

Skilled Immigration and Firm-Level Innovation: The U.S. H-1B Lottery

Andy Wu*

Strategy Unit
Harvard Business School

February 20, 2018

Abstract: The growth of the global technology industry drives the migration of skilled labor towards countries like the United States that can utilize it, but the U.S. limits the immigration of skilled workers that are employed domestically by U.S. firms. Proponents argue that skilled immigration allows firms to access technical skills that unavailable domestically and promote innovation, but there is little evidence of whether this firm-level effect exists. We evaluate the impact of skilled immigration on innovation in U.S. firms by exploiting a random lottery in the H-1B visa program, which allows us to estimate precise causal effects of marginal skilled immigration on firm-level innovation as measured by patenting productivity. Our set of empirical models compare firms that applied for the same number of lottery-subject immigrants but won different numbers of immigrants because of the lottery. Our results suggest that winning an H-1B immigrant does not significantly increase patent applications or grants at the firm level. Consistent with our main empirical findings, further analysis of the composition of participating immigrants and employers suggests that the current utilization of the H-1B program is not conducive to generating firm-level innovation; we find pervasive use of the program in industries, occupations, and firms where patenting is low.

Keywords: Immigration, Migration, Innovation, H-1B Visa, Intellectual Property

*We are grateful to Vikas Aggarwal, Santosh Anagol, Mark Duggan, Todd Gormley, Andrew Johnston, David Hsu, Laura Huang, Joseph Harrington, Harbir Singh, and Olivia Mitchell, and seminar participants at the Strategic Management Society and the Academy of Management. Priyanka Aurora, Rachel Hong, Cliff Kao, Sarah Kho, Rajan Sheth, and especially Charles Du and Arthur Halim provided excellent research assistance. The authors acknowledge financial support from Harvard Business School, the Mack Institute for Innovation Management at the Wharton School, and the Kauffman Foundation. Views expressed are those of the author only. Address: 15 Harvard Way, Morgan Hall 242, Boston, MA 02163 USA. Telephone: (781) 266-8088. Email: awu@hbs.edu.

1 Introduction

The growth of the global technology industry drives labor migration pressures towards developed countries and regions with firms that can efficiently utilize skilled labor for knowledge production and technological development (Bresnahan and Gambardella, 2004). As a result, skilled labor is more attracted to some markets—such as Silicon Valley in the United States—where demand for this labor is high, and away from other markets—such as India—where the supply of skilled labor is plentiful and cannot be fully utilized. The resulting effect is a net imbalance in global mobility trends (Alnuaimi et al., 2012). In response, organizations now actively compete in the international market for skilled labor (Lepori et al., 2015).

This trend is often viewed as a negative from the perspective of the countries from which the labor is emigrating away. This process, which has become known as “brain drain” (Bhagwati and Hamada, 1974; Docquier and Rapoport, 2012), refers to the perceived loss of human capital from a home country when its skilled human capital leaves to pursue potential greater opportunity elsewhere (Davenport, 2004; Janger and Nowotny, 2016).¹ For example, top Indian students are more likely to emigrate away than worse students, and the ones leaving appear to be the best and from the educated labor pool (Qin, 2015).² Indeed, highly skilled emigrants can generate more novel output and impact because they have access to more resources in their host country than their home country (Gibson and McKenzie, 2014).

Since this migration is viewed as a negative for the departing country, we might expect something of a beneficial “brain inundation” on the country receiving skilled immigrants. In addition to the additional human capital offered by the immigrants, we might expect the

¹These immigrants do occasionally return to their home country, and these returnees bring back knowledge which contributes to innovation in the home country (Filatotchev et al., 2011; Baruffaldi and Landoni, 2012; Kahn and MacGarvie, 2016).

²There is negative selection on the return migration of Indian immigrants back to India along measures of ability, post-migration education, and income (Qin, 2015).

immigrants to bring a diverse set of external knowledge to their hiring organizations, where this external knowledge is an increasingly important source of innovation for organizations seeking sustained competitive advantage (Filatotchev et al., 2011). Accordingly, technology leaders in the United States have argued that the skilled immigrants increase domestic technological innovation and productivity, both through established and entrepreneurial organizations³: Laszlo Bock, the former Senior Vice President of People Operations at Google argues that “talented foreign-born individuals have played and will continue to play a vital role at Google and throughout our economy” and that restrictions on “our immigration policies are stifling innovation” (Bock, 2013).

In the arena of U.S. public opinion and policy, the argued benefit is not so obvious. While visa limitations on immigration pursuing academic study⁴ and intra-organizational transfers⁵ are relatively loose in the United States, the hiring of skilled immigrants to fill open positions at private sector firms—predominately through the H-1B program—continues to be limited by the government over the last decade. This constraint binds more and more tightly as U.S. firm demand for skilled labor and foreign demand for US jobs increases while the limit remains constant. In 2014, more than half of H-1B applications were denied to satisfy the statutory caps, representing the rejection of a substantial number of presumably mutually beneficial transactions between foreign workers and domestic employers. Demand

³Immigrant entrepreneurs have a visible and documented role in the growth of US technological clusters. (Saxenian, 1996). Some 42% of Fortune 500 firms were started by immigrants or their children; these 211 firms produce over \$5 trillion in revenue annually (Ballmer et al., 2011).

⁴The movement and sharing of scientists across academic research institutions exhibits knowledge transfer benefits in scholarly work by facilitating “brain circulation” (Edler et al., 2011) and collaboration opportunities (Jonkers and Cruz-Castro, 2013; Scellato et al., 2015). The United States offers the F visa program for academic training. The J-1 visa is issued to research scholars (e.g. post-doctoral fellows), professors, and culture exchange visitors. The H-1B program is used for in limited cases to support immigrants pursuing academic study, but it is uncapped for non-profits such as universities.

⁵International intracompany transfers can transmit knowledge from one office to another and help coordinate activities across offices, including between the headquarters and the lower offices (Criscuolo, 2005; Storz et al., 2015). In the United States, there is also a different category of visa, the L visa program, for intracompany transfers of employees who worked for the firm for over a year within the three preceding years. This program is does not have as strict limits as the H-1B program for new hires.

for highly skilled immigrants has steadily risen, possibly reflecting a dwindling technical capacity of U.S. domestic graduates ([Bound et al., 2009](#)). In contrast, many developed countries explicitly recruit highly skilled immigrants without placing limits on the number of skilled immigrants that may come, while the United States allows only a limited number, primarily through its H-1B program.⁶

Opponents of skilled immigration expansion, particularly from organized labor, argue that skilled immigration, particularly through the H-1B program, crowds out domestic employees and suppresses domestic wages. Richard Trumka, President of the AFL-CIO labor union, writes that “High-tech companies say there are ‘too few’ American high-tech workers, but that’s not true. . . 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” ([Trumka, 2013](#)). If we cannot find that the skilled immigrants are productively superior to domestic labor, such as in generating technological innovation, it becomes harder to justify the H-1B program, particularly if foreign labor is just a direct substitute for domestic labor and thus depressing domestic wages (e.g. [Samuelson, 1964](#)).

The prior literature on skilled immigration of new workers and technological innovation (as measured by patent counts) has generally found a positive relationship between the them ([Kerr and Lincoln, 2010](#); [Hunt, 2011](#)). However, the prior literature is dominated by empirical work using data at the regional level, instead of at the level of individual firms and individual immigrants. We argue that research on immigration must delve into the details of specific visa programs because their unique regulations have distinct impacts on the composition of immigrant flows. From a policy perspective, the practical reality for the United States is that adjustments are considered at the level of each specific visa program, either through presidential executive order or congressional bill, and rarely at the level of full

⁶Australia, Canada, Japan, South Korea, the United Kingdom, France, the Netherlands and Germany place no cap on their high-skills immigration ([Ochel, 2000](#)).

immigration reform across programs; indeed, comprehensive immigration reform has failed several times just in the last 10 years.⁷ Moreover, a direct adjustment to the already binding numerical limitation of visas offered through the H-1B visa program would have immediate action in increasing the number of skilled immigrants in the coming year.

To offer a precise causal estimate of the main effect of interest that has direct implications for domestic innovation policy, we focus exclusively on the firm-level effects of additional marginal skilled immigrants on domestic firm-level innovation by examining a single program, namely the H-1B worker visa program. We exploit a random lottery in the H-1B visa program that allows us to obtain causal estimates of the impact of immigrants in this program. We can compare firms that applied for the same number of lottery-subject applicants but won different numbers of immigrants because of the lottery, using patent output as a measure of innovation across these firms to identify the impact of these immigrants.

Across a robust set of empirical models, we find that winning an H-1B immigrant does not significantly increase patent applications or grants at the firm level. Our results suggest the existing literature, which shows a positive correlation between the spatial distributions of skilled immigrants and patents, are not driven by a direct firm-level effect. We then argue that our results are consistent with the current utilization of the H-1B program in circumstances not conducive to generating firm-level patenting gains; we find pervasive use of the program in industries, occupations, and firms where patenting is low.

This paper contributes to a small but growing literature on the impact of skilled immigration on innovation and firm productivity. Our paper is the first to use patent application records in conjunction with H-1B administrative data to estimate the causal impact of an H-1B immigrant on firm-level innovative productivity. We also document the utilization patterns of firms and immigrants participating in the H-1B program.

⁷Major failed immigration reform bills in the United States include the *Secure Borders, Economic Opportunity and Immigration Reform Act* (S. 1348) of 2007 and the *Border Security, Economic Opportunity, and Immigration Modernization Act* (S. 744) of 2013, among others.

Our paper proceeds as follows. In [Section 2](#), we describe the institutional details of the H-1B visa program and the unique circumstances around the H-1B lottery that enable our empirical methodology. In [Section 3](#), we elaborate more on the lottery that forms the basis of our empirical design and explain the associated econometrics. In [Section 4](#), we detail our data collection and construction. In [Section 5](#), we present our main results. In [Section 6](#), we provided detailed data on the composition and utilization of the H-1B program that are consistent with our main empirical findings. [Section 7](#) concludes.

2 H-1B Visa Program

2.1 Background

Firms apply for H-1B visas on behalf of skilled foreign workers they would like to employ domestically. The H-1B visa program—administered by the United States Citizenship and Immigration Services (USCIS)—enables U.S. employers to seek temporary foreign workers in a *specialty occupation*, an occupation that “requires theoretical and practical application of a body of highly specialized knowledge in fields of human endeavor.”⁸ The visa application fees are paid for by the employer, and the fee and accompanying legal services generally cost several thousand dollars per visa petition filed. The visa lasts for three years but can be extended to six years by the employer. It is officially classified as a non-immigrant visa—i.e., for those not seeking long-term permanent residency—but it works as a “dual intent” visa, meaning it enables its holder to seek lawful permanent resident status through a green card. Administratively and throughout this paper, the H-1B sponsoring firm is referred to as the *petitioner*, while the foreign worker is known as the *beneficiary*. A visa application is referred to as a *petition*. The visa ties workers’ legal status to their continued employment at the

⁸As stated in the U.S. Code of Federal Regulations “Special requirements for admission, extension, and maintenance of status” in 8 CFR§214.2(h).

firm; if the worker quits or is fired, the worker must secure another visa or may be required to leave the country. The H-1B program contains a number of measures designed to protect the employment and wages of domestic workers, codified in the Labor Condition Application (LCA), which is self-attested by the petitioners.⁹

2.2 The Lottery

Our empirical design exploits an idiosyncratic property of the H-1B program that led to randomization of visa issuance when the number of petitions filed is greater than the legal limit of visas that can be issued. The Immigration Act of 1990 established a cap on the number of new H-1B visas that can be issued each year. Congress sets the cap for each year, which normally exists at 65,000 visas; there is a separate cap, known as the “advanced-degree cap exemption,” for 20,000 immigrants with a master’s degree or higher. In this paper, we focus on only the 65,000 non-advanced degree visas that are subject to the regular cap and which are the part of the program subject to the most policy debate. At the beginning of the 1990s, the number of available visas exceeded the number of petitions, but the number of petitions rose until the mid-1990s when the cap became binding. The cap was raised in 1998 and 2000, reaching a high of 195,000 visas and becoming non-binding again in many years. However, this high cap was not renewed, and by 2004, the cap returned to the original 65,000 and became binding again for non-advanced degree holders. These 65,000 visas are referred to as being “cap-subject,” because they are limited by the cap.

The lottery is a direct consequence of ambiguity in the order of priority of visas to be

⁹The Labor Condition Application (LCA) requires that employers attest to the following conditions. First, the immigrant’s wage must meet or exceed the prevailing wage for the majority of employees in their area of employment. Second, the hiring of the immigrant must not adversely affect working conditions of workers similarly employed. Third, the immigrant cannot be employed in an occupation and place of employment where there is currently a strike, lockout or work stoppage. The American Recovery and Reinvestment Act of 2009 added a number of other restrictions, including that employers must take good-faith steps to recruit U.S. workers for the open position and that they must not have laid-off and will not lay-off any U.S. worker in an equivalent job.

issued when the program limit is exceeded. In theory, the program operates on a first-come, first-serve basis, so the petitions that arrive earlier in the mail take priority. USCIS begins accepting petitions for H-1B visas for the next year on the first business day of April in the current year. When firms (and their immigration attorneys) anticipate that the limit in a given year will be reached, firms apply earlier in the cycle. As demand for the visas increase, firms apply earlier and earlier, eventually applying as early as the earliest day possible, the first business day in April. In 2005, the available cap-subject visas ran out in 132 days. By 2006, the visas ran out in 56 days, and in 2007 and 2008, the entire supply of H-1B ran out within days of the application cycle opening. In 2006, 2007, and 2008, they received far more petitions than they had available visas in final days before the program limit was reached. Because USCIS cannot distinguish which petitions had arrived in the mail earlier if they all arrive on the same days, they subjected H-1B petitions received on the days the program limit was reached¹⁰ to a “computer generated random selection process”¹¹ to determine which petitions were approved and which were not. We refer to the random selection process as the *lottery*.

In 2006, the quota was exceeded on the final day, which was 56 days into the application cycle, and a lottery was conducted on petitions received that day. In 2007, the quota was exceeded, almost twice over, in the first two days of the application cycle: by the second day, 123,480 applications for H-1B visas were submitted to USCIS, and a lottery was used to randomly allocate 65,000 visas among the many petitions. In 2008, the quota was exceeded by April 7, and winners were randomly drawn from the application pool of those who filed for H-1B visas from April 1 to April 7. We exploit the use of these lotteries for our empirical design.

¹⁰This day is referred to as a “final receipt date”.

¹¹This terminology is explicitly used in official USCIS press releases.

3 Empirical Design

The empirical design leverages the H-1B lottery to precisely estimate the causal effect of skilled immigration on patenting as a measure of innovation. Patent output has widely been used to measure of the innovation output and productivity of immigrants (Chellaraj et al., 2008; Hunt and Gauthier-Loiselle, 2010; Kerr and Lincoln, 2010; Hunt, 2011).¹²

3.1 Petition Bin Specification

We base our main regression model on Black et al. (2003), who leverage random assignment in the unemployment insurance system, and Angrist et al. (2012), who study charter school lotteries in Boston. To isolate random variation from the lottery, Angrist et al. (2012) condition on “risk set” indicators designating the combination of schools each student applied to. Conditional on the risk set, winning in the lottery is random.

In the simplest analysis, we could compare the patent outcomes of firms who petitioned for the same number of H-1B immigrants but won different number of H-1B visas because of the lottery. This approach has limited statistical power due to the small number of firms that applied for the exact same number of immigrants. Instead we pool these regressions using a fixed effect for each observed number of applications for a firm, thus making comparisons only among firms that applied for the same number of visas.

Firms petition for one or multiple H-1B visas for foreign workers. The number of H-1B visas a firm wins in the lottery is not unconditionally random because they apply for different

¹²Hunt and Gauthier-Loiselle (2010) find a 1 percentage point increase in immigrant college graduates’ population share increases patents per capita by 9-18 percent. Hunt (2011) finds that immigrants entering on a student/trainee or temporary work visa have higher patent productivity, publishing productivity and company founding rates than American-born workers. Kerr and Lincoln (2010) find that higher aggregate H-1B admissions increase immigrant employment and patenting in their study of city-year variation in H-1B dependence. Chellaraj et al. (2008) study variation in enrollment in U.S. graduate programs by international versus domestic students and find that a 10% increase in foreign graduate students leads to 6.8% more patent grants to universities and 5.0% more to non-universities. The prior literature is dominated by empirical work using data at the regional level, instead of at the level of individual firms. This type of regional analysis may be susceptible to various confounding factors, such as macroeconomic shocks.

numbers of visas. However, the number of H-1Bs a firm wins is random conditional on how many applications the firm submits. Thus, we specify a regression model to compare the mean patenting of firms winning more lotteries within a group of firms who petitioned for the same number of H-1B visas. In all our analysis, we study three separate lottery events which occurred in 2006, 2007, and 2008.

We refer to our baseline model as the ‘‘Petition Bin’’ model. For firm s , year t , petition count j , lottery year Y , and patent lag k , we estimate:

$$y_{s,t+k} = \alpha + \beta\theta_{st}^Y + \sum_j \delta_j d_{sj}^Y + \epsilon_{st}$$

$y_{s,t+k}$ represents the number of patents a firm applies for or is granted for the k years after the lottery, θ_{st}^Y represents the number of H-1B immigrants an employer won in year Y , and d_{sj}^Y represents a dummy that equals 1 for firms with petition count j and 0 otherwise for firms in the Y lottery. β captures the average effect of one additional H-1B immigrant on a firm’s patenting behavior in one year. This model represents the most parsimonious use of the risk set, which pools data across firms with different numbers of petitions to increase statistical power. The Petition Bin model is estimated using ordinary least squares (OLS).

Conditional on the number of petitions a firm submits, θ_{st}^Y approximates a binomial random variable. The proportion of available visas relative to the total number of petitions defines the probability p that a particular petition will be successful and be awarded a visa. The number of petitions a particular firm files is denoted as n . A firm with full information and competent paperwork ability will expect np successes with variance $np(1 - p)$.

3.2 Differences-in-Differences Specification

To check robustness, we compare the results of this petition bin specification to those of a generalized differences-in-differences with firm fixed effects to control for pre-existing

differences in firms between lottery winners and losers and improve efficiency by absorbing more residual variation. If the lottery is not truly random—if winning visas in the lottery is somehow correlated with the dependent variable of patent productivity—the generalized differences-in-differences model controls for a baseline patent rate and addresses that issue. This model controls for underlying time-invariant firm quality. For firm s , year t , and lottery Y , we estimate:

$$y_{st} = \alpha + \beta\omega_{st}^Y + \zeta\pi_{st}^Y + \gamma_s + \lambda_t + \epsilon_{st}$$

We refer to this as the “Diff-in-Diff with Firm FE” model. The variable y_{st} represents the patent count of firm s at time t . The variable ω_{st}^Y represents the number of lottery wins the firm won in Y if $t \geq Y$; for all $t < Y$, ω_{st}^Y is equal to 0. π_{st}^Y represents the number of H-1B petitions made by the firm in Y if $t \geq Y$. For all $t < Y$, π_{st}^Y is equal to zero. The variables γ_s and λ_t represent firm and year fixed effects respectively. The firm fixed effects still function as risk sets, but they are a more generalized control than the risk sets based upon petition bins.

Finally, we implement another model that replaces the firm effects with petition bins (which would otherwise be collinear with the firm fixed effects). For firm s , year t , petition count j , and lottery Y , we estimate:

$$y_{st} = \alpha + \beta\omega_{st}^Y + \sum_j \delta d_{sj}^Y + \lambda_t + \epsilon_{st}$$

We refer to this as the “Diff-in-Diff with Petition Bins” model. d_{sj}^Y represents a dummy that equals 1 for firms with petition count j and 0 otherwise. As before, the variable y_{st} represents the count of patents for firm s at time t . The variable ω_{st}^Y represents the number of lottery wins the firm won in Y if $t \geq Y$; for all $t < Y$, ω_{st}^Y is equal to 0. This model is less general than the firm FE model—and it may underperform if the lottery is not truly random—but it would control for non-randomness of lottery if the non-randomness is predicted by the

number of petitions made, which can be thought of as a proxy for firm size.

We estimate our differences-in-differences models with ordinary least squares (OLS), OLS with a logged dependent variable (e.g. $\ln(\text{Patent Count} + 1)$), and negative binomial (NBR) with conditional firm fixed effects. The OLS model is the most transparent and best linear unbiased estimator, but inference on the parameters requires a normality assumption, and the patent count dependent variables are clearly non-negative and heavily skewed with long right tail: this non-normal distribution causes inefficiency in the basic OLS model. Logging the dependent variable of patent count is a rough improvement to the performance of OLS as it makes the dependent variable more normal.¹³ Finally, the negative binomial distribution allows for count data with different means and variances, and it is commonly used to analyze patent data (Hausman et al., 1984; Allison and Waterman, 2002). We provide robust standard errors for all our regression models.

In parallel to our study, Doran et al. (2015) also implement the H-1B lottery for identification of firm-level impact of skilled immigration, but they use a different econometric strategy involving instrumental variables and focus on patent grants and the lottery years of 2006 and 2007.¹⁴ Similar to our findings, they report an insignificant correlation between H-1B lottery wins and patenting using a different identification strategy. Across various specifications, we also find insignificant and near negative effects of H-1B immigration on firm-level patent grants and patent applications.

¹³We add 1 to the patent count before we take the \ln to address values of 0 in the dependent variable that would otherwise be undefined when logged. This modification makes ex-post interpretation of the coefficient slightly more nuanced, as it doesn't easily fit the percentage change interpretation of logged regression variables that is traditionally used.

¹⁴We also include the 2008 lottery, the largest of the 3 years. Beyond just patent grants, we also examine patent applications, giving us a longer time window for analysis: because of a significant lag between patent application and patent grants, patent grants may be insufficient to observe immigrant effects on patents in the relatively recent time period of 2006–2008, making the patent application data useful for enlarging the sample and ruling out patent approval lags as an explanation for the null effects.

4 Data

We construct our dataset from U.S. Citizenship and Immigration Services (USCIS) administrative records and U.S. Patent and Trademark Office (USPTO) patent records.

4.1 Immigration Records

The administrative H-1B immigration records were obtained through five Freedom of Information Act (FOIA) requests made from 2012 to 2014. Our original dataset contains the universe of H-1B petitions from 1999 to 2012, consisting of over 3.6 million petitions, with information on the final decision regarding each petition and other characteristics of the immigrant and the employer which are further detailed in [Section 6](#). We manually check firm names to correct for errors and to aggregate petitions clearly made by subsidiaries of larger entities.

We replicate the lottery sample for the years of 2006, 2007, and 2008 from the full set of petitions. First, we limit the sample to the set of petitions with receipt dates before or on the “final receipt date” as announced by the USCIS in their press releases. Second, we eliminate petitions from non-profit entities. Third, we drop beneficiaries with technical masters and Ph.D. degrees as they would be exempt from the regular cap and thus would not be subject to the lottery. Fourth, we retain only regular filings made for new employment and exclude filings related to changing visa status, extending the immigrant’s stay, or amending immigrant’s stay. Finally, we keep only petitions on behalf of beneficiaries not currently in the U.S. to create the strictest sample possible. While this sample is largely representative of the full H-1B lottery, one drawback of this method of data construction is that we cannot completely distinguish between petitions declined because of bureaucratic issues (such as failing to correctly fill out an employer identification number) and those declined because they were not chosen in the lottery. The number of these bureaucratically declined applica-

tions is small relative to the lottery because the petition is almost universally completed by expert legal help; thus, rejections should be stochastic and not reflect the intrinsic quality of the petitioning firm. As a robustness check, we implement the generalized differences-in-differences model described in the last section to control for underlying time-invariant firm quality.

4.2 Patent Records

We utilize patenting outcomes—patent grants and patent applications—as our main dependent variables as measures of innovation. Our patent grant data is from the IQSS Patent Network database (Li et al., 2014). The IQSS Patent Network database extends until the end of 2013. We also include patent applications as a dependent variable in our study to give us a longer time window after the lottery event. Our patent application data is from the USPTO/Google Patent Application Publication dataset, available from 2001. Patent grants may understate the effect of H-1B immigration on innovative productivity as a long review period truncates what we can observe. Between 1976 and 1996, patent applications took anywhere from 1 to 1,143 months to be granted, with a mean of 28.4 months (Popp et al., 2004). The variation in patent grant time is driven mostly by idiosyncratic factors, although there is systematic variation in patent grant lag across technological classes. Assuming the lottery is uncorrelated with the patent grant lag across technological classes, our results are not biased by different lag lengths in different technological classes. Given that our patent grant data only extended until the end of 2013, based upon a two to three year patent grant lag, we have a four year window of observation for our 2006 lottery (2007–2010) and a three year window of observation for our 2007 lottery (2008–2010), with many patent grants missing at the tail end of the observation window as some patent applications have not been granted yet. We do not use the 2008 lottery in our study of patent grants because of the short post-lottery observation window, but we do study the impact of the 2008 lot-

tery on patent applications. The patent applications provide a longer observation window, up until 2012. Patent applications of course do not fully translate into patent grants, and some applications are rejected in the process. The grant rate is fairly high though, as 72.3% of patent applications filed in January 2001 were published before April 2006 (Lemley and Sampat, 2012), so it is a reasonable proxy for patenting outcomes and at minimum a clear measure of patenting activity.

We fuzzy match petitioning firm names with patent assignees using the Microsoft Research and Microsoft Business Intelligence algorithm at a 0.85 level (Arasu et al., 2011),¹⁵ and we then manually verify every single one of the generated matches. We do so separately for the IQSS and Google Patent data. We were unable to use citation-weighted patent counts (Trajtenberg, 1990), commonly considered to be a better measure of innovation impact, because there is not enough of a time window after the lotteries. Our data is limited to patents granted up to 2013, which realistically means that we can only observe patents up to those applied for in late 2010. We do not have a complete sample of the most recent patent grants, because patents applied for in 2010 have not been granted yet if the process took longer than three years.

4.3 Descriptive Statistics and Randomization Check

The descriptive statistics are presented in Table 1. The data summarized here is structured as a balanced firm-year panel from 2005 to 2012. Patent grant data are available only up until 2010, and patent application data are available up until 2012. To gain a sense of patent productivity of the average firm in our sample and their average interaction with the lottery, in 2008 the average firm in our sample of the lottery submitted 1.7 patent applications and received 1.3 patent grants. In the same year, the average firm in our lottery sample

¹⁵Arasu, Chaudhuri, Chen, Ganjam, Kasushik, and Narasayya at Microsoft developed the fuzzy matching technology that made this project possible. We leave it to the reader to recognize and appreciate the irony of this.

submitted petitions for 6.3 and won 4.4 visas in the lottery.

—————Insert [Table 1](#)—————

To assess whether the lottery was random, we implement a placebo test in which we regress lagged patents granted on the lottery win share; results of this test are in [Table 2](#). The regression demonstrates that there is a precise zero correlation between the firm’s pre-lottery patent rate and the share of applications the firm won in the lottery. The consistent zero-results suggest the lottery is likely random, although we do implement robustness checks that perform regardless of the assumption of randomness.

—————Insert [Table 2](#)—————

5 Results

Across all our specifications, we find little to no evidence that an additional H-1B visa allocated to a firm is associated with gains in firm-level patenting productivity.

5.1 Petition Bin Results

Results for our baseline specification of the Petition Bin model are presented in [Table 3](#). Winning an H-1B petition has no statistically significant effect on patents granted to the firm. The coefficients are small, insignificant and fall on either side of zero. The top row shows the estimated effect of an H-1B win on patents granted in that column’s year. The estimates to the left of the vertical bar function as placebo estimates. The estimate reflects the average effect of one additional H-1B worker on the number of patents a firm has been granted in that year.¹⁶ The last column reflects the average per-year effect of an H-1B

¹⁶While the results are statistically insignificant, the coefficient on *Wins in 2007* in model (12) would imply that an average firm that won one additional H-1B worker would receive 0.003 additional patent grants three years later in 2010 (but this is not statistically significant and should not be interpreted as such).

worker after the lottery. Because there is a significant lag between patent application and patent grant, the grant measure may understate the effect of H-1B workers since a patent the immigrant contributed to would not even be reviewed within the observable patent data. We incorporate data on patent applications, which does not have the review lag.

—————Insert **Table 3**—————

Table 4 presents a similar exercise as **Table 3**, but the dependent variable is patent applications rather than patent grants. Again, our results are primarily insignificant, with the notable exception of the 2008 lottery. The effect is very close to zero in the 2006 lottery (model (9)), grows to a statistically insignificant 0.08 in the 2007 lottery (model (17)), and becomes statistically significant and positive in the 2008 lottery. However, we find that the patent-application productivity is higher for firms that win the 2008 lottery before the lottery, shown in models (18) and (19) of **Table 4**. In 2008, the lottery selected higher productivity firms; the difference in patent productivity before and after the lottery are positive and consistently significant at least the 10% level. The most likely case is that the lottery randomly selected firms with a higher pre-lottery patent-application productivity, which indeed could happen in one of ten draws. It is also possible but unlikely that the lottery was not fully random and influenced by the extant political environment in the 2008 election year. Given our results for 2008 in **Table 4**, there is some concern that the lottery could be random while also selecting for high performing firms, which would affect the identification assumption of this model.

—————Insert **Table 4**—————

5.2 Differences-in-Differences Results

To control for the implications of lottery draws skewed towards a high performing set of firms, we implement a generalized differences-in-differences design with firm fixed effects and

with petition bins that would yield unbiased results even if the lottery were non-random, namely if some firms are better at winning the lottery than others and that this ability is correlated with firm patenting productivity. The results of this alternative model are presented in [Table 5](#), with a dependent variable of patent grants, and [Table 6](#), with a dependent variable of patent applications. The results from this specification are largely the same, but the coefficient estimates are even smaller, usually insignificant, and sometimes negative and significant. We find some significant results on our OLS Log DV models, but these coefficients demonstrate a negative relationship between visa wins and firm patenting and thus there is still no evidence for a positive relationship. The coefficient on the logged dependent variable represents the approximate percentage change in patents, so [Table 5](#) model (2) is interpreted as showing a 2% reduction in patenting for an additional H-1B visa won, significant to the 5% level. The 2008 results shift downward in the differences-in-differences framework, suggesting that the lottery randomly or non-randomly chose more productive firms in 2008.

————— [Insert Table 5](#) —————

————— [Insert Table 6](#) —————

5.3 Year-to-Year Adjustment of Petition Volume

We further document an aspect of firm strategy with regards to the visa program: firms may adjust their H-1B petition filing strategy based upon the previous year’s success or failure. Our estimated effects would be understated if, for instance, firms that lose the lottery apply for more visas the next year; in this case, the firm is only without the productive capital of an H-1B worker for a single year. Randomly winning (losing) a visa in 2007 would decrease (increase) the number of applications submitted in the next year 2008. On the other hand, our estimates may be overstated if there were returns to scale in hiring skilled immigrants; in

this case, randomly winning a visa in 2007 would be positively correlated with the number of petitions a firm submits in 2008, or in other words, winning (losing) one visa in 2007 increases (reduces) the firm’s H-1B petition filings in 2008. We find that an additional H-1B visa received in 2007 causes firms to petition for more visas in 2008, as shown in [Table 7](#); one additional H-1B visa won in the 2007 lottery causes the average firm to apply for about 0.2 more H-1B visas the following year 2008, significant to a 5% level. Since firms apply for more visas in later years when the win in a previous year, this year-to-year petition volume effect does not attenuate any possible positive effect of visa wins.

—————**Insert [Table 7](#)**—————

5.4 Results Summary

The null results are striking in the context of the existing literature, the majority of which demonstrates a positive relationship between the H-1B program and patenting. Some caution is needed in interpreting our particular results. The estimates reflect the average effect of a non-master’s degree H-1B immigrant on patent applications and grants. The results do not, for instance, speak directly to the impact of the subset of skilled immigrants educated at elite American universities, as many of the immigrants were educated at foreign universities of variable quality. More importantly, our study focuses only on the effect of H-1B visas subject to the lottery, i.e. those without a relevant master’s or doctorate degree; petitions for immigrants with that higher level of education beyond the bachelor’s degree are considered in a different pool of visas, and thus not subject to the lottery.

6 Implications of Program Composition

The results prompt a major question: why would an H-1B win have no positive effect on patent productivity at the level of the petitioning firm? We use our detailed data on

petitions, covering both characteristics of the potential immigrants and their employers, to shed light on why the effects are plausibly zero. We argue that the nature and composition of the skills of these particular immigrants and the industries of their employers make it unsurprising that we would find no effect on patenting productivity. The figures presented in this section contain data from our sample of petitions from 1999 to 2012, including both lottery and non-lottery petitions; the conclusions are the same when we isolate to just the sample of lottery petitions.

We first document the distribution of employer industries and the occupation types of the immigrants. [Figure 1](#) shows the distribution of firm types participating in H-1B program. 48.8% of petitions originate from computer-related firms. The next largest categories are Architecture & Engineering (12.2%), Education (9.6%) and Administrative Specializations (8.4%). Life Science firms account for 2.71% of the petitions. Typically, we do not expect any impact on firm-level patenting from educational or administrative firms because they do not engage in research and development that would lead to patents.

—————**Insert [Figure 1](#)**—————

[Figure 2](#) shows the distribution of occupation types (as defined in the Dictionary of Occupational Types by the US Department of Labor) in H-1B petitions. Diving further into the occupations that beneficiaries are being placed into, we find that 42.9% of the occupations are in “Systems Analysis and Programming,” which represents information technology support roles, which is clearly unlikely to generate intellectual property, and software engineers, which as we discuss next, are also unlikely to generate substantial intellectual property.

—————**Insert [Figure 2](#)**—————

The H-1B program is clearly weighted towards computer-related firm and employment in software, which matters because of heterogeneity in patenting along these dimensions. A patent’s ability to internalize the social benefits of innovation varies widely. Software

has been known to have historically weak patent protection (Bessen and Maskin, 2009). In industries such as life sciences, chemicals, and semiconductors, patents have been more effective at protecting intellectual property arising from R&D, generating a greater value to patenting (Arora et al., 2008). Furthermore, two recent US Supreme Court cases, *Bilski v. Kappos* in 2010 and *Alice v. CLS Bank* in 2014, generated ambiguity around the value of software patents and further reduced the interest in patenting in this space.

In addition, software patents are not predominantly generated by software firms. Bessen and Hunt (2007) construct a dataset of software patents and find that software patents tend to be assigned to firms in industries known to accumulate large patent portfolios and to pursue patents for strategic reasons (computers, electrical equipment, and instruments), particularly in manufacturing and not to actual software publishers, who only hold 5% of their sample of software patents. The manufacturing sector acquires 75% of software patents but employs only 11% of programmers and analysts. These patent portfolios are often part of “patent thickets,” a set of overlapping patents used by practicing firms for defensive purposes and for offensive purposes by non-practicing firms such as patent assertion entities.

In contrast to software, patenting is much more common in life sciences occupations and firms because it is perceived to be a more effective tool for value capture. To study this further, we run our main analyses (Petition Bin Model & Differences-in-Differences Model) for the set of petitions by firms applying for immigrants specializing in life sciences occupations. We find no significant effects.¹⁷ This result may be due to a lack of statistical power given the small sample size. But another possibility is that the set of cap-subject immigrants in life sciences may indeed have low patenting productivity. As legally defined, the set of cap-subject immigrants we are studying are those without master’s degrees and doctorates, because they wouldn’t be cap-subject in these years if they had those degrees. It is difficult to presume that a life science employee with just an undergraduate degree would

¹⁷The results from this analysis are available from the author but omitted here for brevity.

be very effective at generating patents in a field where a Ph.D. is common and necessary training to conduct independent research. The more likely reality is that the cap-subject life science occupation immigrants are taking support roles in laboratories and not leading independent research.

We find that the patent classes represented by the firms participating in the H-1B program do not reflect the full distribution of patent classes. [Figure 3](#) shows the 10 most popular patent classes among patents by H-1B firms and there is a clear lean toward software and information technology patents, which are less common in the full set of all USPTO patents.

—————**Insert [Figure 3](#)**—————

Furthermore, the distribution of the technology skill set and education levels of the immigrants is connected to their regional origin. This distribution of H-1B beneficiary country of origin is shown in [Figure 4](#). The distribution of H-1B immigrants is heavily skewed towards India, which is a net exporter of labor to the United States. India has been historically known to produce a large amount of engineers, particularly in software ([Arora et al., 2001](#); [Banerjee, 2008](#)).¹⁸ In consideration of the variance in immigrant education across countries of origin, European immigrants are more likely to immigrate with advanced degrees than say Asian immigrants, which are more likely to have only Bachelors degrees ([Kenney and Patton, 2015](#)). This pattern is consistent with what we observe: these immigrants are heavily skewed towards Asian countries (Canada, Mexico, and the United Kingdom are the only non-Asian countries in the top 10 H-1B origin countries), and the component of the program we study is tailored for those with bachelor’s degrees.

While not captured in our data, it has been noted that H-1B immigrants are younger than domestic employees who might otherwise take the occupation ([Kerr et al., 2015](#)). Greater age is found to be beneficial to innovative impact ([Jones, 2010](#)).

¹⁸On the Indian end, work has shown that inventors hired from foreign firms generate higher impact inventions than those hired from domestic Indian firms ([Alnuaimi et al., 2012](#)).

—————**Insert Figure 4**—————

Finally, we reflect on the distribution of visas that would occur in an expansion of the program by looking more closely at the largest users of the H-1B visa program. [Figure 5](#) shows the top 10 largest filers of H-1B petitions in 2007. With the exception of Microsoft Corporation, the other nine companies are all what would be classified as information technology consulting firms, best known for outsourcing services. These firms provide services from foreign labor based located in a foreign country, which for these firms is usually India, and they are experts in navigating immigration procedures. Once they obtain a visa for these workers, the firms place the worker with a client. Because these firms are functionally working for another firm, the visa-sponsoring firm likely does not own patents resulting from the immigrant. We cannot ascertain the identity of their clients and contracted placement of the immigrants as this data is not broadly collected by the U.S. government.

—————**Insert Figure 5**—————

To summarize, our main empirical findings are explained by the observation that the majority of firms do not hire H-1B immigrants to generate innovation that is captured by patents. The median H-1B immigrant is a young Indian software engineer or information technology support specialist and not a patent-generating researcher or scientist, and many of these immigrants are not even employed by their hosting firm.

We cannot rule out that these immigrants may have impacts on their neighboring firms and communities via knowledge spillovers ([Jaffe et al., 1993](#)) and other benefits to regional productivity. Regional data on H-1Bs and patenting behavior would be able to capture spill-over effects, which we would not observe here, but it would be captured in prior studies that study the regional impacts of H-1B immigration. [Kerr and Lincoln \(2010\)](#) show that there were large increases in the rate of Indian and Chinese patenting in cities and firms that depend on H-1B visas when H-1B visa availability expanded.

7 Conclusion

We estimate the causal impact of H-1B workers on the patent production of American firms by exploiting a lottery that *randomly* issues visa to petitioning firms. We find that traditional H-1B immigrants (no relevant master’s or Ph.D.) have no observable impact on a firm’s patent applications or patent grants. Using a rich dataset on the universe of H-1B petitions, we demonstrate that the H-1B program is primarily employed in firms and occupations that do not contribute to innovation as captured by patents. Our results provide new evidence on the impact and role of skilled immigration. Our result is surprising given a substantial prior literature demonstrating a robust positive correlation between H-1B visa use and patent production. These results are particularly important in light of an on-going debate surrounding the use of immigration to meet domestic labor demand.

The present work leaves much to be explored in the context of understanding the H-1B program and skilled immigration more generally, both in the United States and in other countries. Our analysis focuses on the impact of H-1B immigrants without a relevant education beyond a bachelor’s degree, and our results do not speak directly to the innovative value-add of highly skilled immigrants trained in prestigious American universities or the H-1B visas offered to immigrants with a master’s degree. The role of the program for master’s and Ph.D. educated immigrations needs to be further studied as the cap on that portion of the H-1B program has also been reached in the years after the program years in our study. There are also other categories of the U.S. skilled immigration that have yet to be well studied, including the specialized H-1B1 program for Chile and Singapore, the H-1B2 program related to U.S. Department of Defense R&D, O-1 visa for individuals with “extraordinary ability or achievement”, the F visa for academic training, and the J-1 visa for research scholars and cultural exchange visitors. Beyond the visa programs, there is also a similar set of open questions related to the impact of lawful permanent residency programs,

informally known as a green card: permanent residency provides a much longer time window and a greater breadth of employment options, including entrepreneurship, to the immigrant. This paper should therefore be viewed as just a step toward characterizing the impact of high-skilled immigration on American innovation.

References

- Allison, Paul D., Richard P. Waterman. 2002. Fixed-effects negative binomial regression models. *Sociological Methodology* **32**(1) 247–265.
- Alnuaimi, Tufool, Tore Opsahl, Gerard George. 2012. Innovating in the periphery: The impact of local and foreign inventor mobility on the value of Indian patents. *Research Policy* **41**(9) 1534–1543.
- Angrist, Joshua D., Susan M. Dynarski, Thomas J. Kane, Parag A. Pathak, Christopher R. Walters. 2012. Who benefits from KIPP? *Journal of Policy Analysis and Management* **31**(4) 837–860.
- Arasu, Arvind, Surajit Chaudhuri, Zhimin Chen, Kris Ganjam, Raghav Kaushik, Vivek Narasayya. 2011. Towards a domain independent platform for data cleaning. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* .
- Arora, Ashish, V. S. Arunachalam, Jai Asundi, Ronald Fernandes. 2001. The Indian software services industry. *Research Policy* **30**(8) 1267–1287.
- Arora, Ashish, Marco Ceccagnoli, Wesley M. Cohen. 2008. R&D and the patent premium. *International Journal of Industrial Organization* **26**(5) 1153–1179.
- Ballmer, Steven A., Michael R. Bloomberg, Julian Castro, Phil Gordon, Robert Iger, William Marriott, James McNerney, Rupert Murdoch, Michael Nutter, Antonio Villaraigosa. 2011. The “New American” Fortune 500. Tech. rep., Partnership for a New American Economy.
- Banerjee, Payal. 2008. Flexible hiring, immigration, and Indian IT workers’ experiences of contract work in the United States. Marjorie L. DeVault, ed., *People at Work: Life, Power, and Social Inclusion in the New Economy*, 1st ed., chap. 5. New York University Press, New York, 97–111.
- Baruffaldi, Stefano H., Paolo Landoni. 2012. Return mobility and scientific productivity of researchers working abroad: The role of home country linkages. *Research Policy* **41**(9) 1655–1665.
- Bessen, James, Robert M. Hunt. 2007. An empirical look at software patents. *Journal of Economics & Management Strategy* **16**(1) 157–189.
- Bessen, James, Eric Maskin. 2009. Sequential innovation, patents, and imitation. *RAND Journal of Economics* **40**(4) 611–635.
- Bhagwati, Jagdish, Koichi Hamada. 1974. The brain drain, international integration of markets for professionals and unemployment. *Journal of Development Economics* **1**(1) 19–42. ID: 271688271688.

- Black, Dan A., Jeffrey A. Smith, Mark C. Berger, Brett J. Noel. 2003. Is the threat of reemployment services more effective than the services themselves? Evidence from random assignment in the UI system. *American Economic Review* **93**(4) 1313–1327.
- Bock, Laszlo. 2013. Google supports high-skilled immigration reform.
- Bound, John, Sarah Turner, Patrick Walsh. 2009. Internationalization of U.S. doctorate education. Richard B. Freeman, Daniel Goroff, eds., *Science and Engineering Careers in the United States: An Analysis of Markets and Employment*, chap. 2. University of Chicago Press, Chicago, IL, 59–97.
- Bresnahan, Timothy, Alfonso Gambardella. 2004. *Building high-tech clusters: Silicon Valley and beyond*. Cambridge University Press, Cambridge, UK.
- Chellaraj, Gnanaraj, Keith E. Maskus, Aaditya Mattoo. 2008. The contribution of international graduate students to us innovation. *Review of International Economics* **16**(3) 444–462.
- Criscuolo, Paola. 2005. On the road again: Researcher mobility inside the R&D network. *Research Policy* **34**(9) 1350–1365.
- Davenport, Sally. 2004. Panic and panacea: brain drain and science and technology human capital policy. *Research Policy* **33**(4) 617–630.
- Docquier, Frdric, Hillel Rapoport. 2012. Globalization, brain drain, and development. *Journal of Economic Literature* **50**(3) 681–730.
- Doran, Kirk, Alexander Gelber, Adam Isen. 2015. The effects of high-skilled immigration policy on firms: Evidence from H-1B visa lotteries. *NBER Working Papers* .
- Eidler, Jakob, Heide Fier, Christoph Grimpe. 2011. International scientist mobility and the locus of knowledge and technology transfer. *Research Policy* **40**(6) 791–805.
- Filatotchev, Igor, Xiaohui Liu, Jiangyong Lu, Mike Wright. 2011. Knowledge spillovers through human mobility across national borders: Evidence from Zhongguancun Science Park in China. *Research Policy* **40**(3) 453–462.
- Gibson, John, David McKenzie. 2014. Scientific mobility and knowledge networks in high emigration countries: Evidence from the Pacific. *Research Policy* **43**(9) 1486–1495.
- Hausman, Jerry, Bronwyn H. Hall, Zvi Griliches. 1984. Econometric models for count data with an application to the patents-R&D relationship. *Econometrica* **52**(4) 909–938.
- Hunt, Jennifer. 2011. Which immigrants are most innovative and entrepreneurial? Distinctions by entry visa. *Journal of Labor Economics* **29**(3) 417–457.
- Hunt, Jennifer, Marjolaine Gauthier-Loiselle. 2010. How much does immigration boost innovation? *American Economic Journal: Macroeconomics* **2**(2) 31–56.

- Jaffe, Adam B., Manuel Trajtenberg, Rebecca Henderson. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics* **108**(3) 577–598.
- Janger, Jrgen, Klaus Nowotny. 2016. Job choice in academia. *Research Policy* **45**(8) 1672–1683.
- Jones, Benjamin F. 2010. Age and great invention. *Review of Economics and Statistics* **92**(1) 1–14.
- Jonkers, Koen, Laura Cruz-Castro. 2013. Research upon return: The effect of international mobility on scientific ties, production and impact. *Research Policy* **42**(8) 1366–1377.
- Kahn, Shulamit, Megan MacGarvie. 2016. Do return requirements increase international knowledge diffusion? evidence from the Fulbright program. *Research Policy* **45**(6) 1304–1322.
- Kenney, Martin, Donald Patton. 2015. Gender, ethnicity and entrepreneurship in initial public offerings: Illustrations from an open database. *Research Policy* **44**(9) 1773–1784.
- Kerr, Sari Pekkala, William R. Kerr, William F. Lincoln. 2015. Skilled immigration and the employment structures of US firms. *Journal of Labor Economics* **33** S147–S186.
- Kerr, William R, 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.
- Lemley, Mark A., Bhaven Sampat. 2012. Examiner characteristics and patent office outcomes. *Review of Economics and Statistics* **94**(3) 817–827.
- Lepori, B., M. Seeber, A. Bonaccorsi. 2015. Competition for talent. country and organizational-level effects in the internationalization of European higher education institutions. *Research Policy* **44**(3) 789–802.
- Li, Guan-Cheng, Ronald Lai, Alexander DAmour, David M. Doolin, Ye Sun, Vetle I. Torvik, Amy Z. Yu, Lee Fleming. 2014. Disambiguation and co-authorship networks of the U.S. patent inventor database (1975–2010). *Research Policy* **43**(6) 941–955.
- Ochel, Wolfgang. 2000. Recruitment of highly skilled labour abroad: The practice of selected OECD countries. *CESifo Forum* **1**(2) 48–52.
- Popp, David, Ted Juhl, Daniel KN Johnson. 2004. Time in purgatory: Examining the grant lag for US patent applications. *Topics in Economic Analysis & Policy* **4**(1).
- Qin, Fei. 2015. Global talent, local careers: Circular migration of top Indian engineers and professionals. *Research Policy* **44**(2) 405–420.
- Samuelson, Paul A. 1964. Theoretical notes on trade problems. *Review of Economics and Statistics* 145–154.

- Saxenian, AnnaLee. 1996. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Harvard University Press, Cambridge, MA.
- Scellato, Giuseppe, Chiara Franzoni, Paula Stephan. 2015. Migrant scientists and international networks. *Research Policy* **44**(1) 108–120.
- Storz, Cornelia, Federico Riboldazzi, Moritz John. 2015. Mobility and innovation: A cross-country comparison in the video games industry. *Research Policy* **44**(1) 121–137.
- Trajtenberg, Manuel. 1990. A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics* **21**(1) 172–187.
- Trumka, Richard. 2013. AFL-CIO: American high-tech workers not in shortage.

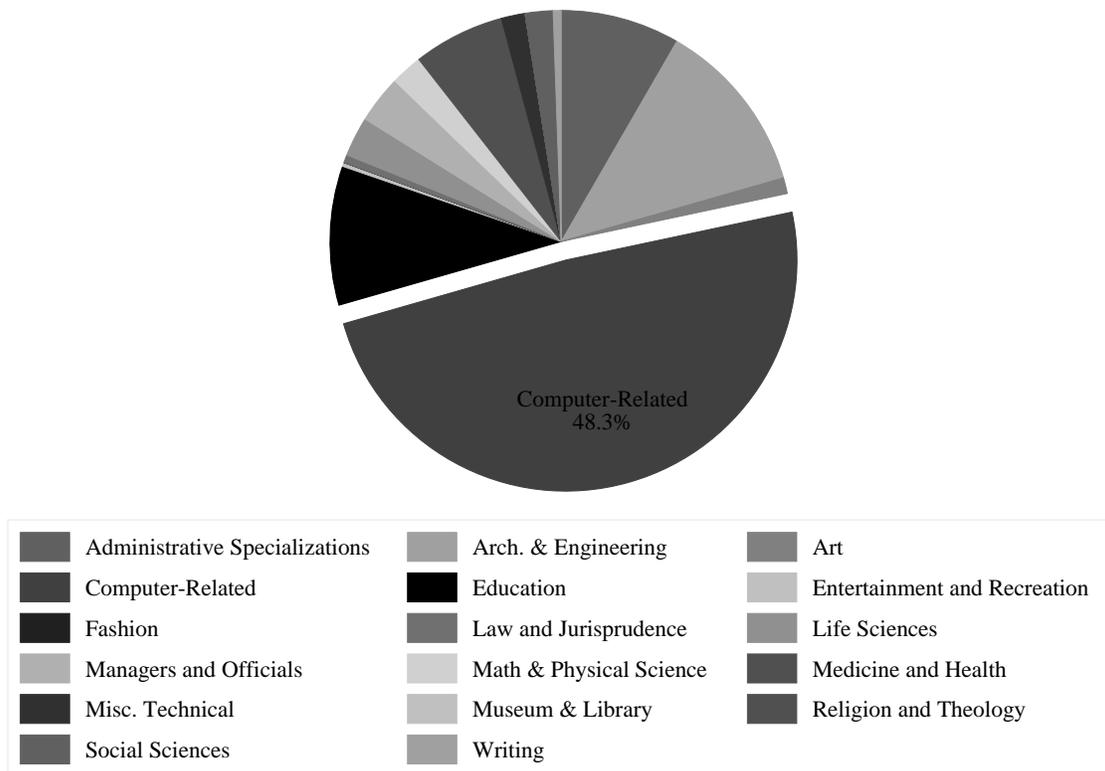


Figure 1: **Firm Industries.** This figure depicts distribution of firm industries who filed petitions in the H-1B program. The data is from United States Citizenship and Immigration Services (USCIS) and represents the full sample of H-1B petitions filed from 1999 to 2012.

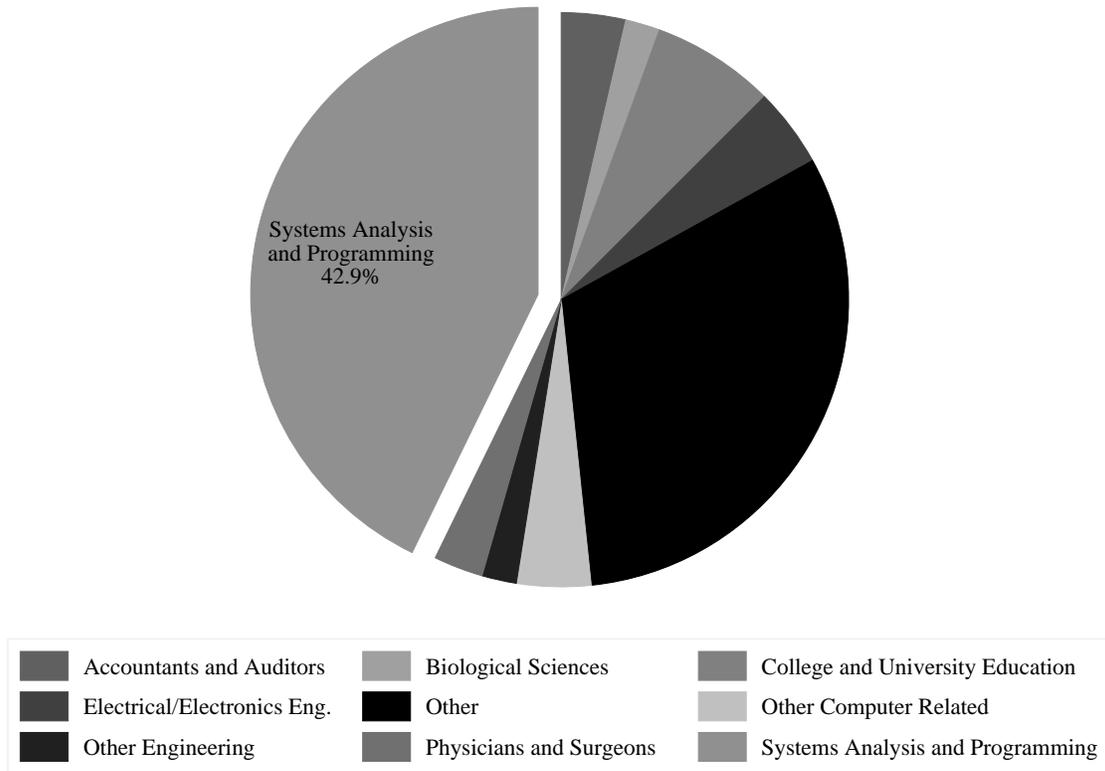


Figure 2: **Occupation Types.** This figure depicts distribution of beneficiary (immigrant) occupation types for petitions filed in the H-1B program. The occupation types are defined in the *Dictionary of Occupational Types* by the U.S. Department of Labor. The largest category, "Systems Analysis and Programming", represents information technology support roles, such as software developers and technical support. The data is from United States Citizenship and Immigration Services (USCIS) and represents the full sample of H-1B petitions filed from 1999 to 2012.

Class	Class Title	H-1B Firms	USPTO
705	Data Processing: Financial, Business Practice, Management, or Cost/Price Determination	4.2%	0.7%
707	Data Processing: Database and File Management or Data Structures	3.1%	1.2%
514	Drug, Bio-Affecting and Body Treating Compositions	2.4%	1.3%
435	Chemistry: Molecular Biology and Microbiology	2.4%	1.8%
370	Multiplex Communications	2.4%	1.7%
709	Electrical Computers and Digital Processing Systems: Multicomputer Data Transferring	2.1%	1.5%
455	Telecommunications	2.0%	1.7%
424	Drug, Bio-Affecting and Body Treating Compositions	1.8%	1.0%
73	Measuring and Testing	1.5%	0.9%
340	Communications: Electrical	1.5%	1.0%

Figure 3: **Top Patent Classes.** This figure shows the top 10 most popular classes for patents filed by firms who filed at least one petition in the H-1B program from 1999 and 2012. The columns *Class* and *Class Title* are the official class number and name as defined by the United States Patent and Trademark Office (USPTO). The column *H-1B Firms* represents the percentage of all the patents filed by H-1B participating firms in that patent class. The column *USPTO* represents the percentage of all patents in that patent class in patents applied for in the year 2007. The data is from United States Citizenship and Immigration Services (USCIS) and the United States Patent and Trademark Office (USPTO).

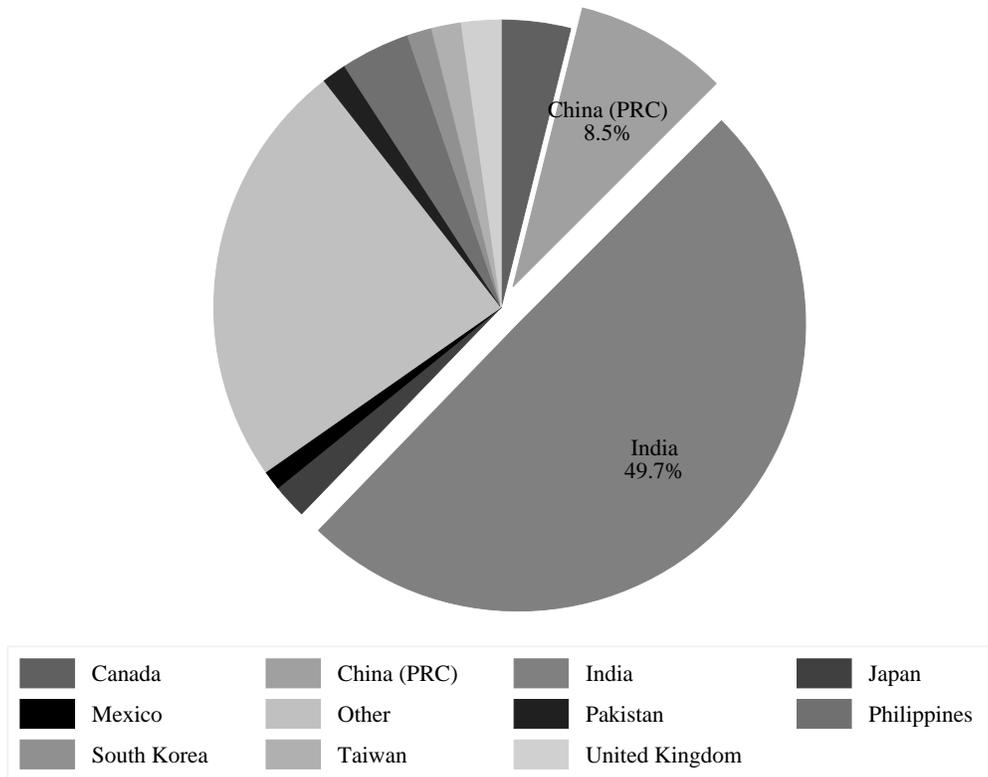


Figure 4: **Country of Origin.** This figure depicts country of origin of beneficiaries (immigrants) who had H-1B petitions filed on their behalf. The data is from United States Citizenship and Immigration Services (USCIS) and represents the full sample of H-1B petitions filed from 1999 to 2012.

Firm	Petitions
Infosys Technologies Limited	4175
Wipro Limited	2253
Satyam Computer Service Limited	1131
Cognizant Technology Solutions Corporation	979
Tata Consultancy Services Limited	576
Patni Computer Systems Limited	435
Microsoft Corporation	331
US Technology Resources, LLC	322
Accenture PLC	320
Larsen & Toubro Infotech Limited	257

Figure 5: **Top H-1B Petitioners.** This figure shows the top ten largest filers of H-1B petitions in 2007. With the exception of *Microsoft Corporation*, the other nine of top ten are information technology consulting firms, commonly associated with outsourcing services. The data is from the administrative records of the United States Citizenship and Immigration Services (USCIS) and represents the full sample of H-1B petitions filed in 2007.

Table 1: **Summary Statistics.** This table presents summary statistics for our firm-year data composed of firms filing H-1B petitions in the 2006, 2007, and 2008 lotteries. H-1B petition data is from U.S. Citizenship and Immigration Services (USCIS), drawn from the complete set of H-1B petitions from 1999 to 2012 obtained via FOIA request. Patent grant data is from the IQSS Patent Network database (Li et al., 2014), and it is only available up until 2010. Patent application data is from the USPTO/Google Patent Application Publication system, and it is only available up until 2012.

2006	Mean	Std. Dev.	Min	Max	Obs.
Patent Grants	3.40	71.93	0	3152	6024
Patent Applications	5.79	107.39	0	3085	8032
H-1B Wins	2.26	5.40	0	90	10040
H-1B Petitions	2.88	6.27	1	96	10040
H-1B Share Won	0.77	0.37	0	1	10040
2007	Mean	Std. Dev.	Min	Max	Obs.
Patent Grants	1.45	32.39	0	3152	45720
Patent Applications	1.93	45.63	0	3085	60960
H-1B Wins	3.83	38.97	0	2242	76200
H-1B Petitions	6.02	58.96	1	4175	76200
H-1B Share Won	0.69	0.40	0	1	76200
2008	Mean	Std. Dev.	Min	Max	Obs.
Patent Grants	1.27	28.50	0	3152	42150
Patent Applications	1.74	38.50	0	3085	56200
H-1B Wins	4.36	50.93	0	2735	70250
H-1B Petitions	6.30	74.18	1	4778	70250
H-1B Share Won	0.73	0.39	0	1	70250

Table 2: **Lottery Placebo Test.** This table presents the results of the lottery placebo test. We regress patent grants filed by a firm in the year before the lottery on its share of wins for a given lottery in 2006, 2007, or 2008. This model is estimated using ordinary least squares (OLS). Statistical significance is represented by $*p < 0.10$, $**p < 0.05$, and $***p < 0.01$. Robust standard errors are shown in parentheses.

	(1)	(2)	(3)
Petition Win Share	2006	2007	2008
Patents	0.000 (0.000)	0.000 (0.000)	0.001 (0.000)
Constant	0.766*** (0.012)	0.691*** (0.005)	0.730*** (0.005)
Observations	1004	7620	7025

Table 3: **Petition Bin Model (Patent Grants)**. This table presents the results of the *Petition Bin* model with the dependent variable of *Patent Grants*. *All* contains all available post-treatment (post-lottery) years. This model is estimated using ordinary least squares (OLS). Statistical significance is represented by $*p < 0.10$, $**p < 0.05$, and $***p < 0.01$. Robust standard errors are shown in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Patent Grants	2005	2006	2007	2008	2009	2010	All
H-1B Wins in 2006	0.808 (1.246)	0.440 (0.841)	-0.037 (0.501)	0.036 (0.24)	-0.050 (0.151)	-0.005 (0.024)	-0.014 (0.225)
Petition Bins	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1004	1004	1004	1004	1004	1004	4016

	(8)	(9)	(10)	(11)	(12)	(13)
Patent Grants	2006	2007	2008	2009	2010	All
H-1B Wins in 2007	0.149 (0.112)	0.090 (0.067)	0.039 (0.026)	0.016 (0.010)	0.003* (0.002)	0.019 (0.012)
Petition Bins	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7864	7864	7864	7864	7864	23592

Table 4: **Petition Bin Model (Patent Applications)**. This table presents the results of the *Petition Bin* model with the dependent variable of *Patent Applications*. *All Years* contains all available post-treatment (post-lottery) years. This model is estimated using ordinary least squares (OLS). Statistical significance is represented by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$. Robust standard errors are shown in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Patent Applications	2005	2006	2007	2008	2009	2010	2011	2012	All Years
Wins in 2006	0.359 (0.493)	0.500 (0.927)	0.096 (1.278)	-0.053 (1.349)	-0.071 (1.003)	0.164 (1.367)	0.022 (0.952)	-0.096 (0.343)	0.010 (1.008)
Petition Bins	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1004	1004	1004	1004	1004	1004	1004	1004	6024

	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)
Patent Applications	2006	2007	2008	2009	2010	2011	2012	All Years
Wins in 2007	-0.001 (0.048)	0.046 (0.070)	0.045 (0.079)	0.055 (0.076)	0.088 (0.096)	0.138 (0.110)	0.062 (0.040)	0.078 (0.076)
Petition Bins	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7864	7864	7864	7864	7864	7864	7864	39320

	(18)	(19)	(20)	(21)	(22)	(23)	(24)
Patent Applications	2007	2008	2009	2010	2011	2012	All Years
Wins in 2008	0.156* (0.083)	0.192** (0.091)	0.174** (0.084)	0.203** (0.096)	0.190** (0.085)	0.059** (0.030)	0.157** (0.071)
Petition Bins	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7525	7525	7525	7525	7525	7525	30100

Table 5: **Differences-in-Differences (Patent Grants)**. This table presents the results of the *Diff-in-Diff with Firm FE* model—in the first 3 columns under the header *Firm FE*—and the *Diff-in-Diff with Petition Bins* model—in the last 2 columns under the header *Petition Bins*—with the dependent variable of *Patent Grants*. The *Diff-in-Diff with Firm FE* model is estimated with ordinary least squares (OLS) with firm fixed effects, OLS with a logged dependent variable $\ln(\text{Patent Count}+1)$ and firm fixed effects, and negative binomial (NBR) with conditional firm fixed effects. The *Diff-in-Diff with Petitions Bins* model is estimated with ordinary least squares (OLS) with firm fixed effects and OLS with a logged dependent variable $\ln(\text{Patent Count}+1)$ and firm fixed effects. Statistical significance is represented by $*p < 0.10$, $**p < 0.05$, and $***p < 0.01$. Robust standard errors are shown in parentheses.

	Firm FE			Petition Bins	
	OLS	OLS Log	NBR	OLS	OLS DV
2006 Lottery (FY 2004–2010)					
Patent Grants	(1)	(2)	(3)	(4)	(5)
Wins * Post (2006)	-7.181 (6.518)	-0.020** (0.009)	-0.055 (0.090)	-2.114 (2.328)	-0.007** (0.003)
Petitions * Post (2006)	4.346 (3.712)	0.011* (0.007)	0.042 (0.083)		
Observations	6024	6024	498	6024	6024
2007 Lottery (FY 2005–2010)					
Patent Grants	(6)	(7)	(8)	(9)	(10)
Wins * Post (2007)	-0.167 (0.214)	-0.001* (0.000)	0.017 (0.013)	0.040 (0.046)	-0.001*** (0.000)
Petitions * Post (2007)	0.087 (0.112)	0.000 (0.000)	-0.017 (0.012)		
Observations	47184	47184	4500	47184	47184
Pooled 2006–2007 Lotteries (FY 2004–2010)					
Patent Grants	(11)	(12)	(13)	(14)	(15)
Wins * Post (All)	-0.089 (0.112)	-0.000 (0.000)	0.014 (0.011)	-0.019 (0.021)	-0.000*** (0.000)
Petitions * Post (All)	0.049 (0.062)	0.000 (0.000)	-0.014 (0.010)		
Observations	77592	77592	6846	77592	77592
Firm FE	Yes	Yes	Yes	No	No
Petition Bins	No	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Table 6: **Differences-in-Differences (Patent Applications)**. This table presents the results of the *Diff-in-Diff with Firm FE* model—in the first 3 columns under the header *Firm FE*—and the *Diff-in-Diff with Petition Bins* model—in the last 2 columns under the header *Petition Bins*—with the dependent variable of *Patent Applications*. The *Diff-in-Diff with Firm FE* model is estimated with ordinary least squares (OLS) with firm fixed effects, OLS with a logged dependent variable $\ln(\text{Patent Count} + 1)$ and firm fixed effects, and negative binomial (NBR) with conditional firm fixed effects. The *Diff-in-Diff with Petitions Bins* model is estimated with ordinary least squares (OLS) with firm fixed effects and OLS with a logged dependent variable $\ln(\text{Patent Count} + 1)$ and firm fixed effects. Statistical significance is represented by $*p < 0.10$, $**p < 0.05$, and $***p < 0.01$. Robust standard errors are shown in parentheses.

	Firm FE			Petition Bins	
	OLS	OLS Log	NBR	OLS	OLS Log
2006 Lottery (FY 2004–2012)					
Patent Applications	(1)	(2)	(3)	(4)	(5)
Wins * Post (2006)	0.823 (1.021)	0.006 (0.006)	-0.064 (0.100)	0.338 (0.352)	0.003 (0.002)
Petitions * Post (2006)	-0.414 (0.612)	-0.002 (0.004)	0.056 (0.092)		
Observations	8032	8032	568	8032	8032
2007 Lottery (FY 2005–2012)					
Patent Applications	(6)	(7)	(8)	(9)	(10)
Wins * Post (2007)	-0.036 (0.026)	-0.000 (0.000)	-0.000 (0.001)	0.002 (0.008)	0.000** (0.000)
Petitions * Post (2007)	0.026* (0.014)	0.000*** (0.000)	0.000 (0.001)		
Observations	62912	62912	4960	62912	62912
2008 Lottery (FY 2006–2012)					
Patent Applications	(11)	(12)	(13)	(14)	(15)
Wins * Post (2008)	-0.092 (0.094)	-0.001 (0.000)	-0.001 (0.001)	0.010 (0.018)	0.000 (0.000)
Petitions * Post (2008)	0.058 (0.053)	0.000* (0.000)	0.001 (0.001)		
Observations	52675	52675	3976	52675	52675
Pooled 2006–2008 Lotteries (FY 2004–2012)					
Patent Applications	(16)	(17)	(18)	(19)	(20)
Wins * Post (All)	-0.031 (0.026)	-0.000 (0.000)	-0.001 (0.000)	-0.000 (0.006)	0.000 (0.000)
Petitions * Post (All)	0.021 (0.015)	0.000** (0.000)	0.000 (0.000)		
Observations	103456	103456	7688	103456	103456
Firm FE	Yes	Yes	Yes	No	No
Petition Bins	No	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Table 7: **Firm Behavior.** This table presents the results of an ordinary least squares (OLS) regression with a dependent variable of *Petitions in 2008*, the number of H-1B petitions filed by a firm in 2008, and an independent variable of *Wins in 2007*, H-1B petitions awarded to the firm in 2007. *Petition Bins* fixed effects, representing petitions filed in 2007, are included. Statistical significance is represented by $*p < 0.10$, $**p < 0.05$, and $***p < 0.01$. Robust standard errors are shown in parentheses.

	(1)
	Petitions in 2008
Wins in 2007	0.197** (0.091)
Petition Bin FE	Yes
Observations	2805