The Effects of High-Skilled Immigration Policy on Firms: Evidence from Visa Lotteries

Kirk Doran

University of Notre Dame and Institute for Labor Economics

Alexander Gelber

University of California San Diego and National Bureau of Economic Research

Adam Isen

Office of Tax Analysis, US Department of the Treasury

We compare winning and losing firms in lotteries for H-1B visas, matching administrative data on these lotteries to administrative tax data on US firms and to approved US patents. Winning one additional H-1B visa crowds out about 1.5 other workers at the firm. Additional H-1Bs have insignificant and at most modest effects on firm innovation. More general evidence from the universe of US firms and the universe of H-1B visas using alternative estimation strategies is consistent with these results. Firms that hire H-1Bs grow faster and innovate more because they are different in other ways from firms that do not.

I. Introduction

A key issue in several fields of economics is how easily firms can substitute one type of worker for another. Substitutability has implications for firms'

We thank US Citizenship and Immigration Services for help with the H-1B lottery data. We thank Gita DeVaney and Sunil Vidhani for outstanding research assistance. We thank

Electronically published September 8, 2022

Journal of Political Economy, volume 130, number 10, October 2022.

https://doi.org/10.1086/720467

production functions, predicting the demand for labor, and the incidence of policies that affect the price or availability of different types of labor. However, it is rare to observe a setting in which the availability of one or more types of labor is truly random, allowing a sharp test of hypotheses.

In this paper, we explore what happens to firms that win lotteries for an important type of labor in the United States: highly skilled foreign workers in the H-1B visa program.[1] Firms often argue that H-1B workers have exceptional skills that they cannot otherwise obtain and that obtaining these unique skills is necessary for them to grow and innovate. Others argue that H-1Bs have skills that firms could otherwise obtain and thus that additional H-1Bs would generally crowd out other workers and have more muted effects on firm outcomes such as innovation. While firms that hire H-1B visa workers grow faster and innovate more than other firms, a key question is whether this is due to the causal effect of H-1Bs on the firms or because of other differences. We answer this question by using administrative data on the entrants in these lotteries, matched to their tax filings and patenting. We find that winning a lottery for a skilled foreign worker crowds out otherwise available workers and does not increase firm patenting.[2] Winning firms neither grow faster nor innovate more than losing firms.

Notre Dame, the University of California Berkeley Institute for Research on Labor and Employment, and the Wharton School of the University of Pennsylvania for research support. We are grateful to George Borjas, John Bound, Sean Farhang, Richard Freeman, Hilary Hoynes, Jenny Hunt, Ben Jones, Damon Jones, Larry Katz, Bill Kerr, Norman Matloff, Ankur Patel, Dina Pomeranz, Jesse Rothstein, Fabian Waldinger, and seminar participants at the Fed Board, Harvard Business School, the Labor and Employment Relations Association, the National Bureau of Economic Research, the Center of Excellence for Migration and Integration Research, the US Treasury, and the Workshop on Innovation in the Global Economy for helpful comments. We thank 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 and the authors of Bell et al. (2019) for sharing their data. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the US Treasury Department, or of any other person associated with the US Treasury Department. All errors are our own. Replication information is available in a zip file. This paper was edited by Michael Greenstone.

[1] In 2010, immigrants accounted for 16% of the US adult population with at least a bachelor's degree, and high-skilled immigrants represented 24% of workers in occupations closely tied to innovation (Pekkala Kerr, Kerr, and Lincoln 2015). In recent years, prominent voices from government, business, labor, and academia have discussed significant changes to US immigration law. Many proposals have envisioned changes to the largest US high-skilled immigration program: H-1B visas for temporary immigration, which allow US firms to employ foreign workers for 3 years. The path of high-skilled immigration into the United States is unusual by international standards: the H-1B program is built around written requests from individual firms for access to specific workers with ostensibly unique skills. How H-1B workers affect the firms that have applied for them is the subject of much public discussion but little empirical work.

[2] We also find that it decreases the probability of hiring contractors and has no effect on the amount of the research and experimentation (R&E) credit claimed.

This result applies more generally than the setting of these lotteries. Regression discontinuity estimates based on application submission dates that compare firms that just did and did not miss out on being rationed, estimates based on the winners of large lotteries compared to a proxy control group, and shift-share estimates based on variation in visa caps all reveal that H-1B visa workers crowd out otherwise available workers at the firm level, leaving measures of firm innovation roughly constant. Likewise, adding firm fixed effects to naïve regressions of firm employment on new firm H-1B visas among the universe of US firms produces similar results.

In particular, lottery point estimates show that winning one additional H-1B visa worker crowds out approximately 1.5 otherwise available workers. Regression discontinuity point estimates show crowd-out of roughly 1.5 other workers, point estimates from the winners of a larger quasi lottery show crowd-out of about 0.7 other workers, shift-share point estimates show crowd-out of 1.25 other workers, and firm fixed effect point estimates show crowd-out of 0.9 other workers.[3]

Overall, our results are not supportive of the narrative that H-1Bs have unique skills that firms cannot otherwise obtain, in contrast to what is suggested by lawmakers.[4] Rather, they are more supportive of a narrative in which marginal H-1Bs crowd out other workers, are paid less than alternative workers, and increase the firm's profits—despite having little effect on measures of firm innovation. Indeed, we find some evidence from our lottery variation that additional H-1Bs increase profits and some evidence that additional H-1Bs decrease payroll costs per employee.

Several advances over the previous literature make these analyses possible, including administrative H-1B data on both lottery winners and losers and linkage of this and other H-1B data to the tax filings and patent data for the universe of US firms. Thus, relative to other studies on H-1B visas and other immigration programs, ours is the only one, to our knowledge, to leverage randomized variation or a discontinuity to estimate the effect of immigration on outcomes in the receiving economy as well as to focus on the universe of affected US firms. Our paper relates to previous

[3] Note that we generally cannot reject that any particular strategy's estimates are different from another's.

[4] The Senate Judiciary Committee reports that accompanied legislation to expand the H-1B program in 1998 and 2000 exemplify the narrative in which H-1Bs help firms address "shortages" of special skills. The former report noted that "companies across America are faced with severe high skill labor shortages that threaten their competitiveness" (Senate Judiciary Committee 1998, sec. 2 of amendment). The latter states, "America faces a serious dilemma when employers find that they cannot grow, innovate, and compete in global markets without increased access to skilled personnel. . . . [E]ven apart from shortages in particular fields, in our increasingly global economy, highly skilled foreign workers are certain to be in a position to make unique contributions to the U.S. economy. A person from another country may simply be a uniquely talented individual with unique knowledge and skills. . . . The country needs to increase its access to skilled personnel immediately in order to prevent current needs from going unfilled" (Senate Judiciary Committee 2000, sec. I).

work on the effects of immigration on the labor market (e.g., Card 1990, 2001; Borjas, Freeman, and Katz 1997; Friedberg 2001; Borjas 2003; Edin, Fredriksson, and Åslund 2003; Lubotsky 2007; Borjas, Grogger, and Hanson 2012; see surveys in Borjas 1994; Friedberg and Hunt 1995; Freeman 2006; Dustmann, Glitz, and Frattini 2008; Hanson 2009; and Pekkala Kerr and Kerr 2011) as well as on measures of innovation (e.g., Borjas and Doran 2012; Foley and Kerr 2013; Moser, Voena, and Waldinger 2014; Grogger and Hanson 2015; see the Kerr 2013 survey). Previous studies on the labor market or innovation impacts of the H-1B program or similar programs include Hunt and Gauthier-Loiselle (2010), Kerr and Lincoln (2010), Hunt (2011), Peri, Shih, and Sparber (2013), Bound et al. (2015), and Pekkala Kerr, Kerr, and Lincoln (2015). Regression analysis in the literature has found no clear evidence of crowd-out of other employment and in some cases has found crowd-in.[5] The literature has also found that H-1Bs lead to positive effects on patenting.

The paper is structured as follows. Section II describes the policy environment. Section III describes the data. Section IV discusses simple statistics and overall comparisons among firms. Section V discusses our empirical specifications. Section VI describes the full set of results, and section VII concludes.

II. Policy Environment

H-1Bs are sponsored by firms, which apply to the US Citizenship and Immigration Services (USCIS) and pay a fee to obtain a visa for each H-1B worker they wish to hire. In its application for each visa, a firm must specify the identity of the worker it wishes to hire. An H-1B visa allows a skilled foreigner to enter the United States for 3 years. The H-1B is considered a "nonimmigrant" visa because it allows those with H-1Bs to stay in the United States only temporarily. After these 3 years, the worker may leave the United States or a firm may seek to renew the worker's H-1B visa. Firms may also sponsor the worker to be a permanent resident. The H-1B worker may also move to another firm before the initial 3 years are up, though several frictions pose barriers to a move: the new firm must pay USCIS application and legal fees; upon moving, an H-1B goes to the "back of the line" for gaining permanent residency; some H-1Bs may not know that they can change jobs; and in the years we study, the

[5] Kerr and Lincoln (2010) find no evidence that H-1Bs crowd out other workers. Pekkala Kerr, Kerr, and Lincoln (2015) find mixed evidence on the effect of H-1Bs on total firm size. Peri, Shih, and Sparber (2013) find that H-1Bs increase native employment. However, the simulations of Bound et al. (2015) show that the ability to hire foreign computer scientists should reduce equilibrium employment and wages of natives while increasing equilibrium aggregate employment and output.

worker had to wait for several months until the new firm's H-1B application was approved, but a gap of only 2 weeks was allowed between jobs.[6]

The firm submitting the H-1B application (Form I-129) 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 nonimmigrants does not adversely affect working conditions of workers similarly employed in the area of intended employment." Firms are required to pay H-1Bs comparably with workers in one of four skill categories (defined by experience, education, and level of supervision).[7]

A. *Caps and Lotteries on H-1Bs*

The total number of H-1B visas awarded to for-profit firms in a given year is subject to a maximum number, or "cap." This cap is different for visas given to workers who have a master's degree or higher from a US institution (the "advanced degree exemption" [ADE] H-1B visa) and those without such a degree (the "regular" H-1B visa). In recent years, including the years for which we have lottery data, the cap for ADE visas has been 20,000 and the cap for regular visas has been 65,000. Much of the literature has identified the effects of H-1Bs through variation in these caps over time, especially the large increase in the caps in 1999, followed by a large decrease in 2004.

The number of H-1B visa applications in any given year has not always exceeded the cap, but in the two lottery years we study, the cap was reached for each of the two types of H-1B visa. Visa applications were accepted on a rolling basis once the application season began, and USCIS allocated visas by lottery only for applications submitted on the date when the total number of applications received exceeded the remaining available slots. 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.

Our main strategy makes use of lotteries for H-1B visas in fiscal year (FY) 2006 and FY2007. (While lotteries were run in subsequent years, USCIS did not keep data on which firms won or lost the lottery.) The caps for

[6] Depew, Norlander, and Sorensen (2013) study a single multinational information technology firm and find that from 2003 to 2011, 22% of its H-1Bs quit and moved to another firm while on the H-1B visa.

[7] Employers who are "H-1B dependent"—whose workforce has a sufficiently large fraction of H-1B employees—face additional requirements to attempt to recruit, and not displace, US workers. Firms may legally hire an H-1B in lieu of a worker who would have been at a higher skill level.

the FY2006 regular visas, the FY2006 ADE visas, the FY2007 regular visas, and the FY2007 ADE visas were reached on August 10, 2005, January 17, 2006, May 26, 2006, and July 26, 2006, respectively. That those were the dates the cap was reached was known only ex post—making it effectively impossible for firms to game the system by applying on the (unknown) lottery date for more visas than they desired, on the basis of the (unknown) probability of selection. Even across the four lotteries we study, the probability that an application won varied widely. Indeed, these were the first two years USCIS used a lottery to allocate H-1Bs, and it was not announced in advance that lotteries were going to be run.[8] Each lottery was conducted within a month of reaching the relevant cap.

If a firm is denied a capped H-1B, it has several alternatives to hiring no one. Other than hiring US citizens or foreigners who are permanent residents, firms can hire foreigners on other visas, including L-1 temporary work visas, optional practical training (OPT) extensions of F-1 student visas, or H-1Bs not subject to the cap. L-1s allow multinational firms to bring a worker at a foreign branch to the United States temporarily. Visa lottery losers would likely not resort to bringing the same worker to the United States on an L-1, since a firm would have typically applied for an L-1 rather than an H-1B if the L-1 were feasible (as the L-1 is more advantageous to the firm than the H-1B). Only 11% of lottery participants are multinationals, further limiting the importance of the L-1 in our context. In FY2006 and FY2007, OPT extensions allowed F-1s to extend their stays for only 12 months, limiting substitutability with H-1Bs, and the majority of H-1B applicants are not eligible for an F-1 visa.

B. *What Job Tasks Do H-1Bs Do?*

Before describing the data, it is worthwhile to consider what tasks H-1Bs typically do, which can help contextualize the results that follow. Data on I-129 applications show that H-1B visa workers work in a variety of occupations doing a variety of tasks, from working as researchers to working in tech support call centers. However, despite this diversity, the majority of I-129 applications from firms list one occupation: "systems analysis and programming."[9]

We break down this broad occupation category into specific occupations that reveal worker tasks by using the Department of Labor's 2006 Occupational Employment Statistics by state, containing mean annual

[8] One to two weeks before each lottery, USCIS publicly announced the number of applications it had received. Thus, firms may have been able to anticipate approximately when the cap might be reached, but they could not reasonably predict the exact day it would be reached or probability of selection on this day.

[9] While this occupation constitutes only a plurality among just ADE applications, those H-1B applications total less than one-quarter of all capped H-1B slots.

earnings for the 16 suboccupations within the “computer and mathematical occupations” category. The mean annual earnings of these different suboccupations vary widely within each state. Thus, the 2006 salary and state of a person in the computer occupations can determine whether they are being paid like someone in a creative research role or like a worker in tech support.

We report in appendix table 1 differences between the mean annual salaries in each subcategory and the actual annual salaries paid to regular H-1Bs in “systems analysis and programming,” reported in their I-129s. The occupations with the closest annual salaries to H-1Bs, adjusted by time and place, are not the creative professions of software engineers or even the higher-level technical professions such as programmers. Rather, they are lower-level supporting jobs, such as computer support specialists, technicians, and systems administrators. The wages are so much lower for these H-1Bs than the time-and-place-adjusted wages for creative professions that it is very unlikely that many of these H-1Bs are doing creative work, at least if prevailing-wage restrictions are met.

III. Data

This paper combines, for the first time, data on the universe of US firms with data on H-1B immigrants at the firm level. It combines these data with the universe of US patents to paint a thorough picture of highly skilled immigrants and their output at US firms.

A. *Data Sources*

1. Internal Revenue Service (IRS) Data

We use IRS tax data on the universe of US firms. The holder of every unique Employer Identification Number (EIN) that employs workers must generally report firm employment on a quarterly basis (Form 941). We also rely on additional tax forms for other outcomes, including business income tax returns (1120, 1120S, and 1065), W-2s, and 1099-MISC.

2. USCIS Data

We use USCIS administrative data on the H-1B submissions for FY2006 and FY2007. The data contain the following information on each H-1B visa application in each of these years: EIN and name of the firm applying, the date the firm applied for a visa, the type of H-1B (regular or ADE), how many of each firm’s applications won or lost the lottery, whether each application was approved by USCIS, and firm-reported worker characteristics from the I-129, such as highest degree completed.

These data are used for our main lottery analysis as well as our regression discontinuity (RD) design.

For the remaining ancillary identification strategies, we use USCIS data on the number of newly approved capped H-1Bs at each firm in the United States for each FY from 2003 through 2008 (not including renewals). In each FY, these data contain approximately 30,000 firm names (but not EINs), each with an associated number of approved H-1Bs for that FY.

3. Patenting Data

We obtained the Patent Dataverse on the universe of granted US patent applications from 1975 to 2013 at each firm (identified by name) based on US Patent and Trade Office (USPTO) data organized by year of application.[10] We also observe total patent citations. For ancillary identification strategies, we use the data in Bell et al. (2019) that link individuals in the IRS data with those listed on patent applications.

B. Match between USCIS, IRS, and Patenting Data

Using EINs, we merged firms from the USCIS lottery data to their IRS records.[11] The IRS data give us firm-level quarterly employment and annual net income ("profit"), wage bill, and R&E credit for research and development expenses. For the patent merge, as explained in appendix 1, we performed an intentionally liberal automatic string-matching procedure between the USCIS lottery and patenting data sets to obtain all plausible matches between firms and patents. We then searched through the matches by hand to detect and remove all matches that appeared spurious.

For estimation strategies that rely on USCIS data on the number of newly approved H-1Bs at each firm for FYs 2003–8, we match them to IRS data, using a fuzzy match on firm names. (Since there are approximately 30,000 separate firms for each FY, we must apply a fuzzy matching technique that does not involve any by-hand work, as explained in app. 1.) The resulting matched data contain over 10 million for-profit firms (the

[10] See https://dataverse.harvard.edu/dataverse/patent. Granted patents are classified by the calendar year a firm applied for the patent. For example, our measure of the number of patents at a firm in year 0 refers to patents the firm applied for in year 0 that were approved by 2013.

[11] We drop the 2.0% of firms in the USCIS data that did not match to the IRS EIN master list. Of the remaining firms, 4.5% did not match to the quarterly firm employment data; we treat these data as missing for employment analyses. Of the rest, 17.9% have missing employment data in year −1. We try two options for these firms: first, dropping these data for the purpose of the employment specifications, and second, using a dummy variable to indicate missing preperiod employment and assigning mean preperiod employment to firms with missing data. We separately test for balance across each selection point.

universe of for-profit firms in the United States), of which 68,092 firms have at least one H-1B between the years 2003 and 2008.[12] We also construct a match between the universe of firms in the IRS data and the universe of US patenting data. We start with an existing link between individuals on approved patent applications and their taxpayer IDs (Bell et al. 2019) and then use W-2 data to determine whether each individual was an employee of a firm in the year a patent application was submitted and, if so, at which firm the individual was working. This allows us to infer patent counts by firm-year. We explain in further detail in appendix 1.

IV. Simple Statistics and Comparisons

A. Summary Statistics and Sample

Table 1 shows summary statistics for the universe of US firms. It is clear that firms that have ever had an H-1B have higher employment, experience larger increases in employment each year, patent more, have higher profits, and pay their workers higher salaries. It is possible that these differences in outcomes are due to positive causal effects of H-1Bs on firms. It is also possible that these differences in outcomes are artifacts of other differences between the types of firms that happen to hire H-1Bs and the types of firms that do not. Our lottery identification strategy and ancillary identification strategies can address this question.

Column 3 of table 1 shows summary statistics for the FY2006 and FY2007 lottery sample. We use data on 2,750 firms (i.e., EINs) in the full lottery sample. In 300 cases (9.84%), firms apply for at least one visa in both FY2006 and FY2007. Thus, over both lottery years, there are 3,050 firm–lottery year observations, where "year" refers to a year of the lottery, rather than a year when an outcome is observed. The mean and standard deviation of the number of employees during Q1–Q4 in the full sample are very large. In firms with 30 or fewer, or 10 or fewer, employees in year −1 (two representative cutoffs we use), the mean and standard deviation of Q1–Q4 employment are lower but still large. Median employment is lower than mean. Winsorizing also reduces the mean and standard deviation.

In the FY2006 regular lottery, the vast majority of applications lost the lottery, and in the FY2007 regular lottery, the vast majority won. 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 our estimates; such effects on precision will be reflected in the confidence intervals.

[12] To make the regressions tractable, we restrict to a 1% random sample of firms without an H-1B between 2003 and 2008, assigning them a weight of 100.

TABLE 1
SUMMARY STATISTICS FOR FULL MERGED IRS DATA

Variable	All Firms, Mean (SD) (1)	H-1B Firms, Mean (SD) (2)	Lottery Firms, Mean (SD) (3)
No. of employees:			
All	22.23 (827.55)	422.88 (7,503.84)	1,877.84 (39,721.31)
≤30	6.09 (11.36)	12.45 (94.91)	43.09 (1,904.34)
≤10	3.78 (9.32)	7.11 (105.75)	9.64 (55.63)
Median employees:			
All	4	20	31
≤30	4	8	10
≤10	3	5	6
Winsorized employment first difference:			
All	.21 (3.57)	3.20 (14.66)	27.28 (92.39)
≤30	.13 (1.58)	1.01 (3.04)	4.35 (9.43)
≤10	.12 (1.13)	.75 (1.96)	3.22 (6.84)
No. of patents:			
All	.070 (5.15)	2.17 (39.38)	4.52 (56.11)
≤30	.011 (.33)	.17 (1.93)	.23 (8.59)
≤10	.007 (.26)	.10 (1.27)	.023 (.49)
IHS of patents:			
All	.012 (.173)	.23 (.81)	.15 (.80)
≤30	.005 (.099)	.06 (.37)	.017 (.22)
≤10	.003 (.081)	.04 (.28)	.010 (.14)
IHS of R&E:			
All	.049 (.788)	.70 (2.98)	1.55 (4.74)
≤30	.014 (.399)	.24 (1.69)	.15 (1.39)
≤10	.007 (.290)	.14 (1.31)	.14 (1.22)
Fraction with R&E:			
All	.004 (.063)	.054 (.225)	.099 (.30)
≤30	.001 (.036)	.020 (.141)	.013 (.11)
≤10	.001 (.026)	.012 (.110)	.013 (.11)
Median payroll per employee ($):			
All	16,800	37,944	49,332
≤30	16,433	33,930	42,281
≤10	16,049	31,075	38,657
Median firm profits ($):			
All	11,308	45,844	
≤200			80,250
≤30	9,822	27,928	43,301
≤10	8,196	20,454	30,398
New H-1Bs as share of employees:			
All	.00079 (.08)	.075 (.75)	
≤30	.00081 (.08)	.12 (.95)	
≤10	.00085 (.09)	.16 (1.19)	
Fraction winning lottery:			
2006 regular			.04
2006 ADE			.17
2007 regular			.98
2007 ADE			.55

TABLE 1 (*Continued*)

Variable	All Firms, Mean (SD) (1)	H-1B Firms, Mean (SD) (2)	Lottery Firms, Mean (SD) (3)
Fraction in NAICS code 54:			
All	.11 (.31)	.32 (.47)	.56
≤30	.11 (.32)	.36 (.48)	.66
≤10	.12 (.32)	.37 (.48)	.65

Note.—The data are from IRS, USCIS, and USPTO administrative sources. Column 1 ("All Firms") refers to the universe of firm-year observations in the IRS data. Column 2 ("H-1B Firms") refers to firms listed as hiring at least one new H-1B during FY2003–8. The "all" rows refer to firm-year observations of all employee sizes; the "≤30" ("≤10") rows refer to those firms with 30 (10) or fewer employees in a given year. Number of patents refers to approved patents in each year from year 0 to 2013. Employment data are observed in Q1–Q4, the first four quarters when the H-1B worker may work at the firm. R&E, payroll per employee, and firm profits are measured in years 0–3, the duration of the H-1B visa. We pool and stack time periods. For profits, we use the size category with ≤200 employees; our regressions did not converge for higher thresholds. NAICS code 54 is professional, scientific, and technical services. For R&E, the sample size is also smaller because the data measure only the R&E credit for C-corporations. The fraction patenting or with the R&E refers to the mean of a yearly patenting dummy in years 0–8 or to the mean of a yearly dummy for taking the R&E credit in years 0–3. Here and throughout the paper, dollar amounts (e.g., the R&E credit) are measured in real 2014 dollars.

The lottery 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 chance lottery wins (defined below) is 0.33, and its range runs from −2.65 to 2.96. Over half of firms are in North American Industry Classification System (NAICS) code 54, representing professional, scientific, and technical services. The H-1B application data show that across all lotteries, applicant average age is around 30 (app. table 2).

B. *Comparison of Lottery Firms to Other Firms That Applied before the Last Day*

Our primary identification strategy comes from comparing firms that randomly received H-1Bs to those that randomly did not. This comparison is comprised of firms that applied on the day the cap was reached, which addresses effects for marginally changing the number of H-1Bs allowed, a question of great relevance, as visa cap changes are contemplated. Nonetheless, it is worthwhile to compare this sample to the broader sample of firms applying for H-1B visas in these years. In appendix table 3, we regress characteristics of the firms or workers on a dummy for applying on the last day and lottery fixed effects. Applications on the last day tend to be from larger firms and those that are more likely to be in professional, scientific, and technical services industries and to have patented more in the past, compared to the set of firms that applied earlier. On the last day, firms disproportionately submit applications for workers with higher educational degrees, for those with higher intended worker salaries, for "systems analysis and programming" jobs, and for younger workers. If H-1Bs hypothetically have more positive innovation effects in firms that patented more in the past and/or are in scientific industries, or among workers with advanced degrees or higher salaries, then our sample will arguably be primed to find a positive effect on innovation relative to using the full set of H-1B firms. Likewise, in appendix table 5, we also show that firm growth rates are not inferior for firms applying on the last day or on a day approaching it. Results from other empirical strategies will speak more directly to the generalizability of the lottery experiment.

V. Empirical Strategies

Our main outcomes of interest are number of employees and patenting. We also consider the effect on the R&E credit, the firm's wage bill per employee, and profits. Below, we first discuss our primary empirical strategy, which exploits the random assignment of H-1B visas among firms in the FY2006 and FY2007 H-1B lotteries, before turning to secondary strategies.

Our lottery 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 applications won a visa and W is the random number of H-1B visas given to the firm, then the average number of H-1B visas given to the firm in expectation is $\mathrm{E}[W] = pn$. If w is the random realization of W, then the number of "chance lottery wins" or "chance visas," $u = w - pn$, is the random realization of the net number of wins relative to the ex ante statistical expectation conditional on p and n and will be exogenous in the regression we specify below. Thus, our main independent variable is the random variable U, the net number of a firm's chance lottery wins, which by construction has a mean of 0 and whose realization is u.

To find the causal effect of U on an outcome Y, we estimate versions of

$$Y_{itT} = \beta_0 + \beta_1 U_{iT} + \varepsilon_{itT}, \tag{1}$$

where t is the number of calendar years since the lottery in question occurred; for example, $t = 0$ corresponds to year 0.[13] The term T indexes the year of the lottery in question, that is, FY2006 or FY2007; U_{iT} is the

[13] For a given lottery year (i.e., FY2006 or FY2007), we refer to the calendar year 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," etc. 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 FY2006 lottery) as "Q1," the next quarter as "Q2," etc., which is relevant for employment, the only outcome we can observe quarterly.

number of chance H-1B visa lottery wins for firm i in the lottery in year T;[14] ε_{itT} is an error term. The term β_1 represents the intent-to-treat (ITT) effect of an additional chance H-1B visa win, which is relevant because firms and policy makers are interested in the raw effects on firms of allowing a marginal capped visa to the firm. To increase statistical power, we pool the regular and ADE lotteries for a given fiscal year T and stack data from FY2006 and FY2007 in the same regression.[15] We cluster the standard errors at the firm level.

After a firm wins an H-1B lottery, its application may be approved, denied, or withdrawn. The total number of capped H-1B visas approved for a firm in any given year is potentially endogenous because it depends on the fraction of those that win the lottery that are approved. In practice, most are approved. Still, we can use lottery wins as an instrument for approved capped H-1B visas in a two-stage least squares (2SLS) model:

$$A_{iT} = \alpha_0 + \alpha_1 U_{iT} + \nu_{iT}, \tag{2}$$

$$Y_{itT} = \gamma_0 + \gamma_1 A_{iT} + \eta_{itT}, \tag{3}$$

where A_{iT} represents the number of capped H-1B visas approved for firm i in the lottery that occurred in year T. In the first stage (eq. [2]), we regress A_{iT} on U_{iT} using ordinary least squares (OLS). In the second stage (eq. [3]), we regress Y_{itT} on A_{iT} (instrumented with U_{iT}) using OLS. The coefficient γ_1 represents the local average treatment effect of an extra approved capped H-1B visa among the compliers (i.e., those induced by winning the lottery to increase their number of approved capped H-1B visas); ν_{it} and η_{it} are error terms. This specification is most relevant in the employment context, where we are testing whether additional H-1Bs crowd in or crowd out other employment.

Our employment variable measures a firm's total employment, including any H-1B workers, so we adjust our employment coefficients by 1 to subtract the H-1B worker and present effects on the employment of other workers. The key question of interest is a two-sided test of whether the coefficient on H-1B visas is significantly different from 0, as theory is ambiguous about such a relationship. If the coefficient is positive and significant, it would indicate that the extra H-1B visa increases employment

[14] Although the randomization implies that U_i should be exogenous in eq. (1), it is also possible to control for various predetermined covariates. We generally control for a lagged value of an outcome variable at the firm (e.g., when the dependent variable is the number of employees, we can control for $Y_{i,\text{preperiod},T}$, the number of employees in firm i observed in the period preceding the lottery) and for the expected number of lottery wins pn.

[15] Similarly, whenever we examine an outcome across multiple time periods t, we stack the data across these periods. For the 2.69% of firms that participate 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 chance lottery wins across both of the lotteries that the firm enters in year T.

of other workers at the firm—as opposed to crowding out workers that the firm would have otherwise hired, in which case the coefficient would be negative. One-for-one crowd-out would yield a coefficient of −1. Finally, it is possible that the new H-1B worker works more hours or works harder than others (perhaps to secure another visa or green card for continued employment in the United States) and therefore crowds out more than one other worker, for a coefficient of less than −1.

We investigate the effects of H-1Bs on each outcome for the three-year duration of the initial visa for consistency, but we also focus on shorter and longer durations that are most appropriate for our main employment and patenting outcomes, respectively. In particular, we focus on quarterly employment from Q1 to Q4 of the visa, when the H-1B worker is almost always working at the firm and when a coefficient's relation to zero will therefore most reliably indicate crowd-in or crowd-out, and on patents from year 0 to the latest year available in the data, year 8, given the sometimes substantial time taken to develop patents.

To adjust for the long right tail of the employment distributions and the relatively small sample sizes in the lottery sample, we take two approaches. First, in our baseline specification, we run median regressions. Second, we run mean regressions where we either include firm-by-lottery fixed effects and winsorize the dependent variable at the 95% level or let the dependent variable be the winsorized first difference of employment. The first difference ΔY_{itT} is taken from before the lottery to period t after the lottery. Such regressions would not capture large effects on employment outcomes, but below we explore this issue in depth.

Because of the long right tail of the distribution of patents, previous literature has typically examined transformations of the number of patents. We approximate the log of the number of patents using the inverse hyperbolic sine (IHS), which is defined at zero and approximates the log for larger values of its arguments (e.g., Burbidge, Magee, and Robb 1988; Pence 2006). The IHS of patents Y is defined as

$$\mathrm{IHS}(Y) = \ln\left(Y + \sqrt{1 + Y^2}\right).$$

We alternatively deal with the distributional challenges posed by patents by using the negative binomial regression. This regression takes into account that patenting is a count variable.

We expect our results to be most statistically distinguishable in small and medium-sized firms, which, in the aggregate, contribute in important ways to US employment and innovation (Acs and Audretsch 1990) and comprise a substantial fraction of all H-1B lottery applicants (e.g., 19% of H-1B visas come from firms with 10 or fewer employees; 34% come from firms with 30 or fewer employees). To evaluate how the effects vary across firms of different sizes, we investigate firms with 10 or fewer employees in year −1 (roughly the 25th percentile of firm size in our sample), those with 30 or fewer employees in year −1 (roughly the 50th percentile), many other firm size cutoffs, and firms of all sizes.

We now turn to discuss our secondary strategies, which, while relying on stronger identification assumptions, complement our main approach by providing two main benefits. First, the alternative approaches provide substantial improvements in statistical power over our lottery results, affording us the ability to more regularly recover meaningful confidence intervals from unwinsorized mean regressions and for full samples of firms.[16] Second, they enable us to examine contexts that sometimes provide larger-scale variation affecting more H-1B applications.

Our first alternative approach adopts an RD design that exploits how the probability that a firm had a successful application fell from nearly 100% if submitted before the cap was reached to much lower if submitted on the day a cap was reached in FY2006 and FY2007. Importantly, it was not known in advance which days the caps were going to be reached, eliminating firms' ability to sort around the cutoff. We implement our RD approach using days relative to the cap being reached as the running variable and the optimal bandwidth of 20 days calculated by the method of Calonico, Cattaneo, and Titiunik (2014). Specifically, we tweak equations (2) and (3) by instrumenting for the number of H-1B visas approved with whether the applications were submitted before the day the cap was received and including a quadratic running variable to fit the relationship between the number of accepted H-1B applications and the number of days before the cap was reached and visa type–by-year fixed effects.[17]

We then consider the large FY2008 lottery, in which 123,480 H-1B applications were received in the first two days of the application season for only 85,000 capped slots. In this lottery, all applications were entered into the lottery, making this setting an interesting contrast to the moderately binding years of FY2006 and FY2007. The setting is difficult to analyze, however, because USCIS did not record information for the firms that lost the lottery.[18] Therefore, we must compare approved lottery winners with a control group constructed to proxy lottery losers. We construct this comparison group by determining the subset of Labor Condition Application (LCA) submissions that were likely to be eligible for submitting an H-1B application that would have been subject to the lottery (a prior LCA submission is required before an H-1B application can be submitted, but

[16] As a result, we exclusively present unwinsorized employment regressions for these strategies.

[17] Some firms submit applications on multiple days, but keeping firms that show up in both the treatment and control groups greatly attenuates any effect on visa approval. Thus, in order for there to be a first stage, we must limit our sample to firms that submitted all their applications for a particular visa category–year on one given day (76.4% of firms).

[18] Additionally, we can observe only approved winners.

it does not guarantee an H-1B application submission).[19] Ultimately, we cannot rule out that the results are influenced by attrition from the LCA sample into the lottery, nor can we determine which losing applications were in the regular and ADE lotteries, which had different selection probabilities.[20] With this proxy control group constructed and approved wins, we calculate "chance wins" and estimate equation (1).

We also utilize a shift-share strategy, which is one of the most common strategies employed when studying the impact of H-1B workers or migrants more generally (e.g., Pekkala Kerr, Kerr, and Lincoln 2015).[21] This strategy exploits cross-sectional reliance on H-1Bs and changes in the national H-1B cap over time, including those that occurred after 2003, when the cap decreased substantially, from 195,000 in 2003 to 65,000 in 2004, while modestly increasing to 85,000 for the remaining years for which we have H-1B data (2005–8). The shift-share strategy requires a stronger identification assumption, namely, that of parallel trends. To implement the shift-share strategy, we calculate the share of new H-1Bs as a function of lagged firm employment in the first year (2003) for which we have H-1B counts. Firms that relied on many new H-1Bs in 2003 as a share of the firm's 2002 overall employment were more likely to be affected by post-2003 visa cap changes than firms with a smaller share of H-1Bs. We run a first-stage regression in which a firm's new H-1Bs in a given year between 2003 and 2008 as a share of previous-year employment are regressed on the firm's 2003 H-1B share interacted with each year's annual visa cap. We also control for year and firm fixed effects and an interaction of a time trend with the preperiod value of each dependent variable.

Finally, we compare H-1B firms to a large random sample of the universe of remaining US firms. To do so, we augment equation (3) with firm fixed effects in which the independent variable of interest is the number of new capped H-1Bs approved for a given firm in a given year between 2003 and 2008, ranging from zero (most firms) to several thousand (the most heavy H-1B users).

[19] This consisted of all LCAs submitted by for-profit firms between March and April of 2007 in which the proposed H-1B worker would begin work between September and October of 2007. These restrictions are based on the USCIS rules that determine which I-129 applications would be subject to the FY2008 lottery, but we also tested and found robustness to variation in these restrictions, and we determine for-profit status through a merge to the IRS business entity database.

[20] The ratio of LCA submissions to H-1B applications was approximately 1.63, raising the possibility of large amounts of bias. Relatedly, Peri, Shih, and Sparber (2015) exploit these large lotteries at the metropolitan area level, which likely implies, as a result of the law of large numbers, that nearly all of the variation being exploited stems from attrition and differences in reliance on ADE vs. regular H-1Bs across areas.

[21] The literature on the impact of H-1Bs on firms has made use of smaller samples (such as portions of the COMPUSTAT database or large firms from census data). We are the first to analyze the results from this strategy applied to the universe of US firms.

VI. Results

A. Validity

Appendix table 6 verifies the validity of the randomized design by regressing variables that should not be affected by the lottery on chance lottery wins. The table confirms that none of the lagged dependent variables is significantly (or jointly) related to chance lottery wins.[22] Employee characteristics, as shown in appendix table 2, are also individually and jointly insignificantly related to lottery wins ($p = .31$ in the joint test).

B. Employment Results

Table 2 shows our main results from estimating equations (1)–(3). Across all specifications for small and medium-sized firms, the coefficients are negative and statistically significant, indicating that H-1B workers decrease the number of other workers employed at the firm. Column 1 of table 2 displays the β_1 coefficients from median regressions of the form outlined in equation (1), while columns 2 and 3 show the results from 2SLS mean regressions shown in equations (2) and (3).[23] Column 1 shows that for firms with 10 or fewer employees, an additional chance H-1B visa win crowds out approximately one-and-a-half other workers in the year after the lottery. The results are similar for firms with 30 or fewer employees and the full sample of firms. We then repeat these exercises for mean regressions with firm fixed effects (col. 2) and firm differences (col. 3) and similarly find evidence of crowd-out for firms with 10 or fewer or 30 or fewer employees, albeit with somewhat more negative but less precise

[22] We investigate the effects on year -2 outcomes and control for the dependent variable measured in year -1, which is the same control as in our regressions in later tables. By investigating year -2 outcomes, we can also determine the firm size cutoffs by measuring employment in year -1, yielding the same firms in each size category as in our later regressions. When we investigate year -1 outcomes as the dependent variable, controlling for year -2 observations and using firm size cutoffs from year -2, the regressions are insignificant for all but one of the 27 dependent variables, consistent with random chance (results not shown).

[23] Because instrumental variable quantile regressions typically did not converge, we use the ITT for median regression. These regressions do not reflect that some H-1B lottery winners' applications are rejected, but our first-stage coefficient presented in app. table 9 is extremely precise and quite close to 1 (ranging from 0.88 to 0.89), such that scaling our estimates by the first stage would not alter our findings. Of course, instrumental variable quantile regressions do not rely on a Wald estimate, but in practice, in the rare median instrumental variable median regressions that converged, the coefficients on approved H-1B visas were very similar to the ITT median coefficient divided by the OLS or median first stage—i.e., only around 10% larger than in the ITT median regressions. Additionally, after their visas are approved by USCIS, some workers may not show up for their jobs in the United States. However, North (2011) estimates that around the time we study, nearly all (95%) of those approved for H-1Bs ended up being admitted. Thus, (further) scaling our scaled ITT or treatment-on-the-treated estimates a bit more to reflect this would have negligible effects on our results and would not affect our conclusions.

TABLE 2
Effect of H-1B Lottery Wins on Employment of Other Workers

	Median Regressions (1)	2SLS Mean Regressions: With Firm Fixed Effects (2)	2SLS Mean Regressions: With Differences (3)
	Employment in First Year (Q1–Q4)		
A. ≤10 employees	−1.52***	−2.09**	−2.10**
	[−2.15, −.89]	[−3.85, −.33]	[−3.88, −.32]
B. ≤30 employees	−1.36***	−2.13**	−2.26**
	[−2.09, −.63]	[−4.21, −.05]	[−4.25, −.29]
C. All	−2.05**	−12.39	−3.41
	[−3.67, −.43]	[−65.60, 40.82]	[−18.76, 11.94]
	Employment in First 3 Years (Q1–Q12)		
A. ≤10 employees	−1.64***	−2.60*	−2.75**
	[−2.30, −.88]	[−5.38, .19]	[−5.42, −.07]
B. ≤30 employees	−1.74***	−2.56*	−2.98**
	[−2.61, −.87]	[−5.46, .34]	[−5.66, −.30]
C. All	−2.95***	−29.74	−11.00
	[−5.37, −.52]	[−98.57, 39.09]	[−31.68, 9.67]

Note.—The table shows coefficients on a chance H-1B lottery win, minus the H-1B visa worker the firm won in the lottery, with 95% confidence intervals in brackets. Column 1 shows median regressions of firm employment on chance lottery wins, defined as actual wins minus the expectation of wins conditional on number of applications and the probability each application wins. In the rare case that median regressions were unstable, we searched within a 0.5% grid to find consistently more stable estimates. Columns 2 and 3 show 2SLS (mean) regressions where the dependent variable in col. 2 is firm employment in the quarter of question winsorized at the 95th percentile (with firm-lottery fixed effects and the observations from the first quarter of year −1) and that in col. 3 is the difference of firm employment from the first quarter of year −1 to the quarter in question, winsorized at the 5th and 95th percentiles. We pool and stack observations across quarters. The first panel examines employment from Q1 to Q4, while the second panel examines employment from Q1 to Q12. All specifications control for employment in this preperiod and expected lottery wins (equal to number of H-1B applications entering a lottery multiplied by the probability of winning the lottery). The 5th and 95th percentiles of the first difference in employment are −109 and 352, respectively, in the full sample; −9 and 30, respectively, among those with 30 or fewer employees; and −6 and 22, respectively, among those with 10 or fewer employees. In these regressions, the instrument is chance lottery wins, and the endogenous variable is approved capped H-1B visas. The regressions include controls for employment from the first quarter of year −1 and the number of expected lottery wins. See table 1 for other notes and sample sizes. If the H-1B worker works at the firm, a coefficient of 0 corresponds to neither crowd-out nor crowd-in of other employment, and a coefficient of −1 corresponds to one-for-one crowd-out of other employment.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

estimates.[24] (Note that there is little difference between the firm fixed effects and differenced estimates.) However, the mean regression estimates for the full sample of firms are imprecise, which is unsurprising because of

[24] We find no evidence that winsorizing employment in mean regressions is biasing our findings. Namely, an extra H-1B visa has an insignificant effect on the probability that the

the high variance of large employers. Across all specifications, the full set of results remains similar whether we pool Q1–Q4 or Q1–Q12, with point estimates slightly increasing in absolute value in the longer time horizon.

To see more comprehensively how the estimates vary across the employer size distribution, we estimate our baseline specification from column 1 of table 2, varying the size threshold from 10 or fewer employees to 500 or fewer in increments of 10.[25] Figure 1*A* displays the results, showing that the results remain stable across the employer size distribution, with most point estimates hovering around −1.5, as seen in table 2. In all cases, the coefficients are statistically significant, ruling out no effect on the employment of other workers at the 1% level. Figure 2 displays the dynamics of these effects by augmenting equation (1) and allowing the impact of chance H-1B visa wins to vary by years since the lottery. Before the lottery, the effect of a chance H-1B visa win has a point estimate that hovers around zero, in line with the random assignment to treatment. Following the lottery, the impact on the employment of others is negative, consistent with crowding out. Appendix 2 describes additional robustness exercises,[26] and appendix 3 describes efforts to estimate the effect on employment of foreigners and nonforeigners separately. While there are limitations with this exercise, the results suggest that at least some of the crowd-out may be of other non-US citizens.

C. *Innovation and Other Outcomes*

To examine the effect that H-1B visas have on other firm outcomes, we estimate regressions of the form outlined in equation (1), varying the dependent variable. We present the results in table 3. We first examine the effect of chance H-1B visa wins on the innovative activity of the firm. Column 1 shows the results with the IHS of patents over the 8 years following the lottery as the dependent variable. We obtain a precisely estimated null effect of chance H-1B visa wins on the IHS of patents. For firms with 10 or fewer employees, the point estimate is a 0.026% increase in patents. With 95% confidence, we bound the effect between −0.42% and 0.47%, ruling out any material impact of an additional chance H-1B on patenting activity. Column 2 uses negative binomial regressions and recovers similar null effects, where with 95% confidence we can bound the impact to be between −1.08% and 0.21% for the smallest firms.

change in employment is outside the 95th (or higher) percentile or fixed size cutoffs (app. table 9), and we get similar results (though not as precise) without winsorizing (app. table 10).

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

[26] In particular, we show negative effects on employment of other workers quarter by quarter, where firms missing prelottery employment are included, across a variety of subgroups and contractors.

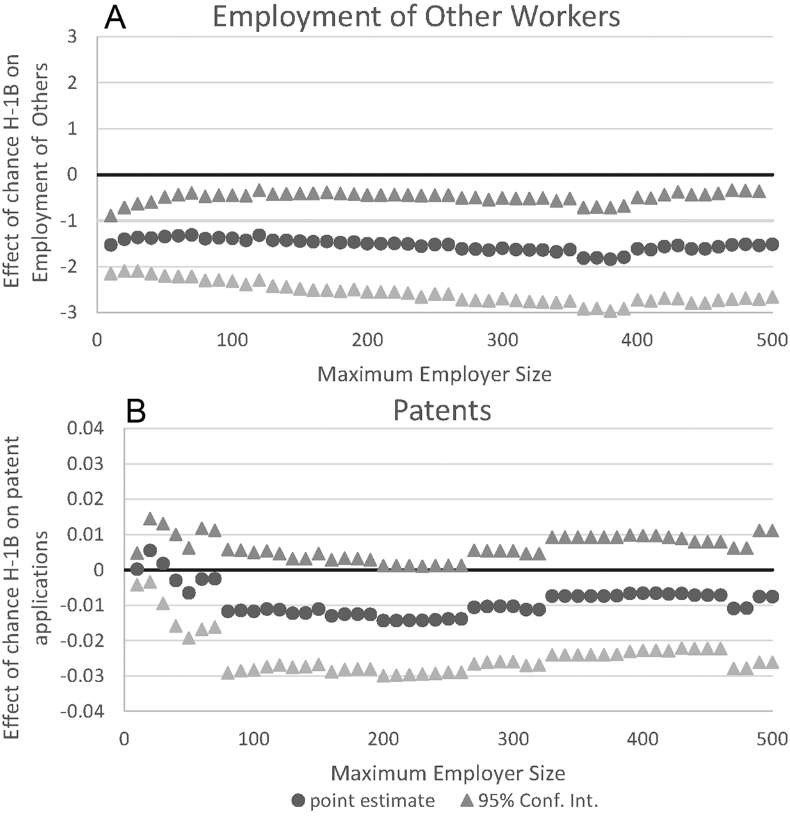

FIG. 1.—Effect of H-1B visas on firm outcomes, by employer size. *A*, Coefficient and 95% confidence interval (Conf. Int.) on a chance lottery win, minus the chance H-1B visa worker the firm won in the lottery, from median regressions, pooling together quarters 1–4 of the first fiscal year that an employee can work at the firm in the regression, among employers of the indicated size or smaller in year −1 (where the maximum employer size in each case is shown on the *x*-axis). The horizontal line at 0 on the *y*-axis corresponds to the case where hiring an extra H-1B visa worker leaves other employment unchanged. The horizontal line at −1 on the *y*-axis corresponds to the case where hiring an extra H-1B visa worker crowds out other workers one for one. *B*, Coefficient and 95% confidence interval on chance H-1B visas when the dependent variable is the IHS of patents in each year over years 0–8, among employers of the indicated size or smaller in year −1 (where the maximum employer size in each case is shown on the *x*-axis). After multiplication by 100, the coefficient should be interpreted as the approximate percentage increase in firm patenting due to a chance H-1B visa lottery win. We show the coefficient for employers of each size ranging from 0–10 to 0–500, with the upper bound of the size range in increments of 10. Note that the samples overlap across different regressions; for example, firms with 10 or fewer employees are included in the samples in all 50 regressions shown. We use the baseline employment specification, in which we control for the lagged dependent variable and expected lottery wins.

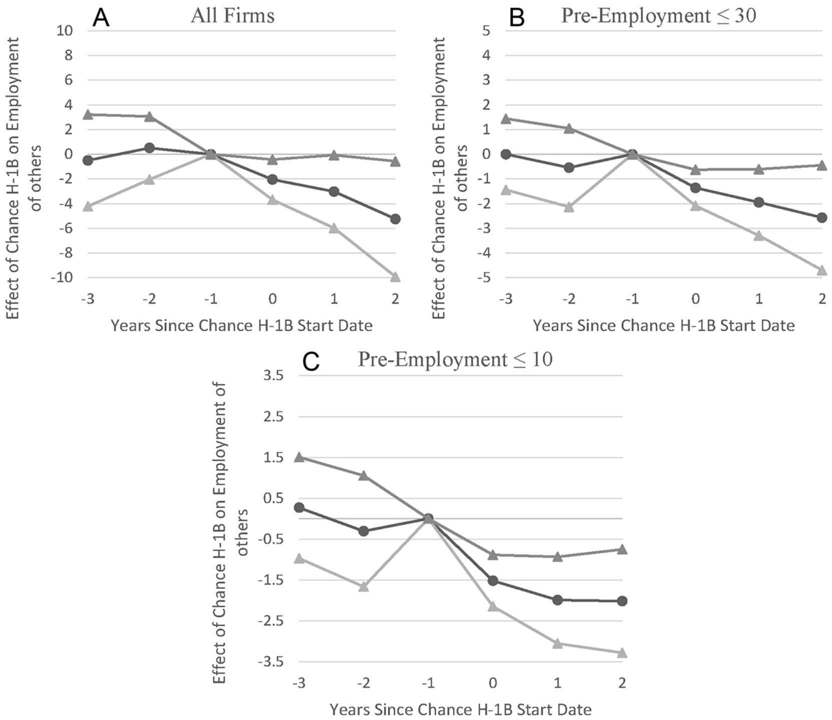

Fig. 2.—Effect of chance H-1B visa on employment of other workers, by years since H-1B lottery and by employer size. This figure shows the effect of a chance H-1B on firm employment from median regressions, by the years since the H-1B lottery, minus in years 0–2 the chance H-1B worker the firm won in the lottery. Because the employment data are quarterly, year 0 includes all quarters in the first year of the visa, year 2 reflects the third year of the visa, and all coefficients are estimated relative to base year −1. Firm employment in years 0–2 represents average quarterly employment each year. Each point estimate is a coefficient on chance H-1Bs from the baseline specification in table 2, with the dependent variable varying from year −3 to year 2. The dark lines represent the point estimates, and the lighter lines represent 95% confidence intervals. For more information, see thetable 2 note.

Examining robustness of the patenting results, figure 1*B* again varies the employer size threshold from 10 or fewer employees to 500 or fewer, in increments of 10. At the lower end of the firm size distribution, the point estimates hover around zero. As employer size increases to 80, we see a decline in the point estimates, yet we still cannot reject the null hypothesis of no effect at the 5% size, but a near-zero effect lies toward the top of the 95% confidence interval. Figure 3 repeats an exercise similar to that in figure 2, but with patents as the dependent variable. Effects are insignificant in the years before the lottery, and the dynamics in response to a chance H-1B visa do not display a clear pattern, in line with our precisely estimated "pooled" zero from before. Together, the evidence indicates that an additional chance H-1B visa win results in at most

TABLE 3

EFFECT OF H-1B LOTTERY WINS ON OTHER FIRM OUTCOMES

	IHS of No. of Patents (1)	No. of Patents (negative binomial) (2)	Amount of R&E Credit (IHS) (3)	Claiming R&E Dummy (4)	Payroll per Employee (5)	Profits (6)
A. ≤10 employees	.00026	−.0044	−.12	−.011	−4,861**	6,519
	[−.0042, .0047]	[−.0108, .0021]	[−.27, .041]	[−.025, .0041]	[−9,553, −168]	[−6,943, 19,979]
B. ≤30 employees	.0018	−.0073	−.065	−.0061	−2,725*	11,469**
	[−.0094, .013]	[−.0239, .0093]	[−.15, .018]	[−.014, .0016]	[−5,977, 527]	[201, 22,736]
C. All	−.0089	−.067	.19	.016	80	2,527
	[−.037, .019]	[−.18, .04]	[−.33, .72]	[−.018, .049]	[−1,348, 1,509]	[−32,169, 37,222]

NOTE.—The table shows regressions of patents in each year on chance H-1B lottery wins, with 95% confidence intervals in brackets. Specification 1 is an OLS regression with IHS of patents as the dependent variable. Specification 2 is a negative binomial regression with the number of patents as the dependent variable. Controlling for "prior patents" refers to controlling for the IHS of the total number of patents in year −1. The coefficients in the IHS specifications should be interpreted as the approximate percent effect on the number of patents. In specification 3, the dependent variable is the IHS of the amount of the R&E credit claimed in each year over years 0–3. In specification 4, the dependent variable is a dummy variable for whether the firm claimed any R&E credit in each of years 0–3, so that the coefficient reflects the effect on the fraction of years claiming any R&E. The "prior R&E" control refers to controlling for the amount (in cols. 1 and 2) or presence (in cols. 3 and 4) of the R&E credit in year −1. The IRS data measure only the R&E credit for C-corporations; other firms are excluded from the regressions. We find comparable results at other size thresholds, no significant interactions with covariates, and no significant differences across groups. The coefficients in the IHS specifications should be interpreted as the approximate percent effect on the amount of R&E taken. Specification 5 shows a median regression of payroll costs per employee in years 0–3 on chance H-1B visas and controls, pooling and stacking years. Years 0–3 cover the duration of the H-1B visa. The effect on payroll per employee in years 0 and 1 is comparable to the estimates shown. Payroll costs per employee in a given year are measured as total firm payroll costs in that year (in real 2014 dollars) divided by the total number of employees in the firm in that year. We use W-2 data because median regressions using Form 941 data generally did not converge. Specification 6 shows median regressions of profits in years 0–3 on chance H-1B visas and controls, pooling and stacking years. Profits are measured in real 2014 dollars. 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, e.g., in the sample of firms of all sizes. See table 1 for additional notes. Standard errors are clustered by firm.

* $p < .10$.

** $p < .05$.

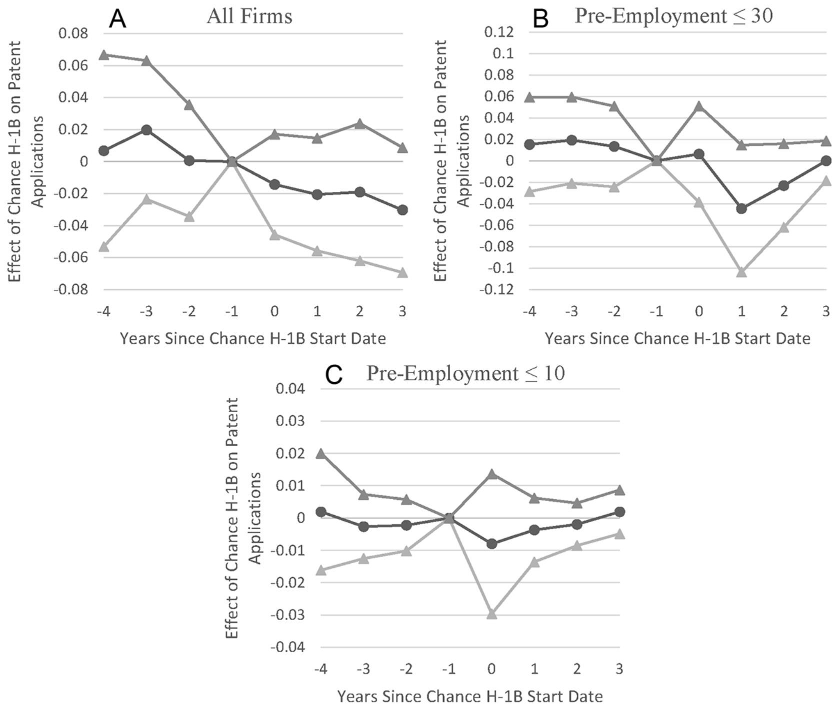

FIG. 3.—Effect of chance H-1B visa on patents, by years since H-1B lottery and by employer size. This figure shows the effect of a chance H-1B on the IHS of firm patenting, by the years since the H-1B lottery. Because the patent data are by calendar year and the visa is on a fiscal year basis, year 0 includes only one treated quarter. All coefficients are estimated relative to base year −1. Each point estimate is a coefficient on chance H-1Bs from the baseline specification in table 3, with the dependent variable varying from year −4 to year 3. The dark lines represent the point estimates, and the lighter lines represent 95% confidence intervals. For more information, see the table 3 note.

a very modest increase in patenting and, in the most negative case, a small decline in patenting. In appendix 2, we also explore and confirm robustness to our patenting findings, including when examining patents weighted by citations or limited to only firms in the more highly skilled ADE visa lotteries (app. tables 16, 18).

While patenting is one measure of the innovative output of a firm, measuring inputs into the production of innovative output is another way of capturing the firm's innovative activity. To do this, we estimate how a chance H-1B visa win affects intensive- and extensive-margin claims of the R&E credit during the duration of the visa in columns 3 and 4, respectively. For firms with 10 or fewer employees, the point estimate is negative, although the 95% confidence interval is large, allowing us to rule out only increases of greater than 4.1%. The point estimate for extensive-margin use of the R&E credit is negative, with the 95% confidence interval ruling out an impact greater than 0.41 percentage points. We can rule out even smaller increases for firms with 30 or fewer employees, but when all firms are used to estimate the effect, the results are less precise.

Focusing on the remaining firm outcomes, column 5 examines the impact that a chance H-1B visa win has on median payroll per employee during the duration of the visa. For smaller firms, we find declines in payroll per employee, with point estimates of several thousand dollars. Column 6 shows the results for firm profits. While the results are statistically significant only for firms with 30 or fewer employees, the positive point estimates in all three cases provide suggestive evidence for a moderate positive effect of chance H-1B wins on profits ranging from \$3,000 to \$11,000.

D. *Results from Secondary Strategies*

We now turn to estimating the impact of H-1Bs with our secondary empirical strategies. The results for all strategies, along with some reproduced experimental estimates, are presented in table 4 across the main dependent variables of interest. Starting with the RD design, appendix figure 1 visualizes the discrete and dramatic drop, relative to prior days, in the number of H-1B applications accepted that occurs on the day the cap is reached for all visa type–years but regular 2007, which is roughly in line with their respective win rate probabilities. In column 1, we report unwinosrized mean employment results for all secondary strategies. For small and medium-sized firms (and for the full set of firms in median regressions, as shown in app. fig. 2), the RD employment estimates are precisely measured and demonstrate crowd-out of other workers. These estimates are in line with the results from our experimental strategy, but, as expected, with tighter confidence intervals than the mean winsorized lottery estimates (and similar confidence intervals to the median lottery results).[27] While the effect of an H-1B on patenting is less precisely measured than that for our experimental design, the effect on patenting and use of the R&E credit is negative and generally significant (cols. 2 and 3). We take this as suggestive evidence that the RD approach finds little effect of H-1Bs on the innovative activity of the firm, with small negative effects possible. Results for pay per employee and profits are mixed and imprecise (cols. 4 and 5).

Turning to winners of the large FY2008 lottery and the comparison group of nonwinners presented in the third row of each panel in table 4, the estimates we recover are precise, including for the full sample of firms,

[27] The findings are robust to variation in bandwidth, polynomial order of the running variable, firm samples, and use of median regressions, with app. fig. 2 showing the full set of median estimates from varying the bandwidth from as low as 10 days through as high as 100 days.

and substantially more so than the prior strategies. We find that the effect of an H-1B is to similarly crowd out other workers, although the magnitude (around 0.75 workers) is smaller than what is seen in the experimental and RD results. The point estimates for patenting estimates are small and negative. The tight confidence interval allows us to rule out modest changes to the patenting activity of the firm. The results for the R&E credit rule out moderate changes in the use of the tax credit. Together, these results confirm the picture painted by the experimental and RD strategies: the effect of an H-1B on the innovative activity of a firm is small. While the results for profits and payroll per employee are less robust to changes in firm size, we see some evidence of higher profits and decreases in payroll per employee, which is in line with the experimental evidence.

For the shift-share strategy, as expected, the results indicate that the variation in the visa cap over time interacted with a firm's preperiod H-1B employment share has a statistically significant impact on the firm's H-1B employment share, with first-stage *F*-statistics of at least 66. As presented in the fourth row of each panel in table 4 and starting with employment, the results are in line with our previous findings; namely, they show a decrease in the employment of other workers on the order of 1.25 employees, and the effects are precisely estimated for all samples of firms. We find the effect of an H-1B on the patenting of the firm to be insignificant, with point estimates around zero, while effects on use of the R&E credit are negative. We also find negative but insignificant effects on payroll per employee and, somewhat surprisingly, negative effects on profit. Outside of the latter estimate, these results align with the prior results and more generally contribute to the story that H-1Bs have somewhere between a small negative and a zero impact on the innovative activity of a firm as measured by patenting activity and use of the R&E credit.[28]

Finally, the OLS regressions, presented in the fifth row of each panel in table 4, indicate that with firm fixed effects, an additional H-1B visa is associated with a decrease in the employment of other workers on the order of nearly 1 employee, regardless of firm size. Thus, the naïve difference between firms that hire and those that do not hire H-1B workers, as seen in table 1, is an artifact of other differences. H-1Bs are associated with very small increases in patenting and no changes in the R&E credit,

[28] That the prior literature finds positive effects on patenting could be driven by a failure to adequately capture time-specific differences that correlate with H-1B availability, by the earlier increases in the cap not contained within our sample, or by the cap failing to bind in many years. Using our specific years of data, sample of firms, and specification, we estimate positive effects on patenting when we do not control for interactions of a time trend with the preperiod values of the dependent variables. However, this sensitivity may not apply to other analyses. Relatedly, the negative effects on profits we recover could in part be because shift-share designs, which compare relative changes between the independent and dependent variables, are not well suited to dependent variables that are frequently negatively signed (in this case net profit) or could hint at violations in the parallel-trend assumption.

TABLE 4
COMPARISON OF IDENTIFICATION STRATEGIES: ESTIMATES OF EFFECT OF H-1BS ON FIRM OUTCOMES

	Employment of Other Workers (1)	IHS Patents (2)	IHS R&E Tax Credit (3)	Pay per Employee (4)	Profits (5)
			A. ≤10 Employees		
Experimental	−1.52***	.00026	−.12	−4,861**	6,519
	[−2.15, −.89]	[−.0042, .0047]	[−.27, .041]	[−9,553, −168]	[−6,943, 19,979]
Regression discontinuity	−1.46***	−.002	−.43*	−281	−3,585
	[−2.38, −.54]	[−.036, .031]	[−.89, .04]	[−3,117, 2,554]	[−15,451, 8,280]
2008 lottery winners vs. proxy control group	−.78***	−.0005**	−.0047**	−42.79	5,506**
	[−.83, −.73]	[−.00010, −.00001]	[−.0082, −.0012]	[−158.41, 72.82]	[640, 10,373]
Shift-share	−1.24***	.0046	−.14	−.033	−1.22**
	[−1.41, −1.08]	[−.013, .022]	[−.38, .09]	[−.082, .016]	[−2.33, −.18]
Fixed effects	−.83***	.00059**	.00062	76.90*	8,442
	[−.96, −.70]	[.00025, .00094]	[−.0019, .0032]	[−.12, 153.92]	[−6,801, 23,685]
			B. ≤30 Employees		
Experimental	−1.36***	.0018	−.065	−2,725*	11,469**
	[−2.09, −.63]	[−.0094, .013]	[−.15, .018]	[−5,977, 527]	[201, 22,736]
Regression discontinuity	−1.60**	−.003	−.32	417	−1,080
	[−2.90, −.29]	[−.033, .026]	[−.76, .12]	[−1,451, 2,286]	[−14,029, 11,868]
2008 lottery winners vs. proxy control group	.73***	−.0015***	.0016	−55.38	49,008
	[−.79, −.69]	[−.0025, −.0005]	[−.0090, .0122]	[−157.94, 47.19]	[−36,216, 134,233]
Shift-share	−1.27***	.0027	−.10	−.028	−1.05**
	[−1.41, −1.13]	[−.013, .018]	[−.24, .04]	[−.068, .013]	[−1.93, −.17]
Fixed effects	−.87***	.00077**	.00045	−5.57	−11,027
	[−.98, −.75]	[.00041, .00024]	[−.0015, .0024]	[−58.55, 47.40]	[−56,701, 34,647]
			C. All		
Experimental	−2.05**	−.0089	.19	80	2,527
	[−3.67, −.43]	[−.037, .019]	[−.33, .72]	[−1,348, 1,509]	[−32,169, 37,222]

Regression discontinuity	−29.89	−.020	.18	529	10,689*
	[−82.04, 22.25]	[−.049, .009]	[−.30, .66]	[−518, 1,575]	[−2,023, 23,402]
2008 lottery winners vs. proxy control group	−.64*	−.000015	.012**	−.20	−5,958,330*
	[−1.42, .14]	[−.00058, .00055]	[.0030, .021]	[−5.81, 5.42]	[−1.19E+07, 28,514]
Shift-share	−1.29***	−.022**	−.34***	−.018	−.90**
	[−1.41, −1.17]	[−.040, −.004]	[−.59, −.09]	[−.053, .017]	[−1.65, −.15]
Fixed effects	−.97**	.00012*	.00095	−.86	1,957,148
	[−1.30, −.64]	[−.00001, .00024]	[−.0018, .0037]	[−2.6, .88]	[−1,631,617, 5,545,912]

NOTE.—The table shows estimates of the effect of new H-1Bs on firm outcomes using various identification strategies, with 95% confidence intervals in brackets. The rows labeled "Experimental" are taken from tables 2 and 3. For the experimental identification strategy, specification 1 refers to col. 1 of table 2, panel A, while specifications 2–5 come from table 3. The dependent variables for all other strategies are the same as for the base lottery estimates, except that we use unwinsorized means, given the sufficient statistical power afforded by the other approaches, and use W-2 data for measures of employment, because the quarterly Form 941 data are not available for all of the years of study. The employment regression coefficients all subtract 1 to represent the effect on employment of other workers. For the rows labeled "Regression discontinuity," the number of new H-1Bs is instrumented by whether the firm submitted its visa application before the day the cap was reached interacted with visa category–year fixed effects, and the (quadratic) running variable is the day in which the firm submitted its visa application(s). The effect of a firm applying before the day the cap is reached on the number of H-1Bs approved for that firm is 0.601 (0.050) for the 2006 ADE lottery, 0.976 (0.025) for the 2006 regular lottery, and 0.459 (0.057) for the 2007 ADE lottery ($p < .01$ for all three), with no first stage for the excluded 2007 regular lottery. Similar first stages are observed when examining firms with up to 30 or up to ≤10 employees. We restrict the sample to the 76.4% of all firms (73.32% of applications) that submitted all their applications in any given visa category–year on one day (in the absence of this restriction, there is no discontinuity at the firm level). All regressions use the optimal bandwidth of 20 days. All specifications include visa category–by-year fixed effects, lagged dependent variable, and a constant. The number of observations varies from 1,881 (panel A, col. 6) to 58,452 (panel C, col. 3). The rows labeled "2008 lottery winners vs. proxy control group" are estimated via regressions that control for expected wins on the basis of a control group proxy of LCAs, lagged dependent variable, and a constant. The number of observations varies from 4,977 (col. 1, panel A) to 79,870 (col. 2, panel C). The rows labeled "Shift-share" show the results from the shift-share strategy, which controls for firm fixed effects, indicators for year, an interaction of a time trend with the preperiod value of each dependent variable, and a constant and uses the IHS of the dependent variable for interpretation as a percent change on a percent change. Given the shift-share setup, the dependent variables are contemporaneous to the number of approved H-1B. F-statistics in the first stages range from 66 to 143. The number of observations varies from 9,826 (col. 4, panel A) to 77,930 (col. 2, panel C). The rows labeled "Firm fixed effects" show OLS regressions of the number of new H-1Bs in that firm on firm outcomes. All regressions include year and firm fixed effects and a constant. The number of observations varies from 247,759 (col. 5, panels A and B) to 2,882,375 (col. 2, panel C). Standard errors are clustered at the firm level in all regressions.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

with confidence intervals that rule out even small effects for both measures. Finally, we find little evidence of a relationship between H-1Bs and other outcomes at the firm.

In summary, we find that alternative strategies, with a variety of different strengths, weaknesses, and degrees of generalizability, validate our experimental finding that the effect of an H-1B is to reduce the employment of other workers at the firm and that H-1Bs have no meaningful impact on the patenting activity of firms.[29]

E. *Discussion of Results*

Our employment results consistently show decreases in the employment of other workers, thus implying that additional H-1Bs robustly crowd out other individuals who would have worked at the firm. Taken at face value, most lottery and RD point estimates in fact imply that an H-1B worker crowds out more than one worker, although we cannot reject a null hypothesis of one-for-one crowd-out; said differently, we do not find statistically significant evidence that overall employment (including the H-1B) at the firm goes down. In considering our other strategies, the shift-share point estimates, while smaller in size than the lottery and RD estimates, similarly show greater than one-for-one crowd-out (statistically significantly so), while the remaining strategies point to crowd-out close to, but below, one for one, and their confidence intervals can generally rule out more than that one-to-one crowd-out. Yet while there is some variation in the point estimates across strategies, it is also almost always the case that we cannot reject the null hypothesis that any strategy's estimates are the same as another set's. Altogether, our results are consistent with the possibility that new H-1B workers and otherwise available workers are perfect substitutes.[30]

Still, given the size of the coefficients for some strategies, it is worthwhile to ask, What underlying feature of H-1B employment might cause H-1Bs to crowd out multiple other workers at once? We offer several hypotheses that would be consistent with this. First, it is possible that H-1Bs work harder than alternative employees do. This could arise as a result of employer monopsony power over opportunities for a green card or for reasons intrinsically related to the quality of the workers themselves.

[29] Appendix figs. 3 and 4 show that for these identification strategies, our findings on the effect of H-1Bs on employment of other workers and patenting are robust across the firm employment distribution.

[30] They do not necessarily imply perfect substitutability, as it would also depend on the degree of substitutability or complementarity of H-1B labor with capital. See Lewis (2011), which studies the interaction of immigration with capital. And if a firm faces frictions in finding a new employee that limit the degree of crowd-out of other workers, the amount of crowd-out we detect is all the more notable.

TABLE 5
Effect of H-1Bs on Patents in Large or Innovative Firms

FY2006/2007 Lotteries (1)	FY2006/2007 RD (2)	FY2008 Quasi Lottery (3)	Shift-Share (4)	OLS Firm Fixed Effects (5)
A. Firms with at Least One Patent in the Preperiod				
−.057	−.17	−.00064*	−.034	.000070
[−.30, .19]	[−.46, .12]	[−.0013, 1.49E−06]	[−.11, .04]	[−.000063, .000203]
B. Large Firms (Preemployment > 30)				
−.023	−.030	−.00038	−.15	.000088
[−.073, .027]	[−.074, .013]	[−.00094, .00019]	[−.41, .10]	[−.000021, .000196]

Note.—The table shows the results of five identification strategies applied to firms that are either innovative (panel A, with at least one patent in the preperiod) or large (panel B, with preemployment > 30), with 95% confidence intervals in brackets. The specification in each column corresponds to the baseline specification in each case; see table 4 for more details. The number of observations varies from 1,557 (col. 1, panel A) to 636,605 (col. 5, panel B). The standard errors are clustered at the firm level.
* $p < .10$.

Second, (full-time) H-1B workers could be replacing multiple part-time workers, though our results on payroll per employee are not consistent with this hypothesis. Third, H-1Bs may indeed have special skills, but firms may use these skills to facilitate changes in production technology that lead to a different mix of outputs and inputs and, in particular, to a decrease in demand for otherwise available workers. The evidence in section II about the types of tasks H-1B workers likely perform is not very consistent with this hypothesis.

Turning to patenting, we find generally small and insignificant effects.[31] This result is not especially surprising, given the likely occupations reported in section II. While our lottery results are not an artifact of focusing on small noninnovative firms hiring workers with relatively low education—firms applying on the day the cap is reached are more likely than other applicants to have patented in the past, to be in scientific industries, and to apply for workers with higher educational degrees—the ability to focus on lottery effects in large or especially innovative firms where most patenting is concentrated is limited by statistical power. However, this is not the case for our other strategies.

In panel A of table 5, we restrict the sample to innovative firms by dropping those that have never patented in prior periods and examining effects using all of our identification strategies. The results confirm that several strategies (in particular, the 2008 quasi-lottery, shift-share, and OLS

[31] The extremely small positive estimates from firm fixed effects regressions are offset by the modest decreases in patenting found in some 2008 quasi-lottery and shift-share estimates.

firm fixed effect regressions) are indeed sufficiently powered to produce precise estimates for the small subsample of innovative firms, showing no increase in patenting due to an additional H-1B. In panel B, we focus only on firms with over 30 employees in the preperiod, given that patenting is more prevalent among larger firms, and we find similar results. Thus, our findings apply even when just considering more innovative firms. Additionally, the R&E credit results tend to be clustered around zero or somewhat negative. While not always as precise as the results for patenting, the combined R&E and patenting results paint the picture that both firm inputs into the innovative process and innovative output do not see meaningful increases from additional H-1Bs.

Finally, we find some evidence of a decrease in payroll per employee and increase in profits in our lottery setting. The impact on payroll per employee could be coming from multiple mechanisms. First, employers may pay H-1Bs less than the average wage. This could be due to monopsony power that the employer has because of the individual's H-1B status. Second, employment of H-1Bs could reduce wages paid to H-1B workers, which, combined with lower employment of other workers, suggests that the effect of H-1Bs is to lower the demand for other workers. In turn, the positive indication on profits, along with payroll-per-employee results, suggest that the firm is able pay its employees less without a similar-sized drop in output; that is, the firm is able to extract (additional) rents from labor. That said, the results on profits and payroll per employee are less precise and robust across strategies than our main outcomes and thus more speculative.

VII. Conclusion

The effect of highly skilled immigrants on firms is one of the centrally important US immigration policy questions. We examine the impact of the United States' largest high-skill immigration program, the H-1B program, and find that new H-1Bs crowd out other workers associated with similar observable levels of innovation. This result is informative for understanding the effects of these individual workers on firms and for understanding the labor market for high-skilled technical workers in the United State and the nature of labor market substitution. We bring several new advances to the literature, including randomized visa lotteries and IRS data on the universe of US firms. We apply these IRS data not only to the new randomized visa lotteries strategy but also to a new RD design, as well as other strategies that the prior literature has relied upon that we now apply to the firm level, drawing on the universe of firms.

Taken together, these identification strategies run the gamut from precise causal identification on small subsets of firms to more general correlative statements about the universe of US firms, and from an analysis of years in which demand for H-1Bs moderately outstripped supply to an analysis of years in which demand for H-1Bs greatly outstripped supply. The primary finding that across all identification strategies one additional H-1B leads to significant crowd-out of other workers at the firm level holds whenever the sample we study is sufficiently powered to produce precise estimates, sometimes even when including very large firms. This is relevant in light of frequent claims that H-1Bs have unique skills that cannot easily be obtained elsewhere.

Likewise, across all identification strategies and all years, one additional H-1B does not cause any meaningful increase in patenting. In contrast, point estimates of the effects of H-1Bs on patenting are typically negative, and standard errors are typically small enough to rule out more than a small percentage or absolute effect. These findings are not an artifact of small numbers of marginal H-1Bs being denied because of the cap, as they also hold even for identification strategies such as the 2008 quasi lottery and the 2004 cap reduction, which involve substantial numbers of H-1Bs at stake in the economy as a whole and in each firm.

Consistent with firm profit maximization, we find some evidence that extra H-1B visas increase firm profits. We also find some evidence that extra H-1B visas lead to a decrease in earnings per employee, especially in the lottery identification strategy. If these findings reflect higher economic profits and/or lower pay for H-1Bs than for alternative workers, then this would suggest the existence of market frictions, such as firm labor market monopsony power.

Overall, our results are more supportive of the view that H-1Bs crowd out alternative workers, are paid less per unit of effort than the alternative workers whom they crowd out, and thus increase the firm's profits despite no measurable effect on innovation. Prima facie, these results appear at odds with a chief goal of the program, as articulated by policy makers in legislation, of providing firms with skilled workers who have unique, innovative skills that the firms cannot otherwise obtain. Even though firms attest that hiring H-1Bs does not adversely affect similarly employed workers, our results raise this possibility.[32] Future research should investigate whether H-1Bs' pay is consistent with prevailing-wage regulations and whether employers meet the test of being unable to hire a comparable worker. And while we find little effect on firms' quantity of innovation, assessing impacts on productivity should be a priority for further research.

Our study estimates only partial-equilibrium effects, not general-equilibrium effects. In particular, we isolate the effect of additional H-1B visas allocated to a given firm on outcomes at that firm (holding all else

[32] Our results do not necessarily imply that firms' behavior is inconsistent with their attestations, e.g., because the congressional intent may have been to prevent harm to US citizens specifically.

equal). However, these should be a key determinant of the general-equilibrium effects. If the crowded-out workers become employed in other firms (i.e., assuming that labor demand is not perfectly inelastic) and innovate there, aggregate employment and patenting could still increase, as long as this boost to innovation does not crowd out innovation elsewhere. These or other mechanisms could reconcile positive aggregate effects with small firm-level effects. However, this mechanism for raising employment and innovation would be very different from the hypothesis that H-1Bs directly raise employment and innovation at the firm level as well, as both the business community and policy makers have claimed, and suggests smaller aggregate effects than would otherwise be expected.

References

Acs, Zoltan J., and David B. Audretsch. 1990. *Innovation and Small Firms.* Cambridge, MA: MIT Press.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019. "Who Becomes an Inventor in America? The Importance of Exposure to Innovation." *Q.J.E.* 134 (2): 647–713.

Borjas, George J. 1994. "The Economics of Immigration." *J. Econ. Literature* 32 (4): 1667–717.

———. 2003. "The Labor Demand Curve *Is* Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *Q.J.E.* 118 (4): 1335–74.

Borjas, George J., and Kirk B. Doran. 2012. "The Collapse of the Soviet Union and the Productivity of American Mathematicians." *Q.J.E.* 127 (3): 1143–203.

Borjas, George J., Richard B. Freeman, and Lawrence F. Katz. 1997. "How Much Do Immigration and Trade Affect Labor Market Outcomes?" *Brookings Papers Econ. Activity* 1997 (1): 1–90.

Borjas, George J., Jeffrey Grogger, and Gordon H. Hanson. 2012. "Comment: On Estimating Elasticities of Substitution." *J. European Econ. Assoc.* 10 (1): 198–210.

Bound, John, Breno Braga, Joseph M. Golden, and Gaurav Khanna. 2015. "Recruitment of Foreigners in the Market for Computer Scientists in the United States." *J. Labor Econ.* 33 (S1): S187–S223.

Burbidge, John B., Lonnie Magee, and A. Leslie Robb. 1988. "Alternative Transformations to Handle Extreme Values of the Dependent Variable." *J. American Statis. Assoc.* 83 (401): 123–27.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82 (6): 2295–326.

Card, David. 1990. "The Impact of the Mariel Boatlift on the Miami Labor Market." *Indus. and Labor Relations Rev.* 43 (2): 245–57.

———. 2001. "Immigrant Inflows, Native Outflows, and the Local Market Impacts of Higher Immigration." *J. Labor Econ.* 19 (1): 22–64.

Depew, Briggs, Peter Norlander, and Todd A. Sorensen. 2013. "Flight of the H-1B: Inter-Firm Mobility and Return Migration Patterns for Skilled Guest Workers." IZA Discussion Paper No. 7456, Institute for the Study of Labor, Bonn.

Dustmann, Christian, Albrecht Glitz, and Tommaso Frattini. 2008. "The Labour Market Impact of Immigration." *Oxford Rev. Econ. Policy* 24 (3): 477–94.

Edin, Per-Anders, Peter Fredriksson, and Olof Åslund. 2003. "Ethnic Enclaves and the Economic Success of Immigrants—Evidence from a Natural Experiment." *Q.J.E.* 118 (1): 329–57.

Foley, C. Fritz, and William R. Kerr. 2013. "Ethnic Innovation and U.S. Multinational Firm Activity." *Management Sci.* 59 (7): 1529–44.

Freeman, Richard B. 2006. "People Flows in Globalization." *J. Econ. Perspectives* 20 (2): 145–70.

Friedberg, Rachel M. 2001. "The Impact of Mass Migration on the Israeli Labor Market." *Q.J.E.* 116 (4): 1373–408.

Friedberg, Rachel M., and Jennifer Hunt. 1995. "The Impact of Immigrants on Host Country Wages, Employment and Growth." *J. Econ. Perspectives* 9 (2): 23–44.

Grogger, Jeff, and Gordon H. Hanson. 2015. "Attracting Talent: Location Choices of Foreign-Born PhDs in the United States." *J. Labor Econ.* 33 (S1): S5–S38.

Hanson, Gordon H. 2009. "The Economic Consequences of the International Migration of Labor." *Ann. Rev. Econ.* 1:179–208.

Hunt, Jennifer. 2011. "Which Immigrants Are Most Innovative and Entrepreneurial? Distinctions by Entry Visa." *J. Labor Econ.* 29 (3): 417–57.

Hunt, Jennifer, and Marjolaine Gauthier-Loiselle. 2010. "How Much Does Immigration Boost Innovation?" *American Econ. J. Macroeconomics* 2 (2): 31–56.

Kerr, William R. 2013. "U.S. High-Skilled Immigration, Innovation, and Entrepreneurship: Empirical Approaches and Evidence." Working Paper No. 14-017, Harvard Bus. School.

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

Lewis, Ethan. 2011. "Immigration, Skill Mix, and Capital Skill Complementarity." *Q.J.E.* 126 (2): 1029–69.

Lubotsky, Darren. 2007. "Chutes or Ladders? A Longitudinal Analysis of Immigrant Earnings." *J.P.E.* 115 (5): 820–67.

Moser, Petra, Alessandra Voena, and Fabian Waldinger. 2014. "German Jewish Émigrés and US Invention." *A.E.R.* 104 (10): 3222–55.

North, David. 2011. "Estimating the Size of the H-1B Population in the U.S." Memorandum, Center for Immigration Studies, Washington, DC.

Pekkala Kerr, Sari, and William R. Kerr. 2011. "Economic Impacts of Immigration: A Survey." *Finnish Econ. Papers* 24 (1): 1–32.

Pekkala Kerr, Sari, William R. Kerr, and William F. Lincoln. 2015. "Skilled Immigration and the Employment Structures of US Firms." *J. Labor Econ.* 33 (S1): S147–S186.

Pence, Karen M. 2006. "The Role of Wealth Transformations: An Application to Estimating the Effect of Tax Incentives on Saving." *B.E. J. Econ. Analysis and Policy* 5 (1): 20.

Peri, Giovanni, Kevin Shih, and Chad Sparber. 2013. "The Effects of Foreign Skilled Workers on Natives: Evidence from the H-1B Visa Lottery." Working paper, Univ. California, Davis.

———. 2015. "STEM Workers, H-1B Visas, and Productivity in US Cities." *J. Labor Econ.* 33 (S1): S225–S255.

Senate Judiciary Committee. 1998. "Report to Accompany American Competitiveness Act." United States Senate, 105th Congress. S. Rep. No. 105-186.

———. 2000. "Report to Accompany American Competitiveness in the Twenty-First Century Act." United States Senate, 106th Congress. S. Rep. No. 106-260.