

NBER WORKING PAPER SERIES

COPING WITH H-1B SHORTAGES:  
FIRM PERFORMANCE AND MITIGATION STRATEGIES

Anna Maria Mayda  
Francesc Ortega  
Giovanni Peri  
Kevin Y. Shih  
Chad Sparber

Working Paper 27730  
<http://www.nber.org/papers/w27730>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
August 2020

The authors gratefully acknowledge funding from the National Science Foundation, Award 1535723, "The Effect of H-1B Workers on the Innovation and Productivity in US Firms." The grant was functional in developing the data used in this paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Anna Maria Mayda, Francesc Ortega, Giovanni Peri, Kevin Y. Shih, and Chad Sparber. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Coping with H-1B Shortages: Firm Performance and Mitigation Strategies  
Anna Maria Mayda, Francesc Ortega, Giovanni Peri, Kevin Y. Shih, and Chad Sparber  
NBER Working Paper No. 27730  
August 2020  
JEL No. F22,J61

**ABSTRACT**

The United States' H-1B visa program, which allows private firms to hire highly skilled foreign workers, was so severely over-subscribed in the years since 2014 that H-1B status was distributed by lotteries to a subset of applicants. Using data on H-1B applications and on a range of outcomes for publicly traded companies, we find that employers using the H-1B program experienced reduced employment, sales and profits, compared to non-users in the years since 2014. We also find that some employers anticipated the rationing of H-1Bs and retained a larger share of H-1B workers, mitigating the damaging effects of H-1B rationing on their performance.

Anna Maria Mayda  
Department of Economics and  
School of Foreign Service  
Georgetown University  
37th and O Streets, NW  
Washington, DC 20057  
amm223@georgetown.edu

Kevin Y. Shih  
Queens College CUNY  
Economics Department  
300O Powdermaker Hall  
65-30 Kissena Blvd  
Flushing, NY 11367  
kshih@qc.cuny.edu

Francesc Ortega  
Queens College CUNY  
Economics department,  
300A Powdermaker Hall,  
65-30 Kissena Blvd,  
Flushing, NY 11367  
fortega@qc.cuny.edu

Chad Sparber  
Department of Economics,  
Colgate University,  
13 Oak Drive,  
Hamilton, NY, 13346.  
csparber@colgate.edu

Giovanni Peri  
Department of Economics  
University of California, Davis  
One Shields Avenue  
Davis, CA 95616  
and NBER  
gperi@ucdavis.edu

# 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 competition for American jobs. In a recent iteration 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.<sup>1</sup>

Though details of the program have changed over time, there has always been an annual cap on the number of new H-1B visas issued to employees of private firms.<sup>2</sup> Initially set at 65,000 per year in the early 1990s, the cap was raised to 115,000 for fiscal years (FYs) 1999 and 2000, peaked at 195,000 for FYs 2001-2003, then reverted to 65,000 in FY 2004. Since FY 2006, the cap has restricted new inflows to 65,000 general H-1B workers per year, plus 20,000 H-1B visas for workers with a Master’s Degree (or higher degree) from a U.S. university. H-1B status is valid for up to three years and can generally be renewed for an additional three years. H-1B workers seeking to change employers or renew their status are exempt from the cap.

Although the cap of 85,000 H-1B applications has been in place since 2006 and has been reached every year, the constraints it imposes on hiring have been particularly binding in times of tight labor markets. Firms competing to hire foreign labor face smaller windows for doing so when labor markets are tight. **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 2007-2009 and 2014-2020 that the cap was reached in the very first week of the application period. As a result, instead of approving applications on a first-come/first-served basis, the government selected

---

<sup>1</sup>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.

<sup>2</sup>Employees of universities, government, and non-profit research institutions are exempt from the annual H-1B cap.

among new H-1B applications using a random lottery. H-1B workers were clearly in high demand and firms applying for them were heavily rationed during these “lottery years.”

The goal of this paper is to analyze the effect of the rationing of H-1B visas in lottery years on firm outcomes and on firm behavior. To do so, we build a data set that merges the universe of H-1B applications between fiscal years 1999 and 2018 with Compustat data on a wide range of outcomes for all publicly traded companies in the United States. To avoid the depth of the great recession, our analysis focuses on fiscal years 2010-2018 and examines a broad set of firm-level indicators including employment, sales, profits, market value, and research and development (R&D) expenditures.<sup>3</sup> The period of tight H-1B rationing include FYs 2014-2018, while FYs 2010-2013 are the pre-rationing period. In a difference in difference approach, 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.

Although the heavy rationing was generated by rising aggregate demand for the H-1B program, from the perspective of an individual firm, the “lottery years” constitute a sudden reduction in the probability of obtaining an H-1B hire. Hence, looking at this event through the lens of a “search model”, in lottery years each individual firm had a lower probability of successfully matching with the right H-1B worker. For given firm characteristics and productivity, the treatment we are considering is a negative shock to the matching probability of that firm with an H-1B worker.

Our empirical analysis reveals several important results. First, firms that employed H-1B workers experienced lower total employment, sales, profits, market value, and R&D expenditures in “lottery years”. This difference in outcomes is unlikely to be driven by firm-specific labor-demand shocks for several reasons. First, firms hiring H-1B workers are relatively larger, more productive and more technologically advanced, and those are firms that tend to grow faster in booms. Second, in all specifications we include both industry-by-year and firm fixed-effects, which will capture demand changes at the industry level and firm heterogeneity in productivity. Finally, we find no evidence of diverging trends between H-1B employers and non-employers in the pre-treatment period. Hence, the difference in outcomes during treatment years is consistent with rationing of H-1B visas being the cause of those.

---

<sup>3</sup>Results from robustness checks using data from FY1999-FY2009 are available upon request. This longer period includes only one earlier episode of heavy rationing and complete lottery allocation in FY2008-FY2009.

Next, we document a strategy that a subset of companies developed in order to mitigate the effects of H-1B rationing. These firms hired more continuing H-1B workers who renewed their status as the hiring of new H-1B workers was becoming harder. We divide H-1B employers into two groups – high users and low users – based on the new H-1B workers as share of employment prior to the lottery years. We find that high users applied this mitigation strategy by renewing existing H-1B workers at higher rates than low-users and possibly attracting existing H-1B workers from low-using employers in lottery years. Our data do not allow us to distinguish the within-firm renewal channel from the poaching of continuing H-1B workers from other companies, however the evidence seems more consistent with the former behavior. Somewhat counter-intuitively, we find that after FY 2014, high-use firms experienced smaller reductions in employment, sales (and other performance measures) compared to low-use firms and the adoption of mitigation strategies could be the explanation for this phenomenon.

We provide further support for these strategies by showing that high-use employers increased their applications for new H-1B workers in the years leading to the lotteries (2012 and 2013) when applications were still being approved in a first-come, first-served basis, and built a “buffer” of H-1B visas, relative to low-use firms.

The remaining of the paper proceeds as follows. [Section 2](#) summarizes academic literature on high skilled migration to provide context for its importance to the US economy, while [Section 3](#) provides background on the H-1B program. [Section 4](#) describes the data and presents descriptive statistics. [Section 5](#) examines the effects of H-1B rationing on company outcomes by comparing the trajectories of users vs. non-users of the program before and after FY2014. [Section 6](#) provides evidence of the H-1B buffering that occurred immediately before lottery years. [Section 7](#) concludes.

## 2 Related Literature

This paper contributes to a growing body of literature on the economic impacts of the H-1B visa program. Most studies focus on the labor market outcomes of native-born workers, and use variation from aggregate changes to the H-1B cap or from H-1B lotteries. [Kerr and Lincoln \(2010\)](#) leverage aggregate variation in the H-1B cap, and find that growth in the skilled immigrant workforce generates sizable increases in innovation, with no evidence that natives are crowded out of jobs. [Peri et al. \(2013\)](#) complement this by also using H-1B cap variation to show that city-level increases in foreign STEM workers are associated with higher wages paid to native skilled workers

(both in STEM and non-STEM occupations). [Mayda et al. \(2018\)](#) implement difference-in-differences specifications around reductions in the H-1B cap, and find that the native-born employment did not rise in response. Furthermore, H-1B cap reductions altered the composition of new H-1B workers, increasing those in computer-related occupations, those from India, and those working at firms that had previously relied heavily on the program. [Kerr et al. \(2015b\)](#) use US Census Bureau’s worker-firm linked records and exploit cap-driven variation in aggregate Skilled immigrant employment. Results indicate that firm-level increases in skilled immigration raise overall employment, but that these increases are disproportionately concentrated amongst young skilled natives, while the employment share of older natives falls.

Other recent studies have utilized H-1B lotteries to identify causal impacts on workers and innovation. ([Peri et al., 2015](#)) leverage variation from the H-1B lotteries in fiscal years (FYs) 2008 and 2009 that produce negative supply shocks to H-1B workers. Results indicate such negative shocks are associated with reductions in foreign-born computer-related employment, with no corresponding increases in native-born computer employment. [Doran et al. \(2014\)](#) provide an important recent contribution by using individual patent and tax records to study the partial H-1B lotteries that took place in FYs 2005 and 2006.<sup>4</sup> They find no evidence that the lotteries affected patenting, and a moderate negative effect on overall employment in the firm. [Dimmock et al. \(2019\)](#) focus on start-up firms and show that those with a higher win-rate during the H-1B lotteries of FYs 2008, 2009, 2014, and 2015, were more likely to receive additional venture capital funding and/or a successful exit via IPO or other acquisition. They also examine patenting and find H-1B visa lottery winners receive more patents and patent citations.

Our paper expands this literature by bringing attention to firm behavior. Importantly, understanding the labor market outcomes of workers 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.<sup>5</sup> 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](#)

---

<sup>4</sup>In FYs 2005 and 2006, USCIS only performed a random lottery on applications received on the final receipt date – i.e. the date on which the number of applications received exceeded the cap. In FYs 2008, 2009, and 2014-2020, USCIS awarded all cap-subject H-1B applications via random lottery.

<sup>5</sup>For example, [Glennon \(2020\)](#) finds that H-1B restrictions leads multinational firms to offshore more jobs. Outside of the H-1B program, [Bloom et al. \(2019\)](#) show that part of the decline in employment associated with Chinese imports found in many influential studies, was actually due to firms changing industries.

et al., 2015a).

We contribute to the literature in several ways. First we do an original data construction, by matching the universe of I-129 petitions to all Compustat companies with positive employment from FYs 1999-2018, assembling a balanced panel of 1600 firms. As our data set contains the universe of publicly traded firms in the U.S., we speak to impacts on large employers. Second, relative to the studies cited above, we broaden the scope of the analysis by examining a wider set of outcomes, including employment, sales, profits, market value, and R&D expenditures. Third, while much of the literature has focused on earlier reductions in the cap (circa FY 2004) or initial lotteries (FYs 2006-2009), we focus on a more recent period of heightened H-1B demand (FYs 2010-2018). Finally, we document for the first time that firms respond differently to rationing of new H-1B, with a group of them increasing the retention/renewal of existing H-1B workers. Let us emphasize that the last finding has important implications for methodologies relying on cap or lottery-driven variation, as firms may respond differently in anticipation and in response of these events. Even when using lottery-based variation one must account for differential mitigation strategies to identify correctly the causal effect of H-1B workers change on firms' outcomes.

### 3 Hiring H-1B Workers

Firms wanting to use the H-1B program to hire a highly-skilled foreign worker begin the process by filing 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, including the occupation title, wage, starting date, and ending date. Firms must attest that they will pay the foreign worker the local prevailing wage or the employer's actual wage (whichever is higher), and that they will not displace American workers. Approved LCAs serve as a permission slip of sorts for firms to search for qualified foreign workers. Although important analyses, including [Kerr and Lincoln \(2010\)](#), have relied upon LCA data to approximate the demand for H-1B workers, this approach faces important limitations. Chief among them is the fact that the number of LCAs filed exceeds, systematically and by far, the number of foreign workers who receive job offers through the H-1B program. LCA data therefore provides an inaccurate measure of the H-1B applications and of the corresponding workforce.

Once a firm has identified a foreign worker it would like to hire, it must file an I-129 petition seeking H-1B status on the worker's behalf. Unlike the LCA, an I-129

is tied to a specific candidate. 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, an individual cannot apply for H-1B status on his or her own; a firm with an approved LCA must petition the government for H-1B status on his or her behalf. Moreover, the H-1B provides dual-intent status that allows firms to sponsor their H-1B employees for permanent residency (i.e. a green card). Employed H-1B workers might not want to jeopardize this possibility by participating in subsequent job searches.<sup>6</sup> 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. limits the number of new H-1B issued to employees of for-profit firms each year. Since FY 2006, this cap has been set at 65,000 general H-1B workers per year, plus 20,000 to workers who have obtained a Master’s Degree (or higher) from a US university. Employees of colleges, universities, and other non-profit research organizations are exempt from this cap. This paper focuses only on H-1B workers at publicly-traded for-profit firms.

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. As we show, this provides firms with an important margin 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. Such allocation process 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

---

<sup>6</sup>See [Sparber \(2019\)](#) or [Depew et al. \(2017\)](#) for a review.



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 of visa rationing, with the resulting tight hiring constraints, as “lottery years”. From the perspective of individual firms, a lottery year represents an exogenous drop in the probability of successfully matching with the chosen H-1B worker.

Use of the H-1B program varies substantially across firms. One expects that lottery years would particularly constrain the hiring for firms that use the H-1B program intensively, negatively affecting their growth and performance. To examine this hypothesis, we classify employers into users and non-users of the program. Users are firms that hired at least one new H-1B employee in the years prior to the onset of the rationing period (i.e. in FYs 2010-2013). Furthermore, we sub-divide this latter group of firms into two subgroups, *low users* and *high users*, according to whether they hired an above or below median number of H-1B workers relative to total company employment.

The variation over time in the severity of H-1B hiring constraints in “lottery years” and firms’ heterogeneous participation in the H-1B program suggest the use of a difference-in-difference estimation strategy. For our purposes, lottery years constitute a treatment period in which the probability of hiring (i.e. matching with the desired) H-1B workers was exogenously reduced for all firms. Firms that participate in the H-1B visa program (by hiring at least one new H-1B employee during the pre-treatment period) are considered the treatment group; firms that did not participate serve as a control.

We restrict our main analysis to FY 2010-2013, as the pre-treatment period, and FY 2014-2018, as the treatment period. We discard the first episode of rationing (FY 2008-2009) because it overlapped with both the Great Recession and with other important immigration policy changes that might have affected the demand for H-1B workers.<sup>7</sup> Our first outcome of interest is employment, which is directly affected by the rationing of H-1B visas. We also examine several other outcomes, such as sales, worker productivity,

---

<sup>7</sup>For example, 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. Regression results for this first episode of rationing are available upon request and support many of the results in this paper but are not informative.

profits, market value and R&D expenditures.

## 4 Data Sources and Description

### 4.1 USCIS Approved H-1B Petitions: FY 1999-2018

Our original data covers two decades of the H-1B program, namely FY 1999-2018. The main empirical results, however, are based on the period starting in FY 2010, which excludes the Great Recession and some years with important immigration policy changes that would introduce confounding factors. 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 perspective 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.<sup>8</sup>

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 year from FY 2009-2018. This dataset also includes the total number of H-1B petitions by request type for each firm and fiscal year. This allows us to build a complete firm-level longitudinal dataset tracking I-129 petitions through FY 2018.

Panel (a) and (b) of [Figure 2](#) 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 Datahub source in the red dashed line. Three things are worth noting. First, in the overlapping years the totals do not exactly coincide, but the discrepancy is small and the two series move in tandem. This provides reassurance that the data are consistent across sources. In the regression analysis, for the overlapping years (i.e. FY 2009-2012), we use FOIA data. From FY 2013 onward, we use the USCIS Data Hub.

Second, the data also reflect well-known stylized facts—the approved new H-1B issuances increased around the turn of the millennium, reflecting the increases in the H-1B cap and possibly also the Dot-Com boom. New H-1B issuances sharply contract follow-

---

<sup>8</sup>The approval/denial classification is based on USCIS' first decision but is not updated to reflect appeals and later changes. It does not include pending petitions.

ing the Dot-Com bust, and also begin to decline after the reduction in the cap circa 2004. New H-1Bs also follow a V-shaped pattern closely tracking the Great Recession subsequent recovery. From around 2013-onward, new H-1B issuances appear to level off, possibly as growing labor demand meets the fixed cap. Finally, the number of continuing H-1B issuances (panel b) appear to continuously grow over time, except for a slight dip during the Great Recession. This notably accelerates after 2013, during the lottery years, creating a sizable divergence in the number of new (roughly flat or declining) and continued (increasing) H-1B.

## 4.2 Compustat

We use Compustat data to measure firm-level outcomes. These data include information on all publicly traded firms covering the two decades between FY 1999 and 2018. Our sample retains only those firms which have positive employment in each year during this period, while those that ever report 0 employment or missing employment are dropped. The available outcome variables we study include total employment, sales, worker productivity (i.e. sales per employee), (EBITDA) profits, market value, and R&D expenditure.

The only firm identifying information available in both the H-1B and Compustat datasets is firm name. In order to merge the datasets, we performed a fully manual concordance of I-129s for all Compustat companies. In particular, we retained the sample of Compustat firms that had positive sales and employment over the FY 2009-2018 period. We then manually matched H-1B records by firm name.

As can be seen in [Figure 3](#) and [Table 1](#), 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 (that is, 2 in 3 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 exploded in FY2018. In this year there was also an uptick in new-employment denials.

Already from the raw data we can see an interesting new trend that will play an important role in our analysis. [Figure 3](#) shows the number of new (in blue solid line) and continuing (red dashed line) H-1B applications by firms in the Compustat database, that we use for our analysis. We see that since FY 2014, there has been a substantial diver-

gence between new H-1B employees, whose number declines and continuing-employment H-1B workers, whose number increases.<sup>9</sup> One interpretation of this divergence is that as new H-1B became harder to obtain as the market became tight and each firm faced a decline in probability of getting them, those that could increased the rate of renewal of existing H-1B generating an increase in their number.

Panels (a) and (b) of [Figure 4](#) show trends in the number of new and continuing H-1B petitions separately for the top four H-1B employers over this period (Infosys Ltd., Tata Group, Cognizant Tech Solutions, and Wipro Ltd.) in Panel (a), and in the remaining Compustat firms in Panel (b). Several points are worth noting. First, the reduction in new H-1B employment between 2012 and 2018 observed in [Figure 3](#) is largely driven by the reduction in approvals for the top firms. New approvals for employment at those firms dropped from about 20,000 in FY2014 to a few thousand in FY2018. This may be a consequence of the larger scrutiny in the application of these firms, resulting in less approvals. Second, the bottom figure shows only a small reduction in new approvals since FY 2014 for companies excluding these top four receivers. Third, both top users and other Compustat firms experienced a large rise in approvals for cap-exempt workers continuing on H-1B status. These approvals have dropped drastically for the large companies in FY 2018, possibly suggesting an extension of the scrutiny to renewals of H-1B petitions for this group of firms.

### 4.3 Descriptive Statistics

Our final sample consists of 1600 firms with consistently positive employment in each fiscal year from 1999 through 2018. Hence, our firm sample is a panel of established firms with at least two decades of continuous history. Company closings, acquisitions, or start-ups that occur during our time period are therefore excluded. Thus, our analysis does not change in composition over time.<sup>10</sup>

We classify companies into users and non-users of the program: Users are defined as firms that employed at least one new H-1B worker in the years immediately prior to the lottery rationing period, namely in FY 2010-2013. Then we divide users into two subgroups – low users and high users – based upon whether the firm employed an above or below median number of H-1B workers relative to its total employment in the same pre-treatment period.

---

<sup>9</sup>This tendency is also noticeable, but less extreme, in the brief rationing period FY2007-2008.

<sup>10</sup>Issues of closures in response to H-1B program variations are analyzed in [Dimmock et al. \(2019\)](#) and [Glennon \(2020\)](#).

**Table 2** summarizes firm characteristics of non-users, low users, and higher-users of the H-1B program. Notice that the sample is almost equally split between users (52% of total employers considered) and non-users (48% of total). Users, in turn, are divided between high and low users evenly (50-50) by construction. On average non-users are much smaller than high users, which are in turn smaller than low users, along several dimensions. Employment, for example, averages 6,600 at non-using firms compared to 21,100 for high users and 54,300 for low users. The same picture emerges in terms of average sales: \$2.0 billion for non-users, \$10.9 billion for high users, and \$17.4 billion for low users. Interestingly, profits and market value are much higher for firms that hire H-1B workers compared to those that do not. However, differences between low and high use firms are small.

We also observe important differences in industry composition. About 45% of companies using the H-1B program are manufacturing firms, compared to only 29% among the non-users. Additionally, non-users are much more likely to be in finance & real estate than users of the program. Importantly, high users are much more likely to belong to the computer industry (10% versus 2% for the low or no users in Compustat).<sup>11</sup>

By construction, new-employment approvals (obtained in the period FY2010-2013) vary across the three groups. While non-existing for the non-users, new-employment approvals are 0.3% of the stock of employees for the low users and four times higher (1.2%) among the high users. In interpreting the magnitudes, it is important to note that we are comparing a flow variable (new H-1B hires) to a stock variable (employment in the company). Thus the inflows of new H-1B hires is significant for the high users. While the characteristics of the treated and control groups are different, we will show in the next section that the event analysis does not show any significant pre-treatment differential trend in outcomes between the two groups, implying that the identification assumption of our strategy are not violated.

## 5 Econometric Specification and Main Results

We perform an event study comparing the evolution of (the logarithm of) employment and other firm-level variables between firms that participated and those that did not participate in the H-1B program during the pre-rationing period. The first question we address is whether severe rationing led companies that use the H-1B program to

---

<sup>11</sup>The computer industry is an aggregation of several NAICS industry codes.

experience slower employment growth relative to non-users. This would happen if firms could not find replacements for H-1B workers, or if H-1B workers are complementary to other workers in the firm. As a second step, we analyze whether employment growth was different for very high users compared to very low users of the program.

## 5.1 Event Study and Employment Results

We start by estimating the following two models, where  $y_{ijt}$  is the outcome for firm  $i$  in industry  $j$  and fiscal year  $t$ :

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

and

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

In both models the dependent variable is the log of employment (or other firm-level variables), the term  $\gamma_i$  captures a set of firm-level fixed effects accounting for time-invariant heterogeneity across firms, and the term  $\gamma_{jt}$  is a set of sector-by-year effects, capturing sector-specific shocks over time.<sup>12</sup> The coefficients of interest are  $\beta_t$  in Equation (1), capturing yearly differences in outcomes between users and non-users.  $\beta_t^{Low}$  and  $\beta_t^{High}$ , in Equation (2), capture the relative performance of low and high users, respectively, relative to non-users in each year. Fiscal year 2013 is the omitted category, such that other coefficients are interpreted relative to 2013. The inclusion of firm fixed-effects implies that  $\beta_t$ ,  $\beta_t^{Low}$ , and  $\beta_t^{High}$  will be identified by within-company employment changes among users, low users, and high users (respectively), relative to non-users.

Figure 5 displays the estimates of  $\beta_t$  for Equation (1), 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 years FYs 2010-2013. Qualitatively, the series perhaps shows a slight upward trend, suggesting that during the recovery from the Great Recession, H-1B users may have hired at the same pace, or faster, as non-users. Second, during the

---

<sup>12</sup>When we examine firm-level outcomes for which 0 is a common value, we employ a different functional form that is well defined for zero values, namely the inverse hyperbolic sine, which is defined as  $\sinh^{-1}(x) = \ln(x + \sqrt{x^2 + 1})$ .

lottery years total employment grew significantly less at firms using the H-1B program. By FY 2018, H-1B users ended up with significantly lower employment relative to non-users, as percent of their 2013 levels—about a 10 log-point deficit (roughly 10 percent), implying a slower growth by about 2 percent each year. This is larger than the direct effect of losing H-1B workers as the H-1B new hires of program users during FYs 2010-13 averaged 0.6 percent of their employment. This steep decline in employment growth is consistent with complementarity between H-1B workers and other types of workers. Our results imply that participating firms suffered employment losses beyond the reduction in their H-1B workforce.

**Figure 6** presents the estimates of the coefficients  $\beta_t^{Low}$  and  $\beta_t^{High}$  from **Equation (2)**. The top panel (a) shows the estimates and 95% confidence interval for the set of the coefficients  $\beta_t^{Low}$  that imply a comparison of non-users to low users. Analogously, the bottom Panel (b) shows the estimates of coefficients  $\beta_t^{High}$  which imply a comparison between non-users and high users.

An interesting pattern emerges from these Figures. For the low users we see clear evidence that H-1B rationing had a negative effect on employment growth, relative to non-users. In fact, the point estimates imply that the peak effect is equal to a 15 log-point reduction in employment four years into the rationing period. The 95% confidence interval includes a 10 log-point loss, which was the estimated average effect. Low users appear to be strongly affected by rationing, and lose more than just their pre-existing H-1B workforce, where new H-1B hires averaged 0.3 percent of employment during the pre-rationing period.

More surprisingly, the bottom figure suggests that employment among high users was less negatively affected. Although the pattern depicted in the figure is similar to that for low users, the point estimates are 1/3 to 1/2 of their size, with the reduction in employment peaking at 5 log-points. While the 95% confidence interval also includes a peak loss of 10 log-points, the average loss in employment for this group seems equal to, or smaller, than the direct loss in H-1B hires (equal to about 1.2 percent of employment during pre-rationing years). It appears that low H-1B users suffered a significant drop in employment as a consequence of rationing, presumably because they were unable to replace H-1B workers and may have also lost other workers due to complementarity effects. In contrast, high users found a way to attenuate such losses. Before turning to possible explanations for these findings, we show formal difference-in-differences estimates for employment and a wider range of firm-level outcomes.

## 5.2 Difference-in-Difference Estimates. All Outcomes

Table 3 reports difference-in-difference estimates (similar to Equation (1) and Equation (2)) for a set of company outcomes including employment, sales, profits, market value and R&D spending. For the sake of brevity, we only report the relevant coefficient estimates, namely, the interaction of the indicator for “Lottery Years” (rationing period) and the indicator for H-1B users (in the top panel), or indicator variables for “low users” and “high users” (in the lower panel).

These estimates provide an average estimate of the effects of rationing (lottery years) on various outcomes for firms that employ H-1B workers.<sup>13</sup> Point estimates indicate a 6 log-point reduction, after rationing during lottery years, in the total sales of H-1B users relative to non-users. This effect is similar to the drop in employment, implying stability in sales per worker, shown in column (3). Furthermore, the profits and market value of these firms fell by about 9 log points and R&D expenditures fell by 11 log points.<sup>14</sup> The H-1B rationing produced declines in profitability, market valuation and R&D that appear to be larger proportionally than the employment decline. This finding is consistent with the interpretation that H-1B workers were associated with higher than average value-added and were more often involved in R&D activities than the average worker.

The bottom part of Table 3 further distinguishes users by low and high intensity; non-users continue to serve as the comparison group. We observe that low users experienced large reductions in all outcomes. Effects on profits, market value and RD are more than proportional to the effects on employment. Relative to non-users, low users had 9-10 log-point reductions in employment and sales after 2013, and 11-13 log-point lower profits, market value and R&D spending. In comparison, the effects of rationing on high users of the H-1B program, are more muted, especially in terms of employment and sales. Qualitatively, the point estimates suggest negative effects of rationing for high users, but the effects are either small or imprecisely estimated.

One potential margin of adjustment to the reduced hiring of skilled workers would be for firms to substitute H-1B workers with machinery/technology (Lewis, 2011; Clemens et al., 2018). If the tasks performed by H-1B workers could be easily substituted by computers or other capital equipment, the drop in employment should be accompanied

---

<sup>13</sup>Note that this coefficient is an estimate of the average effect after 2013, not of the peak effect (estimated in the event study for year 2018)

<sup>14</sup>Only about 40% of the companies in our sample report positive R&D expenditures. To accommodate the large prevalence of zeros, we apply the inverse hyperbolic transformation to this variable. The point estimates can be interpreted similarly to the regular log transformation.



by an increase in capital investment. On the other hand, because H-1B workers are somewhat skilled and may perform tasks that are difficult to automate, substitution by new technology may not be feasible and it may even be the case that H-1B workers and technologically intensive capital equipment are complementary (Autor, 2015). In this case, visa rationing may reduce both H-1B employment and capital investments. To test these hypotheses, we extend our analysis by including two new dependent variables: annual capital expenditure flows and values of (the stock of) property, plant and equipment (PPP) of the firm.

Table 4 reports the estimates. Column (1) suggests that H-1B rationing caused about a 6% reduction in the stock of capital for companies using the H-1B program. Once again, the reduction was larger for low-intensity users than for high-intensity users. Column (2) shows that these reductions were driven by declines in investment (capital expenditures) among companies participating in the H-1B program. In this case we see a more significant 12 percent drop in investment associated with rationing, and less of a difference between high and low users of the program<sup>15</sup> Overall, these estimates suggest that capital and H-1B workers are complementary. Rationing of H-1B visas appears to have triggered reductions in capital investments.

We conclude this section with an important check. As documented earlier, a few companies account for a large share of all H-1B employment in many years.<sup>16</sup> Thus, we are concerned that the negative effects of rationing may stem primarily from these few firms. To evaluate this concern, we eliminate the ten companies that received the highest number of total (new and continuing) H-1B workers over our sample period. As shown in Table 5, the exclusion of these firms has no effect on our estimates and implies that the rationing of H-1B visas significantly affected hiring and outcomes of all users, not just the largest ones.

---

<sup>15</sup>Columns 3 and 4 confirm the findings when we use the inverse hyperbolic sine transformation of the corresponding dependent variable.

<sup>16</sup>Figure 4 plotted H-1B approvals at the top 4 receivers (Infosys Ltd., Tata Group, Cognizant Tech Solutions, and Wipro Ltd.) in our dataset. In FY 2014, these firms accounted for about 60% of all H-1B approvals.

## 6 Mitigating Mechanism: H-1B renewal intensity and anticipatory “buffering”

The results in the previous sections reveal that H-1B rationing affected the outcomes of high and low H-1B users non-monotonically. In particular, during the visa rationing period, companies defined as high-intensity H-1B users experienced smaller relative declines in employment than low users. To better understand this surprising finding, we next explore mitigating mechanisms that can account for the non-monotonic pattern. We argue that high H-1B users have a greater need for H-1B workers, or they possess a more sophisticated understanding of the intricacies related to obtaining H-1B visas. As a result, they may have developed the following mitigation strategy: these employers created “buffer stocks” of new-employment H-1B workers that could be renewed (for a second 3-year period) in the event of rationing. To the extent that these companies may have changed their behavior in *anticipation* of the onset of the rationing period, the causal interpretation of our difference-in-difference estimates may need to be qualified. Next, we discuss the buffering behavior in the pre-rationing period.

### 6.1 Continuing/Renewal of H-1B workers

In order to understand the mechanisms of adjustment at the firm level, it helps to partition H-1B workers in two groups: (1) new hires (in their initial 3-year appointment) and (2) continuing workers (past their initial 3-year appointment and renewed by their previous employer or a different one). For each company and year we define the stock of (new and continuing) H-1B workers as the sum of the flows of approved applications of *both* types of workers over the previous 3 years.<sup>17</sup>

Specifically, we define the estimated stock as:

$$StockH1B_{ijt}^{New} = NewH1B_{ijt} + (1 - \delta)NewH1B_{ijt-1} + (1 - \delta)^2NewH1B_{ijt-2}$$

$$StockH1B_{ijt}^{Cont} = ContH1B_{ijt} + (1 - \delta)ContH1B_{ijt-1} + (1 - \delta)^2ContH1B_{ijt-2}$$

$$StockH1B_{ijt} = StockH1B_{ijt}^{New} + StockH1B_{ijt}^{Cont},$$

where  $NewH1B$  represent the annual flows of approved H-1B petitions for newly em-

---

<sup>17</sup>For example, the estimated stock in 2010 is the sum of the number of approved new workers in 2010, 2009, and 2008, and the number of renewed workers in 2010, 2009, and 2008.

ployed workers and *ContH1B* denote annual approvals of renewals for continuing H-1B workers. This procedure assumes that those workers who renew their H-1B status do so in their third year. We also allow for a yearly attrition rate for both new and continuing H-1B workers,  $\delta$ , that we set at 6%, based on the estimates reported by [Depew et al. \(2017\)](#). This attrition rate reflects people separating from employers or returning to their home country before exhausting the 3-year initial or final appointment.<sup>18</sup> Measurement error in the stock may be generated by the fact that not all new-employment H-1Bs are used for the whole 3-year period and that renewals can take place earlier than 3 years after issuance and, occasionally, extend beyond the standard 3-year period.

We examine the evolution of these stock variables before and during the lottery years using an event study specification that distinguishes between high and low H-1B users. To do so we restrict ourselves to H-1B program users only, and we estimate a flexible model where the dependent variable,  $StockH1B_{ijt}^c$  is the stock of category  $c \in \{New, Cont\}$  H-1B workers at company  $i$  in sector  $j$  and year  $t$ , normalized by average company employment over the period FY 2010-2013 ( $Emp_i$ ). That is,

$$\frac{StockH1B_{ijt}^c}{Emp_i} = \alpha + \sum_{r \neq 2013} \beta_r^{High} (I_i^{High} \times \mathbb{1}(year_r = t)) + \gamma_i + \gamma_{jt} + \varepsilon_{ijt} \quad (3)$$

$\gamma_i$  are firm-level fixed-effects and  $\gamma_{jt}$  are industry-by-year indicators. Since we only include firms participating in the H-1B program, indicator variable  $I^{High}$  captures the difference between high and low users. We interact the high user indicator with dummies for each year, omitting year 2013 (whose coefficient is normalized to zero). Similar to previous specifications, coefficient  $\beta_t^{High}$  represents the difference between high and low users in the stock of H-1B workers in the corresponding year.

**Figure 7** plots the coefficients  $\beta_t^{High}$  (and the corresponding 95% confidence intervals) for the stock of new (Panel a) and continuing (Panel b) H-1B workers. The vertical scale measures the number of H-1B per thousand employees, and the coefficients represent the difference over time in the stock of high users relative to low users, after controlling for firm fixed-effects and sector-year trends. Three findings stand out. First, during the lottery years (2014-18) a large and growing gap in continuing H-1B workers emerges, as high users accumulate larger stocks of this type of workers relative to low users (panel a). By 2018, the gap is substantial, at 40 H-1B workers per thousand employees. In the pre-rationing years 2010-2013, the companies defined as high users had relatively smaller, but quite similar, stocks of continuing H-1B workers relative to low users. While the

---

<sup>18</sup>We also construct an alternative measure assuming no attrition ( $\delta = 1$ ).

figure suggests a possible slightly steeper trend for high users already before 2013, the gap grew intensely after 2013, consistent with the idea that high users increased renewal rates to a larger degree than low users, or succeeded in poaching (cap-free) continuing H-1B workers from other firms during the rationing period. As we discuss next, the small pre-2013 differential trend may be part of the “anticipation” strategy that we describe below.

Second, in terms of stocks of new H-1B, the rationing seems to affect high and low users in the same way, as one would expect given that visa petitions for new-employment were allocated by means of lotteries. Coefficient estimates in panel (b) are never significantly different from 0 and are very close to zero in lottery years 2014-2018. This confirms that the lottery years constituted a uniformly proportional rationing on all companies hiring new H-1B workers, consistent with a simple and uniform reduction of matching probability. The Figure also reveals a differential pre-2013 trend in the stock of new H-1B workers (relative to employment), suggesting that high users may begin to build a “visa buffer stock” by increasing the hiring of new-employment H-1B workers (relative to low users) that could be renewed in the event of a future rationing episode.

We will discuss in the next section how this mechanism may have worked and if it affects the interpretation and identification of the difference-in-differences analysis of outcomes, which relies on similar trends of hiring between H-1B users and non-users before rationing. We also note that the mechanism of increased renewal (or poaching from other firms) of continuing H-1B workers provides only temporary relief for the effects of visa rationing due to the one-time renewal limit for H-1B visas (with the exception of workers waiting for permanent residence).

## **6.2 Anticipatory H-1B Excess Hiring**

The increased ability of high users to hire continuing H-1B workers after 2013, could imply that they hire some of them away from low users or that they are more successful in retaining/renewing visa holders within the company. Close employer-employee ties might make hiring from other firms difficult in practice. H-1B workers might not want to jeopardize the possibility that their employers will sponsor their permanent residency green card, for example. The existence of firm-specific human capital may also make such a move less attractive.

High-intensity users of the H-1B program may have been sophisticated enough to monitor the gradual increase in the demand for visas (e.g. by keeping track of the

evolution of the final receipt date as in [Figure 1](#)), anticipating the arrival of a period of severe rationing one or two years in advance. For instance, the trend towards earlier final receipt dates in FYs 2012 and 2013, could have easily allowed a company to predict the rationing that began in FY 2014, and begin preparations to mitigate the effects of the upcoming H-1B visa shortages. Such a firm, could have hired an excess number of new H-1B workers in the years immediately before 2014, when visas were still being allocated on a first-come, first-served basis.

To examine the presence of anticipatory “excess hiring” in the 2010-2013 period by high users relative to low users, we proceed as follows. First, we build a firm-level predictor of the demand for new-employment H-1B workers. Namely, we construct a [Bartik \(1991\)](#)-style predictor interacting the firm’s base-year H-1B intensity with the overall total employment in firm  $i$  in year  $t$ . Choosing 2000 as the base year, the product of these terms is given by  $\left(\frac{NewH1B2000_i}{Emp2000_i} \times Emp_{it}\right)$ . Second, we estimate the difference between the observed flow of newly employed H-1B workers in the company (normalized by total company employment),  $h_{it}$ , and the predicted flow as the residual of the following regression model:

$$h_{it} = \alpha + \beta \left( \frac{NewH1B2000_i}{Emp2000_i} \times Emp_{it} \right) + u_{it}. \quad (4)$$

The residual of this regression ( $\hat{u}_{it} = h_{it} - \hat{h}_{it}$ ) is an estimate of the number of new H-1B workers hired by the company in a specific year in excess of the baseline “usage intensity” of these workers based on previous years. Accordingly, we refer to  $\hat{u}_{it}$  as the firm’s excess H-1B hiring.

We then turn to a short panel of companies, that includes all the H-1B users in the sample, covering the pre-rationing period (FY 2010 through 2013). To test whether high use companies engaged in H-1B excess hiring in 2012 and 2013 (right before the rationing period), we estimate:

$$\hat{u}_{it} = \alpha_i + \beta_1 I(2012 \leq t \leq 2013) + \beta_2 I(2012 \leq t \leq 2013) I_i^{High} + \gamma Emp_{it} + \varepsilon_{it}. \quad (5)$$

If anticipatory hiring behavior occurred for high use firms, we expect we expect  $\beta_2$  – the coefficient on an indicator for FY 2012 and 2013 interacted with the indicator of High H-1B use – to be positive. If such behavior took place, it would imply that high users would exhibit a different hiring behavior, relative to low users, in the years immediately prior to the rationing period because they would be building their “visa

buffer” to cushion the impact of the upcoming rationing event.

[Table 6](#) reports the estimate for  $\beta_2$ , which will capture whether high users were building visa buffer stocks in FY 2012 and 2013, relative to low users. The estimates in column 1 indicate that high H-1B users did in fact accumulated an unusually high number of new-employment H-1B visas in the years immediately prior to the rationing episode. This is true even after we control for current company employment (Column 2). The magnitude of the estimates say that high H-1B users hired 12 to 13 extra h-1B workers (per 1000 employees) annually, relative to low users, which amounted to an increase in H-1B workers equal to 1.2 percent of their employment. The small estimated coefficient of the 2012-13 dummy also implies that low users did not undertake any excess hiring in the pre-2013 period, relative to predicted values. As a result, during the rationing period, low users of the H-1B program had more limited access to the mitigation strategy consisting of increasing renewal rates of their H-1B workers.

The anticipatory nature of the actions taken by the high users implies that some caution is needed in interpreting causally the difference-in-differences estimates presented in [Section 5.2](#). It is our view that, since low-use firms did not engage in anticipatory behavior, the estimates for that group (the only one with negative significant impacts) can still be given a straightforward causal interpretation. Moreover, the anticipatory excess hiring for high users is in line with the slight differential upward trend visible in [Figure 7](#) for high users (panel b).

## 7 Conclusions

The H-1B program provides an important channel through which U.S. firms can hire skilled foreign labor. Though there has always been a cap on the number of new H-1Bs issued to employees of private for-profit firms each year, that cap has been particularly binding since fiscal year 2014. All cap-subject H-1B visas have been allocated by lottery since that year. Due to this rationing, employers that use the H-1B visa program have experienced a uniform decline in the probability of matching with a H-1B employee.

We find that the rationing episode that began in 2014 led to a reduction in employment, sales, profits, market value, R&D, and capital expenditures for the firms that were using the H-1B program (prior to 2010) relative to firms that did not (conditional on firm fixed-effects and industry trends).

A bit surprisingly, the effects of rationing are non-monotonic in the intensity of use of the visa program. Specifically, firms that had used the the H-1B program more

intensively experienced fewer reductions in performance than less intensive users. We argue that more intensive users of the program developed strategies to attenuate the effects of the rationing of new H-1B visas. In particular, these firms increased the renewal of H-1B workers already in the company (or successfully poached them from other firms), taking advantage of the fact that these status extensions are effectively automatic and do not count toward the annual cap on (new) H-1B visas. Secondly, these firms successfully anticipated the arrival of the rationing period and prepared for it by creating a “buffer” of H-1B workers, that is, they hired new H-1B workers, in excess of the predicted value, in the years right before the rationing, when cap-bound, new-employment visas could be obtained on a first-come, first-served basis.

This mitigation strategy proved effective in limiting the economic impact of H-1B rationing on these companies. However, the effectiveness of this strategy is only temporary because, typically, foreign workers can hold H-1B status for a maximum of six years. Thus, the attenuating effects of the “visa buffer” should be practically disappearing around 2020.

## References

- Autor, David H, “Why are there still so many jobs? The history and future of workplace automation,” *Journal of Economic Perspectives*, Summer 2015, 29 (3), 3–30.
- Bartik, Timothy J, “Who benefits from state and local economic development policies?,” 1991.
- Bloom, Nicholas, Kyle Handley, André Kurmann, and Philip Luck, “The impact of chinese trade on us employment: The good, the bad, and the apocryphal,” in “in” 2019.
- Clemens, Michael A, Ethan G Lewis, and Hannah M Postel, “Immigration restrictions as active labor market policy: Evidence from the mexican bracero exclusion,” *American Economic Review*, 2018, 108 (6), 1468–87.
- CompeteAmerica, “324 Signatory Business Letter to President and Secretaries-Nonimmigrant Visa Facts,” May 2020.
- Depew, Briggs, Peter Norlander, and Todd A. Sørensen, “Inter-firm mobility and return migration patterns of skilled guest workers,” *Journal of Population Economics*, April 2017, 30 (2), 681–721.
- 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.
- 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.
- Lewis, Ethan, “Immigration, skill mix, and capital skill complementarity,” *The Quarterly Journal of Economics*, 2011, 126 (2), 1029–1069.
- 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.
- Peri, Giovanni, 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.

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 in the H-1B Application Period



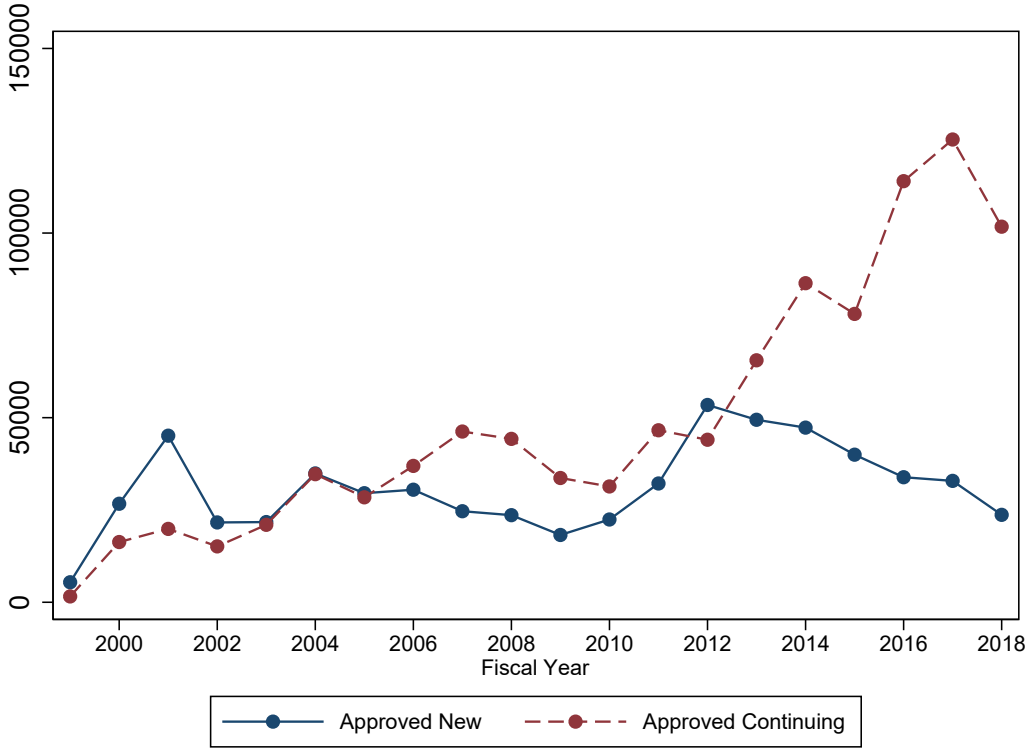
Notes: Number of days between the first week of April (when application window opens) and final receipt day (when USCIS has received enough applications to meet the statutory cap). Zero values correspond to years when USCIS allocated all cap-bound H-1Bs by lottery. In other years, USCIS allocated H-1Bs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

Figure 2: Aggregate I-129s. FOIA and DataHub comparison



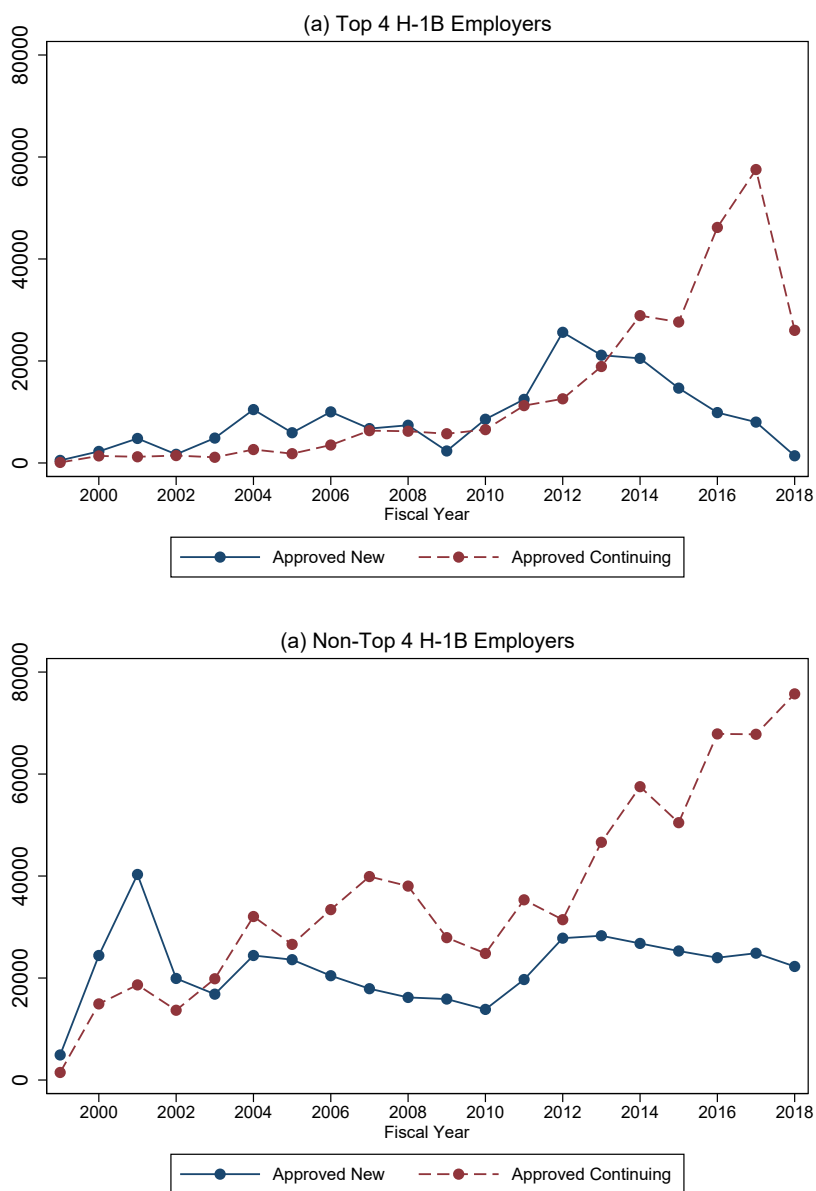
Note: Figures show aggregate totals of approved new H-1B petitions in panel (a), and approved continuing H-1B petitions in panel (b), by fiscal year. Blue series shows data from USCIS I-129 FOIA data. Red series displays data from the USCIS Datahub.

Figure 3: 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 4: Approved H-1B Petitions at Four Largest H-1B Employers (Top Panel) and Other Compustat Firms (Bottom)



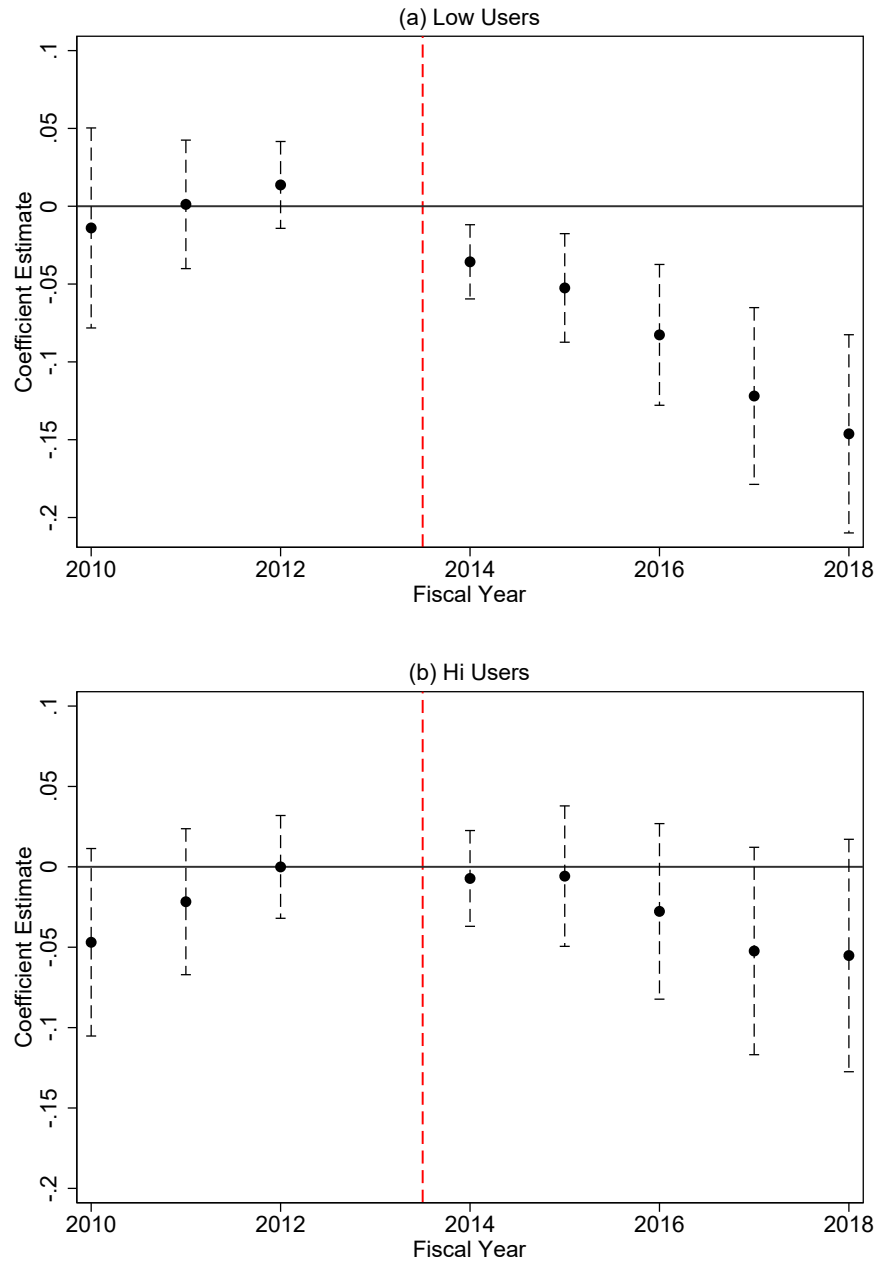
Notes: Figures show aggregate totals of approved new and continuing H-1B petitions for the sample of Compustat firms, by fiscal year. Blue series shows approved new H-1B petitions, while the red dashed series shows approved continuing H-1B petitions. Totals for the top four H-1B employers, in panel (a),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

Figure 5: 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.

Figure 6: Event Study Comparing Total Employment at Firms that Employ H-1B Workers to Those that Do Not. Results by H-1B Usage Intensity.



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. Top panel compares low H-1B use firms to no-use firms. Bottom panel compares high use firms to no-use firms. Point-estimates are relative to FY 2013, the omitted category. 95% Confidence intervals are shown in dashed lines.

Figure 7: Event Study Comparing the Ratio of the Stock of H-1B Workers at High use vs Low use Firms



30

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.



Table 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.

Table 2: Firm Characteristics by Intensity of Firms' H-1B Employment.

	(1)	(2)	(3)
	Non-Users	Low Use	High Use
Employment (thousands)	6.6	54.3	21.1
Sales (\$ Million)	2,042.0	17,396.2	10,893.4
Profits (\$ Million)	561.0	5,863.4	4,675.7
Market Value (\$ Million)	1,716.1	13,754.9	14,200.5
Agr. Mining (%)	4.27	4.76	6.33
Construction (%)	1.78	1.32	0.53
Manufacturing (%)	29.06	43.65	45.65
Trade & Trans. (%)	22.3	28.31	8.18
Fin. & Real Estate (%)	32.98	11.11	14.78
Bus. Services (%)	6.88	8.2	13.46
Other (%)	0.12	0.79	0.79
Computer Ind. (%)	2.61	1.85	10.29
USA HQ (%)	92.53	87.83	93.14
Approvals New 2010-13	0	10.5	167.8
Approvals Cont 2010-13	0.6	35.8	219.8
New Approvals/Emp201013 (%)	0	0.3	1.2
Number of Firms	843	378	379

Notes: Employment in thousands of people. Nominal sales, profits, and market values are in millions of dollars.

Table 3: The Effect of Rationing on Company Outcomes. DiD Estimates.

Treatment: FY2014-2018						
Control: FY2010-2013						
Dep Var: ln of	(1) Emp	(2) Sales	(3) Sales/Emp	(4) Profits	(5) MV	(6) <i>R&amp;D</i>
Lottery Years $\times$ User	-0.05 (0.02)	-0.06 (0.03)	-0.01 (0.02)	-0.09 (0.03)	-0.09 (0.04)	-0.11 (0.05)
Lottery Years $\times$ LowUser	-0.09 (0.03)	-0.11 (0.03)	-0.01 (0.02)	-0.11 (0.03)	-0.12 (0.04)	-0.13 (0.07)
Lottery Years $\times$ HighUser	-0.01 (0.03)	-0.01 (0.04)	0.00 (0.03)	-0.08 (0.04)	-0.07 (0.05)	-0.09 (0.05)
Y Mean	1.08	7.04	5.92	5.55	7.28	1.49
Firms	1600	1598	1598	1501	1512	1600
N	14400	14301	14301	12809	12865	14400
Firm FE	yes	yes	yes	yes	yes	yes
Industry-year FE	yes	yes	yes	yes	yes	yes

Notes: The top panel reports results from a regression model that classifies companies as users or non-users of the H-1B program. Users are companies that employed at least one cap-subject H-1B worker in FY 2010-2013. The regression model in the bottom panel further partitions H-1B employers into Low and high users of the program based upon whether the firm employed an above or below median number of new H-1B workers. The dependent variable is the log of the corresponding outcome, except in column 6 where we apply the inverse hyperbolic sine transformation to account for the large prevalence of zero values (about 60%). Firm fixed-effects included (defined by *gvkey*). *Rationing* is an indicator for the treatment fiscal years 2014-2019. During this rationing period, all H-1B petitions for new employment were allocated through lottery. The sample only contains companies with positive employment in every fiscal year for period FY1999-2018. Firms that hired zero cap-subject new H-1B workers prior to the rationing period serve as the control group and reference category in each regression. Standard errors clustered at the firm level.

Table 4: The Effect of Rationing on Capital Expenditures. DiD Estimates.

Treatment: FY2014-2018				
Control: FY2010-2013				
Dep Var:	(1)	(2)	(3)	(4)
	ln PPE	ln CapEx	asinh PPE	asinh CapEx
Lottery Years $\times$ User	-0.06 (0.03)	-0.13 (0.04)	-0.04 (0.04)	-0.10 (0.04)
Lottery Years $\times$ LowUser	-0.10 (0.04)	-0.13 (0.04)	-0.08 (0.04)	-0.11 (0.05)
Lottery Years $\times$ HighUser	-0.01 (0.05)	-0.12 (0.05)	-0.00 (0.05)	-0.09 (0.05)
Y Mean	5.55	3.75	6.00	4.28
Firms	1539	1522	1600	1600
N	13740	13520	14400	14400
Firm FE	yes	yes	yes	yes
Industry-year FE	yes	yes	yes	yes

Notes: The top panel reports results from a regression model that classifies companies as users or non-users of the H-1B program. Users are companies that employed at least one cap-subject H-1B worker in FY 2010-2013. The regression model in the bottom panel further partitions H-1B employers into Low and high users of the program based upon whether the firm employed an above or below median number of new H-1B workers. The dependent variables are the natural log ( $ln$ ) or inverse hyperbolic sine ( $asinh$ ) of property, plant and equipment (PPE, a stock variable) or capital expenditures ( $CapEx$ , an annual flow variable). Firm fixed-effects included (defined by  $gvkey$ ). *Rationing* is an indicator for the treatment fiscal years 2014-2019. During this rationing period, all H-1B petitions for new employment were allocated through lottery. The sample only contains companies with positive employment in every fiscal year for period FY1999-2018. Firms that hired zero cap-subject new H-1B workers prior to the rationing period serve as the control group and reference category in each regression. Standard errors clustered at the firm level.

Table 5: The Effect of Rationing on Company Outcomes. DiD Estimates. Excludes Top Users.

Treatment: FY2014-2018						
Control: FY2010-2013						
Dep Var: ln of	(1)	(2)	(3)	(4)	(5)	(6)
	Emp	Sales	Sales/Emp	Profits	MV	<i>R&amp;D</i>
Lottery Years $\times$ User	-0.05 (0.02)	-0.06 (0.03)	-0.01 (0.02)	-0.10 (0.03)	-0.09 (0.04)	-0.08 (0.06)
Lottery Years $\times$ LowUser	-0.09 (0.03)	-0.11 (0.03)	-0.01 (0.02)	-0.11 (0.03)	-0.12 (0.04)	-0.17 (0.07)
Lottery Years $\times$ HighUser	-0.02 (0.03)	-0.01 (0.04)	0.00 (0.03)	-0.08 (0.04)	-0.07 (0.05)	-0.02 (0.07)
Y Mean	1.07	7.03	5.92	5.54	7.26	3.71
Firms	1595	1593	1593	1496	1507	629
N	14355	14256	14256	12764	12820	5382
Firm FE	yes	yes	yes	yes	yes	yes
Industry-year FE	yes	yes	yes	yes	yes	yes

Notes: The top panel reports results from a regression model that classifies companies as users or non-users of the H-1B program. Users are companies that employed at least one cap-subject H-1B worker in FY 2010-2013. The regression model in the bottom panel further partitions H-1B employers into Low and high users of the program based upon whether the firm employed an above or below median number of new H-1B workers. The dependent variable is the log of the corresponding outcome, except in column 6 where we apply the inverse hyperbolic sine transformation to account for the large prevalence of zero values (about 60%). Firm fixed-effects included (defined by gvkey). *Rationing* is an indicator for the treatment fiscal years 2014-2019. During this rationing period, all H-1B petitions for new employment were allocated through lottery. The sample only contains companies with positive employment in every fiscal year for period FY1999-2018. We omit the top ten H-1B employers, five of which recorded positive employment for the entire sample period. Firms that hired zero cap-subject new H-1B workers prior to the rationing period serve as the control group and reference category in each regression. Standard errors clustered at the firm level.

Table 6: H-1B Buffering: De-Trended Approved H-1B Petitions for New Employment.

Treatment: FY2012-2013		
Control: FY2010-2011		
	(1)	(2)
D201213	0.119 (0.229)	-1.441 (1.386)
D201213 $\times$ HighUser	12.200 (6.004)	13.207 (6.593)
Employment		0.772 (0.592)
Obs.	3,028	3,028
Firms	757	757

Notes: The dependent variable is the de-trended number of approved H-1B petitions relative to company employment, computed as the residual from a regression of new H-1Bs (relative to company employment in year 2000) on a Bartik variable measuring the interaction between a firm's year-2000 H-1B dependence and year  $t$  total employment.  $D201213$  is an indicator for years 2012 and 2013 – years immediately preceding the H-1B rationing period. Regressions do not include non-users, and thus only compare high users to low users. Standard errors are clustered at the firm level.