

The Effects of High-Skilled Immigration on Firms: Evidence from H-1B Visa Lotteries¹

Kirk Doran
University of Notre Dame

Alexander Gelber
Goldman School of Public Policy, UC Berkeley, and NBER

Adam Isen
Office of Tax Analysis, U.S. Department of the Treasury

March 2015

Abstract

We study the effect of a firm winning an additional H-1B visa on the firm's outcomes, by comparing winning and losing firms in the Fiscal Year 2006 and 2007 H-1B visa lotteries. We match administrative data on the participants in these lotteries to the universe of approved U.S. patents, and to IRS data on the universe of U.S. firms. Winning additional H-1B visas has an insignificant effect on patenting within eight years, with confidence intervals that rule out moderate-sized or larger effects. H-1Bs substantially crowd out employment of other workers. We find some evidence that additional H-1Bs lead to lower average employee wages while raising firm profits.

¹ This is a greatly revised version of NBER Working Paper 20668, previously titled "The Effect of High-Skilled Immigration on Patenting and Employment: Evidence from H-1B Visa Lotteries." We thank U.S. Customs and Immigration Services for help with the H-1B lottery data. We thank Sunil Vidhani for outstanding research assistance. We thank Notre Dame and the Wharton School of the University of Pennsylvania for research support. We are grateful to George Borjas, John Bound, David Card, Sean Farhang, Hilary Hoynes, Larry Katz, Bill Kerr, Jesse Rothstein, and Ankur Patel for helpful comments, and to Lee Fleming for sharing the patent data with us. We thank Danny Yagan for sharing his code to probabilistically identify natives and foreigners in the Treasury data. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the U.S. Treasury Department, or of any other person associated with the U.S. Treasury Department. All errors are our own.

1. Introduction

What are the effects of skilled immigration on the economy receiving the immigrants? This debate has reached a fever pitch in the last several years, with prominent voices from government, the business community, the labor community, and academia discussing major changes to U.S. immigration law. High-skilled immigrants, in particular, represent around a quarter of U.S. workers in occupations closely tied to innovation (Pekkala Kerr, Kerr, and Lincoln forthcoming). Many proposals have envisioned changes to the largest high-skilled immigration program in the U.S.: H-1B visas for temporary immigration.

One common narrative often argues that H-1Bs have exceptional skills that firms cannot otherwise easily obtain, and so giving H-1Bs to a firm could increase the firm's rate of innovation or patenting. Moreover, if H-1Bs have special skills, H-1Bs generally would not replace other workers who otherwise would have worked at the firm—consistent with firms' legal obligation that the employment of H-1Bs “will not adversely affect the working conditions of workers similarly employed.”² In a particularly positive scenario, the firm could even increase employment of other workers. This is exemplified by former Microsoft Chairman Bill Gates' congressional testimony (Gates 2008), arguing that H-1Bs have special, innovative skills and that technology firms hire five additional employees to support each new H-1B worker (based on National Foundation for American Policy 2008).

In a competing, frequently encountered narrative, H-1Bs have more muted effects on firm outcomes like patenting or employment.³ If H-1Bs displace employment of other workers, and the worker displaced from the firm otherwise would have patented or innovated at the firm as much as the H-1B, then we would not expect these outcomes to improve at the firm that received the H-1B. Moreover, many H-1Bs are not in scientific industries, and many H-1B workers perform jobs (*e.g.* technical support) that might be expected not to lead to patenting in the great majority of cases. Economic theory predicts

² Immigration and Nationality Act (INA) §212(n)(1)(A)(ii).

³ We contrast these two competing narratives not because they cover all economically possible combinations of effects of H-1Bs on patenting, employment, profits, wages, and other outcomes, but to contrast two common narratives espoused by the business community, the labor community, policy-makers, the media, or academics.

that firms should apply to hire an H-1B worker as long as this increases the firm's profit in expectation. H-1Bs could increase the firm's profit even if they have no effect on the firm's patenting and/or displace other workers to some extent, as in the case studies described in Matloff (2003) or Hira (2010)—for example, if the H-1B is substitutable with other workers and the firm pays the H-1B less than the other worker who was displaced.⁴ Firms submit legal attestations that they will pay the H-1B a “prevailing wage” comparable to other similar workers, but it is possible that these regulations are ineffective in some cases. Indeed, profit-maximizing firms willingly apply for H-1Bs even though they must pay a fee to apply, underscoring the possibility that the regulations are ineffective.

Our paper addresses these issues by estimating the causal impact of receiving extra H-1B visas on the receiving firm's outcomes, examining outcomes suited to assess the narratives above. We use randomized variation from the Fiscal Year (FY) 2006 and FY2007 H-1B visa lotteries. In these years, when firms submitted H-1B visa applications precisely on the date when U.S. Citizenship and Immigration Services (USCIS) reached the maximum number of H-1B visa applications allowed for a given year and visa type, the applications submitted on these dates were subject to a lottery. Some visa applications were randomly chosen by USCIS to win the lottery, while the remaining visa applications were randomly chosen to lose the lottery. Across both years and across visa lotteries for those with and without advanced degrees, 3,050 firms applied for 7,243 visas, of which 4,180 applications won the lottery. We use administrative data from USCIS on the entrants in these lotteries, matched to U.S. Patent and Trademark Office (USPTO) data on the universe of patents at U.S. firms, and matched to Internal Revenue Service (IRS) microdata on the universe of employment at U.S. firms. In the context of the FY2006 and FY2007 lotteries, our results speak specifically to the effects of marginally increasing the H-1B visa cap on firms' outcomes. Legislation and many commentators have proposed changes in the cap, and our results therefore speak to a scenario that is of great interest to firms and policy-makers.

⁴ Profit could also increase if the H-1B increases the firm's productivity but not its employment of other workers. The simulations of Bound, Braga, Golden, and Khanna (forthcoming) show that the ability to hire foreign computer scientists reduces employment and wages of natives (while at the same time increasing aggregate employment and output).

Across our patenting specifications, which examine the impacts of an additional H1-B visa win on the firm's approved patents over the eight years following the start of the visa, the estimated effects cluster around zero and are never significantly positive. Our confidence intervals rule out moderate-sized or larger effects, and in many cases they are even more precise. This holds true even when we exclude firms that likely provide temporary technical support services, such as Infosys, Wipro, or Tata. Firms that apply on the date of the cap are more likely than firms applying on other dates to be in scientific industries, and are more likely to have patented prior to the year of application, arguably making it more striking that we find little patenting effect in our sample.

Parallel to the patenting results, we robustly find that new H-1Bs cause no significant increase in firm employment. New H-1Bs substantially and statistically significantly crowd out median employment of other workers. More suggestive evidence (based on probabilistic determination of which workers are foreigners) shows that H-1Bs crowd out employment of other foreigners to some extent, and rules out the scenario in which H-1Bs replace natives one-for-one.

Consistent with the presumption that H-1Bs should increase firm profits, we find some degree of evidence that additional H-1B visas increase median profits. We also find some evidence that H-1Bs decrease median payroll per employee, which may be related to the increase in profits.

Our paper is the first we know to isolate the effect of an additional H-1B visa given to a particular firm on outcomes at that firm (holding constant H-1Bs given to other firms).⁵ This is relevant to firms and policy-makers seeking information on the firm-level effects of granting firms additional H-1Bs. We demonstrate that H-1Bs given to a firm on average do not raise the firm's patenting and/or other employment, contrary to firms' frequent claims. Overall our results are more consistent with the second narrative, in which H-1Bs replace other workers to some extent, are paid less than alternative workers, and increase the firm's profits (despite little, if any, effect on firm patenting).

⁵ Kerr and Lincoln (2010) and Pekkala Kerr, Kerr, and Lincoln (forthcoming) examine the effect of giving an additional H-1B to a firm by interacting firm characteristics with the H-1B visa cap, and as such are among the first to examine the role of firms. Changes in the aggregate H-1B cap could affect outcomes at a given firm through general equilibrium effects, including effects of the cap increase on other firms. Thus, this previous work addresses a different question of interest than ours does.

Relative to previous studies on H-1Bs and other immigration programs, ours is also the first to our knowledge to leverage true randomized variation to estimate the effect of immigration on outcomes of the receiving economy,⁶ and ours is one of the first that has used administrative data. Our paper relates to previous work on the effects of immigration on innovation or productivity (*e.g.* Borjas and Doran 2012; Stuen, Mobarak, and Maskus 2012; Foley and Kerr 2013; Grogger and Hanson 2013; Moser, Voena, and Waldinger 2014; see the Kerr 2013 survey). Previous studies specifically of the innovation or labor market impacts of the H-1B program or similar programs include Kerr and Lincoln (2010), Hunt and Gauthier-Loiselle 2010, Hunt (2011), Peri, Shih, and Sparber (2013), and Pekkala Kerr, Kerr and Lincoln (forthcoming). These papers have found that H-1Bs have large positive impacts on innovation and productivity and have found no clear evidence of displacement of other employment.⁷ In preliminary work, Peri, Shih, and Sparber (2014) examine the implications of winners of H-1B visa lotteries, but because they do not have access to the list of lottery losers their paper does not leverage randomized variation (and also does not examine the firm level).⁸ Finally, our work relates to papers on the labor market impacts of immigration not specifically in the H-1B context (*e.g.* Card 1990; Borjas, Freeman, and Katz 1997; Card 2001; Friedberg 2001; Borjas 2003; Edin, Fredriksson, and Åslund 2003; Lubotsky 2007; Borjas,

⁶ Edin, Fredriksson, and Åslund (2003) and Åslund, Edin, Fredriksson, and Grönqvist (2011) use variation that appears quasi-random.

⁷ Kerr and Lincoln (2010) find no evidence that H-1Bs displace other workers. Pekkala Kerr, Kerr, and Lincoln (forthcoming) find mixed evidence on the effect of H-1Bs on total firm size. Peri, Shih, and Sparber (2014) find that H-1Bs increase employment of natives.

⁸ Specifically, Peri, Shih, and Sparber (2014) examine the effects of H-1B visas on local labor markets using the FY2008 and FY2009 H-1B visa lotteries. However, in these years, USCIS did not record the list of lottery losers (personal correspondence with USCIS, 2009). The paper attempts to reconstruct the list of lottery losers by using Department of Labor (DOL) records on Labor Condition Applications (LCA), which must be submitted before firms can submit an H-1B application to USCIS. The identification strategy assumes that conditional on having an LCA application that is approved by DOL, selection for an H-1B is random. However, many approved LCA applications end up *not* being subject to the H-1B lottery. When a firm is no longer interested in hiring the worker for which the firm had previously submitted the approved LCA application, the firm does not submit an H-1B application to USCIS. In FY 2008 and 2009, at least 20 percent of LCA applications are contaminated by these companies that chose to not apply for an H-1B visa (*e.g.* USCIS 2008, DOL 2014). This raises the concern that the analysis of that paper is confounded by demand shocks; for example, firms in areas that experience negative shocks might be less likely to submit H-1B applications to USCIS (conditional on having an approved LCA application), and one would expect that this negative shock would be correlated with subsequent economic outcomes.

Grogger, and Hanson 2012; Cortes and Pan 2014; see surveys in Borjas 1994; Friedberg and Hunt 1995; Freeman 2006; Dustmann *et al.* 2008; and Pekkala Kerr and Kerr 2011).

As our results speak to the impact of additional H-1B visas given to a particular firm on that firm's outcomes, our findings are consistent with the possibility that an aggregate increase in H-1Bs increases firm or aggregate patenting and/or employment, as in previous literature cited above. For example, at the firm level, our results show that the H-1B worker replaces other workers; the displaced workers may find employment elsewhere (unless demand is perfectly inelastic), and they could increase patenting in this other firm relative to the counterfactual (and these increases in patenting could further lead to positive spillovers, as in *e.g.* Bloom, Schankerman, and van Reenen 2013). Our results demonstrate that if H-1Bs do indeed have large positive effects on aggregate patenting or employment, as previous economics literature has found, then this is not occurring because an extra H-1B at a given firm leads to increases in these outcomes at the firm level—in contrast to the first narrative above.

The paper is structured as follows. Section 2 describes the policy environment. Section 3 describes our empirical specification. Section 4 describes the data we use. Section 5 demonstrates the validity of the randomization. Section 6 describes our empirical results on patenting. Section 7 describes our results on employment. Section 8 discusses effects on payroll per employee and profits. Section 9 concludes.

2. Policy Environment

The H-1B visa is the largest program for temporary skilled migration to the United States. H-1Bs are sponsored by firms, which apply to the U.S. government to obtain a visa for each H-1B worker they wish to hire. In their application for each visa, firms must specify the identity of the worker they wish to hire. An H-1B visa allows a skilled foreigner to enter the U.S. for three years. The H-1B may stay at the initial sponsoring firm or move to another firm. The H-1B is considered a “nonimmigrant” visa because it allows those with H-1Bs to stay in the U.S. only temporarily, rather than more permanently. After these three years, the worker may leave the U.S., or a firm may seek to renew the worker's H-1B visa or sponsor the worker to be a permanent resident.

The firm submitting the H-1B LCA to DOL must attest, among other things, that: “(a) H-1B nonimmigrants will be paid at least the actual wage level paid by the employer

to all other individuals with similar experience and qualifications for the specific employment in question or the prevailing wage level for the occupation in the area of employment, whichever is higher”; and “(b) The employment of H-1B non-immigrants does not adversely affect working conditions of workers similarly employed in the area of intended employment.”⁹

We study the lotteries for H-1B visas that were conducted for certain visas granted in FY2006 and FY2007. We study these lotteries because for other years we have sought, USCIS did not keep data on which firms won and lost the lottery (personal communication with USCIS, 06/01/2011). Because USCIS ran this lottery on its own, we are evaluating an existing government program (as opposed to evaluating a randomized experiment designed by researchers). Visas given for FY2006 allowed a worker to work from October 2005 to October 2008, and visas given for FY2007 allowed a worker to work from October 2006 to October 2009.

The total number of H-1B visas awarded to firms in a given year is subject to a maximum number or “cap.” This cap is different for visas given to workers who have only a B.A. (the “Regular” H-1B visa) and for visas given to workers who have a masters degree or higher from a U.S. institution (the “Advanced Degree Exemption” (ADE) H-1B visa). The cap for regular H-1B visas was 65,000 in each year for FY2006 visas and for FY2007 visas, and the cap for ADE H-1B visas was 20,000 in each year for FY2006 visas and for FY2007 visas. In each year and for each of the two types of H-1B visa, visas are allocated by lottery to firms that applied on the date when the total number of applications reached the cap. In a given lottery, firms are allowed to apply for multiple visas. In those cases in which firms applied for multiple visas in a given lottery, the probability that the firm won each visa was independent and equal to the number of lottery winners divided by the number of lottery entrants. The lotteries were conducted by USCIS. In each of these lotteries, the total number of applications that won the lottery was equal to the number of remaining visas necessary to reach the cap.

The cap does not apply to a number of visa categories, which are therefore excluded from the lotteries. Visas given for work at non-profit firms, including U.S.

⁹ Employers who are “H-1B dependent”—whose workforce is comprised of a sufficiently large fraction of H-1B employees—face additional requirements to attempt to recruit, and not displace, U.S. workers.

educational institutions, are not subject to the cap. Citizens of five countries (Australia, Canada, Chile, Mexico, and Singapore) are in effect not subject to H-1B limits. Finally, those applying for extension of an existing visa, or those who have an existing H-1B visa and are changing jobs during the period the visa covers, are not subject to the cap. Our results therefore do not speak to the effects of such un-capped visas, so it is difficult to compare our results to studies that have examined student/trainee or temporary work visas in general (*e.g.* Hunt 2011).

Firms did not know in advance the date when the cap would be reached, and they did not know the probability that firms applying on this date would be selected for an H-1B. For the FY2006 regular visa, the cap was reached on August 10, 2005; for the FY2007 regular visa, the cap was reached on May 26, 2006; for the FY2006 ADE visa, the cap was reached on January 17, 2006; and for the FY2007 visa, the cap was reached on July 26, 2006 (personal correspondence with USCIS, 2011). These dates were determined by the number of applications that were received on different dates in these years, which was unknown to firms at the time—making it effectively impossible for firms to successfully game the system and apply for more visas than they desire.¹⁰ Each of the lotteries was conducted within a month of reaching the cap for that lottery.

Firms pay fees for filing a visa application for initial H-1B status, ranging from total fees of \$1,575 to \$3,550 depending on firm size and whether the firm requests expedited processing. Fees for applications that lost the lottery were refunded to firms.

For a given lottery year (*i.e.* FY2006 or FY2007), we refer to the calendar year that the lottery occurred (*e.g.* 2005 in the case of the FY2006 lottery) as “Year 0.” The year before this calendar year is “Year -1”; the year after Year 0 is “Year 1”; and so on. We refer to the first quarter when an H-1B employee would begin work at a firm (*e.g.* the first quarter of FY2006 in the case of the FY 2006 lottery) as “Q1”; we refer to the next quarter as “Q2”; and so on. A fiscal year begins in October of the previous calendar year; for example, Q1 of FY2006 corresponds to October to December of calendar year 2005.

¹⁰ Such a hypothetical strategy would be hampered by the fact that firms must submit visa applications for specific workers and pay a fee to apply, implying significant costs of applying for each visa. These were also the first two years USCIS used a lottery to allocate H-1B slots, and it was not announced in advance that lotteries were going to be run for FY2006. Executives at firms hiring H-1Bs have indicated to us that they apply for the number of H-1Bs they desire, rather than gaming the system by applying for more.

3. Empirical Strategy

Our empirical strategy exploits the random assignment of H-1B visas in the lotteries. Thus, we consider only the sample of firms that entered the FY2006 or FY2007 H-1B lotteries. Our main outcomes of interest are patenting and number of employees. We also consider the effect on the firm's wage bill per employee and profits.

Our strategy must accommodate firms that applied for multiple H-1B visas. If a firm submits n visa applications to a lottery in which p percent of lottery applications won an H-1B visa, and W is the number of H-1B visas awarded to the firm, the expected number of H-1B visas awarded to the firm is $E[W]=pn$. If the actual number of visas won is w , then the number of "unexpected" wins $u=w-pn$ reflects the random realization of the net number of *unexpected wins* (or losses) and will be orthogonal to the error in the regression we specify below.¹¹ Thus, our main independent variable is the random variable U , the net number of unexpected wins (or losses) (whose realization is u).

To determine the causal effect of an unexpected H-1B visa win on an outcome Y , we run the "reduced form" (*i.e.* intent-to-treat (ITT)) regression:

$$Y_{iT} = \beta_0 + \beta_1 U_{iT} + \varepsilon_{iT} \quad (1)$$

Here t is defined as the number of calendar years since the lottery in question occurred; for example, $t=0$ corresponds to Year 0, *i.e.* 2005 in the case of the FY2006 lottery, or 2006 in the case of the FY2007 lottery. We often run this regression separately for different choices of t . T indexes the year of the lottery in question, *i.e.* FY2006 or FY2007. U_{iT} represents the number of unexpected H-1B visa lottery wins for firm i in the lottery that occurred in year T . ε_{iT} is an error term. This is our primary specification.¹²

After a firm wins an H-1B lottery, its application may be approved, denied, or withdrawn. For example, the application may not have met the eligibility criteria, leading to a denial, or the applicant firm may go out of business, leading to a withdrawal. In some cases it is of interest to examine the effect of an approved H-1B visa on firm outcomes, in addition to examining the effect of an H-1B lottery win. The total number of H-1B visas approved for a firm in any given year is potentially endogenous, because it depends on

¹¹ By "unexpected," we are not necessarily referring to firms' actual expectations, which are unobserved.

¹² This specification makes a linearity assumption: moving from no visa to one has the same effect as moving from one to two, etc. We estimate insignificant coefficients on higher-order terms in visa wins.

the fraction of those that win the lottery that also are approved. We can exploit the lottery to provide an instrument for approved H-1B visas in a two-stage least squares model:

$$A_{iT} = \alpha_0 + \alpha_1 U_{iT} + v_{iT} \quad (2)$$

$$Y_{iT} = \beta_0 + \beta_1 A_{iT} + \eta_{iT} \quad (3)$$

Here A_{iT} represents the number of H-1B visas approved for this firm in the lottery that occurred in year T . In the first stage (2), we regress approved H-1B visas A_{iT} for firm i in lottery T on unexpected wins U_{iT} . Thus, the first stage is the same no matter what the time period t when the outcome is observed. In the second stage (3), we regress Y_{iT} on A_{iT} (instrumented using U_{iT}). We interpret the coefficient β_1 as a local average treatment effect of an extra approved H-1B visa among the compliers (*i.e.* those induced by winning the lottery to have an extra approved H-1B visa). v_{it} and η_{it} are error terms.

The ITT and IV estimates represent different empirical objects, both of which are of interest. The ITT estimates show the effects of granting another visa to a given firm. This is practically relevant, in the sense that firms and policy-makers are interested in the effects on firms of allowing a marginal capped visa to the firm. For example, policy-makers have often considered the effects of marginally expanding the number of capped visas. Thus, for both patenting and employment we show our main “reduced form” regression (1). In addition, the IV estimates are particularly relevant when we are testing the hypothesis that new H-1Bs crowd out other employment, because in this context we are interested in comparing the coefficient on approved H-1Bs to a specific level, namely to the coefficient in no-crowdout scenario (*i.e.* a coefficient of 1). Thus, for employment we additionally show IV specifications. (The NBER Working Paper version, Doran, Gelber, and Isen (2014), shows the IV estimates of the effect of approved H-1B visas on patenting.) In practice, the first stage regressions corresponding to equation (2) that we show later are extremely strong, with first stage coefficients near 1 (ranging from 0.86 to 0.88), and with first stage F-statistics ranging from 239.94 to 993.51. Thus, in practice there is generally little difference between the OLS coefficient on unexpected lottery wins and the IV coefficient on approved H-1B visas.

In those cases in which a firm participates in more than one lottery in a given fiscal year T (*e.g.* a firm participates in both the 2006 regular and ADE lotteries), we calculate U_{iT} by summing the total number of unexpected wins across both of the lotteries

that the firm enters in year T (except for specifications in which we run separate regressions for the Regular and ADE lotteries).¹³ We seek as much statistical power as possible, and so we pool the FY2006 and FY2007 ADE and regular lotteries in our baseline. In these pooled regressions, for a given firm, we stack data corresponding to the FY2006 lottery and data corresponding to the FY2007 lottery, so that we can capture the effects of winning the lottery in Year 0 on employment in each subsequent year (measured consistently as number of years since the lottery in question occurred). We cluster our standard errors at the level of the firm.

Although the randomization implies that U_i should be orthogonal to the error in (1), it is also possible to control for various pre-determined covariates. For example, we can control for a lagged value of an outcome variable at the firm (*e.g.* in the case in which the dependent variable is the number of patents, we can control for $Y_{i,-3 \text{ to } -1,T}$, the number of patents in firm i observed from Year -3 to Year -1, where year is measured relative to lottery T); for the expected number of lottery wins pn ; or other covariates.

Previous literature has not examined the level of patenting due to the volatility of this variable; instead, it has examined transformations of the number of patents that reduce volatility. Given the approximate lognormality of patents, one may wish to run a specification in which log patents forms the dependent variable (as in, for example, Kerr and Lincoln 2010). However, in our context, estimating exactly this specification would lead to a problem: we would like to include firms in the regressions that have zero patents, as a large fraction of firms have zero patents in our context, but the log of zero is undefined.¹⁴ Thus, we approximate the log of the number of patents using the inverse

¹³ In both the case of FY2006 and FY2007 visas, the Regular visa lottery chronologically occurred on a date before the ADE cap was reached. We verified that winning a slot in one lottery does not affect the probability of applying for subsequent H-1B visas. We also verified that unexpected wins in earlier lotteries have no significant effect on the probability of applying for or obtaining subsequent H-1B visas. For example, when we pool FY2006 and FY2007 and regress total ADE H-1B visa approvals in a given year on unexpected lottery wins in the Regular lottery in that year, we find a coefficient on unexpected lottery wins of -0.20, with a standard error of 0.18 (insignificant at conventional levels, $p=0.26$). We also find that unexpected lottery wins in 2006 have no effect on approved 2007 visas; for example, when regress total FY2007 Regular and ADE approvals (summed) on unexpected lottery wins in the FY2006 Regular and ADE lotteries combined, we obtain a coefficient on unexpected lottery wins of -0.05, with a standard error of 1.45 ($p=0.97$). Finally, we verified that winning one lottery also does not affect the probability of winning a subsequent lottery conditional on entering the subsequent lottery.

¹⁴ This is not a problem in the context of Kerr and Lincoln (2010); they examine patents at the city level, where patents are greater than zero.

hyperbolic sine (IHS) of the number of patents. The IHS approximates the log function but is defined at zero and negative values (see related work in Burbidge, Magee, and Robb 1988, Pence 2006, or Gelber 2011). The IHS of patents Y is defined as:

$$\sinh^{-1}(Y) = \ln(Y + \sqrt{1 + Y^2})$$

In the specifications in which the IHS of patents is the dependent variable, the coefficient on H-1B visas reflects the approximate percent increase in patents caused by an extra unexpected H-1B visa.

A binary outcome, specifically a dummy for whether the firm patented, is also less volatile than the level of patenting. When this dummy is the outcome, we control for a dummy for whether the firm patented between Year -3 and Year -1. When we investigate binary outcomes in our panel data, we run a linear probability model to avoid an incidental parameters problem.¹⁵

A third way of ensuring that we are examining a sample where the lottery variation is substantial relative to the variance of the error term is to investigate the effects in smaller firms, where the impact of one additional employee is likely to be most clearly distinguishable from the baseline in a statistical sense. To evaluate how the effects vary across firms of different sizes, we investigate the effect in the sample of firms with 10 or fewer employees in Year -1 (which represents roughly the 25th percentile of firm size in our sample); in those with 30 or fewer employees in Year -1 (which represents roughly the 50th percentile); and in the sample of firms of all sizes (as well as a variety of other firm size cutoffs).

In the case of the employment outcome, we run a related set of specifications across all of these firm size categories. As in the patenting context, previous literature has not examined the level of employment as a dependent variable, but has instead examined transformations employment, such as the log, that reduce volatility (*e.g.* Pekkala Kerr, Kerr, and Lincoln forthcoming). As we show, the employment outcome is much more volatile (*i.e.* has a much larger standard deviation) than the patenting outcomes we investigate. As a result, noise in the dependent variable is an especially important issue in

¹⁵ We would run into an incidental parameters problem with logits or probits in the case of binary outcomes, or with negative binomial or Poisson regressions in the case of count outcomes.

the employment context. Our main way of addressing this is by running median regressions in our baseline specification in the employment context. (The median value of patents is zero, so it does not make sense to run median regressions in this context.) In these median regressions, we are unable to run quantile instrumental variables regressions because of a practical consideration: they typically did not converge. Instead we run “reduced form” median regressions corresponding to model (1) above.

Our second method of addressing noise and reducing the influence of outliers in the employment variable involves running a two-stage least squares regression as in (2)-(3), where the dependent variable is the winsorized first difference of employment. The first difference ΔY_{it} is taken from before the lottery, in Year -1 (*i.e.* the first quarter of 2005 for FY2006 visa applicants and the first quarter of 2006 for FY2007 visa applicants), to time period t after the lottery. Winsorization is common in administrative data (*e.g.* Chetty *et al.* 2011) and in survey data (*e.g.* the topcoding in the Current Population Survey).¹⁶ Of course, winsorized regressions would not capture large effects on employment outcomes. However, in practice when we run our IV regressions without winsorizing, the point estimate of the effect is negative and insignificant, which lessens the concern that winsorization dulls an actual positive effect. We also find that an extra H-1B visa has an insignificant effect on the probability that the change in employment is outside the 95th percentile. Nonetheless, because of these issues, the quantile regressions serve as our primary specification in the employment context.

A third way of addressing noise in the employment variable is to estimate the effect on the (first-differenced) IHS of employment. In the case of this IHS specification, before testing whether the coefficient on unexpected H-1B visas is equal to 1 (reflecting a

¹⁶ Of course, if we did not winsorize, estimating the effect of unexpected H-1B visas on the first difference of employment while additionally controlling for Year -1 employment (as we often do) is equivalent to simply controlling for Year -1 employment with the Year t level (rather than first difference) of employment as the dependent variable, since the coefficient on Year -1 employment mechanically changes by exactly 1 from the specification with the Year -1 control to the specification without. However, given that we do winsorize the dependent variable, our regressions give different results than those obtained from controlling for Year -1 with the year t level of employment as the dependent variable. We winsorize the first difference of employment and control for lagged employment, rather than winsorizing the level of employment in period t after the lottery and controlling for lagged employment, again because in the context of examining firms of all sizes, winsorizing the first difference is more effective in removing large outliers than is winsorizing the level of employment. When we limit the sample to smaller firms, the two specifications show very similar results, with similar point estimates and confidence intervals.

scenario with no crowdout), we transform the coefficient from the regression (which reflects the approximate percentage increase in employment, rather than the increase in the absolute level of employment) by multiplying it by the mean level of employment. We then test whether this transformed coefficient is equal to 1. We apply the coefficient to the mean level of employment because it is illustrative, but this strategy is subject to the limitation that the coefficient could also be applied to other employment levels. Thus, we present the IHS employment results in the Appendix, rather than in the main tables. (In the patenting context, our interest is instead in the mean effect of H-1Bs on patents, as opposed to testing whether this effect is different than a fixed specific number—as in the employment context, where we test for a difference from unity.)

We verify that when we run exactly parallel specifications in the employment and patenting contexts, we obtain comparable results.

Importantly, our measure of total employment *includes* the H-1B worker if the H-1B worker is at the firm; thus, if the H-1B worker works at the firm, then the effect of an additional H-1B visa on total employment will mechanically be equal to *one plus* the effect on employment of individuals *other* than the new H-1B. One test of interest is a two-sided test of whether the coefficient on unexpected H-1B visas is significantly different from 0. If a coefficient were positive and significant, it would indicate that the extra H-1B visa lottery win increases total employment at the firm—as opposed to simply replacing a worker that the firm would have otherwise hired, in which case the coefficient would be 0. In principle, an extra H-1B visa could even decrease employment at the firm, for example if the new H-1B worker works more hours or works harder than others (for example, to secure another visa for continued employment at the firm, or for another reason) and therefore replaces more than one other worker. Another question of interest is a two-sided test of whether the coefficient on unexpected H-1B visas is significantly different from 1. If the coefficient were greater than 1, this would indicate that an additional H-1B visa leads to employing a greater number of other workers. If the coefficient is less than one, this can indicate that an extra H-1B worker at least partially crowds out other worker(s) who would otherwise have worked at the same firm.

4. Data

Match between USCIS Data and patenting data

We merge several administrative datasets. First, we use USCIS administrative data on the H-1B lotteries for FY2006 and FY2007. The data contain information on each H-1B visa application that entered the lottery in each of these years, for both regular and ADE H-1Bs. These data contain information on Employer Identification Number (EIN); the exact date the firm applied for a visa; the type of H-1B (regular or ADE); the name of the firm that applied for the H-1B; whether the H-1B application won or lost the lottery; and whether or not the H-1B application was approved by USCIS.

We obtained data on U.S. patents from the Patent Dataverse from 1975 to 2013.¹⁷ This database contains data on the universe of U.S. patents granted within in these years, based on USPTO data. We use data from the Patent Dataverse on firm name and the number of patents granted. (The Patent Dataverse does not contain data on firm EINs.) Granted patents are classified by the calendar year when a firm applied for the patent. For example, our measure of the number of patents at a given firm in Year 0 reflects the number of patents the firm applied for in Year 0 that were approved by 2013.

The time to develop a patent can range from months to years, with substantial variance. The mean approval time reported by USPTO for patents filed in FY2008 is 32.2 months (USPTO 2012), although there is again substantial variance. Our data will allow us to estimate the effect on an important set of patents, namely those within eight years of the initial H-1B visa period, but the effect on subsequent patents is unobserved.¹⁸

Since the Patent Dataverse does not contain EIN, but does contain firm name, we matched data from the Patent Dataverse to the USCIS lottery data using firm names. As we describe in greater detail in Appendix 1, to match firms between these two datasets, we performed an intentionally liberal automatic matching procedure between these datasets to obtain all plausible matches between companies and patents. We then searched through the matches by hand to detect and remove all matches that appeared spurious. We classified firms into three categories: (1) 392 firms that definitely matched

¹⁷ We thank Lee Fleming for sharing these data with us. These data build upon the Harvard Business School Patent Dataverse, which contains data from only 1975 to 2010, by updating the sample to 2013.

¹⁸ The majority of H-1B petitions are for workers aged 25 to 34, whereas noted innovations peak around age 40 (*e.g.* Jones 2010), raising the possibility that some H-1B workers who stay will innovate more beyond our sample period. However, Jones (2010) finds that innovation in the 25-to-34 age range is well over half of its level at its peak. We leave examination of effects at longer time horizons to future research.

between the datasets; (2) 63 firms for which it was ambiguous whether they matched; and (3) the remaining 2,595 firms that definitely did not match. In our main results, we exclude the 63 possible matches from the list of matched companies. In the Appendix, we show that the results are robust to assuming that the possible matches were in fact true matches. In general, our results are robust to similar alternative matching procedures. A firm would not match between the datasets if it did not patent during this time period; thus, the non-matching firms appear in our data as having zero patents.

Match between USCIS data and IRS data

Using firms' EIN, we also merged the USCIS lottery data to IRS data on the universe of U.S. firms. These IRS data contain information on overall firm employment for each EIN, among other outcomes. Employment at a firm in a given quarter is taken from IRS form 941. Our measure of employment in Q1 (which reflects the first quarter of the fiscal year, *i.e.* the last quarter of the preceding calendar year) reflects employment as measured in mid-December of that quarter.¹⁹ Thus, between the time when a firm learned that it won or lost the lottery in June to August of Year -1, and the end of Q1, when workers generally begin working at the firm and after which employment is measured, firms had a number of months to react by potentially hiring other worker(s). For example, firms were notified of the FY2007 regular visa lottery results in June of 2006, which gave firms over six months until the last month of the first quarter of FY2007, in December of calendar year 2006. However, in the sole case of the FY2006 ADE lottery, the lottery was held on January 17, 2006, *after* Q1 of FY2006 ended. Thus, in the employment regressions, we drop data corresponding to Q1 of the FY2006 ADE lottery, since firms' hiring decisions in Q1 could not have been influenced by the results of the lottery.

We use data from 2004 to 2013. In the IRS data, the first data available from form 941 are in the first quarter of calendar year 2004. The form 941 data are missing the second through fourth quarters of calendar year 2004, and thus we measure employment in calendar year 2004 using the data on the first quarter of calendar year 2004.

We drop the 2.0 percent of the firms in the USCIS data that did not match to the EIN master list in the IRS data. Pooling over all quarters, 4.5 percent of the remaining

¹⁹ See <http://www.irs.gov/pub/irs-pdf/f941.pdf> (accessed October 16, 2014).

firms in the USCIS data did not match to the quarterly firm employment in the IRS data; we treat this data as missing. We make additional restrictions in the employment data: of the remaining firms, 17.9 percent have missing employment data in Year -1, which makes it impossible to run our specifications (in which we control for Year -1 employment), and we drop these data for the purposes of the employment specifications. Of the remaining observations, pooling over Q1-Q4, 2.2 percent are missing in a given quarter. We verify in Appendix Tables 8 and 9 that going out of business (conditional on the other restrictions) is unrelated to unexpected lottery wins, and we verify in Table 2 that the other sample restrictions are also unrelated to this exogenous variation in H-1Bs.

The USCIS data do not contain identifying information (*e.g.* Tax Identification Numbers) on individual H-1B applications that can be linked to the IRS data.²⁰ Thus, we observe overall employment, but we are not able to distinguish the employment of a particular H-1B worker whose application entered the lottery from employment of others. The data also do not allow us to distinguish H-1Bs in general (whether lottery winners or other H-1Bs) from non-H-1Bs; thus, the employment effects we estimate may include effects on employment of H-1Bs, including H-1Bs other than the lottery winners.

As a further analysis, we also investigate the effect on the firm's yearly net income (its "profit") and wage bill (per employee), both as reported to the IRS.

Summary statistics

Table 1 shows summary statistics. We use data on 3,050 firms (where "firm" refers to an EIN). The mean (37.74) and standard deviation (390.95) of patents are very large, primarily due to a small number of firms that patent in large numbers. The mean (0.33) and standard deviation (1.28) of the IHS of the number of patents are much lower. The means and standard deviations are smaller among the 1,276 firms with 30 or fewer employees, and smaller still among the 749 firms with 10 or fewer employees. As a

²⁰ We were given the lottery data to link firms, not workers. The LCA information on the salary intended for a worker cannot usefully be used to link USCIS applications to the IRS data, as there is significant measurement error: (a) the employer could pay the employee more than the stated amount on the LCA, *e.g.* due to overtime; or (b) the employer could pay the employee less than the stated amount on the LCA, *e.g.* because the employee arrives at the firm at a later date than stated on the application or because of fraud. A link would be further complicated because multiple employees at the firm could be paid the same amount, *e.g.* under a prevailing wage. Finally, identification of the H-1B would be hampered because the H-1B need not be a new employee of the firm if the firm previously employed the H-1B under a different visa.

result, in many of our results we focus on smaller firms, in which we might *a priori* expect that an extra H-1B might have a more noticeable effect on the outcomes. Only a modest fraction of the sample patents (*e.g.* 9.3 percent in the full sample of firms).

The mean (1,877.84) and standard deviation (39,721.31) of the number of employees over Q1-Q4 in the full sample are very large. In firms with 30 or fewer, or 10 or fewer, employees in Year -1, the mean and standard deviation of Q1-Q4 employment are much lower but still quite large. Median employment is much lower than the mean. Winsorizing also reduces the mean and standard deviation greatly.

In the FY2006 regular lottery the vast majority of applications lost the lottery, and in the FY2007 regular lottery the vast majority of applications won the lottery, whereas the ADE lotteries have a more even fraction of winners and losers. The fact that the vast majority either won or lost the regular lotteries will not directly pose an issue for us: as long as we estimate the standard errors correctly, the estimates will show whether we estimate precise results. Hypothetically excluding data on uneven lotteries should lead to a loss of statistical efficiency. Other estimates in randomized contexts have also relied on uneven fractions of wins and losses (*e.g.* Imbens, Rubin, and Sacerdote 2001).

The sample contains 7,243 visa applications, with an average of 2.37 H-1B applications per firm summing over both years. The average firm in our sample won 0.57 H-1B visas when aggregating across both years. The standard deviation of the number of unexpected lottery wins (defined above) is 0.33, and its range runs from -2.65 to 2.96.

Comparison of lottery firms to other firms

As our regressions will only investigate the effect on firms that applied on the day the cap was reached and therefore are subject to the lottery, it is relevant to compare this sample to the broader sample of firms. Table 2 shows regressions where we regress characteristics of the firms on a dummy for applying on the last day (*i.e.* a dummy for being subject to the lottery) and lottery fixed effects.²¹ Firms applying on the last day are *more* likely to have patented in the past, and patented more in the past. Similarly, firms applying on the last day are quite a bit (17 percentage points) more likely to be in

²¹ In our context, we pool data across four different lotteries. It is not informative to compare summary statistics (*e.g.* means) of variables of interest between firms that applied on the last day and other firms, because the number of firms applying in each year and visa type could be correlated with the outcomes in question, confounding such a comparison of means if we pooled data from all four lotteries together.

scientific industries (NAICS=54). If H-1Bs hypothetically have bigger positive patenting effects in firms that patented more in the past and/or are in scientific industries, then our sample will arguably be primed to find a particularly positive effect on patenting.

Applications on the last day are 22 percentage points more likely to be for occupations in “systems analysis and programming,” and they tend to be from larger firms.

5. Validity of Randomization

Table 3 verifies the validity of the randomized design by regressing pre-determined variables that could not be affected by the lottery on unexpected lottery wins. The table confirms that none of the pre-determined variables is significantly related to unexpected lottery wins: the lagged dependent variables (various measures of patenting, employment, wage bill per employee, and profits); dummies for whether firms from the USCIS lottery data match to other datasets; and a dummy for whether the firm has North American Industry Classification System (NAICS) code 54—representing professional, scientific, and technical services, which comprises 56.43 percent of the sample.

6. Patenting Results

We estimate the effect of an unexpected H-1B visa win on patenting outcomes in Table 4. We focus on the effect on the IHS of patents as our baseline. We also estimate the effect on a dummy for whether the firm patented.

We investigate each of these outcomes separately over Years 0 to 7 (inclusive). For each of our outcomes, we show the results with two alternative sets of controls: (a) controlling for the number of patents from Year -3 to Year -1; and (b) controlling for the number of patents from Year -3 to Year -1, as well as the expected number of lottery wins (conditional on the number of H-1B applications and the probability of winning the lottery in question). We take specification (b) as our baseline, though the results are similar either way. The results are nearly identical when we control for additional or alternative controls, such as controlling additionally for the two-digit NAICS code of the firm, controlling for the firm’s number of H-1B lottery applications n , and/or controlling for dummies for each of the four lotteries considered. The results are similar when pre-period patenting is measured over other time periods rather than Year -3 to -1.

In Table 4, row A shows the results for firms with 10 or fewer employees. We estimate precise, insignificant effects in all specifications. The upper end of the 95

percent confidence interval rules out more than a moderate effect. When the dependent variable is the IHS of the number of patents, the upper end of the 95 percent confidence interval enables us able to rule out an increase in patents of more than 1.8 percent, relative to a “base” mean number of patents of 0.027 per year.²² When the dependent variable is the dummy for whether the firm patented, the upper end of the 95 percent confidence interval in the baseline is 0.0087, indicating that an extra H-1B lottery win does not raise the probability of patenting over the full period (*i.e.* Years 0 to 7) by more than 0.87 percentage point. All of these results are similar when controlling only for prior patents. While the point estimates are sometimes below zero, we do not conclude that extra H-1B wins actually decrease patenting, as the 95 percent confidence intervals can never rule out a decrease of zero at any standard significance level; of course, this is why confidence intervals are useful in determining what we can rule out with a standard degree of statistical certainty.²³

Row B shows the results for firms with 30 or fewer employees. These results also show small coefficients with narrow confidence intervals, although the confidence intervals are somewhat wider than in row A (which is unsurprising given the much larger standard deviation of patents among these firms). In the baseline, we can bound the increase in patents below 3.0 percent, relative to a yearly mean number of patents of 0.27. Row C shows the results for firms of all sizes. In the baseline, the upper end of the 95 percent confidence interval rules out an increase greater than 1.1 percent. Since the yearly mean of patents is 4.87, this implies a maximum yearly increase in patents of 0.05. When the dependent variable is the probability of patenting, the upper end of the 95 percent confidence interval rules out an increase greater than 2.5 percentage points.

Our choices of the number of employees in our size thresholds (*i.e.* 10 or fewer, or 30 or fewer) could be varied. Figure 1 plots the coefficient and confidence interval on unexpected H-1B visas when the dependent variable is the IHS of number of patents over Years 0 to 7, as a function of the size of the employer. We show the results for employers

²² We calculate the “base” mean by taking the mean number of yearly patents in Years 0 to 7 in a “control group,” specifically firms whose number of unexpected wins was less than or equal to zero.

²³ Throughout Table 4 and our other patenting and employment regressions, our results also show no evidence of a positive impact on patenting when we weight firms by their number of H-1B applications, by the number of applications on the date of the cap, or by the expected number of lottery wins.

of each size from under 10 to under 500, in increments of 10.²⁴ Our main results of interest relate to the upper end of the 95 percent confidence interval, which ranges from near 0 to around 0.05; across all 50 choices of the employer size threshold shown, in the most positive case we are able to rule out an increase in patents more than around 5 percent (and typically the upper bound is substantially smaller). The point estimate is positive in only one out of 50 cases (though it is insignificant in this case). While several of the estimates are negative and barely significant at the 5 percent level, this is not a robust finding as the estimates are typically insignificant at the 10% (or 5%) level.²⁵ Overall, we robustly rule out more than a modest percentage increase in patenting.

Appendix Table 2 shows that the effect on the number of patents is small relative to the baseline variation; for example, in employers with 10 or fewer employees, the upper end of the 95 percent confidence interval rules out an increase in the number of patents greater than 1.8 percent of a standard deviation. To investigate the effects over different time horizons, natural grouping of calendar years is to separate them into (1) years covered by the initial H-1B visa (Years 0 to 3) and (2) subsequent years (Years 4 to 7).²⁶ If developing a patent often takes a few years, we might expect effects to be larger in the later period. Appendix Table 3 shows that we estimate comparable results when we limit the period over which we observe the outcome to Years 0 to 3, or to Years 4 to 7.²⁷ Appendix Table 4 shows that the results are similar when we assume that those companies that possibly matched between the USCIS and patenting database in fact did actually match, rather than assuming that they did not match as in our baseline.²⁸

²⁴ The necessity of keeping a sufficiently large number of firms in each category, to prevent the potential identification of any given firm, prevents us from going beyond 500 employees in increments of 10.

²⁵ When we investigate the patenting dummy, the results are comparable to those shown across the entire set of firm sizes from 0 to 500.

²⁶ Because patents are measured in each calendar year, whereas H-1B visas cover fiscal years, Year 0 refers to a calendar year when the H-1B worker worked at the firm for typically only one-quarter of the calendar year (*i.e.* October to December, corresponding to the first quarter of the fiscal year). Similarly, three-quarters of calendar Year 3 (*i.e.* January to September of calendar Year 3) is covered by the H-1B visa.

²⁷ The most recent firm-level patent citation data we were able to access were through 2010, substantially limiting the period since the visa lotteries. Examining the effect on patent cites just through 2010 likewise shows no significant impact, with small confidence intervals (results available upon request).

²⁸ When we examine the sample of all firms in Appendix Table 4, the estimated effect for the full period is negative and significant at the 10 percent level when we include the full set of controls, although we do not consider this a robust finding: (a) it is not significant at more conventional levels (*i.e.* 5% or 1%); (b) it is not robust to other specifications such as removing the “possible” matches; and (c) it is not matched by any significant estimate when we investigate other dependent variables, including the patenting dummy.

Heterogeneity

Table 5 examines whether there is heterogeneity in the effect on patents across type of lottery or type of industry, using our baseline specification and examining effects on patents in Years 0 to 7 combined. Row A examines the Regular H-1B lottery. The results are comparable to those in the full sample—with point estimates that cluster near zero, and the upper end of the 95% confidence interval ruling out large effects—which should not be surprising since 85.96% of the full sample participates in the Regular lottery. Row B examines the ADE lottery. The point estimates are all negative and insignificantly different from zero. Among firms with fewer than 10 or fewer than 30 employees, the confidence intervals rule out large effects. However, the upper end of the 95% confidence interval is larger than in the case of the Regular lottery—consistent with the loss of statistical power due to the much smaller sample in the ADE lotteries.²⁹

The effect on patenting is particularly relevant in professional, scientific, and technical services (NAICS code 54), since the bulk of patents occur in this industry. We find no evidence of an effect on patenting in this group, with negative point estimates and confidence intervals that are in the same range as those in Table 4. In firms outside NAICS code 54, the point estimates are more positive, though insignificant.

Many H-1Bs are given for workers in firms like Infosys, Tata or Wipro that primarily offer outsourcing for temporary support services (often temporary technical support services). By contrast, other H-1Bs are given to companies like Intel or Google that do not specialize in such services. While it is not possible to determine with certainty which visas fall in the broadly-defined “temporary support services” category, it is illuminating to investigate the effects in firms that likely specialize in such services. To probabilistically identify such firms, we first compiled a list of those firms among the largest 100 H-1B sponsors and that had “outsourcing services” or “IT support services” in the description of the company on its website. We found that these firms were in only six, six-digit NAICS categories.³⁰ We then ran our regressions only in firms in these

²⁹ When we investigate the effect separately in each year of the lottery (*i.e.* separating the FY2006 lotteries from the FY2007 lotteries), or separately in each of the four lotteries (FY2006 Regular, FY2006 ADE, FY2007 Regular, and FY2007 ADE), we again estimate insignificant effects in each year separately, with comparable point estimates to those in the full sample, though again with larger confidence intervals.

³⁰ The NAICS codes are 541511, 541519, 541600, 541330, 519100, 423600, and 541512.

industries, and separately ran the regressions only among firms in other industries. The point estimates and top end of the 95 percent confidence intervals are smaller in “temporary support services” industries, although the estimates are insignificant in both sets of industries (and insignificantly different across the two different samples).

Table 6 shows interactions of unexpected H-1B visas with continuous covariates. In principle, it is possible that the H-1B visa could tend to have more (or less) positive effects on firms that apply earlier for the visas. For example, the visas have the largest positive effects in such firms, motivating their earlier applications. In Table 6 Panel A we interact the number of unexpected H-1B visas with the number of days taken to reach the cap in each lottery (which ranges across the four lotteries from 55 to 291). We find no significant interaction in Column 1: the point estimate is nearly zero, and the 95 percent confidence interval rules out a substantial interaction. However, this evidence is merely suggestive: heterogeneity across the lotteries in the effect of H-1Bs visas that happens to be correlated with the time taken to reach the cap would confound our estimate of the interaction. In Panel B Column 1 we show that the interaction of the IHS of prior patents (Years -3 to -1) with unexpected visa lottery wins is positive but insignificant. The interaction of unexpected visas with prior firm size is also insignificant.

7. Effect on employment

Table 7 shows estimates of the effect of extra H-1B visa lottery wins on total firm employment. We show median regressions where the dependent variable is total employment (model (1) above), as well as IV regressions where the dependent variable is the first difference in employment winsorized at the 95th percentile (model (2)-(3)). In our baseline, we show the regressions for employment from Q1 to Q4. It is most straightforward to compare the coefficient on unexpected wins to 1 in this time period, when the worker is almost always working at the firm.

We begin with the median regressions in firms with 10 or fewer employees. In the baseline specification with the more extensive set of controls, the top end of the 95% confidence interval is 0.11, indicating that an extra H-1B visa win leads to an increase in total employment of at most 0.11 workers. While the point estimate is below zero, we do not conclude that extra H-1B visa lottery wins actually decrease employment, because our confidence interval is compatible with an increase. Similarly, in this specification in

firms with 30 or fewer employees, the top end of the confidence interval is 0.37. In the full sample of firms, we can rule out an increase greater than 0.57.

Turning to the two-stage least squares regressions, we again find that the point estimates are under one and mostly find that the coefficient is significantly different from 1, at the 1% level. In the baseline specification among firms with 10 or fewer employees, the top end of the 95% confidence interval is 0.68, indicating that we cannot rule out a moderate positive effect. With 30 or fewer employees, we can rule out a coefficient of 0.71 or greater. In the full sample of firms, the results are extremely imprecise, and we are unable to rule out a coefficient above 1.

Figure 2 plots the coefficient and confidence interval on unexpected lottery wins when we run median regressions in the baseline specification (i.e. from Q1 to Q4) and the dependent variable is the number of employees in the firm, as a function of the size of the employer. The point estimates are always negative and insignificantly different from zero. We focus on the upper end of the 95 percent confidence interval; across all 50 choices of the employer size threshold, in the most positive case we are able to rule out an increase in total employees of more than 0.6. In all cases, the estimate is significantly less than 1 at the 1% level. Thus, we are robustly able to rule out an increase in employment due to extra H-1B visa wins that is substantially less than one-for-one.

Appendix Table 5 shows that in each quarter from Q1 to Q4 individually, we are typically able to rule out a coefficient of 1, particularly in smaller firms. Appendix Table 6 shows that a number of other specifications yield comparable results: winsorizing at the 99th percentile, rather than at the 95th percentile as in Table 7; letting the dependent variable be the IHS of the first difference in employment (as we run IHS specifications in the patenting context); winsorizing the IHS of the first difference in employment at the 99th percentile (to address occasional outliers that appear even in the IHS); winsorizing the IHS of the level of Q1-Q4 employment at the 99th percentile; and running median regressions when the dependent variable is the first difference of employment and we include no controls.

The rationale for the discrepancy between the specifications run in the patenting context and those run in the employment context is described in our Empirical Specifications section, but it is worth additionally describing further results when we run

other parallel specifications in both contexts: In the patenting context, we omit the median patenting regressions because the median number of patents is zero, but regressions at higher quantiles in the patenting context continue to yield no positive impacts. When the first-difference (or level) of the number of patents (or the IHS of patents) is the dependent variable and we winsorize at the 95th (or 99th) percentile, parallel to those in the employment context, our results are very similar to those shown in Appendix Table 2 but are more precise and allow us to bound the maximum increase in patenting at a still lower level. For example, when we winsorize the first difference of the level of patenting from Years 0 to 7 at the 95th percentile in the baseline specification in the sample of firms of all sizes, the 95 percent confidence interval runs from -0.092 to only 0.037. When we run the two-stage least squares employment regressions but do not winsorize the dependent variable, the results are extremely imprecise among firms of all sizes or among firms with 30 or fewer employees in Year -1, which is unsurprising given the very large standard deviation of employment and large outliers. However, when we do not winsorize and run this specification among firms with 10 or fewer employees in Year -1, the top end of the 95 percent confidence interval is 0.31, and we are able to rule out a coefficient of 1 ($p=0.015$). In sum, running parallel specifications in the two contexts does not change any of our conclusions, except that our results are unsurprisingly imprecise when we examine employment in a two-stage least squares regression and do not winsorize.

We separately estimate the effect on employment of foreigners and non-foreigners in Appendix Table 7 and describe our data and results in detail in Appendix 2.³¹ Foreigners constitute a majority of the workforce in the average firm in our sample, and we find that H-1Bs displace employment of other foreigners at least to some extent. The point estimates suggest essentially no crowdout of U.S. natives, and also suggest essentially no crowdout of U.S. citizens; the confidence intervals rule out one-for-one crowdout of U.S. natives/citizens, but are still consistent with a moderate effect on employment of U.S. natives/citizens. One caveat is that our two measures of the number

³¹ Ottaviano and Peri (2012) raise the possibility that new immigrants are more substitutable with existing immigrants than with natives.

of foreigners and non-foreigners are both imperfect—though the concordance of our results across two separate measures increases our confidence in the results.

Other time periods

Table 8 shows the results at other aggregations of time. Rows A and B show Q5 to Q8, and Q9 to Q12, respectively—*i.e.* each of the remaining two of the three years covered by the H-1B visa in question (beyond Q1 to Q4). The table shows that we are generally able to rule out a coefficient of 1 at the 5 percent significance level in these years. Nearly all of the point estimates are negative, and they are never significantly positive. Row D shows results for Q13 through Q32 (the latest quarter in the sample) pooled, when we estimate less significant results. We conclude that in all years of data available, the preponderance of evidence shows that H-1Bs displace other workers, though our findings unsurprisingly become less precise in later years (and the interpretation as crowding out other workers becomes weaker as the fraction of H-1Bs no longer at the initial firm grows).³²

Heterogeneity across samples

Table 9 investigates whether there is heterogeneity in the employment results across samples, using our baseline employment specification in Q1 to Q4 with median regressions and the more extensive set of controls. (Other specifications show similar results.) The point estimates are more negative for the Regular lotteries than for the ADE lotteries, and they are more negative for scientific services (*i.e.* NAICS code 54) than for other industries. In fact, the point estimates are often positive and substantial in the case of the ADE lotteries, and in the case of non-scientific services—particularly when we examine firms of all sizes. The point estimates are negative in likely “temporary support services” employers but positive in other six-digit industries (though the estimates are insignificantly different across the two samples), and among “temporary support services” the coefficient estimate can be distinguished from unity in more firm size categories than in other industries. However, there are no significant differences across the different samples, including when we compare the 2006 and 2007 lotteries.

³² In Q13 to Q32, the H-1B worker has typically left the firm, so the test of a difference in the coefficient from 1 does not indicate displacement of other workers.

In Table 6 Column 2 shows that the estimated interaction of unexpected wins with the number of days taken to reach the cap is positive but insignificant. It also shows that the interaction of unexpected visas with the IHS of prior patents is extremely imprecise. The interaction of unexpected visas with prior firm size is also insignificant.³³

Interpreting the estimates

Our “reduced form” effects are relevant for firms and policy-makers interested in better understanding the average employment effects of a policy granting additional H-1B visas firms by marginally relaxing the cap. These results show employment on average will increase by less than one worker for every additional capped H-1B visa, and the estimates show no indication that employment will rise at all on net.

Moving beyond the policy-relevant “reduced form” effects, institutional features of this labor market are relevant to determining whether new H-1Bs crowd out other workers. In principle, a limitation of our results is that we do not observe if the worker actually ended up at a firm (as opposed to having an approved H-1B visa, which we do observe). For example, after being approved by USCIS, some workers may die before being admitted to the U.S. to start their job, or the State Department may not approve their visa. However, in practice this is likely to affect our employment results only negligibly. In the employment context, we examine (among other things) the immediate impact on employment in the first quarter of Year 0; North (2011) estimates that 95% of those approved for H-1Bs end up being admitted.³⁴ This would not pose an issue for ruling out that employment of the H-1B causes a one-for-one increase in employment, because in the employment context we are typically able to rule out coefficient in the initial quarters that is under 0.6 (*i.e.* well under 0.95). North (2011) also estimates that 82% of those allowed the H-1B are still at their initial firm for the full three years, as some workers return home or depart for another reason. Similarly, in the National Survey of College Graduates 2003 and 2010 waves, 79 percent of temporary visa holders who had been previously employed in the U.S. for two years were at the initial employer. This

³³ More generally, there are many factors that theoretically could influence the size of the impacts, but we tend to find that the estimates are similar across groups.

³⁴ In North (2011), the fraction admitted is calculated by including those who were already in the U.S. and apply for a renewal of their H-1B. Excluding these individuals would not materially change our conclusions.

is relevant to interpreting our patenting results and our longer-term employment results. Note, however, that our results would be similar if they were scaled up by 22% ($=1/0.82$), and that the top end of the 95 percent confidence interval usually rules out coefficients of 0.82 or greater even in later time periods like Q5-Q8 or Q9-Q12.³⁵

In the case of the median regressions in the employment context, which are “reduced form” regressions, the coefficients do not take account of the fact that some H-1B lottery winners do not have their applications approved. However, our first stage coefficient is extremely precise and quite close to 1 (specifically, it ranges from 0.86 to 0.88), so this consideration is also unlikely to change our conclusion that H-1B workers at least partially replace other workers at the firm.³⁶ Moreover, the two-stage least squares specifications in this context show comparable results.

Even taking a “worst case” scenario in which these factors worked together, we would still generally conclude that H-1Bs displace other workers at least to some extent. For example, the top end of the 95 percent confidence interval is 0.37 among firms with 30 or fewer employees, and it is 0.12 among firms with 10 or fewer employees. Even after scaling these estimates to reflect the modest attrition issues described in this section, it is clear that the top end of the 95 percent confidence interval will still be far below one.

If firms respond to an extra H-1B visa by reducing contracting work or outsourcing to other firms or countries—neither of which appears in our measure of employment at the firm itself—then by examining only employment at the firm, new H-1Bs will appear to be *less* substitutable with other potential employees than they actually are. Thus, it is all the more notable that we are able to rule out a coefficient on unexpected H-1Bs of one or greater. Fraud has also been alleged in the context of H-1Bs;³⁷ this could lead to a larger coefficient on unexpected H-1Bs (if firms fraudulently

³⁵ In rare cases, workers start working at the firms after the first quarter of the first year. We use USCIS administrative data on the proposed start dates of each H-1B application that won the lottery in our sample to calculate that 91.87 percent of H-1Bs started working at the firms under this H-1B in Q1, and 100 percent had started working at the firms by Q2. Thus, nearly everyone had started working at the firms, and this does not represent a major issue. Our Q1 estimates would be little affected by scaling our estimates to account for this (*i.e.* multiplying by $1/0.9187$).

³⁶ In the rare quantile instrumental variables median regressions that did converge, the coefficients on unexpected lottery wins were only around 10 percent larger than in the “reduced form” median regressions.

³⁷ For example, see <http://www.bloomberg.com/news/2013-10-30/infosys-settles-with-u-s-in-visa-fraud-probe.html> (accessed September 16, 2014).

obtain other types of visas for the workers who would have been H-1Bs if the firm had been awarded an H-1B) or a smaller coefficient (if the firm responds to not receiving an H-1B by hiring a worker off the books).

8. Effects on profits and payroll per employee

Firms attest to paying H-1Bs a prevailing wage, but it is possible that H-1B sponsoring firms could pay H-1Bs less relative to other comparable workers, for example if the sponsoring firm has a greater degree of monopsony power/bargaining power with respect to the H-1B than with alternative workers. Table 10 shows the effect of unexpected H-1B visas on median firm yearly payroll per employee during the duration of the H-1B visa (stacking Years 0 to 3), calculated by dividing total firm payroll in a given year by the total number of employees appearing in W-2s in that year.³⁸ In firms with 10 or fewer, or 30 or fewer, employees, we find some evidence that the additional H-1B reduces median payroll per employee ($p < 0.05$ in one estimate, and $p < 0.10$ in two estimates, of the four total). The point estimates suggest substantial decreases in payroll per employee in these firms (with larger point estimates in the smaller firm size category); however, the confidence intervals do not rule out much smaller effects. In the full sample of firms, the additional H-1B worker typically reflects only a small percentage of total employment and would be expected to influence payroll per employee little, and unsurprisingly we estimate no measurable effect in these firms.³⁹ Note that H-1Bs could in principle decrease payroll per employee not only if the firm pays the H-1B less than an alternative worker, but also if the H-1B negatively affects the earnings of other employees at the firm.

Many have suggested that firms are able to increase profits by paying H-1Bs less, as in Table 10. Table 11 examines the effect of an unexpected H-1B visa on the firm's reported profit in Years 0 to 3, using median regressions.⁴⁰ The point estimate is always positive across all the firm size cutoffs considered and is sometimes significant. The point

³⁸ In the employment context, we are more interested in comparing the coefficient on unexpected visas to a specific level (*i.e.* 1); thus, it makes sense to limit the primary time period to only Q1 to Q4, when the employee is almost always at the firm.

³⁹ When investigating other firm size thresholds, we typically continue to find negative effects, though they unsurprisingly become increasingly attenuated at larger firm size thresholds. We find no significant interactions of the effect with covariates, and no significant differences in the effect across groups.

⁴⁰ Note the caveat that we observe profits in the IRS data, not economic profits.

estimates generally cluster around showing an increase in profits of five to ten thousand dollars per year, though in all cases the 95 percent confidence interval does not rule out a substantially smaller effect on profits. The median regressions do not converge for many firm size cutoffs, including for the sample of firms of all sizes; the largest firm size cutoff we show is less than or equal to 200 employees, as the regressions did not reliably converge above this threshold. Across thresholds between 30 and 200 for which the regressions did converge, the regressions generally continue to cluster around showing a positive effect of approximately five to ten thousand dollars per year. Overall, we have some evidence of a positive effect on profits, though it is not robustly significant.

Our results on profits and payroll per employee suggest the existence of market frictions that allow such profits, and that allow lower wages to be paid to H-1B employees—such as regulations restricting the free flow of workers across borders and/or firm labor market monopsony power. Our results on payroll per employee suggest that this effect may be related to at least some of the increase in profits (as opposed to profits increasing, for example, because H-1B workers increase firm gross income).⁴¹ While profits and payroll per employee are important outcomes, we consider these results to be secondary because our results are not fully dispositive: the profits results often did not converge, and the results on profits and payroll per employee are less precise.⁴²

Firm productivity is also of interest. Appendix Table 10 shows that the effect on revenue per employee, or net income per employee, is insignificant (though imprecise).⁴³ The effect on total firm payroll is also insignificant and imprecise.

9. Conclusion

The effect of raising the H-1B visa cap is one of the centrally important U.S. immigration policy questions. We examine the marginal impact of allowing an extra

⁴¹ Note the important caveat that medians do not “add,” as the effect on median profits may refer to different firms than the effect on median payroll per employee.

⁴² It is also possible that an unexpected H-1B lottery win affects a firm’s competitors. We find no significant impact of unexpected H-1B lottery wins on any of the outcome variables among all other firms in that firm’s 6-digit NAICS code, which is unsurprising given the large size of a six-digit industry.

⁴³ Likewise, median or OLS regressions in which firm gross income or non-payroll costs are the dependent variables are imprecise. For example, in the median regression in which total yearly firm gross income over years 0 to 3 is the dependent variable, the coefficient on unexpected H-1B visas is -\$19,197.43, with a very large standard error of 76,438.43. Ghosh, Mayda, and Ortega (2015) study effects of H-1Bs on productivity, firm size, profits, and other outcomes.

capped H-1B visa to a firm on the firm's outcomes. We find an insignificant effect of additional H-1B visas on patenting, and across a variety of specifications the preponderance of evidence allows us to rule out moderate-sized or larger effects. Parallel to these patenting results, we find that H-1B workers at least partially replace other workers, with the estimating generally indicating substantial crowdout of other workers. More suggestive evidence indicates that new H-1Bs crowd out other foreigners at least to some extent, and that H-1Bs do not replace U.S. citizens/natives one-for-one. The results hold when we exclude firms that likely specialize in temporary help and often provide technical support services, and we find no evidence that the effects vary based on prior firm size. It is arguably striking that we find little patenting or employment effect even among firms applying on the day the cap is reached, which are more likely than those in the full sample to have patented in the past and to be in scientific industries. We find some evidence that H-1Bs cause a decrease in median earnings per employee, and consistent with the presumption of firm profit maximization, we find some evidence that H-1B visas increase median firm profits.

Our results are consistent with the narrative about the effects of H-1Bs on firms in which H-1Bs are paid less than alternative workers whom they replace, thus increasing the firm's profits. Profits may increase despite no measurable effect on patenting (though we do not rule out that productivity or measures of innovation respond). The results raise the possibility that firms' behavior is inconsistent with two aspects of the law—firms' legal obligation not to adversely affect similarly employed workers, and their attestations that they pay a prevailing wage—in at least some cases that drive our estimates.⁴⁴

Our results are consistent with the possibility that H-1B and non-H-1B workers are perfect substitutes. This is notable in light of frequent claims that H-1Bs have unique skills that cannot easily be obtained elsewhere. Our study focuses on estimating the causal impacts of additional H-1Bs, which could provide some of the building blocks for estimating parameters such as the elasticity of substitution between new H-1Bs and other

⁴⁴ If the H-1B receives the job rather than a similar worker, then the similar worker is harmed as long as the alternative job she may find yields lower utility. There is an *a priori* case that the workers crowded out by the H-1B were “similarly employed”: arguably, if H-1Bs displace anyone, they should displace the most closely substitutable, similar workers. Note also that the Congressional intent may have been to prevent harm to U.S. workers specifically.

workers in future work. However, we would be limited by having one instrument—H-1B lottery wins—but multiple parameters of interest. The degree of crowdout of other workers should depend not only on the nature of the substitutability or complementarity of additional H-1B and other workers (and/or labor and capital), but also factors including the nature of the process that matches firms with workers (possibly including search frictions).⁴⁵ If the firm faces frictions in finding a new employee that *limit* the degree of crowdout of other workers (consistent with *e.g.* Isen 2014), it would be all the more notable that we find that an H-1B worker *does* partially replace other workers, and that we cannot rule out that an H-1B worker has no effect on total employment. Because the degree of crowdout of employment of other workers depends on a multiplicity of factors, one cannot interpret our estimates as necessarily *implying* that H-1Bs are perfect substitutes with other technical workers.

Our study is different from previous work on the effects of H-1Bs on the receiving economy in several notable ways. First, we are the first to use randomization, and we are one of the first to use administrative data. Second, we examine the effects of H-1Bs given to a particular firm on that firm's outcomes (holding constant H-1Bs at other firms). Third, some previous literature examines the effects of temporary visas in general, not specifically those subject to the H-1B cap. Fourth, our results are estimated from the FY2006 and 2007 lotteries, where the results may differ from other environments.⁴⁶ At first pass, our results also apparently differ starkly from those in the previous economics literature, as previous work has found very large positive effects of H-1Bs on patenting and employment.⁴⁷ Though it is outside the scope of this paper, which focuses on estimating the effects of H-1Bs on firm outcomes, future work could try to clarify further

⁴⁵ Estimating such parameters would be difficult in our data because the additional H-1B workers may not represent the same quantity of labor as other workers do; the additional H-1B workers could work a greater or smaller number of hours than other workers do. Hours worked is unobserved in our data, as in most administrative datasets. Lewis (2011) studies the interaction of immigration with capital.

⁴⁶ Kerr and Lincoln (2010) exploit variation in the cap in other years (including the expansion of the cap in the late 1990s and subsequent contraction).

⁴⁷ For example, Kerr and Lincoln (2010) find that a 10 percent growth in a city's H-1B population corresponded with a 0.3 percent to 0.7 percent increase in total patenting for each standard deviation growth in "city dependency," a measure of H-1B applications per capita in each city. Given the standard deviation of city dependency in their sample, this would imply an increase in patenting at least 10 times as large as the maximum effect allowed by our 95% confidence interval. However, as noted, the estimates are not directly comparable for several reasons.

whether the divergence in substantive results—*i.e.* large positive effects *vs.* no evidence of a positive effect—relates to the difference in identification strategy, the difference in the type of outcome examined (*i.e.* aggregate *vs.* firm-level), the type of visas in question, or the particular contexts examined. To address the effects of H-1Bs in other years, USCIS could begin regularly saving the data on H-1B lottery winners and losers.


References

- Åslund, Olof, Per-Anders Edin, Peter Fredriksson, and Hans Grönqvist.** “Peers, Neighborhoods and Immigrant Student Achievement: Evidence from a Placement Policy.” *American Economic Journal: Applied Economics* 3.2 (2011): 67-95.
- Bloom, Nicholas, Mark Schankerman, and John van Reenen.** “Identifying Technology Spillovers and Product Market Rivalry.” *Econometrica* 81.4 (2013): 1347-1393.
- Borjas, George.** “The Labor Demand Curve Is Downward-Sloping: Reexamining the Impact of Immigration on the Labor Market.” *Quarterly Journal of Economics* 118 (2003): 1335-374.
- Borjas, George.** “The Economics of Immigration.” *Journal of Economic Literature* 32 (1994): 1667–1717.
- Borjas, George, and Kirk Doran.** “Cognitive Mobility: Native Responses to Supply Shocks in the Space of Ideas.” *Journal of Labor Economics* (forthcoming 2015).
- Borjas, George, and Kirk Doran.** “The Collapse of the Soviet Union and the Productivity of American Mathematicians.” *Quarterly Journal of Economics* 127.3 (2012): 1143-203.
- Borjas, George, Richard Freeman, and Lawrence Katz.** “How Much Do Immigration and Trade Affect Labor Market Outcomes?” *Brookings Papers on Economic Activity* (1997): 1-90.
- Borjas, George, Jeffrey Grogger, and Gordon Hanson.** “Comment: Substitution between Immigrants, Natives, and Skill Groups.” *Journal of the European Economic Association* 10.1 (2012): 198-210.
- Bound, John, Breno Braga, Joseph Golden, and Gaurav Khanna.** “Recruitment of Foreigners in the Market for Computer Scientists in the U.S.” Forthcoming, *Journal of Labor Economics*.
- Burbidge, John, Lonnie Magee, and A. Leslie Robb.** “Alternative Transformations to Handle Extreme Values of the Dependent Variable.” *Journal of the American Statistical Association* 83 (1988): 123-127.
- Card, David.** “The Impact of the Mariel Boatlift on the Miami Labor Market.” *Industrial and Labor Relations Review* 43.2 (1990): 245-57.
- Card, David.** “Immigrant Inflows, Native Outflows, and the Local Market Impacts of Higher Immigration.” *Journal of Labor Economics* 19.1 (2001): 22-64.
- Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan.** “How does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR.” *The Quarterly Journal of Economics* 126.4 (2011): 1593-1660.
- Cortes, Patricia, and Jessica Pan.** “Foreign Nurse Importation and the Supply of Native Nurses.” *Journal of Health Economics* 37 (2014): 164-180.
- Doran, Kirk, Alexander Gelber, and Adam Isen.** “The Effect of High-Skilled Immigration on Patenting and Employment: Evidence from H-1B Visa Lotteries.” NBER Working Paper 20668.
- Dustmann, Christian, Albrecht Glitz, and Tommaso Frattini.** “The Labour Market Impacts of Immigration.” *Oxford Review of Economic Policy* 24.3 (2008): 477-94.

- Edin, Per-Anders, Peter Fredriksson, and Olof Åslund.** “Ethnic Enclaves and the Economic Success of Immigrants—Evidence from a Natural Experiment.” *The Quarterly Journal of Economics* 118.1 (2003): 329-57.
- Foley, Fritz, and William Kerr.** “Ethnic Innovation and U.S. Multinational Firm Activity.” *Management Science* 59.7: 1529-1544.
- Freeman, Richard.** “People Flows in Globalization.” *Journal of Economic Perspectives* 20.2 (2006): 145–70.
- Friedberg, Rachel M.** “The Impact of Mass Migration on the Israeli Labor Market.” *The Quarterly Journal of Economics* 116.4 (2001): 1373-408.
- Friedberg, Rachel, and Jennifer Hunt.** “The Impact of Immigrants on Host Country Wages, Employment and Growth.” *Journal of Economic Perspectives* 9.2 (1995): 23-44.
- Gates, William H.** “Testimony before the Committee on Science and Technology.” U.S. House of Representatives, Washington D.C. 12 March. 2008. Address.
- Gelber, Alexander M.** “How Do 401(k)s Affect Saving? Evidence from Changes in 401(k) Eligibility.” *American Economic Journal: Economic Policy* 3.4 (2011): 103-22.
- Ghosh, Anirban, Anna Maria Maybe, and Francesc Ortega.** “The Impact of Skilled Foreign Workers on Firms: an Investigation of Public Traded U.S. Firms.” CUNY Working Paper (2015).
- Grogger, Jeff, and Gordon Hanson.** “The Scale and Selectivity of Foreign Born Ph.D. Recipients in the U.S.” *American Economic Review Papers and Proceedings* 103.3 (2013): 189-192.
- Hira, Ron.** “The H-1B and L-1 Visa programs: Out of Control”, EPI Policy Paper (2010).
- Hunt, Jennifer.** “Which Immigrants Are Most Innovative and Entrepreneurial? Distinctions by Entry Visa.” *Journal of Labor Economics* 29.3 (2011): 417-57.
- Hunt, Jennifer, and Marjolaine Gauthier-Loiselle.** “How Much Does Immigration Boost Innovation.” *American Economic Journal: Macroeconomics* 2.2 (2010): 31-56.
- Isen, Adam.** “Dying to know: Are Workers Paid their Marginal Product?” University of Pennsylvania working paper (2013).
- Jones, Benjamin.** 2010. “Age and Great Invention.” *Review of Economics and Statistics* 92(1): 1-14.
- Imbens, Guido, Donald Rubins, and Bruce Sacerdote.** “Estimating the Effect of Unearned Income on Labor Supply, Earnings, Savings and Consumption: Evidence from a Sample of Lottery Players.” *American Economic Review* 91.4 (2001): 778-94.
- Kerr, William R.** “U.S. High-Skilled Immigration, Innovation, and Entrepreneurship: Empirical Approaches and Evidence.” HBS Working Paper (2013): 14-17.
- Kerr, William R., and William F. Lincoln.** “The Supply Side of Innovation: H-1B Visa Reforms and US Ethnic Invention.” *Journal of Labor Economics* 28.3 (2010): 473-508.
- Lewis, Ethan.** “Immigration, Skill Mix, and Capital Skill Complementarity.” *Quarterly Journal of Economics* 126.2 (2011): 1029-1069.


- Lowell, B. Lindsay.** “H-1B Temporary Workers: Estimating the Population.” UCSD Center for Comparative Immigration Studies Working Paper No. 12 (2000).
- Lubotsky, Darren.** “Chutes or Ladders? A Longitudinal Analysis of Immigrant Earnings.” *Journal of Political Economy* 115.5 (2007): 820-867.
- Matloff, Norman.** “On the Need for Reform of the H-1B Non-immigrant Work Visa in Computer-Related Occupations.” *University of Michigan Journal of Law Reform* 36.4 (2003): 815-914.
- Miano, John.** “H-1B Visa Numbers: No Relationship to Economic Need.” *Center for Immigration Studies Policy Brief* (2008).
- Moser, Petra, Alessandra Voena, and Fabian Waldinger.** “German Jewish Émigrés and U.S. Invention.” *American Economic Review* 104: 3222-3255.
- National Foundation for American Policy.** “H-1B visas and job creation.” *Policy Brief*, Arlington, VA (2008).
- North, David.** “Estimating the Size of the H-1B Population in the U.S.” *Center for Immigration Studies Memorandum* (2011).
- Ottaviano, Gianmarco, and Giovanni Peri.** “Rethinking the Effects of Immigration on Wages.” *Journal of the European Economic Association* 10.1” 152-197.
- Pekkala Kerr, Sari, and William R. Kerr.** “Economic Impacts of Immigration: A Survey.” *Finnish Economic Papers* 24.1 (2011): 1-32.
- Pekkala Kerr, Sari, William R. Kerr, and William F. Lincoln.** “Skilled Immigration and the Employment Structures of U.S. Firms.” *Journal of Labor Economics* (forthcoming).
- Pence, Karen.** “The Role of Wealth Transformations: An Application to Estimating the Effect of Tax Incentives on Saving.” *The B.E. Journal of Economic Analysis & Policy* 5.1 (2006): 1-24.
- Peri, Giovanni, Kevin Shih, and Chad Sparber.** “STEM Workers, H-1B Visas, and Productivity in US Cities.” *Journal of Labor Economics* (forthcoming).
- Peri, Giovanni, Kevin Shih, and Chad Sparber.** “The Effects of Foreign Skilled Workers on Natives: Evidence from the H-1B Visa Lottery.” UC Davis Working Paper (2014).
- Stuen, Eric, Ahmed Mushfiq Mobarak, and Keith E. Maskus.** “Skilled Immigration and Innovation: Evidence from Enrolment Fluctuations in U.S. Doctoral Programs.” *Economic Journal* 122 (2012): 1143-1176.
- U.S. Customs and Immigration Services.** “Change in H-1B Procedures Trims Weeks Off Final Selection Process.” Web.
<http://www.uscis.gov/sites/default/files/files/pressrelease/H1Bfy08CapUpdate041907.pdf> (accessed October 7, 2014).
- U.S. Department of Labor.** Office of Foreign Labor Certification Data. Web.
<http://www.foreignlaborcert.doleta.gov/performance/cfm#stat> (accessed October 7, 2014).
- U.S. Patent and Trade Office.** “Performance and Accountability Report: Fiscal Year 2012.” Washington, D.C.: U.S. Government Printing Office.
- Yagan, Danny.** “Moving to Opportunity? Migratory Insurance over the Great Recession.” UC Berkeley Working Paper (2014).

Figure 1. *Effect of Unexpected H-1B Visas on Patents, by Employer Size*



Notes: The figure shows the coefficient and 95 percent confidence interval on unexpected H-1B visas when the dependent variable is the inverse hyperbolic sine of the total number of patents over Years 0-7, among employers of the indicated sizes or smaller in Year -1 (where employer size is shown on the x-axis). We show the coefficient for employers of each size range from 0-10 to 0-500, with the upper bound of the size range in increments of 10. We use the baseline specification, in which we control for lagged number of patents and expected lottery wins. After multiplying by 100, the coefficient should be interpreted as the approximate percentage increase in total firm employment associated with an unexpected H-1B visa lottery win.

Figure 2. *Effect of H-1B Visas on Total Firm Employment, by Employer Size*



Notes: The figure shows the coefficient and 95 percent confidence interval on unexpected lottery wins from median regressions when the dependent variable is the *total* number of employees in a firm, pooling together Quarters 1-4 of the first fiscal year that an employee can work at the firm in the regression, among employers of the indicated size or smaller in Year -1 (where employer size is shown on the x-axis). The horizontal line at +1 on the y-axis corresponds to the case where hiring an extra H-1B visa worker leaves other employment unchanged (so that total employment would increase by exactly one). The horizontal line at 0 on the y-axis corresponds to the case where hiring an extra H-1B visa worker precisely crowds out other workers one-for-one (so that total employment would increase by zero). We show the coefficient for employers of each size range from 0-10 to 0-500, with the upper bound of the size range in increments of 10. We use the baseline specification, in which we control for lagged employment and expected lottery wins.

Table 1. Summary Statistics

<u>Variable</u>	<u>Mean (SD)</u>	<u>N</u>
Fraction Patenting (all)	0.093 (0.29)	3,050
Fraction Patenting (≤ 30)	0.033 (0.18)	1,276
Fraction Patenting (≤ 10)	0.025 (0.16)	749
Number of Patents (all)	37.74 (390.95)	3,050
Number of Patents (≤ 30)	1.92 (61.74)	1,276
Number of Patents (≤ 10)	0.19 (2.87)	749
IHS of patents (all)	0.33 (1.28)	3,050
IHS of patents (≤ 30)	0.064 (0.37)	1,276
IHS of patents (≤ 10)	0.048 (0.34)	749
Median employees in Q1-Q4 (all)	31	2,281
Median employees in Q1-Q4 (≤ 30)	10	1,183
Median employees in Q1-Q4 (≤ 10)	6	712
Winsorized emp. first diff. in Q1-Q4 (all)	27.28 (92.39)	2,281
Winsorized emp. first diff. in Q1-Q4 (≤ 30)	4.35 (9.43)	1,183
Winsorized emp. first diff. in Q1-Q4 (≤ 10)	3.22 (6.84)	712
Median payroll per employee (all)	49,331.89	2,326
Median payroll per employee (≤ 30)	42,280.76	1,118
Median payroll per employee (≤ 10)	38,656.64	630
Median firm profits (≤ 200)	80,249.73	1,520
Median firm profits (≤ 30)	43,300.70	1,033
Median firm profits (≤ 10)	30,397.45	615
Fraction winning lottery		
2006 Regular	0.038	2,687
2006 ADE	0.17	306
2007 Regular	0.98	3,954
2007 ADE	0.55	296

Notes: The source of the data is IRS and USCIS administrative data, and the Patent Network Dataverse. “All” refers to the full sample of firms that enter the lottery; “ ≤ 30 ” refers to those firms that have 30 or fewer employees in Year -1; “ ≤ 10 ” refers to those firms that have 10 or fewer employees in Year -1. Employment data are observed in Q1-Q4, the first four quarters when the H-1B worker may work at the firm (which are the same four quarters we investigate in our main employment results in Table 7). Payroll per employee is measured in Years 0 and 1 to parallel the period over which employment is observed in our baseline, as the H-1B is most likely to be at the firm in this period (so it is most relevant to measure outcomes relating to employees during this period). The number of patents refers to approved patents from the year of the lottery (2006 or 2007) and the subsequent seven years. Firm profits are measured in years 0 to 3 to capture the longer-term effect on profits. For profits, we investigate the additional size category with ≤ 200 employees, because our regressions did not converge for larger size thresholds. “N” refers to the number of firms in the sample, except in the final rows reporting the fraction winning the lottery, where we report the number of applications that entered the lottery. The number of observations varies across outcomes because the number of missing observations in the IRS data varies across outcomes; here and everywhere else, the results are extremely similar when we restrict the sample to be the same across all outcomes.

Table 2. Comparison of Firms Applying on Day of Lottery to Other Applicants

<u>Dependent Variable</u>	<u>Coefficient (SE) on “Last Day” Dummy</u>	<u>N</u>
IHS of patents from Year -3 to Year -1	0.072 (0.023)***	51,483
Fraction patenting from Year -3 to Year -1	0.017 (0.0066)***	51,483
Fraction of applications in “systems analysis and programming” occupations	0.22 (0.0090)***	51,483
IHS of employment in Year -1	0.10 (0.052)**	41,849
Fraction in NAICS=54	0.17 (0.0097)***	46,706

Notes: The table compares characteristics of firms that applied on the day the cap was reached (so they are subject to the lottery) to all firms whose applications reached USCIS (including others that applied before the cap was reached). We report the coefficient and standard error on the dummy for applying on the last day, from an OLS regression of the dependent variable (shown in the first column) on a dummy for applying on the last day, plus dummies for each of the four lotteries (FY06 Regular, FY06 ADE, FY07 Regular, FY07 ADE). Observations on (rare) firms that applied on both the last day and prior to the last day are included in both the sample of firms applying on the last day and the sample applying prior; thus, the table effectively compares firms that applied only on the last day to firms that applied only on days before the last day. Sample sizes differ across regressions because some outcomes are missing in some cases (for example, Year -1 employment is missing in some cases because the firm did not exist in Year -1). The sample size is also far below the number of total visa applications received across these lotteries primarily because a small number of firms apply for many visas, with a very skewed distribution. The regressions treat applications from a given firm on different days as separate observations. Standard errors are clustered by firm. N’s refer to the number of firms. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Table 3. Validity of the Randomized Design

<u>Dependent Variable</u>	<u>Coefficient (SE) on Unexpected Wins</u>
Lottery data has firm information	0.0028 (0.0032)
Whether match to tax master file	0.0080 (0.0079)
Whether match to quarterly employment data	-0.0031 (0.0096)
IHS of patents from Year -3 to Year -1 (all)	0.079 (0.060)
IHS of patents from Year -3 to Year -1 (≤ 30)	-0.036 (0.025)
IHS of patents from Year -3 to Year -1 (≤ 10)	-0.012 (0.0087)
Patented from Year -3 to Year -1 (all)	-0.0039 (0.021)
Patented from Year -3 to Year -1 (≤ 30)	-0.026 (0.019)
Patented from Year -3 to Year -1 (≤ 10)	-0.0032 (0.0097)
Employment in Year -2 (all, quantile)	0.56 (0.62)
Employment in Year -2 (≤ 30 , quantile)	-0.55 (0.45)
Employment in Year -2 (≤ 10 , quantile)	-0.31 (0.44)
Employment in Year -2 (all, winsorized)	0.082 (9.71)
Employment in Year -2 (≤ 30 , winsorized)	0.56 (0.89)
Employment in Year -2 (≤ 10 , winsorized)	-0.091 (0.57)
Employment in Year -1 (all, quantile)	2.91 (4.41)
Employment in Year -1 (all, winsorized)	30.35 (104.55)
Payroll per employee in Year -2 (all, quantile)	91.01 (594.95)
Payroll per employee in Year -2 (≤ 30 , quantile)	1,591.82 (1,519.61)
Payroll per employee in Year -2 (≤ 10 , quantile)	1,645.07 (3,141.91)
Profits in Year -2 (≤ 200 , quantile)	-6,268.96 (4,528.82)
Profits in Year -2 (≤ 30 , quantile)	-8,027.92 (5,498.00)
Profits in Year -2 (≤ 10 , quantile)	-20,306.35 (19,756.56)
Dummy for NAICS=54 (all)	0.007 (0.03)
Dummy for NAICS=54 (≤ 30)	-0.033 (0.043)
Dummy for NAICS=54 (≤ 10)	0.010 (0.058)

Notes: The table illustrates the validity of the randomized design by performing regressions of placebo outcomes on unexpected H-1B lottery wins. We run OLS regressions for those dependent variables for which our main regressions in later tables are OLS (*i.e.* patenting and winsorized employment, plus the NAICS=54 dummy and the match dummies in the first three rows), and we run median regressions for the dependent variables for which our main regressions in later tables are median regressions (*i.e.* employment, earnings per employee, and profits). In the first three rows, the dependent variables are the following (in order of appearance): a dummy for whether the USCIS lottery data have information on the firm's EIN; a dummy for whether a firm's EIN in the USCIS data matches to the EIN of a firm in the IRS file on the universe of U.S. EINs; and a dummy for whether a firm's EIN in the USCIS data matches to the EIN of a firm in the IRS quarterly employment data. Dummies for whether profits or payroll match are also uncorrelated with treatment. When we regress Year -2 measures of patenting on unexpected wins and the parallel Year -1 measures of patenting, we also estimate insignificant coefficients on unexpected wins. In the specifications where employment, earnings per employee, or profits in Year -2 are the dependent variables, we control for employment, earnings per employee, or profits (respectively) in Year -1 to parallel the control for Year -1 employment in our regressions in later tables. To determine the firm size cutoffs, employment is measured in the prior year. We investigate the profits regressions in the sample with 200 or fewer employees, rather than the full sample, because the regressions did not converge for the full sample. In the specifications in which employment in Year -1 is the dependent variable (which we analyze to examine a period closer to Year 0), we have no controls (as we clearly cannot control for Year -1 employment in this context), and we only investigate the results in the "All" sample because selecting this sample based on Year -1 employment would lead to biased and inconsistent results. When we investigate the pre-period in the employment context, we examine only Years -1 and -2, rather than examining a longer pre-period such as all years from Year -3 to Year -1 (as in the case of the patenting data), because the IRS quarterly employment data begin in Year -2 (corresponding to the first quarter of year 2004). When Year -1 employment is the dependent variable and we control for Year -2 employment (not shown), we estimate an insignificant effect with precision similar to the employment regressions in which we investigate the effect on Year -2 employment and control for Year -1 employment. Separately, we also regressed lottery wins on dummies for all two-digit NAICS codes and perform an F-test for joint significance of these dummies; this test showed insignificant results (for example, when using the sample of all firms, $p=0.96$). "Winsorized" means that we winsorize at the 5th and 95th percentiles. Standard errors are clustered by firm. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Table 4. Effect of Unexpected H-1B Lottery Wins on Patenting from Years 0 to 7

	IHS of number of patents		Patenting dummy	
A) ≤ 10 employees	-0.0039 [-0.027, 0.019]	-0.0045 [-0.027, 0.018]	-0.014 [-0.036, 0.0089]	-0.014 [-0.036, 0.0087]
B) ≤ 30 employees	-0.015 [-0.061, 0.030]	-0.015 [-0.060, 0.030]	0.0086 [-0.020, 0.037]	0.0088 [-0.020, 0.037]
C) All firm sizes	-0.053 [-0.12, 0.012]	-0.051 [-0.11, 0.011]	-0.0044 [-0.031, 0.022]	-0.0043 [-0.030, 0.022]
Prior patents	X	X	X	X
E[wins]		X		X

Notes: The table shows OLS regressions of patenting over Years 0 to 7 on unexpected H-1B lottery wins, defined as actual wins minus the expectation of wins conditional on number of applications and the probability each application wins. The table shows coefficients on unexpected H-1B visas, with 95 percent confidence intervals in brackets. The “prior patents” specifications control for the total number of patents from 2000 to Year -1. The “prior patents, E[wins]” specifications control for patents in the pre-period and expected lottery wins (equal to number of H-1B applications considered in a lottery multiplied by the probability of winning the lottery). See Table 1 for additional notes and sample sizes. Standard errors are clustered by firm. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Table 5. Effect of Unexpected H-1B Lottery Wins on Patenting in Subgroups

	(1) ≤ 10 employees	(2) ≤ 30 employees	(3) All firm sizes
A) Regular	0.011 [-0.0053, 0.026] {681}	-0.011 [-0.070, 0.047] {1,136}	-0.019 [-0.067, 0.028] {2,540}
B) ADE	-0.050 [-0.12, 0.024] {68}	-0.023 [-0.090, 0.044] {140}	-0.087 [-0.25, 0.074] {510}
C) Professional, sci., and tech. services	-0.012 [-0.041, 0.017] {484}	-0.022 [-0.078, 0.034] {837}	-0.066 [-0.14, 0.0089]* {1,721}
D) Industries other than professional, sci., and tech. services	0.019 [-0.0075, 0.045] {265}	0.012 [-0.036, 0.060] {439}	-0.014 [-0.12, 0.092] {1,329}
E) “Temporary support services” industries	-0.017 [-0.054, 0.019] {410}	-0.0011 [-0.048, 0.045] {698}	-0.058 [-0.14, 0.024] {1,398}
F) Non-“temporary support services” industries	0.022 [-0.0067, 0.051] {339}	-0.029 [-0.11, 0.050] {578}	-0.038 [-0.14, 0.060] {1,652}

Notes: The table shows OLS regressions of the IHS of patents over Years 0 to 7 on unexpected H-1B lottery wins. All specifications control for patents in the pre-period and expected lottery wins, as in the baseline. The results are comparable when we investigate the patenting dummy or the number of patents as the dependent variable. “Temporary consulting industries” refers to 6-digit NAICS codes 541511, 541519, 541600, 541330, 519100, 423600, and 541512; “non-temp industries” refers to all others. “Professional, scientific, and technical services” refers to NAICS code 54. The number of observations is in {curly brackets} below the confidence intervals in [square brackets]. See Tables 1 and 4 for additional notes. Standard errors are clustered by firm. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Table 6. Interactions of Unexpected Visa Lottery Wins with Covariates

Outcome:	(1) IHS of patents, Years 0 to 7	(2) Employment in Q1 to Q4
Panel A: Interaction of unexpected visas with days to reach cap	-0.000053 [-0.00094, 0.00084]	0.038 [-0.030, 0.11]
Panel B: Interaction of unexpected visas with IHS of patents in Years -3 to -1	0.0062 [-0.066, 0.078]	-5.94 [-31.48, 19.58]

Notes: The table indicates that there is no apparent significant difference in the effects on patenting or employment by employer size or by time taken to reach the visa cap. In Column 1, the dependent variable is the IHS of the number of patents from Years 0 to 7 (inclusive), and the specification is an OLS regression of the IHS of number of patents on the independent variables (as in the baseline). In Column 2, the dependent variable is the number of employees in Q1 through Q4 (pooled and stacked, with each quarter as a separate observation), and the specification is a median regression (again as in the baseline). In Panel A, the main independent variables are the number of unexpected H-1B visas; the number of days taken to reach the visa cap in the year and lottery in question; and the interaction of these two variables. In Panel B, the main independent variables are the number of unexpected H-1B visas; the IHS of total patents from Year -3 to Year -1; and the interaction of these two variables. The table shows coefficients and 95% confidence intervals on the interactions. All specifications additionally control for expected lottery wins, as well as patents in the pre-period (in Column 1) or employment in the pre-period (in Column 2) as in the baseline specifications. The time taken to reach the visa cap was 291 days in FY2006 Regular lottery, 131 days in the FY2006 ADE lottery, 116 days in the FY2007 Regular lottery, and 55 days in the FY2007 ADE lottery. When we allow the time taken to reach the cap to have a different impact in the two ADE lotteries together and the two Regular lotteries together, we also find no significant interaction in each set of lotteries taken together. Standard errors are clustered by firm. See Tables 4 and 7 for sample sizes. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Table 7. Effect of H-1B Visa on Employment in Q1 to Q4

	Median Regressions		Two-stage least squares	
	(1)	(2)	(3)	(4)
A) ≤ 10 employees (n=712)	-0.53 [-1.18, 0.12]***	-0.52 [-1.15, 0.11]***	-0.54 [-1.95, 0.88]**	-1.10 [-2.88, 0.68]**
B) ≤ 30 employees (n=1,183)	-0.44 [-1.16, 0.28]***	-0.36 [-1.09, 0.37]***	-0.97 [-2.96, 1.01]*	-1.26 [-3.25, 0.71]**
C) All (n=2,281)	-1.27 [-3.08, 0.55]***	-1.05 [-2.67, 0.57]**	-20.37 [-230.99, 190.24]	-2.41 [-17.76, 12.94]
Prior employment	X	X	X	X
E[wins]		X		X

Notes: The table shows point estimates and 95% confidence intervals. The first two columns show median regressions of total firm employment, in Q1 to Q4 pooled and stacked, on unexpected lottery wins. The next two columns show two-stage least squares regressions where the dependent variable, the difference of total firm employment from the first quarter of Year -1 to the quarter in question from Q1 to Q4, has been winsorized at the 95th percentile. The 5th and 95th percentiles of the first difference in employment are -109 and 352, respectively, in the full sample; are -9 and 30, respectively, among those with 30 or fewer employees; and are -6 and 22, respectively, among those with 10 or fewer. In these regressions, the instrument is unexpected lottery wins and the endogenous variable is approved H-1B visas. The “prior employment” specifications control for employment from the first quarter of Year -1, and the “prior employment, E[wins]” specifications additionally control for the number of expected lottery wins. “n” refers to the total number of firms used in each regression. See Tables 1 and 4 for other notes. *** denotes estimates that are significantly different from 1 at the 1% level; ** at the 5% level; * at the 10% level. If the H-1B worker works at the firm, a coefficient of 1 corresponds to no crowd-out or crowd-in of other employment, while coefficient of 0 corresponds to one-for-one-crowdout of other employment. None of the estimates is significantly different from 0 at any conventional significance level.

Table 8. Effect of Unexpected Lottery Wins on Later Employment

Outcome	(1) All	(2) ≤ 30 employees	(3) ≤ 10 employees
A) Q5-Q8	-2.03 [-4.97, 0.90]**	-0.95 [-2.29, 0.39]***	-0.99 [-2.05, 0.065]***
N	2,213	1,142	682
B) Q9-Q12	-1.97 [-5.46, 1.52]*	-1.57 [-3.70, 0.56]**	-1.02 [-2.28, 0.25]***
N	2,120	1,087	647
C) Q13-Q32	-3.24 [-7.14, 0.67]**	-0.0096 [-2.26, 2.25]	0.92 [-1.31, 3.14]
N	2,048	1,045	618

Notes: The table shows the effect of unexpected lottery wins on employment in later time periods, displaying point estimates and 95% confidence intervals in square brackets for median regressions of employment on unexpected lottery wins. The regressions pool and stack observations from different quarters. All specifications control for employment in the pre-period and expected lottery wins, as in the baseline. N’s refer to the number of firms included in each regression. See Table 7 for additional notes. The sample sizes fall in later years because fewer firms are still in business. In Q13 to Q32, the H-1B worker has typically left the firm, so the test of a difference in the coefficient from 1 does not suggest displacement of other workers. *** denotes estimates that are significantly different from 1 at the 1% level; ** at the 5% level; * at the 10% level. None of the estimates is significantly different from zero at any conventional significance level.

Table 9. Effect of Unexpected Lottery Wins on Employment in Subgroups

	(1) ≤ 10 employees	(2) ≤ 30 employees	(3) All firm sizes
A) Regular	-0.41 [-1.10, 0.27]*** {651}	-0.59 [-1.46, 0.28]*** {1,069}	-1.26 [-3.33, 0.81]** {1,969}
B) ADE	-0.0000002 [-1.36, 1.36] {67}	0.52 [-1.51, 2.55] {134}	1.38 [-5.63, 8.39] {400}
C) Professional, sci., and tech. services	-0.58 [-1.54, 0.39]*** {456}	-0.72 [-1.92, 0.48]*** {759}	-1.46 [-3.60, 0.67]** {1,275}
D) Industries other than professional, sci., and tech. services	0.36 [-0.50, 1.22] {257}	0.65 [-0.36, 1.65] {426}	1.16 [-2.74, 5.05] {1,015}
E) “Temporary support services” industries	-1.56 [-5.70, 2.57] {384}	-0.68 [-2.09, 0.73]** {628}	-1.54 [-4.03, 0.95]** {4,738}
F) Non-“temporary support services” industries	0.65 [-0.42, 1.72] {330}	0.00 [-0.95, 0.95]** {560}	0.14 [-2.46, 2.74] {1,265}

Notes: The table shows the effect of unexpected lottery wins on employment, displaying point estimates and 95% confidence intervals in [square brackets] for median regressions of employment in Q1-Q4 on unexpected lottery wins. n’s in {curly brackets} show the total number of firms. Sample sizes occasionally differ from totals in other tables because of missing observations; as in all other contexts, the results change negligibly when restricting the sample identically across all regressions. All specifications have the baseline controls: employment in the pre-period and expected lottery wins. *** shows $p < 0.01$ for the test of difference from 1; ** $p < 0.05$; * $p < 0.01$. None of the estimates is significantly different from zero.

Table 10. Effect of Unexpected Lottery Wins on Payroll per Employee

	(1) Fewer controls	(2) More controls
(A) ≤ 10 employees (n=636)	-4,527.58 [-9,258.68, 203.52]*	-4,860.54 [-9,552.97, -168.12]**
B) ≤ 30 employees (n=1,223)	-2,618.66 [-6,200.56, 963.24]	-2,725.03 [-5,976.60, 526.54]*
C) All firm sizes (n=2,191)	26.64 [-1,277.42, 1,330.69]	80.21 [-1,348.07, 1,508.50]
Prior payroll per employee	X	X
E[wins]		X

Notes: The table shows median regressions of payroll per employee in Years 0 to 3 on unexpected H-1B visas and controls. We pool and stack the four years. Years 0 to 3 cover the duration of the H-1B visa. The table shows coefficients and 95% confidence intervals on unexpected H-1B visas. The effect on payroll per employee in Years 0 to 1 is comparable to the estimates shown. Payroll per employee in a given year is measured as total firm payroll in that year divided by the number of employees in the end-of-year (December) “snapshot” of employment from W-2 data. When the dependent variable is instead the first-difference in payroll per employee from the pre-period, the results are broadly similar. Standard errors are clustered by firm. n’s refer to the number of firms. *** refers to $p < 0.01$; ** $p < 0.05$, and * $p < 0.10$.

Table 11. *Effect of Unexpected Lottery Wins on Profits*

	(1) Fewer controls	(2) More controls
(A) ≤ 10 employees (n=615)	8,163.43 [-4,724.93, 21,051.79]	6,518.156 [-6,942.69, 19,979.00]
B) ≤ 30 employees (n=1,033)	3,970.10 [-6,583.254, 14,523.46]	11,468.61 [200.86, 22,736.37]**
C) ≤ 200 employees (n=1,520)	11,538.41 [-1,490.03, 24,566.86]*	2,526.67 [-32,168.54, 37,221.88]
Prior profits	X	X
E[wins]		X

Notes: The table shows median regressions of profits in Years 0 to 3 on unexpected H-1B visas and controls. The table shows coefficients and 95% confidence intervals on unexpected H-1B visas. In Row C we investigate firms with 200 or fewer employees because regressions above this firm size cutoff did not reliably converge; they did not converge, for example, in the sample of firms of all sizes. Years 0 to 3 cover the duration of the H-1B visa. We do not show the effect on median profits in Years 4 to 7 because it is unstable and often did not converge. Estimated effects on profits in the shorter term are comparable. We find no significant interactions of the effect with covariates, and no significant differences in the effect across groups. When the dependent variable is instead the first-difference in profits from the pre-period, the results are broadly similar. Standard errors are clustered by firm. n's refer to the number of firms. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Appendix 1 (for online publication). Description of matching procedure

As described in the main text, we performed an intentionally liberal automatic matching procedure between the USCIS and patenting datasets to obtain all plausible matches between companies and patents. We then searched through the matches by hand to detect and remove all matches that appeared spurious.

The automatic matching procedure proceeded as follows. First, we assigned clearly related firm names to single categories (i.e., “Sony”, “Sony Co.”, “Sony Corporation”, etc.). Then we searched for complete string matches between the name categories in the patenting data and the name categories in the USCIS H1-B visa lottery data, and we classified these as matches between the datasets. After all such matches were made, we then searched for complete string matches between these two sets of name categories with all spaces in the names removed and also classified these as matches. Finally, we performed a “fuzzy” match between USPTO and USCIS firm names. The fuzzy matching procedure calculated a “distance” between words in each list by determining how many characters in the words need to be edited to transform a word from one list into a word in the other. This is necessary to identify all matches because, for example, firm names are occasionally misspelled. Pairs of words in firm name categories were classified as non-matching if the number of characters that differed between the words was more than one for words with six or fewer characters, or when the number of characters that differed between the words was more than two for words with seven or more characters (using the word as spelled in the USCIS data to determine the number of characters in the word). Otherwise, this pair of words was classified as a possible match. If at least 75% of the pairs of words in the firm name were possible matches, then the entire firm name was classified as a possible match.

We intentionally designed this “liberal” procedure so that it is liable to classify many non-matches as matches (but not the reverse); thus, if a firm did not match at all between the two datasets according to the fuzzy match, we can be rather certain that it was not granted any US patents between 1975 and 2013. The goal of this automatic matching procedure was to generate a list of *all* potential matches, which we could then winnow by hand in the next step.

Once this automatic matching procedure was complete, all of the resulting matches were checked by hand to determine whether they appeared to be a possible match. Of the 668 companies in the USCIS lottery list that obtained at least one automatic match in the patenting data, we identified 208 cases in which all of that company’s matches were clearly incorrect through by-hand inspection. We further identified 392 cases in which all of that company’s matches were clearly correct (legitimate variations on the correct company name) through by-hand inspection. Finally, we identified 63 cases in which the matches were ambiguous; in our judgment the match is possibly correct, but we cannot be fully confident that it is correct. We assume that both unmatched companies and those that received clearly incorrect matches did not patent at all between 1975 and 2013.

In the results that we report in the main tables, we exclude the 63 possible matches from the list of matched companies. In the Appendix, we show that the results are robust to assuming that the possible matches were in fact matches. The results are also robust to alternative assumptions and similar alternative matching procedures.

A firm would not match between the datasets if it did not patent during this time period; thus, under any of our ways of determining which companies were non-matches, we code the non-matching firms as having zero patents.

Appendix 2 (for online publication). Estimating effects on employment of foreigners and non-foreigners

Measure of foreigners and non-foreigners

In an exploratory analysis, we investigate how additional H-1Bs affect employment of other foreigners, and separately affect employment of non-foreigners. Although citizenship status is available through IRS data on W-2 forms, these data only have information on the individual's citizenship status *most recently reported* to the Social Security Administration (SSA), as opposed to always being measured in the year in question in our regressions (*e.g.* Year 0 or Year 1). Thus, one way to measure citizenship status is through this measure, which will probabilistically identify those who were citizens and non-citizens around the time of the lotteries (though with measurement error). The data on past citizenship status is not directly available, which is a relevant limitation because a substantial fraction of H-1Bs go on to become permanent residents and in many cases citizens (Lowell 2000). At the same time, for many of those who go on to become permanent residents or citizens, the SSA data will *not* reflect their updated citizenship status, for example because the Tax Identification Number under which they filed taxes as a non-citizen no longer applies once they become a citizen and gain a Social Security Number; thus, our measure of citizenship status is likely to estimate citizenship status at the time of being admitted to the U.S. with only modest error. Using this measure of citizenship of each employee, we measure a firm's employment of citizens and non-citizens from the end-of-year (December) "snapshot" of employment from the W-2.

Given the limitation of the first measure, it is desirable to use a second, unrelated method to probabilistically determine whether individuals are natives or non-natives. Using an algorithm developed in conjunction with Yagan (2014), we identify individuals as natives or non-natives on the basis of individuals' Social Security Numbers (SSNs) in the data. Prior to 2011, SSNs were assigned in a way that makes it possible to determine with a high degree of confidence whether a given individual is an immigrant to the U.S. or a native. SSNs consisted of: 1) a three-digit "Area Number" representing the area where an individual applied for the SSN; 2) a two-digit "Group Number" that is assigned in a specified sequence *within* each area number; 3) a four-digit "Serial Number" that is assigned sequentially within each Group Number.⁴⁸

⁴⁸ See <http://www.ssa.gov/history/ssn/geocard.html>

Thus, within a given geographic area associated with the Area Number, it is possible to determine on the basis of the Group Number and the Serial Number whether the individual applied for the SSN at an earlier or a later date. A majority of H-1Bs arrive when they are aged in their late 20s and early 30s. Thus, if they eventually apply for an SSN, they will do so well later in life than natives whose applications are typically submitted very early in their lives. Individuals whose SSNs indicate that they applied for the SSN late in life have a substantial probability of being an immigrant, while those whose SSNs indicate that they applied early in life have a much smaller probability of being an immigrant. We follow Yagan (2014) in probabilistically classifying individuals as immigrants when their SSNs indicate that they were in the oldest 10% of a given set of SSNs applicants within an Area Number. Our results are robust to choosing other thresholds, as well as to assigning different cutoffs (*e.g.* 15% rather than 10%) in different geographic areas with different percentages of immigrants as identified in Census data (results available upon request).⁴⁹

Estimated effect on employment of foreigners and non-foreigners

We estimate the effect on employment of foreigners vs. natives in Appendix Table 7. To make the time period investigated with these yearly W-2 data as comparable as possible to the quarterly data shown elsewhere (where we investigate Q1 through Q4 of the first fiscal year, corresponding to observations from both calendar years straddled by Q1 through Q4), we pool the W-2 end-of-year snapshot from Year 0 with the snapshot from Year 1.⁵⁰ We investigate our baseline specification across the three employer size categories we investigate elsewhere, though our results hold robustly across other employer size thresholds and other specifications.

In Rows A and B, we measure citizenship using the most recent measure of citizenship in the IRS data. When the dependent variable is the number of non-citizens employed at the firm, in all cases we are able to rule out a coefficient of one or higher—suggesting that H-1Bs do at least partially replace other non-citizens. We are unable to rule out that there is no effect of unexpected lottery wins on the median number of citizens, but we are always able to rule out that the median number of citizens decreases by one. Thus, we find evidence for crowdout of non-citizens, do not find evidence for crowdout of U.S. citizens, and are able to rule out one-for-one crowdout of citizens (though our results are consistent with substantial crowdout of citizens).

⁴⁹ Even if both were perfectly measured, citizenship at the time of the lotteries (or in the most recent IRS data) could be different than whether an individual is a native—namely, in those cases in which a non-native became a citizen prior to the time of the lotteries. Thus, there is no presumption that regressions with number of natives as the dependent variable should show the same results as regressions in which the dependent variable is the number of citizens at a later point in time.

⁵⁰ When the dependent variable is overall employment, the W-2 data show comparable results to the form 941 data. Of course, in interpreting the median regressions, we must recognize that the effects across separate regressions for foreigners and non-foreigners do not “add” to the median effect on overall employment.

One caveat to the results in Rows A and B is that because the IRS data measure most recent citizenship status rather than citizenship status at the time of application, these results could mean that H-1Bs do not displace citizens, but could also mean that H-1Bs sometimes go on to become citizens later. Likewise, the results could indicate that H-1Bs displace other non-citizens, or they could mean that H-1Bs sometimes become citizens later.

To address this ambiguity of interpretation, we also show results in Table 9 (rows C and D) where we probabilistically identify natives and non-natives using their SSNs as in Yagan (2014). Just as when we use the baseline employment specification, we find evidence for crowdout of non-natives (*i.e.* can rule out a coefficient of 1), do not find definitive evidence for crowdout of natives (*i.e.* the coefficient is insignificantly different from zero in this case), and are able to rule out one-for-one crowdout of natives (*i.e.* can rule out a coefficient of -1)—though the results are also consistent with substantial crowdout of natives. This concordance of results between two very different methods (in Rows A and B vs. C and D) increases our confidence that H-1Bs at least partially replace other foreigners.

Appendix Tables (for online publication)

Appendix Table 1. *First stage regressions*

Sample	Coefficient (SE) on Unexpected Lottery Wins	First-stage F- statistic
All	0.87 (0.03)***	993.51
≤ 30	0.88 (0.04)***	420.25
≤ 10	0.86 (0.06)***	239.94

The table shows the first stage regression of the number of approved H-1Bs on the number of unexpected wins. We show the first stage regression for the baseline specification; the first stage in other specifications is extremely similar. See Table 1 for other notes and sample sizes. *** denotes $p < 0.01$; ** denotes $p < 0.05$; * denotes $p < 0.10$.

Appendix Table 2. Effect of Unexpected H-1B Lottery Wins on the Number of Patents

	# Patents	
Panel A: ≤10 employees		
A) Years 0 to 7	0.0035 [-0.046, 0.053]	0.0027 [-0.044, 0.050]
B) Years 0 to 3	0.0055 [-0.036, 0.047]	0.0047 [-0.035, 0.045]
C) Years 4 to 7	-0.0020 [-0.016, 0.012]	-0.00020 [-0.015, 0.011]
Panel B: ≤30 employees		
D) Years 0 to 7	-0.20 [-0.50, 0.11]	-0.20 [-0.50, 0.11]
E) Years 0 to 3	-0.17 [-0.43, 0.093]	-0.17 [-0.43, 0.093]
F) Years 4 to 7	-0.026 [-0.098, 0.046]	-0.026 [-0.098, 0.046]
Panel C: All		
G) Years 0 to 7	8.98 [-13.19, 31.14]	8.77 [-13.12, 30.67]
H) Years 0 to 3	7.62 [-7.80, 23.04]	7.57 [-7.77, 22.91]
I) Years 4 to 7	1.36 [-6.17, 8.88]	1.20 [-6.19, 8.59]
Prior patents	X	X
E[wins]		X

Notes: See notes to Table 4. The table is identical to Table 4, except that in Appendix Table 2 the dependent variable is the number of patents over Years 0 to 7 combined. When we consider the number of patents in the sample of firms of all sizes, the results are extremely imprecise, which is unsurprising since the standard deviation of patents in this sample is so large, and since an extra H-1B worker represents only a small fraction of mean employment in the full sample of firms. The positive point estimate in this context is very sensitive to outliers; for example, when we winsorize the number of patents at the 99th percentile in the sample of all firms, we obtain negative point estimates, but the estimates are similarly imprecise and insignificant. In the smaller firm size categories, the effects are far more precise and the 95 percent confidence intervals bound the maximum increase in patents at a low level. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Appendix Table 3. Effects of Unexpected H-1B Lottery Wins on Patenting over Different Time Horizons

	IHS of # patents		Patenting Dummy	
Panel A: ≤10 employees				
A) Years 0 to 3	-0.0025 [-0.023, 0.018]	-0.0032 [-0.023, 0.016]	-0.014 [-0.037, 0.0085]	-0.014 [-0.037, 0.0084]
B) Years 4 to 7	0.0017 [-0.0092, 0.013]	0.0015 [-0.0087, 0.012]	-0.00075 [-0.0093, 0.0078]	-0.00090 [-0.0091, 0.0073]
Panel B: ≤30 employees				
C) Years 0 to 3	-0.013 [-0.057, 0.031]	-0.013 [-0.057, 0.031]	0.0085 [-0.020, 0.037]	0.0086 [-0.020, 0.037]
D) Years 4 to 7	-0.0074 [-0.032, 0.017]	-0.0073 [-0.032, 0.018]	-0.0073 [-0.026, 0.011]	-0.0073 [-0.026, 0.011]
Panel C: All				
E) Years 0 to 3	-0.044 [-0.10, 0.016]	-0.042 [-0.098, 0.015]	-0.0061 [-0.032, 0.019]	-0.0058 [-0.031, 0.020]
F) Years 4 to 7	-0.014 [-0.064, 0.037]	-0.014 [-0.064, 0.036]	-0.0014 [-0.024, 0.022]	-0.0016 [-0.025, 0.022]
Prior patents	X	X	X	X
E[wins]		X		X

Notes: The table shows the effect of an extra H-1B visa on patent outcomes over the indicated years. The table is identical to Table 4, except that the dependent variable measured patents over Years 0 to 3 (rows A, C, and E) or Years 4 to 7 (rows B, D, and F), rather than over Years 0 to 7 as in Table 4. See Tables 1 and 4 for additional notes and sample sizes. Standard errors are clustered by firm. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Appendix Table 4. Effect of Unexpected H-1B Lottery Wins on Patenting, using Alternative Matching Procedure

	Inverse hyp. sine of # patents		# Patents		Patenting Dummy	
Panel A: ≤10 employees						
A) Years 0 to 7	-0.0023 [-0.025, 0.021]	-0.0027 [-0.025, 0.020]	0.049 [-0.049, 0.15]	0.047 [-0.047, 0.14]	-0.016 [-0.039, 0.0075]	-0.016 [-0.038, 0.0075]
B) Years 0 to 3	-0.00074 [-0.021, 0.020]	-0.0012 [-0.021, 0.019]	0.052 [-0.043, 0.15]	0.050 [-0.042, 0.14]	-0.017 [-0.040, 0.0064]	-0.016 [-0.039, 0.0064]
C) Years 4 to 7	0.0024 [-0.0091, 0.014]	0.0023 [-0.0085, 0.013]	-0.0029 [-0.017, 0.011]	-0.0027 [-0.016, 0.011]	-0.0090 [-0.010, 0.0085]	-0.00097 [-0.010, 0.0080]
Panel B: ≤30 employees						
D) Years 0 to 7	-0.018 [-0.064, 0.028]	-0.018 [-0.063, 0.028]	-0.14 [-0.48, 0.20]	-0.14 [-0.48, 0.20]	0.0066 [-0.022, 0.035]	0.0068 [-0.022, 0.035]
E) Years 0 to 3	-0.015 [-0.059, 0.029]	-0.015 [-0.059, 0.029]	-0.13 [-0.41, 0.14]	-0.13 [-0.41, 0.14]	0.0057 [-0.023, 0.034]	0.0059 [-0.023, 0.035]
F) Years 4 to 7	-0.0080 [-0.033, 0.017]	-0.0078 [-0.033, 0.017]	-0.0046 [-0.089, 0.080]	-0.0049 [-0.089, 0.079]	-0.0086 [-0.027, 0.010]	-0.0085 [-0.027, 0.010]
Panel C: All						
G) Years 0 to 7	-0.078 [-0.16, 0.0015]*	-0.076 [-0.15, 0.000015]*	3.05 [-21.26, 27.36]	2.74 [-21.35, 26.83]	-0.0082 [-0.035, 0.018]	-0.0082 [-0.035, 0.018]
H) Years 0 to 3	-0.066 [-0.14, 0.0074]*	-0.063 [-0.13, 0.0061]	4.29 [-11.96, 20.54]	4.18 [-11.99, 20.35]	-0.011 [-0.037, 0.015]	-0.011 [-0.037, 0.015]
I) Years 4 to 7	-0.033 [-0.091, 0.026]	-0.033 [-0.091, 0.025]	-1.24 [-10.17, 7.69]	-1.44 [-10.30, 7.41]	-0.0041 [-0.029, 0.020]	-0.0042 [-0.029, 0.021]
Prior patents	X	X	X	X	X	X
E[wins]		X		X		X

Notes: See notes to Table 4. The table is identical to Table 4, except that in defining which firms match between the USCIS data and the Patent Dataverse, Appendix Table 4 includes those firms that are “possible” matches (whereas Table 4 excludes those firms). Also, Appendix Table 4 includes information on the effect on number of patents. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.

Appendix Table 5. Employment regressions by quarter in Q1 to Q4

	Median Regressions		Two-stage least squares	
	(1)	(2)	(3)	(4)
Panel A: ≤10 employees (n=712)				
A) Q1	-0.00 [-1.28, 1.28]	-0.031 [-1.64, 1.58]	0.072 [-1.24, 1.39]	-0.15 [-2.15, 1.86]
B) Q2	-0.00 [-0.68, 0.68]***	-0.41 [-1.17, 0.36]***	-0.80 [-2.34, 0.75]**	-1.46 [-3.29, 0.36]***
C) Q3	-0.78 [-1.78, 0.23]***	-0.53 [-1.42, 0.36]***	-0.66 [-2.40, 1.08]*	-1.33 [-3.47, 0.80]**
D) Q4	-0.76 [-2.05, 0.51]***	-0.61 [-1.79, 0.57]***	-0.90 [-3.12, 1.31]*	-1.72 [-4.52, 1.08]*
Panel B: ≤30 employees (n=1,183)				
E) Q1	-0.35 [-1.41, 0.72]***	-0.32 [-1.38, 0.73]**	-1.05 [-3.17, 1.06]*	-1.31 [-3.47, 0.85]**
F) Q2	-0.22 [-1.08, 0.65]***	-0.17 [-1.11, 0.78]**	-0.73 [-2.57, 1.10]*	-0.95 [-2.90, 1.00]*
G) Q3	-0.95 [-2.17, 0.27]***	-0.76 [-1.83, 0.31]***	-1.00 [-3.23, 1.23]*	-1.33 [-3.62, 0.96]**
H) Q4	-0.53 [-1.82, 0.76]***	-0.53 [-1.85, 0.79]**	-0.92 [-3.51, 1.67]	-1.25 [-3.99, 1.49]
Panel C: All (n=2,281)				
I) Q1	-1.41 [-3.40, 0.58]***	-1.67 [-3.89, 0.54]**	-62.10 [-768.40, 644.19]	-9.40 [-22.73, 3.92]
J) Q2	-1.35 [-3.72, 1.02]*	-1.00 [-3.11, 1.12]*	-17.32 [-180.09, 145.44]	-2.75 [-18.09, 12.58]
K) Q3	-0.055 [-3.15, 3.03]	0.25 [-2.33, 2.83]	4.76 [-72.71, 82.24]	4.43 [-15.97, 24.83]
L) Q4	1.36 [-4.80, 2.07]	-0.31 [-3.64, 3.01]	-13.70 [-191.01, 163.60]	0.04 [-21.57, 21.64]
Prior employment	X	X	X	X
E[wins]		X		X

Notes: See Table 7. *** denotes estimates that are significantly different from 1 at the 1% level; ** at the 5% level; * at the 10% level. n's refer to the total number of firms (which is identical to the total number of observations). None of the estimates is significantly different from 0 at any conventional significance level.

Appendix Table 6. Additional employment specifications

	(1) Winsorize at 99%	(2) IHS	(3) IHS of difference, winsorized at 99%	(4) IHS of level, winsorized at 99%	(5) First difference of employment, no controls
A) ≤ 10 employees	-1.86 [-4.34, 0.62]**	-0.18 [-0.43, 0.066]**	-0.18 [-0.43, 0.067]**	-0.18 [-0.42, 0.068]**	-0.53 [-1.37, 0.31]***
B) ≤ 30 employees	-1.69 [-4.55, 1.17]*	-0.16 [-0.35, 0.035]*	-0.15 [-0.34, 0.034]**	-0.16 [-0.35, 0.037]**	-0.69 [-1.68, 0.31]***
C) All	1.06 [-73.91, 76.03]	0.034 [-0.15, 0.22]	0.045 [-0.14, 0.23]	0.032 [-0.14, 0.21]	-1.07 [-3.05, 0.92]**

Notes: Columns 1-4 of the table show the baseline two-stage least squares regressions of employment outcomes on approved H-1B visas, where unexpected lottery wins are the instrument for approved H-1B visas. (The corresponding “reduced form” OLS regressions show very similar results.) In Column 1, the dependent variable is the difference of employment from the first quarter of Year -1 to employment in Q1, Q2, Q3, or Q4 (pooled), and winsorized at the 1st and 99th percentiles. The 1st and 99th percentiles of the first difference in employment are -5,559 and 2,430, respectively, in the full sample; are -20 and 62, respectively, among those with 30 or fewer employees; and are -10 and 53, respectively, among those with 10 or fewer employees. In Column 2, the dependent variable is the IHS of the difference in employment over the same periods. In Column 3, the dependent variable is the IHS of the difference in employment over the same periods, winsorized at the 99th percentile. In Column 4, the dependent variable is the IHS of the level of employment in Q1 through Q4 (pooled), winsorized at the 99th percentile, and the results are nearly identical to those in Column 3. All specifications in Columns 1, 2, 3, and 4 control for prior employment and the number of expected lottery wins, as in the baseline; the results are similar with other controls. In Column 5, we run median regressions (as in Table 7) and the dependent variable is the first difference of employment (from the first quarter of calendar Year -1 to a given quarter of Year 0, and pooling this measure from Q1 to Q4), but we do not include any controls. In all columns, we pool across Q1 to Q4, as in the baseline; the results are comparable (though typically slightly less precise) when we examine each quarter separately. None of the estimates is significantly different from 0 at any conventional significance level. In the case of these IHS specifications, before testing whether a coefficient is equal to 1, we transform the coefficient from the regression (which reflects the percentage increase in employment, rather than the increase in the absolute level of employment) by multiplying it by the mean level of employment. We then test whether this transformed coefficient is equal to 1. The test results that are reported above refer to this test. *** denotes estimates that are significantly different *from 1* at the 1% level; ** at the 5% level; * at the 10% level. See Table 7 for other notes and sample sizes.

Appendix Table 7. Effect of Unexpected Lottery Wins on Employment of Foreigners and non-Foreigners

Outcome	(1) All (n=2,143)	(2) ≤30 employees (n=1,198)	(3) ≤10 employees (n=723)
A) U.S. citizen employment, IRS measure	-0.012 [-0.41, 0.39]***	0.00 [-0.15, 0.15]***	0.00 [-0.19, 0.19]***
B) Non-U.S. citizen employment, IRS measure	-0.55 [-1.89, 0.79]***	-0.12 [-0.97, 0.72]***	-0.26 [-1.14, 0.62]***
C) Native employment, SSN-based measure	-0.073 [-0.72, 0.58]***	0.11 [-0.47, 0.69]***	0.018 [-0.41, 0.44]
D) Non-native employment, SSN-based measure	-0.37 [-1.32, 0.59]***	-0.065 [-0.80, 0.67]***	-0.16 [-1.34, 1.03]*

Notes: The table shows the effect of unexpected lottery wins on employment of foreigners or non-foreigners, displaying point estimates of the coefficient on unexpected lottery wins and 95% confidence intervals from median regressions. “IRS measure” refers to a specification in which we measure employment using current IRS data on the most recent measure of citizenship (the only measure of citizenship immediately available in the data). “SSN-based measure” refers to a measure of nativity using an algorithm developed in conjunction with Yagan (2014), identifying individuals as natives and non-natives on the basis of individuals’ Social Security Numbers (SSNs) in the data. The table shows that the results are similar under both measures. All specifications control for employment in the pre-period and expected lottery wins, as in the baseline. The measure of a firm’s employment is taken from the end-of-year (December) “snapshot” of employment from the W-2. The results are similar when we measure employment as the total number of employees observed at the firm over the year from the W-2 data. To make the time period investigated as comparable as possible to the quarterly data shown elsewhere (where we investigate Q1 to Q4), we pool the snapshot from Year 0 with the snapshot from Year 1. Results are similar when only using the snapshot from either of these years separately. n’s refer to the number of firms. See Table 7 for additional notes. For Rows A and C (regressions for non-foreigners), *** denotes estimates that are significantly different *from -1* at the 1% level; ** at the 5% level; * at the 10% level. For Rows B and D (regressions for foreigners), the number of stars instead denotes the significance test for difference *from 1*. The reason for the difference is that in the case of foreigners, we are primarily interested in testing whether the additional H-1B crowds out employment of other foreigners—which corresponds to the test of a difference from 1 because if the H-1B works at the firm, the coefficient should be 1. In the case of non-foreigners, we are interested in testing whether the H-1B crowds out non-foreigners one-for-one—which corresponds to the test of whether the coefficient is different from -1. None of the estimates is significantly different from zero at any conventional significance level.

Appendix Table 8. Effect of H-1B Visa on Being out of Business

Panel A: ≤10 employees (n=719)		
A) Q1 to Q4	0.024 [-0.016, 0.063]	0.033 [-0.022, 0.088]
B) Q1	0.016 [-0.020, 0.052]	0.023 [-0.030, 0.077]
C) Q2	0.017 [-0.033, 0.066]	0.022 [-0.051, 0.095]
D) Q3	0.032 [-0.014, 0.079]	0.046 [-0.015, 0.11]
E) Q4	0.029 [-0.017, 0.076]	0.041 [-0.022, 0.10]
Panel B: ≤30 employees (n=1,134)		
F) Q1 to Q4	0.010 [-0.019, 0.040]	0.012 [-0.024, 0.047]
G) Q1	0.0033 [-0.028, 0.034]	0.0033 [-0.034, 0.040]
H) Q2	0.0030 [-0.035, 0.041]	0.0029 [-0.043, 0.049]
I) Q3	0.015 [-0.020, 0.050]	0.017 [-0.023, 0.058]
J) Q4	0.020 [-0.013, 0.052]	0.023 [-0.014, 0.060]
Panel C: All (n=2,292)		
L) Q1 to Q4	0.0050 [-0.068, 0.078]	0.0024 [-0.014, 0.019]
M) Q1	-0.032 [-0.39, 0.32]	-0.0053 [-0.022, 0.011]
O) Q2	-0.013 [-0.13, 0.11]	-0.0024 [-0.024, 0.019]
P) Q3	-0.015 [-0.10, 0.13]	0.0054 [-0.014, 0.025]
Q) Q4	0.037 [-0.21, 0.28]	-0.011 [-0.0084, 0.031]
Prior employment	X	X
E[wins]		X

Notes: The table shows point estimates and 95% confidence intervals on unexpected lottery wins, from OLS (linear probability) regressions a dummy for whether the firm is “out of business” is regressed on unexpected lottery wins and controls. We define a firm as being “out of business” if it has either zero employees or is missing the number of employees. The results are similar with other definitions of being out of business. The “prior employment” specifications control for employment from the first quarter of Year -1, and the “prior employment, E[wins]” specifications additionally control for the number of expected lottery wins. None of the estimates is significantly different from 0 at any conventional significance level. “n” refers to the total number of firms in the regressions. See Tables 1 and 7 for other notes. *** denotes estimates that are significantly different *from 1* at the 1% level; ** at the 5% level; * at the 10% level.

Appendix Table 9. Effect of H-1B Visa on Being out of Business

Panel A: ≤10 employees (n=719)		
A) Q5 to Q8	0.020 [-0.088, 0.13]	0.018 [-0.090, 0.13]
B) Q8 to Q12	0.016 [-0.020, 0.052]	0.023 [-0.030, 0.077]
C) Q13 to Q32	0.065 [-0.041, 0.17]	0.068 [-0.039, 0.17]
Panel B: ≤30 employees (n=1,191)		
D) Q5 to Q8	-0.014 [-0.081, 0.054]	-0.013 [-0.081, 0.054]
E) Q8 to Q12	-0.022 [-0.092, 0.048]	-0.022 [-0.092, 0.047]
F) Q13 to Q32	0.023 [-0.053, 0.099]	0.024 [-0.052, 0.10]
Panel C: All (n=2,289)		
G) Q5 to Q8	-0.00025 [-0.033, 0.033]	0.00092 [-0.032, 0.034]
H) Q8 to Q12	-0.015 [-0.053, 0.024]	-0.012 [-0.050, 0.026]
I) Q13 to Q32	-0.0097 [-0.052, 0.033]	-0.0079 [-0.050, 0.034]
Prior employment	X	X
E[wins]		X

Notes: The table shows point estimates and 95% confidence intervals on unexpected lottery wins, from OLS (linear probability) regressions a dummy for whether the firm is “out of business” is regressed on unexpected lottery wins and controls. We define a firm as being “out of business” if it has zero employees or is missing number of employees. The results are similar with other definitions of being out of business. The “prior employment” specifications control for employment from the first quarter of Year -1, and the “prior employment, E[wins]” specifications additionally control for the number of expected lottery wins. None of the estimates is significantly different from 0 at any conventional significance level. “n” refers to the total number of firms in the regressions. See Tables 1 and 7 for other notes. *** denotes estimates that are significantly different *from 1* at the 1% level; ** at the 5% level; * at the 10% level.

Appendix Table 10. *Effect of Unexpected Lottery Wins on Revenue or Net Income per Employee*

	(1) Revenue per Employee	(2) Net Income per Employee
(A) ≤ 10 employees (n=615)	8,376.40 [-6,483.59, 23,236.40]	6,191.85 [-9,414.77, 21,798.48]
B) ≤ 30 employees (n=1,033)	8,326.45 [-2,194.90, 18,845.80]	5,220.35 [-2,660.61, 13,101.31]
C) ≤ 200 employees (n=1,520)	2,600.74 [-1,985.04, 7,186.51]	2,730.51 [-1,426.81, 6,887.82]

Notes: The table shows median regressions of revenue per employee (Column 1) or net income per employee (Column 2) in Years 0 to 3 on unexpected H-1B visas and controls. The table shows coefficients and 95% confidence intervals on unexpected H-1B visas. In Row C we investigate firms with 200 or fewer employees because regressions above this firm size cutoff did not reliably converge; they did not converge, for example, in the sample of firms of all sizes. Years 0 to 3 cover the duration of the H-1B visa. Estimated effects in the shorter or longer term are comparable. We find no significant interactions of the effect with covariates, and no significant differences in the effect across groups. When the dependent variable is instead the first-difference in the dependent variable from the pre-period, the results are broadly similar. Standard errors are clustered by firm. n's refer to the number of firms. *** refers to significance at the 1% level; ** at the 5% level, and * at the 10% level.