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

March 3, 2022

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-serve 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 propensity-score matching, we estimate that the switch in the visa allocation system has 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. As such, it receives significant attention among academics, policy-makers, and business leaders. Supporters argue that the U.S. has high and growing demand for skilled labor, driven by the success of its high tech firms, and that restrictions on H-1B visas inhibit firms' innovation and growth. Opponents contend that the program creates unfair competition for American workers. In recent iterations of this debate, Donald Trump issued a presidential proclamation on June 22, 2020, suspending 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]

Though details of the program 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)–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 new H-1B status in each fiscal year. Prior to FY 2021, April 1 had marked the start of the application period for work to begin on October 1 (the first day of the US fiscal year). Demand for new H-1B workers was so high in FYs

[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 for 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.

2007-2009 and 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 use a lottery to select H-1B applications during these "lottery years".

To be clear, what led to the rationing of H-1B visas was a widespread increase in the aggregate demand for highly skilled foreign workers, which resulted in employers filing applications earlier in time and possibly submitting a larger number of petitions. Because the annual quota on new-employment H-1B visas was kept fixed, the overall number of new visas awarded on an annual basis did not change throughout our sample period. What changed was the method to allocate them, presumably giving rise to a very different allocation of visas among the set of companies that submitted applications. Potentially, some high-surplus employer-employee matches could not take place due to the widespread use of the lotteries. From the employer's viewpoint, the lottery-based allocations reduced the probability of forming the desired match with a specific worker, and created uncertainty that might have distorted other investments complementary to the hiring of the foreign skilled worker.

Our main research question is: what have been the effects of the change in visa allocation system on the companies that use the H-1B program? To investigate this question, we build a data set that merges 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 United States. While these companies are not a random sample of all private-sector employers, publicly traded companies play a crucial role in determining the growth and volatility of the domestic macroeconomy (Gabaix (2011)) and international trade flows (di Giovanni et al. (2018)).

To avoid capturing the economic disruption of the Great Recession, 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 develop-

ment (R&D) expenditures. We consider FYs 2014-2018 as the period of H-1B rationing, and FYs 2010-2013 as the pre-rationing period. Employing a DiD strategy and propensity score 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 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.

Importantly, understanding how workers adjust is incomplete without considering firm-level responses. Adjustments that materialize as changes to native-born wages, employment, or invention, for example, are likely the culmination of various firm-level choices to expand or contract hiring, or relocate, alter, or expand operations. For example, Glennon (2020) finds that H-1B restrictions leads multinational firms to offshore more jobs. 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). Relative to other studies that have looked at firm employment, we broaden the scope of by examining a wider set of important company 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 FY2004), this allows us to focus on a more recent period of heightened demand for H-1B visas (FY2010-2018).

The remaining 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 propensity-score matching analysis. Section 6 concludes.

2 Hiring H-1B Workers and Empirical Strategy

To hire H-1B workers, firms must first file a Labor Condition Application (LCA) with the U.S. Department of Labor. LCAs are not linked to specific individuals, but rather contain basic information about the job. Approved LCAs serve as a permission slip of sorts for firms to search for qualified foreign workers. Several studies (e.g. Kerr and Lincoln (2010)), have relied upon LCAs to approximate the demand for H-1B workers, but this approach faces important limitations. Chiefly, the number of LCAs systematically exceeds the number of petitions for H1-B visas.

Once a firm has identified a foreign worker it would like to hire, it must file an I-129 petition with USCIS seeking an H-1B visa on the worker's behalf. These petitions are tied to a specific candidate and firms cannot transfer H-1B status from one worker to another (nor can they prevent a person with H-1B status from moving to a new company that possesses an approved LCA). Nonetheless, other program features create strong employer/employee ties. For example, individuals cannot apply for H-1B status on their own; they need to be sponsored by an employer (with an approved LCA). Moreover, H-1B visas have dual-intent status, allowing firms to sponsor these employees for permanent residency. Employed H-1B workers might not want to jeopardize this possibility by participating in subsequent job searches (e.g. Sparber, 2019; Depew et al., 2017). Altogether, this implies that I-129 data on approved H-1B petitions are a much

more accurate measure of H-1B employment at the firm level.

The U.S. government grants H-1B status to a foreign worker for up to three years and firms can renew a worker's status for another 3-year period. Moreover, H-1B visa holders waiting for permanent residency may hold H-1B status beyond six years. Importantly, the renewals and extension 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.

The annual caps pertain to the corresponding fiscal year, which begins on October 1. Though processes changed somewhat beginning in FY 2021, historically, a firm could petition for a worker to receive H-1B status beginning six months prior to the work start date (thus, April 1 at the earliest). USCIS allocates H-1B status to qualified workers on a first-come, first-served basis. The last date of receipt occurs when USCIS has received enough applications to meet the cap. USCIS then uses a lottery to randomly select among the applications received on that day. 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.

Consider again Figure 1, which plots the number of days until USCIS determined the final receipt day in each fiscal year. Prior to FY 2008, firms were able to secure their desired number of new H-1B workers so long as they submitted their petitions at the beginning of the application period. 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 FY 2008-09, USCIS had received more than

enough applications for new H-1B visas within a week of accepting applications. 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 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 match with 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 users of 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 of the H-1B program in FY2014 (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 suggest the adoption 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 propensity score matching, where we utilize a subset of non-users that most closely resembles treated companies as the control group.

3 Data Sources and Descriptive Statistics

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

Our original data covers two decades of the H-1B program (FY1999-2018). The main empirical analysis, however, is based on the period starting in FY 2010, which excludes the Great Recession and important changes to immigration policy that might have affected the demand for H-1B visas.[2] To create a consistent firm-level dataset on approved H-1B applications (more precisely, 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. approved, 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 blue solid line and the 2009-2018 data from the Data Hub source in the

[2] 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.

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 covering the two decades between FYs 1999 and 2018. Our sample retains only those firms which 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.[3]

The only firm identifying information available in both the H-1B and Compustat datasets is firm name. In order to tally the individual visa petitions (I-129s) to the company level and to merge the resulting data with Compustat, we relied on string matching techniques and created harmonized 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), our sample of firms accounts for around

[3]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.

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 FY2018 and about 10,000 in FY2019. Approved petitions for continuing employment increased from FY1999 through FY 2017, peaking at 125,000, and then dropping to 102,000 in FY2018. Denials for continuing-employment increased in FY2016 and FY2017 and surged in FY2018. In this year there was also an uptick in new-employment denials.

The raw data highlights an interesting new trend that will play an important role in our analysis. Figure 2 shows the number of new (blue solid line) and continuing (red dashed line) H-1B applications by firms in the Compustat database. 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 FY2007-2008. 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 the visa allocation system 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 1,600 firms with positive employment in each fiscal year from 1999 through 2018. Hence, we have a balanced panel of established firms with at least two decades of continuous history.[4] 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 H-1B petition between FYs 2010 and 2013. We also build measures of the *intensity* of use of the program, partitioning users based on whether they received a number of approved petitions for H-1B visas that was above or below the median number (among 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 normalize by average firm employment in the 2010-2013 period (*ratio dependence*).

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. On average, users are much larger than non-users along several dimensions. 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 magnitude are observed along the other outcomes as well, with the exception of R&D expenditures where users of the H-1B program spend 37 times more than the average non-user. In comparison, low-intensity and high-intensity users of the H-1B program are much more similar on average. 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 double on R&D than the average non-user.

We also observe important differences in industry composition.[5] About 43% of com-

[4]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.

[5]We aggregate SIC industry codes into 8 mutually exclusive groups: Agriculture, Construction,

panies 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 the 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 rationing on the performance of the company.

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.

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. Specifically, we predict annual approved petitions as follows. First, we approximate the employment share of the total stock of H-1B workers in baseline year 2005 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 assume that the 'normal' *total stock* of H-1B workers in the company in years $t \geq 2003$ grows at the same rate as overall employment, leaving its share in employment constant at the baseline 2005 level. Last, we use linear regression to build predictors for the annual *flow* of new and continuing-employment visas as a function of the contemporaneous predicted *total stock* for users of the program in the baseline year.

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

Our identification strategy in this section is based on DiD estimation. Our main goal is to uncover the effects of the switch in the visa allocation system that took place in FY 2014 on the performance of companies using the H-1B program. We will 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 foreign skilled workers they have previously identified. 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 $t \geq 2010$:

$$y_{ijt} = \alpha + \beta Rationing_t \times User_i + \gamma_i + \gamma_{jt} + \varepsilon_{ijt}, \tag{1}$$

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.[6]

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

[6]We use 8 mutually exclusive industry groups, aggregated from 4-digit SIC codes: Agriculture, Construction, Manufacturing, Trade, Finance/Real Estate, Services, Computer, and Other.

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 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 permit valid inference if there is within-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 zero 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 (as in 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.[7] In particular, these reductions in capital investment may be a reflection of outsourcing of the activities that would have been undertaken by the visa holders to other domestic or foreign-based companies (as in 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.

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

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 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. normalized 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 prediction 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 the earlier discussion, 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:

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

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 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 normalizes 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 2013 and 2014) 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 of the program 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 a larger blow whereas the most highly dependent users did not experience any adverse effects from rationing.

Given the disagreement between the two measures of dependence employed here, we choose to remain agnostic in regards to the relationship between the intensity of use of the H-1B program and the effects of rationing on company outcomes.

5 Propensity Score Matching

This section adopts an alternative approach to the estimation of the effects of visa rationing: propensity score matching. More specifically, our goal here is to estimate the average treatment effect on the treated (ATET) where users of the program in 2014 are considered as the treatment group and non-users are used to construct a control group.

Implementation of propensity score matching requires addressing two important issues. First, we need to determine which companies were users of the program in FY 2014, when the rationing episode began. To maintain comparability with our earlier analysis, we again define the treatment indicator on the basis of the approved visa petitions received between FYs 2010 and 2013. Secondly, we need to choose the variables to match each company in the treatment group with one (or more) non-users, which will compose the control group. Fortunately, our dataset contains a rich set of variables, increasing the probability that our estimates uncover causal effects, which requires that, after conditioning on covariates, any remaining influences on the treatment not be related to the potential outcomes.

Our main analysis in this section is based on the comparison of the mean treatment effect among users of the program relative to the corresponding value for the set of

matched non-users (determined by the estimated propensity scores). Specifically, for each firm i in the treatment group, we estimate the (individual) treatment effect as the difference between the outcome observed for i and the outcome of the matched non-user. Then we compute the ATET by averaging over all firms in the treatment set. In addition, we also perform our primary DiD estimation using inverse propensity score reweighting. We demonstrate that both analyses are consistent with each other, and also in line with the non-matching DiD estimates. Because of its greater simplicity, our discussion in this section focuses on the difference-in-means analysis.

5.1 Definition of Treatment

We define a firm to be a user of the H-1B program in 2014 (treated) when the number of total approved petitions it received over the 2010-2013 period is above a threshold. For robustness, we consider two thresholds for the number of total approved petitions: 1 and 6 approved petitions.

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.

5.2 Propensity scores

In order to match each user of the H-1B program in 2014 to a comparable non-user, we estimate a *probit* model for the probability of using the program ($User1418$ indicator). The explanatory variables include previous use of the H-1B program (defined as having received at least one approved petition between 2003 and 2006), employment growth between 2007 and 2010, an indicator for belonging to computer industries, and a comprehensive set of (coarser) industry fixed-effects. In addition, we also include the firm's employment size, profitability and R&D intensity in 2010 as control variables.

Let us begin by examining the marginal effects entailed by the estimates of the *probit* model. The top panel in Table 7 refers to the model where any firm receiving at least one approved petition in 2010-2013 is considered a user of the H-1B program (treatment). The estimates show that three variables are strong predictors of program participation: previous participation, employment size and previous employment growth.[8] Among these, we highlight the large marginal effect of previous usage, which has been the common choice in the existing literature (see Kerr and Lincoln (2010)). Our analysis here shows that other variables are also important predictors. In addition, the bottom panel of the Table confirms the robustness of these findings to the choice of the threshold

[8]It is perhaps surprising that the Computer industry dummy variable (constructed using 3-digit industry codes) does not turn up as a significant predictor of program participation. This is due to the inclusion of (coarser) industry fixed-effects. When these are not included, the Computer industry dummy is large and highly significant, entailing a 12.6 percentage-point increase in the probability of using the program in 2014.

for participation.

5.3 Average Treatment Effect on the Treated (ATET)

Next, we use the estimated coefficients to predict the probability of treatment (user of the program in 2014) for each firm in our sample. Then we assign to each firm in the treatment group the non-user with the closest predicted probability (propensity score), and compute the individual treatment effect as the difference in outcomes. Last, we average the individual treatment effects over all users of the program to obtain the estimated ATET.

Table 8 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. We first report the unmatched mean difference in the 2014-2018 period for users of the program net of the corresponding value for non-users. Column 1 shows that employment growth during the rationing period was approximately 6.1% lower for users of the program. Next, we turn to the ATET, which is based solely on the users and the corresponding matched non-users. We find that the change in the allocation system of visas in 2014 reduced the employment growth of users by approximately 6.2% (1.5% annually) between FYs 2014 and 2018. Columns 2 and 3 consider two sub-periods and 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-employment visas) declined over time. Last, the bottom panel repeats the analysis using a higher participation threshold (of 6 approved petitions), and the estimated ATET remains largely unchanged.

5.4 Re-weighting Difference-in-Difference Estimation

We now consider an alternative use of the estimated propensity scores. Specifically, we perform our main difference-in-differences specification Equation (1) using inverse propensity score reweighting. The results are reported in Table A.3. The top panel in the Table presents estimates based on propensity scores that were estimated *without* industry trends, whereas the bottom employs propensity scores that did include industry-year dummy variables. In both cases a firm was considered a user of the H-1B program if it received at least one approved petition between 2010 and 2013. Both qualitatively and quantitatively, the estimates are highly consistent with those reported in Table 2.

To sum up, this section has adopted an alternative approach to the estimation of the effects of rationing on company outcomes. The estimates obtained here (on the basis of propensity score matching) largely confirm the main finding of a negative impact of visa rationing on the growth of users of the H-1B program.

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 comple-

mentarities 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.

While we feel confident about our findings, we encourage researchers to revisit the analysis of the effects of visa rationing on the outcomes of users of the program using other identification strategies. In particular, analyses based on randomized assignment of visas would be enormously helpful.

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)). For the sake of comparison, allocating NSF grants randomly across applications could be considered fair. However, it would not lead to an efficient use of taxpayer funds. In an ideal world, a rigorous evaluation of each company's petitions would lead to a more efficient allocation. Obviously, because of the large scale of the H-1B visa program, the detailed review of all applications would be very costly. However, auctioning visas for private-sector employers, a points-based system, or other alternative allocation mechanisms should be considered (Peri (2012)). These mechanisms should retain an important role for employers in identifying the foreign workers possessing the skills needed in their companies (Kerr et al. (2015a)). Additionally, the annual cap on petitions filed by private-sector employers should be flexible, for example taking into account labor market tightness in specific occupations.

References

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

Depew, Briggs, Peter Norlander, and Todd A. Srensen, "Inter-firm mobility and return migration patterns of skilled guest workers," *Journal of Population Economics*, April 2017, *30* (2), 681–721.

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.

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.

Krusell, Per, Lee E. Ohanian, Jos-Victor Ros-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, "STEMWorkers, 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.

Sparber, Chad, "Substitution between groups of highly-educated, foreign-born, H-1B workers," *Labour Economics*, 2019, *61*, 101756.

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

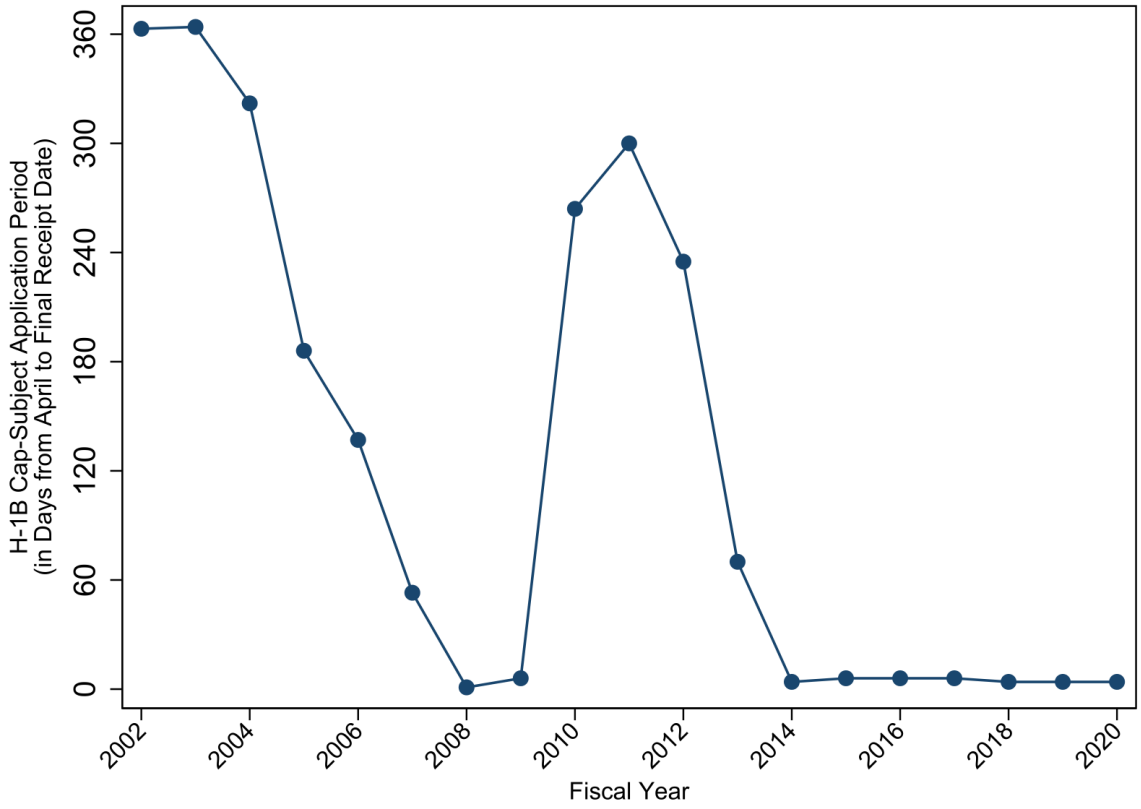

Notes: Number of days between the first week of April (when application window opens) and final receipt day (when USCIS has received enough applications to meet the statutory cap). Zero values correspond to years when USCIS allocated all cap-bound H-1Bs by lottery. In other years, USCIS allocated H-1Bs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

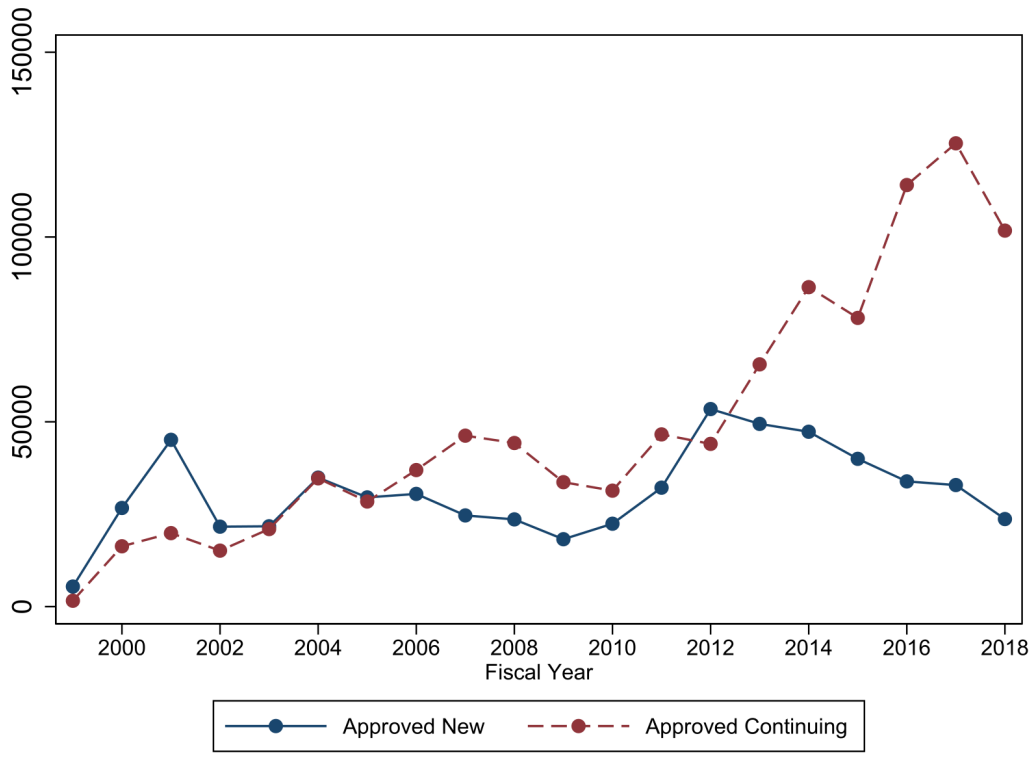

Note: Figures show aggregate totals of approved new and continuing H-1B petitions for the sample of Compustat firms, by fiscal year. Blue series shows approved new H-1B petitions, while the red dashed series shows approved continuing H-1B petitions.

Figure 3: Hoarding

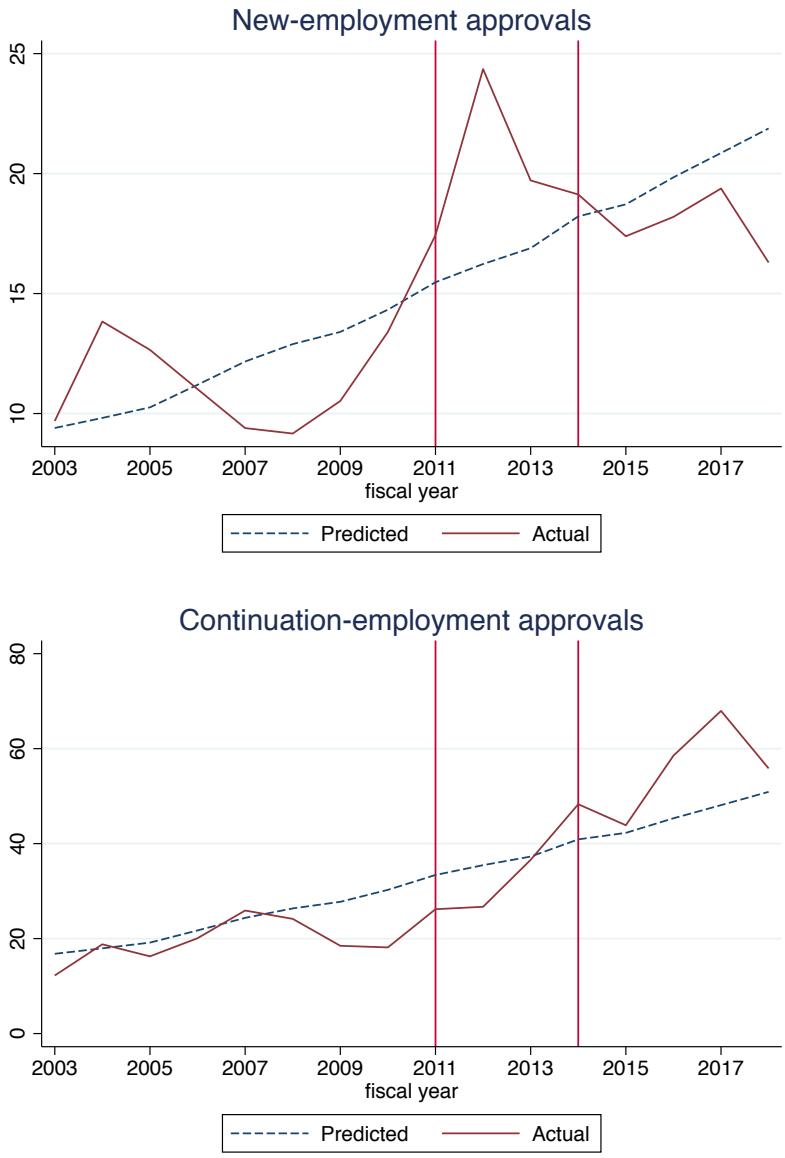

Notes: Dependent variable is the ratio of the stock of (new or continuing) H-1B workers normalized by a firm's average total employment from FY 2010-2013. Model includes company fixed-effects and industry-year fixed-effects. Rationing period covers fiscal years 2014 through 2018. Regressions compare high use to low use firms. Point-estimates are relative to FY 2013, the omitted category. 95% Confidence intervals are shown in dashed lines.

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

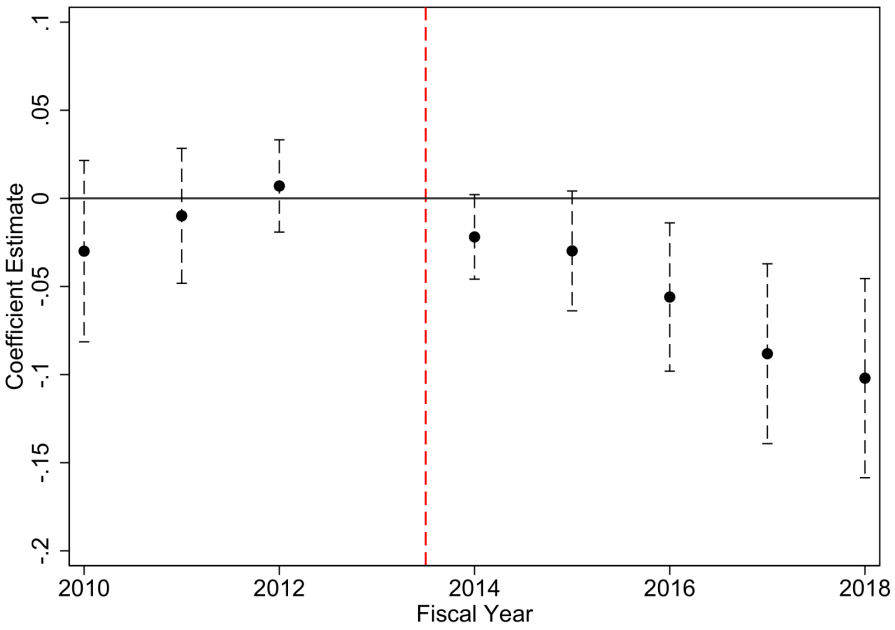

Notes: Dependent variable is the log of employment. Model includes company fixed-effects and industry-year fixed-effects. Rationing period covers fiscal years 2014 through 2018. Point-estimates are relative to FY 2013, the omitted category. 95% Confidence intervals are shown in dashed lines.

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
Employment (thousands)	5.8	33.3	45.1	21.5
Sales ($ million)	1831.3	12335.4	13311.3	11359.5
Profits - EBITDA ($ million)	523.4	4532.2	4388.8	4675.6
Market Value ($ million)	1432.1	12202.7	10581.8	13823.6
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
Approvals 2010-13				
New Employment	0	74.7	7.2	142.2
Continuation	0	107.7	17.8	197.5
Total (New+Cont.)	0	182.4	25.0	339.7
Total / AvEmp (%)	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.

Dependent Variable: ln of	(1) Employment	(2) Sales	(3) Profits	(4) MktVal	(5) CapEx	(6) R&D
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	14400	14301	12793	12856	13519	5414
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	14400	14301	12793	12856	13519	5414
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	14400	14301	12793	12856	13519	5414
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	14400	14301	12793	12856	13519	5414

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. User's 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) 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.

Dependent Variable: ln of	(1) Employment	(2) Sales	(3) Profits	(4) MktVal	(5) CapEx	(6) R&D
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						
Rationing1416 × 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)
Rationing1718 × 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	14400	14301	12809	12865	13520	5418
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) 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

Dependent Variable: ln of	(1) Employment	(2) Sales	(3) Profits	(4) MktVal	(5) CapEx	(6) R&D
3 groups						
Rationing × UserLow	-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 × UserHigh	-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	14400	14301	12809	12865	13520	5418
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						
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	14400	14301	12809	12865	13520	5418
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 (UsersLow) or above (UsersHigh) 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) 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

Dependent Variable: ln of	(1) Employment	(2) Sales	(3) Profits	(4) MktVal	(5) CapEx	(6) R&D
2 groups						
Rationing × UserLow	-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 × UserHigh	-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)	.0765	.0765	.0765	.0765	.0765	.0765
Avg. Dependence (High)	2.22	2.22	2.22	2.22	2.22	2.22
5 groups (quintiles)						
Rationing × Users 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 × Users 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 × Users 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 × Users 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 × Users 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	14400	14301	12809	12865	13520	5418
Avg Dependence (Q1)	.025	.025	.025	.025	.025	.025
Avg Dependence (Q2)	.0869	.0869	.0869	.0869	.0869	.0869
Avg Dependence (Q3)	.2	.2	.2	.2	.2	.2
Avg Dependence (Q4)	.561	.561	.561	.561	.561	.561
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 normalized 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 (UsersLow) or above (UsersHigh) 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) 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

Total Approvals1013 Variable	None All (1600)	Threshold = 1 NoUsers (696)	Threshold = 1 Users (904)	Threshold = 6 NoUsers (999)	Threshold = 6 Users (601)
Ch. lnEmp1418	0.084	0.116	0.060	0.109	0.043
User dummies					
User0306	0.520	0.148	0.806	0.277	0.923
User1013 (th1)	0.565	0	1	0.303	1
User1013 (th6)	0.376	0	0.665	0	1
Approvals1013					
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
Ch. lnEmp0710	0.006	0.032	-0.014	0.033	-0.040
lnEmp2010	0.963	-0.225	1.878	0.037	2.503
Profits/Sales2010	-0.256	-0.580	-0.008	-0.417	0.009
R&D/Sales2010	0.283	0.423	0.177	0.337	0.194

Notes: Total *Approvals1013* is the combined approved petitions received between 2010 and 2013. *Ch. lnEmp1418* is the 2014-2018 change in log employment. $User0306$ takes a value of 1 if the firm received at least one approved petition between 2003 and 2006. Similarly, $User1013$ is an indicator for receiving approvals above the threshold of 1 ($th1$) or 6 ($th6$). *Ch. lnEmp0710* is the 2007-2010 change in log employment. Profits are measured by EBITDA.

Table 7: Probit model for probability of using H-1B program in 2014 (*Use1418*). Average marginal effects

	Avg. Mg. Eff.	SE Delta M.	tstat
Prob (H1Btotal1013>0)			
Use0306	0.33**	0.01	23.86
lnEmp2010	0.05**	0.00	9.69
DlnEmp0710	-0.04 **	0.02	-1.98
Computer Ind.	0.00	0.06	0.01
Profits/Sales2010	0.00	0.00	0.08
R&D/Sales2010	0.00	0.00	0.14
Prob (H1Btotal1013>5)			
Use0306	0.31**	0.02	15.49
lnEmp2010	-0.07**	0.02	-2.87
DlnEmp0710	0.09**	0.00	19.31
Computer Ind.	0.05	0.05	0.9
Profits/Sales2010	0.02	0.02	0.81
R&D/Sales2010	0.03	0.03	0.9

Notes: 2-digit industry codes included in the estimation. $H1Btotal1013$ is the total number of approved petitions received between 2010 and 2013. Standard errors marginal effects computed using the delta method. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: Estimation propensity score matching. Average Treatment Effects on the Treated (ATET)

Difference in Means of Change lnEmp in: (percentage)	(1) 2014-18	(2) 2014-16	(3) 2016-18
Threshold Approvals = 1			
Unmatched	-6.1**	-4.2**	-1.9
ATET	-6.2**	-2.4	-3.8**
Threshold Approvals = 6			
Unmatched	-7.2**	-3.3**	-3.9**
ATET	-5.8	-2.1	-3.7**

Notes: The Table reports two sets of estimates. The *Unmatched* difference in means is the difference between the average 2014-2018 change in log employment among users of the H-1B program in 2014 and the corresponding value for non-users. The *ATET* is the individual treatment effect averaged over all users of H-1B program in 2014. The individual treatment effect is the difference between the 2014-2018 change in log employment for a treated firm (i.e. a user of the H-1B program in 2014) and the corresponding value for the matched non-user (on the basis of the estimated propensity scores). The estimated ATET takes into account that propensity scores are also estimated and was implemented using the *teffects psmatch* command in *Stata16*. The Probit models used to estimate the propensity scores (for both thresholds) include 2-digit industry fixed-effects. The sample used to match individual observations excludes 34 companies for which the predicted probability was either zero or one.

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) 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.	(1)	(2)	(3)	(4)	(5)	(6)
ln of	Emp	Sales	Profits	MktVal	CapEx	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 the Probit model presented in Table 7 (top panel). 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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

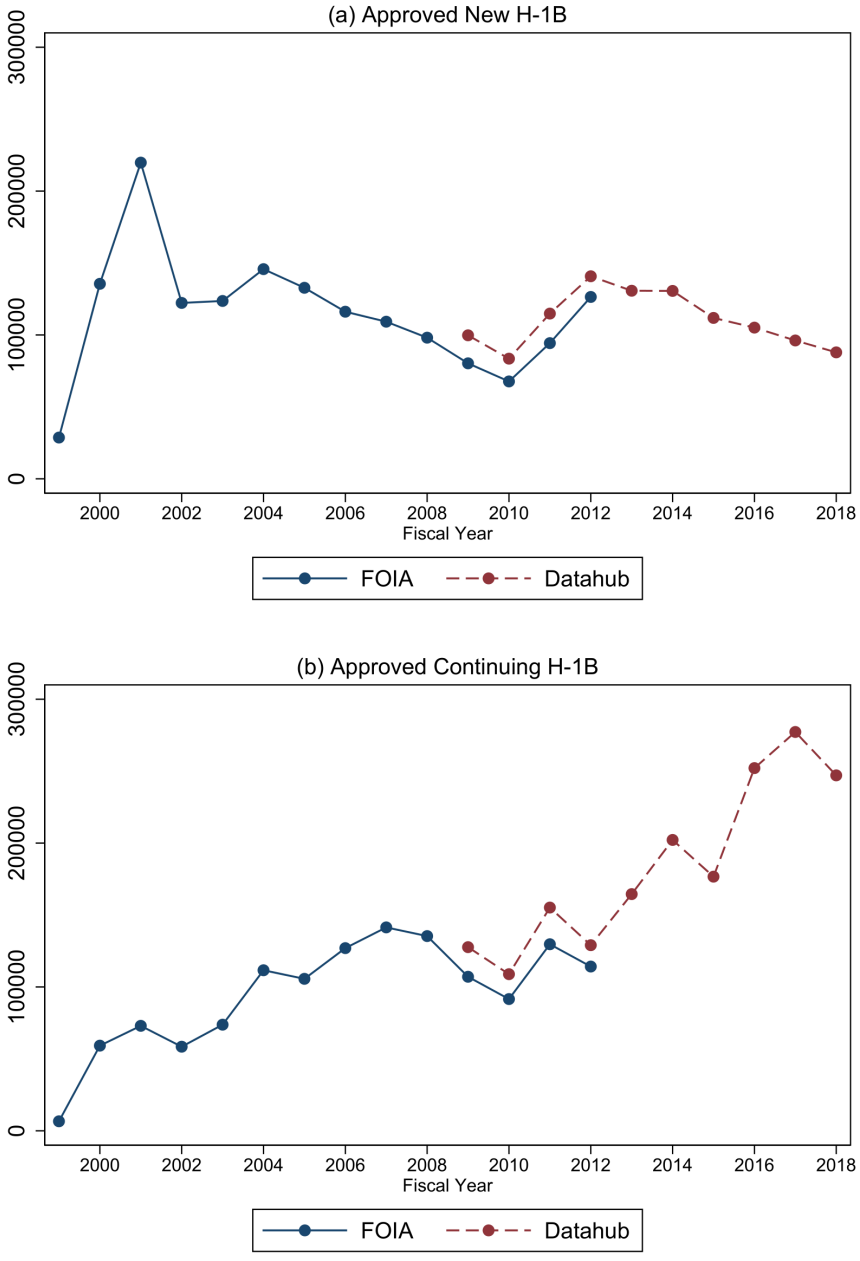

Note: Figures show aggregate totals of approved new H-1B petitions in panel (a), and approved continuing H-1B petitions in panel (b), by fiscal year. Blue series shows data from USCIS I-129 FOIA data. Red series displays data from the USCIS Data Hub.

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

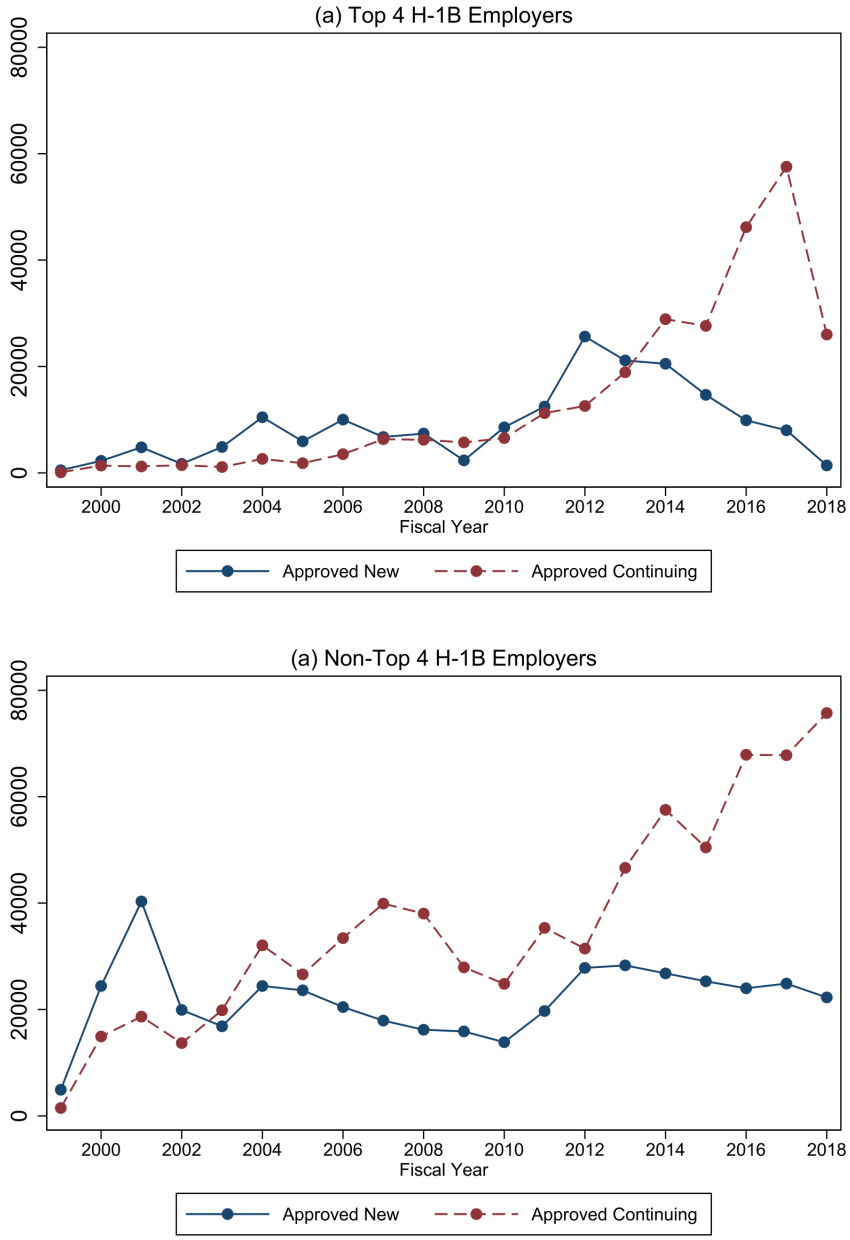

Notes: Figures show aggregate totals of approved new and continuing H-1B petitions for the sample of Compustat firms, by fiscal year. Blue series shows approved new H-1B petitions, while the red dashed series shows approved continuing H-1B petitions. Totals for the top four H-1B employers, in panel (a), 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).