13 period) by ratio and level dependence measures. High/Low level users are those with H-1B approvals in 2010-13 above/below the median of all users). High/Low ratio users are split at the median of their total H-1B approvals as a share of employment in 2010-13. Cells show firm counts.

Table A.7: Intensity of Dependence and the Effect of Rationing, Level Dependence X Ratio Dependence

Dep. Var. *Ln* of:	(1) Employment	(2) Sales	(3) Profits	(4) Market Value	(5) Capital Expenditures	(6) R&D Spending
Rationing × Low-Level & Low-Ratio	-0.046	-0.070	-0.037	-0.125*	-0.060	-0.105
	(0.036)	(0.038)	(0.037)	(0.058)	(0.062)	(0.078)
Rationing × Low-Level & High-Ratio	0.039	0.076	0.087	-0.029	0.058	0.062
	(0.048)	(0.046)	(0.075)	(0.082)	(0.095)	(0.101)
Rationing × High-Level & Low-Ratio	-0.124**	-0.141**	-0.104**	-0.179**	-0.135**	-0.216**
	(0.038)	(0.043)	(0.042)	(0.070)	(0.058)	(0.085)
Rationing × High-Level & High-Ratio	-0.039	-0.056	-0.094**	-0.077	-0.149**	-0.089
	(0.031)	(0.035)	(0.039)	(0.043)	(0.048)	(0.072)
Observations	14,400	14,301	12,809	12,865	13,520	5,418

Notes: Results from a regression model that classifies companies as users or non-users of the H-1B program. Users are further split into 4 groups: (1) low level dependence and low ratio dependence, (2) low level dependence and high ratio dependence, (3) high level dependence and low ratio dependence and (4) high level dependence and high ratio dependence. The omitted category is always non-users. Firm fixed-effects (defined by gvkey, the unique firm identifier in Compustat) and industry-year fixed-effects included in all specifications. The sample only contains companies with positive employment in every fiscal year throughout the sample period. Two-way clustered standard errors by firm (gvkey) and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.