I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 11345

The Effect of the H-1B Quota on the Employment and Selection of Foreign-Born Labor

Anna Maria Mayda
Francesc Ortega
Giovanni Peri
Kevin Shih
Chad Sparber

FEBRUARY 2018

I Z A Institute
of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 11345

The Effect of the H-1B Quota on the Employment and Selection of Foreign-Born Labor

Anna Maria Mayda
Georgetown University and IZA

Francesc Ortega
Queens College CUNY and IZA

Giovanni Peri
UC Davi, NBER and IZA

Kevin Shih
Rensselaer Polytechnic Institute

Chad Sparber
Colgate University and IZA

FEBRUARY 2018

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Effect of the H-1B Quota on the Employment and Selection of Foreign-Born Labor*

The H-1B program allows skilled foreign-born individuals to work in the United States. The annual quota on new H-1B issuances fell from 195,000 to 65,000 for employees of most firms in fiscal year 2004. This cap did not apply to new employees of colleges, universities, and non-profit research institutions. Existing H-1B holders seeking to renew their visa were also exempt from the quota. Using a triple difference approach, this paper demonstrates that cap restrictions significantly reduced the employment of new H-1B workers in for-profit firms relative to what would have occurred in an unconstrained environment. Employment of similar natives in for-profit firms did not change, consistent with a low degree of substitutability between H-1B and native workers. The restriction also redistributed H-1Bs toward computer-related occupations, Indian-born workers, and firms using the H-1B program intensively.

JEL Classification: J61, F22 , O33, R10
Keywords: skilled workers, H-1B, natural experiment

Corresponding author:
Anna Maria Mayda
Georgetown University
Department of Economics
37th and O Streets, NW
Washington DC, 20057
USA
E-mail: amm223@georgetown.edu

* We are grateful for funding from the National Science Foundation, Award 1535723, "The Effect of H-1B Workers on the Innovation and Productivity in US Firms."

1 Introduction

The H-1B program provides a pathway for foreign-born skilled professionals to work in the United States. Proponents argue that the H-1B program increases the skilled labor force, boosts innovation and productivity, and generates long-run economic gains. Opponents instead argue that H-1B workers displace Americans and reduce wages, especially in the Information Technology (IT) sector. Economists and policy-makers continue to debate these issues and several other important aspects of the H-1B program such as whether it actually selects workers of high ability. The program's design might unintentionally favor some firms, workers, or occupations. This paper informs these issues by examining how a large reduction in the number of H-1B visas available beginning in fiscal year 2004 affected the employment and characteristics of new H-1B workers in for-profit firms, whether any effect was observed in the employment of similar natives in for-profit firms, and how this altered the types of firms hiring H-1B labor.

In fiscal year (FY) 2004 the annual quota (or cap) on H-1B visas available for new foreign-born employees of most firms fell drastically from 195,000 to 65,000 per year.[1] While the reduction was substantial, it was not applied uniformly. In particular, two groups of H-1B workers were unaffected by this policy change. First, employees of colleges, universities, and non-profit research institutions have been exempt from the cap since 2001 (hereinafter, "non-profit" or "cap-exempt" firms). Second, the quota only applies to new H-1B applicants. Established H-1B workers seeking to extend their employment or move from one for-profit firm to another have never been subject to any cap. Thus, the quota reduction in 2004 generated a quasi-experiment by creating a sudden discontinuity in the maximum supply of H-1B workers in the "treated" group of newly-hired employees of for-profit firms, relative to those in the "control" groups of newly-hired non-profit employees and experienced H-1B workers.

We empirically study the effects of the H-1B quota reduction using a triple difference strategy. We compare the hiring of new and established H-1B workers across for-profit and non-profit sectors before and after 2004. The extent to which the cap affects employment depends on the difference between H-1B demand and the cap. For this reason fiscal years 2008 and 2009 are particularly notable as aggregate demand for H-1B visas spiked and more than 150,000 applications were received within the first week of the filing period in each year. The US government responded by using a lottery to distribute all cap-subject permits for these years. We examine whether the heightened intensity of rationing during the lottery

[1] The current quota includes these 65,000 general visas, plus 20,000 permits for new H-1B workers who have obtained an advanced degree from a US university.

years had stronger impacts on hiring and other outcomes.

Our analysis exploits a unique dataset on approved H-1B applications acquired through a Freedom of Information Act (FOIA) request. Individual-level information allows us to ascertain the education, experience, and occupation of each approved H-1B worker. We utilize variation in H-1B hiring within nationally-defined skill cells from FYs 2002 through 2009. These cells group workers according to education, experience, and occupation. Following in the tradition of Borjas (2003) and Ottaviano and Peri (2012), we consider each cell as a somewhat separate labor market since workers are closer competitors within rather than across cells. Information on the hiring firm and applicant allow us to parse H-1B approvals and skill cells into four groups defined by whether employers were for-profit or non-profit and whether the application was for a new or established H-1B worker.

Our regressions address two broad issues. First, we quantify the decline in new H-1B hiring in for-profit firms relative to what would have occurred in the absence of hiring restrictions. We find that cap-subject skill cells experienced an approximate 20-50% decline in new H-1B employment relative to what it would have been if firm demand – rather than a legislative limit – determined hiring outcomes. New H-1B employment fell by an additional 3-8% relative to its demand-driven level during fiscal years 2008 and 2009. Thus, the binding cap combined with the surge in applications led to a tremendous rationing in skilled, foreign-born, labor-market entry through the H-1B program.

Key to the identification, our triple difference design effectively removes differential pre-trends between for-profit and non-profit sectors and between new and existing hires for each skill-specific labor market. Furthermore, our results remain stable across a variety of specifications, set of control variables, and robustness checks. We perform an additional test demonstrating that other contemporaneous shocks were unlikely to have generated our central results and that the employment effects for new H-1B workers were not offset by employment effects for natives. Namely, we analyze the market for native-born workers using the same triple difference framework and find no evidence for a change in native-born employment, in occupations similar to those of H-1B workers, at for-profit firms after 2004. These results have an important implication. Since the reduced cap caused H-1B hiring to fall without generating an offsetting rise in the stock of native employment, the findings show no evidence for short-run native and H-1B labor substitutability.

Second, we assess whether the quota affected the selection of new H-1B employees and the types of firms participating in the H-1B program. We find evidence of a decline in H-1B recipients at the tails of the wage-offer distribution. This is particularly concerning at the high end of the wage distribution as it indicates that the cap reduced the inflow of foreign workers most likely to be highly productive and innovative. We also show that the reduced

cap shifted the composition of workers toward computer-related occupations and Indian-born workers. Moreover, the policy change redistributed H-1B labor toward firms that employ 50 or more H-1B workers each year and away from employers that use the program less intensively, thus increasing the concentration of H-1B workers in fewer firms. These results could be an indication that past experience and existing networks became more relevant in securing a share of the constrained supply of new H-1B workers.

The results help to inform both economic and policy debates. In terms of the economics literature, this paper is most directly related to the growing research on the economic impacts of the H-1B program. Several papers have shown that H-1B workers bring positive benefits to innovation, productivity, and labor markets (e.g. Kerr and Lincoln 2010; Peri, Shih and Sparber 2015; Ghosh, Mayda, and Ortega 2016). In contrast, Doran, Gelber, and Isen (2014) provide recent evidence from H-1B pilot lotteries showing that H-1B workers do not increase firm productivity or innovation.

By estimating how the cap has affected the inflow of H-1B workers, native employment, and the types of workers and firms participating in the program, this paper can improve our understanding of the economic impact of this program. In this way we complement past work utilizing the change in the H-1B cap to examine unintended consequences on the international undergraduate student population. For example, Kato and Sparber (2013) show that applicant quality declined after the quota was reduced. Shih (2016) finds that the fall in the cap led to a decline in undergraduate enrollment that was especially pronounced among students from countries with higher expected returns to working in the United States. Amuedo-Dorantes and Furtado (2016) show that the cap altered the career choices of foreign students as they substituted away from private sector firms and toward research institutes.[2]

Beyond direct questions about immigration, the H-1B quota also provides a rare opportunity to understand how labor markets respond to quantity constraints. Much work has been done on limits in the form of a minimum wage floor (e.g. Card and Krueger 1994, 2015), but since the US government does not normally impose limits on the types of workers firms may hire, there is little empirical evidence on the impact of quantity restrictions, and in particular, how they affect wages and the selection of workers. Our work improves understanding about the effects of quantity restrictions in the labor market.

In terms of policy debates, H-1B reform is part of the US government's current agenda. Concerns over skilled immigrants' potential to negatively affect American workers have been heightened by several cases receiving significant media attention in which companies replaced their domestic IT staff with foreign H-1B workers.[3] Descriptive statistics indicate

[2]Note that we run a robustness test which shows that the results in Amuedo-Dorantes and Furtado (2016) are not an issue for our identification and findings.

[3]High profile cases include Southern California Edison, Disney, and University of California – San Fran-

an increasingly large share of H-1B visas are awarded to a very small subset of IT firms.[4] These IT firms have been accused of using the H-1B program to outsource American jobs by hiring Indian computer workers at low wages. Our analysis shows that reductions in the cap exacerbate such issues. Moreover, the H-1B quota has changed quite dramatically over its history, and currently-competing proposals have called for both increases and decreases in the quota.[5] Given the regular use of the quota as a policy lever, it is important to evaluate its efficacy and consequences. Fortunately – and unlike other policy proposals[6] – prior experience provides the empirical means to study implications of changing the cap.

The rest of the paper is organized as follows. Section 2 describes the H-1B program in greater detail, including the drop in its quota that provides the natural experiment of interest. Section 3 provides a simple theory deriving demand for H-1B workers and illustrating how a binding supply constraint would affect the employment of specific groups of workers. Section 4 describes the data. Section 5 discusses our empirical identification strategy, while Section 6 shows the empirical results. Section 7 concludes.

2 Skilled Workers and The H-1B Program

The H-1B program has been the primary channel for foreign-born college-educated professionals to enter the US workforce since its inception in 1990. H-1Bs are distributed to individuals, not firms, but the employment-based nature of the program creates a strong link between an employee and his or her employer. Moreover, its dual intent allows firms to sponsor its H-1B workers for an employment-based green card. Many who have entered on an H-1B visa have remained permanently in the country. H-1B recipients are concentrated among science, technology, engineering, and mathematics (STEM) occupations: roughly 50% of the growth in the college-educated STEM workforce since 1990 is attributable to H-1B workers (Peri, Shih, and Sparber 2015).

Several steps are required for an individual to acquire an H-1B. First, she must obtain an employment offer from a US firm. This firm must file a Labor Condition Application (LCA)

cisco. See Anderson (2015), Preston (2016), and Fernandez Campbell (2016) for press coverage.

[4] In a testimony to Congress on February 25, 2016, Ron Hira cites that 30% of new visas subject to the quota were distributed to only 10 large IT services firms. Transcripts of the testimony are available at`https://www.judiciary.senate.gov/`.

[5] Proposed quota expansions include the SKILLS Visa Act – which would raise the cap to 155,000 – and S.744 – which would have allowed the cap to fluctuate according to market conditions within the 115,000-180,000 band. Proposed contractions include H.R. 3163 – which would have created additional restrictions for large employers – and the S.2365 Protecting American Jobs Act – which sought to reduce the quota by 15,000.

[6] Aside from the quota, other policy proposals have included enforcing a minimum wage for H-1B workers, increasing application fees, auctioning off visas, or intensifying monitoring for compliance.

with the Department of Labor outlining the nature of the job and attesting compliance with H-1B regulations. There is no limit to the number of LCAs a firm can file. After receiving LCA approval from the Department of Labor, the employer and worker combine efforts to complete the H-1B application. This includes evidence of a formal job offer, the firm's approved LCA, and an I-129 form that documents further information about the individual and the job. If awarded an H-1B, the foreign applicant may then move to the United States and begin work on the start date indicated in the H-1B application for a duration of 3 years. H-1B status can be renewed for a total duration of 6 years. Further stay in the United States beyond 6 years requires either a successful petition for permanent residency or a transfer to a different visa.

US Citizenship and Immigration Services (USCIS) oversees the H-1B adjudication process. April 1 marks the opening of the filing period for H-1B applications to work on the first day of the fiscal year beginning on October 1. For example, April 1, 2007 marked the opening of the filing period for H-1B petitions under the fiscal year 2008 cap. October 1st, 2007 was the earliest date on which an H-1B recipient could begin work.

New issuances have always been subject to an annual cap. This quota – initially set at 65,000 per fiscal year – has changed significantly over time. USCIS has traditionally used a first-come first-served approach to allocate permits until available permits are exhausted. The end of the application period is demarcated by the "final receipt date" when USCIS has received enough applications to fill the remaining H-1Bs available. Applications submitted after the final receipt date are not processed.

This paper examines firms' H-1B hiring behavior during three eras. The first covers FYs 2002 and 2003. The American Competitiveness in the 21st Century Act (AC21) – signed and put into effect on October 2000 – temporarily raised the limit on new H-1B issuances to employees of for-profit firms to 195,000 for FYs 2001 through 2003. This cap was never reached. AC21 also permanently exempted universities and non-profit research institutions from the cap. Thus, the stock of people with H-1B experience grew during this period as issuances for both the for-profit and non-profit sectors were effectively unconstrained.[7]

The temporary nature of this high cap implies that it would have taken an act of Congress for it to continue. By October 2003, however, Congress declined to issue a renewal and the policy expired, reducing the quota back to the original limit of 65,000. Thus, the second era of our analysis begins in FY 2004 when the temporarily high quota on new issuances at for-profit firms expired. Our primary interest is in this sudden event that generated a sector-specific shock that persists. Though the H-1B Visa Reform Act in December 2004

[7] We begin the analysis in fiscal year 2002, in part, because the timing of AC21's passage complicates how we assign H-1B applications to particular years.

generated a small expansion by allocating an additional 20,000 permits to workers who earned a master's degree or Ph.D. from a US institution, the general H-1B quota has been reached before the end of each fiscal year since 2004. This year clearly marks the beginning of a period characterized by for-profit firms' limited ability to hire new H-1B workers. However, such firms could substitute toward established H-1B workers who were in ample supply due to the non-binding caps of the preceding era.

The third era covers FYs 2008 and 2009. These years require special examination for two reasons: First, demand for H-1B workers grew as the economy expanded prior to the dawn of the US Financial Crisis and Great Recession. USCIS faced overwhelming demand such that the number of cap-bound H-1B applications exceeded the number of available permits within just the first week of the application period. USCIS responded by allocating all cap-subject H-1Bs by a computerized random lottery. Second, the pool of established H-1Bs would have necessarily declined by this point, many long-established H-1B workers would have exhausted their six-year limit, and inflows of new H-1B workers in preceding years were more restricted. Thus, compared to earlier years, for-profit firms wishing to hire H-1B workers would have been less able to substitute away from new workers and toward established ones.

Our paper will empirically assess the consequences of the H-1B cap reduction. We will use a national-level approach that groups workers into different labor markets defined by skill. One difference from previous papers using this approach is that we add dimensions separating groups according to employer type (for-profit versus non-profit) and H-1B issuance type (new versus established H-1B workers). We assume that demand for comparably skilled workers evolves similarly across sectors and is driven by technological change. We also allow for overall productivity and demand to differ across sectors. We track the evolution in demand for new H-1B workers and evaluate the effect of introducing a binding cap in 2004 in for-profit firms by comparing them to similar non-profit hires. Before turning to the empirical analysis, we provide a brief theoretical sketch illustrating the determinants of demand for skilled foreign-born workers across experience groups and sectors. The framework helps justify the empirical specification.

3 Theoretical Framework

To motivate our empirical design and its assumptions, we present a simple model of production and relative labor demand. We consider two sectors (j): one made of firms producing goods and services for-profit (PR), and the other made of non-profit firms (NP). Each of these sectors produces one good or service with price p_j. Prices can change over time in

response to demand. Non-profit firms earn a profit margin equal to zero, but this need not be true at for-profit firms. In the long-run, firms compete for workers of each skill type and experience level. Workers earn a wage equal to their marginal productivity. As workers are mobile, the wage for each specific type of worker is equalized across sectors.

Consider the following production function for output (Y) in sectors j and year t:[8]

$$Y_{jt} = A_{jt}\left(K_{jt}\right)^{\eta}\left(\sum_{k}\delta_{jkt}\left[\left(L_{jkt}^{New}\right)^{\gamma}+\theta_{jkt}^{Exp}\left(L_{jkt}^{Exp}\right)^{\gamma}\right]^{\frac{\rho}{\gamma}}\right)^{\frac{1-\eta}{\rho}} \tag{1}$$

In Equation (1) A_{jt} is total factor productivity (TFP) and K_{jt} is a combination of factors (such as physical capital, unskilled labor, and US-born skilled workers) used in production. The summation term in parenthesis, henceforth represented by Γ_{jkt}, is a nested CES combining foreign-born workers (L) of different skills k defined by occupation, education, and potential work experience. These workers are distinguished by their work experience in the United States. Let L_{jkt}^{New} denote the labor input of foreign-born workers with no previous US working experience. In our context, these workers are on their first H-1B term and hence have 0 to 3 years of working experience in the US. The coefficients δ_{jkt} capture the skill-specific relative productivity for workers in sector j, while θ_{jkt}^{Exp} captures the productivity of experienced workers relative to new ones within a skill group. The parameter ρ determines the elasticity of substitution between workers of different skill groups, $\sigma = 1/(1-\rho)$. The parameter γ determines the elasticity of substitution between new and experienced workers, $1/(1-\gamma)$.

The key to our empirical strategy is that employment of each skill group is determined by relative demand when the H-1B quota is large and non-binding (as it was prior to 2004). Utility maximizing workers will arbitrage away wage differences by moving between the for-profit and non-profit sectors. Employment of new workers of each skill type satisfies the equalization of wages between for-profit and non-profit firms, $w_{kt}^{New,NP} = w_{kt}^{New,PR} = \omega_{kt}^{New}$, in each period [9]. Equalization of wages for experienced H-1B workers also occurs since they are not subject to any quota: $w_{kt}^{Exp,NP} = w_{kt}^{Exp,PR} = \omega_{kt}^{Exp}$. Wages equal to marginal productivity implies:

$$\omega_{kt}^{New} = (1-\eta)p_t^j A_t^j\left(K_t^j\right)^{\eta}\Gamma_{jkt}^{\frac{1-\eta-\rho}{\rho}}\delta_{jkt}\left[\left(L_{jkt}^{New}\right)^{\gamma}+\theta_{jkt}^{Exp}\left(L_{jkt}^{Exp}\right)^{\gamma}\right]^{\frac{\rho-\gamma}{\gamma}}\left(L_{jkt}^{New}\right)^{\gamma-1} \tag{2}$$

[8] Note this can be considered as an extension of the production function adopted in Ottaviano and Peri (2012) or Card and Lemieux (2001)

[9] Wages for a type of workers may be different between profit and non-profit sector, due to a percentage premium or differential. The main results from this model follow even in the presence of such differential, as long as it does not change over time

$$\omega_{kt}^{Exp} = (1-\eta) p_t^j A_t^j \left(K_t^j\right)^{\eta} \Gamma_{jkt}^{\frac{1-\eta-\rho}{\rho}} \delta_{jkt} \left[\left(L_{jkt}^{New}\right)^{\gamma} + \theta_{jkt}^{Exp}\left(L_{jkt}^{Exp}\right)^{\gamma}\right]^{\frac{\rho-\gamma}{\gamma}} \theta_{jkt}^{Exp}\left(L_{jkt}^{Exp}\right)^{\gamma-1} \quad (3)$$

To obtain the basic estimating equation when quotas do not bind, we first take the logarithm of both sides of (2) and (3) in each sector. Then we calculate the difference in log-wages and rearrange to obtain:

$$\ln\left(L_{jkt}^{New}\right) - \ln\left(L_{jkt}^{Exp}\right) = \frac{1}{\gamma-1}\left(\ln\theta_{jkt}^{Exp} + \ln(\omega_{kt}^{New}) - \ln(\omega_{kt}^{Exp})\right) \quad (4)$$

Finally, we difference Equation (4) across sectors $j = \{FP, NP\}$. Since new-workers' and experienced-workers' wages are common across the for-profit and non-profit sectors, this difference simplifies to:

$$\begin{aligned} & \left[\ln\left(L_{PR,kt}^{New}\right) - \ln\left(L_{PR,kt}^{Exp}\right)\right] - \left[\ln\left(L_{NP,kt}^{New}\right) - \ln\left(L_{NP,kt}^{Exp}\right)\right] \quad (5) \\ = & \frac{1}{\gamma-1}\left(\ln\left(\theta_{PR,kt}^{Exp}\right) - \ln\left(\theta_{NP,kt}^{Exp}\right)\right) \end{aligned}$$

Equation (5) holds if firms are on their demand curve and workers are free to move across sectors. We allow the experience-specific productivity for each skill group, $\ln\left(\theta_{jkt}^{Exp}\right)$, to have a component that is unique to the NP or PR sector. This is separable from a skill-specific component common between sectors that evolves over time and varies across skills due to differences in relative demand, relative prices, or TFP. The double differencing removes skill- and sector-specific productivity (and prices). Only the relative productivity of new and experienced workers remains. If we assume that experience-specific productivity can be decomposed into $\ln\left(\theta_{jkt}^{Exp}\right) = \ln\left(\theta_{kt}^{Exp}\right) + \ln\left(\theta_{j}^{Exp}\right)$, then the double difference in relative demands expressed in equation (5) will be represented by a constant factor common to all skills, $\alpha = \frac{1}{\gamma-1}\left(\ln\theta_{PR}^{Exp} - \ln\theta_{NP}^{Exp}\right)$.

Combining the above assumptions, allowing for classical measurement error in the employment variables, and rearranging equation (5) provides us with Equation (6).

$$\left[\ln\left(L_{PR,kt}^{New}\right) - \ln\left(L_{PR,kt}^{Exp}\right)\right] = \alpha + \left[\ln\left(L_{NP,kt}^{New}\right) - \ln\left(L_{NP,kt}^{Exp}\right)\right] + \varepsilon_{kt} \quad (6)$$

The above relation holds in the absence of quotas, when relative labor demand and free mobility across sectors combine to determine relative employment. Equation (6) establishes that the demand for new relative to experienced workers in the for-profit sector follows the same variation over time as relative demand in the non-profit sector up to a constant (α) and a zero-mean random error (ε_{kt}). One way of interpreting equation (6) is that the relative

employment of new and experienced H-1B workers in the non-profit sector correlates with that of the for-profit sector when supply responds to demand.

Now consider when the fiscal year 2004 binding quota for new H-1B workers, $\overline{L^{New}}$, was introduced in the for-profit sector only. This rationing causes the difference $\ln\left(L^{New}_{PR,kt}\right) - \ln\left(L^{Exp}_{PR,kt}\right)$ to depart from what is predicted in Equation (6) and a negative gap will arise. The absolute value of this gap is large when the quota is small and/or when the underlying relative demand for new relative to experienced workers is large. This latter condition is captured by relative non-profit sector demand, $\ln\left(L^{New}_{NP,kt}\right) - \ln\left(L^{Exp}_{NP,kt}\right)$.

The quota for new H-1B workers in the for-profit sector has been binding since 2004 when it was significantly reduced by 66% (from 195,000 to 65,000). In FYs 2008 and 2009 an exceptionally high level of demand led the number of H-1B applications to exceed the cap in the first week of the filing period. Hence, we can augment expression (6) with the possibility of a rationing gap opening after 2004 and becoming larger in 2008 and 2009 as follows:

$$\begin{aligned}\left[\ln\left(L^{New}_{PR,kt}\right) - \ln\left(L^{Exp}_{PR,kt}\right)\right] &= \alpha + \left[\ln\left(L^{New}_{NP,kt}\right) - \ln\left(L^{Exp}_{NP,kt}\right)\right] + \\ &+\beta_1 * I(Year \geq 2004) + \beta_2 * I(Year = 2008, 2009) + \varepsilon_{kt}\end{aligned} \tag{7}$$

The indicator $I(Year \geq 2004)$ equals one starting in 2004 so that the coefficient β_1 captures the departure of new H-1B employment among for-profit relative to non-profit firms due to the binding cap. The size of this coefficient represents the estimated percentage deviation of capped H-1B workers relative to unconstrained demand. Similarly, $I(Year = 2008, 2009)$ is an indicator equal to one in FYs 2008 and 2009 so that $\beta_1 + \beta_2$ captures the further distance between the cap and demand when demand increased. Equation (7) will be the basis of our empirical analysis.

4 Data

Our data comprises individual-level information from I-129 forms of processed H-1B applications between fiscal years 2002 and 2009, obtained through a Freedom of Information Act (FOIA) request. This is an ideal period to analyze the effects of a binding cap. First, it corresponds to an expansionary cycle ending just prior to the full onset of the Great Recession. Second, the fall in the H-1B cap occurs towards the middle of this time period. Third, hiring in the for-profit sector in the years prior to 2004 was essentially unconstrained as the quota did not bind. Only after the cap fell did the quota bind, growing progressively more

imposing as the economy expanded before the Great Recession.

This dataset contains information on the employer, employee demographics, wages, occupation, and other characteristics that allow us to determine whether the permit would count toward the H-1B cap. We impose several sample selection criteria. We remove applications that were denied, allowing us to identify H-1B recipients rather than simply the count of processed applications. We retain H-1B recipients age 21-65 (inclusive), born outside of US territories, who have obtained at least a bachelor's degree, and who are working in identifiable occupations. To remove outliers and reduce measurement error, we trim workers in the top and bottom 0.5% of the H-1B wage distribution.

We separate individuals into whether they were employed at for-profit or non-profit firms. The latter group consists of colleges, universities, and non-profit research institutions that were always exempt from H-1B quotas during the period of analysis. We also separate individuals into applicants for new H-1B employment versus established H-1B workers. This latter group consists of H-1B employees who receive an extension of their visa with their current firm as well as those who changed employers.[10]

We then aggregate workers into skill cells representing distinct labor markets. We take two approaches. The first identifies a skill cell according to a worker's occupation, education, and experience. The second uses occupation and education only. While potential experience is often used to differentiate skill groups (e.g. Card and Lemieux 2001 or Borjas 2003), education and occupation more strongly characterize the specific skills used in jobs. Cells based only on those two features are likely to be strongly differentiated and represent well-defined labor markets.

We use 16 broad occupational groupings.[11] The largest, Computer-Related Occupations, accounts for 40% of new H-1B issuances. Not surprisingly, occupational composition varies substantially across the for-profit and non-profit sectors. Computer occupations represent 50% of new issuances among for-profit firms but only 2.5% of new non-profit employment (though they rank among the top five occupations among new hires in each sector). Occupations in education account for over half of new non-profit H-1B employment but just 1% of new for-profit employment. Other important for-profit occupations for new hires include Managers and Administrative Officials[12] (17%), Engineers (14%), Medical and Health Professionals (4.6%), and Social Scientists (2.9%). Important non-profit occupations include

[10] As a technical point, an established H-1B worker at non-profit firm who then seeks for-profit employment would count toward the cap.

[11] Aggregation is based upon the USCIS Department of Homeland Security Form M-746, I-129 Dictionary of Occupational Titles (DOT) Codes.

[12] This is a broadly-defined grouping that includes not just managers but accountants, auditors, and related occupations as well (Dictionary of Occupational Title codes 160-189).

Medical and Health Professionals (18%), Life Scientists (12%), and Mathematicians and Physical Scientists (4.2%).

Individuals are grouped into four education levels: bachelor's (45% of new H-1B recipients), master's (35%), professional (6%), and doctorate (14%). Again, representation varies across sector. Doctorates comprise 48% of new non-profit employment, whereas Bachelor's and Master's degree holders comprise 53% and 40% of new for-profit H-1Bs, respectively.

Consistent with the past literature, we deduce total potential experience through an intersection of age and educational attainment. However, our dataset presents unique challenges. First, as a temporary work permit program for foreign-born labor, one can correctly infer that most H-1B recipients have limited work experience within the US labor market. About half of new H-1Bs are awarded to people under age 30, and 90% are awarded to people under age 40. Hence there is not a lot of experience variation within the program. Second, whereas the national-approach usually defines cells according to educational attainment ranging from high school drop outs to college graduates, our data is from a much narrower range of educational attainment that is limited to college graduates. This may reduce the differentiation across skill groups.

Nonetheless, we proceed as follows: We first define years of experience as the difference between a person's age and their expected age upon completing their highest degree. Expected ages at graduation are 23 for bachelor's, 25 for master's, 26 for professional, and 28 for doctorate. Next we place workers into (approximately) five-year experience groups: those with 5 or fewer years of expected experience (including observations with negative values), 6-10 years, 11-15 years, and so on. The eighth and most experienced group has more than 35 years of expected experience. Not surprisingly, less-experienced workers are highly represented. Workers in experience groups one (57%), two (24%), and three (10%) account for more than 90% of new H-1B issuances. These values do not vary much across for-profit and non-profit sectors.

The year (8), occupation (16), education (4), experience (8), and firm-type (2) groupings imply 8,192 potential jkt (sector-skill-year) observations. We are interested in how a policy change proportionally affected employment of H-1B recipients in each skill group, so we measure employment quantities in natural logs. This leaves the dataset with 4,556 observable first-differences between new and established employment, $\ln\left(L_{jkt}^{New}\right) - \ln\left(L_{jkt}^{Exp}\right)$, and 3,512 observations in which first differences are defined for both $j = PR$ and $j = NP$ (hence, 1,756 of the potential 4,096 skill-k period-t cells have defined values for the double-difference between the for-profit and non-profit sectors). The average value of log H-1B issuances to new employees is 2.34 with a standard deviation of 2.02. The distribution is heavily skewed to the right. Log wages paid to these workers have an average value of 11, a standard

deviation of 0.41, and are normally distributed. Since many experience groupings result in few observations, we also perform regressions that eliminate the eight experience distinctions. This results in 1,024 potential jkt observations, 938 observable first-differences, and 894 first-differences defined for both sectors. Log new H-1B issuances have an average of 4.09 and standard deviation of 2.32. Log wages average 10.93 with a standard deviation of 0.33.

Given the variation in importance of occupations and education levels across sectors, there could be some concern that the unique labor markets identified in our skill cells exclusively exist in only one sector or the other so that comparisons of skill cells between for-profit and non-profit firms cannot be done. Figure 1 helps to alleviate this concern. When new H-1B issuances for the entire period are aggregated to the cell level, there is a high degree of positive correlation between the number of for-profit and non-profit issuances. Cells that contain large numbers of new for-profit workers tend to contain large numbers of non-profit workers as well. Robustness checks in Section 6 will further explore whether large differences in cell size across sectors drives results.

The empirical design compares changes in new relative to experienced H-1B employment between the for-profit and non-profit sector. Figure 2 provides a sense of trends in H-1B issuances from FYs 2002-2009 in the for-profit and non-profit sector. The figure shows the average cell differences in log-issuances between new and established H-1Bs (described further in the next section) for the non-profit (solid line) and for-profit (dashed line) sectors. New issuances were fairly constant relative to established ones prior to the 2004 cap restriction. Additionally, there is no evidence of differential pre-trends across sectors. The difference declined – a regularity that holds for both the for-profit and non-profit sectors. This fall, however, was more severe in the for-profit sector than it was for cap-exempt employers. Our empirical methodology described in the next section will take a more systematic approach to assessing whether differences in these two trends appear after the cap was reduced for fiscal year 2004 and beyond, and whether differences further increased during lottery years.

5 Empirical Methodology: Estimating the Effect of a Binding Cap

Our empirical strategy takes advantage of variation in H-1B hiring across time, for-profit versus non-profit sector, and H-1B experience. AC21 segmented the H-1B market into for-profit and non-profit employee groups in 2001, but the economic effects of that policy change would not be felt until 2004 when the high quota on new for-profit hiring expired. Standard difference-in-difference estimation should be well-suited for analyzing the for-profit versus

non-profit employment effects of this exogenous shock. However, our preferred identification strategy employs triple-difference estimation that exploits not only the difference in employment across sectors, but across new versus established workers as well. Though only new employment experiences a clear shock, potential substitutability suggests that the shock could affect the relative and total employment of these two groups. By exploiting this additional source of variation we can also clean for a host of potential demand shifters – our model includes a rigorous set of fixed effects to account for any trends differentially affecting the for-profit and non-profit sectors.

We formalize our triple difference strategy around the estimating equation in (8), which represents a simple transformation of the model derived in (7).

$$\begin{aligned} \ln\left(H1B_{jkt}^{New}\right) - \ln\left(H1B_{jkt}^{Exp}\right) \quad = \quad & \\ & \alpha + \beta_1 \cdot t_{\geq 2004} \cdot d_{FP} + \beta_2 \cdot t_{\geq 2008} \cdot d_{FP} \\ & + d_{FP} + \sigma_k + \sigma_t + b\varepsilon_{kt} \end{aligned} \tag{8}$$

The dependent variable represents the difference in the natural log of H-1Bs awarded to new workers relative to those awarded to experienced workers within a particular sector (j), skill group (k), and year (t). The indicator $t_{\geq 2004}$ takes a value of one in fiscal year 2004 and after, while $t_{\geq 2008}$ is equal to one for years 2008 and 2009. The dummy d_{FP} is equal to one for observations in the treatment for-profit sector; it equals zero for observations in the control non-profit sector. The model also includes skill-group fixed effects (σ_k) and year dummies (σ_t) to absorb variation in productivity across skill groups and time. The coefficient β_1 represents the proportional reduction of new H-1Bs issued to workers in the for-profit sector due to the decrease in the quota in FY 2004. β_2 represents the additional rationing when firms faced particularly large H-1B demand in FYs 2008 and 2009.

Causal inference in the triple-difference framework requires several assumptions. First, the policy must not have been a response to differentially changing conditions in the labor market for new foreign workers at for-profit firms. In this sense, studying the fall in the H-1B cap in 2004 is arguably advantageous as it was not a new policy, but rather a reversion to a pre-existing policy. The high H-1B quotas in the early 2000s were contingent on Congressional renewal. Congress had once previously renewed the high quota, and even went on to raise the cap further in FY 2001 (Shih 2016). This ex-ante uncertainty of Congressional action circa FY 2004 likely made the fall in the quota unpredictable.

Second, the triple difference should not exhibit a trend pre-2004. Though the period with a non-binding cap is short, Figure 2 suggests no visual pre-trend.[13] Additionally, our

[13] Appendix Figure A1 provides supporting evidence by using Current Population Survey (CPS) data to

analysis will employ a variety of controls to assess robustness. The baseline empirical specification controls for skill cell and year fixed effects. This should account for any productivity difference specific to particular skill cells and for yearly changes between the for-profit and the non-profit sector. Other specifications using year dummies and skill-group by sector fixed effects deliver robust results.

Third, the estimated impacts must not be attributed to any other confounding shocks that differentially affected hiring of new relative to experienced H-1B workers at for-profit versus non-profit firms. To this end, we perform a test by examining native-born workers who are not subject to the quota. Assuming a similar evolution for the relative demand of native and foreign workers in the same skill cell, evidence of a systematic employment difference arising after 2004 for natives might signal that other factors rather than H-1B policy could be at work. Our results do not indicate any differential hiring response among new native workers relative to experienced ones at for-profit firms after 2004. This suggests that our estimates are credibly driven by H-1B policy and not by other confounding economic shocks.

6 Empirical Results

6.1 H-1B and Native Employment

We estimate specification (8) to assess the effect (β_1) of reducing the cap on new H-1B hiring relative to what would have occurred under a non-binding constraint. Results are reported in Table 1. Columns 1-3 (left panel) use skill groups defined by occupation, education, and potential experience. Columns 4-6 (right panel) use skill groups defined by occupation and education only. The columns within each panel represent specifications in which we progressively add controls. The first includes skill-cell fixed effects and indicators for the collective groups of cap-years and lottery-years. The second replaces the cap-year and the lottery-year dummies with fixed effects for each individual year of the dataset. The third column balances the panel so that both sectors $j = PR, NP$ are observed in a kt cell. It also replaces sector and skill indicators with sector-by-skill ($j * k$) fixed effects.[14] Standard errors are clustered by skill-cell in the first two columns and by skill*sector dyad in the third column.

illustrate that for-profit firms and universities exhibited similar college-educated employment trends among native-born workers from the late 1990s through the implementation of the restricted H-1B cap in fiscal year 2004

[14] This third specification delivers coefficient estimates equal to those of a simple transformation of (7) that would regress $\left[\ln\left(L^{New}_{PR,kt}\right) - \ln\left(L^{Exp}_{PR,kt}\right)\right] - \left[\ln\left(L^{New}_{NP,kt}\right) - \ln\left(L^{Exp}_{NP,kt}\right)\right]$ on a constant, $t_{\geq 2004}$ and $t_{\geq 2008}$ indicators, and skill cell (k) fixed effects.

Our interest is in the coefficient estimates of β_1 from the interaction Cap Year*Treated (i.e. $t_{\geq 2004} \cdot d_{FP}$) and β_2 from the interaction Lottery Year*Treated (i.e. $t_{\geq 2008} \cdot d_{FP}$). The estimates of β_1 indicate that H-1B restrictions introduced in 2004 reduced new H-1B hiring at for-profit firms between 18 and 40 log points (between 20 and 50%). All estimates are statistically significant at the 5% level and most are significant at the 1% level. As restrictions became more intense in FYs 2008 and 2009 because firm demand for new H-1B workers was so large that all cap-bound permits were allocated by lottery, H-1B hiring declined by an additional 3 to 8 log points, though estimates are imprecise.

Triple-difference estimation is our preferred identification strategy since it controls for a host of potential demand shifters and exploits variation across employment sectors and H-1B experience. Nonetheless, double-difference estimation is an informative tool that can decompose triple-difference results into its contributing components. Table 2 performs this exercise for the models of Columns 3 and 6 of Table 1.

By decomposing this effect into its principle components we learn that new H-1B employment fell and established H-1Bs grew in both the for-profit and non-profit sectors. The growth in established employment was the same in both sectors, resulting in positive but insignificant point estimates (0.050 for occupation*education*experience cells; 0.016 for occupation*education cells). However, the decline in new employment was much larger at for-profit firms, resulting in double-difference point estimates ranging from a 0.19 to 0.28 log-point employment loss. Thus, it appears that both for-profit and non-profit firms increasingly substituted toward established H-1B employment beginning in 2004, but the disproportionate decline in new H-1B employment at for-profit firms is responsible for driving the triple-difference relative effect.

Lottery-years coefficients reveal that for-profit firms experienced an additional decline in both new and established H-1B hiring relative to non-profit firms. The similarity in magnitude for new and established H-1Bs implies that the triple-difference estimate is insignificant. One interpretation is that new H-1B workers simply did not see a relative decline in employment at for-profit firms in lottery years beyond what they experienced as a result of the H-1B cap reduction in 2004. Another is that for-profit firms responded to the 2004 policy change by substituting away from new toward established H-1Bs when possible. By 2008, the pool of established H-1Bs from which a firm could hire had also diminished so this substitution was no longer possible. Thus, total H-1B employment had to decline.

Consider Table 3 for support of this view. This table examines the effect of H-1B cap reductions by performing double-difference estimation of log-total H-1B issuances between for-profit and non-profit firms. Columns 1 and 4 use all available observations and additive fixed effects for cell, sector, and time period. Columns 2 and 4 balance the panel for the

double-difference estimation and includes year and cell*sector multiplicative fixed effects. Columns 3 and 4 further reduce the sample by including only the observations that would have been available in the triple difference analysis. The table provides evidence for a small, generally negative, and almost always insignificant decrease in total non-profit employment immediately after the cap restriction. During lottery years, however, total H-1B employment at for-profit firms dropped an additional 0.23 to 0.30 log-points. Thus, it appears that for-profit firms did initially respond to restrictions by substituting toward established H-1B workers. Total for-profit H-1B employment declined only later when both the pool of established H-1B workers and potential new H-1B employment were restricted. As is common with a policy change, one never knows what the effect will be even on the most directly targeted variable (e.g., foreign skilled employment) since firms always try to find ways to get around constraints. In this case, it appears that the quota was binding for the new visas in the first few years, but not for overall skilled foreign employment (i.e. the sum of new and experienced issuances) until later years.

The results above point to a need to explore whether conclusions are robust to potential temporal issues. We address this by restricting the sample to just fiscal years 2002 through 2005: two years before and two years after the policy change. Lottery years are excluded. This short panel helps to ensure that the pre-determined potential supply of established H-1Bs is fairly constant across years. Moreover, it has the advantage of avoiding potential changes in individual and firm behavior. For example, one might be concerned that for-profit firms could respond to H-1B restrictions by partnering with non-profit university research affiliates. Similarly, foreign-workers might change their preferences in favor of non-profit work and the guarantee of securing an H-1B (as in Amuedo-Dorantes and Furtado 2016). By exploring effects across a short time period, we reduce the possibility of effects driven adaptive behavior.

The layout of the triple-difference results in Table 4 mirrors that of Table 1. Though the triple-difference term is more muted than in baseline specifications, it remains negative and generally significant. H-1B cap restrictions caused new H-1B hiring to decline between 0.20 and 0.30 log-points when using occupation*education*experience cells, and between 0.13 and 0.14 log-points when using occupation*education cells.

As a final robustness check for the H-1B results, Table 5 considers the role of outliers. The central concern is that the number of H-1B issuances can vary substantially across skill-cells and sectors. We do not want extreme values to be driving results. Table 5 explores the consequences of removing outliers. Columns 1 and 5 reproduce the baseline triple-difference estimates from above for the respective methods of cell-construction. Columns 2 and 6 remove cells in which the logarithm of the combined total H-1B issuances in the

pre-policy period fall more than two standard deviations from the mean. Columns 3 and 7 are instead concerned that H-1B issuances for particular skill-cells might be quite different for the for-profit and non-profit sectors. Thus, we remove observations in which the sectoral difference in log total H-1B issuances in the pre-policy period fall more than two standard deviations from the mean. Finally, columns 4 and 8 simply drop observations for which the dependent variable – log H-1B issuances in a particular cell and year – fall more than two standard deviations from the mean. Resulting coefficients exhibit strikingly little variation across specifications. The triple-difference estimates find large and significant new H-1B employment declines once the cap is reduced. Point estimates for an additional effect during lottery years are negative but generally not statistically different from zero.

Next we analyze whether similar changes were observed in the relative stock of native-born employees after 2004. These results serve as a check on the assumption that there was no deviation in relative demand across sectors before or after 2004. They also serve as test of substitutability between H-1B and native workers. Using data from the American Community Surveys, we define new, native-born, college-educated workers in each skill cell as those having three or fewer years of work experience. This is done to mirror the initial three-year duration of the initial visa period for new H-1B workers. We then use two separate criteria to define experienced workers. First, we use natives who have 4-6 years of experience, which aligns with H-1B workers who have renewed their visa for a second three-year period. Second, we use all native workers with four or more years of experience. Because we cannot differentiate between general work experience and work experience in the United States for natives, we are only able to use skill cells defined by occupation and education.

Table 5 presents results from regressions of (8) using differences in log native employment as the dependent variable. The results are somewhat noisily estimated, but they do not suggest any change in relative demand for native workers after 2004. The gap between new and established native employment in for-profit relative to non-profit sectors did not change significantly after 2004. The point estimate on the Cap-Year*Treated interaction equals about plus or minus ten log points with an equivalent standard error. This evidence supports the idea that the decrease in H-1B employment was due to the fall in the quota and was not attributable to contemporaneous shocks and/or omitted variables. These results also indicate that the fall in H-1B employment did not produce a spillover benefit whereby for-profit firms facing restrictions on foreign workers began hiring natives instead.

In summary, the reduction in the H-1B cap from 195,000 to 65,000 per year in 2004 had its intended impact. The hiring of new foreign-born H-1B workers at for-profit firms fell by roughly 20-50% relative to new foreign-born H-1B workers at non-profit. While not surprising, we are the first to empirically assess the efficacy of the cap reduction. Further,

we find such effects were not accompanied by changes in native employment at for-profit firms. The fact that large declines in H-1B employment were not offset by gains in native employment suggests low substitutability between native-born and H-1B workers in the same skill groups. The reduction in new H-1B workers after the fall in the quota sets the stage for analyzing the effect of policy on H-1B composition, which we turn to next.

6.2 The Composition of H-1B Workers

Our empirical model assumed workers within skill groups and firms were homogeneous. In reality, immigrants and the firms that hire them are heterogeneous. Changes in the characteristics of selected H-1B workers and the types of firms participating in the program could all be important and unintended consequences of immigration policy.

Much of the literature on selection models migration as an individual choice made under constraints, thereby revealing that those who benefit most are the ones more likely to migrate.[15] In practice, however, selection in the H-1B program occurs within a narrow range of college-educated workers and individual choice is not the sole factor – firms play a strong role in selecting workers. Moreover, differential costs in hiring H-1B workers across firms and sectors implies differential levels of efficiency in dealing with the constraints imposed by the quota.

These factors make it difficult to theoretically predict the type of selection generated by the policy change. Analyses of how policy alters immigrant selectivity are, in general, rare.[16] We add to this literature by exploring the potential for the H-1B cap reduction to have altered the selection of H-1B workers that ultimately arrive in the US. We do this without a clear prior of what the direction of selection should be.

6.2.1 Worker Quality

Critics have argued that the H-1B program does not attract the highest ability workers into the country.[17] This section uses wage information to assess how the quota reduction impacted worker quality.

A priori, the effect of restrictive policy on the ability of H-1B workers is ambiguous. On the one hand, the quota reduction generates greater scarcity of foreign labor. Tighter caps could also lead to a selection of higher quality workers if firms choose to only to seek out

[15] Borjas (1987) provided the earliest the formalization of the Roy (1951) model to study migration. More recent work by Chiswick (2000), Grogger and Hanson (2011), Belot and Hatton (2012), and Brücker, Capuano, and Marfouk. (2013) argue that migrants are, in general, positively selected. See Ortega and Peri (2013) or Bansak, Simpson, and Zavodny (2015) for a review.

[16] See Kato and Sparber (2013) and Chen (2005) for examples.

[17] See Matloff (2008) or Hira (2007), for example.

foreign workers with very high marginal productivity and ability. In either case the policy shock in 2004 would cause wages to rise. Efforts to further reduce the quota might even be justifiable on efficiency grounds. On the other hand, high quality workers likely have an abundance of employment opportunities outside the US. Coveted workers might elect to work elsewhere when faced with employment uncertainty due to rationing. In turn, firms may pursue a safer strategy of hiring lower quality workers, helping them to save on costs, rather than spending resources pursuing high achievers. This scenario implies that tighter caps will disproportionately reduce the number of high ability workers, wage offers will fall, and further contractions will exacerbate concerns that H-1B caps inhibit US productivity.

Table 7 performs regressions of (8) using differences in log average wages as the dependent variable. New H-1B workers at for-profit firms continue to constitute the treatment. The format of Table 7 follows that of Table 1. The reduced H-1B cap beginning in 2004 has no statistically significant relationship with wages. Point estimates of the Cap-Year*Treated interaction are small, ranging from a 0.3% wage drop to a 1.1% wage increase. When using cells defined by occupation, education, and experience, however, we do see evidence for a 6.2 to 7.8% wage decline in lottery years for the treated group of new H-1Bs in for-profit firms.

Altogether, the average wage results are rather imprecise and non-robust. A potentially more troubling limitation, however, is that average effects could mask heterogeneous migration effects along the wage distribution. While we know from Table 1 that restrictive H-1B caps reduced total new H-1B employment at for-profit firms, neither those quantity results nor the average wage results in Table 7 are informative about whether the quota had differential effects for the inflows of low, medium, and high wage workers.

To examine this question, we first divide H-1B workers into quintiles of the H-1B wage distribution in which they would have fallen in 2002. The lower-end wage cutoffs for each quintile in real 2010 dollars are {\$0; \$41,088; \$51,660; \$66,550; and \$86,100}. We then return to the triple difference model (8), simply replacing the aggregate sums of H-1B workers in the dependent variable ($H1B_{jkt}^{New}$ and $H1B_{jkt}^{Exp}$) with the sums of H-1B workers from a given quintile ($H1B_{jkt}^{q,New}$ and $H1B_{jkt}^{q,Exp}$, for $q = 1...5$). We then estimate five separate regressions, one for each quintile of the wage distribution.

Figure 3 displays the pertinent estimates from regressions analogous to Columns 3 and 6 from Table 1 that incorporate year and sector-by-skill cell fixed effects. The vertical axis of Figure 3 measures the combined effect of the cap and lottery ($\beta_1 + \beta_2$) on the gap in log job offers between new and established workers at for-profit firms relative to non-profit employers. Markers represent the point estimates, and two-standard-error confidence intervals are provided for reference. The horizontal axis indicates the quintile of the wage distribution. The top panel uses skill cells defined by occupation, education, and experience.

The bottom uses the occupation and education cell definition.

This exercise exposes heterogeneity that is masked when looking at aggregate employment or average wages. New for-profit H-1B employment from the middle of the wage distribution exhibits no statistically significant change arising from the restrictive H-1B quota. Employment losses are concentrated at the tails of the distribution indicating that H-1B restrictions most strongly reduced the number of workers coming from the top and bottom 20% of the H-1B wage distribution. The insignificant average wage estimates of Table 7 arise, in part, because these employment losses offset each other.

In terms of policy ramifications, Figure 3 is perhaps most troubling in that it reveals a particularly sharp decline in employment from the top-end of the wage distribution even after conditioning on occupation, experience, and education. The binding H-1B cap reduced the number of workers who were likely to have been among the most talented and productive foreign individuals seeking US employment. Since employment among the lowest wage workers also fell, it implies that policy has caused the composition of H-1B workers to become more concentrated among workers with middle-levels of skill or ability.

6.2.2 Other Worker Characteristics

A restricted quota could generate changes in the characteristics of hired H-1Bs beyond their productive skill level. For example, it could have altered the cost of hiring some types of workers due to network effects, connections to particular firms, or firm productivity and expertise. Two groups that have garnered significant attention are Indian-born (45% of all H-1B issuances in this period) and computer-related (44%) workers who jointly comprise one third of H-1Bs. Insight into whether cap restrictions alter the proportion of these workers entering the US is important for responding to H-1B critics who argue that the program does not provide any meaningful contribution to innovation or productivity but instead provides a vehicle for firms to import large numbers of low-wage Indian IT workers that displace American computer workers.[18] This section explores how the Indian-born and computer-related occupational composition of H-1B recipients has evolved in response to the cap.

Our regression strategy uses a simple modification of the triple-difference specification in (8). We replace the gap in log employment ($\ln\left(H1B_{jkt}^{New}\right) - \ln\left(H1B_{jkt}^{Exp}\right)$) in the dependent variable with the gap in the share of H-1B employment meeting characteristic c ($\frac{H1B_{jkt}^{c,New}}{H1B_{jkt}^{New}} - \frac{H1B_{jkt}^{c,Exp}}{H1B_{jkt}^{Exp}}$). Specifications include sector-by-skill cell ($j * k$) fixed effects and year indicators, analogous to Columns (3) and (6) of the baseline Table 1 specifications.

Table 8 displays the results. Specifications in columns (1) and (2) define c as employees

[18] See Hira (2007, 2016).

born in India, a group that comprises 44% of the entire sample. The regressions identify whether the reduced H-1B cap affected the proportion of Indian-born workers hired by for-profit firms. Column (1) defines skill cells by occupation, education, and experience. Column (2) defines skill cells only by occupation and education.

In both cases, the effect is small. The first point estimate finds a 2.3 percentage-point increase in the Indian-born share of new H-1B employment at for-profit firms. However, the estimate does not differ from zero at conventional significance levels. Some of this imprecision might be due to a lack of data variation: about two-thirds of the observations have Indian-share values that equal zero. The higher level of aggregation defined by occupation and education cells in Column (2) helps decrease the preponderance of zero values. The estimate rises to a 3.9 percentage-point increase in the Indian-born share of new H-1B hiring at for-profit firms. The coefficient is significant at the 10% level. These results indicate that quota restrictions may have also had the unintended effect of shifting the composition of new H-1B workers towards having Indian-origin and away from other source countries.

Column (3) examines changes in occupational composition and defines c as computer-related workers. This forces us to leave our usual skill cell definitions in favor of one that defines k according to education and experience only, thereby greatly reducing the number of available observations. Nonetheless, regressions uncover important significant effects. The proportion of new job offers extended to computer-related workers at for-profit firms grew by 5.5 percentage-points as a result of the H-1B cap, and it grew another 4.6 percentage points during the period of heightened H-1B demand in fiscal years 2008 and 2009. Again, this suggests an unintended consequence in that policy did not affect all sectors equally.

These results are fascinating, in part, due to their irony: many opponents of the H-1B program who advocate stricter limits lament the number of issuances to Indian-born computer-related workers. But those same limits have led to a compositional shift favoring those workers. These findings are also interesting in their implications for underlying economic behavior. Namely, job search costs and uncertainty in the H-1B application process could favor firms with past H-1B employment experience and new workers with existing labor market networks. The next section examines heterogeneous effects across firms with regressions that identify results consistent with this possibility.

6.3 Firm Participation

The reduction in the number of H-1Bs available altered the composition of H-1B workers, but the effects might not be limited to recipients. The policy change could alter the types of firms participating in the H-1B program as well. Firms' costs of participating in the H-1B

program can be quite significant even in the absence of a quota. In addition to application fees that can range in the thousands of dollars, firms often need legal assistance to deal with the adjudication process. This can incur additional costs for firms that need to outsource such legal tasks. Moreover, firms cannot escape search costs in the hiring process, or costs associated with the uncertainty regarding the true productivity of a worker.

Costs rise in the presence of a restrictive quota. USCIS's allocation procedure requires firms to act quickly to extend offers to desired candidates before permits are exhausted. Uncertainty also rises. Even after identifying its desired workers, the lotteries meant that there was no guarantee that firms would be legally allowed to hire them. These costs could vary across firms for a variety of reasons. Large firms with in-house legal teams might have an easier time navigating the bureaucratic hiring process. They might also be better equipped at absorbing labor shocks that arise when their actual number of job offers winning the H-1B lottery is greater or less than what the firm had expected. Thus, it is conceivable that the increased uncertainty induced by strict H-1B quotas has had a differential effect across employers.

Our dataset includes the firm name of each individual's employer but with an important caveat: Data on this variable is subject to a large degree of measurement error created by alternative firm name spellings and typos. For example, while "Microsoft Corp" accounts for 99% of all firm names beginning with "Microsoft" in Redmond, Washington for fiscal year 2007, there are also observations for "Microsoft Coporation (sic)," "Microsoft Coprporation (sic)," and "Microsoft Inc," among several other related entries. Though performing manual or automated routines to group names of likely-identical firms would be possible, this too could generate measurement error. For the exercise in this section, we take only minimal steps to harmonize firm names and rely primarily on the information provided by USCIS, although we do drop individuals whose employer names are missing or unknown. The heaviest for-profit users of the H-1B program during this period were Infosys Technology Limited, Microsoft Corporation, Cognizant Technology Solutions US Corporation, Wipro Limited, and Intel Corporation. The heaviest non-profit users were Yale University, University of Michigan, Stanford University, and Columbia University. Hospitals and medical clinics are among the most significant non-university cap-exempt H-1B employers.

As a first step in empirically assessing whether the H-1B quota led a shift in the types of firms hiring new H-1B workers, we consider the unequal distribution of H-1Bs across firms over time. The four panels of Figure 4 plot Lorenz curves for four separate H-1B groups: (1) new hires at for-profit (cap-bound) firms, (2) new hires at non-profit (cap-exempt) firms, (3) experienced hires at for-profit firms, and (4) experienced hires at non-profit firms. Each panel plots curves for the pre-binding-cap period (2002-2003), the binding-cap pre-lottery period

(2004-2007), and the lottery years (2008-2009). The vertical axes measure the cumulative percentage of H-1B issuances. The horizontal axes measure the cumulative percentage of firms. A higher concentration of H-1Bs among a smaller number of employers would be indicated by a shift of the Lorenz curve to the lower-right; a Lorenz curve of complete equality would fall along a 45-degree line.

As can be seen in the graphs, the curves for pre- and post-cap periods are similar in all cases except for that of new hires at for-profit firms. Prior to the cap reduction 80% of for-profit firms accounted for 40% of new H-1B hires. This concentration grows after the fall in the cap when 80% of firms account for only 20% of new hires. This growth in the concentration of H-1B visas among firms does not appear in any of the other groups.

Figure 5 provides another visual representation of the rising concentration of H-1B employment among for-profit firms. First we compute two separate indices of inequality for each of the four groups of interest. The first is a Gini coefficient arising from annual Lorenz curves analogous to those in Figure 4. The second is a Herfindahl Index of concentration that measures the sum of squared H-1B employment shares across firms in each of the four groups. We then calculate the double-difference of each index, measuring the gap in inequality between new and established workers at cap-bound versus cap-exempt employers. Consistent with Figure 4, both indices rise after 2004, suggesting new H-1B hires are increasingly concentrated among fewer for-profit firms.

To formalize these concepts in a regression framework, we define firm characteristic c as "large" H-1B employers that hired 50 or more total H-1Bs in a given year.[19] As stated earlier, large firms possess economies of scale that likely allow them to maintain lower costs of hiring H-1B workers. Firms with fewer H-1B workers and less experience with the program often need to outsource the cost of hiring H-1B workers to legal firms. A fall in the ex-ante payoff to hiring a foreign skilled worker due to declines in H-1B limits might induce small firms to reduce participation.

Analogous to our exploration of heterogeneous implications across worker characteristics, our regression strategy in this section measures the difference in the proportion of new versus established H-1B workers hired by large H-1B firms. The regression in Column (4) of Table 8 uses occupation*education*experience skill cells; Column (5) uses occupation*education cells. We again see that the level of aggregation plays a role in determining the significance of coefficients. In Column (4), the point estimate suggests large H-1B firms see a small and insignificant 1.5 percentage-point rise in their proportion of new, for-profit, H-1B workers. However, two-thirds of the share values equal zero. With the higher level of aggregation

[19] We do not use a static definition based on a pre-period base-year to avoid measurement error. We do not want the regressions to be confounded by firm names that might change over time.

in the cell construction used in Column (5), the point estimate rises substantially and is significant at the 5% level. The binding cap causes large firms to account for a 5.9 percentage point increase in the share of new H-1B employment at for-profit firms. Thus, despite the decrease in total new, for-profit, H-1B employment, large firms have an advantage in hiring the workers they seek.

In terms of composition, declines in new H-1B hiring are concentrated among firms that use the program sparingly. The change in H-1B policy appears to have shifted foreign labor resources away from firms that employ H-1B workers less-intensively toward larger ones that might be better able to provide legal services for hiring and/or are more capable of absorbing employment shocks generated by the lottery. The change in the types of firms participating might also be related to changes in the types of workers that were hired. The evidence for compositional changes of firms and workers is suggestive of important network effects. Policy changes restricting inflows of H-1B workers are more punitive to nativities, occupations, and employers with less experience with the H-1B program.

7 Conclusion

In the early 2000s, the United States temporarily raised the annual cap on new H-1B issuances to employees of most firms to 195,000 per year while permanently exempting employees of colleges, universities, and non-profit research institutions from the quota. The cap on new employment at for-profit firms reduced to 65,000 in fiscal year 2004 when Congress declined to renew legislation maintaining the high limit. With the addition of 20,000 H-1Bs available to foreign-workers who have obtained advanced degrees from US universities, the cap increased to 85,000 in 2005 where it has remained ever since. Rising H-1B interest among foreign workers and US firms that wish to hire them has led all new H-1Bs for employees of for-profit firms to be allocated by lottery in recent years.

We presume that by letting the 195,000 quota lapse, policy-makers intended to reduce new H-1B employment at most firms. Initially this caused firms to simply substitute away from new H-1B workers and toward established ones. Over time, however, total hiring did decline. We estimate that new H-1B employment at cap-bound firms declined roughly 20 to 50% compared to what it otherwise would have been.

In addition to these direct effects, the cap restriction also generated consequences that were presumably less-intended. Perhaps most troubling, H-1B declines are concentrated at the lowest and highest ends of the wage distribution. In the latter case, this suggests that it is the highest ability workers with the highest earnings potential who are most likely to be deterred from entering the US labor market as a result of H-1B restrictions. Given the

potential for productivity-enhancing technological gains generated by H-1B workers, this loss could reverberate throughout the economy.

Other important effects are distributional. H-1B restrictions have led to a compositional shift in new H-1B employment favoring Indian-born workers, computer-related occupations, and firms that use the H-1B program heavily. One interpretation of these results is that acute visa restrictions increase the importance of labor networks, economies of scale in hiring foreign-labor, and skill in navigating the H-1B program.

Expanded work on the consequences of immigration policy on the selection of workers and firms remains crucial. As current efforts turn toward reforming immigration policy, legislators have placed a renewed emphasis on the H-1B program. Frustration with the current state of the H-1B program – the large amount of Indian computer scientist, the domination by huge IT firms, and the potential displacement of native workers – has generated increased pressure for further quota restrictions. Our work indicates that further reducing the H-1B cap is likely to skew the characteristics of the H-1B program further in this direction.

References

[1] Amuedo-Dorantes, Catalina and Delia Furtado (2016), "Settling for Academia? H-1B permits and the career choices of international students in the United States."

[2] Anderson, Stuart (2015), "Ted Cruz's Bad Idea: Expel international students from America" *The Daily Caller*. 15 December 2015.

[3] Bansak, Cynthia, Nicole B. Simpson, and Madeline Zavodny (2015), *The Economics of Immigration*. Routledge, New York, NY.

[4] Belot, Michèle VK, and Timothy J. Hatton (2012), "Immigrant Selection in the OECD." *The Scandinavian Journal of Economics* 114.4: 1105-1128.

[5] Borjas, George J. (1987), "Self-Selection and the Earnings of Immigrants." *The American Economic Review*: 531-553.

[6] Borjas, George J. (2003), "The Labor Demand Curve is Downward Sloping: Reexamining the impact of immigration on the labor market." *Quarterly Journal of Economics*, 118 (4), 1335-1374.

[7] Brücker Herbert, Stella Capuano, and Abdeslam Marfouk (2013), "Education, Gender and International Migration: Insights from a panel-dataset 1980-2010." Mimeo.

[8] Card, David, and Alan B. Krueger (1994), "Minimum Wages and Employment: A case study of the fast-food industry in New Jersey and Pennsylvania." *American Economic Review* 84.4: 772-93.

[9] Card, David, and Alan B. Krueger (2015), *Myth and Measurement: The new economics of the minimum wage*. Princeton University Press.

[10] Card, David, and Thomas Lemieux (2001), "Can Falling Supply Explain the Rising Return to College for Younger Men? A cohort-based analysis" *The Quarterly Journal of Economics* 116, no. 2: 705-746.

[11] Chen, Yiu Por (2005), "Skill-Sorting, Self-Selectivity, and Immigration Policy Regime Change: Two surveys of Chinese graduate students' intention to study abroad" *AEA Papers and Proceedings*, Vol. 95(2), 66-70.

[12] Chiswick, Barry R. (2000), "Are Immigrants Favorably Self-Selected? An economic analysis" IZA Discussion Paper No. 131.

[13] Doran, Kirk, Alexander Gelber, and Adam Isen (2014), "The Effects of High-skilled Immigration Policy on Firms: Evidence from H-1B visa lotteries". Working Paper No. w20668. National Bureau of Economic Research.

[14] Fernandez Campbell, Alexia (2016), "There's a Clear Way to Fix the H-1B Visa Program." *The Atlantic*. December 6, 2016.

[15] Ghosh, Anirban, Anna Maria Mayda, and Francesco Ortega (2016), "The Impact of Skilled Migration on Firm-level Productivity: An investigation of publicly traded US firms." IZA Working Paper, IZA Discussion Paper 8684.

[16] Grogger, J. and G.H. Hanson (2011), "Income Maximization and the Selection and Sorting of International Migrants." *Journal of Development Economics*, Vol 103: 182-201.

[17] Hira, Ron (2007), "Outsourcing America's Technology and Knowledge Jobs: High-Skill guest worker permits are currently hurting rather than helping keep jobs at home." Economic Policy Institute Briefing Paper no. 187.

[18] Hira, Ron (2016), "The Impact of High-Skilled Immigration on U.S. Workers." Congressional Testimony February 25th, 2016. Economic Policy Institute `http://www.epi.org/publication/congressional-testimony-the-impact-of-high-skilled-immigration-on-u-s-workers-4/`

[19] Kato, Takao, and Chad Sparber (2013), "Quotas and Quality: The effect of H-1B visa restrictions on the pool of prospective undergraduate students from abroad." *Review of Economics and Statistics* 95.1: 109-126.

[20] Kerr, William and William F. Lincoln (2010), "The Supply Side of Innovation: H-1B visa reforms and US ethnic invention." *Journal of Labor Economics* 28(3): 2010, pp. 473-508.

[21] Matloff, Norman (2008), "H-1Bs: Still not the best and the brightest," Center for Immigration Studies Backgrounder.

[22] Ortega, Francesc, and Giovanni Peri (2013), "The Effect of Income and Immigration Policies on International Migration," *Migration Studies* 1, no. 1: 47-74.

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

[24] Peri Giovanni, Kevin Shih and Chad Sparber (2015), "STEM Workers, H-1B Visas and Productivity in US Cities," *Journal of Labor Economics* 33, no. S1 Part 2 (2015): S225-S255.

[25] Preston, Julia (2016), "Judge Says Disney Didn't Violate Visa Laws in Layoffs." *New York Times*. 13 October 2016.

[26] Roy, Andrew Donald (1951), "Some Thoughts on the Distribution of Earnings." Oxford Economic Papers 3, no. 2: 135-146.

[27] Shih, Kevin (2016), "Labor Market Openness, H-1B Visa Policy, and the Scale of International Student Enrollment in the United States." *Economic Inquiry* 54.1: 121-138.

[28] US Citizenship and Immigration Services (2007). "USCIS Updates Count of FY 2008 H-1B Cap Filings," April 10, 2007.

[29] US Citizenship and Immigration Services (2008). "USCIS Runs Random Selection Process for H-1B Petitions," March 31, 2008.

8 Appendix

Valid double and triple-difference estimates require that treatment and control groups do not exhibit differential trends prior to the policy shock. Figure 2 attempts to illustrate the absence of differential pre-trends in H-1B issuances, but the period prior to the policy shock is short. As an alternative source of illustrative evidence, we turn to Current Population Survey (CPS) data.

Recall that although October marks the beginning of a fiscal year, H-1B filings begin six months earlier. To mirror this behavior, we aggregate college-educated, native-born, employment data into quasi-fiscal years that begin in April. Given characteristics of the H-1B program, we are interested in two groups: employees at private for-profit firms and employees of private and public colleges and universities. We normalize each group's total employment so that Figure A1 illustrates within-group total employment relative to 2000.

This exercise demonstrates that college-educated, native-born, employment growth was fairly constant and comparable across the for-profit and university sectors. Employment was 8 to 10 percent lower in 1997 than in 2000, and 7 to 9 percent higher in 2004. These parallel employment trends help to give us confidence that sectoral differences prior to the policy shock in fiscal year 2004 are not driving our H-1B employment results, and that our empirical methodology is valid.

Figure 1: Aggregate New H-1B Issuances by Skill-Cell

Panel (a)

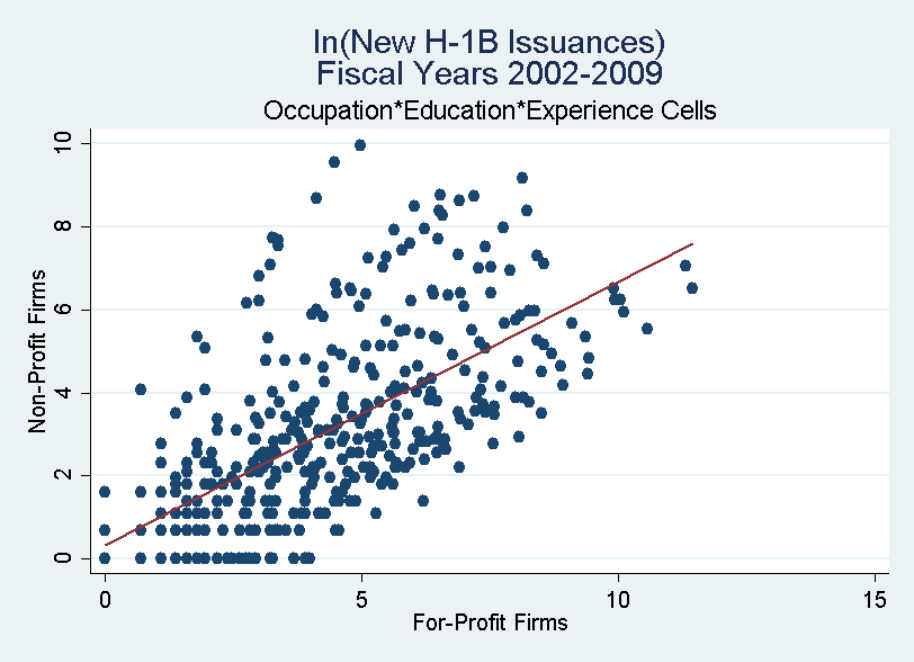

Panel (b)

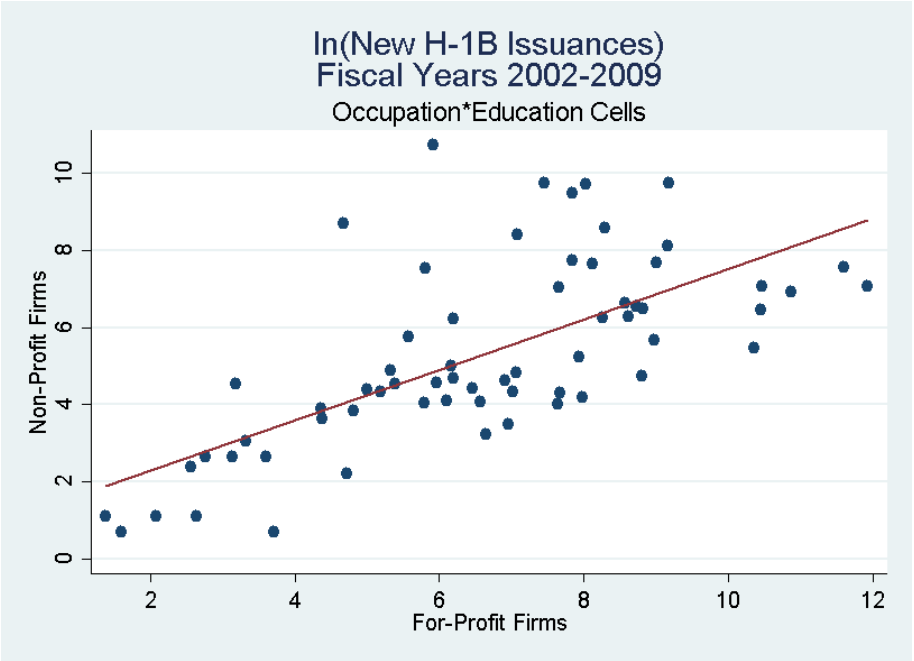

Note: Each point represents the number of new H-1Bs issued for a skill-cell in the whole 2002-2009 period. Panel (a) defines skill-cells by occupation, education, and experience, while Panel (b) uses education and occupation cells. The line shows the linear correlation between non-profit and for-profit new H-1B employment across cells.

Figure 2: Average Difference in New Versus Established Log-H-1B Issuances.

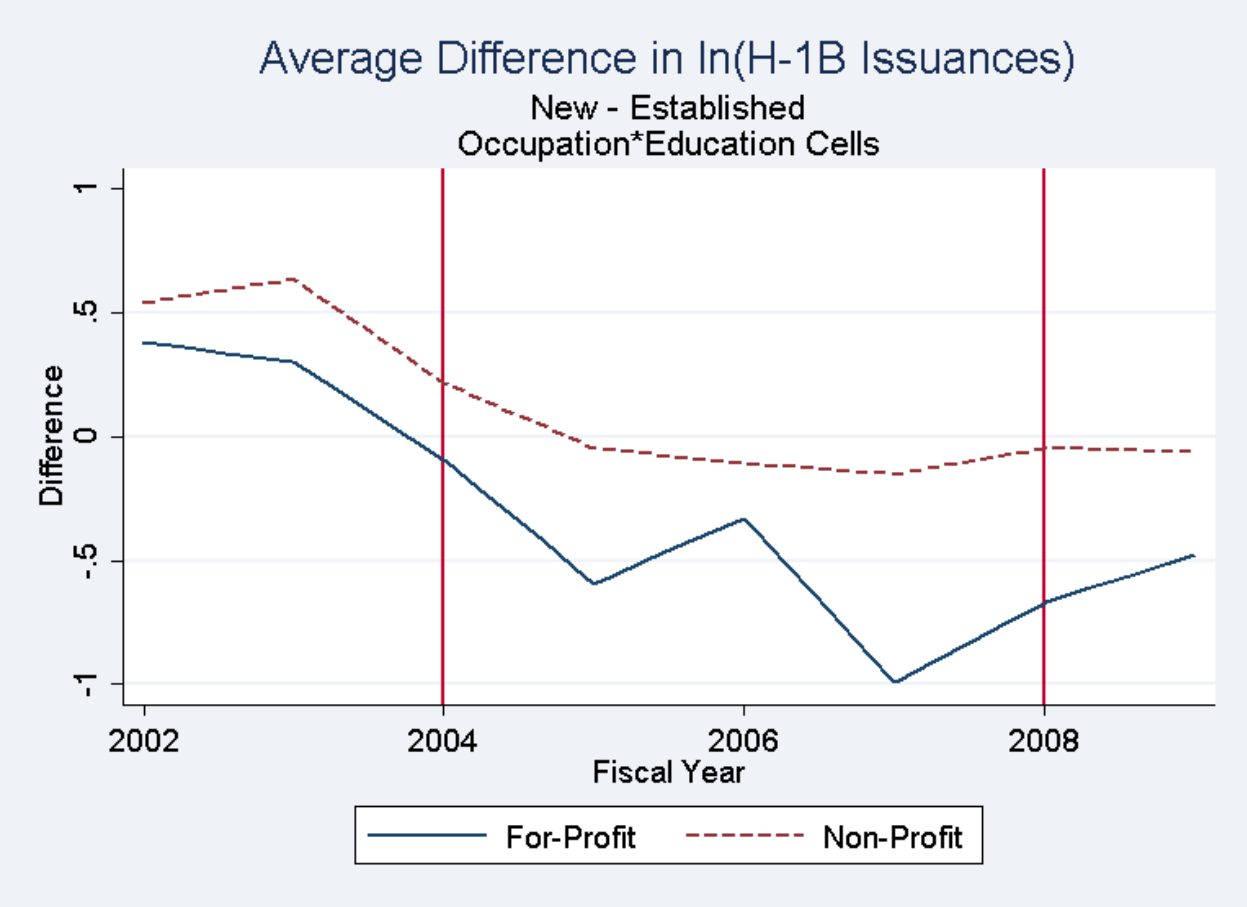

Note: The solid line represents the average difference in log issuances of new versus established H-1Bs for the non-profit sector across skill cells. The dashed line represents the same difference for the for-profit sector. Cells are defined by occupation and education.

Figure 3: Change in For-Profit H-1B Employment at Different Quintiles of the Wage Distribution

*Sum of coefficients on the Cap-Year*Treated plus Lottery*Treated Interactions*

Panel (a)

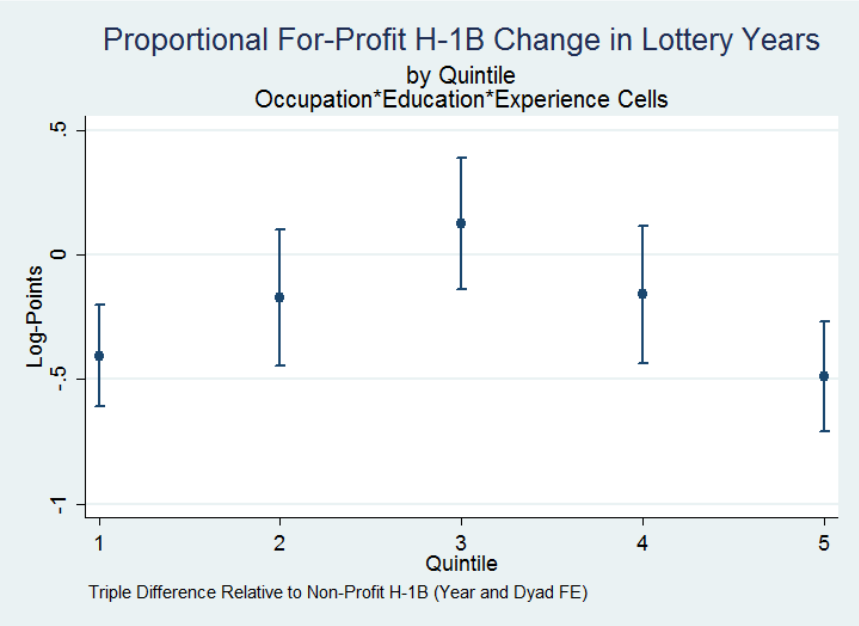

Panel (b)

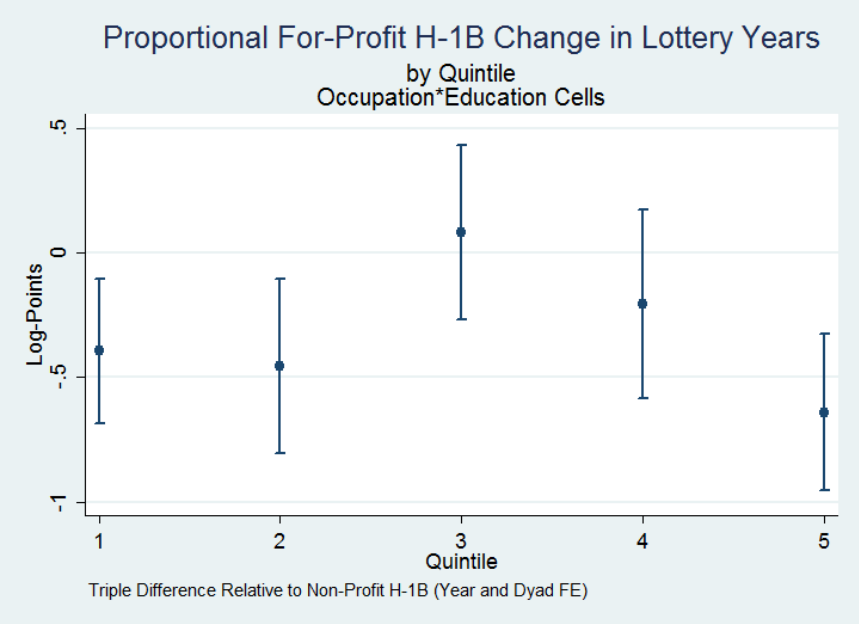

Note: The graphs show the sum of estimated coefficients on the Cap-Year and Lottery-Year interactions with Treated when the dependent variable is the difference in log-issuances in for-profit relative to non-profit firms at the specified quintile of the wage distribution. The vertical bar shows the 95% confidence interval. Panel (a) uses estimates from skill-cells constructed using occupation, education, and experience groups. Panel (b) uses estimates from skill-cells defined only by education and occupation.

Figure 4: H-1B Concentration in Firms

Lorenz Curves for the Inequality of H-1B Issuances across Firms

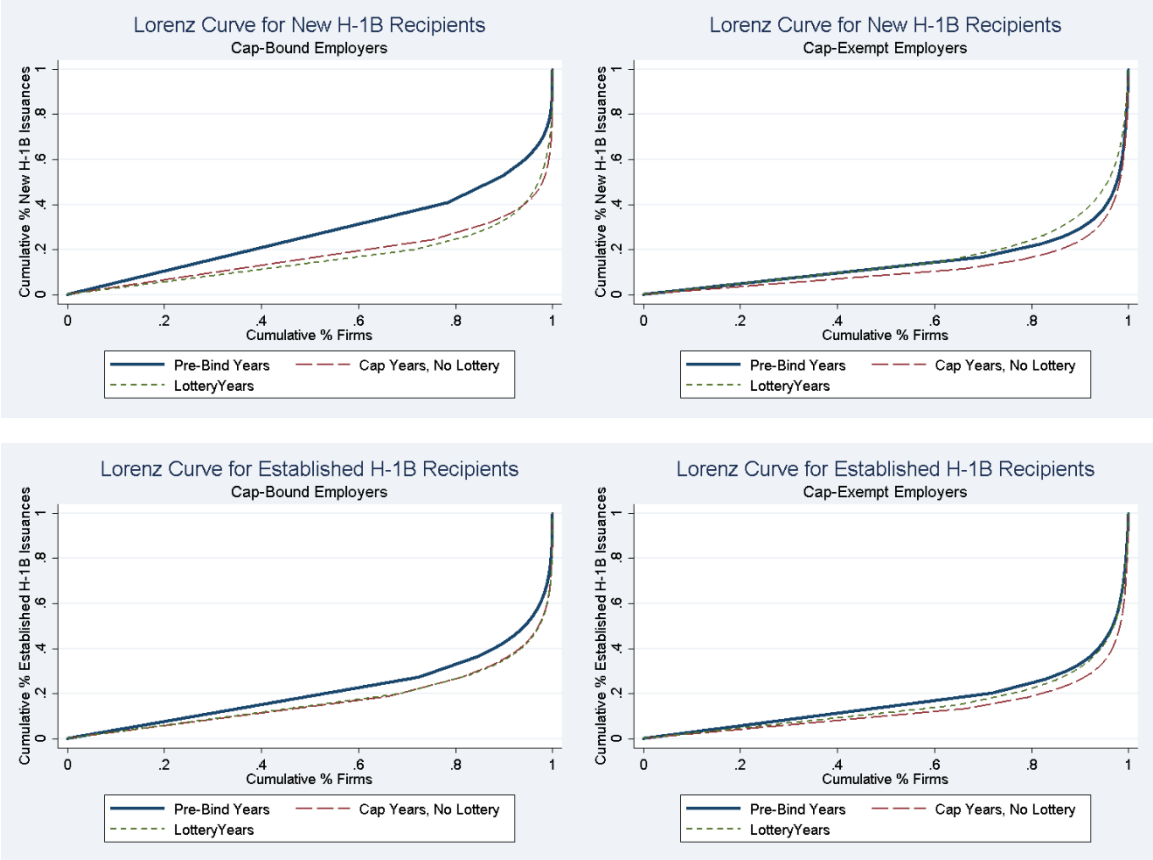

Note: Each panel show the Lorenz curve of H-1B concentration for a group (new and established) and a sector (for profit and non-profit). The three line correspond to the curve pre-2004 when the cap was high (blue, solid, thick), fiscal years 2004-2007 with lower cap (red, long-dash), and 2008-09 with lower cap and Lottery (green, short-dash).

Figure 5: H-1B Concentration in Firms, Gini Coefficient and Herfindahl Index

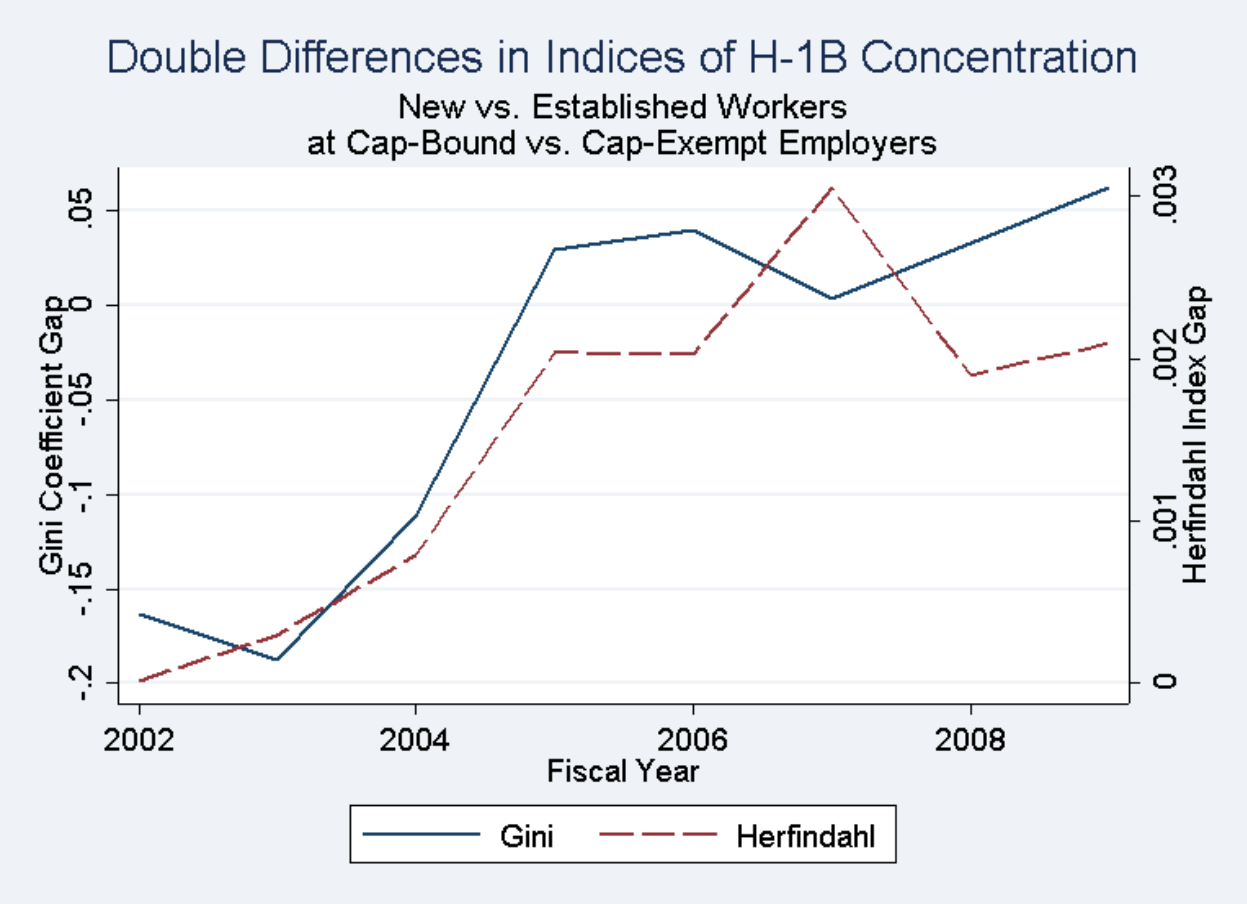

Note: We calculate indices of H-1B concentration within firms for four groups: New workers at for-profit firms, established workers at for-profit firms, new workers at non-profit firms, and established workers at non-profit firms. We then calculate the double-difference in these indices. The figure indicates a proportional increase in H-1B concentration within firms for new H-1B workers at for-profit firms relative to other groups. This is true using both the Gini Coefficient of inequality (blue, sold, left axis) and the Herfindahl (red, dashed, right axis) index of concentration.

Table 1: Baseline Triple-Difference Results for H-1B Issuances

	(1)	(2)	(3)	(4)	(5)	(6)
Cell Definition:	OCC*ED*EXP			OCC*ED		
Panel Fixed Effects:	Skill-Cell + Sector	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell + Sector	Skill-Cell + Sector	Skill-Cell * Sector
Time Fixed Effects:	Cap Year + Lottery Year	Year	Year	Cap Year + Lottery Year	Year	Year
Balanced Panel:	No	No	Yes	No	No	Yes
Cap Year * Treated	-0.389***	-0.403***	-0.331***	-0.181***	-0.187***	-0.202**
	(0.047)	(0.047)	(0.060)	(0.066)	(0.065)	(0.078)
Lottery Year * Treated	-0.052	-0.036	-0.061	-0.078	-0.069	-0.034
	(0.044)	(0.045)	(0.048)	(0.072)	(0.071)	(0.074)
Treated	-0.141***	-0.141***		-0.280***	-0.279***	
	(0.041)	(0.041)		(0.069)	(0.069)	
Observations	4,556	4,556	3,512	938	938	894
R-squared	0.574	0.621	0.707	0.476	0.542	0.632

Note: Dependent variable is the logarithm of new H-1B issuances minus the logarithm of established H-1B issuances in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2009. Each column represents an estimated specification. Columns 1-3 define skill groups by occupation, education, and experience, while columns 4-6 use occupation and experience. The variable "Cap Year" is equal to one in fiscal years 2004-2009 and zero otherwise. The variable "Lottery Year" is equal to one in fiscal years 2008 and 2009 and zero otherwise. The treated variable is a dummy equal to one for cells in the for-profit sector. In columns 1 and 4 skill-cell and sector fixed effects are included plus cap year and lottery year dummies. In columns 2 and 5 skill-cell and sector fixed effects plus year effects are included. In columns 3 and 6 we include skill-cell by sector effects and year effects. The method of estimation is least squares with standard errors clustered at the skill cell level. *** p<0.01, ** p<0.05, * p<0.1

Table 2: Double-Difference Source of Triple-Difference Effect

Cell Definition:		OCC*ED*EXP			OCC*ED		
		New	**Established**	**Difference**	**New**	**Established**	**Difference**
For-Profit	Cap-Years	-0.489***	0.492***	-0.981***	-0.393***	0.432***	-0.826***
		(0.031)	(0.025)	(0.043)	(0.052)	(0.032)	(0.058)
	Lottery-Years	-0.335***	-0.180***	-0.155***	-0.290***	-0.222***	-0.068*
		(0.027)	(0.026)	(0.030)	(0.046)	(0.037)	(0.039)
Non-Profit	Cap-Years	-0.207***	0.442***	-0.650***	-0.208***	0.416***	-0.624***
		(0.030)	(0.031)	(0.042)	(0.048)	(0.045)	(0.051)
	Lottery-Years	-0.039	0.054*	-0.093**	0.000	0.034	-0.034
		(0.032)	(0.029)	(0.038)	(0.060)	(0.040)	(0.064)
Difference	Cap-Years	-0.282***	0.050	-0.331***	-0.186***	0.016	-0.202***
		(0.037)	(0.031)	(0.049)	(0.053)	(0.056)	(0.071)
	Lottery-Years	-0.295***	-0.234***	-0.061	-0.290***	-0.256***	-0.034
		(0.040)	(0.038)	(0.049)	(0.068)	(0.051)	(0.071)
Number of Skill-Cells*Years:				1,756			447

Note: Dependent variable is the logarithm of H-1B issuances in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2009. Cells contain single-difference estimates for Cap-Years (2004-2009) and Lottery-Years (2008-2009). "Difference" columns represent the double-difference estimate for new versus established H-1B issuances within sector. "Difference" row represents double-difference estimates for for-profit versus non-profit H-1B issuances within H-1B experience. Corner cells represent triple-difference estimates for new H-1B issuances at for-profit firms. The left grouping defines skill groups by occupation, education, and experience, while the right uses occupation and experience. Standard errors are clustered at the skill cell level and are in parenthesis. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Double-Difference Results for Total H-1B Issuances

	(1)	(2)	(3)	(4)	(5)	(6)
Cell Definition:		OCC*ED*EXP			OCC*ED	
Panel Fixed Effects:	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector
Time Fixed Effects:	CapYear + LotteryYear	Year	Year	CapYear + LotteryYear	Year	Year
Balanced Panel:	No	Double-Difference Panel	Triple-Difference Panel	No	Double-Difference Panel	Triple-Difference Panel
CapYear * Treated	-0.067**	-0.043	-0.036	-0.013	0.007	-0.022
	(0.031)	(0.032)	(0.030)	(0.044)	(0.051)	(0.049)
LotteryYear * Treated	-0.230***	-0.272***	-0.250***	-0.275***	-0.299***	-0.285***
	(0.032)	(0.033)	(0.032)	(0.049)	(0.059)	(0.053)
Treated	1.428***			1.544***		
	(0.111)			(0.272)		
Observations	6,024	5,122	3,512	985	954	894
R-squared	0.774	0.974	0.982	0.813	0.986	0.989

Note: Dependent variable is the logarithm of total H-1B issuances in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2009. Each column represents an estimated specification. Columns 1-3 define skill groups by occupation, education, and experience, while columns 4-6 use occupation and experience. The variable "Cap Year" is equal to one in fiscal years 2004-2009 and zero otherwise. The variable "Lottery Year" is equal to one in fiscal years 2008 and 2009 and zero otherwise. The treated variable is a dummy equal to one for cells in the for-profit sector. In columns 1 and 4 skill-cell and sector fixed effects are included plus cap year and lottery year dummies. Columns 2, 3, 5, and 6 include skill-cell by sector effects and year effects. The method of estimation is least squares with standard errors clustered at the skill cell level. *** p<0.01, ** p<0.05, * p<0.1

Table 4: Short Panel Triple-Difference Results for H-1B Issuances

	(1)	(2)	(3)	(4)	(5)	(6)
Cell Definition:		OCC*ED*EXP			OCC*ED	
Panel Fixed Effects:	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector
Time Fixed Effects:	CapYear	Year	Year	CapYear	Year	Year
Balanced Panel:	No	No	Yes	No	No	Yes
CapYear * Treated	-0.300***	-0.300***	-0.199***	-0.129*	-0.133*	-0.141
	(0.054)	(0.054)	(0.066)	(0.076)	(0.076)	(0.092)
Treated	-0.130***	-0.128***		-0.278***	-0.277***	
	(0.043)	(0.043)		(0.072)	(0.072)	
Observations	2,315	2,315	1,770	472	472	450
R-squared	0.559	0.590	0.722	0.472	0.523	0.682

Note: Dependent variable is the logarithm of new H-1B issuances minus the logarithm of the established H-1B issuances in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2005. Each column represents an estimated specification. Columns 1-3 define skill groups by occupation, education, and experience, while columns 4-6 use occupation and experience. The variable "Cap Year" is equal to one in fiscal years 2004-2005 and zero otherwise. The treated variable is a dummy equal to one for cells in the for-profit sector. In columns 1 and 4 skill-cell and sector fixed effects are included plus a cap year dummy. In columns 2 and 5 skill-cell and sector fixed effects plus year effects are included. In columns 3 and 6 we include skill-cell by sector effects and year effects. The method of estimation is least squares with standard errors clustered at the skill cell level. *** p<0.01, ** p<0.05, * p<0.1

Table 5: Triple-Difference Results for H-1B Issuances, Outliers Removed

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cell Definition:	OCC*ED*EXP				OCC*ED			
Cells Removed if Extreme Values of:	N/A	ln(Total H-1B Issuances in 2002 & 2003)	For-Profit and Non-Profit Difference in ln(Issuances in 2002 & 03)	The Dependent Variable	N/A	ln(Total H-1B Issuances in 2002 & 2003)	For-Profit and Non-Profit Difference in ln(Issuances in 2002 & 03)	The Dependent Variable
CapYear * Treated	-0.331***	-0.352***	-0.345***	-0.353***	-0.202**	-0.231***	-0.199**	-0.250***
	(0.060)	(0.059)	(0.062)	(0.060)	(0.078)	(0.073)	(0.080)	(0.077)
LotteryYear * Treated	-0.061	-0.072	-0.080	-0.094*	-0.034	-0.040	-0.046	-0.028
	(0.048)	(0.049)	(0.050)	(0.048)	(0.074)	(0.073)	(0.076)	(0.068)
Observations	3,512	3,380	3,358	3,240	894	854	862	830
R-squared	0.707	0.710	0.712	0.711	0.632	0.645	0.630	0.648

Note: Dependent variable is the logarithm of new H-1B issuances minus the logarithm of the established H-1B issuances in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2009. Each column represents an estimated specification. Columns 1-4 define skill groups by occupation, education, and experience, while columns 5-8 use occupation and experience. The variable "Cap Year" is equal to one in fiscal years 2004-2009 and zero otherwise. The variable "Lottery Year" is equal to one in fiscal years 2008 and 2009 and zero otherwise. The treated variable is a dummy equal to one for cells in the for-profit sector. All columns include skill-cell by sector effects and year effects. Columns 1 and 5 repeat previous estimates from balanced panels. Remaining columns remove outliers identified by column headers and in the main text. The method of estimation is least squares with standard errors clustered at the skill cell level. *** p<0.01, ** p<0.05, * p<0.1

Table 6: Triple-Difference Results for Native-Born Employment

	(1)	(2)	(3)	(4)	(5)	(6)
Experience Definition	4 - 6 Years			More than 3 Years		
Panel Fixed Effects:	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector
Time Fixed Effects:	Cap Year + Lottery Year	Year	Year	Cap Year + Lottery Year	Year	Year
Balanced Panel:	No	No	Yes	No	No	Yes
Cap Year * Treated	-0.011	-0.009	0.106	-0.116	-0.115	-0.109
	(0.102)	(0.102)	(0.103)	(0.093)	(0.093)	(0.087)
Lottery Year * Treated	-0.050	-0.053	-0.110	-0.023	-0.022	-0.077
	(0.092)	(0.093)	(0.085)	(0.071)	(0.071)	(0.063)
Cap-Bound Employer	0.134	0.134		0.518***	0.516***	
	(0.081)	(0.082)		(0.095)	(0.095)	
Observations	759	759	686	798	798	718
R-squared	0.326	0.331	0.425	0.656	0.660	0.755
All Regressions Define Skill Cell by Occupation*Education						

Note: Dependent variable is the logarithm of new native employment minus the logarithm of the experienced native employment in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2009. Each column represents an estimated specification. The variable "Cap Year" is equal to one in fiscal years 2004-2009 and zero otherwise. The variable "Lottery Year" is equal to one in fiscal years 2008 and 2009 and zero otherwise. The treated variable is a dummy equal to one for cells in the for-profit sector. In columns 1 and 4 skill-cell and sector fixed effects are included plus cap year and lottery year dummies. In columns 2 and 5 skill-cell and sector fixed effects plus year effects are included. In columns 3 and 6 we include skill-cell by sector effects and year effects. The method of estimation is least squares with standard errors clustered at the skill cell level. *** p<0.01, ** p<0.05, * p<0.1

Table 7: Triple-Difference Results for H-1B Wages

	(1)	(2)	(3)	(4)	(5)	(6)
Cell Definition:	OCC*ED*EXP			OCC*ED		
Panel Fixed Effects:	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector	Skill-Cell + Sector	Skill-Cell * Sector	Skill-Cell * Sector
Time Fixed Effects:	CapYear + LotteryYear	Year	Year	CapYear + LotteryYear	Year	Year
Balanced Panel:	No	No	Yes	No	No	Yes
Cap Year * Treated	0.009	0.009	0.005	0.011	0.011	-0.003
	(0.022)	(0.022)	(0.023)	(0.040)	(0.040)	(0.036)
Lottery Year * Treated	-0.071***	-0.071***	-0.082***	-0.009	-0.009	-0.035
	(0.026)	(0.026)	(0.025)	(0.044)	(0.045)	(0.038)
Cap-Bound Employer	-0.047**	-0.047**		-0.059*	-0.059*	
	(0.019)	(0.019)		(0.032)	(0.032)	
Observations	4,556	4,556	3,512	938	938	894
R-squared	0.181	0.181	0.294	0.137	0.139	0.288

Note: Dependent variable is the average logarithm of New H-1B wage minus the logarithm of the established H-1B wage in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2009. Each column represents an estimated specification. Columns 1-3 define skill groups by occupation, education, and experience, while columns 4-6 use occupation and experience. The variable "Cap Year" is equal to one in fiscal years 2004-2009 and zero otherwise. The variable "Lottery Year" is equal to one in fiscal years 2008 and 2009 and zero otherwise. The treated variable is a dummy equal to one for cells in the for-profit sector. In columns 1 and 4 skill-cell and sector fixed effects are included plus cap year and lottery year dummies. In columns 2 and 5 skill-cell and sector fixed effects plus year effects are included. In columns 3 and 6 we include skill-cell by sector effects and year effects. The method of estimation is least squares with standard errors clustered at the skill cell level. *** p<0.01, ** p<0.05, * p<0.1

Table 8: Triple-Difference Results for H-1B Compositional Changes

	(1)	(2)	(3)	(4)	(5)
Share of Group:	Indian-Born Workers		Computer-Related Workers	Large H-1B Firms	
Cell Definition:	OCC*ED*EXP	OCC*ED	ED*EXP	OCC*ED*EXP	OCC*ED
CapYear * Treated	0.023	0.039*	0.055**	0.015	0.059**
	(0.014)	(0.021)	(0.025)	(0.018)	(0.023)
LotteryYear * Treated	-0.005	0.002	0.046*	-0.000	-0.036*
	(0.015)	(0.025)	(0.025)	(0.018)	(0.021)
Observations	3,512	894	506	3,506	894
R-squared	0.309	0.259	0.431	0.231	0.308

Note: Dependent variable is the share of workers with a characteristic (identified in the top row) among new H-1B workers minus the same share among established H-1B workers in the skill-cell and sector (for-profit and nonprofit) in each year 2002-2009. Columns 1 and 4 define skill groups by occupation, education, and experience. Columns 2 and 5 use occupation and experience. Column 3 uses education and experience cells. The variable "Cap Year" is equal to one in fiscal years 2004-2009 and zero otherwise. The variable "Lottery Year" is equal to one in fiscal years 2008 and 2009 and zero otherwise. The treated variable is a dummy equal to one for the cells in the for-profit sector. All regressions use skill-cell by sector effects and year effects. The method of estimation is least squares with standard errors clustered at the skill-cell level. *** p<0.01, ** p<0.05, * p<0.1

Figure A1: Native College-Educated Employment Trends Prior to Fiscal Year 2004

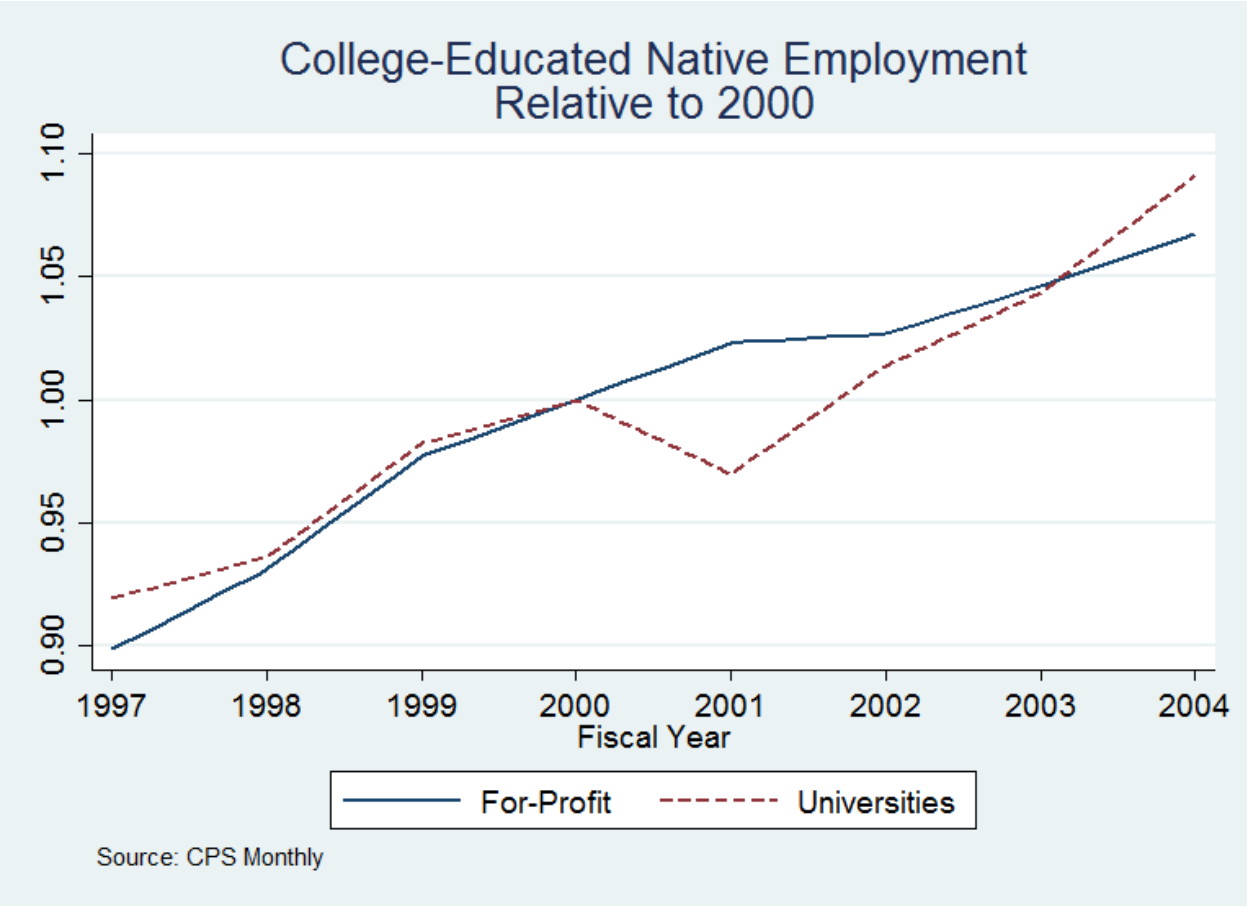

Note: Figure displays For-Profit and University employment trends for college-educated native-born Americans prior to the reduction of the H-1B cap in Fiscal Year 2004. Figures are calculated for months heading into the stated fiscal year and are normalized to equal one in 2000.