CREAM

**Centre for Research &
Analysis of Migration**

Discussion Paper Series

CPD 42/14

- ▶ **The Impact of Skilled Foreign Workers on Firms: an Investigation of Publicly Traded U.S. Firms**
- ▶ Anirban Ghosh, Anna Maria Mayda and Francesc Ortega

Centre for Research and Analysis of Migration
Department of Economics, University College London
Drayton House, 30 Gordon Street, London WC1H 0AX

www.cream-migration.org

The Impact of Skilled Foreign Workers on Firms: an Investigation of Publicly Traded U.S. Firms *

Anirban Ghosh
Georgetown University

Anna Maria Mayda
Georgetown University and CEPR

Francesc Ortega
Queens College, CUNY

November 17, 2014

Abstract

Many U.S. businessmen are vocally in favor of an increase in the number of H-1B visas. Is there systematic evidence that this would positively affect firms' productivity, sales, employment or profits? To address these questions we assemble a unique dataset that matches *all* labor condition applications (LCAs) – the first step towards H-1B visas for skilled foreign-born workers in the U.S. – with firm-level data on publicly traded U.S. firms (from Compustat). Our identification is based on the sharp reduction in the annual H-1B cap that took place in 2004, combined with information on the degree of dependency on H-1B visas at the firm level as in [Kerr and Lincoln \(2010\)](#). The main result of this paper is that if the cap on H-1B visas were relaxed, a subset of firms would experience gains in average labor productivity, firm size, and profits. These are firms that conduct R&D and are heavy users of H-1B workers – they belong to the top quintile among filers of LCAs. These empirical findings are consistent with a heterogeneous-firms model where innovation enhances productivity and is subject to fixed costs.

*We thank Chad Sparber, Thijs van Rens, Aysegul Sahin, and seminar participants in Bologna, Queens College and SNF Sinergia - CEPR conference on Economic Inequality, Labor Markets and International Trade, for helpful comments. The authors thank NORFACE/TEMPO and the Institute for the Study of International Migration (ISIM) at Georgetown University for financial support. Corresponding authors: Francesc Ortega (fortega@qc.cuny.edu), and Anna Maria Mayda (amm223@georgetown.edu).

“I want to emphasize that to address the shortage of scientists and engineers, we must do both – reform our education system and our immigration policies. If we don’t, American companies simply will not have the talent they need to innovate and compete.” (Bill Gates, Testimony at the U.S. House of Representatives, Committee on Science and Technology on March 12, 2008)

1 Introduction

As our opening quote illustrates, CEOs of large U.S. corporations often advocate passionately in favor of an increase in the number of H-1B visas. They argue that there is a shortage of skilled workers in the U.S. labor market in some fields and that, unless the cap on H-1B visas is raised, their firms will not be able to grow, and their innovation efforts and R&D activities may be at risk.¹ The goal of this paper is to evaluate the effects of H-1B visas on firms’ productivity, sales, employment and profits.

There is abundant anecdotal evidence that the contribution of immigrants to innovation, entrepreneurship and education is substantial in the U.S.. Immigrants account for about one quarter of U.S.-based Nobel Prize recipients between 1990 and 2000, of founders of public-venture-backed U.S. companies in 1990-2005, and of founders of new high-tech companies with at least one million dollars in sales in 2006 (Wadhwa et al. (2007)). These authors also report that 24 percent of all patents originating from the U.S. are authored by non-citizens.²

In addition several studies have established a connection between skilled immigration and patenting activity (Hunt and Gauthier-Loiselle (2010), Kerr and Lincoln (2010) and Parrotta et al. (2014b)). What is less clear is whether the increase in patenting activity has translated into innovation with direct effects on firm’s outcomes.³ By focusing on relevant outcomes, such as sales, productivity, employment, and profits, we can examine whether

¹Several large American technology companies, such as Microsoft, Amazon or Facebook, have recently established offices in Vancouver (Canada). One of the main reasons seems to be the difficulty in obtaining H-1B visas for new hires in the United States and the larger abundance of (foreign) skilled labor in Vancouver. Karen Jones, Microsoft’s deputy general counsel, put it in the following words: “The U.S. laws clearly did not meet our needs. So we have to look to other places.” (Bloomberg Businessweek 2014).

²In addition Borjas (2005) shows that foreign students receive over fifty percent of all doctorates granted in the field of engineering. As a share of college-educated employment, the foreign-born population in the United States has increased from 7 percent in 1980 to over 15 percent in 2010 (Peri et al. (2013)). In comparison the foreign-born share in overall employment increased from 6.4 to 16 percent over the same period. The original data source is the U.S. Census, population ages 18 to 65.

³Since the early 1990s there has been an explosion in the number of patents. As argued by Hall and Ziedonis (2001), in some industries (e.g. semiconductors) the increase in patenting may be due to strategic use by firms. Namely, obtaining patents may allow firms to exercise hold-up or prevent competitors from doing so. A very interesting account of the abuse of patents by some firms (*patent trolls*) is examined in a 2011 episode of ‘This American Life’ (episode 441: When Patents Attack!).

the positive effect of skilled immigration on patenting activity documented in the literature in fact translates into improved firm performance. Our results suggest that this is indeed the case. Thus, H-1B workers appear to increase innovation and improve firm performance.

The main goal of our paper is to evaluate the claim that the current quota on H-1B visas is hampering the innovation and growth of U.S. firms. To do so we assemble a unique dataset that matches the universe of labor condition applications (LCAs) – the first step towards H-1B visas for skilled foreign-born workers in the U.S. – with firm-level data from Compustat, which includes all publicly traded U.S. firms. The merged data set allows us to estimate, at the firm level, the impact of skilled migration on the sales, productivity, employment, profits, and R&D expenditure of publicly traded U.S. firms.

Our identification is based on a sharp change in policy, namely the 2004 reduction in the annual H-1B cap – which offers an opportunity to identify the *causal* effect of H-1B workers on firm outcomes.⁴ Since we do not observe the actual number of H-1B workers in the firm and, even if we did, this variable would be clearly endogenous, we estimate a difference-in-difference specification where the impact of the treatment (the exogenous change in policy) is compared across different categories of firms. Specifically, we follow [Kerr and Lincoln \(2010\)](#) and use the number of LCAs in 2001 to measure each firm’s degree of dependency on H-1B visas. We compare the change in outcomes before and after the policy change across firms, within the same industry, that are more dependent on H-1B visas (the “treatment group”) and less dependent firms (“the control group”). Since the overall H-1B cap was sharply reduced between the two years, we expect firms that were initially more dependent on H-1B visas to display *worse* outcomes, namely, lower growth in productivity, firm size (employment and sales), and profits.

We have two main results. First, we find that increases in the number of H-1B workers lead to growth in productivity, firm size, and profits. However, not all firms benefit equally from an increase in the cap on H-1B visas. According to our estimates, the relationship between H-1B workers and measured firm outcomes is highly non-linear. Only firms that file a relatively large number of applications for H-1B visas appear to benefit from increases in the cap and, conversely, are more negatively affected when this cap is reduced. Secondly, we find that the effects are driven by firms that consistently conduct R&D activity, in addition to relying on H-1B workers.

Our empirical findings can be rationalized with a simple monopolistic competition model

⁴Between 2001 and 2003 the annual cap on H-1B visas was 195,000. In year 2004 it was reduced to 65,000 visas. The cap was raised again to 85,000 visas in 2006.

where firms are heterogeneous in initial productivity. These firms can choose to set up a lab to conduct R&D, which requires hiring skilled workers. To be viable, a lab requires a minimum size, in terms of the number of skilled workers, which amounts to a fixed cost. In addition there is a shortage skilled native workers and some firms need to rely on H-1B visas. Firms that are successful in setting up a lab experience an improvement in their productivity, which in turn has positive effects on sales, employment and profits. One of the strengths of our paper is that we are able to provide evidence for positive effects of H-1B workers on all these firm-level outcomes. In contrast, most of the literature has focused on one of these outcomes at a time, often relying on less comprehensive data.

Our paper is related to the large body of literature trying to measure the effects of different dimensions of globalization, such as trade and immigration, on firms, with an emphasis on implications for productivity and innovation. We discuss each of these in turn. We begin by briefly reviewing the relevant literature on (skilled) immigration and innovation. A seminal study in this literature is [Kerr and Lincoln \(2010\)](#) who focus on the effects of H-1B visas on patenting activity. While closely related, our analysis departs from theirs in important ways. First, our dataset contains the universe of publicly traded firms in the U.S. After merging it with the data on labor-condition applications (LCAs), we end up with almost four thousand firms. In contrast, the firm-level analysis in [Kerr and Lincoln \(2010\)](#) is based on a much smaller sample size (of only 77 firms).⁵ Second, we broaden the scope of the analysis by examining a broader set of outcomes, which includes firm productivity, sales, profits, R&D expenditures, and TFP. Third, as noted above, in terms of identification, we rely on the time-variation arising from a single policy event, the large reduction in the national cap for H-1B visas that took place in year 2004 ([Figure 1](#)). In contrast Kerr and Lincoln exploit year-to-year variation in the stock of H-1B visas.⁶ Our paper is also related to the work of [Hunt and Gauthier-Loiselle \(2010\)](#). Exploiting cross-state variation for the United States, these authors find that a one percentage-point increase in the share of immigrant college graduates in the population leads to an increase in patents per capita of 9 to 18 percent, and the main reason is that they disproportionately hold STEM (Science, Technology, Engineering and Mathematics) degrees. [Parrotta et al.](#)

⁵Moreover these firms have been selected on the basis of high patenting activity, their outcome of interest, or of heavy use of labor condition applications. In contrast in our data the majority of firms did not file any LCAs, although heavy users are also part of our dataset. Thus our median firm is very different from theirs, and representative of all public firms in the U.S.

⁶We decided not to rely on year-to-year variation in the national cap (or stock of H-1B visas). Because of implementation delays, deferrals, or reporting inaccuracies, the data are fairly noisy at an annual frequency. In addition computing the stock of visa holders requires a number of assumptions and imputations ([Lowell \(2000\)](#)), introducing further noise.

(2014b) analyze the connection between worker diversity within a firm and its patenting activity using data for Denmark. Their results suggest that ethnic diversity leads to more patenting. Along similar lines, Chellaraj et al. (2008) document that the presence of foreign graduate students has a positive impact on future patents. More similar to our paper, Peri et al. (2013) use variation in the H-1B cap to try to identify the effects of increases in the population of STEM workers in a city on the wages of skilled and unskilled workers in the same city. They find that H-1B-driven increases in STEM workers are associated with increases in the wages paid to skilled workers (both in STEM and non-STEM occupations), and find no evidence of effects on the wages of unskilled workers.

Let us now turn to the literature on economic openness and productivity. Many studies have explored the effects of international trade on productivity (see, for example, Melitz (2003), Pavcnik (2002), Bernard et al. (2007), Melitz and Ottaviano (2008) or Bustos (2011)), which is closely related to our theoretical framework and our empirical strategy. Our paper is also related to the small literature that studies the effects of immigration on productivity. However, many of the migration studies rely on aggregate data and, as a result, require strong identification assumptions and are more vulnerable to omitted-variable bias.⁷ These studies typically correlate the immigrant share in a country or region (a U.S. state or metropolitan statistical area) with aggregate productivity levels (as in Quispe-Agnoli and Zavadny (2002), Peri (2012), and Ortega and Peri (2014)). Several of these studies find that foreign workers have a positive effect on productivity.⁸ However, these studies typically are not able to disentangle which part of the effect arises from spillovers, which requires different identification strategies (Moretti (2004), Greenstone et al. (2010)).

Very few papers have empirically analyzed the effects of immigration on firm-level productivity.⁹ Paserman (2013) exploits cross-firm and cross-industry variation in the concentration of skilled immigrants and finds evidence of a negative correlation between the immigrant share and output per worker in low-tech industries, whereas the relationship becomes positive for high-tech industries. Parrotta et al. (2014a) analyze the effects of diversity within firms on total factor productivity, using a rich matched employer-employee

⁷For an overview of the findings in the literature on the economic effects of skilled immigration see, for example, Bertoli et al. (2012).

⁸Many studies only use a general measure of immigration. One exception is Ortega and Peri (2014) who distinguish between the effect of overall immigration and the effect of the diversity of immigrants by country of birth.

⁹Teruel-Carrizosa and Segarra-Blasco (2008) explore the effects of immigration on firm profits using data for Spain. While their dependent variable is profitability at the firm level, they only measure immigrant density at the city level. Thus identification is still based on cross-city variation. Dustmann and Glitz (2011) also analyze the effects of immigration into a region on the distribution of firms in that region.

dataset for Denmark. Their estimates point toward a negative association between ethnic diversity and firm-level productivity. In ongoing work, [Trax et al. \(2013\)](#) use German establishment data to estimate the effect of cultural diversity on total factor productivity. Their results suggest that higher immigrant concentration in a firm does not lead to higher TFP. However, they find that higher ethnic *diversity* in the firm or in the region where the firm is located do appear to have positive effects on TFP at the firm level, consistent with the findings in [Ortega and Peri \(2014\)](#). Thus the sign and magnitude of the effects of skilled immigration on firm-level productivity is still an open question.

The rest of the paper is organized as follows. [Section 2](#) presents a theoretical framework that guides the empirical analysis. [Section 3](#) presents the data sources and describes the matching algorithm. [Section 4](#) presents summary statistics. [Section 5](#) discusses our empirical strategy. [Section 6](#) presents our main estimates and sensitivity analysis. [Section 7](#) discusses the theoretical implications of our empirical findings, and [Section 8](#) concludes.

2 Theoretical framework

Consider a standard monopolistic competition setup where producers vary in their level of productivity ([Jovanovic \(1982\)](#), [Hopenhayn \(1992\)](#), or [Melitz \(2003\)](#), among many others) and production is subject to fixed and variable costs. The goal of the model is to derive predictions for the effects of adding skilled workers on all relevant firm outcomes. For now we shall assume that skilled workers and, in particular, foreign skilled workers hired through H-1B visas lead to higher firm productivity and derive a number of implications. Later on we shall provide empirical evidence based on our data supporting the connection between skilled labor and firm productivity.

2.1 The economy

Consider a standard monopolistic competition setup where and producers vary in their level of productivity. Specifically, we assume there is a representative consumer with CES preferences. The utility maximization problem is given by

$$\begin{aligned} \max \left[\int_J y(k)^{\frac{\sigma-1}{\sigma}} dk \right]^{\frac{\sigma}{\sigma-1}} \\ s.t. \\ \int_J p(k) y(k) dk = X, \end{aligned}$$

where J is the set of goods available for consumption, X is income, and σ is the elasticity of substitution. As is well known, the solution to this problem gives rise to the familiar demand functions where the spending share on each good k is a function of its relative price:

$$p(k)y(k) = \left(\frac{p(k)}{P} \right)^{1-\sigma} X, \quad (1)$$

where P is the price level in the country:

$$P = \left[\int_J p(k)^{1-\sigma} dk \right]^{\frac{1}{1-\sigma}}. \quad (2)$$

Let us now turn to profit maximization. Each firm wishing to produce is required to pay a fixed cost f (units of labor). There is one factor of production (“labor”) that firms hire at wage w . Since each firm produces its own unique variety, it faces a downward-sloping demand curve. In addition, by definition, in the monopolistic-competition model, each producer is sufficiently small that it takes P as given. As is well known, profit maximization implies that the price is a constant markup over the marginal cost, given by $a(k)w$, which differs across producers.

The profit maximization problem is as follows. The firm with marginal cost $a(k)$ solves:

$$\max [p(k)y(k) - a(k)wy(k) - fwI_{\{y(k)>0\}}], \quad (3)$$

where demand for its variety is given by:

$$y(k) = \left(\frac{p(k)}{P} \right)^{1-\sigma} \frac{X}{p(k)}. \quad (4)$$

2.2 Key predictions

It is straightforward to show that, provided profits are non-negative, the optimal price, quantity, sales, employment and profits,¹⁰ are given by

$$p[a] = \frac{\sigma}{\sigma-1} aw = \bar{m}aw \quad (5)$$

$$y[a] = \frac{X}{P^{1-\sigma}} (\bar{m}aw)^{-\sigma} \quad (6)$$

$$py[a] = \frac{X}{P^{1-\sigma}} (\bar{m}aw)^{1-\sigma} \quad (7)$$

$$\ell[a] = \frac{X}{P^{1-\sigma}} (\bar{m}w)^{-\sigma} a^{1-\sigma} \quad (8)$$

$$\pi[a] = \frac{1}{\sigma} \frac{X}{P^{1-\sigma}} (\bar{m}w)^{1-\sigma} a^{1-\sigma} - fw. \quad (9)$$

¹⁰To lighten notation we drop the k subindex.

Thus, more productive firms (low a) charge lower prices, produce more, have higher sales, hire more workers, and obtain higher profits. We also note that a firm can always choose zero output, which delivers zero profits. Thus low productivity (high- a) firms will not operate. Given X , P and w , the productivity threshold is

$$\pi(\hat{a}) = 0. \quad (10)$$

To close the model it is customary to impose a free entry condition.¹¹ Furthermore, if we assume that *productivity* $1/a$ is distributed Pareto, then we can solve for the productivity threshold, the price level, and other equilibrium values analytically. However, we are solely interested in the following comparative static exercise.

2.3 Comparative static: exogenous increase in productivity

Suppose that a firm's productivity $1/a$ increased. How would this affect the firm's size (in terms of sales and employment) and profits, assuming no change in aggregate variables (X, P, w)? By virtue of equations (6) - (9), we have that

$$\Delta \ln py[a] = (\sigma - 1) \Delta \ln a^{-1} \quad (11)$$

$$\Delta \ln \ell[a] = (\sigma - 1) \Delta \ln a^{-1} \quad (12)$$

$$\Delta \ln \pi[a] = \frac{\sigma - 1}{\sigma} \Delta \ln a^{-1}. \quad (13)$$

Thus, provided that $\sigma > 1$, which is the relevant range for this parameter, an increase in firm productivity, whatever its source, will lead to increases in firm's sales (along with employment and output) and profits. These expressions will be the basis for our empirical specifications.

2.4 Endogenous productivity

All the previous predictions follow regardless of the nature of the increase in firm productivity. However, in the context of our paper we focus on skilled labor as a potential source of increases in firm productivity. There are several channels through which the skills of a firm's workforce can lead to higher productivity (Moretti (2004)). One such channel may be that skilled workers allow a firm to innovate, better tailoring its products to the changing needs of the market, or developing more efficient production processes for the firm's products.

¹¹In the formulation of Melitz (2003), there is an ex-ante stage where potential firms pay a fixed fee, also denominated in units of labor, in order to have the right to obtain a productivity draw. Entry into the lottery occurs up to the point where expected profits are driven to zero.

For now we simply postulate that H-1B visas allow a firm to increase its stock of human capital.¹² And, in turn, this allows that particular firm to increase its productivity, which, as we have shown earlier, should lead to an expansion in firm size and an increase in profits. Later on we shall use our data to shed light on the mechanisms linking the availability of skilled workers in a firm and its productivity level.

3 Data

In this paper we use administrative records on labor condition applications (LCAs), the first step towards H-1B visas for foreign-born workers. We aggregate these data by employer and link them to firm-level variables obtained from the Compustat Industrial Annual data set, which covers publicly traded firms in the United States. To the best of our knowledge, this is the first attempt to link the LCA records and data for all public firms in the U.S. Our analysis builds on [Kerr and Lincoln \(2010\)](#) who also linked LCAs and firm outcomes (patenting activity). However, their sample was small (only 77 firms) and selected to include only firms with high patenting activity. Thus one of the contributions of our paper is to provide a much larger dataset (the universe of LCAs and of publicly traded firms in the US) and to expand the number of firm outcomes that we explore to productivity, sales, employment, profits, and R&D expenditures. The following subsections introduce our two data sources and describe how we match them. We also present summary statistics of the data.

3.1 The LCAs data

H-1B visas are used to employ a foreign worker in a “specialty occupation” which, in general, requires the applicant to hold at least a bachelor’s degree. The H-1B visa is typically a 3-year visa, which can be renewed for a second three-year term. An employer who intends to hire a foreign worker under the H-1B program must first submit a labor condition application (LCA) to the U.S. Department of Labor. Each LCA has a case number, the employer’s name and address, information on whether the application was certified (i.e. processed) or denied, the occupation code of and the wage offered to the immigrant worker, the prevailing wage and also an indicator of the source of the prevailing wage data. Importantly, the employer must document that the prospective H-1B visa holder will receive a wage that is no lower than the prevailing wage for the same position in the relevant geographic area or the wage

¹²It is often argued that there is severe shortage of graduates with specific skills or with degrees in some particular occupations, e.g. STEM.

actually paid by the employer to individuals with similar workplace characteristics. The employer must also attest that the working conditions of U.S. workers similarly employed will not be adversely affected.

LCA records are available online in the website of the Foreign Labor Certification Data Center.¹³ Both first time applications as well as renewals require an application. However, the data provided by the Foreign Labor Certification Data Center does not distinguish between the two types. LCAs can be filed both by fax and electronically (e-file). Our data set includes 2001 fax filings and, for 2006, both e-file and fax data. The e-file option was available starting in 2002 and, by 2004, 90 percent of the LCAs were filed under the e-filing system. Our data for 2001 and 2006 covers 100 percent of all the LCAs submitted.

Once the LCA has been certified by the U.S. Department of Labor, which happens in the vast majority of cases, the employer files a petition to the United States Citizenship and Immigration Services (USCIS).¹⁴ It is at this point in the process that the H-1B quota applies, i.e. the total number of approved petitions by USCIS is no greater than the cap (except for visas in exempt categories). Finally, if the USCIS approves the petition, a visa will be issued by the State Department if the individual lives abroad. If instead the individual is already living in the United States, the USCIS will convert the visa status to H-1B.

Not every LCA results in a H-1B visa. There is a cap on the number of H-1B visas that are granted by the U.S. every year (see [Figure 1](#) which plots the annual cap on H-1B visas between 2000 and 2008). Some institutions, such as universities, are exempt from the H-1B cap since 2001. In addition some countries, such as Chile, have their own quotas which do not fall under the cap. The H-1B cap imposed by Congress was 115,000 visas in 2000, 195,000 in 2001 through 2003, and 65,000 in 2004.¹⁵ In 2006 an additional 20,000 visas were made available through the H-1B program for those individuals who had a Master's degree or higher from a U.S. institution, effectively raising the cap from 65,000 to 85,000. The quotas were not binding from 2000 until 2003 (included) and became binding thereafter.

¹³<http://www.flcdatacenter.com/CaseH1B.aspx>

¹⁴The certification process by the Department of Labor only checks for obvious errors, which are occasionally found. In 2006, out of 385,235 applications, only 8,088 were denied (i.e., 2.1 percent).

¹⁵In year 2001 the number of certified LCAs was higher than the H-1B cap. Three reasons explain this apparently inconsistent pattern in the data. First, not all H-1B visas are subject to a cap. As mentioned above since 2001 some institutions, such as universities, do not fall under the H-1B cap although they are still required to apply for a LCA. Second, H-1B renewals do not fall under the cap but still require a LCA. Lastly, even in non-binding years, a certified LCA may not translate into a H-1B visa due to attrition. For instance, some foreign workers may choose not to seek employment with a particular employer between the time of the application and the H-1B processing.

As we explain in detail later, the sharp reduction that took place in 2004 will be at the core of our identification strategy.

A very time-consuming part of this project has been creating the firm-level LCAs data using the raw LCA records. As mentioned previously, each individual LCA includes the employer name and address. The main problem we encountered is that the name of the employer is not consistent across LCAs in different years. An example of this is that the name “AGFIRST FARM CREDIT BANK” appears in an application in one year, and “AGFIRST FCB” is the name reported in an application for another year. To recognize whether an employer was the same across LCAs, we first located the employers from the same city whose name was very similar and assumed they were the same employer. Next, we focused on employers from different cities whose names were similar. We used a word-match software, which assigns a measure of how close two strings of characters match. If more than 75 percent of the characters matched between two applications then those applications were considered to belong to the same firm. In addition we inspected the top applicants – i.e. firms characterized by the largest numbers of LCAs – and manually checked that they had been correctly assigned. Once we determined which employer names corresponded to the same firm, we assigned a unique identifier to that firm. To check the accuracy of our data, we also compared our totals for LCAs for the top applicants with data from other sources and found consistent results.

Table 1 presents the list of the top 5 users of LCAs for each year between 2001 and 2006. The rankings are fairly stable across these six years, with Microsoft, Oracle, IBM, Infosys, Patni Computer Systems, and Satyam Computer Services topping the rankings. But we also note how the number of LCAs filed by the top firms increased importantly over the 5-year period, ranging from 564-1,736 to 1,262-4,406. In 2006 Microsoft was the firm that filed the most LCAs (4,406).¹⁶ At the other end of the spectrum, almost 80 percent of the firms did not file any LCAs in year 2001 and only 5.3 percent filed exactly one LCA. It is worth noting that some of these firms may be small start-ups that, if successful, may grow to be large and successful firms. These firms may be particularly constrained in hiring highly skilled workers in some occupations. We view the inclusion of these firms as an important strength of our dataset, and an important distinction with the data used by [Kerr and Lincoln \(2010\)](#), whose sample included only firms that were chosen for their high patenting activity and their large number of LCAs.

It is also interesting to examine which sectors have the highest average LCAs. As we

¹⁶In that year 3 out of the top 5 were firms incorporated in India, compared to only one in 2001.

can see in [Table 9](#), based on the 2001 data, the sectors with the highest LCAs are Manufacturing (33), Construction (23), Media and Telecommunications (51), Transportation (49) and Professional, scientific and technical services (54). As a caveat, we note that some employers with large numbers of LCAs are not in the (Compustat) data since these firms are not publicly-traded firms, such as Ernst and Young or Deloitte consulting, both of which are partnerships, or all universities.

3.2 Compustat and the matching process

Compustat is a dataset containing the balance sheet information of all publicly traded firms in the United States. We restricted the sample to firms appearing both in years 2001 and 2006. In addition we dropped firms with missing or zero values for some of the key variables (employees, sales and capital). We also dropped firms with negative sales or negative R&D expenses. We then proceeded to match firms in the LCAs dataset with the firms in our Compustat sample. At this stage of the process, we carried out further manual matching, especially for top applicants, to make sure that they were not missed in Compustat. If a firm appeared in Compustat but not in the LCAs dataset in a given year, we assumed that it did not file LCAs in that year and assigned a value of zero to the LCA variable for that observation.¹⁷

3.3 Firm-level variables

The resulting matched dataset contains almost 4,000 firms, which we shall describe in detail in the next section. Here we introduce the main variables and definitions. We employ two measures of firm size: sales and the total number of employees.¹⁸ Besides firm size we also focus on sales per employee, which we use as a simple measure of average labor productivity in the firm, gross profits, and R&D expenditures as important outcomes of interest. Of course, we also know the number of LCAs filed by each firm in our two years of interest, 2001 and 2006. We focus on the log changes for all the outcomes of interest. Essentially, the midpoint in the interval between years 2001 and 2006 is 2004, the year where the national cap on H-1B visas was sharply reduced.

¹⁷Recall that our LCAs data contains the universe of applications. Thus, under the assumption that our matching has been successful, it is entirely correct to make the assumption that an unmatched firm-year observation across the two datasets corresponds to a firm that filed zero applications in that year. In practice, our matching algorithm was not perfect but we believe our assumption simply introduces some noise in the estimation, but no biases.

¹⁸This variable refers to the total number of employees (usually at year-end) corresponding to consolidated subsidiaries, including both domestic and foreign. Unfortunately, we cannot disaggregate a firm's workforce by skill level of the employees.

As we shall see, R&D expenditures are missing for almost half of the firms. This is problematic because expenditures in R&D are our proxy for innovation. This data challenge has been discussed by several authors. [Bound et al. \(1984\)](#) and [Hirschey et al. \(2012\)](#) recommend setting to zero all missing values of the variable R&D expenditures in Compustat.¹⁹ Following this guideline we create a version of the R&D expenditures variable that makes an imputation similar to the one suggested by these studies.²⁰ However, we are not certain that imputing a (constant) zero value in the two periods to all firms with missing values will not affect our results. Thus, to be safe, we also create a subsample restricted to firm-year observations with positive values for R&D expenses in the raw data.

We also build firm-level measures of Total Factor Productivity (TFP). To do this we exploit the longitudinal dimension of our data with over 40,000 firm-year observations.²¹ Building firm-level estimates of TFP is a two-step process. In the first step we need to estimate the coefficients of the production function, assumed to be Cobb-Douglas but with factor shares that are allowed to vary by industry. The second step then uses these estimates to build a residual TFP term. In our analysis we present a naive estimation of the firm-level production functions where we assume exogenous regressors and estimate the production function by OLS. In the second case we address endogeneity concerns following [Levinsohn and Petrin \(2003\)](#). However, their method requires data on intermediate inputs (materials), which are missing for almost half of the sample, undermining the performance of their estimator in the context of our application. We provide more details in the Appendix.

4 Summary statistics

At the end of our matching procedure, we are left with almost four thousand firms with complete data on sales, factor usage, R&D expenditures, and LCAs for both 2001 and 2006.²² About 20 percent of the firms in our sample filed at least one LCA in years 2001 and 2006.

Table 2 presents summary statistics for our matched LCA-Compustat dataset, which contains 3,945 firms. The table is divided in three panels which focus on, respectively,

¹⁹See the discussion in [Bound et al. \(1984\)](#), page 25, and the table in the previous page.

²⁰We replace missing values by a small positive number (10 dollars) and we also add the same amount to all other firms in order not to distort the distribution. As a result, the firms with missing values in both years will thus display a zero log change in R&D expenditures.

²¹In our main analysis we only use data for years 2001 and 2006 and restrict the sample to firms that are present in the dataset in both years. To estimate TFP we use also data for the years in between.

²²The number of public firms in the U.S. between 2001 and 2006 was fairly constant, at about 6,000 firms. Hence, our sample restrictions leave us with about 66 percent of the universe of publicly traded firms.

firm outcomes in levels for years 2001 and 2006, 2001-2006 changes in firm outcomes, and the distribution of LCAs. We note that sales, sales per employee, capital expenditures and profits are higher on average in 2006 than in 2001, respectively they increased by 46, 30, 20, and 50 log points. These changes partly reflect the fact that these variables are expressed in current dollars. However, note that also average employment and TFP increased, respectively, by 17 and 21 log points, and these variables should be immune to price inflation.²³

Let us now turn to *R&D* expenditures, our proxy for the innovation activity of a firm. In year 2001 the average expenditure was around 112 thousand dollars. However, it is worth noting that only about half of the sample reported this variable to Compustat (1,923 firms in 2001 and 1,962 in 2006). These were essentially the same firms in the two years (1,804 firms reported R&D expenditures in both years). As discussed earlier, we report two variations of the change in the log of R&D expenditures: the raw variable (available for less than half of the sample) and an imputed one (where we assigned a small positive value to firm-year observations with a missing value). Between 2001 and 2006 the average R&D expenditure increased by about 20 log points.²⁴ However, the increase was not uniform across firms and many firms reduced R&D expenditures during this period.

We now turn to LCAs. In 2001 the average firm in our sample submitted slightly over 5 applications whereas in 2006 the average had risen to almost 10. We now describe the distribution of firms on the basis of 2001 applications, which will be the basis for our measure of dependence on H-1B workers. According to our data almost 80 percent of the firms in our sample did not file any applications in 2001. Hence, even when we restrict to publicly traded firms, only a minority of firms attempt to hire H-1B workers. To continue our exploration of the distribution of LCAs, among the firms with at least one application in 2001 we build quintiles but disaggregate the top two quintiles into four deciles to obtain higher resolution at the top of the distribution. These cutoffs will later be used in our non-parametric analysis. The resulting breakdown is as follows: 5.3 percent of all firms in our sample filed one application, 5 percent filed 2-3 applications, 3.3 percent filed 4-7

²³Obviously, these increases are not due to compositional changes since we are considering the same exact set of firms in both years 2001 and 2006.

²⁴Later, when we refer to the R&D subsample, we mean all firms with positive expenditure in the original data. As a result of our imputation all firms with missing data on R&D expenses in both 2001 and 2006 will have a value of zero (no change) for the change in the log of imputed R&D expenditures. That is, we assume that these firms had the same level of R&D in 2001 and 2006. Note that this assumption allows these firms to have any arbitrary (constant) *level* of R&D expenditures and that these may differ across firms. Reassuringly, the average log change is very similar in the raw and imputed variables, respectively, 19 and 20 log points.

applications, 1.6 percent filed 8-10 applications, 1.3 percent filed 11-18 applications, 1.8 percent filed 19-59 applications, and 2.1 percent filed 60 or more applications.²⁵

We also created a subsample for firms that conducted R&D in both 2001 and 2006, which contains about one quarter of the whole sample (969 firms).²⁶ **Table 3** presents the summary statistics. These firms are larger, as measured by sales and employment (by 93 and 53 percent, respectively), than the firms in the whole sample and their mean R&D expenditures are twice as large than the average firm in the whole sample. For this subsample, between 2001 and 2006, sales, employment, and profits increased, respectively, by 49, 13, and 49 log points, which is very similar to the increases reported in the previous table for the whole sample. Turning now to TFP and R&D expenditures, the average increases in the R&D subsample were 31 and 28 log points, respectively, which are noticeably larger than the corresponding increases for the whole sample (21 and 19 log points). Turning now to LCAs, we note that the average number of LCAs for the firms in the R&D subsample was 13 in 2001 and 23 in 2006, more than twice the number of applications in the sample of all firms. In the R&D subsample only 61 percent of firms did not file any LCAs in 2001 (compared to 80 percent in the whole sample). At the other end of the spectrum, 5.5 percent filed 60 applications or more in 2001, compared to only 2.1 percent in the whole sample. Thus the firms conducting R&D file many more LCAs, which may partly be due to their larger size. At any rate, since these firms use H-1B workers more intensively (as proxied by the number of LCAs), we expect them to be more affected by changes in the overall cap for H-1B visas.

5 Empirical strategy

5.1 Specifications

We are interested in estimating the impact of the number of H-1B visa workers in a firm on several outcomes pertaining to that firm. The theory is silent on the functional form so our main estimates will be based on a flexible, non-parametric specification. At this point it helps to start with a more restrictive, but simpler, linear model.

Ideally, we would like to estimate a specification of the following type:

$$\ln y_{ijt} = \alpha_i + \gamma_t + \beta H1B_{ijt} + \delta_j \times t + \epsilon_{ijt}, \quad (14)$$

²⁵Two firms (Microsoft and Satyam Computer Services) filed over 1,000 LCAs in year 2001. In 2006 these two firms filed over 4,000 applications each.

²⁶Specifically, a firm is included in this subsample if its expenditures in research and development in both years were at least \$5,000.

where $\ln y_{ijt}$ represents the outcome of firm i , in industry j , in year t , and α_i and γ_t are, respectively, firm and year fixed effects. In addition, $H1B_{ijt}$ is the actual number of H-1B visa workers in firm i in industry j , at time t . Finally, $\delta_j \times t$ captures the time trend specific to industry j . The key coefficient of interest is β , which we interpret as the effect of adding one H-1B worker. To a first-order approximation, coefficient β is also informative more generally about the effects of skilled labor (native or immigrant) on the firm's outcomes.

There are two serious challenges in implementing an estimation of this model. First, we lack data on the actual number of H-1B visa workers in each firm, which often differs from the number of LCAs for the reasons noted earlier. Thus the estimation of this model is not feasible. Second, even if we had those data, the specification above would lead to a biased estimate of β because H-1B visas are not randomly distributed across firms.

For these reasons the model that we estimate is the following:²⁷

$$\ln y_{ijt} = \alpha_i + \gamma_t + \beta_1 H1B_t \times LCA2001_{ij} + \delta_j \times t + \epsilon_{ijt}. \quad (15)$$

The crucial difference between the two equations is that in equation (15) we do not need data on the number of H-1B visas awarded to each firm in any given year. Our identification strategy exploits the sharp decrease in the annual cap in H-1B visas in year 2004. To allow for implementation delays we pick our pre and post dates to be 2001 and 2006. In 2001 the annual quota was 195,000 H-1Bs and in 2006 it was 85,000, which substantially altered the policy environment by making it much harder to obtain an H-1B visa. This policy change provides us with time-variation that is arguably exogenous from the point of view of an individual firm. We combine the change in policy with the approach proposed by [Kerr and Lincoln \(2010\)](#), which postulates that changes in the annual cap should have larger effects on firms that rely to a greater extent on H-1B visas. The *key distinction* between the approach by [Kerr and Lincoln \(2010\)](#) and ours is that their identification is based on year-to-year changes in policy, whereas our time-variation arises from a single policy change that took place in 2004.²⁸

Another difference with the analysis in [Kerr and Lincoln \(2010\)](#) is that they built the measure of H-1B dependency on the basis of LCAs *divided by employment in the firm*. In contrast, we do not adopt this normalization. While normalizing by firm size is reasonable

²⁷This specification is similar to the one used by [Kerr and Lincoln \(2010\)](#) in their firm-level analysis, although our measure of dependency on H-1B visas is slightly different, as we discuss below.

²⁸Furthermore, the measure of policy chosen by [Kerr and Lincoln \(2010\)](#) is the national stock of H-1B visas (as estimated by [Lowell \(2000\)](#)) whereas we rely on the annual cap on H-1B visas. In practice this makes a difference because Lowell's estimate of the stock of H-1B visas was practically the same in years 2001 and 2006, which would not reflect the change in the policy environment illustrated in [Figure 1](#).

from an empirical point of view, we use the level of LCAs in year 2001 as our measure of dependency, which is more consistent with our theoretical model. In addition we note that using aggregate employment at the firm level is a rough way to normalize and imposes a number of implicit assumptions. For instance, it does not take into account that employment differs systematically across sectors for technological and other reasons, or that the share of the workforce consisting of skilled workers also differs widely across firms and industries. Unfortunately, this information is not available in Compustat.²⁹

Thus our goal is to investigate whether the exogenous change in the national H-1B cap between 2001 and 2006 affected firms differently, according to their pre-existing dependency on H-1B visas. In particular, we measure the degree of dependency by using the number of LCAs filed by each firm in year 2001.³⁰ We expect firms that were more dependent on H-1B visas in 2001 to be more adversely affected by the reduction in the cap than less dependent firms. In other words, more H-1B-dependent firms should exhibit worse outcomes between 2001 and 2006 – namely lower growth in sales, productivity, employment, and profits – than less dependent firms.

Our setup is potentially subject to a problem of reverse causality because firms that grow faster may want to hire more skilled foreign workers. However, note that our measure of dependency is based on LCAs in year 2001 alone. Thus, to the extent that firms are not able to anticipate their growth rate for the 5-year ahead period, reverse causality is not a major concern. Furthermore we will show in the next section that the pre-treatment trends in the main outcomes of interest (for the period 2000-2003) are not correlated with the number of LCAs filed by firms in 2001.

By differencing equation (15) between years 2001 and 2006, we obtain:

$$\Delta \ln y_{ij} = \alpha + \beta_1 \Delta H1B \times LCA2001_{ij} + \delta_j + \varepsilon_{ij}, \quad (16)$$

where the constant captures the aggregate time trend and the industry dummy variables allow for industry-specific trends.³¹ We consider a number of firm outcomes y_{ij} : sales per employee, sales, employment, and profits. We expect β_1 to be positive, i.e. *relaxing* the

²⁹We also note that the sample of firms in [Kerr and Lincoln \(2010\)](#) is effectively restricted to firms that rely heavily on H-1B visas and that account for a substantial number of patents on a regular basis (footnote 26, page 501). For a similar subsample in our data, we show that whether we normalize or not the measure of dependency does not alter the results.

³⁰Since in 2001 the annual quota was not binding, the number of LCAs in that year is a good approximation for the unconstrained number of H-1B visas obtained by the firm in that year.

³¹To ease notation we have dropped the time subindices, which are now unnecessary. The data (in changes) is now a cross-section.

national cap on H-1B visas should benefit more those firms that are more dependent on H-1B visas. Conversely, a reduction in the cap (as in year 2004) should reduce the growth of the more dependent firms.

In addition, we also estimate the following (closely related) specifications:

$$\Delta \ln y_{ij} = \alpha + \beta_2 LCA2001_{ij} + \delta_j + \varepsilon_{ij} \quad (17)$$

$$\Delta \ln y_{ij} = \alpha + \beta_3 \Delta RelH1B \times LCA2001_{ij} + \delta_j + \varepsilon_{ij}. \quad (18)$$

In specification (17), we have dropped the size of the change in the annual cap. Given that our data (in changes) is now a cross-section, this solely entails a re-scaling of the main coefficient of interest. Importantly, the sign of the coefficient changes since the annual cap was reduced between 2001 and 2006 (by 110,000 visas).³² In specification (18), we build an alternative measure of the change in the policy environment meant to account for the large increase in the overall number of applications, which coupled with the sharp reduction in the annual cap made it even harder to obtain a H-1B visa in a non-exempt sector. Thus, in place of the change in the H-1B cap we now use $\Delta RelH1B$, where $RelH1B$ is defined as the annual H-1B cap divided by the overall number of non-exempt LCAs in the corresponding year.³³

5.2 The parallel trends assumption

Our empirical specification can be interpreted as a difference-in-difference estimator, where the treatment and control groups are categories of firms with different levels of LCA dependency and the treatment is the sharp reduction in the annual H1-B cap around year 2004. This estimation strategy requires that the pre-treatment trends in the outcomes of interest be the same for the treatment and control groups (parallel trends). As we argue next, our measure of H-1B dependency is not correlated with pre-treatment *changes* in the outcomes of interest, consistent with the parallel trends assumption.

We begin by relating the *levels* of sales, employment, profits and R&D expenditures with our measure of dependency on H-1B visas. The lack of correlation between the outcome variables in *levels* and our measure of dependency is not a necessary condition for the difference-in-difference estimator but a useful starting point. The top panel in Table 4 presents the estimates of the following relationship:

$$\ln y_{ij,2002} = \alpha + \beta LCA2001_{ij} + \delta_j + \epsilon_{ij}. \quad (19)$$

³²Hence, coefficient $\beta_2 = -\beta_1 \times 110,000$.

³³Aggregate data on LCAs by sector are publicly available.

Clearly, firms that filed a higher number of LCAs in year 2001 tend to be larger in size (by sales and employment), and to have higher profits and spend more in R&D activity. This is not surprising since firms with large sales, employment, profits and R&D will tend to file a large number of applications for H-1B visas.

We next consider a specification relating our measure of dependency to the *change* in outcomes (in logs) at the firm level, which nets out these level differences between the treatment and control groups. This is now in line with our linear specification. We estimate this model using data for the pre-treatment period, that is, using changes between years 2000 and 2003. The results are presented in the bottom panel of the table. We now find no evidence of a systematic relationship between our measure of dependency on H-1B visas and changes in firm-level outcomes. Thus our results are consistent with the parallel trends assumption.

6 Estimates

6.1 Linear relationship

Let us begin with the estimates of the linear model, where we also experiment with different versions of the main explanatory variable, as in equations (16) through (18). In this model the dependent variables are, in turn, the 2001-2006 change in the logs of sales per employee, sales, employment, and profits. We estimate these models on three subsamples: all firms, firms with at least one LCA filed in 2001, and the R&D subsample.

Table 5 presents the estimates. The top panel displays the estimates corresponding to equation (16), the specification that most closely resembles the one used by Kerr and Lincoln (2010). The explanatory variable here is the change in the national cap on H-1B visas between years 2001 and 2006 interacted with the number of applications filed by the firm in year 2001, our firm-level measure of dependency on H-1B visas. Recall also that the cap on H-1B visas fell between 2001 and 2006 so the change in the cap is a negative number. Thus if firms' sales also fell, the estimated coefficient of the interaction between the change in the cap and the measure of dependency on H-1B visas should be *positive*. In columns 1-4 we present estimates corresponding to our key four dependent variables for the whole sample (3,943 firms).³⁴ As hypothesized, the positive coefficients in columns 1-4 suggest that increases in the national cap on H-1B visas are associated with larger increases in the

³⁴We exclude two firms that filed over 1,000 applications in year 2001 (Microsoft and Satyam). Exploratory analysis shows that those two firms are outliers since they behave very differently than the rest of firms, biasing the estimates of the linear model.

outcomes that we examine. However, we also note that we can only reject the zero null for sales (marginally) and sales per employee in columns 1-2. Qualitatively, the results are very similar in the middle and bottom panels.³⁵ Consistent with the earlier findings, in columns 1-4 we find significant coefficients for average labor productivity (sales per employee) and, marginally, also for sales.

In columns 5-8 we display the estimates for the same models but now we restrict the sample to firms that filed at least one LCA in year 2001 (801 firms). The pattern of estimates is fairly similar to the one found for the whole sample, although the point estimates tend to be larger in absolute value. Hence, the previous findings appear robust to identifying the effects solely off of the variation across firms that filed LCAs in year 2001, namely, along the intensive margin.

Finally, the estimates in columns 9-12 are based on the subsample of firms conducting R&D activities during both years (of at least \$5,000). The pattern of estimates is now more striking. Focusing on the top panel, we now find positive and significant estimates for the four outcomes: sales per employee, sales, employment, and profits. Moreover the coefficients for sales, employment and profits are much higher than for the whole sample. These estimates strongly suggest that one of the channels through which H-1B workers help improve firms' outcomes may be by increasing innovation activity, a hypothesis that we investigate further in the sections to come.

6.2 Flexible specification

We believe that the specifications estimated above might be too restrictive since they impose a linear relationship between growth in firm outcomes and the number of H-1B workers in the firm. This seems particularly problematic in the context of the recent literature emphasizing firm-level heterogeneity and a non-linear relationship between productivity and outcomes such as sales, employment, profits or exports and firm productivity (Melitz (2003)). Thus we adopt a non-linear approach, as in the analysis in Kerr and Lincoln (2010) at the state and city levels.³⁶ Specifically, we divide firms into groups according to their LCAs dependence in 2001. We build *eight* categories of firms: those with zero applications

³⁵Naturally, the three specifications are essentially equivalent. Recall that our estimation is based on a cross-section (of changes). Thus multiplying a regressor by a constant only re-scales the estimated coefficients. In the middle panel we do not interact the measure of dependency by the change in the cap for H-1B visas. Thus we now expect negative coefficients. The reason is that we hypothesize that more dependent firms will experience *worse* outcomes.

³⁶Their firm-level analysis was carried out using only the linear model due to the limited number of firms in their sample (77 firms).

in 2001 plus five quintiles for the distribution of firms conditional on at least one LCA in 2001, where the top two quintiles are subdivided into two deciles each.³⁷ In particular, we estimate the following specification:

$$\Delta \ln y_{ij} = \alpha_j + \Delta H1B \times \sum_p \beta_p \times D\{a_p \leq LCA2001_i \leq b_p\} + \varepsilon_{ij}, \quad (20)$$

where a_p and b_p refer to the lower and upper bounds of each of the 7 brackets of LCA use in 2001. We expect *positive* coefficients for firms that rely on H-1B visas and the coefficients should be higher for the groups that exhibit a larger dependency on H-1B visas.

Table 6 presents the estimates. Columns 1-4 present the results for the whole sample for our four main outcomes (sales per employee, sales, employment and profits). The omitted category are firms with zero applications in 2001. Let us begin by focusing on our measure of average labor productivity (sales per employee) and on overall sales (columns 1 and 2). For the first five brackets of firms (fewer than 18 LCAs in 2001) we do not find a consistent pattern for any of the outcomes. However, we find large, positive, and significant coefficients for the top bracket (60 or more applications in year 2001). Turning to columns 3 (employment) and 4 (profits), we observe a similar pattern, with larger coefficients for the top two brackets than for the brackets with lower use of LCA applications, although we can only reject the zero null for the specification on firms' profits.

Next, we turn to columns 5-8, which provide estimates for the same outcomes but restricting the estimation to the R&D subsample. The pattern that emerges is similar, with significant coefficients for the category of firms containing the top filers of LCAs in 2001. There are, however, two points worth noting. First, the estimated coefficients are substantially larger than for the whole sample of firms, over 50 percent larger for sales, employment and profits. Secondly, despite the much smaller sample size, the estimates are more significant, including the one for employment that is now statistically significant. We interpret these findings as further evidence that one of the channels through which skilled immigration improves firms' outcomes is through innovation, as measured by R&D expenses. In order to assess this interpretation further columns 9-12 report estimates on the subsample of firms that either did *not* report R&D expenditures or had very small amounts (below \$5,000).³⁸ In this case the pattern observed in columns 1-8 vanishes, providing additional support for our interpretation.³⁹ In a nutshell, these estimates reveal that the

³⁷By construction the quintiles have similar size in terms of the number of firms.

³⁸Our guess is that the firms that did not report expenditures in R&D were firms that did not conduct any meaningful R&D activity in these years.

³⁹Specifically, despite the large sample, we cannot reject the null hypothesis of a zero effect for sales,

effects detected for the whole sample are, in fact, driven by the subsample of firms that carry out R&D.

In light of our results in this section, we now learn that the positive association between H-1B workers and firm outcomes uncovered using the linear specification is driven fundamentally by the heavy users of H-1B visas. An important implication is that if the cap on H-1B visas were to be relaxed, only the heaviest users of H-1B visas would benefit in terms of productivity, firm size, and profits. Specifically, our results suggest that the threshold can be found at around 18 LCAs per year.⁴⁰

Our finding that the results are driven fundamentally by the subsample of firms that conduct R&D activity resonates with the findings in [Kerr and Lincoln \(2010\)](#), who argue that H-1B workers lead to increases in innovation, as measured by patenting activity. Furthermore our results show that the increase in innovation is accompanied by effects on other relevant firm outcomes, such as average labor productivity, firm size, and profits. It is also interesting to compare our finding with the results in [Paserman \(2013\)](#). Using data for Israeli firms, this author found that immigration was associated with reductions in output per worker in low-tech industries, but with increases in this variable in high-tech industries. Our results for sales follow this same pattern. However, we find similar effects on sales per employee for firms in the subsamples that do and do not carry out R&D activities.

6.3 The innovation channel: TFP and R&D

In this section we ask two questions. First, we examine whether the earlier finding of an effect on average labor productivity (sales per employee) is driven by an increase in Total Factor Productivity. Second, we explore further the connection between the availability of H-1B workers and innovation activity, the focus of [Hunt and Gauthier-Loiselle \(2010\)](#) and [Kerr and Lincoln \(2010\)](#). In particular, we are interested in knowing whether the R&D activity of firms is constrained by the limited availability of skilled workers.

6.4 Total Factor Productivity

We begin with the question of whether H-1B workers increase TFP in their host firms.⁴¹ As discussed earlier, we need to build firm-level measures of TFP. We do so in two ways.

employment and profits. We do find a positive effect on sales per employee. However, this effect is based on a non-significant positive effect for sales and a non-significant negative effect for employment, which combined give rise to a significant effect on sales per employee.

⁴⁰This finding is well aligned with the statements by Bill Gates and other senior management of large firms advocating for raising the annual cap on H-1B visas.

⁴¹It is also possible that there may be effects that are external to the firm ([Moretti \(2004\)](#)).

The first, and simplest approach, consists in estimating the production function for each industry by OLS and then constructing TFP for each firm as a residual (denoted by TFP^1). The second approach (Levinsohn and Petrin (2003)) relaxes the assumption of exogenous regressors in the estimation of the production function (denoted by TFP^2). However, it requires data on intermediate inputs, which is missing for the majority of firms in our sample.⁴²

Table 7 presents the estimates. Columns 1-4 are estimated on the sample of all firms. For convenience, the first column reproduces earlier estimates for our measure of average labor productivity (sales per employee). As noted earlier, we found that increases in the cap for H-1B visas appear to increase exclusively the labor productivity of heavy users of H-1B visas. In column 2 we replace the dependent variable by the change in the log of TFP^1 , estimated following the first approach described in the previous paragraph. The results are very similar to those presented in the first column, suggesting that H-1B workers help firms increase their TFP. Column 3 now estimates the same model but makes use of TFP^2 , the measure constructed using the Levinsohn-Petrin approach. It is important to note that the number of firms has fallen by 45 percent, from 3,789 to 2,115. In this case we do not find any significant coefficients.

Skipping for now column 4, we turn to columns 5-7 that report the estimates of the same models but for the R&D subsample. The pattern of estimates for sales per employee and TFP is very similar, both in terms of significance and coefficients. As before, we find a significant effect on TFP, but only for the first measure of TFP. For the second measure the sample size falls by almost 50 percent, which may be the reason for the lack of significant results.⁴³ Alternatively, it is also possible that the relationship between immigrant skilled labor in a firm and TFP is more complicated than we have contemplated here. For instance, using German establishment data, Trax et al. (2013) find no evidence of the impact of higher immigrant concentration in a firm on the firm's TFP. However, they do find evidence of a positive productivity effect of ethnic diversity within the firm.⁴⁴

6.5 Research and Development

We now ask whether H-1B workers allow firms to increase their innovation efforts, as measured by R&D expenditures. To do so we use the change in the log of R&D expenditures

⁴²For more details on the construction of these TFP measures, see the Appendix.

⁴³For the TFP built on the basis of the Levinsohn-Petrin estimation the sample size falls to 580.

⁴⁴The findings in Parrotta et al. (2014a) point toward a negative relationship between ethnic diversity and TFP at the firm level in Denmark.

as the dependent variable in columns 4 and 8 in [Table 7](#). Beginning with column 4, which refers to the full sample and required an imputation for about half of the firms in the sample, we do not find evidence of a differential effect on firms with a greater dependency on H-1B visas. However, when we restrict to the R&D subsample (column 8) we do find a positive and significant effect for the category of firms with the highest degree of dependency (60 or more LCAs in 2001). These estimates, thus, suggest that there may be a scarcity of foreign skilled workers that limits the innovation activity of U.S. firms.

6.6 Industry results

Our identification strategy is based on the idea of exploiting the different consequences (for more versus less H-1B dependent firms) of the reduction in the national cap for H-1B visas that took place in year 2004. This argument is more powerful when it is applied to firms within a single industry so that we can rule out confounding factors that vary at the industry level. One way in which we have partially addressed this question is by including industry fixed-effects in all our econometric models. In this section we now re-examine this question by estimating our main models on subsamples of increasingly homogeneous firms in terms of the industry they belong to.

To warm up we first collect some summary statistics in [Table 8](#). This table provides the names of each 2-digit industry along with our estimated capital and labor shares for each industry. For instance, the highest labor shares are found in the Construction industry (23), with 0.92.⁴⁵ [Table 9](#) presents average values for LCAs, sales, employment and R&D expenditures for each industry for years 2001 and 2006. The industries with higher average LCAs (in 2006) are Professional, scientific and technical services (54) and Finance and Insurance (51).⁴⁶ Most relevant for our purposes, we note that the sectors where firms have the highest average R&D expenditures are 32 and 33, both containing firms in Manufacturing. This offers an alternative approach to test our hypothesis that the channel connecting the use of H-1B workers and firm outcomes operates, at least in part, through R&D activity. Specifically, we shall estimate our models on the subsample of firms in Manufacturing. If we find larger effects for the Manufacturing subsample then this will strengthen the link between H-1B workers and innovation efforts. In restricting to subsets of industries, our analysis in this section is based on smaller samples. Hence, we adapt our flexible specifica-

⁴⁵These labor shares have been estimated on the basis of the first approach outlined in the previous section and explained in detail in the Appendix.

⁴⁶In contrast in 2001 the industries with higher average LCAs were Professional, scientific and technical services (54), and Transportation (49). According to our data, the widespread use of H-1B workers in the financial sector is a fairly recent phenomenon.

tion by switching to a coarser partition of firms in terms of H-1B dependency. Instead of 8 groups, we now consider only 4: non-users (zero LCAs in 2001), low users (LCAs 1-18, corresponding to quintiles 1-4), medium users (19-59 LCAs, corresponding to the first half of quintile 5), and heavy users (60 or more LCAs in 2001).

Table 10 presents the results. The top panel replicates the earlier findings using the coarser partition, providing a robustness check. As before, we confirm that heavy users of LCAs, and to some extent also moderate users, are the ones that would benefit the most from an increase in the cap on H-1B visas. This can be seen clearly for productivity, sales, profits, and TFP.⁴⁷ The second panel presents estimates for the R&D subsample. A clearer pattern emerges, with substantially larger coefficients. Moreover, we now find statistically significant effects on employment and R&D expenditures, which provides evidence in favor of the innovation channel as an explanation for the improved outcomes in terms of sales, employment and profits. So far these two panels simply provide a robustness check on our earlier results by employing a less demanding specification (four brackets only).

The bottom panel presents estimates for firms in the manufacturing sector (NAICS 31-33), which is a large category containing industrial equipment, computers, semiconductors, transportation, and so on. Importantly, many of these sub-sectors invest heavily in R&D. We now find large and significant effects for the middle and top brackets (19 LCAs or above) for firm size (measured both by sales and employment) and for profits. In fact, the coefficients are very similar for sales and employment so, not surprisingly, we now do not find a significant effect on sales per employee (column 1). Again we now find evidence suggesting that H-1B workers allow heavy users of H-1B visas to increase the scale of their R&D activities.

It is also interesting to examine the magnitudes implied by our estimates. Consider, for instance, the estimated coefficient for firms in the top bracket (i.e. 60 or more applications in 2001) for the profits outcome (column 4 in **Table 10**). This coefficient is 1.38 for the whole sample. Suppose now that the cap on H-1B visas were set back to its value in 2001, namely, it was increased by 110,000 visas to bring the annual cap back to 195,000 visas. Our estimates imply that profits for firms in this bracket would increase by about 16 percent (0.15 log points).⁴⁸

In conclusion, our estimates reveal that the effects of immigration on firms' outcomes

⁴⁷It is not as clear for employment and R&D activities, where we do not find statistically significant results. We note, however, that the pattern of the point estimates for employment does line up as expected. This is a common finding across all our tables.

⁴⁸This increase is relative to the increase in profits for firms that did not apply for any LCAs in year 2001. Presumably, profits for these firms would be unaffected by a change in the cap for H-1B visas.

are the result of increased innovation that allows the firm to grow in size and to become more profitable. In addition, for the manufacturing industry we find evidence of positive effects also for the category of firms containing firms with a lower degree of dependency on H-1B visas (19 or more applications in 2001).

6.7 Sensitivity Analysis: Headquarters in the United States

A striking feature revealed by the data on LCAs ([Table 1](#)) is that several firms at the top of the ranking of applications have their headquarters outside of the United States (e.g. Satyam Computer Services and Infosys Technologies).⁴⁹ Naturally, these firms will tend to rely more heavily on foreign-workers, typically from the countries where their headquarters are located, than otherwise similar firms in the same industry. These firms therefore contribute to the identification of our coefficients of interest. One concern we may have is that perhaps our identification relies solely on these firms. To address this point we estimate our models excluding all firms with headquarters located outside of the United States, which is the case for about 10 percent of all firms in our dataset.

[Table 11](#) reports our findings. Let us begin with the top panel, which contains all firms in our dataset with headquarters in the United States. The estimates that we obtain are very similar to those reported in [Table 10](#) for the whole sample. The one noteworthy difference is that we now have a significant coefficient for the top bracket regarding the employment outcome. This is quite intuitive since US-based firms that experience an increase in productivity will be more likely to expand their workforce in the United States, close to their headquarters, than firms with headquarters in a foreign country. The estimates in the middle and bottom panels, R&D and Manufacturing, respectively, are very similar to those obtained with the full sample. Thus we conclude that our results are robust to excluding firms with headquarters outside of the United States from our analysis.

7 Discussion: fixed costs of innovation

We have provided evidence of a connection between changes in the national cap on new H-1B visas and several important outcomes of firms that file a significant number of applications for these visas. Furthermore, we have shown that the results arise from the industries with higher R&D (like manufacturing) and, within these industries, from the firms that consistently report R&D expenditures. In this section we take a closer look at the nature of the connection between H-1B workers, R&D activity, and firms' outcomes.

⁴⁹Often these firms are ADR, American Depository Receipts, as is the case for Satyam Computer Services.

To shed light on the channels at play, we begin by sketching a setup that could be formally incorporated in the simple model we presented in [Section 2](#). Assume that there is a shortage of skilled workers in the economy and, as a result, firms are not able to fill up all their vacancies for this type of workers. These skilled workers are employed in labs that conduct R&D activities. Assume also that there is a minimum size, S (in terms of the number of scientists), for a lab to be viable, which amounts to a fixed cost. Firms that are successful in hiring enough scientists and setting up their lab are then able to enhance their productivity, similar to [Bustos \(2011\)](#).⁵⁰

As described in [Section 2](#), firms are ex-ante heterogeneous in productivity, and thus make different choices regarding R&D. Lower productivity firms will not find it worthwhile to pay the fixed cost to set up a lab and, thus, they do not hire scientists. Higher productivity firms will want to set up a lab and will seek to hire skilled workers. Among these firms, some will be successful in matching with skilled workers that are already in the country. However, some other firms need to turn to H-1B visas in order to try to fill their needs, perhaps because they are startups that need to fill in several vacancies at once, or because they truly prefer foreign workers for whatever reason. In addition firms lose scientists for random reasons at some exogenous rate, which forces them to post new vacancies at every period. Firms that are successful at establishing a lab will enhance their initial productivity. As shown in [Section 2](#), these firms will then expand their size (both in terms of output and production employees) and will enjoy higher profits.

In the context of this model, firms will differ in their dependency on H-1B visas. Firms that do not wish to pay the fixed cost to set up a lab, or firms that are able to fill their vacancies with scientists already in the country, will not be affected by changes in the national cap on H-1B visas. However, firms that rely on foreign scientists, partly or totally, will be adversely affected by a reduction in the cap. Because of the fixed cost of innovation, the ideal empirical measure of a firm's dependency on H-1B visas is the number of applications they file, given by the difference between S , the optimal size of a lab, and the size of their current staff of scientists. This is precisely the measure of dependency that we use in our empirical analysis.

Note that normalizing by overall employment in the firm, as Kerr and Lincoln do, is not consistent with a *fixed* cost of innovation. Thus comparing our results to those

⁵⁰[Bustos \(2011\)](#) studies technology upgrading during periods of trade liberalization. She extends the model in [Melitz \(2003\)](#) by allowing firms to pay a fixed cost (denominated in units of labor) in order to adopt a higher-productivity technology characterized by a lower marginal cost of production. In equilibrium the firms with the highest initial productivity choose the technology that requires a fixed cost (and also choose to export).

obtained using their measure of dependency allows us to test for the presence of fixed costs of innovation. Next we estimate again our main specifications using the number of LCAs filed in 2001 normalized by total employment in the firm as the measure of dependency on H-1B visas. [Table B.1](#) in the Appendix presents the results. As we see in the table, the clear pattern of results found earlier vanishes, providing indirect evidence in favor of fixed costs of innovation. One may wonder whether the difference in the results is driven by the more flexible specification that we are employing here, compared to Kerr and Lincoln’s linear relationship in their firm-level analysis, or by the different sample of firms. To address this point we estimate linear models using both measures of dependency on a sample of firms that mimics Kerr and Lincoln’s ([Table B.2](#)).⁵¹ However, we again fail to find a significant relationship between this measure of dependency and changes in firm outcomes.

8 Conclusions

The main result of this paper is that, if the cap on H-1B visas were relaxed, only a subset of firms would benefit. These benefits would take the form of gains in average labor productivity, firm size, and profits. These are firms that conduct R&D and are heavy users of H-1B workers – they belong to the top quintile among filers of LCA applications. These firms tend to be large (e.g. Microsoft, Oracle or IBM) and play an important role as engines of innovation in the economy, which is reminiscent of the findings in [di Giovanni and Levchenko \(2013\)](#) and [di Giovanni et al. \(2014\)](#).⁵²

Taken together our results resonate with the findings of [Kerr and Lincoln \(2010\)](#). These authors argue that H-1B workers lead to higher rates of innovation, as measured by patenting activity. Our findings are consistent with this view, even though our proxy for innovation is R&D expenditure rather than patents. In addition we have provided evidence of effects at the firm level on a number of important outcomes. These empirical findings confirm the predictions of a simple monopolistic competition model where heterogeneous firms can choose to incur a fixed cost of innovation, in the form of hiring a minimum number of scientists, which enables them to enhance their productivity.

Future research should try to make further progress by exploiting new identification

⁵¹Recall that the sample in [Kerr and Lincoln \(2010\)](#) consisted of firms that filed a large number of LCAs in 2001 and had high levels of innovation, as measured by patenting activity. Their sample contained 77 firms. Using our data, we restrict to firms that filed at least 60 LCAs in year 2001 and that reported at least \$5,000 in R&D expenditures in both 2001 and 2006, resulting in 99 firms.

⁵²[di Giovanni and Levchenko \(2013\)](#) and [di Giovanni et al. \(2014\)](#) stress the important role of a relatively small number of very large firms in accounting for the consequences of trade liberalization for welfare and for aggregate fluctuations.

strategies and attempting to assemble even more comprehensive datasets. We believe that such a research agenda will be very fruitful, yielding important insights for the guidance of immigration policy in the United States and elsewhere.

References

- Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott, "Comparative Advantage and Heterogeneous Firms," *Review of Economic Studies*, 2007, 74 (1), 31–66.
- Bertoli, Simone, Herbert Brucker, Giovanni Facchini, and Giovanni Peri, "Understanding highly skilled migration in developed countries: The Upcoming battle for brains," in "Brain Drain and Brain Gain. The Global Competition to Attract High-Skilled Migrants," in T. Boeri, H. Bruecker, F. Docquier and H. Rapoport, Oxford University Press, September 2012.
- Borjas, George J., "The Labor-Market Impact of High-Skill Immigration," *American Economic Review*, *American Economic Association*, May 2005, 95 (2), 56–60.
- Bound, John, Clint Cummins, Zvi Griliches, Bronwyn H. Hall, and Adam B. Jaffe, "Who Does R&D and Who Patents?," in "R & D, Patents, and Productivity" NBER Chapters, National Bureau of Economic Research, Inc, National Bureau of Economic Research, Inc, May 1984, pp. 21–54.
- Bustos, Paula, "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review*, February 2011, 101 (1), 304–40.
- Chellaraj, Gnanaraj, Keith E. Maskus, and Aaditya Mattoo, "The Contribution of International Graduate Students to US Innovation," *Review of International Economics*, Wiley Blackwell, 08 2008, 16 (3), 444–462.
- di Giovanni, Julian and Andrei A. Levchenko, "Firm entry, trade, and welfare in Zipf's world," *Journal of International Economics*, 2013, 89 (2), 283–296.
- , Andrei Levchenko, and Isabelle Mejean, "Firms, Destinations, and Aggregate Fluctuations," *Econometrica*, *Econometric Society*, November 2014, 82 (4), 1303–1340.
- Dustmann, Christian and Albrecht Glitz, "How Do Industries and Firms Respond to Changes in Local Labor Supply?," CReAM Discussion Paper Series 1118, Centre for Research and Analysis of Migration (CReAM), Department of Economics, University College London September 2011.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti, "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy*, *University of Chicago Press*, 06 2010, 118 (3), 536–598.
- Hall, Bronwyn H and Rosemarie Ham Ziedonis, "The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995," *RAND Journal of Economics*, Spring 2001, 32 (1), 101–28.
- Hirschey, Mark, Hilla Skiba, and M. Babajide Wintoki, "The size, concentration and evolution of corporate R&D spending in U.S. firms from 1976 to 2010: Evidence and implications," *Journal of Corporate Finance*, *Elsevier*, 2012, 18 (3), 496–518.
- Hopenhayn, Hugo A, "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, *Econometric Society*, September 1992, 60 (5), 1127–50.
- Hunt, Jennifer and Marjolaine Gauthier-Loiselle, "How Much Does Immigration Boost Innovation?," *American Economic Journal: Macroeconomics*, *American Economic Association*, April 2010, 2 (2), 31–56.
- Jovanovic, Boyan, "Favorable Selection with Asymmetric Information," *The Quarterly Journal of Economics*, *MIT Press*, August 1982, 97 (3), 535–39.

- Kerr, William R. and William F. Lincoln, “The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention,” *Journal of Labor Economics*, 07 2010, 28 (3), 473–508.
- Levinsohn, James and Amil Petrin, “Estimating Production Functions Using Inputs to Control for Unobservables,” *Review of Economic Studies*, Wiley Blackwell, 04 2003, 70 (2), 317–341.
- Lowell, Lindsay, “H-1B Temporary Workers: Estimating the Population,” Technical Report 2000.
- Melitz, Marc J., “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, *Econometric Society*, November 2003, 71 (6), 1695–1725.
- and Gianmarco I. P. Ottaviano, “Market Size, Trade, and Productivity,” *Review of Economic Studies*, 2008, 75 (1), 295–316.
- Moretti, Enrico, “Workers’ Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions,” *American Economic Review*, June 2004, 94 (3), 656–690.
- Ortega, Francesc and Giovanni Peri, “Openness and income: The roles of trade and migration,” *Journal of International Economics*, Elsevier, 2014, 92 (2), 231–251.
- Parrotta, Pierpaolo, Dario Pozzoli, and Mariola Pytlikova, “Labor diversity and firm productivity,” *European Economic Review*, 2014, 66 (C), 144–179.
- , —, and —, “The nexus between labor diversity and firms innovation,” *Journal of Population Economics*, April 2014, 27 (2), 303–364.
- Paserman, M, “Do high-skill immigrants raise productivity? Evidence from Israeli manufacturing firms, 1990-1999,” *IZA Journal of Migration*, December 2013, 2 (1), 1–31.
- Pavcnik, Nina, “Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants,” *Review of Economic Studies*, 2002, 69 (1), 245–276.
- Peri, Giovanni, “The Effect Of Immigration On Productivity: Evidence From U.S. States,” *The Review of Economics and Statistics*, February 2012, 94 (1), 348–358.
- , Kevin Shih, and Chad Sparber, “STEMWorkers, H1B Visas and Productivity in US Cities,” Norface Discussion Paper Series, Norface Research Programme on Migration, Department of Economics, University College London 2013009, Norface Research Programme on Migration, Department of Economics, University College London March 2013.
- Quispe-Agnoli, Myriam and Madeline Zavodny, “The effect of immigration on output mix, capital, and productivity,” *Economic Review*, *Federal Reserve Bank of Atlanta*, 2002, (Q1), 17–27.
- Teruel-Carrizosa, Mercedes and Agust Segarra-Blasco, “Immigration and Firm Growth: Evidence from Spanish cities,” Working Papers, Xarxa de Referncia en Economia Aplicada (XREAP) XREAP2008-11, Xarxa de Referncia en Economia Aplicada (XREAP) November 2008.
- Trax, Michaela, Stephan Brunow, and Jens Suedekum, “Cultural diversity and plant-level productivity,” Technical Report 2013.
- Wadhwa, Vivek, AnnaLee Saxenian, Ben A. Rissing, and G. Gereffi, “America’s New Immigrant Entrepreneurs: Part I (,” Technical Report 23, Duke Science, Technology Innovation Papers January 2007.

Table 1: Top filers of LCAs

Rank	Company Name	LCAs
Year 2001		
5	LUCENT TECHNOLOGIES INC	564
4	SATYAM COMPUTR SVC LTD -ADR	1263
3	IBM CREDIT CORP	1482
2	MICROSOFT CORP	1515
1	ORACLE CORP	1736
Year 2002		
5	IBM CREDIT CORP	399
4	ACCENTURE LTD	430
3	ORACLE CORP	511
2	SATYAM COMPUTR SVC LTD -ADR	675
1	MICROSOFT CORP	1316
Year 2003		
5	SUN MICROSYSTEMS INC	720
4	ORACLE CORP	795
3	CISCO SYSTEMS INC	1339
2	SATYAM COMPUTR SVC LTD -ADR	1506
1	MICROSOFT CORP	1926
Year 2004		
5	ORACLE CORP	962
4	IBM CREDIT CORP	1000
3	PATNI COMPUTER SYSTEMS -ADR	1991
2	MICROSOFT CORP	2260
1	SATYAM COMPUTR SVC LTD -ADR	3616
Year 2005		
5	INTEL CORP	1280
4	INFOSYS TECHNOLOGIES -ADR	1742
3	PATNI COMPUTER SYSTEMS -ADR	1778
2	MICROSOFT CORP	2142
1	SATYAM COMPUTR SVC LTD -ADR	3280
Year 2006		
5	ORACLE CORP	1262
4	IBM CREDIT CORP	1577
3	PATNI COMPUTER SYSTEMS -ADR	2033
2	SATYAM COMPUTR SVC LTD -ADR	4258
1	MICROSOFT CORP	4406

Table 2: Descriptive statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Sales 2001	2629.352	10627.054	0.002	218529	3,945
Emp. 2001	10296.696	40721.722	1	1383000	3,945
Capital 2001	1094.102	4542.498	0.008	89602	3,945
<i>R&D</i> 2001	112.646	520.26	0	7400	1923
Sales 2006	3975.286	16704.722	0.002	345977	3,945
Emp. 2006	11447.071	47273.98	1	1900000	3,945
Capital 2006	1522.554	6547.444	0.001	113687	3,945
<i>R&D</i> 2006	152.594	690.894	0	8258	1962
$\Delta \ln Sales$	0.457	0.882	-6.188	7.468	3,945
$\Delta \ln \frac{Sales}{Emp.}$	0.291	0.689	-6.824	6.816	3,945
$\Delta \ln R\&D$	0.186	0.97	-9.683	6.807	1,854
$\Delta \ln R\&D(imputed)$	0.204	1.074	-9.683	10.012	3,945
$\Delta \ln Emp.$	0.166	0.702	-5.27	6.731	3,945
$\Delta \ln Capital$	0.199	1.017	-6.217	7.922	3,945
$\Delta \ln TFP$	0.213	0.641	-7.093	7.019	3,789
$\Delta \ln Materials$	0.461	0.754	-5.776	3.781	2,198
$\Delta \ln Profits$	0.505	0.821	-6.256	6.277	3,648
<i>LCA</i> 2001	5.152	41.492	0	1515	3,945
<i>LCA</i> 2006	9.573	108.71	0	4406	3,945
$D\{LCA_{2001} = 0\}$	0.794	0.407	0	1	3,945
$D\{LCA_{2001} = 1\}$	0.053	0.226	0	1	3,945
$D\{1 < LCA_{2001} \leq 3\}$	0.050	0.22	0	1	3,945
$D\{3 < LCA_{2001} \leq 7\}$	0.033	0.18	0	1	3,945
$D\{7 < LCA_{2001} \leq 10\}$	0.016	0.129	0	1	3,945
$D\{10 < LCA_{2001} \leq 18\}$	0.013	0.117	0	1	3,945
$D\{18 < LCA_{2001} \leq 59\}$	0.018	0.135	0	1	3,945
$D\{LCA_{2001} \geq 60\}$	0.021	0.145	0	1	3,945

Notes: TFP is computed as a Solow residual, where the coefficients are estimated at the 2-digit industry level. ΔX corresponds to the change between years 2006 and 2001 for variable X. *Capital* refers to net capital. Profits refers to gross profits. Imputed $\Delta \ln R\&D$ assigns a value of zero to all firms with a missing value for $\Delta R\&D$. All variables (in levels) are in current thousands of dollars.

Table 3: Descriptive statistics - R&D subsample

Variable	Mean	Std. Dev.	Min.	Max.	N
Sales 2001	4960.564	17160.8	0.004	187510	969
Emp. 2001	17228.753	48172.254	14	484000	969
Capital exp. 2001	1816.759	7047.616	0.01	89602	969
<i>R&D</i> 2001	220.413	716.493	5.001	7400	969
Sales 2006	7686.526	27150.422	0.002	335086	969
Emp. 2006	17492.587	46920.677	2	475000	969
Capital exp. 2006	2417.193	9961.898	0.008	113687	969
<i>R&D</i> 2006	305.557	959.481	5.065	8258	969
$\Delta \ln Sales$	0.488	0.934	-4.962	7.468	969
$\Delta \ln \frac{Sales}{Emp.}$	0.355	0.79	-4.539	6.816	969
$\Delta \ln R\&D$	0.282	0.696	-2.544	2.818	969
$\Delta \ln R\&D$ (imputed)	0.282	0.696	-2.544	2.818	969
$\Delta \ln Emp.$	0.133	0.602	-3.068	2.461	969
$\Delta \ln Capital$	0.054	0.941	-5.174	6.113	969
$\Delta \ln TFP$	0.319	0.784	-4.57	7.019	947
$\Delta \ln Materials$	0.560	0.627	-3.173	3.781	591
$\Delta \ln Profits$	0.497	0.76	-4.305	3.373	856
LCA 2006	23.03	160.24	0	4406	969
LCA 2001	12.908	63.766	0	1515	969
$D\{LCA_{2001} = 0\}$	0.613	0.487	0	1	969
$D\{LCA_{2001} = 1\}$	0.057	0.232	0	1	969
$D\{1 < LCA_{2001} \leq 3\}$	0.071	0.257	0	1	969
$D\{3 < LCA_{2001} \leq 7\}$	0.068	0.252	0	1	969
$D\{7 < LCA_{2001} \leq 10\}$	0.044	0.206	0	1	969
$D\{10 < LCA_{2001} \leq 18\}$	0.044	0.206	0	1	969
$D\{18 < LCA_{2001} \leq 59\}$	0.047	0.213	0	1	969
$D\{LCA_{2001} \geq 60\}$	0.055	0.228	0	1	969

Notes: The sample reported here corresponds to firms that reported R&D expenses of at least \$5,000 in both 2001 and 2006. TFP is computed as a Solow residual, where the coefficients are estimated at the 2-digit industry level. ΔX corresponds to the change between years 2006 and 2001 for variable X. *Capital* refers to net capital. Profits refers to gross profits. Imputed $\Delta \ln R\&D$ assigns a value of zero to all firms with a missing value for $\Delta R\&D$. All variables (in levels) are in current thousands of dollars.

Table 4: Parallel trends

Dep. var. is the ln of	(1) Sales	(2) Emp	(3) Profits	(4) R&D
Panel 1: Level for 2002				
LCA2001	0.77*** [0.20]	0.71*** [0.18]	0.79*** [0.20]	0.71*** [0.22]
Observations	6,831	6,831	6,454	2,807
Panel 2: Change 2000-2003				
LCA2001	-0.01 [0.01]	0.00 [0.01]	-0.01 [0.01]	-0.03 [0.02]
Observations	5,764	5,764	5,306	1,691

Notes: The number of LCA in 2001 has been divided by 100 to rescale the coefficients. 2-digit industry fixed effects included in all regressions. Standard errors are robust to heteroskedasticity.

Table 5: Linear model

Dep. var. is $\Delta \ln$ of	(1) <i>Sales/Emp</i>	(2) <i>Sales</i>	(3) <i>Emp</i>	(4) <i>Profits</i>	(5) <i>Sales/Emp</i>	(6) <i>Sales</i>	(7) <i>Emp</i>	(8) <i>Profits</i>	(9) <i>Sales/Emp</i>	(10) <i>Sales</i>	(11) <i>Emp</i>	(12) <i>Profits</i>
$\Delta H1B \times LCA2001$	0.72*** [0.18]	0.89* [0.48]	0.17 [0.50]	0.45 [0.42]	0.97*** [0.22]	0.91* [0.54]	-0.07 [0.56]	0.29 [0.46]	0.68*** [0.24]	1.92*** [0.43]	1.24*** [0.39]	1.12*** [0.40]
LCA2001	-0.08*** [0.02]	-0.10* [0.05]	-0.02 [0.06]	-0.05 [0.05]	-0.11*** [0.02]	-0.10* [0.06]	0.01 [0.06]	-0.03 [0.05]	-0.07*** [0.03]	-0.21*** [0.05]	-0.14*** [0.04]	-0.12*** [0.04]
$\Delta Re/H1B \times LCA2001$	0.28*** [0.07]	0.35* [0.19]	0.07 [0.20]	0.17 [0.16]	0.38*** [0.09]	0.35* [0.21]	-0.03 [0.22]	0.11 [0.18]	0.26*** [0.09]	0.75*** [0.17]	0.48*** [0.15]	0.44*** [0.15]
Observations	3,943	3,943	3,943	3,646	809	809	809	749	968	968	968	855
Sample	All	All	All	All	LCA01	LCA01	LCA01	LCA01	R&D	R&D	R&D	R&D

Heteroskedasticity-robust standard errors in brackets

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

Notes: The number of LCA in 2001 has been divided by 100 to rescale the coefficients. Columns 1-4 report estimates based on the whole sample (except for two outliers with LCA in 2001 over 1000). Columns 5-8 are based on the subsample with one or more LCA applications in 2001. Columns 9-12 are based on the subsample with non-missing R&D expenses for years 2001 and 2006. $\Delta H1B = -110,000$ and corresponds to the change in the cap in H1B visas between years 2001 and 2006. $\Delta Re/H1B$ is the change in the annual cap of H1B visas divided by the number of petitions in the same year (excluding petitions that were granted for higher education since these are not subject to the cap). All regressions contain 2-digit sector dummies. (1) This $\Delta R\&D$ variable imputes a (small positive) value to observations with values.

Table 6: Non-parametric

Dep. var. is $\Delta \ln$ of	(1) $\frac{Sales}{Emp}$	(2) Sales	(3) Emp	(4) Profits	(5) $\frac{Sales}{Emp}$	(6) Sales	(7) Emp	(8) Profits	(9) $\frac{Sales}{Emp}$	(10) Sales	(11) Emp	(12) Profits
$\Delta H1B \times D\{LCA01 = 1\}$	-0.41 [0.46]	0.48 [0.56]	0.89* [0.49]	0.93 [0.58]	0.53 [1.23]	1.12 [1.41]	0.59 [0.67]	0.10 [1.10]	-0.70 [0.46]	0.23 [0.58]	0.94 [0.62]	1.20* [0.67]
$\Delta H1B \times D\{1 < LCA01 \leq 3\}$	-1.03** [0.48]	-0.49 [0.57]	0.54 [0.40]	0.82 [0.56]	-1.62* [0.90]	-1.52 [1.07]	0.10 [0.68]	0.84 [0.99]	-0.58 [0.56]	0.27 [0.66]	0.85* [0.51]	1.02 [0.68]
$\Delta H1B \times D\{3 < LCA01 \leq 7\}$	-0.78 [0.70]	-1.10 [0.77]	-0.32 [0.60]	-0.48 [0.61]	-0.44 [1.04]	-1.17 [1.17]	-0.73 [0.66]	-1.33 [0.92]	-0.81 [1.01]	-0.64 [1.06]	0.17 [1.03]	0.46 [0.79]
$\Delta H1B \times D\{7 < LCA01 \leq 10\}$	-0.56 [0.66]	0.40 [0.96]	0.96 [0.66]	0.28 [0.88]	-1.01 [0.78]	0.22 [1.17]	1.23* [0.74]	-0.13 [1.25]	0.42 [1.32]	1.59 [1.82]	1.17 [1.32]	1.15 [1.24]
$\Delta H1B \times D\{10 < LCA01 \leq 18\}$	1.16* [0.63]	0.86 [1.03]	-0.30 [0.90]	-0.72 [1.09]	1.01 [0.81]	0.95 [1.26]	-0.06 [1.09]	-1.06 [1.32]	2.49** [1.23]	2.25 [1.64]	-0.24 [1.04]	0.99 [1.48]
$\Delta H1B \times D\{18 < LCA01 \leq 59\}$	0.60 [0.37]	1.27* [0.67]	0.67 [0.64]	1.46* [0.84]	0.73 [0.57]	1.63* [0.91]	0.90 [0.78]	1.44 [1.14]	0.53 [0.60]	1.49 [1.11]	0.96 [1.16]	1.57 [1.30]
$\Delta H1B \times D\{LCA01 \geq 60\}$	1.14*** [0.38]	2.06*** [0.62]	0.92 [0.61]	1.36** [0.60]	1.07** [0.49]	3.18*** [0.79]	2.12*** [0.67]	2.07*** [0.73]	1.62** [0.77]	0.80 [1.12]	-0.82 [1.21]	0.34 [1.16]
Observations	3,945	3,945	3,945	3,648	969	969	969	856	2,976	2,976	2,976	2,792
R-squared	0.02	0.03	0.03	0.04	0.02	0.03	0.03	0.04	0.02	0.04	0.04	0.05
Sample by $R\&D$ status	All	All	All	All	Yes	Yes	Yes	Yes	No	No	No	No

Heteroskedasticity-robust standard errors in brackets

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

Notes: Columns 1-4 are estimated on the whole sample. Columns 5-8 are estimated on the subsample with non-missing observations on R&D expenses. $\Delta H1B = -110,000$ and corresponds to the change in the cap in H1B visas between years 2001 and 2006. The omitted category are firms that had exactly zero LCA applications in year 2001. All regressions contain 2-digit sector dummies. The distribution of firms by 2001 LCA applications is as follows: missing values for 1,301 firms. Among the non-missing (3,945 firms), 79% had zero applications, 5.5 percent had 1 application, 5.2 percent had 2 or 3 applications, 3.4 percent had 4-7 applications, 1.7 percent had 8-10 applications, 1.44 percent had 11-18 applications, 1.9 percent had 19-59, and 2.2 percent had 60 or more applications in 2001. (1) This $\Delta R\&D$ variable imputes a (small positive) value to observations with values.

Table 7: The channels: TFP and R&D

Dep. var. is $\Delta \ln$ of	(1) $\frac{Sales}{Emp}$	(2) TFP ¹	(3) TFP ²	(4) R&D*	(5) $\frac{Sales}{Emp}$	(6) TFP ¹	(7) TFP ²	(8) R&D
$\Delta H1B \times D\{LCA01 = 1\}$	-0.41 [0.46]	-0.42 [0.45]	-0.21 [0.39]	-0.07 [0.73]	0.53 [1.23]	0.40 [1.26]	0.71 [0.79]	1.25* [0.73]
$\Delta H1B \times D\{1 < LCA01 \leq 3\}$	-1.03** [0.48]	-1.00** [0.49]	0.37 [0.48]	1.14 [0.81]	-1.62* [0.90]	-1.17 [0.96]	0.68 [0.55]	-0.22 [0.71]
$\Delta H1B \times D\{3 < LCA01 \leq 7\}$	-0.78 [0.70]	-0.45 [0.59]	0.12 [0.37]	0.71 [0.57]	-0.44 [1.04]	-0.63 [1.11]	0.54 [0.56]	0.13 [0.71]
$\Delta H1B \times D\{7 < LCA01 \leq 10\}$	-0.56 [0.66]	-0.70 [0.64]	-0.23 [0.53]	1.77** [0.72]	-1.01 [0.78]	-0.93 [0.79]	-0.48 [0.82]	1.81* [0.97]
$\Delta H1B \times D\{10 < LCA01 \leq 18\}$	1.16* [0.63]	0.90 [0.62]	0.64 [0.70]	-0.45 [1.04]	1.01 [0.81]	1.02 [0.79]	0.32 [0.88]	-0.77 [1.26]
$\Delta H1B \times D\{18 < LCA01 \leq 59\}$	0.60 [0.37]	0.46 [0.35]	0.54 [0.37]	0.34 [1.11]	0.73 [0.57]	0.74 [0.55]	0.51 [0.55]	0.88 [1.05]
$\Delta H1B \times D\{LCA01 \geq 60\}$	1.14*** [0.38]	0.97*** [0.35]	0.42 [0.34]	0.61 [0.98]	1.07** [0.49]	1.14** [0.52]	0.13 [0.47]	1.76** [0.86]
Observations	3,945	3,789	2,115	3,945	969	947	580	969
R-squared	0.02	0.03	0.06	0.03	0.02	0.02	0.05	0.04
Sample by R&D status	All	All	All	All	Yes	Yes	Yes	Yes

Heteroskedasticity-robust standard errors in brackets

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

Notes: Columns 1-4 are estimated on the whole sample. Columns 5-8 are estimated on the subsample with non-missing observations on R&D expenses. $\Delta H1B = -110,000$ and corresponds to the change in the cap in H1B visas between years 2001 and 2006. The omitted category are firms that had exactly zero LCA applications in year 2001. All regressions contain 2-digit sector dummies. The distribution of firms by 2001 LCA applications is as follows: missing values for 1,301 firms. Among the non-missing (3,945 firms), 79% had zero applications, 5.5 percent had 1 application, 5.2 percent had 2 or 3 applications, 3.4 percent had 4-7 applications, 1.7 percent had 8-10 applications, 1.44 percent had 11-18 applications, 1.9 percent had 19-59, and 2.2 percent had 60 or more applications in 2001. (1) TFP based on naive OLS estimation of production functions. (2) TFP based on Levinsohn-Petrin estimation of production functions (revenue approach). (*) This $\Delta R\&D$ variable imputes a (small positive) value to observations with missing values.

Table 8: Sector-specific factor shares

Description sector	Sector NAICS2dg	Capital	Employment
Agriculture, Forestry, etc.	11	0.47	0.56
Mining and gas and oil extraction	21	0.72	0.38
Utilities	22	0.48	0.47
Construction	23	0.08	0.92
Manufacturing - Food	31	0.27	0.66
Manufacturing - Wood	32	0.30	0.88
Manufacturing - Many categories	33	0.20	0.84
Wholesale trade	42	0.16	0.79
Retail Trade (1)	44	0.45	0.43
Retail Trade (2)	45	0.23	0.61
Transportation (1)	48	0.28	0.56
Transportation (2)	49	0.12	0.75
Information (media and telecomm.)	51	0.21	0.82
Finance and Insurance	52	0.17	0.89
Real estate	53	0.29	0.57
Professional, Scientific and Technical services	54	0.15	0.88
Administrative and Support	56	0.34	0.53
Educational services	61	0.43	0.27
Health care and social assistance	62	0.20	0.71
Arts and entertainment	71	0.39	0.48
Accommodation and Food services	72	0.28	0.69
Other services (repair and maintenance)	81	0.24	0.53

Notes: In each row we report the OLS estimates of the following linear regression estimated sector by sector: $\ln Sales_{it} = \beta + \beta_k \ln Capital_{it} + \beta_{emp} \ln Emp_{it} + u_{it}$. The regressions do not include an intercept or year dummies to be consistent with a Cobb-Douglas production function. Manufacturing 33 contains many categories: metal, industrial, computers, semiconductors, electrical, motor vehicles, medical equipment, others.

Table 9: Industry averages

Sector	Year		2001		2001		2001		2001		2006		2006		2006		2006	
	NAICS (2dg)	Obs.	LCA	Sales	Emp.	R&D	Obs.	LCA	Sales	Emp.	Obs.	LCA	Sales	Emp.	RnD	Obs.	LCA	Sales
Other services (repair and maintenance)	81	32	0.06	339.45	6263.19	2.9	18	0.33	609.8	10723.72	6.12	0.33	609.8	10723.72	6.12	0.33	609.8	10723.72
Arts and entertainment	71	42	0.07	286.89	4641.02	0.92	35	0.29	443.5	5064.66	0.27	0.29	443.5	5064.66	0.27	0.29	443.5	5064.66
Agriculture and Forestry	11	18	0.17	678.78	8385.39	36.18	12	0.5	1240.44	11721.08	18.07	0.5	1240.44	11721.08	18.07	0.5	1240.44	11721.08
Accommodation and Food	72	137	0.31	968.77	21797.8	0.86	92	0.61	1805.67	34409.79	0.65	0.61	1805.67	34409.79	0.65	0.61	1805.67	34409.79
Utilities	22	247	0.43	4464.27	6891.92	18.94	227	0.77	3624.71	4185.56		0.77	3624.71	4185.56		0.77	3624.71	4185.56
Manufacturing - Food	31	237	0.49	3001.77	13362.35	65.19	163	0.82	3955.12	13878.26	81.76	0.82	3955.12	13878.26	81.76	0.82	3955.12	13878.26
Real estate	53	86	0.51	678.2	4729.1	7.62	83	1.31	970.47	5552.01	13.34	1.31	970.47	5552.01	13.34	1.31	970.47	5552.01
Wholesale trade	42	195	0.56	2839.58	4244.69	3.23	136	0.91	4439.19	5413.26	10.09	0.91	4439.19	5413.26	10.09	0.91	4439.19	5413.26
Mining and extraction	21	178	0.7	1212.39	3585.51	27.67	223	0.83	1863.14	3620.96	39.64	0.83	1863.14	3620.96	39.64	0.83	1863.14	3620.96
Educational services	61	17	0.71	241.59	3128.53	0.46	15	1.2	514.44	5445.2		1.2	514.44	5445.2		1.2	514.44	5445.2
Health care and social assistance	62	111	0.73	1046.83	10324.91	1.84	98	2.16	1487.96	11616.58	1.02	2.16	1487.96	11616.58	1.02	2.16	1487.96	11616.58
Retail trade (2)	45	108	0.9	5371.92	34874.05	3.36	78	6.76	9803.01	52892.94	13.01	6.76	9803.01	52892.94	13.01	6.76	9803.01	52892.94
Transportation (1)	48	137	1.39	3047.55	11673.57	1.03	121	2.34	3362.45	10982.98		2.34	3362.45	10982.98		2.34	3362.45	10982.98
Manufacturing - Wood	32	721	1.97	3025.68	8444.42	133.81	619	4.34	5829.41	8259.57	217.87	4.34	5829.41	8259.57	217.87	4.34	5829.41	8259.57
Retail trade (1)	44	172	2.3	3878.32	24085.79	0.1	125	13.59	7250.65	32969.31	0.16	13.59	7250.65	32969.31	0.16	13.59	7250.65	32969.31
Finance and insurance	52	881	2.73	2420.34	5684.79	6.85	946	5	3288.24	4955.81	5.44	5	3288.24	4955.81	5.44	5	3288.24	4955.81
Administrative and support	56	155	2.92	768.97	14433.84	12.84	109	5.01	1329.84	20681.86	23.76	5.01	1329.84	20681.86	23.76	5.01	1329.84	20681.86
Manufacturing - Many categories	33	1587	4.81	2017.87	8281.66	131.6	1144	9.72	2825.74	8882.46	149.15	9.72	2825.74	8882.46	149.15	9.72	2825.74	8882.46
Construction	23	72	5.93	1684.19	7585.04	124.5	68	5.32	3248.67	7251.31	105.95	5.32	3248.67	7251.31	105.95	5.32	3248.67	7251.31
Information (media and telecom.)	51	881	9.56	1438.5	5978.07	50.87	528	17.89	2246.15	6738.85	83.49	17.89	2246.15	6738.85	83.49	17.89	2246.15	6738.85
Transportation (2)	49	14	11.64	10672.57	126394.07	14.5	9	0.78	7719.23	69021.34		0.78	7719.23	69021.34		0.78	7719.23	69021.34
Professional, scientific and technical	54	366	19.24	860.09	5036.92	73.98	236	35.57	1237.98	7129.43	74.21	35.57	1237.98	7129.43	74.21	35.57	1237.98	7129.43

Notes: Industry averages for LCA, sales, employment, and R&D expenses. We use the 2-digit NAICS classification and we have ranked industries in increasing order by the 2001 LCAs. Specifically, the top three industries by LCA in 2001 are Professional, Scientific and Technical services (54), Transportation (49), and Information (media and telecommunications, 51). The bottom three are Agriculture, Forestry and so on (11), Arts and Entertainment (71), and Other Services (repair and maintenance, 81).

Table 10: Industry results

Dep. var. is $\Delta \ln$ of	(1) $\frac{Sales}{Emp}$	(2) Sales	(3) Emp	(4) Profits	(5) TFP^1	(6) $R\&D$
Sample: All firms (N=3,945)						
$\Delta H1B \times D\{1 \leq LCA2001 \leq 18\}$	-0.56** [0.29]	-0.10 [0.35]	0.46* [0.27]	0.44 [0.33]	-0.52* [0.27]	0.24 [0.49]
$\Delta H1B \times D\{19 \leq LCA2001 \leq 59\}$	0.59 [0.37]	1.27* [0.67]	0.67 [0.64]	1.47* [0.84]	0.45 [0.35]	0.11 [1.02]
$\Delta H1B \times D\{LCA2001 \geq 60\}$	1.13*** [0.37]	2.05*** [0.62]	0.93 [0.61]	1.38** [0.60]	0.96*** [0.35]	0.07 [1.01]
Sample: R&D (N=969)						
$\Delta H1B \times D\{1 \leq LCA2001 \leq 18\}$	-0.41 [0.55]	-0.26 [0.65]	0.15 [0.40]	-0.26 [0.56]	-0.34 [0.56]	0.39 [0.44]
$\Delta H1B \times D\{19 \leq LCA2001 \leq 59\}$	0.72 [0.57]	1.62* [0.91]	0.90 [0.78]	1.47 [1.13]	0.73 [0.54]	0.89 [1.05]
$\Delta H1B \times D\{LCA2001 \geq 60\}$	1.06** [0.49]	3.17*** [0.78]	2.12*** [0.67]	2.09*** [0.73]	1.14** [0.52]	1.76** [0.85]
Sample: Manufacturing (N=1,549)						
$\Delta H1B \times D\{1 \leq LCA2001 \leq 18\}$	-0.93* [0.49]	-0.31 [0.56]	0.63 [0.42]	0.59 [0.50]	-1.15** [0.49]	0.82* [0.48]
$\Delta H1B \times D\{19 \leq LCA2001 \leq 59\}$	0.28 [0.64]	2.39*** [0.92]	2.11*** [0.77]	3.06** [1.37]	0.16 [0.53]	1.12 [1.19]
$\Delta H1B \times D\{LCA2001 \geq 60\}$	0.48 [0.52]	2.90*** [0.87]	2.42*** [0.72]	2.33*** [0.78]	0.42 [0.53]	1.79** [0.87]

Heteroskedasticity-robust standard errors in brackets

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

Notes: The omitted category are firms with zero LCAs in 2001. All regressions include 2-digit sector dummy variables. The first panel contains the whole sample of firms. The second panel is the R&D subsample. The third panel is the subsample of manufacturing sectors (NAICS 31-33).

Table 11: Industry results - US firms only

Dep. var. is $\Delta \ln$ of	(1) $\frac{Sales}{Emp}$	(2) Sales	(3) Emp	(4) Profits	(5) TFP^1	(6) $R\&D$
Sample: All firms (N=3,653)						
$\Delta H1B \times D\{1 \leq LCA2001 \leq 18\}$	-0.66** [0.30]	-0.31 [0.36]	0.35 [0.28]	0.23 [0.34]	-0.58** [0.28]	0.13 [0.50]
$\Delta H1B \times D\{19 \leq LCA2001 \leq 59\}$	0.67* [0.39]	1.01 [0.72]	0.34 [0.69]	1.05 [0.91]	0.57 [0.37]	-0.12 [1.09]
$\Delta H1B \times D\{LCA2001 \geq 60\}$	0.96** [0.39]	2.23*** [0.63]	1.27** [0.61]	1.39** [0.63]	0.79** [0.36]	0.32 [0.99]
Sample: R&D (N=853)						
$\Delta H1B \times D\{1 \leq LCA2001 \leq 18\}$	-0.50 [0.60]	-0.36 [0.70]	0.15 [0.42]	-0.46 [0.59]	-0.32 [0.61]	0.21 [0.46]
$\Delta H1B \times D\{19 \leq LCA2001 \leq 59\}$	0.64 [0.64]	1.29 [0.96]	0.66 [0.81]	0.95 [1.20]	0.78 [0.61]	0.54 [1.10]
$\Delta H1B \times D\{LCA2001 \geq 60\}$	0.88 [0.55]	3.27*** [0.86]	2.39*** [0.72]	1.93** [0.80]	1.04* [0.58]	1.95** [0.90]
Sample: Manufacturing (N=1,418)						
$\Delta H1B \times D\{1 \leq LCA2001 \leq 18\}$	-1.08** [0.51]	-0.49 [0.58]	0.59 [0.44]	0.41 [0.52]	-1.26** [0.51]	0.78 [0.50]
$\Delta H1B \times D\{19 \leq LCA2001 \leq 59\}$	0.18 [0.66]	2.30** [0.96]	2.12*** [0.80]	2.92** [1.42]	0.06 [0.55]	1.12 [1.24]
$\Delta H1B \times D\{LCA2001 \geq 60\}$	0.34 [0.56]	2.86*** [0.91]	2.52*** [0.77]	2.25*** [0.85]	0.31 [0.57]	1.89** [0.91]

Heteroskedasticity-robust standard errors in brackets

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

Notes: All samples here exclude firms with headquarters outside of the United States, which is slightly below 10 percent of all firms in our dataset. The omitted category are firms with zero LCAs in 2001. All regressions include 2-digit sector dummy variables. The first panel contains the whole sample of firms. The second panel is the R&D subsample. The third panel is the subsample of manufacturing sectors (NAICS 31-33).

Figure 1: H1B policies

Notes:.

Appendices

A Construction firm-level TFP

We build several measures of TFP at the firm-level. For all of them we make use of the longitudinal dimension of our data (at the annual level). Our sample period is 2001-2006 and we have over 40,000 firm-year observations.

Building firm-level estimates of TFP is a two-step process. In the first step we need to estimate the coefficients of the production function, typically within the Cobb-Douglas family. The second step then uses these estimates to build a residual TFP term. In our analysis we present a naive estimation of the firm-level production functions and a more sophisticated ones. However, there is a trade-off. While the naive estimation makes strong assumptions – i.e. exogeneity of the quantities of inputs –, the other is more demanding in terms of data and entails a substantial loss in the number of observations (we are left with slightly over 25,000 firm-year observations). We now describe our estimation and results in detail.

A.1 Naive estimation

Let us assume that firms produce by means of Cobb-Douglas production functions, with exponents allowed to vary by industry and not constrained to add up to one. In particular, we assume that output (measured by sales) is produced using capital and labor. In logs, the level of output for firm i in industry j in year t is thus given by

$$y_{ijt} = \beta_0 + \beta_k^j k_{ijt} + \beta_\ell^j \ell_{ijt} + \varepsilon_{ijt}. \quad (21)$$

We estimate this regression model by OLS clustering standard errors at the firm level to allow for autocorrelation in the idiosyncratic unobserved productivity shocks. Note though that consistent estimation of the β coefficients requires the strong assumption that $\{k_{ijt}, \ell_{ijt}\}$ are uncorrelated with unobserved productivity shocks.

A.2 Levinsohn and Petrin (2003) estimator

We now relax the assumption of exogenous regressors in the estimation of the production functions. Intuitively, [Levinsohn and Petrin \(2003\)](#) propose a GMM approach that relies on the use of additional information (on the cost of materials) as well as on some moment conditions, namely, the current capital stock is predetermined and does not respond to contemporaneous productivity shocks, and the previous period's level of material usage is uncorrelated with current productivity shocks.

This estimator comes in two variations. In the first one the estimation of the production function is done using value-added, which is the one we adopt for ease of comparability.⁵³ We first define value added (v_{ijt}) as sales minus the cost of materials.⁵⁴ We then estimate

$$v_{ijt} = \beta_0 + \beta_k^j k_{ijt} + \beta_\ell^j \ell_{ijt} + (\omega_{ijt} + \varepsilon_{ijt}), \quad (22)$$

where ω_{ijt} may be correlated with input usage, introducing a problem of endogenous regressors. The estimation method proposed by [Levinsohn and Petrin \(2003\)](#) delivers, under some assumptions, consistent estimates of β_k^j and β_ℓ^j .⁵⁵

⁵³The results are very similar if we employ the other variation of this method, which is based on sales and includes the use of materials explicitly as a regressor.

⁵⁴Specifically, our Compustat data allows us to estimate the cost of materials as follows. We add up the cost of goods sold (data41) and selling, general and administrative expenses (xsga). We then subtract depreciation and amortization (data14) and wage expenses.

⁵⁵We estimate this model separately for each 2-digit sector.

A.3 Data and estimates production functions

We begin by reporting some summary statistics of the variables involved in these estimations.

Table A.1: Summary statistics - production functions

Variable	Mean	Std. Dev.	Min.	Max.	N
naics2dg	42.31	12.943	11	99	40502
sales	2549.485	11422.374	0.001	345977	40502
materials	1352.935	7728.088	-119640.945	270660	30567
value added	823.752	3698.091	-1054.991	119726.75	30567
icapt	2543.143	13528.148	-25767	597207.209	40416
employment	8859.791	38934.078	11	2545209	40502

The observations in this table correspond to firm-year cells. It spans years 2001-2006 and over 9,000 firms but the panel is unbalanced, with observations for some years missing for a good number of firms. The most important observation is that the variable materials can only be constructed for about half of the industries and little over half of the individual firm-year observations.

Next we warm up by applying our two estimation methods to all firms pooled together, that is, imposing equal coefficients for all industries. The table below summarizes the results:

Table A.2: Estimates Cobb-Douglas production functions (pooled)

Factor of Production	OLS	Levinsohn-Petrin VA
L	0.64	0.64
K	0.42	0.41
M	-	-
Sum coeff.	1.06	1.05
Number Obs.	40,502	26,326
Number sectors	24	13

Several observations are worth noting. First, the number of observations and the number of sectors contained in each of the effective samples differ. For the naive OLS estimation we have over 40 thousand observations whereas the implementation of the Levinsohn-Petrin estimator is based on many fewer observations (26 thousand) because the cost of materials is not available for a number of sectors. The final observation is that the estimated factor shares (for the aggregate production function) in columns 1 and 2 are very similar, at least for the pooled sample.

Now we estimate the production functions for each sector separately, by each of the two estimation methods. The next table presents the results. The top line displays again the estimated coefficients based on the sample that pools all sectors. Below we report the estimates for each 2-digit sector. Clearly, there is a lot of variation across sectors. There is also a positive correlation (0.20) between the (labor) shares obtained by OLS and by the Levinsohn-Petrin GMM estimator. However, there are also important differences. For example the OLS estimation identifies the Construction sector (23) as the one with the highest labor share, in contrast the Levinsohn-Petrin approach implies that Professional, Scientific and Technical services (54) is the industry with the highest labor share. Both findings are reasonable but these discrepancies are likely to lead to differences in the resulting estimates of TFP.

Table A.3: Estimates Cobb-Douglas production functions (by sector)

Estimation	OLS	OLS	Levinsohn-Petrin	Levinsohn-Petrin
Factor	Labor	Capital	Labor	Capital
Pooled	0.64	0.42	0.64	0.40
Sector				
11	0.56	0.47		
21	0.38	0.72		
22	0.47	0.48		
23	0.92	0.08	0.60	0.49
31	0.66	0.27	0.58	0.37
32	0.88	0.30	0.60	0.38
33	0.84	0.20	0.75	0.48
42	0.79	0.16	0.68	0.56
44	0.43	0.45	0.74	0.28
45	0.61	0.23	0.74	0.32
48	0.56	0.28		
49	0.75	0.12		
51	0.82	0.21	0.70	0.38
52	0.89	0.17	0.73	0.31
53	0.57	0.29		
54	0.88	0.15	0.92	0.25
55	0.64	0.42		
56	0.53	0.34	0.75	0.22
61	0.27	0.43		
62	0.71	0.20	0.72	0.28
71	0.48	0.39	0.52	0.59
72	0.69	0.28	0.70	0.26
81	0.53	0.24	0.66	0.47

Several observations are worth noting. First, the table has many gaps. This corresponds to the sectors for which the cost of materials is missing.⁵⁶ We fill all these gaps using the estimated coefficients obtained on the pooled sample, using the appropriate estimation method in each case.

A.4 Construction of TFP

On the basis of the sector-specific estimates for the coefficients of the production functions we now build firm-year TFP measures as a residual. Specifically, for the naive OLS method we do the following:⁵⁷

$$TFP_{ijt}^1 = \widehat{\varepsilon}_{ijt} = y_{ijt} - \hat{\beta}_k^j k_{ijt} - \hat{\beta}_\ell^j \ell_{ijt}. \quad (23)$$

For the Levinsohn-Petrin estimates we compute

$$TFP_{ijt}^2 = \widehat{\varepsilon}_{ijt} = v_{ijt} - \hat{\beta}_k^j k_{ijt} - \hat{\beta}_\ell^j \ell_{ijt}. \quad (24)$$

We next compute the 2001-2006 log difference for each of these measures, which is our focus on interest. The correlation coefficients among the log changes for these variables is around 0.40. The advantage of the first measure (naive estimation) is that we can compute it for all firms in the sample, spanning 24 sectors. The measures based on the Levinsohn-Petrin estimation have a stronger theoretical and econometric foundation. However, they are missing for about half of the firms in our final sample.

B Alternative measure of dependency on H-1B visas

The Tables that follow are part of the discussion in [Section 7](#). Briefly, we experiment with an alternative measure of dependency on H-1B visas. Specifically, we follow [Kerr and Lincoln \(2010\)](#) and use the number of LCAs filed in 2001 normalized by total employment in the firm, whereas, in our main analysis, we have used the total number of LCAs filed by a firm in year 2001.

⁵⁶For sector 55 the naive OLS estimation delivers a negative capital share, which does not make sense so we drop those estimates.

⁵⁷We do not include the estimated intercept in the construction of the TFP value for each firm-year observation. Since our focus is on *changes* in TFP over time for a given firm this will not affect our results.

Table B.1: Industry results - Alternative dependency measure

Dep. var. is $\Delta \ln$ of	(1) $\frac{Sales}{Emp}$	(2) Sales	(3) Emp	(4) Profits	(5) TFP^1	(6) $R\&D$
Sample: All firms (N=3,945)						
$\Delta H1B \times D\{0 \leq LCA01/Emp01 \leq p80\}$	-0.27 [0.25]	0.44 [0.31]	0.71*** [0.25]	0.71** [0.30]	-0.26 [0.23]	0.22 [0.49]
$\Delta H1B \times D\{p80 \leq LCA01/Emp01 \leq p90\}$	-0.11 [0.91]	0.11 [1.09]	0.21 [0.83]	0.51 [0.98]	-0.39 [0.92]	0.66 [1.05]
$\Delta H1B \times D\{LCA01/Emp01 \geq p90\}$	-0.76 [1.00]	-1.51 [1.22]	-0.75 [0.99]	-0.15 [1.49]	-0.45 [0.92]	-0.37 [1.24]
Sample: R&D (N=969)						
$\Delta H1B \times D\{0 \leq LCA01/Emp01 \leq p80\}$	0.06 [0.46]	0.53 [0.57]	0.47 [0.37]	0.51 [0.50]	0.23 [0.47]	0.26 [0.42]
$\Delta H1B \times D\{p80 \leq LCA01/Emp01 \leq p90\}$	0.15 [1.28]	0.74 [1.58]	0.59 [0.85]	1.10 [1.36]	-0.04 [1.30]	1.91** [0.97]
$\Delta H1B \times D\{LCA01/Emp01 \geq p90\}$	-1.50 [1.35]	-0.74 [1.61]	0.76 [1.15]	-2.97 [2.02]	-1.92 [1.46]	1.91 [1.33]
Sample: Manufacturing (N=1,549)						
$\Delta H1B \times D\{0 \leq LCA01/Emp01 \leq p80\}$	-0.55 [0.42]	0.26 [0.49]	0.82** [0.35]	1.02** [0.43]	-0.63 [0.42]	0.56 [0.44]
$\Delta H1B \times D\{p80 \leq LCA01/Emp01 \leq p90\}$	-1.56 [1.69]	0.41 [1.92]	1.97 [1.24]	2.22 [1.79]	-2.37 [1.73]	3.40*** [1.23]
$\Delta H1B \times D\{LCA01/Emp01 \geq p90\}$	-1.16 [1.67]	-0.72 [2.11]	0.43 [2.13]	-0.94 [2.93]	-1.64 [1.74]	1.53 [1.66]

Heteroskedasticity-robust standard errors in brackets

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

Notes: The omitted category are firms with zero LCAs in 2001. All regressions include 2-digit sector dummy variables. The first panel contains the whole sample of firms. The second panel is the R&D subsample. The third panel is the subsample of manufacturing sectors (NAICS 31-33).

Table B.2: Linear Model with Kerr-Lincoln sample - Alternative dependency measure

Dep. var. is $\Delta \ln$ of	(1) $\frac{Sales}{Emp}$	(2) Sales	(3) Emp	(4) Profits	(5) $R\&D$
$\Delta H1B \times LCA2001$	0.09 [0.08]	0.23 [0.35]	0.14 [0.32]	0.11 [0.21]	0.23 [0.39]
$\Delta H1B \times LCA2001/Emp2001$	4.03 [8.44]	-8.02 [23.00]	-12.05 [20.92]	3.31 [31.02]	-13.46 [24.47]
Observations	99	99	99	98	99

Notes: All regressions include 2-digit sector dummy variables. The sample contains firms with at least 60 LCAs in year 2001 with R&D expenditures in both 2001 and 2006 of at least \$5,000.