

How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program

Britta Glennon, University of Pennsylvania
Working Paper: February 2020

Abstract

The decision to encourage or restrict high-skilled immigration has long been controversial. Advocates argue that high-skilled immigration is critical for firm competitiveness and innovation; critics argue that skilled immigrants displace native workers and drive down wages. The debate, however, has largely overlooked the secondary consequences of restrictions on high-skilled hiring of immigrants: multinational firms faced with decreased access to visas for skilled workers have an offshoring option, namely, hiring the foreign labor they need at their foreign affiliates. This paper documents the impact of restrictive high-skilled immigration policies on the offshoring of high-skilled jobs by US multinational companies using a unique matched firm-level dataset of H-1B visas and multinational firm activity and two different identification strategies. Both strategies yield the same result: that restrictions on H-1B immigration caused increases in foreign affiliate employment at both the intensive (US multinationals employed more people at their foreign affiliates) and the extensive (US multinationals opened more foreign affiliates) margins. The effects are concentrated among highly H-1B-dependent firms and R&D-intensive firms operating in offshorable services sectors, and the expansion of activity was concentrated in Canada, India, and China. The findings of the paper have important policy implications; the offshoring of jobs appears to be an unforeseen consequence of restricting skilled immigration flows. Even if H-1B immigrants displace some native workers, any policies that are motivated by concerns about the loss of native jobs should consider that these same policies have the unintended consequence of encouraging firms to offshore jobs abroad.

Acknowledgments: I would like to thank Lee Branstetter, Brian Kovak, J. Bradford Jensen, Giovanni Peri, and Erica Fuchs for helpful comments. The statistical analysis of firm-level data on U.S. multinational companies was conducted at the Bureau of Economic Analysis (BEA), United States Department of Commerce under arrangements that maintain legal confidentiality requirements. I thank Bill Zeile, Jim Fetzer, and Ray Mataloni for helpful discussions on the BEA data. All errors and omissions remain my own responsibility. The views expressed are those of the author and do not reflect official positions of the U.S. Department of Commerce.

1. Introduction

The question of the impact of immigration on the host country has long been controversial, but it has risen to the forefront of political debates in recent years. Unexpected political shifts such as the Brexit vote and the election of President Trump have been attributed to voter concern about the impact of immigration. While the debate surrounding low-skilled immigration has captured headlines in the US, high-skilled legal immigration – and particularly the H-1B visa program – has also been contentious. Critics of the H-1B program argue that skilled immigrants displace native-born workers and drive down their wages.¹ Indeed, H-1B rejection rates have more than tripled since President Trump signed the Buy American and Hire American Executive Order² in early 2017.³ However, business leaders have decried both these recent measures and long-standing restrictions on high-skilled immigration, arguing that the shortage of workers with specialized skills has negatively affected the competitiveness and innovation of high-tech firms and of the US economy.⁴

Policy debates like these have spawned an extensive academic literature evaluating the claims of each side. The debate, however, has largely overlooked the secondary consequences of restrictions on high-skilled hiring of immigrants: multinational companies faced with decreased access to visas for skilled workers have an offshoring option, namely, hiring the foreign labor they need at their foreign affiliates. Understanding the response of multinational companies to these restrictions on skilled immigrants is especially significant because they are the leading employers of skilled immigrants and engage in the vast majority of formal innovative activities in and outside of the US. In addition, unlike other firms⁵, they have the option of responding to skilled immigration restrictions by offshoring their high-skilled activities. Despite the implications of such an option being used by such significant actors in the domestic

¹ Richard Trumpka, the President of the AFL-CIO wrote that: “Clearly, high tech is not looking to bring in H-1B visa holders for a few years at a time because there is a shortage of tech workers. They want a massive expansion of H-1B visa holders because they can pay them less. This is not about innovation and job creation. It is about dollars and cents.” <https://www.usatoday.com/story/opinion/2013/05/28/h1b-visa-high-tech-workers-afl-cio-editorials-debates/2367769/> Critics like the AFL-CIO often cite the work of Hira (2010) and Matloff (2003).

² The Buy American and Hire American Executive Order directed four federal agencies – the departments of State, Labor, Justice and Homeland Security – to crack down on fraud and abuse of the H-1B and other work visa programs. In response, USCIS has increased H-1B inspections and commenced site visits of businesses employing foreign workers holding “specialized knowledge”. USCIS also has significantly increased challenges and requests for more evidence.

³ U.S. Customs and Immigration. Non-Immigrant Worker RFE Data. <https://www.uscis.gov/sites/default/files/USCIS/Resources/Reports%20and%20Studies/Immigration%20Forms%20Data/BAHA/non-immigrant-worker-rfe-h-1b-quarterly-data-fy2015-fy2019-q1.pdf>

⁴ In Congressional testimony in 2008, Bill Gates warned that unless the U.S. expanded its H-1B program, it would be “at risk of losing its position of technological leadership”. Eric Schmidt, speaking at MIT’s Computer Science and Artificial Lab, said that limits on the H-1B visa program “make it more difficult for U.S. companies to remain competitive.” <https://www.cnn.com/2017/05/04/googles-eric-schmidt-h1b-limit-is-stupidest-us-policy.html>

⁵ Unless they choose to internationalize.

and global economy, to the best of my knowledge, no paper has examined whether the offshoring of jobs is a consequence of restricting skilled immigration flows. If US multinationals use this option in response to restrictive H-1B policies – as their public statements suggest⁶ - then such restrictive migration policies are unlikely to have the desired effects of increasing employment and earnings of high-skilled natives, but rather have the effect of offshoring high-skilled jobs.

This project directly examines the impact of restrictive high-skilled immigration policies on the offshoring of high-skilled jobs by US multinational companies (MNCs). I use a unique dataset that combines firm level data on H-1B visas and multinational firm activity to examine whether restrictions on H-1B visas result in (1) an increase in high-skilled foreign affiliate employment, and (2) an increase in the likelihood of opening foreign affiliates in new countries.

The data used in this paper are constructed by combining four different datasets: (1) the Bureau of Economic Analysis (BEA) annual surveys on US Direct Investment Abroad database, which contains detailed microdata on the financial and operating characteristics of both the US parent companies and their foreign affiliates, (2) H-1B visa microdata, (3) Labor Condition Application (LCA) data⁷, and (4) US Patent and Trademark Office (USPTO) patent data. These data allow me to measure exactly how constrained each firm was as the cap grew more restrictive over time by comparing – at the firm-level – LCA requests (demand) and issued H-1B visas (realized supply), and how their foreign affiliate employment responded to these constraints. I analyze the impact of restrictions on H-1B visas on foreign affiliate activity using two identification strategies. The first exploits the 2004 drop in the H-1B visa cap, while the second exploits randomized variations in firm-level excess demand from the H-1B visa lotteries in high demand years. Both strategies yield the same result: that restrictions on H-1B immigration caused increases in foreign affiliate activity at both the intensive margin (US multinationals employed more people at their existing foreign affiliates) and the extensive margin (US multinationals opened foreign affiliates in new countries). The effects are concentrated among highly H-1B-dependent firms and R&D-intensive firms operating in offshorable services sectors. The expansion of foreign affiliate employment has been largely concentrated in three countries: China, India, and Canada.

⁶ Carbonite: “if [we] can’t get them admitted to the United States, [we’ll] staff up at Carbonite offices in Canada and Europe” <https://www.bostonglobe.com/business/2017/04/02/tech-industry-talent-shortage-claims-under-new-scrutiny/EsxYnPpoKBNvliTjRl6LL/story.html>

Amazon: “we are currently assessing alternatives that could include placement in countries other than the United States” <https://www.geekwire.com/2017/trumps-immigration-crackdown-may-force-amazon-microsoft-shift-workers-canada/>

⁷ Labor Condition Applications are the first step towards H-1B visas for skilled foreign-born workers in the U.S.

Two aspects of this paper are novel. First, this paper provides the first empirical evidence to support the hypothesis that restrictions on high-skilled immigration cause the offshoring of skilled jobs. While high-profile cases – like Microsoft’s decision to open an R&D foreign affiliate in Vancouver⁸ – have suggested that restricting skilled immigration flows could lead to the offshoring of jobs, this paper presents the first empirical work both proposing this hypothesis and examining whether this claim is indeed true. Second, it is the first paper to use a matched firm-level dataset of H-1B visas and multinational firm activity.

The findings of the paper have important policy implications; the offshoring of jobs appears to be an unforeseen consequence of restricting skilled immigration flows. Even if H-1B immigrants displace some native workers, any policies that are motivated by concerns about the loss of native jobs should consider that policies aimed at reducing immigration have the unintended consequence of encouraging firms to offshore jobs abroad.

2. The Effects of a Negative Foreign Labor Supply Shock on Offshoring: Literature Review

Highly-skilled workers are crucial and relatively scarce inputs into firms’ productive and innovative processes. An increasingly high proportion of these workers – and particularly STEM workers – in the US were born abroad and immigrated to the US (Bound et al. 2014). This phenomenon has spawned an extensive literature on skilled immigration and a heated policy debate on the appropriate admissions levels of skilled immigrants. The literature has largely focused on the impact of high skilled immigration along three dimensions – (i) the impact on innovation, (ii) the impact on native workers’ outcomes, and (iii) the impact on the source country (brain drain).⁹ The policy debate has largely centered around these same three issues.

⁸ “Microsoft opens Canada center in response to US immigration problems.” <http://workpermit.com/news/microsoft-opens-canada-center-response-us-immigration-problems-20070710>

⁹ Studies of the impact on innovation and entrepreneurship include Chellaraj, Maskus, and Mattoo (2008), Wadhwa et al. (2007), Akcigit, Grigsby, and Nicholas (2017), Burchardi et al. (2019), Doran and Yoon (2019), Ganguli (2015) Agrawal et al. (2018), Borjas and Doran (2015b), Ghosh et al. (2015), Hunt (2011), Hunt and Gauthier-Loiselle (2010), Kerr and Lincoln (2010), Moser et al. (2014), Oettl and Agrawal (2008), and Wu (2017). Docquier and Rapoport (2012) provide an extensive review on the brain drain literature. Studies of the impact on native wages and jobs include Bound, Khanna, and Morales (2017), Turner (2017), Borjas (2005, 2003), Card (2009, 2001), Choudhury and Kim (2018), Doran et al. (2016), R. B. Freeman (2006), Friedberg 2001, Hayes and Lofstrom (2011), Hunt (2011), Kerr, William R Kerr, et al. (2015), Lowell (2001), Matloff (2003), Mayda et al. (2017), Mithas and Lucas (2010), Ottaviano, Peri, and Wright (2012), Peri et al. (2015), Salzman and Lowell (2007), and Tambe and Hitt (2009).

In this extensive immigration literature, there is surprisingly little focus on the role of the firm, as was noted in Kerr, Kerr, and Lincoln (2015).¹⁰ The lacuna is particularly surprising in the U.S. literature, since applications for skilled immigration visas like H-1B visas are made by sponsoring firms. Fortunately, in recent years, a series of careful empirical papers have begun to address this gap. Ashraf and Ray (2017) and Wu (2017) consider the impact of high skilled immigration on firm innovation. Doran et al. (2016), Kerr, et al. (2015), and Mayda et al. (2017) examine the impact on firm structure and employment. Ashraf and Ray (2017), Ghosh et al. (2015), and Xu (2016) examine the impact on other firm outcomes.

However, while these papers have significantly contributed to our understanding of the impact of high-skilled immigration on firm outcomes, they have largely overlooked the multinational nature of many of these firms in their analysis. Multinational companies (MNCs) are the leading employers of skilled immigrants, engage in the vast majority of formal innovative activities, and – unlike other firms¹¹ - have the option of responding to skilled immigration restrictions by offshoring their high-skilled activities. Branstetter, Glennon, and Jensen (2018) argue that skilled labor shortages in the past – a result of a large IT- and software-biased shift in innovation – did drive US MNCs abroad, and particularly drove them to locations with large quantities of STEM workers who possessed IT and software skills. High-skilled immigration provides another way of addressing this shortage, but an increasingly restrictive cap on H-1B admissions that began in 2004 reduced the ability of US firms to use this approach to meet their human capital needs. This line of thinking suggests that, to some degree, immigration and offshoring are substitutes.

Some recent papers in the international trade literature (Olney and Pozzoli 2018; Ottaviano, Peri, and Greg C. Wright 2013) have indeed found that immigration and offshoring are substitutes at the multilateral level, but rely on import data to measure offshoring while I directly observe foreign affiliate employment. Furthermore, these papers largely consider the offshoring of manufacturing and immigration of all types of workers, as opposed to focusing on high-skilled immigration and high-skilled offshoring activity. Global production networks and global R&D networks are fundamentally different in how they operate however; while production has become highly dispersed around the world, most formal research and development and other high-skilled activities remain highly concentrated in a few firms' headquarters in only a few countries. This is at least in part due to substantial frictions in international collaboration (Argote, Mcevily, and Reagans 2003; Audretsch 1998; Patel and Pavitt 1991) that often require researchers to work in physical proximity; tacit knowledge is best transferred in person (Polanyi 1958,

¹⁰ Kerr, Kerr, and Lincoln (2015) provide a vivid example to support this, noting that in the 51 pages of Borjas (1994)'s classic survey of the economics of immigration literature, the word "firm" does not appear once.

¹¹ Unless they choose to internationalize.

1966; Singh 2008; Szulanski 1996; Teece 1977). In keeping with this, the globalization of R&D literature has long recognized that a central driver of globalized R&D is a host country's supply of human capital (Hall 2010; Serapio and Dalton 1999; Siedschlag et al. 2013; Thursby and Thursby 2006), while this is much less important for the offshoring of manufacturing. Hence, the degree to which offshoring and immigration are substitutes could look quite different for high-skilled tasks than for production activities.

There is one avenue through which offshoring and immigration could be complementary rather than substitutable, and this is at the bilateral level; hiring an immigrant of a given ethnicity can actually increase a firm's offshoring in that immigrant's country of origin through ethnic ties. Firms may utilize the expertise and networks of an immigrant when offshoring in that immigrant's country of origin (Arora and Gambardella 2005; Choudhury 2016; Choudhury and Kim 2018; Foley and Kerr 2013; Gould 1994; Hernandez 2014; Iriyama, Li, and Madhavan 2010; Kerr 2008; MacGarvie 2005; Saxenian 2006). They argue that ethnic ties facilitate the disintegration of innovative activity across borders and allow multinationals to more easily form new affiliates abroad. This hypothesis implies that a reduction in immigration from a given country could reduce offshoring, since it would make it more difficult for firms to set up new foreign affiliates there without the diaspora networks.¹²

In short, the sign and magnitude of the effects, if any, of high-skilled immigration restrictions on multinational foreign affiliate employment are still an open question.

3. Overview of the H-1B Program

Firms have multiple ways in which they can hire foreign high-skilled workers: the H-1B, L-1, O, OPT, and TN visas are just a few examples.¹³ The first is the most widely used and is the focus of this paper. The H-1B visa is a nonimmigrant¹⁴ visa that enables firms to hire foreign workers in the US for a three-year period, renewable once for a total of six years. They make up about 50% of temporary work visas, and are used to employ foreign workers in "specialty occupations"¹⁵ which typically means the

¹² This argument is not necessarily incompatible with substitution at the multilateral level; Olney and Pozzoli (2018) find both that immigration substitutes for offshoring at the multilateral level and that it complements offshoring at the bilateral level. In other words, they find that while bilateral offshoring increases with immigration from the host country, it decreases with immigration from other countries.

¹³ I provide some discussion of other high-skilled visa alternatives in the appendix.

¹⁴ They are called "nonimmigrant" visas because they allow those with H-1Bs to stay in the US only temporarily. However, they are also "dual intent" visas, which means that workers can reside in the US with a nonimmigrant status while simultaneously applying for permanent residency.

¹⁵ According to USCIS, "to qualify as a specialty occupation, the position must meet one of the following requirements: (1) a bachelor's or higher degree or its equivalent is normally the minimum entry requirement for the position; (2) the degree requirement is common to the industry in parallel positions among similar organizations or, in the alternative, the position is so complex or unique that it can be performed only by an individual with a degree; (3) the employer normally requires a degree or its

individual must have at least a Bachelor's degree. Firm interviews conducted with the author suggest that U.S. firms typically use H-1B visas to hire international students at domestic universities.

There are five aspects of the H-1B program that are important in the context of this paper.

First, H-1B visas are tied to the firm, so it is possible to directly infer firm hiring responses to quantity constraints. Firms – not foreign workers – determine demand for H-1B visas. Legal and application fees are substantial; depending on the size of the company, the H-1B filing fee alone in 2017 was between \$1,710-\$6,460, not including the attorney fee.

Second, the H-1B application process is a two-stage process. This feature allows examination of both latent demand and realized supply. In the first stage, firms must file a Labor Condition Application (LCA) with the Department of Labor's Employment and Training Administration (DOLETA). This first stage measures demand. There is no limit – beyond cost – on the number of LCAs that a firm can file, so demand is measured independent of whether an H-1B is ultimately issued or not. In the second stage, after the LCA is approved, the firm must file an I-129 petition with USCIS, which makes the ultimate determination about the visa application. Constraints imposed by the H-1B cap are imposed in the second stage, where the final decision is made, so this stage measures realized supply. The two-stage structure of the H-1B application process allows me to measure exactly how constrained each firm was as the cap grew more restrictive over time by comparing – at the firm-level – LCA requests (demand) and issued H-1B visas (realized supply).

The next two features provide sources of identification from an econometric point of view.

Third, variations in the cap on H-1B visa supply provide a source of exogenous variation. The number of new H-1B visas that can be issued to private sector businesses has been subject to a cap since their inception in the Immigration Act of 1990. This cap is set by Congress and the President. Figure 1 plots the cap on the number of H-1B visas by fiscal year. The cap only applies to new H-1B visas issued to private sector businesses; there is no cap for the following categories: (1) those for non-profit firms, universities, and research labs, (2) those that are an extension of an existing H-1B visa, (3) those that have an existing H-1B visa and are changing jobs during the period of the existing visa, and (4) citizens of countries with whom the United States has a relevant free trade agreement.

There are three discrete phases of interest in terms of hiring constraints over time. The first phase is one in which the hiring constraint was not binding: throughout most of the 1990s, the cap was set at

equivalent for the position; or (4) the nature of the specific duties is so specialized and complex that the knowledge required to perform the duties is usually associated with attainment of a bachelor's or higher degree.”

65,000 visas and applications rarely outstripped supply¹⁶. Phase 2 began in 1998-2000, when the cap was increased to 195,000 visas by the American Competitiveness and Workforce Improvement Act of 1998 and the American Competitiveness in the Twenty-First Century Act of 2000 (AC21). During this period, the cap limits were never reached. Phase 3 occurred when trends in increasing visa availability reversed in 2004 and the cap reverted to the initial level of 65,000 visas, although 20,000 additional visas were granted to applicants with a graduate degree in 2006 (to a total of 85,000). Since then, the cap has not changed, and it has been (increasingly) binding in every year since 2004. Due to data constraints, this paper focuses on Phase 2 and 3.

The fourth relevant characteristic of the H-1B program from an econometric point of view is the random variation that is introduced by the process by which H-1B visas are distributed. H-1B petitions are distributed in a first-come-first-served fashion or by lottery in especially high demand years. The process is illustrated in Figure 2. On the first business day of each April, USCIS begins accepting H-1B applications from firms seeking permits that will count towards the following fiscal year. Since the H-1B visa program operates on a first-come, first-served basis, petitions are accepted until the cap hits, at which point no more petitions are processed. The end of the application period is demarcated by the “final receipt date”, which is the date on which they receive enough applications to fill the remaining available permits under the cap. Any cap-subject petitions submitted after the final receipt date were automatically rejected. This date is announced by USCIS in a press release every year, and it varies every year, as shown in Table 1. On the date(s) that the available permits are exhausted, a computer-generated random selection process selects the petitions that will be processed. Firms have no way of knowing in advance what date the cap would be reached. The dates of the lottery are not announced in advance and are in fact unknown in advance; they are determined by the number of applications received on different dates. These dates are only made known to firms after the cap is reached. In April 2007 and 2008, USCIS received so many petitions within the first week that all cap-subject petitions were distributed by lottery for fiscal years 2008 and 2009. The lottery generated a random negative shock in the supply of foreign-born skilled workers to firms; the second empirical approach exploits the lottery-generated randomized variation from the H-1B visa lotteries in those two years, which allows for a causal interpretation of the effect of constrained foreign-born skilled worker supply on the offshoring of skilled jobs.

The final relevant characteristic is significant for measuring demand. The timing of petitions can be used to reveal whether an application is for a cap-subject H-1B visa. As described above, the prerequisite to filing an H-1B petition with USCIS is obtaining an approved LCA from the Department of Labor. An

¹⁶ Fiscal years 1997 and 1998 were the lone instances when the cap was reached.

LCA cannot be filed more than six months prior to the start of employment. In order to apply for a visa for the following fiscal year (beginning in October), one would expect that firm to file an LCA no earlier than April. Furthermore, an LCA is only valid for three years; the earlier the application submitted, the fewer months a foreign-born worker would be eligible to work. In short, without any restrictions on H-1B supply, one would expect all firms to apply for LCAs no earlier than April, and probably much later.

However, the rising demand for H-1B visas and the first-come, first-served nature of the distribution process changed firm behavior. Firms that needed cap-subject H-1B visas wanted to submit their petitions as early as possible (i.e. April) to ensure the submission would be before the final receipt date, which also meant the LCA application had to be submitted prior to April. Figure 3 illustrates the change in the timing of LCA applications; as demand for cap-dependent H-1B visas increased, LCA applications were filed earlier. This feature is relevant for determining which LCA applications were for cap-dependent H-1B petitions.

4. Data

I use a combination of four sources of data to generate a unique dataset that permits the analysis of the link between MNC hiring decisions and how they respond to high-skilled immigration constraints. The first dataset provides information about multinational activity, including employment and R&D expenditures. The second and third are particularly useful because they provide information about both the demand and the realized supply for foreign workers, because they consist of both H-1B visas issued to foreign-born workers at US firms, as well as those requested by US firms. The fourth dataset provides information about the innovative activity of the firms, as measured by patents.

4.1 Multinational activity data

The data that I use to examine multinational activity are firm-level data from the Bureau of Economic Analysis's (BEA) annual surveys on U.S. Direct Investment Abroad. BEA is under a congressional mandate¹⁷ to track investment into and out of the United States, and as such, their data comprise the most comprehensive available data on US multinational activity abroad. Of particular importance is that the data includes foreign affiliate employment, which is the primary variable of interest for this paper. The data are confidential, and only accessible at a restricted site at the Bureau of Economic Analysis in Suitland, Maryland.

¹⁷ By the International Investment and Trade in Services Survey Act. The data are collected for the purpose of producing publicly available aggregate statistics on the activities of multinational enterprises.

I constructed a panel dataset of this activity from 1994 through 2014.¹⁸ Each firm may report on a consolidated basis for multiple affiliates in the same country under certain conditions.¹⁹ Therefore, rather than conducting analysis at the affiliate level, I aggregate all foreign affiliate activity up to the host country level for a given parent firm for a given year. The panel contains 2,263 firms with multinational activity.

4.2 H-1B Data

The second step of the data construction is to measure the firm-level hiring patterns of foreign-born workers. The source of the information is worker-level application records from the U.S. Department of Labor (DOL) and worker-level approved H-1B petition data from the U.S. Customs and Immigration Services.

Measures of firm-level demand for H-1B visas come from the DOL Labor Condition Application (LCA) data. Before a firm can file a petition with US Customs and Immigration Services (USCIS), they must file an LCA with the DOL.²⁰ There is no limit (other than financial constraints) on the number of LCAs that a firm can file. The primary purpose of the LCA is for employers to attest to the employment details of H-1B applicants and affirm that the worker will be employed in accordance with U.S. law.²¹ This data set comprises 6.4 million records between 2001 and 2016, which I aggregate by employer-year and then link to the BEA data.

The LCA data do not contain information on which applications are for H-1B visas that would be cap-subject. This does not matter for the first empirical strategy, but for the second strategy, in order to measure excess demand due to H-1B cap constraints, I infer whether a given LCA application is for a cap-subject H-1B visa by looking at the date of the LCA application. I assume that any LCA filed between January and April with a work start date 5-6 months in the future represents demand for a cap-

¹⁸ The most extensive data are collected in benchmark years: 1994, 1999, 2004, 2009, AND 2014. The reporting requirement threshold varies by year, size of the affiliate, and the parent's ownership stake. BEA estimates values of some variables of some affiliates in non-benchmark years in order to estimate a consistent universe across years. I only use the reported data in this paper.

¹⁹ These conditions are that the affiliates operate in the same country and same industry classification or are integral parts of the same business operation.

²⁰ These applications have been made publicly available by the DOL since 2001, and contain information on the employer's name and address, the occupation code of and the wage offered to the worker, and the geographic location of employment for the position to be filled by the visa recipient.

²¹ There are four main labor conditions that they are required to meet: (1) recipients of the visa must receive the same or better wages and benefits as other similar company employees and as similar employees in the geographic area, (2) working conditions must be similar for all employees, (3) there must not be a "strike, lockout, or work stoppage" at the employment location when the LCA is signed and submitted, (4) any employee bargaining representatives must be notified of every application submitted.

subject H-1B visa for the following fiscal year. Any LCA filed according to a different timeline thus represents demand for non-cap-subject H-1B visas.

Measures of realized H-1B labor supply come from I-129 H-1B visa applications, obtained by FOIA request²². These data are used in the second identification strategy. The original dataset contains I-129 petitions from fiscal years 2004-2014²³, consisting of about 3.3 million petitions, with information on the final decision regarding each petition, the type of visa being requested, the beneficiary's country of birth, the employer name and location, the job code, compensation, and other administrative details. An I129 form is needed for many types of visas, but for the purpose of this paper, the most relevant is the H-1B visa.

Not all H-1B visas were affected by the cap. To identify the visas that were constrained, I first remove (1) those for non-profit firms, universities, and research labs, (2) those that are an extension of an existing H-1B visa, (3) those that have an existing H-1B visa and are changing jobs during the period of the existing visa, and (4) citizens of five countries that were effectively exempt from H-1B limits due to bilateral trade agreements (Australia, Canada, Chile, Mexico, and Singapore). Second, I remove petitions that were submitted after April (the month of the lottery); the cap was reached in April. Finally, I focus on petitions submitted for fiscal years 2008 and 2009 – the two years of the lottery. The remaining petitions comprise the realized H-1B labor supply among cap-dependent petitions. Excess demand for foreign labor due to H-1B restrictions can thus be measured by subtracting realized cap-dependent H-1B petitions from cap-dependent LCA applications.

An examination of the heterogeneity of H-1B petition filings by firm, industry, and country provides some intuition regarding which types of firms might be most impacted by H-1B restrictions, and where the expansion of foreign affiliate activity might be expected to take place. Figures 6 and 7 show the country and industry breakdowns of H-1B petition filings in 2017 and illustrate that the median H-1B visa petition is for an Indian worker in a computer-related occupation. Computer-related occupations accounted for 69% of H-1B petition filings in 2017, and 85% of H-1B petition filings were for workers from India or China. These results suggest that the firms most impacted by constraints on H-1B visas would be firms dependent on computer-related workers. They also suggest that I might expect to see large increases in foreign affiliate employment in India and China. Finally, Figure 7 shows the top 30 H-1B petition filing firms in FY 2016 and illustrates that there is also significant skewness in H-1B visas across firms; this suggests that a non-linear approach may be appropriate. Note that these are not all US

²² Chad Sparber, Kevin Shih, Anna Maria Mayda, and Giovanni Peri also generously shared their 2001 H-1B petition data to supplement the data that I received from the FOIA request. I am very grateful to them for their help.

²³ Not all cap-subject petitions that were rejected by USCIS are included.

multinational companies; the sample of firms in my analysis is somewhat different. Specifically, my data include US multinational firms only. They do not include US domestic firms, foreign multinational firms, or the Indian outsourcing firms who are the largest H-1B applicants in the US. The importance of this distinction is illustrated in Figure 8; based on median compensation alone, Indian outsourcing companies hire a very different set of workers than US multinational firms.

4.3 Patent Data

The final source of data provides information on the innovativeness of the firm. The data are constructed from US Patent and Trademark Office (USPTO) patent data and includes all utility²⁴ granted patent applications through 2017.

4.4 Final Dataset

The final dataset is at the firm-country-year level and contains 2,263 multinationals. Of the 2,263 multinationals, 28% filed at least one LCA in 2001, 29% applied for at least one USPTO patent, and 15% both had at least one LCA application in 2001 and had at least one USPTO patent at some point. Summary statistics are shown in Table 2.

5. Empirical Approach and Results

This section describes my estimation framework; I use two identification strategies. The first exploits the 2004 drop in the H-1B visa cap, while the second exploits randomized variation from the H-1B visa lotteries in high demand years. The results yield the same qualitative result and strongly support the hypothesis that restrictions on high skilled immigration caused increased foreign affiliate employment and made firms more likely to open up foreign affiliates abroad.

5.1 Identification Strategy 1: Exploiting the 2004 policy change

The identification in this strategy is based on a plausibly exogenous shock to high-skilled immigration supply: the sharp reduction in the annual H-1B cap in fiscal year 2004, shown in Figure 1. As described in the “Overview of the H-1B and L-1 Visa Programs” section, the cap was not binding in the years leading up to the reduction (1998-2004) but has been binding in every year since 2004.

²⁴ There are three types of patents: utility patents, design patents, and plant patents. Utility patents cover inventions that have a useful and specific function and make up the majority of all patents. Language referring to a patent typically refers to utility patents. Design patents protect aesthetic appearance, while plant patents are for the discovery or invention of plants that are asexually reproduced.

My empirical specification can be interpreted as a difference-in-difference estimator – similar to that used in Ashraf and Ray (2017), Ghosh et al. (2015), Kerr and Lincoln (2010), and Xu (2016) – where the treatment and control groups are categories of firms with different levels of H-1B dependency in 2001 (when the cap was not binding) and the treatment is the reduction in the annual H-1B cap in 2004. In other words, the difference-in-differences approach relies on pre-existing variation in demand for foreign-born skilled workers to identify how exogenous constraints in supply affected foreign affiliate employment. Accordingly, the regression compares the change in foreign affiliate employment before and after the policy change across multinationals, within the same firm and country, that were more dependent on H-1B visas prior to the policy change (the “treatment” group) and less dependent firms prior to the policy change (the “control” group). Figure 4 provides a graphical version of the strategy, and shows that while foreign affiliate employment growth for non-H-1B dependent firms remained fairly flat after the policy change, extremely H-1B dependent firms experienced rapid growth in foreign affiliate employment after the policy change. Furthermore, the trajectory of foreign affiliate employment growth of both types of firms remained parallel and quite flat prior to the policy change. The regression results confirm the associations in the raw data presented in Figure 4.

In the baseline specifications, H-1B dependency is defined as the total LCA applications for a given multinational divided by that multinational’s US employment in 2001, as in Kerr and Lincoln (2010) and Xu (2016). The dependency measure is calculated in a pre-policy year to help address the problem of reverse causality. The dependency measure is my preferred metric because it measures demand for H-1B visas, and it is measured independent of whether an H-1B visa is ultimately issued or not. Furthermore, because of the high cost of application, the dependency measure can be seen as reflecting real measured demand. Finally, the dependency measure closely mirrors DOL’s own measure of H-1B dependency, namely: “The determination as to whether an employer is H-1B dependent is a function of the number of H-1B nonimmigrants employed as a proportion of the total number of full-time equivalent employees employed in the United States.”²⁵

The regression specification is as follows:

$$\ln(FAemp_{ict}) = \alpha_c + \alpha_i + \alpha_t + \beta_1 H1BDep_i * policy_t + \epsilon_{ict}$$

where i indexes the firm, c indexes the country, and t indexes the year. Country, firm, and year fixed effects are captured by α_c , α_i , and α_t respectively. $FAemp$ is foreign affiliate employment. $Policy$ is a

²⁵ Labor Condition Application for Nonimmigrant Workers. ETA Form 9045CP – General Instructions for the 9035 and 9035E, U.S. Department of Labor.
https://icert.doleta.gov/library/ETA_Form_9035CP_2009_Revised_03.18.09.pdf

dummy variable, equal to one for all years after and encompassing 2004, and zero otherwise. This is interacted with *HIBDep*, which is defined in one of two ways. The first is the continuous version of H-1B dependency described above, where the number of LCA applications (the measure of H-1B demand) by firm *I* in pre-policy year 2001 is normalized by a multinational's employment in the US in pre-policy year 2001, while the second is a new variable equal to one if the firm was in the top group of H-1B dependency²⁶ in 2001, and zero if the firm had zero LCA applications in 2001. All standard errors are clustered at the firm level but are also robust to clustering at the country level. I expect β_1 to be positive; after the 2004 reduction in the H-1B cap, firms that were ex-ante more dependent on H-1B visas should be more affected by the policy change and therefore more likely to expand their foreign affiliate activity. Column 1 of Table 3 shows the results of the continuous version, while column 2 shows the results of the binary version. Both are statistically and economically significant. The coefficient in the binary version tells us that the 2004 policy change caused highly H-1B dependent firms to increase their foreign affiliate employment by 27% more than a non-H-1B dependent firm.

A modification of this specification that interacts year fixed effects with the treatment allows observation of time-varying effects and answers the question: was the firm response to the cap change a one-off effect, or did the effect grow over time as firms felt more constrained? The results are shown in Table 4 and provide evidence that firms accelerated offshoring over time as constraints became tighter.

The measure of foreign affiliate employment is not the ideal measure; I am interested in the substitution of high-skilled immigrants for high-skilled foreign affiliate employment. Fortunately, although aggregate foreign affiliate employment is not broken into more disaggregated pieces in most years, foreign affiliate R&D employment is reported in benchmark years, which are every five years. In this sample, the benchmark surveys are 1994, 1999, 2004, 2009, and 2014. Table 5 provides the results of the same differences-in-differences specification as that reported in Table 3, but with a different dependent variable: logged foreign affiliate R&D employment. These results provide additional evidence that the type of foreign affiliate employment affected by the H-1B restrictions were in fact skilled employment.

An Alternate Specification

An alternative to the more traditional differences-in-differences regression approach is a series of cross-sectional long-differenced regression specifications, as follows:

²⁶ Measured in the same way as in the earlier non-linear specification.

$$\Delta \ln(FA_emp_{ic,t-2001}) = \alpha_j + \alpha_c + \beta_1 \left(\frac{LCAapps}{US_emp} \right)_i^{2001} + \Delta \epsilon_{ic,t-2001}$$

where i indexes the firm, j indexes the industry, c indexes the country, and t is a post-policy year. As before, FA_emp is foreign affiliate employment in country c by firm i , $LCAapps$ is the measure of demand (the number of LCA applications) by firm i in pre-policy year 2001, US_emp is a multinational's employment in the US in pre-policy year 2001, and α_j and α_c capture industry (NAICS 4-digit level) and country time trends respectively. The dependent variable is the logged differenced foreign affiliate employment of firm i in country c , between year t and a pre-policy year (2001). I expect β_1 to be positive in post-policy change years (2005-2014) and null in pre-policy change years (2002-2004). After the 2004 reduction in the H-1B cap, firms that were more dependent on H-1B visas should be more affected by the policy change and therefore more likely to expand their foreign affiliate activity. Before the 2004 cap change, any pre-existing variation in demand for foreign-born skilled workers, as measured by H-1B dependency, should not be correlated with foreign affiliate employment growth. The main advantage of this specification is that I can control for industry time trends or other trends among sets of firms.

The results are shown in Table 6, where each column represents a long difference between 2001 and a later year. The results provide evidence that there were no existing pre-trends in the differences in foreign affiliate employment growth that correspond with the measure of H-1B dependency; β_1 is not statistically significant until 2005, the first year after the policy change. The results also show that firms that were one percentage point more H-1B dependent than average saw a 3-8% larger increase in foreign affiliate employment than average, as a result of increased immigration restrictions resulting from the 2004 cap drop. Since the average foreign affiliate in 2001 had 1,151 employees, the estimated effect at the mean is thus an increase of 35-90 employees at every foreign affiliate.

As with the differences-in-differences specification, the approach requires that pre-treatment trends in foreign affiliate employment were the same for the treatment and control groups (the parallel trends assumption). If one looks simply at how the level of foreign affiliate activity differed between our treatment and control groups, it is apparent that firms that filed more LCAs in 2001 were not identical to firms with fewer LCAs. For example, firms that filed large numbers of LCAs tend to do more R&D abroad but have fewer employees abroad. However, the difference in levels is not in of itself problematic for my specification; the threat to identification would be if my measure of H-1B dependency were correlated with pre-treatment changes in foreign affiliate employment.

To test for this possibility, I include the 1994-1999 pre-trend in the baseline specification. These results are shown in column two of Table 7, with the baseline results shown in column 1 for the purpose of comparison. The coefficient of interest gets slightly smaller, but remains positive and statistically significant, and hence lessening any concerns about endogeneity. Note that I only report the results for the 2013-2001 long difference; they are robust to choosing any post-2004 year, but due to space constraints and constraints from the use of confidential data, I only report the 2013 results here.

Another concern might be that results are driven by systematic growth rate differences across firms of different size, internationalization, or innovativeness to begin with. If, for instance high-patenting firms naturally expand their activity abroad more quickly than non-patenting firms, even within the same industry, then my results could reflect that correlation rather than the effect of the policy change. I therefore tested whether the coefficient of interest changes when including controls for the size or type of firm in terms of their 2001 employment quantile, in terms of their 2001 sales quantile, in terms of their 2001 R&D spending quantile, in terms of their 2001 total patenting quantile, and in terms of the number of foreign affiliates quantile in 2001. These controls are included as group fixed effects. The identification in these regressions, once various controls are added, is now based on the comparison of trajectories within the same industry and the same category of employment size or sales size or R&D size or patenting amount or degree of internationalization, depending on the control. Column 3 of Table 7 represents the results of a specification that contains indicators for each firm's 2001 patenting quantile; the other results can be requested separately but show the same result, which is that the main results remain qualitatively unaffected.

Exploring Firm Heterogeneity

This section examines what types of firms responded most to restrictions on high-skilled immigration flows. Multinational firms differ in their strategic needs and capabilities; these sources of firm heterogeneity will affect how different firms respond to legal restrictions on migration. I examine four sources of firm heterogeneity: R&D intensity, software-orientation, offshorability, and H-1B dependency. I also explore the sensitivity to the model specification. Finally, I examine host country heterogeneity: particularly focusing on differences between countries which possess the necessary raw human capital (such as India or China) versus countries where it is easy to make use of foreign affiliates with easy access to foreign born workers (such as Canada²⁷).

²⁷ Much has been made anecdotally of Canada as a destination for firms struggling with immigration constraints in the United States. See, for example, the Envoy 2019 Immigration Trends Report, where 38% of surveyed firms were thinking about expanding to Canada because their immigration policy is more favorable, and 21% already had at least one office there.

I begin by examining firm's R&D intensity, since I expect that the firms most responsive to restrictions on high skilled immigration would be firms that conduct high-skilled activity and are reliant on high-skilled human capital to do so²⁸. Column 4 of Table 7 shows the results of the long-differenced regression run on firms that were in the top 20% of R&D-intensive firms in 2001. I find that the coefficient is larger than on the entire firm sample suggesting that the results are driven by the firms conducting high-skilled activity, as expected.

I follow this by examining the importance of the H-1B program for firms that hire computer-related occupations, since Figure 6 showed that computer-related occupations account for the vast majority of H-1B petition filings. Specifically, I look at firms that patented extensively in software, where software patents are defined in the same way as Arora, Branstetter, and Drev (2013) and Branstetter, Drev, and Kwon (2018).²⁹ As expected, the coefficient of interest is larger than on the entire firm sample. I also examine in more detail those industries where offshoring is a viable substitute for immigration. As expected, if I restrict the sample to firms that operate within the most offshorable services sectors,³⁰ the results are even more striking; the coefficient nearly doubles. These results are available by request.

Of course, the relationship between growth in multinational foreign affiliate employment and the share of H-1B workers might be nonlinear, as suggested by the skewness in applications shown in Figure 7. I use a non-parametric approach to examine this possibility: I divide multinationals into groups according to their H-1B dependence in 2001. I create 7 categories of multinationals, where the base category is all multinationals with zero LCA applications, and the remainder are divided into five quantiles, with the top category divided into two groups. Again, I expect positive coefficients with especially large coefficients on high-dependency multinationals, and Column 5 of Table 7 shows exactly that pattern. In particular, I find large, positive, and statistically significant coefficients for the top bracket (with an LCA application-US employment ratio of at least 0.0158 in 2001). In other brackets, there is no statistical significance. These estimates suggest that the positive effect of H-1B restrictions on foreign affiliate employment is being driven by the heaviest users of H-1B visas.

²⁸ Where R&D intensity is measured as the ratio of R&D spending to sales.

²⁹ Branstetter, Drev, and Kwon (2018). This methodology uses a set of keywords (e.g. "computer program" or "software") associated with software-based technologies as defined by Bessen and Hunt (2007) to define one set of patents as software, and uses a narrow set of IPC categories as defined in Graham and Mowery (2003) to define another set. The final population of software patents is defined as the union of these two sets of patents. I define the software intensity of a firm as the software stock of all USPTO citation-weighted patents applied for by that firm (in any location) by 2001. Here, I restrict to the top quartile of software-patenting firms.

³⁰ As defined in (Jensen and Kletzer 2010). It includes industries such as software publishing and scientific research and development services.

Exploring Country Heterogeneity

I also explore country heterogeneity. The largest countries of origin for H-1B visa holders are China and India – as shown in Figure 5 – while many of the prominent examples of companies opening foreign affiliates abroad in response to H-1B restrictions are concentrated in Canada³¹. Canada is a special case for US firms; cities like Vancouver and Toronto are geographically close to Silicon Valley and other multinational headquarters, but Canada has much less restrictive high-skilled immigration policies than the US. These characteristics mean that the fixed costs of offshoring are relatively low. These facts suggest that the expansion of foreign affiliate activity could operate through two channels: (1) a direct channel, whereby multinationals expand foreign affiliate activity in countries where the raw human capital they need is located (e.g. India or China), or (2) an indirect channel, whereby multinationals expand foreign affiliate activity in countries like Canada where it is easy to open foreign affiliates housing immigrants from other countries. Canada is also a special case due to geographic proximity. To test the relative effects, I construct two samples - one of foreign affiliates in China and India (the “raw human capital” countries), and one of foreign affiliates in Canada – and run the baseline differences-in-differences regression model on these different samples, shown in Columns 2 and 3 of Table 8. This approach is an imperfect way of measuring direct and indirect flows, but it does provide some sense of the possible operational channels. The effects are statistically significant, positive, and larger than the base sample for both subsets, suggesting that a lot of the expansion was concentrated in these three countries. However, a fourth regression – shown in Column 4 – that excludes Canada, India, and China shows that these three countries are not the only ones impacted. The coefficient is smaller when the three countries are excluded, but it is still economically and statistically significant. Hence, while expansion in response to immigration restrictions was concentrated in Canada, China, and India, they were not the only countries affected.

Extensive Margin

The results thus far show the effect of immigration restrictions on foreign affiliate employment at the intensive margin; because foreign affiliate employment was logged in all regressions to this point, regressions so far have only captured the change in employment at foreign affiliates that existed in both 2001 and the later post-policy period. It does not capture the effect of any foreign affiliates that were opened after the policy change in response to the policy. Both effects are of interest here, and in fact, any impact on the extensive margin represent a much stronger strategic response by firms – and a larger long-term impact – than an impact on the intensive margin.

³¹ <http://www.talenteconomy.io/2017/06/19/tighter-immigration-policy-pushes-firms-open-foreign-satellite-offices/>

To measure the extensive margin effect, I use the long-differenced regression specification, but I change the dependent variable to a binary variable equal to one if the firm has a foreign affiliate in a given country by the post-policy year and equal to zero otherwise. I use a linear probability model so that I can include time trends, but the results are robust to a logit model. The results are shown in Table 9, and illustrate a statistically significant and consistent response on the extensive margin – no matter the year – to the H-1B cap change in 2004. An important caveat to the results must be made here: as noted in the data section, each firm may report on a consolidated basis for multiple affiliates in the same country under certain conditions,³² so I aggregate all foreign affiliate activity up to the host country level for a given parent firm for a given year. This implies that effects on the extensive margin will only be observed when a firm enters a new country; the extensive margin results will not, for example, reflect a firm opening a foreign affiliate in Vancouver if it already has a foreign affiliate in Toronto. Hence, the results are likely smaller than the true effect.

Table 10 drills slightly deeper into the extensive margin results using 2013³³. Column 1 of Table 10 shows that that the 2004 cap decline made a one percentage point more H-1B dependent firm 0.2% more likely to open a foreign affiliate in a new country than the average firm by 2013. Column 2 shows the same regression specification, but instead examines the likelihood of a multinational initiating R&D activity abroad. Here, the coefficient is slightly larger, as expected, since one would expect skilled immigration restrictions to affect the opening of an R&D lab, but not necessarily to affect the opening of a new manufacturing facility.

The coefficients in Columns 1 and 2 are fairly small, even if they are statistically significant. This is in large part due to the measurement issue described above, as well as the preponderance of zeros in the data; the average US multinational firm in the BEA data is active in about two countries, while the dataset includes 50 countries. To counteract this noise, Column 3 shows the same regression but for a subset of the 25 countries with the most activity in the data, and with the binary version of the independent variable. As expected, the coefficients become substantially larger. The results show that the 2004 cap decline made highly H-1B dependent firms 6% more likely to start conducting R&D in a new country than a non-H-1B-dependent firm.

³² These conditions are that the affiliates operate in the same country and same industry classification or are integral parts of the same business operation.

³³ The results are robust to the choice of another year; I chose 2013 in response to space constraints.

Summary and Additional Robustness Checks

In combination, the results show that immigration restrictions had an economically and statistically significant effect on offshoring, both on the intensive and extensive margins. But the effect is not spread throughout all firms; the firms that responded most strongly were firms that (1) were R&D-intensive, (2) operated in industries where services could be easily offshored, and (3) depended heavily on H-1B visas prior to the restrictions taking place. The effect is also not geographically spread out; Canada, China, and India have seen the largest expansion in US multinational foreign affiliate employment.

I implement several additional robustness checks. First, I use 2002 as the base year instead of 2001. Second, instead of normalizing applications by US employment as the H-1B dependency measure, I use the count of applications. The results are robust to both of these robustness checks. A third robustness check addresses the concern that LCAs are an imperfect measure of demand for skilled H-1B visas. In particular, we know that some firms continuously file LCA applications and only utilize some of them. In addition, some of these applications are for continuing and transfer H-1B visa applications. A robustness check that measures H-1B dependency in 2001 – our treatment measure – using H-1B petitions to USCIS in 2001 rather than LCA applications to DOL in 2001 can be found in the Appendix and confirms that the results are not driven by the imperfectness of LCAs as a measure of demand. Finally, the Appendix also includes a version of the baseline differences-in-differences that combines the extensive and intensive margin by using the inverse hyperbolic sine of foreign affiliate employment as the dependent variable rather than the natural log.

Effect on the Share of Foreign Affiliate Employment Abroad

The analysis above aimed to establish that the level of foreign affiliate employment increased in response to skilled immigration restrictions. However, it did not establish whether the share of foreign affiliate employment increased. In other words, did US multinational firms simply increase their foreign affiliate employment while also increasing their domestic employment, or was there a shift in the location of their employees? An increase in the share of employment abroad in response to restrictions on skilled immigration would provide even stronger evidence in support of a substitution effect.

I use a panel difference-in-differences regression approach similar to those used in the analysis above to address this question:

$$Share_{ict} = \alpha_c + \alpha_i + \alpha_t + \beta_1 H1BDep_j * policy_t + \epsilon_{ict}$$

where $Share_{ict}$ is the share of firm i 's employment in country c in time t relative to firm i 's total employment. The H-1B dependency measure is the same measure as used in all prior regressions up to

this point and can be shown in either continuous or binary form. Table 12 shows the results; that while on average there is no clear evidence of any shift in the share of employment, when the sample is constrained to those countries where the expansion of activity as a direct response to skilled immigration restrictions was concentrated, there was a clear shift towards those countries. US multinationals appear to have increased the share – not just the levels – of their total employment and patenting to three main countries: Canada, India, and China.

Estimation Issues

The main threat to identification comes from any shocks correlated with both the timing of the H-1B policy and its effects across firms. In particular, the tech bubble in the late 1990s and early 2000s may have been correlated with increases in the cap. After the bubble burst, the H-1B visa cap was higher than average and the economy experienced a downturn. To the extent that the recession particularly affected H-1B dependent firms, the estimates could be biased. The direction, however, is unclear. They may have been more likely to increase foreign affiliate activity to escape the recession in the US, which would lead to upward bias, or they may have been more likely to shrink their firms, which would lead to a downward bias. However, the robustness of the results to the inclusion of industry time trends in all regressions suggests that this is not problematic; any unobserved demand shocks for highly skilled workers would need to vary across firms within the same industry for there to be any bias.

A separate concern surrounds the parallel trends assumption and any anticipation of the cap decline; namely, did firms behave differently leading up to the policy change in anticipation of soon experiencing immigration constraints, and did these behavioral differences correlate with my constructed measures of H-1B dependency? Figure 4 does not show any clear change in trends prior to 2004, and more rigorous tests, shown in Tables 3-7, further support the case that firms did not change behavior in advance. There is also qualitative evidence to suggest that firms did not change behavior in advance; an immigration lawyer told the author in discussion surrounding this policy change that: “my clients weren’t expecting it...and keep in mind that at that time, the cap wasn’t being met...[the firms thought that even if] the cap is not increased...who cares, we don’t meet it anyway.”

Both concerns are further ameliorated by the results of a second identification strategy, which does not suffer from the same sources of potential bias and yet produces consistent results.

5.2 Identification Strategy 2: Utilizing the Random Lottery Feature of the H-1B Application Process

The difference-in-differences approach relies on pre-existing variation in demand for foreign-born skilled workers to identify how constraints in supply – induced by a change in the cap – affected foreign affiliate employment. It is unable, however, to measure the precise constraints firms faced as the cap grew more restrictive. I therefore also take another approach that does measure exactly how constrained each firm was as the cap grew more restrictive over time by comparing – at the firm-level – LCA requests (demand) and issued H-1B visas (realized supply). A feature of the H-1B allocation system – the lottery – allows for a causal interpretation of the effect of constrained foreign-born skilled worker supply on the offshoring of skilled jobs.

The identification in this strategy exploits random variation in the allocation of H-1B workers across U.S. multinational firms resulting from the H-1B lotteries of 2007 and 2008. In both of those years, the number of cap-dependent H-1B visa petitions submitted within the first month of the filing period far exceeded the annual limit of available permits in those years, as shown in Table 1. In those years, all petitions received by the final receipt date (April 3 and April 8 respectively) were put through a computer-generated random selection process that selected which petitions would be processed. This produced a random shock to the supply of H-1B workers; some firms were successful in the lottery, while others were not. My approach exploits this random H-1B variation.

The dependent variable in this approach is the same as the differences-in-differences approach: the change in foreign affiliate employment. However, instead of regressing the change in foreign affiliate employment on a measure of the firm's pre-policy-change H-1B dependency, I regress it on a measure of excess demand for foreign labor that is driven by exogenous supply shocks. Following Peri, Shih, and Sparber (2015a), I calculate excess demand as the difference between the firm-level demand for new H-1B workers (LCA applications that were filed early) and the firm-level capped supply of H-1B workers (the lottery allocation of permits). I scale this absolute measure of excess demand by the firm's US-based employment in 2007; an equivalent number of H-1B permits denied to two firms will represent a much larger shock for firms with few workers in the US as opposed to for firms with many workers in the US.

There are two mechanisms at work in this approach, both generating variation in normalized excess demand across firms. The first is the same mechanism at work in the differences-in-differences approach: firms that are more H-1B-dependent will feel the effects of H-1B supply constraints more acutely than those that do not hire many H-1B workers. The second is the unexpected supply shock coming from the lottery.

I regress the change in foreign affiliate employment growth between a pre-lottery year (2005) and a post-lottery year (2010-2014) on the firm-level excess demand in the two lottery years combined (2007 and 2008), as shown in the following regression specification:

$$\Delta \ln(FA_emp_{ic,t-2005}) = \alpha_j + \alpha_c + \beta_1 \left(\frac{ExcessDemand_i^{2007+2008}}{US_emp_i^{2007}} \right) + \beta_2 LCA_i^{07+08} + \Delta \epsilon_{ic,t-2005}$$

where as before, i indexes the firm, j indexes the industry, c indexes the country, and t is a post-lottery year. FA_emp is foreign affiliate employment in country c by firm i , $ExcessDemand$ is the measure of excess demand (the number of LCA applications minus the number of H-1B permits received) by firm i in lottery years 2007 and 2008 combined, US_emp is a multinational's employment in the US in 2007, LCA_i^{07+08} controls for the number of LCA applications a given firm submitted, and α_j and α_c capture industry (NAICS 4-digit level) and country time trends respectively. Results are robust to the exclusion or inclusion of the control for number of LCA applications. The dependent variable is the logged differenced foreign affiliate employment of firm i in country c , between a post-lottery year (t) and a pre-lottery year (2005). I expect β_1 to be positive; firms that lost a larger share of their H-1B petitions should be more likely to expand their foreign affiliate activity. All standard errors are clustered at the firm level.

Table 12 shows that β_1 is indeed significantly positive. The columns display foreign affiliate employment growth one, two, three, four, and five years after the lottery and illustrate that there was a persistent positive effect. The results show that a random negative shock to H-1B supply equal to one percentage point of initial employment caused an increase in the foreign affiliate growth rate of between 12 and 16%. This positive effect is consistent with the results from the differences-in-differences approach³⁴, and provides further support for the hypothesis that restrictions on high skilled immigration cause increased offshoring of high-skilled jobs.

Translating these coefficients into the number of jobs offshored, I find that about 0.3 foreign affiliate jobs were created for every unfilled H-1B position. To calculate this number, I divided excess demand among US multinationals during the two relevant lottery years by the predicted change in foreign affiliate employment. To estimate the predicted change in foreign affiliate employment, I multiplied each firm's normalized excess demand by the coefficients identified in Table 12. I then multiplied this value by logged foreign affiliate employment in 2005 and took the exponential to calculate the predicted change in

³⁴ The coefficients are not directly comparable. The two approaches have different samples (the first is much bigger since it includes multinationals that have never applied for an H-1B while the second only includes the subset of multinationals applying for LCAs in one of those two years). Furthermore, the key regressor is measured differently.

foreign affiliate employment for each firm. I did this for each firm and then summed across firms. Table 13 displays these figures for US multinationals in aggregate. Column 1 shows excess demand. Column 2 shows the predicted increase in foreign affiliate employment. The final column displays the number of offshored jobs lost per unfilled H-1B position.

This 0.3 estimate of substitution should be treated with caution and is likely a lower bound for several reasons. First, it relies on calculations on the intensive margin and does not consider increased foreign affiliate employment on the extensive margin. Second, there are likely at least some firms that did not submit their petition(s) in time to be considered for the lottery, which means that their behavior is not captured in the analysis. Third, this estimate only captures how existing firms modified their strategy; new firms born into the visa-restrictive world may have incorporated offshoring into their strategy from the beginning. Finally, the estimate does not include foreign multinational company behavior, and one would expect that foreign MNCs would be even more likely to substitute foreign-born US-based skilled workers for foreign-born skilled workers at headquarters or other foreign affiliates.

One might be concerned that some firms anticipated the lottery and increased their submitted applications to improve their chances of winning the lottery. The two lottery years in the empirical analysis were selected precisely to address this concern; these were the first two years in which a lottery was held to distribute all H-1B visas. To further address this concern, Table 14 duplicates the same analysis, but with only the 2007 lottery, and finds the same qualitative result: that a random negative shock to H-1B supply caused an increase in the foreign affiliate growth rate.

6. Conclusion

This paper explores the relationship between immigration and offshoring by measuring the effect of high-skilled immigration restrictions on increases in foreign affiliate jobs, and finds that they serve as substitutes. I utilize a unique firm-level panel dataset that links firm-level H-1B visa data and firm-level data on the financial and operating characteristics of US multinational firms and their foreign affiliates. I utilize two identification strategies, the first of which exploits the 2004 drop in the H-1B visa cap, while the second exploits randomized variation in firm-level excess demand from the H-1B visa lotteries in high demand years.

Both strategies yield the same result: that foreign affiliate employment increased as a direct response to increasingly stringent restrictions on H-1B visas. This effect is driven on the extensive and intensive side; firms were more likely to open foreign affiliates in new countries in response, and employment

increased at existing foreign affiliates. The effect is strongest among R&D-intensive firms in industries where services could more easily be offshored. The effect was somewhat geographically concentrated: foreign affiliate employment increased both in countries like India and China with large quantities of high-skilled human capital and in countries like Canada with more relaxed high-skilled immigration policies and closer geographic proximity. These empirical results also are supported by interviews with US multinational firms and an immigration lawyer.³⁵

Despite the outsized role that multinational firms play in the economy – for example, US multinational firms are responsible for 80% of US R&D and employ about ¼ of US private employees – policy debates surrounding immigration have largely overlooked the fact that multinational companies faced with decreased access to visas for skilled workers have an offshoring option, namely, hiring the foreign labor they need at their foreign affiliates. This is the first paper to provide evidence that multinational firms do in fact utilize this option – both at the extensive and the intensive margin – and to examine the relationship between foreign affiliate employees and immigrants, in contrast to the relationship between immigrants and native-born workers.

The results have important implications for understanding how multinational firms respond to artificial constraints on resources and how they globally re-distribute those resources. The results also have important policy implications; the offshoring of jobs appears to be an unforeseen consequence of restricting skilled immigration flows. Even if H-1B immigrants displace some native workers, any policies that are motivated by concerns about the loss of native jobs should consider that policies aimed at reducing immigration have the unintended consequence of encouraging firms to offshore jobs abroad.

The finding that skilled foreign-born workers will be hired at foreign affiliates rather than in the US also has important implications for the innovative capacity of the US. Skilled immigrants have been shown to have outsized impacts on innovation in the host country through spillovers (e.g. Hunt et al. 2017; Moser, Voena, and Waldinger 2014). The spatial diffusion of these spillovers disappears with distance (Burchardi et al. 2019) since innovative spillovers are geographically localized (Jaffe 1986; Jaffe, Trajtenberg, and Henderson 1993; Marshall 1920). From a nationalistic perspective, this is problematic; if skilled foreign-born workers are at a US firm's foreign affiliate instead of in the US, the innovative spillovers that they generate will go to another country instead. Furthermore, the finding that immigrants often are not equally innovative outside the United States (Kahn and Macgarvie 2016) has even wider welfare implications. In short, restrictive H-1B policies could not only be exporting more jobs and businesses to countries like Canada, but they also could be causing the U.S.'s innovative capacity to

³⁵ Quotes from these interviews can be found in the Appendix.

fall behind. Exploring the effects on the location and direction of innovation are the subject of ongoing effort.

References

- Agrawal, Ajay, John McHale, and Alex Oetl. 2018. "Does Scientist Immigration Harm US Science? An Examination of Spillovers." *National Bureau of Economic Research Working Paper Series No. 24519*.
- Akcigit, Ufuk, John Grigsby, and Tom Nicholas. 2017. "Immigration and the Rise of American Ingenuity." *NBER Working Paper*.
- Argote, Linda, Bill Mcevily, and Ray Reagans. 2003. "Framework and Review of Emerging Themes Managing Knowledge in Organizations: An Integrative Framework and Review of Emerging Themes." *Management Science* 49(August 2014):571–82.
- Arora, Ashish, Lee G. Branstetter, and Matej Drev. 2013. "Going Soft: How the Rise of Software-Based Innovation Led to the Decline of Japan's IT Industry and the Resurgence of Silicon Valley." *Review of Economics and Statistics* 95(3):757–75.
- Arora, Ashish and Alfonso Gambardella, eds. 2005. *From Underdogs to Tigers: The Rise and Growth of the Software Industry in Brazil, China, India, Ireland, and Israel*. Oxford University Press.
- Ashraf, Rasha and Rina Ray. 2017. "Human Capital, Skilled Immigrants, and Innovation." *Working Paper*.
- Audretsch, David B. 1998. "Agglomeration and the Location of Innovative Activity." *Oxford Review of Economic Policy* 14(2):18–29.
- Bessen, James and Robert M. Hunt. 2007. "An Empirical Look at Software Patents." *Journal of Economics and Management Strategy* 16(1):157–89.
- Borjas, George J. 2005. "The Labor-Market Impact of High-Skill Immigration." *American Economic Review* 95(2):56–60.
- Borjas, George J. 2003. "The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *Quarterly Journal of Economics* (November):1335–74.
- Borjas, George J. and Kirk B. Doran. 2015. "How High-Skill Immigration Affects Science: Evidence from the Collapse of the USSR." *Innovation Policy and the Economy* 15(1):1–25.
- Bound, John, Murat Demirci, Gaurav Khanna, and Sarah Turner. 2014. "Finishing Degrees and Finding Jobs: U.S. Higher Education and the Flow of Foreign IT Workers." *NBER Working Paper 20505*.
- Bound, John, Gaurav Khanna, and Nicolas Morales. 2017. "Understanding the Economic Impact of the H-1B Program on the U.S." *Population Studies Center Research Report* 16–857.
- Branstetter, Lee, Matej Drev, and Namho Kwon. 2018. "Get With the Program: Software-Driven Innovation in Traditional Manufacturing." *Management Science*.
- Branstetter, Lee, Britta Glennon, and J. Bradford Jensen. 2018. "The IT Revolution and the Globalization of R&D." in *Innovation Policy and the Economy, Volume 19*, edited by J. Lerner and S. Stern. University of Chicago Press.
- Burchardi, Konrad B., Thomas Chaney, Tarek A. Hassan, Lisa Tarquinio, and Stephen J. Terry. 2019. "Immigration, Innovation, and Growth." *NBER Working Paper*.
- Card, David. 2012. "Comment: The Elusive Search for Negative Wage Impact of Immigration." *Journal of the European Economic Association* 10(1):211–15.

- Card, David. 2001. "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration." *Journal of Labor Economics* 19(1):22–64.
- Card, David. 2009. "Immigration and Inequality." *American Economic Review* 99(2):1–21.
- Chellaraj, Gnanaraj, Keith E. Maskus, and Aaditya Mattoo. 2008. "The Contribution of International Graduate Students to US Innovation." *Review of International Economics* 16(3).
- Choudhury, Prithwiraj. 2016. "Return Migration and Geography of Innovation in MNEs: A Natural Experiment of On-the-Job Learning of Knowledge Production by Local Workers Reporting to Return Migrants." *Journal of Economic G* 16(3):585–610.
- Choudhury, Prithwiraj and Do Yoon Kim. 2018. "The Ethnic Migrant Inventor Effect: Codification and Recombination of Knowledge Across Borders." *Strategic Management Journal*.
- Docquier, Frédéric and Hillel Rapoport. 2012. "Globalization, Brain Drain, and Development." *Journal of Economic Literature* 50(3):681–730.
- Doran, Kirk, Alexander Gelber, and Adam Isen. 2016. *The Effects of High-Skilled Immigration Policy on Firms: Evidence from Visa Lotteries*.
- Doran, Kirk and Chungun Yoon. 2019. "Immigration and Invention: Evidence from the Quota Acts." *NBER Work*.
- Foley, C. Fritz and William R. Kerr. 2013. "Ethnic Innovation and U.S. Multinational Firm Activity." *Management Science* 59(7):1529–44.
- Freeman, Richard B. 2006. "People Flows in Globalization." *Journal of Economic Perspectives* 20(2):145–70.
- Friedberg, Rachel M. 2001. "The Impact of Mass Migration on the Israeli Labor Market." *Quarterly Journal of Economics* 116(4):1373–1408.
- Ganguli, Ina. 2015. "Immigration and Ideas: What Did Russian Scientists 'Bring' to the United States?" *Journal of Labor Economics* 33(S1):S257–88.
- Ghosh, Anirban, Anna Maria Mayda, and Francesco Ortega. 2015. "The Impact of Skilled Foreign Workers on Firms: An Investigation of Publicly Traded U.S. Firms."
- Gould, David M. 1994. "Immigrant Links to the Home Country : Empirical Implications for U.S. Bilateral Trade Flows." *The Review of Economics and Statistics* 76(2):302–16.
- Graham, Stuart J. H. and David C. Mowery. 2003. "Intellectual Property Protection in the U.S. Software Industry." Pp. 217–58 in *Patents in the Knowledge-Based Economy*, edited by W. M. Cohen and S. Merrill. Washington DC: National Academies Press.
- Hall, Bronwyn H. 2010. "The Internationalization of R&D." 1–36.
- Hernandez, Exequiel. 2014. "Finding a Home Away from Home: Effects of Immigrants on Firms' Foreign Location Choice and Performance." *Administrative Science Quarterly* 59(1).
- Hira, R. O. N. 2010. "The H-1B and L-1 Visa Programs." *Economic Policy Institute Briefing Paper* 1–19.
- Hunt, Jennifer et al. 2017. "How Much Does Immigration Boost Innovation ? 1 ^ Studied Immigration Impact on a Variety of Host Country." 2(2):31–56.
- Hunt, Jennifer. 2011. "Which Immigrants Are Most Innovative and Entrepreneurial? Distinctions by Entry Visa." *Journal of Labor Economics* 29(3):417–57.

- Hunt, Jennifer and Marjolaine Gauthier-Loiselle. 2010. "How Much Does Immigration Boost Innovation?" *American Economic Journal: Macroeconomics* 2(2):31–56.
- Iriyama, Akie, Yong Li, and Ravi Madhavan. 2010. "Spiky Globalization of Venture Capital Investments: The Influence of Prior Human Networks." *Strategic Entrepreneurship Journal* 4(2):128–45.
- Jaffe, Adam B. 1986. "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value." *American Economic Review* 76(5):984–1001.
- Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics* 108(3):577–98.
- Jensen, J. Bradford and Lori G. Kletzer. 2010. "Measuring Tradable Services and the Task Content of Offshorable Services Jobs Chapter." Pp. 309–35 in *Labor in the New Economy*, edited by K. G. Abraham, J. R. Spletzer, and M. Harper. University of Chicago Press.
- Kahn, Shulamit and Megan J. Macgarvie. 2016. "How Important Is U.S. Location for Research in Science?" *The Review of Economics and Statistics* 98(2).
- Kerr, Sari Pekkala, William R. Kerr, and William F. Lincoln. 2015. "Firms and the Economics of Skilled Immigration." in *Innovation Policy and the Economy*, vol. 15, edited by W. R. Kerr, S. Stern, and J. Lerner.
- Kerr, Sari Pekkala, William R. Kerr, and William F. Lincoln. 2015. "Skilled Immigration and the Employment Structures of US Firms." *Journal of Labor Economics* 33(S1):S186.
- Kerr, William R. 2008. "Ethnic Scientific Communities and International Technology Diffusion." *The Review of Economics and Statistics* 90(3):518–37.
- Kerr, William R. and William F. Lincoln. 2010. "The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention." *Journal of Labor Economics* 28(3):473–508.
- Kirkegaard, Jacob Funk. 2005. "IIE Working Paper 05-15 Outsourcing and Skill Imports: Foreign High-Skilled Workers on H-1B and L-1 Visas in the United States." *October* 3(DECEMBEr).
- Lofstrom, Magnus and Joseph Hayes. 2011. *H-1Bs: How Do They Stack up to US Born Workers?*
- Lowell, B. Lindsay. 2001. "Skilled Temporary and Permanent Immigrants in the United States." *Population Research and Policy Review* 20(1):33–58.
- MacGarvie, Megan. 2005. "The Determinants of International Knowledge Diffusion as Measured by Patent Citations." *Economics Letters* 87(1):121–26.
- Marshall, Alfred. 1920. *Principles of Economics*. London: Macmillan.
- Matloff, Norman. 2003. "On the Need for Reform of the H-1B Non-Immigrant Work Visa in Computer-Related Occupations." *University of Michigan Journal of Law Reform* 36(101):1–99.
- Mayda, Anna Maria, Francesc Ortega, Giovanni Peri, Kevin Shih, and Chad Sparber. 2017. "The Effect of the H-1B Quota on Employment and Selection of Foreign-Born Labor."
- Mithas, Sunil and Henry C. Lucas. 2010. "Are Foreign IT Workers Cheaper? U.S. Visa Policies and Compensation of Information Technology Professionals." *Management Science* 56(5):745–65.
- Moser, Petra, Alessandra Voena, and Fabian Waldinger. 2014. "German Jewish Emigres and US Invention." *American Economic Review* 104(10):3222–55.

- Nowrasteh, Alex. 2017. "The Economics of Immigration." *Cato Journal* 57(3):445–47.
- Oettl, Alexander and Ajay Agrawal. 2008. "International Labor Mobility and Knowledge Flow Externalities." *Journal of International Business Studies* 39(8):1242–60.
- Olney, William W. and Dario Pozzoli. 2018. *The Impact of Immigration on Firm-Level Offshoring*.
- Ottaviano, Gianmarco I. P., Giovanni Peri, and Greg C. Wright. 2013. "Immigration, Offshoring, and American Jobs." *American Economic Review* 103(5):1925–59.
- Ottaviano, Gianmarco I. P., Giovanni Peri, and Greg C. Wright. 2013. "Immigration, Offshoring, and American Jobs." *American Economic Review* 103(5).
- Ottaviano, Gianmarco I. P., Giovanni Peri, and Greg C. Wright. 2018. "Immigration, Trade and Productivity in Services: Evidence from U.K. Firms." *Journal of International Economics* 112:88–108.
- Patel, Pari and Keith Pavitt. 1991. "Large Firms in the Production of the World's Technology: An Important Case of 'Non-Globalisation.'" *Journal of International Business Studies* (June):1–21.
- Peri, Giovanni, Kevin Shih, and Chad Sparber. 2015a. "Foreign and Native Skilled Workers: What Can We Learn from H-1B Lotteries?" *NBER Working Paper* (21175).
- Peri, Giovanni, Kevin Shih, and Chad Sparber. 2015b. "STEM Workers, H-1B Visas, and Productivity in US Cities." *Journal of Labor Economics* 33(S1):S225–55.
- Polanyi, Michael. 1958. *Personal Knowledge, towards a Post Critical Epistemology*. Chicago, IL: The University of Chicago Press.
- Polanyi, Michael. 1966. *The Tacit Dimension*. Chicago, IL: The University of Chicago Press.
- Salzman, Hal and B. Lindsay Lowell. 2007. *Into the Eye of the Storm: Assessing the Evidence on Science and Engineering Education, Quality, and Workforce Demand*.
- Saxenian, Annalee. 2006. *The New Argonauts*. Cambridge, MA: Harvard University Press.
- Serapio, Manuel G. and Donald H. Dalton. 1999. "Globalization of Industrial R&D: An Examination of Foreign Direct Investments in R&D in the United States." *Research Policy* 28(2/3):303–16.
- Siedschlag, Iulia, Donal Smith, Camelia Turcu, and Xiaoheng Zhang. 2013. "What Determines the Location Choice of R&D Activities by Multinational Firms?" *Research Policy* 42(8):1420–30.
- Singh, Jasjit. 2008. "Distributed R&D, Cross-Regional Knowledge Integration and Quality of Innovative Output." *Research Policy* 37(1):77–96.
- Szulanski, Gabriel. 1996. "Exploring Internal Stickiness: Impediments to the Transfer of Best Practice Within the Firm." *Strategic Management Journal* 17(Winter Special Issue):27–43.
- Tambe, Prasanna and Lorin M. Hitt. 2009. *H-1B Visas, Offshoring, and the Wages of US Information Technology Workers*.
- Teece, D. J. 1977. "Technology Transfer by Multinational Firms : The Resource Cost of Transferring Technological Know-How." *The Economic Journal* 87(346):242–61.
- Thursby, Marie C. and Jerry G. Thursby. 2006. *Here or There? A Survey of Factors in Multinational R&D Location*. Washington, D.C.: The National Academies Press.
- Turner, Patrick S. 2017. *High-Skill Immigration and the Labor Market : Evidence from the H-1B Visa*

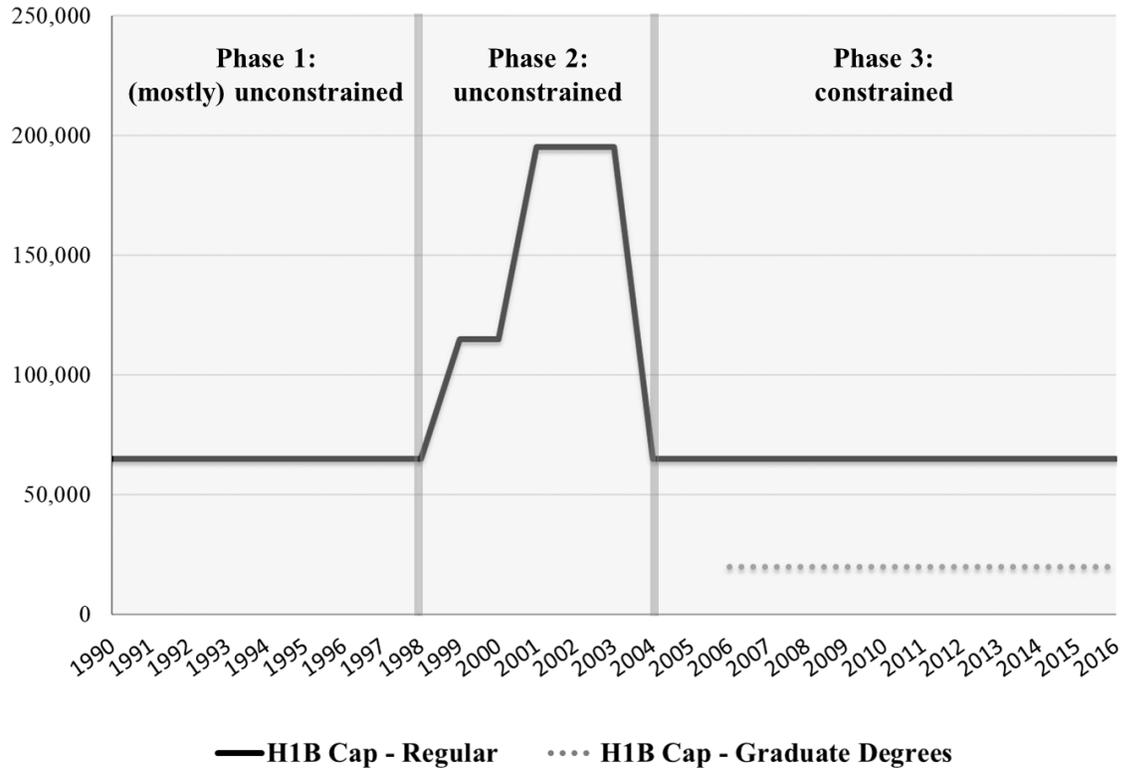
Program.

Wadhwa, Vivek, Annalee Saxenian, Ben Rissing, and Gary Gereffi. 2007. "America's New Immigrant Entrepreneurs."

Wu, A. 2017. *Skilled Immigration and Firm-Level Innovation: The US H-1B Lottery*.

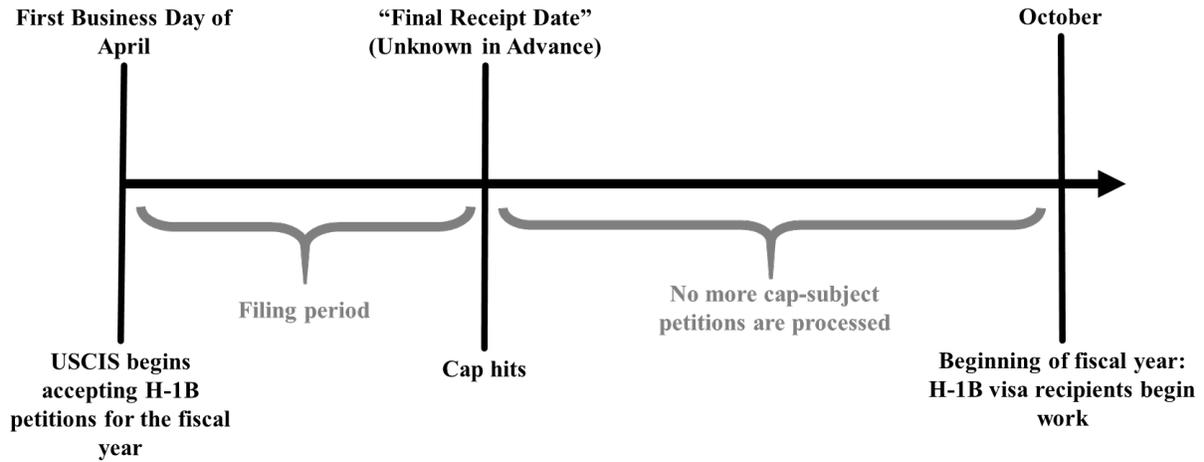
Xu, Sheng-Jun. 2016. "Skilled Labor Supply and Corporate Investment: Evidence from the H-1B Visa Program." *SSRN Electronic Journal* 1–52.

Figure 1: H-1B Cap Change Over Time



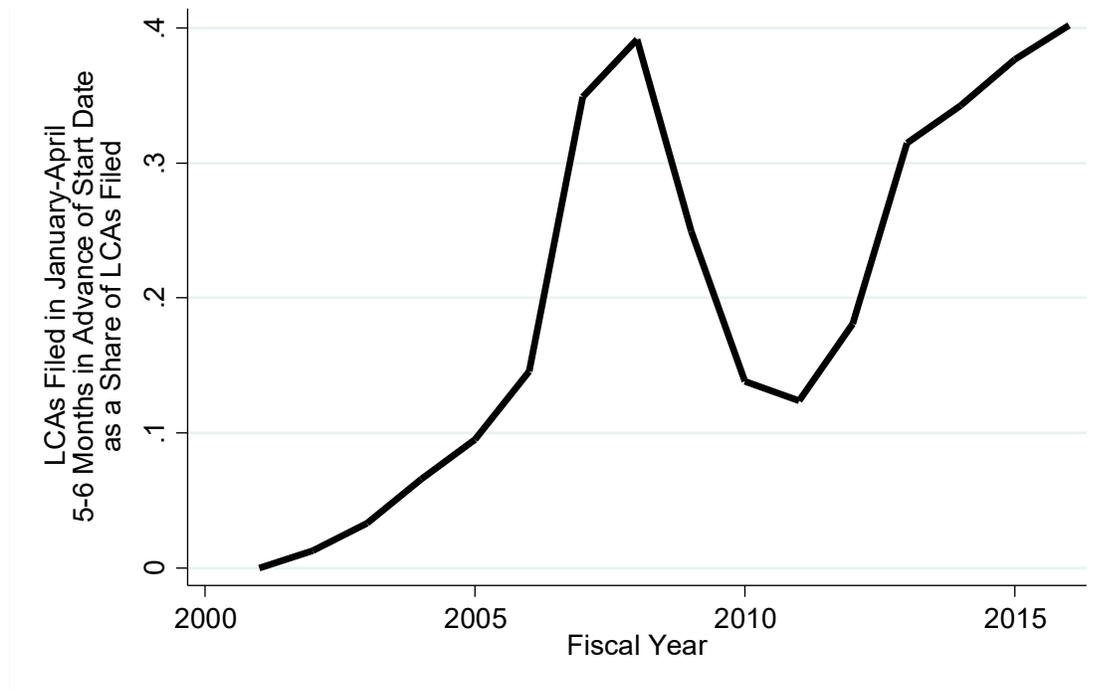
Notes: This figure plots the cap on the number of H-1B visas by fiscal year. Since the Immigration Act of 1990, there has been an annual cap on the number of new H-1B visas that can be issued to private sector businesses. This cap is set by Congress and the President. Throughout most of the 1990s, the cap was set at 65,000 visas and applications rarely outstripped supply. It was increased to 195,000 visas by the American Competitiveness and Workforce Improvement Act of 1998 and the American Competitiveness in the Twenty-First Century Act of 2000 (AC21). During this period, the cap limits were never reached. The AC21 stipulated that this reversion would happen in the absence of any additional legislation, but, despite a trend towards less restrictive labor laws, no legislation was enacted, and the cap level reverted back to 65,000. It was raised by 20,000 in 2006, but those additional 20,000 could only be used for applicants with a graduate degree. Although in the early 2000s, the cap was not binding, since the cap changed in 2004, it has been binding in every year. The identification in this paper exploits the sharp reduction in the annual H-1B cap in fiscal year 2004.

Figure 2: H-1B Visa Submission Timeline



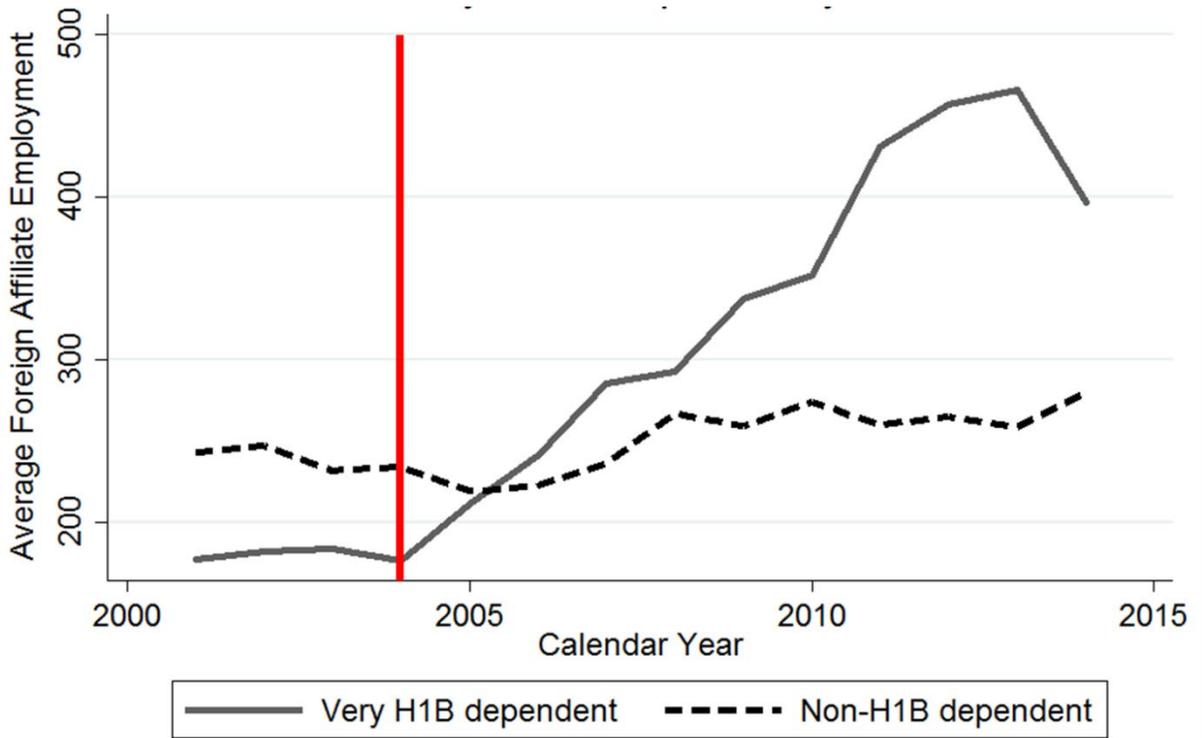
Notes: On the first business day of each April, USCIS begins accepting H-1B petitions from firms seeking permits that will count towards the following fiscal year. Petitions are accepted on a first-come, first-served basis until the cap hits, at which point no more petitions are processed. The end of the application period is demarcated by the "final receipt date", which is the date on which they receive enough applications to fill the remaining available permits under the cap. Any cap-subject petitions submitted after the final receipt date are automatically rejected.

Figure 3: In High-Demand Years, LCA Applications Were Filed Earlier



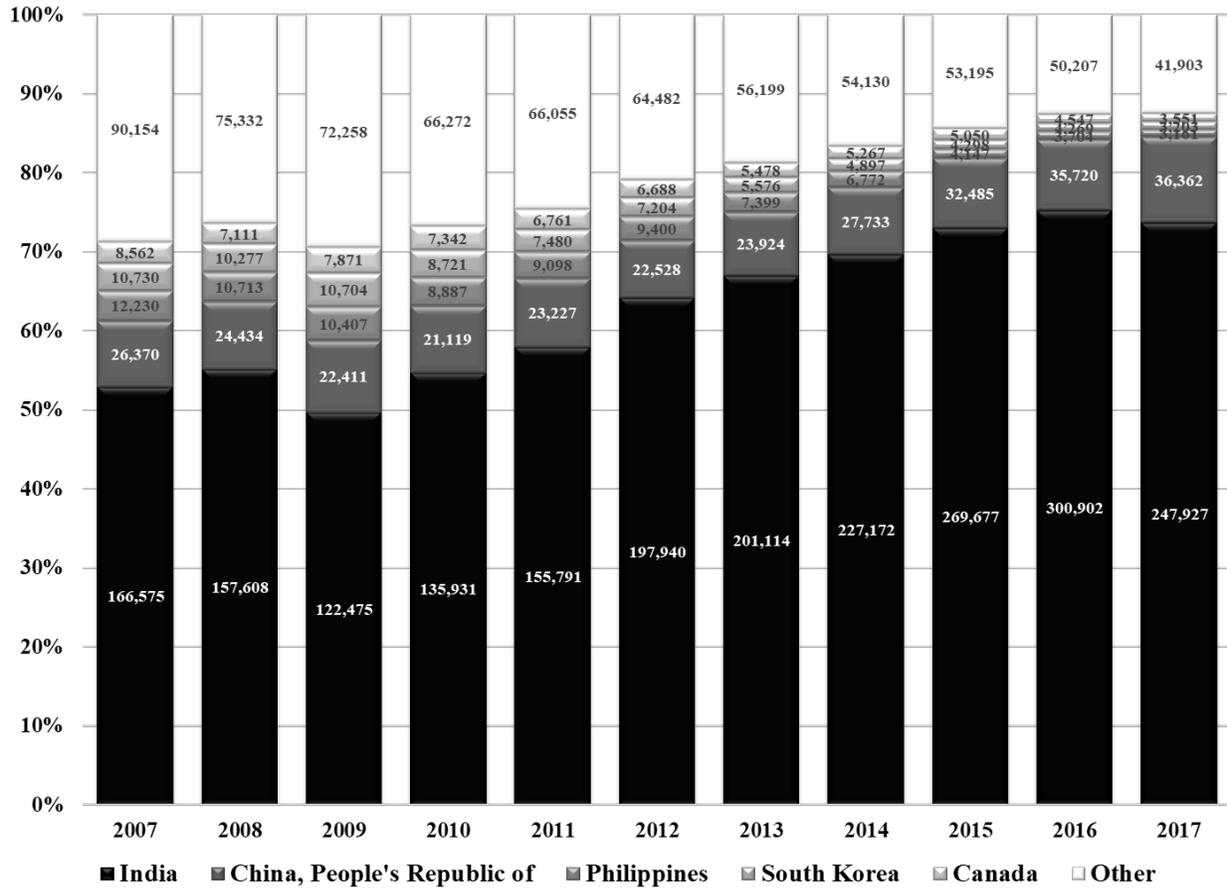
Notes: The prerequisite to filing an H-1B petition with USCIS is obtaining an approved Labor Condition Application (LCA) from the Department of Labor; the LCA is used in this paper as the primary measure of demand. LCA application data does not include information about whether a given LCA is for a cap-subject H-1B visa, but the timing of petitions can be used to reveal whether an LCA application is for a cap-subject H-1B visa. An LCA cannot be filed more than six months prior to the start of employment. In order to apply for a visa for the following fiscal year (beginning in October), one would expect that firm to file an LCA no earlier than April. Furthermore, an LCA is only valid for three years; the earlier the application submitted, the fewer months a foreign-born worker would be eligible to work. In short, without any restrictions on H-1B supply, one would expect all firms to apply for LCAs no earlier than April, and probably much later. However, the rising demand for H-1B visas and the first-come, first-served nature of the distribution process changed firm behavior. Firms that needed cap-subject H-1B visas wanted to submit their petitions as early as possible (i.e. April) to ensure the submission would be before the final receipt date, which also meant the LCA application had to be submitted prior to April. This figure illustrates the change in the timing of LCA applications; as demand for cap-dependent H-1B visas increased, LCA applications were filed earlier. I infer that LCA applications submitted in the first quarter of the calendar year are for cap-dependent H-1B petitions.

Figure 4: Growth in Average Foreign Affiliate Employment by H-1B Dependency



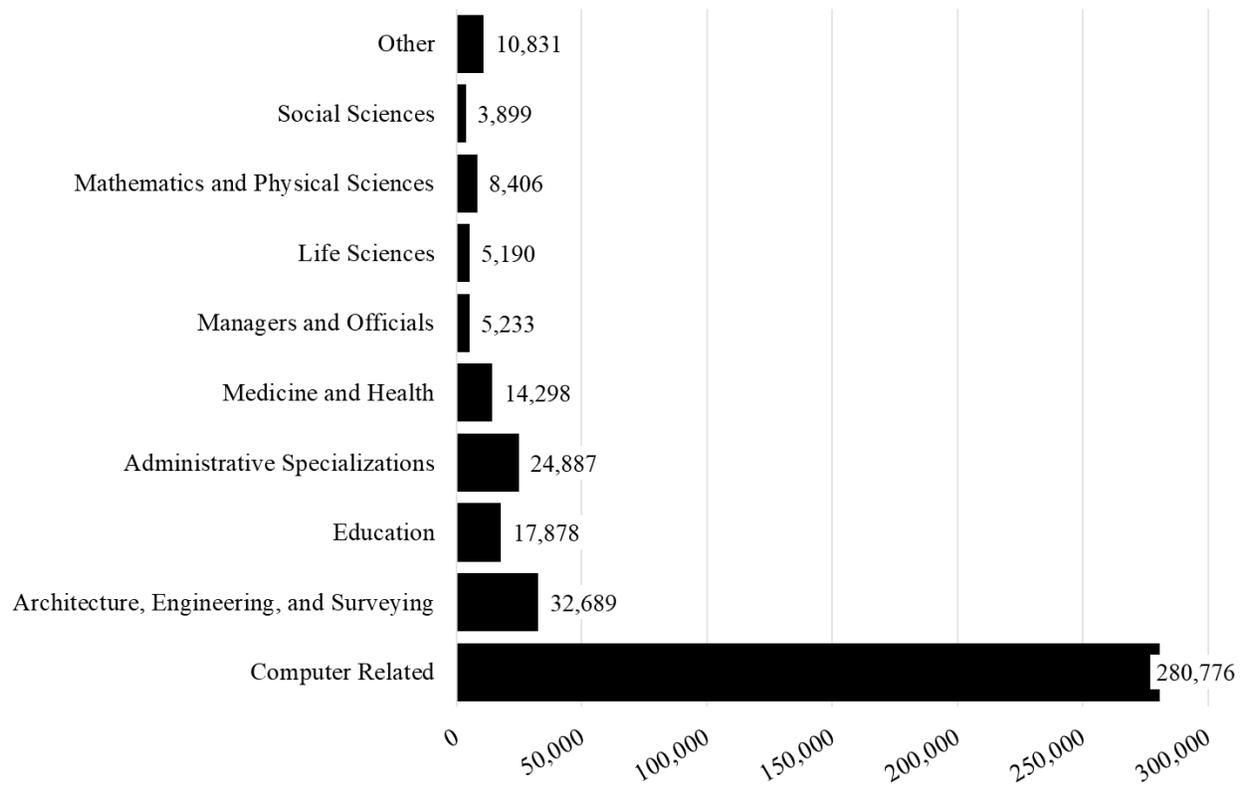
Notes: This figure plots foreign affiliate employment growth for non-H-1B dependent firms relative to very H-1B dependent firms. Non-H-1B-dependent firms had zero LCA applications in 2001, while very H-1B dependent firms were in the top category of H-1B dependency in 2001. H-1B dependency is defined as the total number of H-1B positions requested in LCA applications for a given multinational divided by that multinational’s US employment in 2001. Categories of H-1B dependency are measured throughout the paper by dividing non-zero LCA applications into five quantiles, with the top category divided into two groups. The red line demarcates the year of the 2004 policy change. The figure shows that while foreign affiliate employment growth for non-H-1B dependent firms remained fairly flat after the policy change, very H-1B dependent firms experienced rapid growth in foreign affiliate employment after the policy change. Furthermore, the trajectory of foreign affiliate employment growth of both types of firms remained parallel and quite flat prior to the policy change, in keeping with the parallel trends assumption. The regression results confirm the associations in the raw data.

Figure 5: H-1B Petitions Filed in FY2007-2017, by Beneficiary Country of Birth



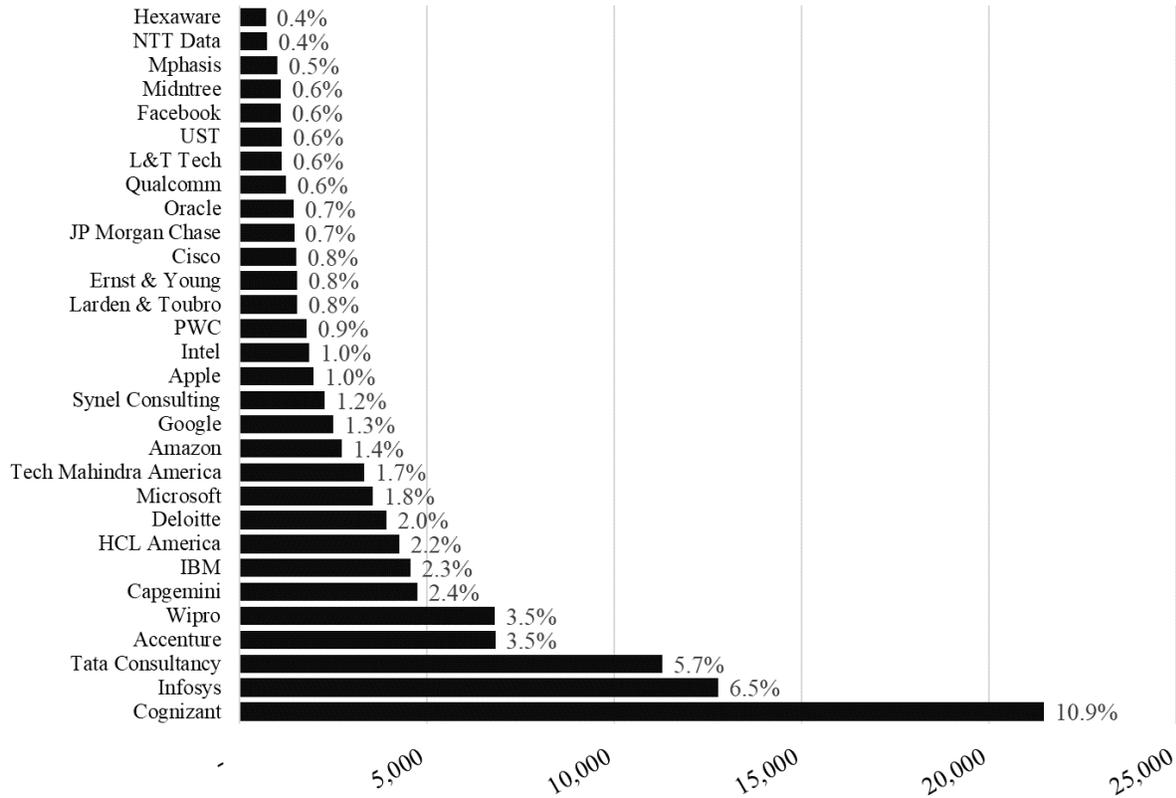
Notes: This figure shows the country breakdowns of H-1B petition filings between FY 2007 and 2017. It illustrates that the distribution of countries of origin is increasingly skewed; by 2017, 85% of H-1B petition filings were for workers from India or China. Note that these are petitions filed not petitions approved. Petition counts include both cap-subject and cap-exempt, initial and continuing employment. The data are publicly available from USCIS.

Figure 6: H-1B Petitions Filed in FY 2017, by Beneficiary Occupation Category



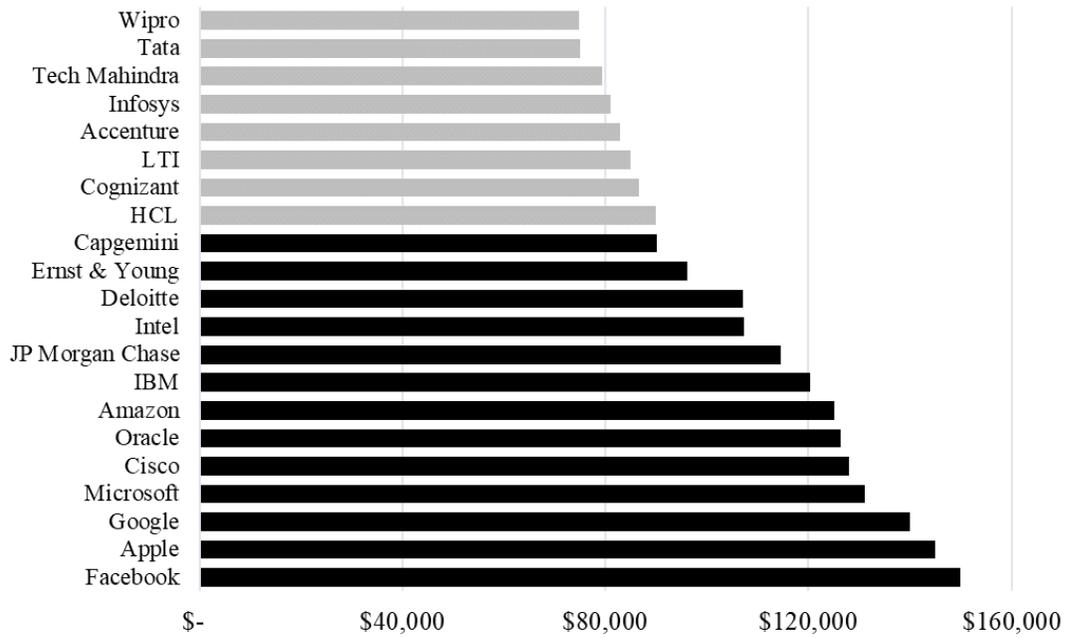
Notes: This figure shows the industry breakdowns of H-1B petition filings in 2017. It illustrates that the distribution of occupations is highly skewed; computer-related occupations accounted for 69% of H-1B petition filings in 2017. Note that these are petitions filed not petitions approved. Petition counts include both cap-subject and cap-exempt, initial and continuing employment. The data are publicly available from USCIS.

Figure 7: Total Number of H-1B Approved Petitions in FY 2016, for the Top 30 Firms



Notes: This figure shows the number of approved H-1B petitions for the top 30 firms in fiscal year 2016. All data are based on approved petitions during the fiscal year. Total number of beneficiaries includes initial, renewal, cap-subject, and cap-exempt petitions. This figure is constructed from publicly available data, which means that these may be underestimates in some cases; while some company names may appear multiple times, USCIS does not combine companies even where the names are the same because in all cases the employer tax identification numbers are different. In the confidential data used in the paper, this is not the case. The figure illustrates that there is also significant skewness in H-1B visas across firms. Note that these are not all multinationals and therefore are not equivalent to the BEA sample of firms used in this paper’s data; the sample of firms in my analysis is somewhat different.

Figure 8: Median Salary of H-1B Visa Holders at Top 20 H-1B Hiring Firms in 2018



Notes: This figure shows the median salary of H-1B visa holders hired by the top 20 H-1B hiring firms in 2018. The light grey bars show Indian outsourcing companies, while the black bars show American firms. The companies shown in black bars are closer to the sample of firms used in this paper.

Table 1: Final Receipt Dates of the Cap-Subject H-1B Petition Filing Period in Each Fiscal Year

Fiscal Year	Final Receipt Date	Days in Filing Period	Number of Lottery-Subject H-1B Petitions Received During the Filing Period
2004	February 17, 2004	323	
2005	October 1, 2004	184	
2006	August 10, 2005	132	
2007	May 26, 2006	56	
2008*	April 3, 2007	3	150,000
2009*	April 7, 2008	7	163,000
2010	December 21, 2009	265	
2011	January 26, 2011	301	
2012	November 22, 2011	236	
2013	June 11, 2012	72	
2014*	April 7, 2013	7	124,000
2015*	April 7, 2014	7	172,500
2016*	April 7, 2015	7	233,000
2017*	April 7, 2016	7	236,000
2018*	April 7, 2017	5	199,000
2019*	April 6, 2018	5	190,098

Notes: On the first business day of each April, USCIS begins accepting H-1B applications from firms seeking permits that will count towards the following fiscal year. Since the H-1B visa program operates on a first-come, first-serve basis, petitions are accepted until the cap hits, at which point no more petitions are processed. The end of the application period is demarcated by the “final receipt date”, which is the date on which they receive enough applications to fill the remaining available permits under the cap. Any cap-subject petitions submitted after the final receipt date were automatically rejected. This date is announced by USCIS in a press release every year, and it varies every year, as shown in the table above. In some years, USCIS received so many petitions within the first week that all cap-subject petitions were distributed by lottery. These years are demarcated with a star (*), and the fourth column shows the total number of petitions on which the lottery was run. For the purposes of this paper, the most relevant are fiscal years 2008 and 2009.

Table 2: Summary Statistics

Summary Statistics of MNCs in 2001

	N	Mean	Std. Dev	10 th percentile	90 th percentile
Reporter employment	2263	7785.2	29915.7		
Ratio of LCA Applications to US Employment in percentage point units	2263	0.176	0.812	0	0.2835
Number of countries active in	2263	2.831	5.646		
Number of LCA applications	2263	11.87	79.85		

Summary Statistics of MNCs in 2013

	N	Mean	Std. Dev
Reporter employment	2263	6300.6	35868.7
Number of countries active in	2263	2.422	5.992

Summary Statistics of Existing Foreign Affiliates in 2001

	N	Mean	Std. Dev
FA Employment	6407	1151.0	3964.3
FA R&D Expenditure	6407	2765.6	26276.7

Summary Statistics of Still-Existing Foreign Affiliates in 2013

	N	Mean	Std. Dev
FA Employment	5482	1559.4	8307.1
FA R&D Expenditure	5482	5741.0	37289.6

Table 3: Effect on Foreign Affiliate Employment for All Firms, 1994-2014

	(1)	(2)
	ln(FA emp)	ln(FA emp)
Continuous Treatment	0.0625*** (0.0183)	
Binary Treatment		0.274*** (0.0906)
Country FE	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	137459	57984
R^2	0.429	0.429
Adjusted R^2	0.420	0.413
F	11.68	9.140

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows the results of the baseline differences-in-differences approach. The timeframe is 1994-2014 and the dependent variable is ln(foreign affiliate employment). I control for country, firm, and year fixed effects in both columns. Both columns utilize an independent variable constructed as the interaction between a dummy variable equal to one in 2004-onwards and zero before, and a measure of H-1B dependency. In the continuous case, H-1B dependency is measured as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001. In the binary case, it is a new variable equal to one if the firm was in the top group of H-1B dependency in 2001, and zero if the firm had zero LCA applications in 2001. Column 1 shows the results of the continuous version, while column 2 shows the results of the binary version. Both are statistically and economically significant. The coefficient in the binary version tells us that the 2004 policy change caused highly H-1B dependent firms to increase their foreign affiliate employment by 27% more than a non-H-1B dependent firm. Note that these results are robust to standard errors clustered at the country level.

Table 4: Effect on Foreign Affiliate Employment for All Firms, 1994-2014, with Time-Varying Effects

	(1) ln(FA emp)
year=2004 # Continuous Treatment	0.0249 (0.0164)
year=2005 # Continuous Treatment	0.0432** (0.0216)
year=2006 # Continuous Treatment	0.0408* (0.0230)
year=2007 # Continuous Treatment	0.0477* (0.0245)
year=2008 # Continuous Treatment	0.0562*** (0.0171)
year=2009 # Continuous Treatment	0.0720*** (0.0208)
year=2010 # Continuous Treatment	0.0522*** (0.0192)
year=2011 # Continuous Treatment	0.0713*** (0.0179)
year=2012 # Continuous Treatment	0.0848*** (0.0172)
year=2013 # Continuous Treatment	0.103*** (0.0179)
year=2014 # Continuous Treatment	0.132*** (0.0222)
Country FE	Yes
Firm FE	Yes
Year FE	Yes
Observations	137459
R^2	0.429
Adjusted R^2	0.420
F	6.334

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows the results of the baseline differences-in-differences approach shown in Table 3, but also includes time-varying effects, where the treatment is interacted with the post-policy year. Continuous treatment is measured as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001, multiplied by a dummy variable equal to zero prior to 2004 and one after. As before, the

timeframe is 1994-2014 and the dependent variable is $\ln(\text{foreign affiliate employment})$, and I control for country, firm, and year fixed effects. The increasing coefficients over time suggest that the firm response to the cap change was not a one-off effect, and in fact provide evidence that firms accelerated offshoring over time as constraints became tighter.

Table 5: Effect on Foreign Affiliate R&D Employment for All Firms, 1994-2014

	(1)	(2)	(3)	(4)
	ln(FA emp)	ln(FA emp)	ln(FA R&D emp)	ln(FA R&D emp)
Continuous Treatment	0.0615*** (0.0176)		0.0614*** (0.0184)	
Binary Treatment		0.274*** (0.0906)		0.283* (0.155)
Country FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	137459	57984	6542	2293
R^2	0.429	0.429	0.408	0.410
Adjusted R^2	0.420	0.413	0.349	0.311
F	12.14	9.140	11.13	3.350

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The first two columns of this table are identical to the first two columns of Table 3, but are shown here for comparison purposes. The third and fourth columns are identical to the first two columns with one important difference: the dependent variable. Columns 3 and 4 measure the effect of the 2004 H-1B policy on not just aggregate foreign affiliate employment, but on foreign affiliate R&D employment specifically. An important caveat is that the R&D employment variable is only available in benchmark years, which in my sample includes 1994, 1999, 2004, 2009, and 2014.

Table 6: Effect on Growth in Foreign Affiliate Employment for All Firms, 2001 Base Year

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014
Ratio of LCAs to US-based Employment in 2001 in percentage point units	0.0113 (0.0103)	0.0084 (0.0120)	0.0140 (0.0148)	0.0316* (0.0178)	0.0397** (0.0186)	0.0498* (0.0262)	0.0451** (0.0201)	0.0443* (0.0235)	0.0330* (0.0195)	0.0578*** (0.0222)	0.0650*** (0.0252)	0.0722*** (0.0248)	0.0628 (0.0434)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5619	5277	4657	4418	4343	4122	3798	4160	3687	3449	3237	3130	3277
R^2	0.052	0.062	0.086	0.085	0.091	0.134	0.143	0.147	0.162	0.181	0.196	0.208	0.198
Adjusted R^2	0.018	0.026	0.046	0.043	0.050	0.092	0.100	0.107	0.119	0.138	0.151	0.162	0.155
F	1.198	0.492	0.895	3.164	4.527	3.606	5.040	3.554	2.852	6.754	6.682	8.481	2.096

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column represents a long difference in logged foreign affiliate employment between 2001 and a later year. We would expect null results in 2002-2003 (pre-policy change), and positive results in 2005-2014 (post-policy change), and this is exactly what we see. The key independent variable is my measure of H-1B dependency, defined as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001. The fixed effects account for industry (NAICS 4-digit level) and country time trends. The results show that firms that were one percentage point more H-1B dependent than average saw a 3-8% larger increase in foreign affiliate employment than average, as a result of increased immigration restrictions resulting from the 2004 cap drop. Note that the differences in observations in each year are driven by the fact that the dependent variable is logged; hence any observations equal to zero are dropped. Thus this regression only measures the effect on the intensive margin.

Table 7: Effect on Growth in Foreign Affiliate Employment for All Firms, 2001-2013

	(1)	(2)	(3)	(4)	(5)
	Baseline	Controlling for Pre-trend	Base year patent control	Most R&D- Intensive	Nonlinear Specification
Ratio of LCA Applications to US Employment in 2001 in percentage point units	0.0722*** (0.0248)	0.0684** (0.0283)	0.0893*** (0.0282)	0.126** (0.0551)	
Change in ln(FA emp), 1994-1999		-0.0818** (0.0326)			
2 nd Quantile of H-1B Dependency in 2001 = 1					0.0509 (0.108)
3 rd Quantile of H-1B Dependency in 2001 = 1					-0.00229 (0.118)
4 th Quantile of H-1B Dependency in 2001 = 1					-0.0518 (0.0990)
5 th Quantile of H-1B Dependency in 2001 = 1					-0.0792 (0.135)
6 th Quantile of H-1B Dependency in 2001 = 1					-0.129 (0.154)
Top Quantile of H-1B Dependency in 2001					0.462** (0.209)
Industry FE	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Patent Group FE			Yes		
Observations	3130	1913	3130	488	3130
R ²	0.208	0.247	0.211	0.261	0.209
Adjusted R ²	0.162	0.187	0.165	0.151	0.163
F	8.481	6.108	2.636	5.250	1.550

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable in every column is the long difference in logged foreign affiliate employment between 2001 and 2013. The key independent variable in columns 1-4 is my measure of H-1B dependency, defined as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001. The first column is the duplicate of Table 2, column 8, and serves as a reference baseline. The second column includes a control for the 1994-1999 pre-trend as a robustness check. Reassuringly, the coefficient of interest gets slightly smaller, but remains positive and statistically significant, further reassuring endogeneity concerns. Column three address the concern that results may be driven by systematic growth rate differences across firms of different size, internationalization, or innovativeness to begin with. If, for instance high-patenting firms naturally expand their activity abroad more quickly than non-patenting firms, even within the same industry, then my results could reflect that correlation rather than the effect of the policy change. I therefore tested whether the coefficient of interest changes when including controls for the size or type of firm in terms of their 2001 employment quantile, in terms of their 2001 sales quantile, in terms of their 2001 R&D spending quantile, in terms of their 2001 total patenting

quantile, and in terms of the number of foreign affiliates quantile in 2001. These controls are included as group fixed effects. The identification in these regressions, once various controls are added, is now based on the comparison of trajectories within the same industry and same category – based on employment size or sales size or R&D size or patenting amount or degree of internationalization, depending on the control. Column 3 results of a specification that contains indicators for each firm's 2001 patenting quantile; the other results can be requested separately but show the same result, which is that the main results remain qualitatively unaffected. Column 4 shows the results of the same baseline regression run on firms that were in the top 20% of R&D-intensive firms in 2001. I find that the coefficients are larger than on the entire firm sample suggesting that the results are driven by the firms conducting high-skilled activity, as expected. Finally, column 5 shows a non-linear approach, whereby I divide multinationals into groups according to their H-1B dependence in 2001. I build 7 categories of multinationals, where the base category is all multinationals with zero LCA applications, and the remainder are divided into five quantiles, with the top category divided into two groups. The table shows large, positive, and statistically significant coefficients for the top bracket (with a LCA application-US employment ratio of at least 0.0158 in 2001). In other brackets, there is no statistical significance. These estimates suggest that the positive effect of H-1B restrictions on foreign affiliate employment is being driven by the heaviest users of H-1B visas. Results are robust to the choice of any post-2004 year, but I only show 2013 results here due to space constraints and constraints due to the confidentiality of the data.

Table 8: Exploring Country-level Heterogeneity in Growth in Foreign Affiliate Employment for All Firms

	(1) Baseline	(2) Canada	(3) China & India	(4) Excluding CA CN & IN
Continuous Treatment	0.0615*** (0.0176)	0.105** (0.0424)	0.223*** (0.0319)	0.0422*** (0.0153)
Country FE	Yes	No	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	137459	12220	5704	119207
R^2	0.429	0.840	0.676	0.315
Adjusted R^2	0.420	0.823	0.643	0.304
F	12.14	6.168	49.14	7.640

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table explores heterogeneity in the location of expansion. The largest countries of origin for H-1B visa holders are China and India while many of the prominent examples of companies opening foreign affiliates abroad in response to H-1B restrictions are concentrated in Canada. Canada is a special case for US firms; cities like Vancouver and Toronto are geographically close to Silicon Valley and other multinational headquarters, but Canada has much less restrictive high-skilled immigration policies than the US. These characteristics mean that the fixed costs of offshoring are relatively low. To test whether these are indeed the locations that see the most significant expansion, I construct two samples: one of foreign affiliates in China and India, shown in column 2, and one of foreign affiliates in Canada, shown in column 3. All four columns show the same regression as Column 1 of Table 3, but with different country subsamples, where as before, the dependent variable is logged foreign affiliate employment and the independent variable of interest multiplies H-1B dependency in 2001 by a post-2004 indicator variable. All specifications include firm and year fixed effects, but columns 2 and 3 do not include country fixed effects since they are country subsamples. The effects are statistically significant, positive, and larger than the base sample in columns 2 and 3, suggesting that a lot of the expansion was concentrated in these three countries. However, a fourth regression – shown in Column 4 – that excludes Canada, India, and China shows that these three countries are not the only ones impacted. The coefficient is smaller when the three countries are excluded, but it is still economically and statistically significant. Hence, while expansion in response to immigration restrictions was concentrated in Canada, China, and India, they were not the only countries affected.

Table 9: Effect of H-1B Policy Restriction on the Likelihood of a US MNC Entering a New Country

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	2006	2007	2008	2009	2010	2011	2012	2013	2014
H-1B	0.00201**	0.00296**	0.00294**	0.00408**	0.00281**	0.00276*	0.00303**	0.00294**	0.00303**
Dependency in 2001	(0.000945)	(0.00123)	(0.00127)	(0.00162)	(0.00129)	(0.00142)	(0.00149)	(0.00149)	(0.00149)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	108623	108623	108623	108623	108623	108623	108623	108623	108623
R^2	0.014	0.023	0.024	0.029	0.026	0.028	0.031	0.032	0.041
Adjusted R^2	0.012	0.021	0.022	0.027	0.023	0.026	0.028	0.029	0.038
F	4.507	5.799	5.391	6.344	4.706	3.757	4.128	3.881	4.152

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable here is a binary variable equal to one if the firm had opened operations in a new country by a post-policy year (that they did not have in 2001) and equal to zero otherwise. I use a linear probability model so that I can include time trends, but the results are robust to a logit model. The independent variable is my measure of H-1B dependency, defined as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001. Results show that that the 2004 cap decline made a one percentage point more H-1B dependent firm 0.2-0.4% more likely to open a foreign affiliate in a new country than the average firm. The results illustrate a statistically significant and consistent response on the extensive margin – no matter the year – to the H-1B cap change in 2004. An important caveat to the results must be made here: as noted in the data section, each firm may report on a consolidated basis for multiple affiliates in the same country under certain conditions, so I aggregate all foreign affiliate activity up to the host country level for a given parent firm for a given year. This implies that effects on the extensive margin will only be observed when a firm enters a new country; the extensive margin results will not, for example, reflect a firm opening a foreign affiliate in Vancouver if it already has a foreign affiliate in Toronto. Hence, the results are likely smaller than the true effect.

Table 10: Effect on the Extensive Margin, Exploring Heterogeneity

	(1) DV: New FA by 2013	(2) DV: New R&D by 2013	(3) DV: New R&D by 2013 Subsample of 25 countries
Ratio of LCA Applications to US Employment in 2001 in percentage point units	0.00282* (0.00151)	0.00317* (0.00172)	
Binary form of H-1B Dependency in 2001			0.0648** (0.0259)
Industry FE	Yes	Yes	Yes
Country FE	Yes	Yes	Yes
Observations	108623	108623	8305
R^2	0.032	0.025	0.083
Adjusted R^2	0.029	0.022	0.060
F	3.467	3.411	6.281

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable here is a binary variable equal to one if the firm had opened new operations in a given country by 2013 (that they did not have in 2001) and equal to zero otherwise. I use a linear probability model so that I can include time trends, but the results are robust to a logit model. The independent variable in the first two columns is my measure of H-1B dependency, defined as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001. In the third column, the independent variable is the binary version of H-1B dependency in 2001, comparing the top quantile to the bottom quantile. Column 1 shows that that the 2004 cap decline made a one percentage point more H-1B dependent firm 0.2% more likely to open a foreign affiliate in a new country than the average firm by 2013. Column 2 shows the same regression specification, but with a look at the start of R&D. Here, the coefficient is even higher, suggesting that immigration restrictions caused US multinationals to both start conducting R&D in more existing foreign affiliates, but also to open more R&D-performing foreign affiliates. The effect is slightly larger on existing foreign affiliates starting to conduct R&D than on new R&D-performing foreign affiliates, but it is positive and statistically significant in both cases. Column 3 shows that the effect is even larger among highly H-1B dependent firms in a subset of 25 countries. The sample of countries in Columns 1 and 2 is fifty. Results are robust to the choice of any post-2004 year, but I only show 2013 results here due to space constraints and constraints due to the confidentiality of the data.

Table 11: Effect on Foreign Affiliate Employment Share of Firm Employment, 1994-2014

	(1) Emp Share, All firms & countries	(2) Emp Share, China, Canada, and India	(3) Emp Share, All firms & countries	(4) Emp Share, China, Canada, and India
Continuous Treatment	-0.0000789 (0.000270)	0.00487*** (0.00116)		
Binary Treatment			-0.00207 (0.00167)	0.0130** (0.00588)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Observations	296002	32899	134598	16773
R^2	0.380	0.709	0.420	0.752
Adjusted R^2	0.375	0.695	0.413	0.735
F	0.0855	17.52	1.545	4.906

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reflects a differences-in-differences approach that utilizes the panel dataset all at once, rather than a series of long-differenced cross-sectional regressions. The timeframe is 1994-2014 and the dependent variable is the share of firm i 's foreign affiliate employment in country c in time t . I control for country, firm, and year fixed effects in all columns. All columns utilize an independent variable constructed as the interaction between a dummy variable equal to one in 2004-onwards and zero before, and a measure of H-1B dependency. In the continuous case, H-1B dependency is measured as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001. In the binary case, it is a new variable equal to one if the firm was in the top group of H-1B dependency in 2001, and zero if the firm had zero LCA applications in 2001. Columns 1 and 2 show the results of the continuous version, while columns 3 and 4 show the results of the binary version. The sample in columns 1 and 3 is all firms in all years, while in columns 2 and 4, it is constrained to the three countries most impacted by foreign affiliate expansion in response to skilled immigration restrictions. On average (in the full sample), there is no clear evidence of any shift in the share of employment. However, when the sample is constrained to those countries where the expansion of activity as a direct response to skilled immigration restrictions was concentrated (Canada, India, and China), there was a clear shift of firm employment share towards those countries.

Table 12: Effect on Growth in Foreign Affiliate Employment for All Firms, 2005 Base Year

	(1)	(2)	(3)	(4)	(5)
	2010	2011	2012	2013	2014
Adjusted Excess Demand in 2007-8	0.0484 (0.0334)	0.137*** (0.0496)	0.121** (0.0499)	0.161** (0.0681)	0.142* (0.0836)
Number of LCAs in 2007-8	-0.0000989 (0.000113)	-0.000342** (0.000139)	-0.000205 (0.000157)	-0.000358* (0.000185)	-0.000295 (0.000216)
Industry FE	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Observations	4308	4029	3776	3685	3834
R^2	0.140	0.162	0.172	0.190	0.179
Adjusted R^2	0.101	0.121	0.131	0.149	0.138
F	1.093	3.891	3.192	2.840	1.493

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column represents a long difference in logged foreign affiliate employment between 2005 and a post-lottery year (2010-2014). The key independent variable is my measure of adjusted excess demand, defined as the total number of cap-subject LCA applications minus cap-subject H-1B petitions issued for a given multinational in 2007 and 2008, divided by that multinational's US employment in 2007. This number is multiplied by 100 for purposes of interpretation. The regression also includes a control for the number of LCA applications, since lottery randomization is at the petition, not firm, level. The fixed effects account for industry (NAICS 4-digit level) and country time trends. The results show that a random negative shock to H-1B supply equal to one percentage point of initial employment caused an increase in the foreign affiliate growth rate of between 12 and 16%.

Table 13: Estimating Offshored Jobs per unit of Excess Demand

Total Excess Demand Among US MNCs in April 2007-08	Predicted Change in Foreign Affiliate Employment, 2005-2014	Foreign Affiliate Jobs per Unit of Excess Demand
14,218	4,449	0.31

Notes: This table shows that about 0.3 jobs were offshored for every unfilled H-1B position. To calculate this number, I divided excess demand among US multinationals during the two relevant lottery years by the predicted change in foreign affiliate employment. To estimate the predicted change in foreign affiliate employment, I multiplied each firm's normalized excess demand by the coefficients identified in Table 9. I then multiplied this value by logged foreign affiliate employment in 2005, and then took the exponential to calculate the predicted change in foreign affiliate employment for each firm. I did this for each firm and then summed across cities. The table displays these figures for US multinationals in aggregate. Column 1 shows excess demand. Column 2 shows the predicted increase in foreign affiliate employment. The final column displays the number of offshored jobs lost per unfilled H-1B position. The 0.31 estimate of substitution should be treated with caution and as a lower bound for the true number of foreign affiliate jobs added per rejected H-1B visa for several reasons. First, it relies on calculations on the intensive margin and does not consider increased foreign affiliate employment on the extensive margin. Second, there are likely at least some firms that did not submit their petition(s) in time to be considered for the lottery, which means that their behavior is not captured in the analysis. Third, this estimate only captures how existing firms modified their strategy; new firms born into the visa-restrictive world may have incorporated offshoring into their strategy from the beginning. Finally, the estimate does not include foreign multinational company behavior, and one would expect that foreign MNCs would be even more likely to substitute foreign-born US-based skilled workers for foreign-born skilled workers at headquarters or other foreign affiliates. This suggests that, if anything, the regression estimates are smaller than the true effects – and so is the estimate of substitution.

Table 14: Effect on Growth in Foreign Affiliate Employment for All Firms, 2005 Base Year

	(1)	(2)	(3)	(4)	(5)
	2010	2011	2012	2013	2014
Adjusted Excess Demand in 2007	0.106* (0.0600)	0.205** (0.0871)	0.172* (0.0962)	0.203** (0.103)	0.119 (0.0921)
Number of LCAs in 2007	-0.000132 (0.000114)	-0.000279** (0.000129)	-0.000140 (0.000156)	-0.000241 (0.000157)	-0.000143 (0.000154)
Industry FE	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Observations	4308	4029	3776	3685	3834
R^2	0.140	0.161	0.172	0.189	0.178
Adjusted R^2	0.101	0.121	0.130	0.148	0.137
F	1.548	2.903	1.739	1.939	0.845

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column represents a long difference in logged foreign affiliate employment between 2005 and a post-lottery year (2010-2014). The key independent variable is my measure of adjusted excess demand, defined as the total number of cap-subject LCA applications minus cap-subject H-1B petitions issued for a given multinational in 2007, divided by that multinational's US employment in 2007. This number is multiplied by 100 for purposes of interpretation. The regression also includes a control for the number of LCA applications, since lottery randomization is at the petition, not firm, level. The fixed effects account for industry (NAICS 4-digit level) and country time trends. The results show that a random negative shock to H-1B supply equal to one percentage point of initial employment caused an increase in the foreign affiliate growth rate of between 10 and 20%. The table differs from Table 9 in that it only includes variation from the 2007 lottery, in an effort to address the concern that some firms anticipated the lottery.

Appendix A1. Other High-Skilled Visas

The H-1B visa is not the only visa option for multinational firms that wish to hire high-skilled foreign-born workers. If other high-skilled visas are easily substitutable for the H-1B visa, then H-1B visa restrictions will not actually constrain multinational firms' ability to hire skilled foreign-born workers. Hunt (2011) provides an extensive discussion of all different types of visa categories, but here I focus on those relevant to skilled workers. The TN visa was created under the North American Free Trade Agreement and allows Canadian and Mexican citizens to come to the United States only if they already have a job offer with a firm and are in a certain set of occupations typically requiring a bachelor's degree. The O visa is for workers of extraordinary ability. Figure A1-1 illustrates that the TN and O visas are used far less frequently than the H-1B visa.

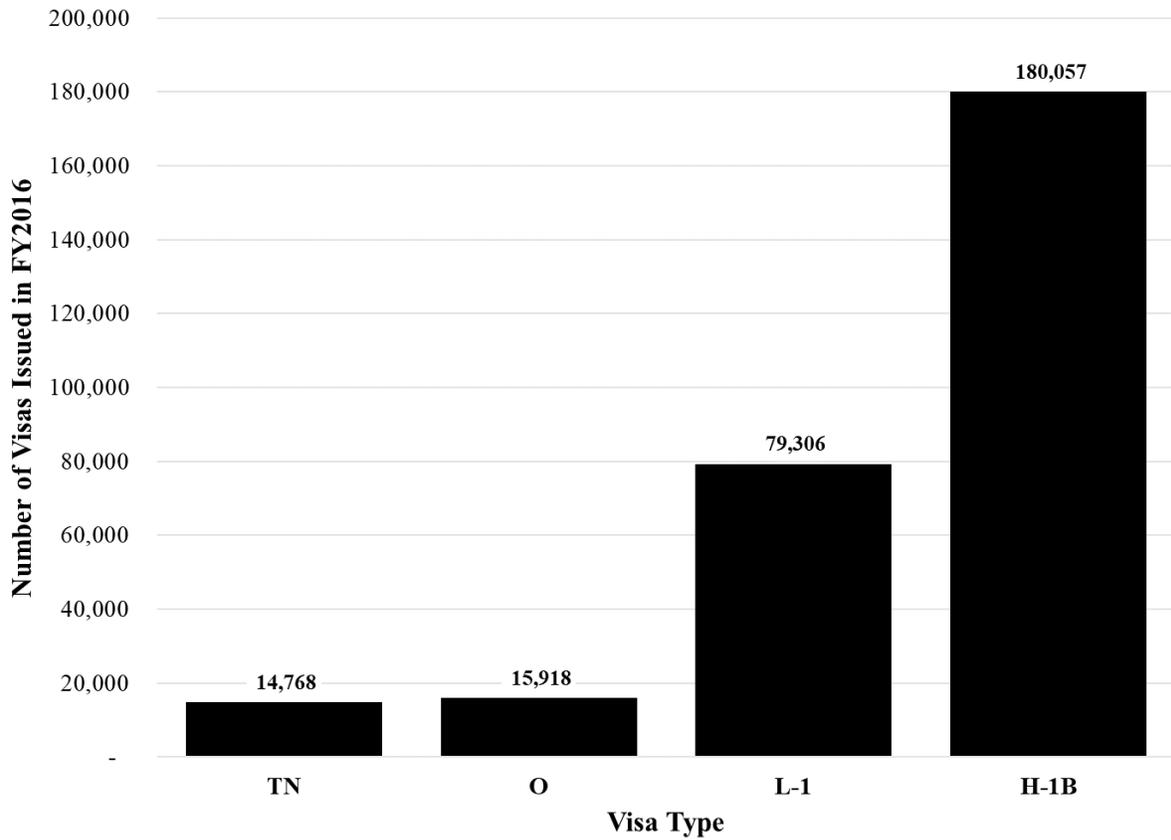
Of particular relevance for multinational firms is the L-1 visa. The L-1 visa is used to permit overseas branches or subsidiaries of MNCs to transfer foreign workers within the company. To be eligible, the employee must have worked abroad for one continuous year in the last three years in a related business entity in a managerial/executive position (L1A) or in a specialized knowledge staff capacity (L1B). The former has a duration of three years and is extendable up to seven years, while the latter has a duration of two years and is extendable by two years. Like the H-1B, the employer files for the L-1 visa with the USCIS on behalf of the worker using the I-129 petition. Unlike H-1B visas, L-1 visas are not capped and do not require employers to pay the prevailing wage; no labor certification is needed. The degree to which multinationals have substituted L-1 (or other) visas for H-1B visas as H-1B visas have become more difficult to attain is unknown – this is a question for future research.

However, the evidence suggests that the two visas are not close substitutes; Figure A1-2 demonstrates the aggregate number of L-1 visa petitions have flattened in recent years. Interviews with multinational firms also suggest that they are not easily substitutable and are used for different categories of people: *“we use both...but they're just different...so if we have someone...already with the firm, they're in a different category...than someone who may be just out of school or in an F1 that's transitioning to an H”*. Furthermore, while L-1 visas are not officially capped, L-1 denial rates have grown substantially over the last ten years, growing from about 5% in 2006 to nearly 70% in 2018. As another interviewee commented, *“you can't have certainty with either the H or the L”*.

There is one other program that skilled immigrants can use to work in the US: the OPT program. Unlike the other visas described above, the OPT program is not a visa but instead allows temporary employment under the F-1 student visa. Only foreign students on an F-1 visa with a higher education degree from a US college or university are eligible for the OPT program, and foreign students do not

require employer sponsorship to apply for the OPT. Furthermore, there is no cap on the number of approvals under the OPT program. However, the program is also shorter than the visas described above: the standard program allows up to 12 months of OPT employment. The program is longer for students in STEM fields. In 2008, an OPT extension was granted for up to twenty-nine months of work authorization for F-1 nonimmigrant students with STEM degrees, and in 2016, this was extended even further to thirty-six months. The OPT program is a popular pathway that foreign students on F-1 visas use to remain in the US for longer, in particular because it gives students more time to apply multiple times for H-1B visas. The STEM extension of the OPT program in 2008 likely helped loosen the constraints US multinational companies faced, and may have been another margin that firms adjusted on. The degree to which firms and students substituted the OPT program for H-1B visas – in addition to substituting between foreign workers at home and abroad – is unknown and is a question for future research, but current data sources on usage of the OPT program make this question difficult to answer.

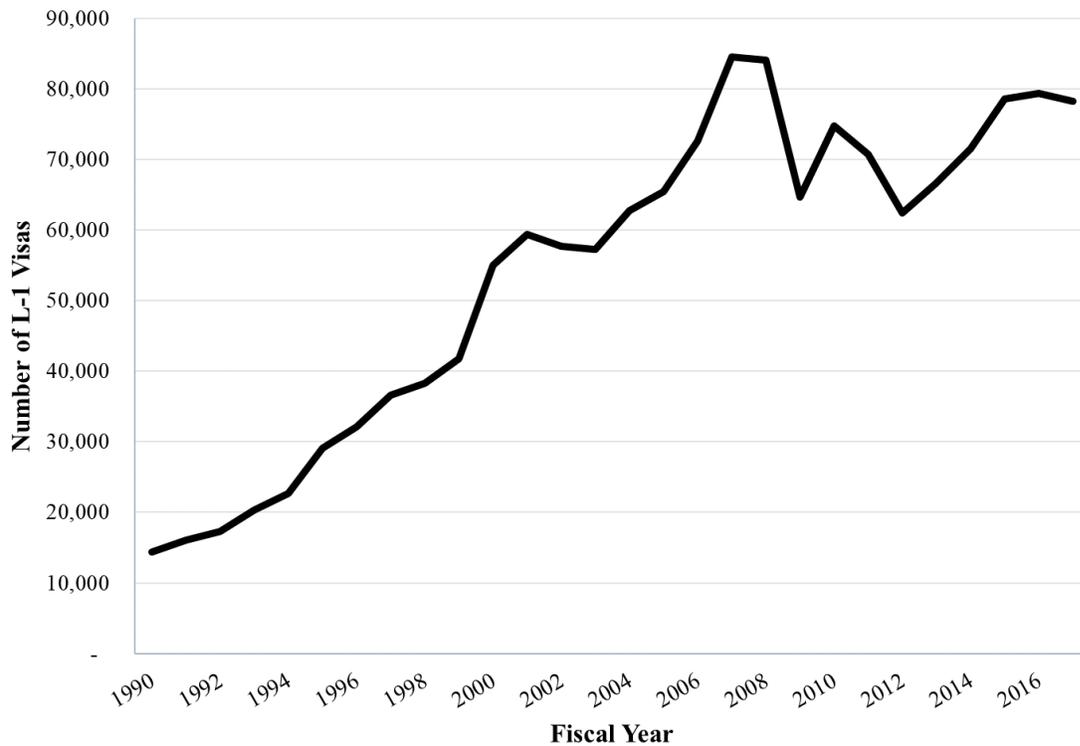
Figure A1-1: Number of High-Skilled Visas Issued in FY2016



Notes: The graph shows that H-1B are by far the most commonly used visa for high-skilled workers. The TN visa is for Canadian and Mexican citizens only under the North American Free Trade Agreement. The O visa is for “the individual who possesses extraordinary ability in the sciences, arts, education, business, or athletics, or who has a demonstrated record of extraordinary achievement in the motion picture or television industry and has been recognized nationally or internationally for those achievements.” The L-1 visa is for intracompany transfers of foreign workers in the managerial, executive, or specialized knowledge categories to the U.S. The data is sourced from the State Department’s Annual Reports on Nonimmigrant Visa Statistics, Table XVI: Nonimmigrant Visas Issued by Classification³⁶.

³⁶ <https://travel.state.gov/content/travel/en/legal/visa-law0/visa-statistics/nonimmigrant-visa-statistics.html>

Figure A1-2: Number of L-1 Visas Issued Over Time



Notes: The graph shows that the time trend of aggregate L-1 visa petitions from FY 1990 through FY 2017. It shows that the aggregate number of L-1 visa petitions have not increased since the H-1B policies started to become more restrictive, suggesting that the two visas are not close substitutes. The 1990-2003 data is from Kirkegaard (2005), while the 2004-2017 data is from the State Department’s Annual Reports on Nonimmigrant Visa Statistics, Table XVI: Nonimmigrant Visas Issued by Classification³⁷. The decline in the number of issued L-1 visas in recent years is consistent with the opinion of the immigration lawyer the author spoke to that there has been a significant informal crackdown on L-1 visa petitions in parallel to tightening H-1B restrictions.

³⁷ <https://travel.state.gov/content/travel/en/legal/visa-law0/visa-statistics/nonimmigrant-visa-statistics.html>

Appendix A2. More Detail on the Matching Procedure

There are no numerical identifiers that exist across all datasets, so I matched each database into the BEA data using firm names. For each match, I conducted several rounds of fuzzy matching between BEA multinationals and the firms in each of the other datasets (patent assignees/H-1B petitioners) using the “reclink2” Stata command at a 0.8 level, followed by manual verification to ensure the generated matches were correct. If a firm appeared in the BEA data but not in the patent data in a given year, I assumed that it did not apply for any patents in that year. Similarly, if a firm appeared in the BEA data but not in the I129 or LCA data in a given year, I assumed the firm did not apply for any visas in that year. Finally, I restrict the sample to the fifty countries with the most foreign affiliate activity to reduce noise, and to the multinationals who existed by 2001. Without this second restriction, multinationals who had zero H-1B applications in 2001 – because they did not exist yet – would be falsely classified as non-H-1B dependent.

Appendix A3. Quotes from Conversations with US Multinationals and an Immigration Lawyer

“The site [in India] ...was set up in response to...growing restrictions on H-1B visas. HQ decided that if we could not bring the talent to the US, then [we] would go where the talent was.”

“this project we want to be staffed by these five people...one is in France and one is in the Czech Republic and one is in India, and we can't bring them here on visas with the H-1B or L-1B, so we are going to house our own internal project in Prague instead of Chicago, and we'll send [the team] to Prague instead of to Chicago.”

“we're not going to be the destination location if we can't figure out ways to move people in and out”

“...an Irish company had a particular proprietary software that was useful to municipalities...and they started getting a lot of contracts in the United States for various local governments...they were recruiting at US campuses and not all of the people they were hiring were US citizens and they couldn't use the H-1B successfully...[because of the cap] ...so they did set up operations in Canada specifically so they could do [the work] close at hand”

Appendix A4. Estimates on US-based Employment

In the main body of this paper, I find that restrictions on high-skilled immigration caused US multinational firms to hire additional foreign workers abroad at their foreign affiliates. I therefore argue that US multinational firms substitute between hiring foreign skilled workers in the US and hiring foreign workers abroad at their foreign affiliates.

To further support this argument, one ideally would be able to observe a corresponding decline in the hiring of foreign skilled workers in the US. However, the data do not allow me to isolate out foreign US-based skilled workers; I only observe total US-based workers. Since I cannot separate out foreign and native employees of US parents, the effect is hard to isolate, as displayed in Table A4-1 below, which shows the estimates resulting from regressing the effect of skilled immigration restrictions on total US-based employment:

$$\Delta \ln(US_emp_{i,t-2001}) = \alpha_j + \beta_1 \left(\frac{LCAapps}{US_emp} \right)_i^{2001} + \Delta \epsilon_{i,t-2001}$$

where i indexes the firm, j indexes the industry, and t is a post-policy year. There are two main differences between this regression and the earlier baseline regression: the dependent variable is now US_emp , or US-based employment (native and foreign) by firm i , and the regression is now at the firm-year level rather than at the firm-country-year level. One would expect β_1 to be negative, but perhaps not statistically significant, if there were a decline in foreign US-based skilled workers; foreign US-based skilled workers make up a small fraction of total US-based employment by US multinational companies. Table A4-1 shows exactly that: a negative, but very noisy, effect on US-based employment, likely due to the difficulty of isolating the effect on foreign workers employed in the US.

Table A4-1: Effect on Growth in US-based Employment for All Firms, 2001 Base Year

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013
Ratio of LCAs to US-based Employment in 2001 in percentage point units	0.00436 (0.00880)	0.00594 (0.0106)	-0.0275** (0.0127)	-0.0157 (0.0159)	-0.00276 (0.0200)	-0.0192 (0.0221)	-0.0154 (0.0254)	-0.0168 (0.0289)	0.00122 (0.0319)	-0.0212 (0.0296)	-0.0199 (0.0332)	-0.00736 (0.0366)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1198	1120	960	944	921	866	832	893	804	749	702	669
R^2	0.138	0.145	0.195	0.160	0.154	0.217	0.224	0.254	0.234	0.219	0.230	0.228
Adjusted R^2	0.026	0.025	0.065	0.025	0.018	0.079	0.086	0.127	0.095	0.075	0.076	0.072
F	0.246	0.316	4.715	0.979	0.0191	0.756	0.369	0.337	0.00146	0.514	0.359	0.0404

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column represents a long difference in logged US-based employment between 2001 and a later year. We would expect null results in 2002-2003 (pre-policy change), and negative (but potentially noisy) results in 2005-2014 (post-policy change), and this is exactly what we see. The key independent variable is my measure of H-1B dependency, defined as the total number of LCA applications for a given multinational divided by that multinational's US employment in 2001. The fixed effects account for industry (NAICS 4-digit level) time trends. The results are unable to detect a statistically significant change in US-based employment in response to the policy change, but this is likely due to the data's inability to parse out foreign US-based employment from native US-based employment in the US.

Appendix A5. Additional Robustness Checks

Table A5-1: Effect on Foreign Affiliate Employment for All Firms, 1994-2014, where H-1B dependency is measured using H-1B petitions to USCIS in 2001 rather than LCA applications to DOL in 2001

	(1) ln(FA emp)	(2) ln(FA emp)
Continuous Treatment using H-1B petitions	0.0620* (0.0335)	
year=2004 # Continuous Treatment using H-1B petitions		0.00240 (0.0304)
year=2005 # Continuous Treatment using H-1B petitions		0.0322 (0.0289)
year=2006 # Continuous Treatment using H-1B petitions		0.0283 (0.0357)
year=2007 # Continuous Treatment using H-1B petitions		0.0508* (0.0273)
year=2008 # Continuous Treatment using H-1B petitions		0.0944*** (0.0295)
year=2009 # Continuous Treatment using H-1B petitions		0.0876** (0.0406)
year=2010 # Continuous Treatment using H-1B petitions		0.0382 (0.0410)
year=2011 # Continuous Treatment using H-1B petitions		0.0636 (0.0534)
year=2012 # Continuous Treatment using H-1B petitions		0.102** (0.0482)
year=2013 # Continuous Treatment using H-1B petitions		0.137*** (0.0479)
year=2014 # Continuous Treatment using H-1B petitions		0.158*** (0.0506)
Country FE	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	137459	137459
R^2	0.429	0.429
Adjusted R^2	0.420	0.420
F	3.425	3.629

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows the results of the baseline differences-in-differences approach shown in Table 3 and 4, but measures continuous treatment differently. As before, a measure of H-1B dependency is multiplied by a dummy variable equal to zero prior to 2004 and one after. But in this table, H-1B dependency is measured using H-1B petitions to USCIS rather than LCAs to DOL. As before, the timeframe is 1994-2014, the dependent variable is $\ln(\text{foreign affiliate employment})$, and I control for country, firm, and year fixed effects.

Table A5-2: Effect on Foreign Affiliate Employment for All Firms, 1994-2014, DV is inverse hyperbolic sine of foreign affiliate employment

	(1) IHS(FA emp)	(2) IHS(FA emp)
Continuous Treatment	0.110** (0.0427)	
year=2004 # Continuous Treatment		0.0817* (0.0460)
year=2005 # Continuous Treatment		0.0889*** (0.0301)
year=2006 # Continuous Treatment		0.0853*** (0.0311)
year=2007 # Continuous Treatment		0.142*** (0.0494)
year=2008 # Continuous Treatment		0.0720* (0.0427)
year=2009 # Continuous Treatment		0.167** (0.0666)
year=2010 # Continuous Treatment		0.0987* (0.0569)
year=2011 # Continuous Treatment		0.155** (0.0629)
year=2012 # Continuous Treatment		0.152** (0.0693)
year=2013 # Continuous Treatment		0.177** (0.0704)
year=2014 # Continuous Treatment		-0.00922 (0.113)
Country FE	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	409794	409794
R^2	0.285	0.285
Adjusted R^2	0.281	0.281
F	6.632	1.491

Standard errors in parentheses

Standard errors clustered at the firm level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows the results of the baseline differences-in-differences approach shown in Table 3 and 4, but shows both the extensive and intensive margin at once by using the inverse hyperbolic sine of foreign affiliate employment as the dependent variable rather than the natural log. The advantage of the inverse hyperbolic sine is that it provides a similar transformation as natural log but does not drop observations equal to zero. As before, the timeframe is 1994-2014, the independent variable is H-1B dependency multiplied by a dummy that equals one after 2004 and zero prior to then, and I control for country, firm, and year fixed effects.