

Foreign and Native Skilled Workers: What Can We Learn from H-1B Lotteries?*

Giovanni Peri (University of California, Davis)

Kevin Shih (Rensselaer Polytechnic Institute) Chad Sparber (Colgate University)

May 8, 2015

Abstract

In April of 2007 and 2008, the U.S. randomly allocated 65,000 H-1B temporary work permits to foreign-born skilled workers. About 88,000 requests for computer-related H-1B permits were declined in each of those two years. This paper exploits random H-1B variation across U.S. cities to analyze how these supply shocks affected labor market outcomes for computer-related workers. We find that negative H-1B supply shocks are robustly associated with declines in foreign-born computer-related employment, while native-born computer employment either falls or remains constant. Most of the correlation between H-1B supply shocks and foreign employment is due to rationing that varies with a city's initial dependence upon H-1B workers. Variation in random, lottery-driven, unexpected shocks is too small to identify significant effects on foreign employment in the full sample of cities. However, we do find that random rationing affects foreign employment in cities that are highly dependent upon the H-1B program. Altogether, the results support the existence of complementarities between native and foreign-born H-1B computer workers.

Key Words: Skilled Workers, H-1B Work Permit, Computer-Related Occupations, Natural Experiment.

JEL Codes: J61, F22, O33, R10

*Address: Giovanni Peri, Department of Economics, UC Davis, One Shields Avenue, Davis, CA 95616, gperi@ucdavis.edu; Kevin Shih, Department of Economics, UC Davis, One Shields Avenue, Davis, CA 95616, kyshih@ucdavis.edu; Chad Sparber, Department of Economics, Colgate University, 13 Oak Drive, Hamilton, NY 13346, csparber@colgate.edu.

1 Introduction

Access to foreign talent has long been considered one of the key ingredients for U.S. scientific and technological success. Many studies have noted that U.S. global leadership in science and technology heavily depends upon foreign-born workers who are responsible for a large share of U.S. innovation (measured in Nobel laureates, patents, or journal articles, for example).¹ But do these foreign workers crowd out natives from desirable jobs and push them into lower paid ones? Mixed results have been presented on this topic. While some recent papers emphasize that foreign skilled workers have improved macroeconomic productivity and the labor market opportunities for natives (Peri, Shih, and Sparber, 2015; Moser, Voena, and Waldinger, 2014) others have argued that high-skilled foreign labor has displaced native-born Americans and reduced their opportunities (Borjas and Doran, 2012).

This paper explores the short-term effects of high-skilled foreign labor flows on labor market outcomes for college-educated natives at the metropolitan area level. We exclusively focus on computer-related workers² and the labor market effects generated by the H-1B program – the primary channel through which computer-related foreign-born workers enter the United States each year. The program is popular among employers and many foreign-workers, but is opposed by individuals concerned that it might harm the job prospects of native-born workers, especially those in computer-related occupations (e.g. Matloff, 2008; Hira, 2007).³ This paper focuses on this issue and makes important progress along two dimensions.

First, our identification strategy uses random variation in the allocation of H-1B workers across metropolitan areas in the U.S. resulting from the H-1B lotteries of 2007 and 2008. The number of new H-1B applications filed in the first month of acceptance far exceeded the 65,000 annual limit of available permits in those years. The U.S. government responded by holding a lottery to allocate permits. Applications submitted by April 3, 2007 and April 7, 2008 were put in a common pool during their respective years. In each year, 65,000 applicants were randomly selected. This produced a random shock to the supply of H-1B workers – some cities were successful in the lottery and received many of the workers they requested, while other cities

¹See for instance Kerr and Lincoln (2010), Hunt and Gauthier-Loiselle (2010), and Peri (2007).

²“Computer-related occupations” are those defined as such by the DOT classification found in the Labor Certification Applications for the H-1B program. They include Census occupations Computer-System Analyts, Computer Scientists, and Software Developers and Programmers of Numerically Controlled machines (IPUMS occ1990 codes 64, 229, and 233).

³See Peri et al. (2014) for a similar discussion in the non-academic literature.

received fewer workers than anticipated. Our methodology exploits this lottery-induced shock.

Second, we recognize that the effects from H-1B restrictions operate through two channels. One is an “average rationing” effect. Cities that are heavily dependent upon foreign-workers might feel the effects of national-level rationing more acutely than those hiring a small percentage of foreign employees. That is, cities in which foreign workers represent a large share of employment will have a large potential flow of workers who are immediately affected by changes in federal policy. The other channel is an “unexpected shock” effect above and beyond this average rationing channel. That is, cities might expect to receive a specific number of foreign-born workers given national rationing rates and city-level dependence upon the H-1B program, but the actual number of H-1B workers deviated from this expectation for random reasons during lottery years.

What makes our paper’s innovations particularly important is that they help to reconcile results from the previous literature in three important ways. First, the majority of scholarly articles analyzing the impact of the H-1B program on employment and wages (Peri, Shih, and Sparber, 2015), innovation (Kerr and Lincoln, 2010), or firm sales and total factor productivity (Ghosh, Mayda, and Ortega 2014) tend to identify the consequences of the first channel’s variation in foreign dependency. They rely on pre-existing variation in demand for foreign skilled workers across labor markets (metropolitan areas) or firms interacted with aggregate flows of skilled immigrants affected by changes in policy (the cap of new H-1B issuances allowed). An unavoidable objection to this approach is that large foreign-dependence could be systematically associated with other features affecting productivity and labor demand. The presence of persistent unobserved demand shocks for highly skilled workers, possibly correlated with dependence on foreign-born workers, could generate a positive spurious correlation between outcomes of foreign and native-born skilled workers that bias the estimates obtained with this method. This limitation therefore muddles clean identification of the effects of policy restrictions. By instead relying upon random variation generated by an H-1B lottery, we come closer to using a truly exogenous source of variation. Importantly, we largely support previous work by finding that negative H-1B shocks are associated with native and foreign-born employment losses.

Second, we improve upon past methods by identifying the impact of the unexpected supply shocks. Such shocks exhibit very small variance at the city level, thus making it difficult to identify effects from this

channel. However, following previous literature that emphasizes how only highly dependent firms and labor markets are affected by H-1B policy variation, we also explore heterogeneity of effects across cities. We find that unexpected losses in H-1B workers reduce foreign and native employment for cities highly dependent upon H-1B workers, but not for the full sample of cities. This result complements work by Ghosh, Mayda and Ortega (2014) who find that H-1B restrictions negatively affect outcomes for firms that are in the top quintile of H-1B employers.

Third, our results are helpful in understanding important related work by Doran, Gelber, and Isen (2014), who also exploit random H-1B variation. Our framework is based on data from a broad cross-section of cities during fiscal years (FYs) 2008 and 2009 when all cap-dependent H-1B permits were distributed by lottery. This ensures external validity by including all firms relying upon the H-1B program in those years. Doran, Gelber, and Isen (2014) exclusively rely on the small cross-section of firms during FYs 2006 and 2007 whose desired employees submitted an H-1B application on the day the cap was reached. Only a small fraction of H-1B permits (about 4-5%) were distributed by lottery in those years. While the authors' firm-level data provide an excellent and important step toward using exogenous data variation, the reliance upon a subset of H-1B participants, small shocks at the firm level, weak power in the impact of shock on foreign-labor supply, and external validity challenges are significant limitations of their approach.

Our analysis is based on a simple empirical model. We regress city-level employment and wage growth rates for foreign and native-born, college-educated, computer-related workers on negative H-1B supply shocks experienced by those cities. The outcome variables are taken from the Census and American Community Survey (ACS). We calculate these labor market outcomes in years 2005 (pre-lottery) through 2011 (3 years after the second lottery). The more innovative aspect of our analysis lies in the calculation of the explanatory variable. It is a measure of excess demand for foreign labor that is driven by exogenous supply shocks, and it is calculated as the difference between the desired number of new H-1B workers and the lottery allocation of permits. This requires the merger of two datasets. The first source comes from publicly-available Labor Condition Applications (LCAs). Firms wishing to hire an H-1B worker must first file an LCA stating the nature of the work. Since this data identifies job vacancies directed by employers to foreign workers, it serves as a measure of foreign labor (H-1B) demand. Foreign individuals who then go on to apply for a

permit must submit an I-129 form. We acquired information from I-129 forms from United States Citizen and Immigration Services (USCIS) through a Freedom of Information Act (FOIA) request. Only the I-129 records for lottery winners – those randomly selected for processing – are available. Thus, they serve as our measure of capped labor supply. We aggregate these supply and demand figures at the city level and calculate the negative H-1B supply shock as the number of unfilled jobs that were subject to the H-1B quota.

To establish the random and exogenous nature of the explanatory variable, we regress foreign-born employment growth between 2000 and 2005/06 (before the lottery) on the negative H-1B shock. We find no statistical evidence for a correlation between our explanatory variable and this pre-policy outcome, but we must caution that point estimates and standard errors are large. We then proceed by examining whether the net growth of foreign-born employment between 2005/06 and 2008/09, 2009/10, or 2010/11 (one, two, or three years after the lottery) was affected by the random negative H-1B supply shock. The presence of a significant correlation confirms that labor markets that were randomly denied the number of desired foreign workers experienced reduced foreign employment growth. The randomness of the lottery facilitates causal interpretation of our regression estimates. For a negative H-1B supply shock equal to one percentage point of initial computer-related employment in the city, the growth rate of college educated foreign-born computer workers was between 0.5 and 2 percentage points slower in that metropolitan area three to five years later. Equivalently, this translates to about 0.7 foreign jobs lost for every unfilled H-1B position.

Our attention then turns to the impact of the H-1B supply shocks on the employment and wage growth rates of college-educated, native-born, computer-related workers. Wage effects are usually small and insignificant. However, we also find that college-educated, computer-related, natives working in cities with larger negative H-1B shocks exhibit slower employment growth. The effects are estimated with some noise, but the point estimate is always negative and usually significant: a random negative H-1B shock of 1 percentage point of initial employment results in an approximate 0.5 percentage point decline in the native college-educated employment growth rate. That is, about 0.2 jobs for college-educated native-born computer workers were lost for every unfilled H-1B position.

Interestingly and importantly, however, this result appears to be driven by differences in average rationing that arises from the interaction of national-level policy and city-level H-1B dependency. The random unex-

pected shock component of H-1B supply across cities is much smaller and its effects are much less precisely estimated. Only by restricting the analysis to cities that have the highest dependence on the H-1B program do we see that random, unexpected, negative H-1B supply shocks reduce foreign-born employment growth and, in turn, might cause a mild decline in the employment growth rate of native-born, college-educated, computer workers.

Taken together, our results contradict the idea that by reducing H-1B permits in a labor market, employers would increase their demand for native skilled workers and/or the wages paid to them. To the contrary, the results suggest that foreign workers are complementary to native workers and/or stimulate the productivity of natives in computer-related occupations. One might speculate about the mechanism at work: H-1B rationing might reduce firm expansion; firms might have to cancel projects when a sufficient number of team members are unavailable; productivity growth might suffer due to diminished skill complementarity. In any case, we find no evidence that H-1B workers displace native-born computer-related workers in a local labor market.

The rest of the paper is organized as follows. Section 2 describes the H-1B program and the 2007 and 2008 lotteries. Section 3 provides a simple framework to understand the effect of negative random supply shocks of foreign workers on the demand for native workers. Section 5 describes the data, identification, and power of the lottery in affecting supply. Section 6 presents our empirical framework and results. Section 7 discusses the advantages and limits of our approach and relates our findings to other recent studies in this area. Section 8 translates our results into foreign and native jobs lost due to restrictive H-1B limits. Section 9 concludes.

2 The H-1B Program and the 2007 and 2008 Lotteries

The H-1B program provides temporary permits to foreign-born skilled-workers in specialty-occupations wishing to work in the United States.⁴ The permit can be approved for up to a three-year period and can be renewed to a maximum of six total years. Five aspects of the H-1B program are important in the context of this paper.

⁴The program has been studied by many papers in the immigration literature. See program descriptions in Clemens (2010), Doran, Gelber, and Isen (2014), Kato and Sparber (2013), Kerr and Lincoln (2010), Peri, Shih, and Sparber (2015), and Peri et al. (2014).

First, the work permit has been subject to a cap (or quota) since its inception in 1990. Initially set at 65,000 new issuances per fiscal year, the quota was not reached until 1997 and again in 1998. Congress responded by raising the cap to 115,000 for 1999 and 2000, then to 195,000 for 2001 through 2003. 2001 also marked the advent of a policy exempting employees of higher-education, non-profit, and government-research organizations from the quota. Trends in increasing permit availability reversed in 2004 when the cap reverted to 65,000, though 20,000 additional permits were granted for workers who earned a master's degree or Ph.D. from a U.S. institution. The H-1B quota on new H-1B issuances to people working at for-profit firms has been reached before the end of each fiscal year ever since.

Second, Table 1 demonstrates that the final receipt date – the date on which USCIS has received more H-1B applications than are available through the quota – occurred increasingly quickly from 2004 through 2009. Applications are accepted beginning on April 1, six months prior to the October 1 start of the federal fiscal year. The cap for FY 2004 was reached almost a year after the first date of acceptance. For FY 2007, the final receipt date was reached after two months. In 2007 and 2008, USCIS received more than 65,000 cap-bound applications within the first week of eligibility for FYs 2008 and 2009.

This popularity is important for the third relevant characteristic of the H-1B program. In principle, new H-1B permits for employees of for-profit firms are distributed on a first-come first-served basis. Then all applications received on the final receipt date are processed randomly until the quota has been met. Doran, Gelber, and Isen (2014) exploit the randomness of the applications received on the August 10, 2005 and May 26, 2006 final receipt dates in their firm-level analysis. However, the speed with which people applied for H-1Bs in April 2007 and 2008 led USCIS to distribute all cap-bound permits by lottery for fiscal years 2008 and 2009. Our methodology effectively uses the randomness in H-1B distributions that occurred in these two years.

To exploit this randomness, we must rely on a fourth characteristic of the H-1B program: Trends in program popularity changed firm behavior. A firm wishing to hire an H-1B worker must first file an LCA with the department of labor describing the nature of the job. The form includes information on the prevailing wage of the occupation, the wage to be paid to the prospective worker, where the work will occur, and the beginning and end dates of the position. A worker wishing to secure an H-1B must submit an I-129 form

that is accompanied by an approved LCA. The permit, if approved, belongs to an employee. However, it is only valid to work for the firm and location specified on the forms.

The rising popularity of the H-1B program altered firm behavior. Since an LCA cannot be filed more than six months prior to intended employment, a firm wishing to hire a worker to begin work at the start of the fiscal year might be expected to file an LCA on April 1 at the earliest. However, trends in H-1B scarcity led firms to “predate” LCAs in 2007 and 2008. That is, firms filed LCAs in January through April that specified a start date five to six months in the future, even though they knew that new H-1B workers would only be eligible to work beginning on October 1. This process is advantageous in that all desired H-1B workers would have an approved LCA in hand when applying for an H-1B on April 1. The disadvantage is that the firm would lose one or more months of job approval at the end of the desired work period since LCAs are applicable only up to a three year period.

Pre-dating behavior is important for our methodology because it allows us to measure demand for foreign-born skilled labor. There is no limit on the number of LCAs that a firm can file. Firms can file LCAs throughout the course of the year because they can use approved LCAs to hire existing H-1B workers who are exempt from quotas. Pre-dated LCAs, however, provide a sense of the demand for new H-1B workers in lottery years. Figure 1 illustrates this behavior. Non-profit firms, whose employees are exempt from H-1B caps, file roughly 13% of their LCAs in January through April with work start dates occurring five to six months after LCA submission. This figure is constant from 2002 through 2011. For-profit firms, in contrast, almost never adopted this behavior when H-1B permits were plentiful. As they became scarcer, firms began pre-dating at such extreme levels that more than 40% of LCAs filed by for-profit firms were pre-dated in 2007 and 2008. Pre-dating was the only way for such firms to hire new H-1B workers in these years. As a result, USCIS received roughly 120,000 and 160,000 cap-dependent I-129 H-1B applications in the first weeks of acceptance in 2007 and 2008, respectively.⁵ USCIS randomly selected applications for processing until the 65,000 available permits were allocated. According to Clemens (2010), records for unselected applications do not exist for these years.

The number of H-1B applications in April 2007 and 2008 was roughly double the number of available

⁵ See “USCIS Updates Count of FY 2008 H-1B Cap Filings,” US Citizenship and Immigration Services, April 10, 2007; and “USCIS Runs Random Selection Process for H-1B Petitions,” US Citizenship and Immigration Services, March 31, 2008 (date on press release is clearly wrong).

permits. There is no data on individuals who would have applied for an H-1B later in the year in the absence of a quota. As noted, USCIS responded to rising H-1B popularity by randomly distributing permits among individuals who applied in April. This lottery generated a random negative shock in the supply of workers to firms and metropolitan areas. Our methodology examines the short term employment and wage consequences for labor markets that randomly received rationing shocks.

The fifth and final characteristic of the H-1B program relevant to our study pertains to the occupational composition of the program displayed in Figure 2. Although there is no statutory preferential treatment for any particular occupation among the specialty-occupations recognized by the program, two thirds of cap-bound permits were for computer-related workers in 2007 and 2008. This is true whether one takes a firm's perspective (pre-dated LCAs filed by for-profit firms) or an individual's (approved I-129 H-1B applications that count toward the quota). This is particularly striking when considering that computer-related workers accounted for just 3% of the urban workforce in 2005. Overall, the importance of the H-1B program for these occupations encourages us to focus exclusively on computer-related work in our analysis. Note, however, that trends indicate a massive rise in the popularity of the program for computer-related work leading up to lottery years. This implies that we need to be careful about considering pre-trend behavior in our analysis. Before thoroughly describing the data and conducting that analysis, however, we begin with a simple supply and demand framework to provide context for the empirical model.

3 The Effects of a Negative Foreign Labor Supply Shock: A Framework

Our empirical analysis organizes data into metropolitan area cells and focuses on computer-related occupations – by far the largest group of H-1B applications. While we can generally think of foreign-born labor supply as constrained in the short-term by the limits and quotas imposed by national policy, different labor markets experience heterogeneous constraints due to their varied degree of reliance upon foreign workers and, in 2007 and 2008, due to the randomness produced by the lotteries across labor markets.⁶

⁶Note that firms experience heterogeneous constraints much like cities do, and would provide an interesting alternative level of analysis. One difference is that firms have a greater opportunity to respond to negative H-1B shocks by hiring other foreign-born workers. For example, if a single firm receives a negative H-1B supply shock but the local market has a relatively large supply

Let us illustrate the potential effects of a negative supply shock on the wages and employment of native workers in the same labor market (computer occupations in a metropolitan area) under two different sets of assumptions. In the first case we assume that foreign and native computer workers are perfect substitutes. Foreign workers have no effect on firm total factor productivity, but they are hired before native workers due to lower cost per unit of productivity. In the second case we assume that foreign computer workers are possibly complementary to native workers and/or that their presence has a positive productivity effect on firms. These two effects are strong enough that increased foreign employment raises native productivity.

Figure 3 illustrates the first case. It represents the demand and supply for native (left panel) and foreign-born (right panel) computer workers in one market (metropolitan area). The two types of workers are perfectly substitutable, and hence the demand for natives shifts one-for-one with foreign labor supply. In particular, the vertical axis measures wages per unit of productivity for each type of worker. The demand and supply for foreign computer workers are represented in the right panel and they cross at the point F_0 . Foreign-born supply is given by the maximum number of H-1Bs that can be obtained in that market. The demand for native computer workers is obtained as the residual demand for a perfectly substitutable foreign worker. Hence at the wage per unit of productivity corresponding to the point F_0 , demand for natives will be zero (reported on the left panel) and will increase as the wage decreases. The equilibrium employment and wage for native workers will be at N_0 – the crossing point with the native supply of computer workers – with native wages per unit of productivity equal to w_N^0 and employment of natives at E_N^0 .

Now consider an identical market in which lottery rationing denies D H-1B workers (permits). That market will have a foreign labor supply equal to $F - D$ and the equilibrium would be at F_1 . This equilibrium implies a corresponding shift up by the amount D of the residual demand for native computer workers, which is now represented by the line DN_1 . The shock implies higher native employment and wages if the labor supply is upward sloping, or only higher native wages if the labor supply is vertical due to tight labor markets close to full employment (as was possibly the case in 2007). Figure 3 implies that the larger is the number of denied permits (D) relative to unconstrained supply, the larger native employment ($E_N^1 - E_N^0$) and wages

of previously hired H-1B workers (or is less constrained by the shock overall), then a firm could respond to lottery losses by hiring existing foreign workers from other firms. In contrast, if a whole local labor market is restricted then it would represent a complete reduction in the local supply of foreign computer-related workers. We focus our attention on this latter city-wide shock.

$(W_N^1 - W_N^0)$ would be relative to a market that did not receive rationing. This stems from the fact that the demand for natives is a residual from the demand for foreign labor in the case of perfect substitution.

A very different scenario is depicted in Figure 4. In this case foreign and native computer workers are not perfect substitutes. Instead, they are differentiated types of labor (different inputs in production) with the native-born labor market characterized in the left panel and the foreign-born labor market depicted in the right panel. In this case the demand for native computer workers, $DN_1(E_F)$, is a function of foreign employment (E_F) but not in a residual way. In particular, it could be an increasing function of foreign labor either because the two workers are complementary or because foreign workers increase the overall productivity of firms and their native-born employees. Alternatively, if native and foreign-born computer workers are neither complements nor substitutes but are instead segmented types of workers, a change in foreign employment might have no effect on native-born employment (no native demand shift).

Figure 4 illustrates the case of complementarity – a negative foreign labor supply shock equal to D produces a decline (left shift) in the demand for natives and possibly in their wages. The equilibrium for labor markets with larger negative demand shocks (D) implies lower native-born employment and wages.

Differences in underlying assumptions cause theoretical models to generate opposing predictions for the native wage and employment consequences of foreign labor supply shocks. However, observed H-1B lottery supply shocks vary substantially across metropolitan areas, and this allows us to empirically identify whether foreign skilled workers substitute for natives or complement them. In the former case, native workers promptly replace foreign ones when H-1B rationing occurs. In the latter, foreign labor enhances the productivity of natives so that rationed H-1B supply does not help, but potentially hurts, native employment and wages.

4 Data and Measuring the H-1B Negative Supply Shock

Our data are from multiple sources. H-1B lottery information comes from 2007 and 2008 LCA and I-129 data. Wage and employment information comes from the Census and American Community Surveys from 2000 through 2011.

Our main explanatory variable is a measure of excess demand for H-1B workers driven by exogenous

supply shocks. Measures of firm demand for foreign labor come from U.S. Department of Labor LCA data on employer name, location (city and zip code), occupation, number of H-1B employees requested, the filing date, work start date, and work end date by year. As noted above, we infer whether the firm is trying to hire a cap-dependent worker in 2007 and 2008 from these data by accounting for the pre-dating behavior. Namely, we assume that LCAs filed by for-profit firms in January through April with a work start date five to six months in the future represent demand for cap-dependent H-1B workers for the following fiscal year.

Measures of realized H-1B labor supply come from I-129 H-1B applications, which we acquired from the USCIS through a FOIA. Available information includes employer name, location (city and zip code), occupation, and several individual characteristics of the applicant. Forms also provide more explicit indicators for whether a person would count toward the H-1B quota. First, we limit our study to individuals who file I-129 petitions in April, and assume that the applicants are to begin work in the following fiscal year (i.e., in October). Second, we classify cap-exempt workers as those who work for colleges, universities, and other non-profit firms; applicants for H-1B extensions, renewals, and job changes; and citizens of countries effectively exempt from H-1B limits due to special bilateral trade agreements (Australians, Canadians, Chileans, Mexicans, and Singaporeans).

We compute city-level (c) excess demand for foreign labor (D_c) by first removing H-1B workers that were not subject to the cap (X_c) from the total LCAs filed (L_c), and then subtracting the cap-dependent H-1Bs awarded to the city (H_c). We combine data for 2007 and 2008 to create a single cumulative measure of excess demand for 238 cities whose scale and variation combine those two lotteries.⁷ Importantly, this excess demand measure varies across cities due to supply shocks caused by different levels of H-1B rationing and the randomness of the lottery.

Clearly, an equivalent number of H-1B permits denied to two cities will represent a much larger shock for small markets than for large ones. In order to better capture the impact and magnitude of this shock on the local market, we scale the absolute measure of excess demand (i.e., the number of jobs unfilled due to rationing) by pre-lottery (2005/06) total computer-related employment ($Empl_{c,2005/06}$; we average the 2005/06 data in the ACS to have a larger and more reliable sample). Hence, our key explanatory variable

⁷As noted, we do this only for computer-related occupations given the importance of the H-1B program for this type of work. Unfortunately, extending the analysis to other fields results in small occupation-by-metropolitan area cells with large measurement error, so we elect not to pursue that avenue.

– defined in Equation (1) – measures normalized excess demand (d_c), or equivalently, the negative supply shock as percentage of initial employment.

$$d_c = \frac{D_c}{Empl_{c,2005/06}} = \frac{L_c - X_c - H_c}{Empl_{c,2005/06}} \quad (1)$$

These normalized negative supply shocks vary significantly across cities, with a standard deviation equal to 0.09 (i.e., 9 percentage points of initial employment). The d_c term captures the total rationing effect of the lottery on different labor markets. Later sections will decompose it into two components: One is linked to average rationing, which varies across cities due to heterogenous dependence upon foregone workers. That is, since cities maintain different levels of foreign labor demand (per unit of employment) the common, national-policy-driven, percentage of job rationing will generate heterogeneous effects across cities. The other component is due to random, unexpected, supply variation attributable to the lottery realization alone.

The dependent variables in our regressions use the 2000 5% Census and 2005-2011 1% American Community Surveys to construct the growth in employment and wages for computer-related workers in a metropolitan area. Employment changes are calculated between 2005/06 and periods one, two, and three years past the lottery (2008/09, 2009/10, and 2010/11, respectively).⁸ We select civilian employees age 18-65 who do not live in group quarters, are not enrolled in school, and reside in metropolitan areas. Only those who worked in the previous year, earned positive wages, and are currently employed full time (more than 39 weeks per year and 34 hours worked per week) are included. By creating this average wage and employment growth at the metropolitan area level and merging it with our LCA and I-129 data, we can assess whether the random negative employment shocks to local labor markets affected outcomes for native and foreign-born workers in those markets.

5 Identification and Power

The primary interest of this paper is in analyzing the effect of foreign-labor shocks on the employment and wage outcomes of native-born, college-educated, computer-related workers. In order for the results to

⁸We merge two years into single periods in order to create a more representative and larger sample of cities.

be interpreted as reasonably causal, we must first establish that our H-1B rationing variable is random and/or unrelated to other shocks that could affect the dependent variables. We turn to both descriptive and analytical data for this task.

One concern is that firms applying early in the H-1B cycle could be fundamentally different from other firms. If, for example, the only firms (and cities) participating in the lottery were those that faced strong economic pressure to act early, then the model will suffer from a selection bias. The 2007 fiscal year offers suggestive evidence both for and against this possibility. The H-1B filing period in that year lasted 56 days (from April 1 through May 26, 2006). Figure 5 illustrates that in the beginning of the filing period, roughly 55% of the cap-bound H-1B permits were awarded to computer-related workers while the rest went to individuals in other occupations. This percentage rises during the first month of the application period and stabilizes around 75-80% until the quota was met. If firms wishing to hire computer-related workers are slower to submit LCAs and make job offers, then the very short filing period for FYs 2008 and 2009 could create a further rationing of those jobs. However the very prominent pre-dating behavior documented above implies that most companies, and certainly those most dependent on the H-1B program, were getting ready to apply earlier due to the shorter period of availability.

Figure 6, on the other hand shows that the type of workers that firms hired, as measured by annual salary, did not change much over approval date. This figure displays the median wage paid to approved computer-related H-1B workers subject to the FY 2007 cap for each day of the filing period. Although a large amount of dispersion is present and the median wage was somewhat higher for those workers hired on the first day of the eligibility, the time profile appears relatively flat throughout the filing period. Moreover, the (nominal) median wage for such workers was \$55,000 in FY 2007. This compares closely to the \$55,000 and \$60,000 (nominal) median wages in FYs 2008 and 2009, suggesting that productivity of employees and employers was comparable across years.

Regressions can provide more rigorous identification tests. One important check involves testing pre-lottery trends. If the negative supply shock is truly exogenous, it should be uncorrelated with the pre-lottery growth of foreign-born college-educated workers in computer-related occupations. Consider the following

specification.

$$\% \Delta \text{Empl}_{c,2000-2005/06}^{\text{foreign-computer-college}} = \beta_0 + \beta_{\text{pre-lottery}} \cdot d_c + \varepsilon_c \quad (2)$$

The explanatory variable is the normalized negative supply shock in a metropolitan area for computer-related occupations (d_c) as defined in (1) and derived from the LCA and I-129 data. The dependent variable is the growth rate of college-educated, foreign-born, computer-related employment in a metropolitan area between 2000 and 2005/06 – that is, employment growth before the lottery years.

Panel A of Table 2 displays estimates of $\beta_{\text{pre-lottery}}$. All regressions throughout this paper are estimated with heteroskedasticity-robust standard errors. Columns are differentiated by weighting conventions. Regressions are unweighted in the first column. The second column weights observations by H-1B dependence (specifically, the number of LCAs filed in a city relative to its 2005/06 total computer-related employment). Column 3 weights observations by the logarithm of this dependence.⁹ The results are encouraging in that they do not find statistical evidence for a correlation between the random H-1B rationing supply shock and pre-lottery foreign-employment growth. Though the point estimates are insignificant, two of them are quite large with large standard errors, indicating some noise in the data.¹⁰

While the regressions of Panel A provide evidence for exogeneity, the regressions in Panel B explore the power of the lottery-based H-1B variable in explaining foreign-born employment growth as specified in equation (3). This model mimics that of equation (2) but replaces the dependent variable with foreign-born, college-educated computer-related, employment growth between 2005/06 and period t (combined years 2008/09, 2009/10, and 2010/11). We also include the lagged value of the dependent variable as a regressor – as a way of controlling for potential lingering correlation with pre-lottery trends across cities.

$$\% \Delta \text{Empl}_{c,2005/06-t}^{\text{foreign-college}} = \beta_0 + \beta_1 \cdot d_c + \beta_2 \cdot \% \Delta \text{Empl}_{c,2000-2005/06}^{\text{foreign-college}} + \varepsilon_c \quad (3)$$

This model serves to verify the power of the lottery by assessing whether the variation in the lottery

⁹The choice to weight observations is motivated by a desire to mitigate problems associated with potential measurement error in LCA data. Logged H-1B dependence weights are less susceptible to influence from potential outliers.

¹⁰Throughout this paper, we rely on regressions of pre-lottery foreign employment growth to examine the exogeneity of our H-1B shock variable. The appendix discusses pre-lottery native employment growth regressions, which do find some correlation that is driven by the dependency channel. This correlation does not appear among highly rationed cities. While this is problematic for a fully causal interpretation of regression results, we will alleviate the problem by including pre-lottery growth variables in all regressions that examine the main outcomes of interest.

results actually corresponds to a change in post-lottery college-educated, foreign-born, computer-related employment growth within a city. Foreign-born substitutes for new H-1B workers exist, so it is possible for firms to respond to H-1B denials by seeking alternative routes for hiring foreign labor. These might include attracting existing H-1B workers from other labor markets, hiring recent foreign graduates from U.S. colleges through the Optional Practical Training (OPT) program, bypassing the quota by sub-contracting to a local non-profit lab, or hiring foreign-born workers with L-visas (provided the company has an international presence). Hence, equation (3) provides a valuable check on whether the H-1B negative supply shocks are significantly and negatively correlated with the post lottery growth of foreign-born college-educated workers in computer-related occupations across cities. If this correlation is not present, then the policy instrument has no power in moving the actual supply of foreign computer-related workers.

Notice also that if an unobservable positive demand shock in a city increases demand for H-1B workers, other foreign workers, and native workers, it would produce a positive bias in the coefficient β_1 in (3) and erroneously argue for substitutability between foreign and native workers.¹¹ This is because increased demand in January-April of the lottery year could raise the average size of rationing d_c (by increasing the variable L_c) and increase native and foreign-born employment after the lottery. In contrast, if the main variation of d_c is driven by supply changes then we would observe a negative value of β_1 as increased H-1B rationing would slow foreign employment growth.

Panel B of Table 2 shows that β_1 is significantly negative. Regressions are based on employment growth one, two, and three years after the lottery. Estimates are presented for the three aforementioned weighting schemes. The point estimate for each regression is negative, and it is significant in all but one case. A negative H-1B shock in a city equal to one percentage point of employment caused a reduction in the college-educated, computer-related, foreign employment growth rate in that city between 0.5 and 2 percentage points. The effect is smallest when weighting cells by H-1B dependence, but every specification demonstrates a significant slowdown in foreign employment growth when cities experience a larger negative shock in the H-1B lottery. Also, the negative effect seems to exhibit some reasonable non-permanent persistence: Two and three years after the lottery, cities that lost the lottery still show a lower than normal cumulative foreign-born

¹¹The same is also true if d_c proxies for the unattractiveness of a city to foreign workers (i.e., the inability to hire such workers). This would imply a negative correlation between d_c and the foreign-born share in pre-lottery years, and no correlation after the lottery.

employment growth.

We recognize that while H-1B allocation was entirely random, the construction of our normalized excess demand variable is subject to criticism that omitted factors could be correlated with $L_c - X_c$. This would call into question the true randomness of the explanatory variable. Despite these concerns, however, we show that our normalized excess demand variable is reasonably exogenous and we control for pre-lottery trends. Overall our estimates suggest that the size of a city’s H-1B shock exhibits no statistical correlation with foreign-born hiring growth before the lottery. In contrast, foreign-born computer-related employment grew at much slower rates after the lottery in cities receiving larger negative H-1B shocks. We now consider whether these negative H-1B shocks (and slower foreign-born employment growth) were associated with improved employment and wage growth for natives.

6 Empirical Specification and Results

We estimate equation (4) to analyze whether labor market outcomes for college-educated, native-born, computer workers were affected by lottery-driven supply shocks (d_c).

$$\% \Delta y_{c,2005/06-t}^{native} = \beta_0 + \beta_1 \cdot d_c + \beta_2 \cdot \% \Delta y_{c,2000-2005/06}^{native} + \varepsilon_c \quad (4)$$

The dependent variable $\Delta y_{c,2005/06-t}^{native}$ uses ACS data to measure the growth rate of outcome y (alternatively employment or wages) for college-educated, native-born, computer-related workers in metropolitan area c . Growth rates are calculated between the base period of 2005/06 and period t (alternatively 2008/09, 2009/10, and 2010/11). Our main explanatory variable, d_c , measures excess demand for H-1B workers as a percentage of pre-lottery total computer-related employment. Its value is driven by exogenous supply shocks, and it can equivalently be described as the negative supply shock as a percentage of initial employment in a metropolitan area for computer-related occupations, or as the total supply shock due to H-1B rationing. The above analysis argued that it is reasonable to assume that d_c is orthogonal to other unobserved determinants of $\% \Delta y_{c,2005/06-t}$ captured by the random measurement error term ε_c . However, our regression also includes the pre-lottery trend of the dependent variable, $\% \Delta y_{c,2000-2005/06}^{native}$ to control for possible persis-

tence arising from past local labor market conditions. Under these assumptions, the Ordinary Least Squares (OLS) estimates of the coefficient β_1 consistently identify the causal effect of the negative supply shock on native outcomes. To account for the fact that the variance of the outcomes and the exogenous regressor can be different in small and large cells, we use robust standard errors and, in some specifications, we weight observations by the size of cities' H-1B dependence.

6.1 Baseline Effects on Native Employment and Wages

Table 3 shows the estimated effect (β_1) of the negative H-1B supply shock on the employment and wage growth rate of college-educated, native-born, computer-related workers in a city. Panel A reports employment outcomes, while Panel B describes wage consequences. Different rows display the lottery impact on the cumulative growth of the dependent variable since pre-lottery period 2005/06 up to years 2008/09 (i.e., one year after the lotteries), 2009/10 (two years), and 2010/11 (three years). Columns are distinguished by alternative weighting mechanisms.

The top panel shows evidence that negative H-1B shocks reduce native employment. A random negative shock to H-1B supply equal to 1 percentage point of initial employment is associated with an employment growth rate reduction between 0.4 and 0.6 percentage points – an estimate between one third and one half of the employment losses for foreign-born workers. No specification finds a positive and statistically significant response to the negative random H-1B shock. All point estimates are negative. Hence these results are not consistent with the hypothesis that H-1B workers displace similar native workers. Rather, the estimates are consistent with the idea that cities that randomly won the opportunity to hire many H-1B workers responded by hiring more native workers as well, potentially because they could guarantee a certain work flow, undertake scheduled projects, or ensure that they could employ workers in specific jobs.

Wage results are in Panel B of Table 3. Weighted regressions uncover no evidence that larger negative H-1B supply shocks are associated with significant changes in the wage growth rate one, two, and three years after the lottery. Negative point estimates are found in unweighted regressions, but the coefficients are small and barely significant. Altogether, negative H-1B shocks seem to have no significant effect on native wages in the subsequent one to three years.¹²

¹²Peri, Shih, and Sparber (2015) uncover a positive relation between H-1B workers and native wages in the long run.

6.2 Decomposing H-1B Shocks

We now turn our attention to the important task of sharpening our identification and separating our explanatory variable d_c into two parts: one that is more strictly random and another that is driven by the interplay between average national-level lottery rationing and city-level reliance upon H-1B workers.

At the heart of the matter is that two different mechanisms generate variation in normalized excess demand across cities. First, national-level policy rationed permits by approving, on average, about 42% of requests. This disproportionately affected cities that employ a large number of H-1B workers as a share of total city employment. This is largely the variation that research by Kerr and Lincoln (2010), Peri, Shih, and Sparber (2015), and Ghosh, Mayda, and Ortega (2014) exploit. They find that cities (or firms) that heavily used the H-1B program in the past and continue to maintain high demand suffer losses when the federal government limits the total number of foreign-workers that can be hired. This rationing mechanism is present every year, and it becomes more stringent when the total number of permits is reduced or when the average demand for computer-related workers increases. Second, cities experience an additional and unexpected shock during the lottery years. For completely random reasons, the lottery causes some cities to receive more or fewer H-1B workers than was expected from national rationing rates. A way of describing this is that local labor markets expect a rationing of permits proportional to the satisfied H-1B quota at the national level, but they might also experience a random unexpected restriction to supply (that might be positive or negative) in addition to this effect.

To separate these channels, first define the national H-1B winning percentage (W) as in equation (5). This expression simply calculates the national win rate as the aggregate sum of cap-bound H-1Bs awarded divided by the sum of LCAs filed net of cap-exempt H-1B employees of for-profit firms.

$$W = \frac{\sum_c H_c}{\sum_c (L_c - X_c)} \quad (5)$$

Second, by substituting this value into the expression defining city-level normalized excess foreign-labor demand (d_c) in equation (1), one can compute the equivalent identity in equation (6) .

$$d_c = \underbrace{\frac{(1 - W) \cdot (L_c - X_c)}{Empl_{c,2005/06}}}_{d_c^{AVG}} + \underbrace{\frac{W \cdot (L_c - X_c) - H_c}{Empl_{c,2005/06}}}_{d_c^u} \quad (6)$$

This expression provides additional insight into the mechanisms through which H-1B shocks affect city-level labor supply. The first term in (6), d_c^{AVG} , represents an average rationing effect in city c from the H-1B distribution in 2007 and 2008. It multiplies a city's dependence upon the H-1B program by the national-level of H-1B rationing (or loss rate). As policy becomes more restrictive at the national level, this rationing effect rises, thus limiting a city's ability to hire foreign workers. As noted, this is effectively the channel of H-1B policy analyzed in papers like Kerr and Lincoln (2010) and Ghosh, Mayda, and Ortega (2014). The second term in (6), d_c^u , represents the unexpected shock. The $W \cdot (L_c - X_c)$ component reflects the number of new H-1Bs a city could reasonably expect to receive given national policy, whereas H_c represents the number of new cap-bound H-1Bs that a city actually receives. The difference therefore represents a random deviation from average induced by variation in lottery success across cities. It is an unexpected shock that, when positive, represents a random negative shock to hiring as many H-1B workers as expected.

The regressions in this subsection alter the specifications in (2), (3), and (4) by replacing d_c with d_c^{AVG} and d_c^u and allowing estimation of separate coefficients for these rationing and unexpected shock terms. Importantly, interpretation of the relative magnitudes of the coefficients has to account for the very different standard deviations of the variables. In Table 4 we report the unweighted mean, standard deviation, and 90/10 percentile range of the total, average, and unexpected rationing variables. While the mean of unexpected rationing is 0 by construction, we notice that the key difference is that most of the variation in the negative supply shock is driven by average rationing. The standard deviation of d_c^{AVG} (0.096) is 12 times larger than the standard deviation for d_c^u (0.008). Similarly, the 90/10 percentile range for the average rationing (0.09) is ten times larger than for the unexpected rationing.

Figure 7 shows the histogram for the whole distribution of d_c^u and we see that most of the mass corresponds to shocks that are very small (between +/-0.005, i.e. half a percentage point of initial employment).¹³ This clearly indicates that the purely random variation of the negative supply shock, while theoretically useful for identifying a causal effect, may be far too small to produce empirically relevant responses (though attenuation

¹³The figure omits four outliers for illustration purposes.

at the aggregate local labor market level should be less severe than one would find doing an analysis at the firm level).

Table 5 shows the coefficients on the variables d_c^{AVG} and d_c^u from specifications that are otherwise identical to those estimated in Tables 2 and 3. Panel A provides an exogeneity test by regressing pre-lottery foreign-born employment growth on the two components of rationing. Interestingly, we find no evidence that the variation of average rationing is correlated with pre-trend foreign-born growth so that even if this variable is not purely exogenous, we can at least argue that it is not significantly correlated with foreign-born employment growth in previous years. As expected, the unexpected lottery shock is not correlated with pre-trends. However, its correlation is estimated very imprecisely due to the very small variation of this term relative to city employment. In particular, an unexpected negative random shock to a city's H-1B supply equal to 0.01 (one percentage point) of computer employment is associated with pre-lottery growth as large as 29% or as low as -8% (this is roughly the 99% confidence interval). Hence, no pre-lottery outcome can be ruled out.

Panel B explores whether shocks are correlated with subsequent foreign-born employment growth (between 2005/06 and 2008/09). Recall that Table 2 confirmed that overall negative lottery shocks were associated with decreased foreign-born employment growth. Table 5 can only detect a statistically significant effect operating through the average rationing channel. This means that cities more heavily dependent upon the H-1B program see foreign-born employment losses when national-level policy reduces the number of permits available. This result complements previous work exploiting the combination between national-level policy and city-variation in foreign-born dependence to examine the effects of immigration. However, the regressions find no statistically significant evidence of an effect from random unexpected cross-city shocks. That is, when a city randomly receives fewer H-1B workers than expected from national rationing behavior, that city and its firms do not respond by hiring fewer foreign workers overall. The magnitude of the estimates and of the coefficients, however, emphasizes how this zero result is really arising from the lack of power rather than from a small point estimate. The coefficient on the random shock has a standard error between 4.8 and 8.1, which is too large to identify any reasonably sized effect. If we imagine for a moment that a reasonable effect of rationing on foreign-born hires could be equal to a one percentage point reduction in foreign employment

growth for each percentage point of H-1B rationing (roughly the effect estimated using average rationing) then even with the correct point estimate we could not rule out effects as large as +10 percentage points or -10 percentage points. These responses are unreasonably large in magnitude.

In the bottom of Panel B we examine the employment and wage consequences for native-born workers. The results again mirror previous findings. The native employment growth losses found in Table 3 appear to occur through the average rationing channel. Comparing a city with zero rationing to one with an average level results in reduced native college-educated computer-related employment growth between 2.1 (unweighted results) and 9.7 (dependence-weighted results) percentage points between 2005/06 and 2008/09 – a large effect that should be interpreted with some caution. Conversely, coefficients on the unexpected component of the shock are again insignificant and their standard errors too large to say anything meaningful. The final panel’s results for native wages fails to uncover any statistical effect, concurring with our earlier results suggesting that the H-1B program does not have immediate wage consequences for natives.

Table 5 is informative as it emphasizes a trade-off in the use of lottery data to identify the impact of H-1B rationing. First, while the city is an appropriate unit of analysis to find effects of changing the supply of H-1B workers, we find that the variation introduced by the pure random realization of the lottery are too small to identify any reasonable effect. Second, we establish that identification based on differences in H-1B dependence across metropolitan areas is consistent with exogeneity of the shock to pre-trends, has predictive power on foreign-born employment growth, and shows no crowding out of native employment or wages. In fact it finds mild evidence of crowding-in. Third, since the random shock has very small variance (one tenth of the variation of average rationing across cities) we might only be able to leverage that small variance on some subset of cities with certain characteristics.

This partially motivates our next stage of analysis. Ghosh, Mayda and Ortega (2014) argue that H-1B restrictions generate particularly strong and negative effects for firms highly-dependent upon H-1B workers. This is consistent with our rationing results and makes it worth examining whether the effect of the unexpected shock also depends on how tightly a city is constrained by the average rationing. Table 6 explores this by only selecting the 30 cities which have the highest levels of average rationing. Table 7 uses a slightly different selection criteria and includes only the cities that have above average computer-related employment

in 2005/06 and above average H-1B demand according to LCA filings. Only 22 cities in the dataset meet these criteria and are referred to as “Large H-1B Cities.”

Table 6 follows the same structure as Table 5 but displays average city-level rationing and the unexpected lottery shock coefficients only for highly rationed cities. The top part of the panel confirms exogeneity of the lottery by regressing pre-trends in foreign-born employment growth on the H-1B variables. The lottery is not associated with pre-lottery behavior, though point estimates and standard errors on the unexpected shock are very large.¹⁴ Results for post-lottery foreign-born employment growth are interesting. In two cases the unexpected negative shock to H-1B supply is now significantly associated with reduced foreign-born employment growth in the computer-sector of the city (though only at the 10% level). The magnitude and significance of the coefficient on average rationing are similar to those in Table 4. Thus, for highly rationed cities, the small unexpected negative H-1B supply shocks seem to be associated with decreases in foreign employment growth.

As the lottery satisfies these exogeneity and power tests, we can look at the impact on native employment and wages. Though no point estimate is significant, we find that even in this case all point estimates are negative. This implies that a reduction in H-1B permits, either via average rationing or randomly, do not produce any increase in native employment or wages, but are rather associated with insignificant declines in both.¹⁵

Examination of large H-1B cities in Table 7 shows similar correlations of average and unexpected rationing with native employment and wage growth. Cities experiencing larger negative supply shocks, either due to dependence or to randomness, do not exhibit any positive change in employment and wages of natives. Rather, they experience insignificant negative effects. However, for these cities the average rationing seems somewhat correlated with pre-lottery foreign employment growth, reminding us that this component of the variation may not be perfectly random. At the same time the unexpected shock is small and not strongly correlated with the post-lottery change in foreign-born workers, hence the identification power of our strategy for these cities is limited.

¹⁴Note also that the standard deviation of unexpected shocks for small cities is just 0.0038, so the point estimate of a one-standard deviation shock to such cities is 0.17.

¹⁵This regression can, of course, be run for other groups of cities as well. For example, unreported regressions – analogous to those of Table 6 – for both the second 30-city tier of rationing and the 30 least rationing cities find insignificant effects on employment growth in all but one case.

7 Discussion: What Can We Learn from Cities, Firms, and H-1B Lotteries?

This paper is the first to combine LCA data on labor demand and I-129 data on H-1B permits awarded in 2007 and 2008 to measure excess demand for temporary foreign workers. The random assignment of all H-1B permits due to the national lottery in those years allows us to exploit random variation in foreign labor supply to analyze consequences for native employment and wages. We aggregate the data on computer-related occupations to the metropolitan level. The resulting 238 cities serve as labor markets and units of observations for our analysis. This methodology helps to improve upon, understand, and reconcile related work in the literature in two important ways.

First, our work is directly comparable to previous papers that use local labor markets as units of analysis. Such work generally finds complementarities between foreign and native-born labor. One advantage of our approach is that we can identify effects through random, lottery-driven, exogenous changes in labor supply (as opposed to using conventional instrumental variables methods). Another advantage is that we recognize that H-1B rationing has two components. The component emphasized by previous studies (e.g., Kerr and Lincoln, 2010; Ghosh, Mayda, and Ortega, 2014; and Peri, Shih, and Sparber, 2015) leverages the change in aggregate work permits interacted with a city's dependence on foreign workers. It is reasonable to expect that a decrease in the total number of available H-1B permits would mostly affect cities (and firms) who rely on foreign workers as a large percentage of employment. The other component is more clearly random and attributable to the lottery realization. It records that some labor markets received a higher (or lower) shock than expected.

Second, our work is comparable to Doran, Gelber, and Isen (2014), which also exploits random variation to identify the effects of H-1B supply shocks. Their work does not find complementarity between foreign and native-born labor. One advantage of our approach is that we rely on variation across cities, whereas they exploit variation across firms. Their approach is potentially problematic if firms are able to compensate for H-1B losses by drawing from local supply of other foreign workers such as existing H-1B workers or those working on OPT status after graduating from local universities. It might not be so easy for cities to make

the same types of substitutions. Another advantage of our approach is that we use 2007 and 2008 H-1B data when all permits were distributed by lottery. Doran, Gelber, and Isen (2014), by contrast, exploit random variation from the 2005 and 2006 lotteries when only the small number of applications submitted on the last day of receipt were subject to a lottery. Taken together, the variation in H-1B availability at the firm level used by Doran, Gelber, and Isen (2014) is small. It is not surprising, therefore, that the authors find that lottery-driven shocks fail to produce significant firm-level variation in the number of foreign workers hired or of innovation output. It appears that losing one new H-1B worker could be easily offset by a firm with a foreign local hire, at least in the context of the small lotteries of 2005 and 2006. That paper address the policy question of whether there would be any significant employment and innovation effect at the firm level of marginally relaxing (by a few thousand) the H-1B quota. Our analysis, instead, considers whether there would be an effect at the local labor market level (on jobs and wages) of significantly relaxing the H-1B quota.

Ultimately, our results suggest that H-1B workers do not crowd out, and may even crowd in, native-born workers. Evidence is supported by our analysis of effects operating through both the average rationing and unexpected shock channels.

From an average rationing perspective, our results clearly support prior evidence that average H-1B rationing reduces foreign-employment and harms labor market opportunities for native-born workers. We must acknowledge that this channel remains subject to the same limitations of past research. Namely, if demand shocks are correlated with employment growth and H-1B dependence, then our model could suffer from omitted variables bias. However, we also found that H-1B variation across cities driven by foreign dependence exhibits no statistical correlation with pre-lottery foreign employment trends. Hence, there is no *prima facie* evidence that an omitted variables bias would exist. Moreover, when controlling for pre-lottery trends in the model, we find that negative shocks to H-1B supply exhibit insignificant or negative correlation with subsequent native employment growth. Thus, crowding out does not exist, but crowding in might.

From an unexpected shock perspective, we find that this more purely random shock component of H-1B rationing exhibits extremely small variation. *A priori* it has the potential to establish causality more convincingly than past work, but the standard errors do not allow for identification of reasonable effects in

the full sample of cities. It is only when we restrict the sample to cities that are highly-dependent upon H-1B workers that we find random, unexpected, negative H-1B shocks reduce foreign-born employment growth. It is still not the case, however, that cities compensate for H-1B losses by hiring more native-born workers. Thus, the model and data have the power to identify first-order foreign-born employment effects, and yet we still find no evidence that H-1B workers crowd out native-born, college-educated, computer workers from employment opportunities.

8 H-1B Rationing and Jobs

Before concluding, and keeping in mind the caveats about causality expressed above when using the normalized excess demand as a measure of the negative supply shock generated by H-1B rationing, we would like to translate our coefficients into number of jobs lost. To do so, we multiply each city's normalized excess demand (d_c) by the coefficients identified in the unweighted regressions of employment growth through 2008/09 in Tables 2 and 3 (-2.235 for foreign-workers; -0.510 for native-born workers). We then multiply this value by pre-lottery employment to calculate the predicted change in the number of jobs for foreign and native-born workers through 2008/09 (a year after the lottery). We do this for each city and then sum across cities to obtain the aggregate effect on jobs. Finally, we divide these values by excess demand in order to calculate the number of jobs lost per denied H-1B permit (that is, an unfilled LCA job opening).

Table 8 displays the figures for the 22 Large H-1B cities and the U.S. as a whole. Column 1 displays excess demand. Columns 2 and 3 show the corresponding negative effect on foreign and native college-educated computer-related employment. The final columns display the foreign and native jobs lost per denied H-1B permit, or more precisely, per unit of excess demand.

The first row (in bold) demonstrates that, economy-wide, each denied H-1B permit corresponded to an average of 0.7 fewer jobs for foreign-born workers by 2008/2009. This is very reasonable. Either because firms file an excess number of LCAs, or because firms might find alternative ways to hire foreign-born workers, the conversion of excess demand to foreign employment is less than one-to-one but still economically meaningful. The resulting effect of rationing on jobs for native-born Americans is also plausible: one denied H-1B costs 0.2 native jobs. In aggregate, H-1B rationing was associated with 207,627 fewer jobs for foreign-born workers

and 63,942 fewer jobs for Americans. This is consistent with the effects found by other papers that leveraged H-1B policy changes and found a positive contribution of the program on innovation, productivity, the size of firms, and native employment (Kerr and Lincoln, 2010; Peri, Shih, and Sparber, 2015; Ghosh, Mayda, and Ortega, 2014). With all the needed caveats about causality, this is also consistent with the idea that skilled foreign workers complement and positively affect the productivity of skilled American workers.

9 Conclusions

This paper used a simple framework to analyze an important question: If we reduce the number of H-1B skilled foreign worker permits, will native workers employed in computer-related jobs experience employment and wage benefits, losses, or no effect? We estimate the consequences of a random rationing of H-1B permits that took place in April 2007 and 2008 due to excess demand relative to the quota. In those years, some foreign-born individuals with job offers who requested a new permit were randomly awarded one while others were not. By corollary, some firms (and the cities in which they are located) were randomly unable to employ the individuals they had intended to hire.

We grouped employers of computer-related occupations by metropolitan area. We then estimated the size of the negative lottery rationing in each city and examined the consequences of these shocks on native-born computer-workers in the same metropolitan area. *A priori*, it is conceivable that native and foreign-born skilled workers are close substitutes. By denying a foreign H-1B worker to a firm, the lottery could prompt that firm to search for a similar native worker. If firms hiring foreign workers do so hoping to depress wages, then the unexpected negative shock should push firms to pay native workers more in order to lure them away from competitors. This would imply that the negative H-1B shock should be positively and significantly related to native employment and wages. In other words, a positive value of the estimated effects on native employment and wages would be a sign that immigrants are close substitute for natives and more immigration would hurt the labor market opportunities of natives.

On the other hand, it is also conceivable that native and foreign-born computer workers are complementary, or that foreign-born workers improve the productivity of firms, thereby leading firms to expand their hiring. In this case denying firms the ability to hire foreign workers reduces the need for complementary

natives as well. It might also reduce productivity growth and hence employment growth. Demand and wages could stagnate or decline for all workers. Additional H-1B workers would cause firms to expand while native computer specialists would not be crowded out and could even experience improved labor market opportunities. In other words, a negative or zero value of the estimated effects on native employment and wages would be a sign that immigrants complement natives and/or improve productivity, and more immigration would not hurt the employment and wage opportunities of natives.

Our estimates support this second conclusion: Employers of computer-related workers did not hire more natives when the foreign workers they intended to hire were denied H-1B permits. Instead the employment and wages of natives in similar occupations was at best unchanged and at worst harmed. This is consistent with the notion that H-1B workers complement native computer workers and/or increase the productivity of the firm, and that skilled foreign-born workers have improved (or at worst left unchanged) the labor market opportunities of natives. In this sense, our results complement work by Peri, Shih, and Sparber (2015) who used decennial Census information and found that for each new foreign STEM worker (driven by H-1B variation), between 0.5 and 0.6 native STEM jobs are created in the longer term.

This paper further advances the literature by decomposing the H-1B lottery into two effects: One is driven by average rationing – the interaction between national-level restrictions and city-specific H-1B dependence. The other arises from unexpected shocks – the difference between H-1B permits expected (given national-level rationing) and the actual number of permits received. We find that most of the negative effects of H-1B restrictions result from the average rationing effect. This largely complements work by Kerr and Lincoln (2010). Unexpected shocks, which are much smaller, are not associated with effects on either foreign or native employment growth for the full sample of cities. However, further analysis reveals that such shocks may be detrimental to cities that were strongly dependent on foreign workers and therefore heavily rationed in 2007 and 2008 – a result consistent with Ghosh, Mayda and Ortega (2014). Combining randomness in the lottery and the non-random degree of dependence on foreign-born labor may be crucial to properly analyze the effect of H-1B policy changes. Overall, the results confirm that restrictive H-1B policy does not benefit, and may harm, native-born, college-educated, computer-related employment, even in the short-run.

References

- Borjas George J. and Kirk B. Doran, (2012). “The Collapse of the Soviet Union and the Productivity of American Mathematicians,” *The Quarterly Journal of Economics*, Oxford University Press, Vol. 127(3), pages 1143-1203.
- Clemens, Michael (2010). “The Roots of Global Wage Gaps: Evidence from Randomized Processing of U.S. Visas.” Center for Global Development, Working Paper 212.
- Doran, Kirk, Alexander Gelber, and Adam Isen (2014). “The Effect of High-Skilled Immigration on Patenting and Employment: Evidence from H-1B Visa Lotteries.” NBER Working Paper 20668.
- Ghosh, Anirban, Anna Maria Mayda, and Francesc Ortega (2014), “The Impact of Skilled Foreign Workers on Firms: an Investigation of Publicly Traded U.S. Firms.” Centre for Research and Analysis of Migration (CReAM) Discussion Paper 42/14.”
- Kato, Takao and Chad Sparber (2013). “Quotas and Quality: The Effect of H-1B Visa Restrictions on the Pool of Prospective Undergraduate Students from Abroad.” *The Review of Economics and Statistics*, Vol. 95(1): 109-126.
- Kerr, William and William F. Lincoln. (2010) “The Supply Side of Innovation: H-1B Visa Reforms and U.S. Ethnic Invention.” *Journal of Labor Economics* 28(3): 2010, pp. 473-508.
- Hira, Ron (2007) “Outsourcing America’s Technology and Knowledge Jobs: High-Skill Guest Worker Visas are Currently Hurting Rather than Helping Keep Jobs at Home.” Economic Policy Institute Briefing Paper no. 187.
- Hunt, Jennifer and Marjolaine Gauthier-Loiselle (2010). “How Much Does Immigration Boost Innovation?” *American Economic Journal: Macroeconomics*, Vol. 2, pp. 31-56.
- Matloff, Norman (2008) “H-1Bs: Still Not the Best and the Brightest,” Center for Immigration Studies Background.
- Moser, Petra, Alessandra Voena, and Fabian Waldinger, (2014). “German-Jewish Emigres and U.S. Invention,” NBER Working Papers 19962, National Bureau of Economic Research, Inc.

Moretti Enrico (2010). “Local Multipliers” *American Economic Review: Papers & Proceedings*, Vol 100, pp 1–7. <http://www.aeaweb.org/articles.php?doi=10.1257/aer.100.2.1>

Peri G. (2007) “Higher Education, Innovation and Growth.” In Giorgio Brunello, Pietro Garibaldi, and Etienne Wasmer eds. *Education and Training in Europe*, Oxford: Oxford University Press

Peri Giovanni, Kevin Shih and Chad Sparber, (2015). “STEM Workers, H-1B Visas, and Productivity in U.S. Cities,” *Journal of Labor Economics*, Forthcoming.

Peri Giovanni, Kevin Shih, Chad Sparber, and Angie Marek Zeitlin (2014). “Closing Economic Windows: How H-1B Visa Denials Cost U.S.-Born Tech Workers Jobs and Wages During the Great Recession,” *The Partnership for a New American Economy*.

U.S. Citizenship and Immigration Services (2007). “USCIS Updates Count of FY 2008 H-1B Cap Filings,” April 10, 2007.

U.S. Citizenship and Immigration Services (2008). “USCIS Runs Random Selection Process for H-1B Petitions,” March 31, 2008.

A Appendix

The main text of this paper relies upon regressions of pre-lottery foreign employment growth trends to examine the exogeneity of the H-1B variables. In each case, we find no statistical evidence for a significant correlation, thereby lending support toward exogeneity.

Other pre-lottery variables, however, might also be worth exploring. One alternative is to consider pre-lottery native-born employment growth. Table A1 displays the results from these regressions, which are analogous to figures presented for foreign employment in the main text. The top panel presents the native-born equivalent to the foreign-born pre-lottery results of Table 2. Unlike with foreign employment, we do find some evidence for a significant negative correlation between native pre-lottery employment growth and the H-1B shock in two of the three weighting schemes. While this is problematic for a fully causal interpretation of regression results in the paper, the problem is alleviated by the inclusion of pre-lottery growth variables in all regressions examining the main outcomes of interest.

Importantly, this relationship is driven by the dependency channel, and our unexpected shock variable (described in Section 6.2) appears to be random. On the one hand, the second panel of Table A2 seemingly questions the validity of this claim, as it uncovers evidence that pre-lottery trends are correlated with unexpected shocks. On the other hand, this correlation is entirely driven by four outliers: Three cities (Davenport, IA; Manchester, NH; and Springfield, IL) lie more than three standard deviations below the mean, and one city (Worcester, MA) lies more than three standard deviations above it. Their removal (in the third panel) attests for the plausible exogeneity and randomness of the shock variable.¹⁶

Also important, the fourth panel indicates that no significant correlation exists among highly rationed cities (the analogue to Table 6). The text notes that our H-1B variables are nonetheless able to identify a first-order effect in which negative supply shocks reduce subsequent foreign employment growth. The final panel indicates that average rationing is, once again, associated with pre-lottery native employment growth, whereas the random shock is not.

The bottom line is that the normalized excess demand variable (d_c) is uncorrelated with pre-lottery

¹⁶Table 5 in the main text includes these four outlying cities. Their removal would cause some coefficients and standard errors to rise, but the results are qualitatively similar. Furthermore, the outlying cities are neither highly rationed nor Large H-1B cities.

foreign-born employment growth but could be correlated with pre-lottery native-born employment growth. We recognize that this creates a challenge for purely causal interpretation of the results in the paper, though the inclusion of pre-lottery trends in the models should help in this regard. This correlation is driven by H-1B dependency; unexpected lottery shocks appear to be random. Such shocks, when negative, reduce subsequent foreign-born employment growth, but have no effect on native employment growth. This finding – like others in the paper – suggests that foreign-born, college-educated, computer-related workers do not crowd out similar native-born workers.

Table 1:
Last Date of Receipt for H-1B Applications in Each Fiscal Year

Fiscal Year	Final Receipt Date	Days in Filing Period
2004	February 17, 2004	323
2005	October 1, 2004	184
2006	August 10, 2005	132
2007	May 26, 2006	56
2008*	April 3, 2007	3
2009*	April 7, 2008	7
2010	December 21, 2009	265
2011	January 26, 2011	301
2012	November 22, 2011	236
2013	June 11, 2012	72
2014*	April 7, 2013	7
2015*	April 7, 2014	7

*= A lottery was held to assign all H-1B permits.

Note: This Table describes the last day in the year in which H-1B permit applications were accepted and the implied number of days available for individuals to file requests. After the last date of receipt, no additional H-1B permits are available to perspective employees of private for-profit firms who do not already have H-1B status.

Table 2
The H-1B supply shock effect on pre and post lottery foreign-born employment
238 MSA level data for computer-related occupations only

Explanatory Variable: Negative H-1B-Supply Shock Relative to Total 2005-06 Employment (Variable d_c in the Text)			
	(1)	(2)	(3)
	Unweighted	Weighted by H-1B Dependence	Weighted by ln(H-1B Dependence)
Panel A			
<i>Dependent Variable:</i>			
Pre-Lottery Growth Rate of Foreign-Born College Educated Employees			
Growth Rate between 2000 and 2005-06	-1.624 (1.031)	-0.225 (0.376)	-0.913 (0.683)
Panel B			
Post-Lottery Growth Rate of Foreign-Born College Educated Employees			
Growth Rate between 2005-06 and 2008-09	-2.235** (1.129)	-0.331 (0.344)	-1.051* (0.622)
Growth Rate between 2005-06 and 2009-10	-1.897* (0.974)	-0.521*** (0.178)	-0.996* (0.574)
Growth Rate between 2005-06 and 2010-11	-1.730* (0.936)	-0.488* (0.253)	-1.166* (0.624)

Note: Each entry in the table is the coefficient of the explanatory variable d_c = “standardized negative-supply shock of H-1B”, described in the text. The dependent variable in the regression is described in the first row. The observations are 238 metropolitan areas in the US. The method of estimation is Least squares, using different weights (as described in each column). Panel A tests the pre-lottery outcomes and Panel B the post-Lottery ones.

Table 3
The H-1B supply shock effect on post-lottery native employment and wages
238 MSA level data for computer-related occupations only

Explanatory Variable: Negative H-1B-Supply Shock Relative to Total 2005-06 Employment (Variable d_c in the Text)			
	(1)	(2)	(3)
	Unweighted	Weighted by Dependence	Weighted by ln(H-1B Dependence)
Panel A			
<i>Dependent Variable: Post-Lottery Growth Rate of Native College-Graduate Employees</i>			
Growth Rate between 2005-06 and 2008-09	-0.510* (0.301)	-0.394*** (0.092)	-0.446* (0.228)
Growth Rate between 2005-06 and 2009-10	-0.528* (0.279)	-0.624*** (0.170)	-0.507** (0.197)
Growth Rate between 2005-06 and 2010-11	-0.337 (0.384)	-0.768*** (0.276)	-0.434 (0.294)
Panel B			
<i>Dependent Variable: Post-Lottery Growth Rate of Wages Paid to Native College-Graduates</i>			
Growth Rate between 2005-06 and 2008-09	-0.067 (0.090)	0.022 (0.080)	-0.014 (0.085)
Growth Rate between 2005-06 and 2009-10	-0.195** (0.089)	-0.075 (0.080)	-0.109 (0.078)
Growth Rate between 2005-06 and 2010-11	-0.191* (0.101)	0.002 (0.035)	-0.095 (0.068)

Note: Each entry in the table is the coefficient of the explanatory variable d_c = “standardized negative-supply shock of H-1B”, described in the text. The dependent variable in the regression is described in the first row. The observations are 238 metropolitan areas in the US. The method of estimation is Least squares, using different weights (as described in each column). Panel A analyzes the impact of the negative H-1B shock on native college graduate employment in computer-related occupations. Panel B analyzes the impact of the negative H-1B shock on native college graduate wages in computer-related occupations.

Table 4
Summary statistics, average and unexpected rationing
238 US Metropolitan areas, combining 2007 and 2008 lotteries

Variable	Obs.	Mean	Std. deviation	90-10 percentile
Normalized Excess demand (total rationing) d_c	238	0.041	0.096	0.09
Average Rationing (dependence-driven) d^{AVG_c}	238	0.042	0.096	0.09
Unexpected Rationing (lottery-driven shock) d^u_c	238	0	0.008	0.009

Note: The Normalized Excess demand or “total rationing” captures the supply shocks of H-1B permits across cities. It is a measure of demand (applications for foreign workers) minus supply (the awarded H-1B) combining the April 2007 and 2007 lotteries, and it is divided by the 2005/06 level of employment in computer-related occupations. The unit of observation is the metropolitan area.

Table 5
Decomposing total rationing into average and unexpected H-1B supply shocks
MSA level data for computer-related occupations only

Explanatory Variables: Average and Unexpected Negative H-1B-Supply Shocks				
		(1)	(2)	(3)
		Unweighted	Weighted by H-1B Dependence	Weighted by ln(H-1B Dependence)
Panel A				
<i>Dependent Variable: Pre-Lottery Growth Rate of Foreign College Educated Employees</i>				
Growth Rate between 2000 and 2005-06	Average Rationing	-1.579 (1.003)	-0.289 (0.365)	-0.924 (0.675)
	Unexpected Shock	11.259 (9.303)	4.805 (3.853)	7.186 (6.432)
Panel B				
<i>Dependent Variable: Post-Lottery Growth Rate of Foreign College Educated Employees</i>				
Growth Rate between 2005-06 and 2008-09	Average Rationing	-2.201** (1.092)	-0.417* (0.242)	-1.078* (0.583)
	Unexpected Shock	13.272 (8.156)	5.870 (4.879)	8.996 (5.842)
<i>Dependent Variable: Post-Lottery Growth Rate of Native College-Graduate Employees</i>				
Growth Rate between 2005-06 and 2008-09	Average Rationing	-0.505* (0.292)	-0.384*** (0.103)	-0.446* (0.237)
	Unexpected Shock	1.301 (6.375)	-1.239 (4.597)	-0.801 (6.041)
<i>Dependent Variable: Post-Lottery Growth Rate of Wages Paid to Native College-Graduates</i>				
Growth Rate between 2005-06 and 2008-09	Average Rationing	-0.069 (0.096)	0.024 (0.080)	-0.013 (0.092)
	Unexpected Shock	-0.663 (1.566)	-0.199 (1.733)	-0.790 (1.529)

Note: The entries in the table are the coefficient on the average rationing and unexpected rationing of H-1B permits as described in the text. The dependent variable in each regression is described in the first row. The observations are 238 metropolitan areas in the US. The method of estimation is Least squares, using different weights (as described in each column).

Table 6
Average and unexpected H-1B supply shocks in the top 30 dependent cities
30 MSA level data for computer-related occupations only

Explanatory Variables: Average and Unexpected Negative H-1B-Supply Shocks		(1)	(2)	(3)
		Unweighted	Weighted by Dependence	Weighted by ln(H-1B Dependence)
<i>Dependent Variable: Pre-Lottery Growth Rate of Foreign College Educated Employees</i>				
Growth Rate between 2000 and 2005-06	Average Rationing	-0.457 (0.970)	-0.220 (0.539)	-0.391 (0.867)
	Unexpected Shock	31.256 (29.638)	0.848 (15.237)	25.143 (26.551)
<i>Dependent Variable: Post-Lottery Growth Rate of Foreign College Educated Employees</i>				
Growth Rate between 2005-06 and 2008-09	Average Rationing	-0.505** (0.242)	-0.448** (0.192)	-0.516** (0.232)
	Unexpected Shock	-10.127* (5.374)	-5.288 (5.121)	-9.245* (5.232)
<i>Dependent Variable: Post-Lottery Growth Rate of Native College-Graduate Employees</i>				
Growth Rate between 2005-06 and 2008-09	Average Rationing	-0.407 (0.278)	-0.200 (0.135)	-0.363 (0.242)
	Unexpected Shock	-6.438 (8.300)	-3.139 (4.743)	-5.706 (7.558)
<i>Dependent Variable: Post-Lottery Growth Rate of Wages Paid to Native College-Graduates</i>				
Growth Rate between 2005-06 and 2008-09	Average Rationing	-0.043 (0.096)	-0.008 (0.063)	-0.035 (0.089)
	Unexpected Shock	-1.531 (1.997)	-0.762 (1.865)	-1.502 (1.901)

Note: The entries in the table are the coefficient on the average rationing and unexpected rationing of H-1B permits as described in the text. The dependent variable in each regression is described in the first row. The observations are the 30 metropolitan areas in the US with the highest dependence on H-1B permits measured as number of cap-bound LCA applications in January-April relative to the city's computer-related employment. The method of estimation is Least squares, using different weights (as described in each column).

Table 7
Average and unexpected H-1B supply shocks in Large H-1B cities
22 MSA level data for computer-related occupations only

Explanatory Variable: Negative H-1B-Supply Shock Relative to Total 2005-06 Employment (Variable d_c in the Text)		(1)	(2)	(3)
		Unweighted	Weighted by Dependence	Weighted by ln(H-1B Dependence)
<i>Dependent Variable: Pre-Lottery Growth Rate of Foreign College Educated Employees</i>				
Growth Rate between 2000 and 2005-06	Average			
	Rationing	-0.543 (0.341)	-0.511** (0.223)	-0.508 (0.315)
	Unexpected Shock	4.479 (9.704)	3.865 (6.311)	4.289 (8.767)
<i>Dependent Variable: Post-Lottery Growth Rate of Foreign College Educated Employees</i>				
Growth Rate between 2005- 06 and 2008-09	Average			
	Rationing	-0.791** (0.328)	-0.502*** (0.164)	-0.720** (0.303)
	Unexpected Shock	1.255 (4.736)	0.864 (3.257)	0.966 (4.511)
<i>Dependent Variable: Growth Rate of Native College-Graduate Employees</i>				
Growth Rate between 2005- 06 and 2008-09	Average			
	Rationing	-0.221 (0.231)	0.024 (0.114)	-0.162 (0.209)
	Unexpected Shock	-5.487 (3.448)	-7.102*** (2.232)	-5.955* (3.254)
<i>Dependent Variable: Growth Rate of Wages Paid to Native College-Graduates</i>				
Growth Rate between 2005-06 and 2008-09	Average			
	Rationing	-0.090 (0.071)	-0.034 (0.037)	-0.077 (0.063)
	Unexpected Shock	-0.029 (1.063)	0.045 (0.854)	0.030 (0.996)

Note: The entries in in the table are the coefficient on the average rationing and unexpected rationing of H-1B permits as described in the text. The dependent variable in each regression is described in the first row. The observations are the 22 metropolitan areas in the US with above average computer-related employment and LCA filings. The method of estimation is Least squares, using different weights (as described in each column).

Table 8
College-educated employment between 2005/06 and 2008/09 and excess demand for H-1B workers

	Excess Demand, H-1B	Predicted Change in Foreign Employment	Predicted Change in Native Employment	Foreign jobs per unit of Excess demand	Native jobs per unit of excess demand
<i>United States (All US Cities)</i>	300,032	-207,627	-63,942	-0.69	-0.21
Atlanta, GA	9,147	-4,232	-2,270	-0.46	-0.25
Austin, TX	1,341	-492	-364	-0.37	-0.27
Baltimore, MD	1,161	-363	-311	-0.31	-0.27
Boston, MA-NH	15,704	-10,746	-4,083	-0.68	-0.26
Chicago, IL	17,576	-9,635	-4,151	-0.55	-0.24
Columbus, OH	1,550	-547	-372	-0.35	-0.24
Dallas-Fort Worth, TX	31,432	-15,202	-6,899	-0.48	-0.22
Detroit, MI	11,104	-6,290	-2,230	-0.57	-0.20
Houston-Brazoria, TX	3,623	-2,305	-706	-0.64	-0.19
Kansas City, MO-KS	1,383	-294	-395	-0.21	-0.29
Los Angeles-Long Beach, CA	5,192	-3,808	-885	-0.73	-0.17
New York-Northeastern NJ	99,931	-87,443	-18,740	-0.88	-0.19
Philadelphia, PA-NJ	5,021	-2,352	-1,218	-0.47	-0.24
Phoenix, AZ	1,197	-376	-254	-0.31	-0.21
Pittsburgh, PA	1,754	-465	-527	-0.27	-0.30
Raleigh-Durham, NC	2,467	-752	-785	-0.30	-0.32
San Diego, CA	1,157	-502	-273	-0.43	-0.24
San Francisco-Oakland-Vallejo, CA	6,179	-5,739	-1,144	-0.93	-0.19
San Jose, CA	14,360	-20,036	-1,959	-1.40	-0.14
Seattle-Everett, WA	2,986	-1,877	-688	-0.63	-0.23
Tampa-St. Petersburg-Clearwater, FL	1,201	-331	-245	-0.28	-0.20
Washington, DC-MD-VA	31,676	-17,938	-7,927	-0.57	-0.25

Figure 1

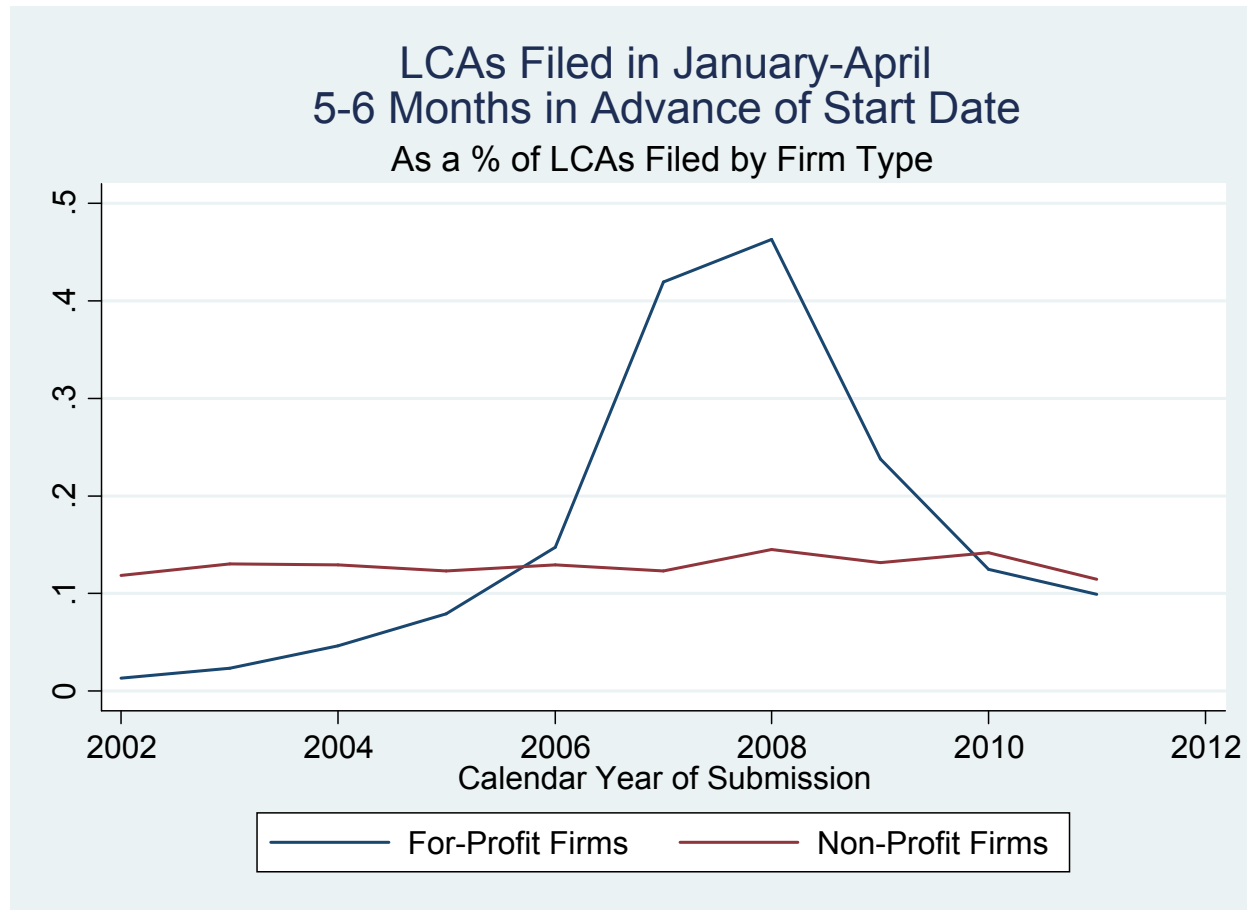


Figure 2

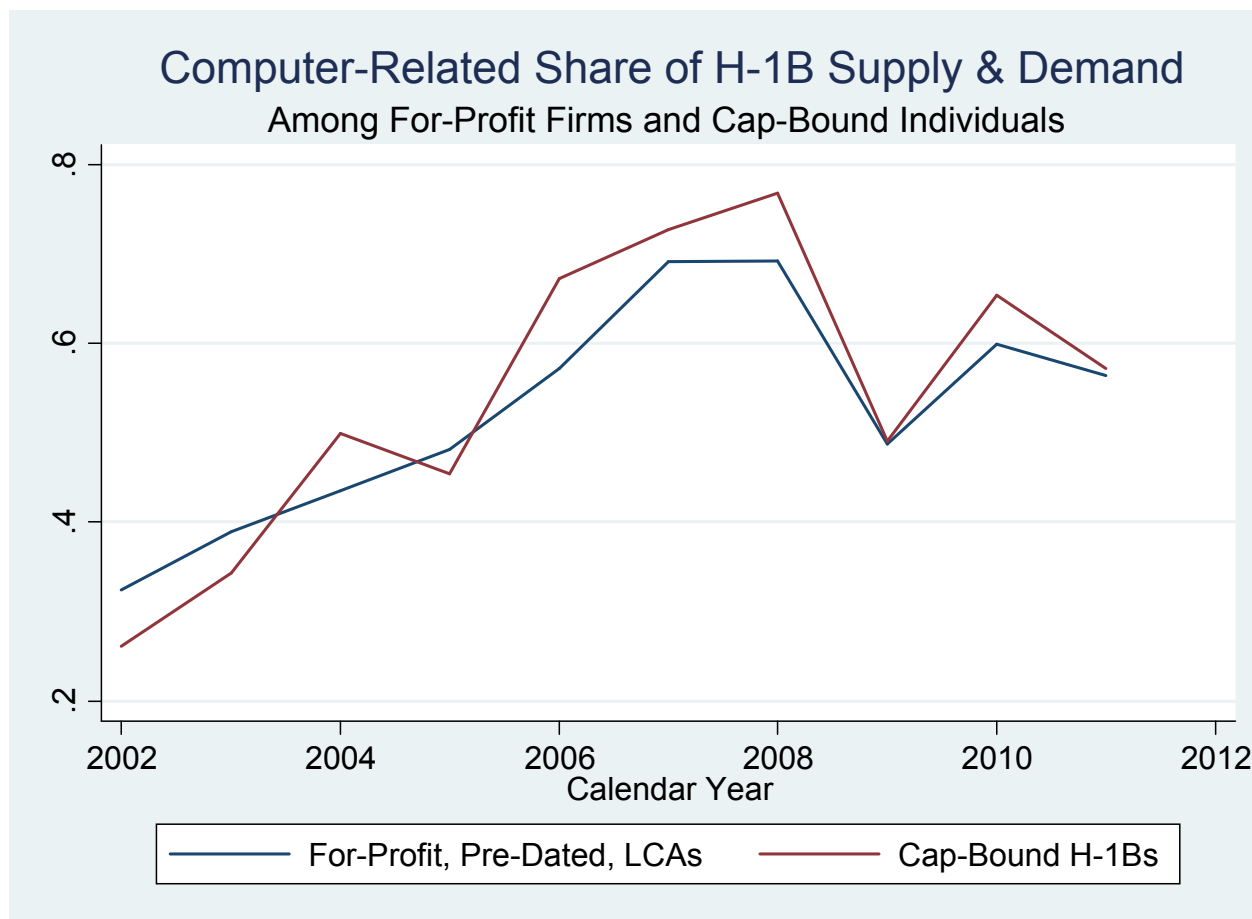


Figure 3

Effect of a negative H-1B supply shock when native and foreign-born computer scientists are perfect substitutes

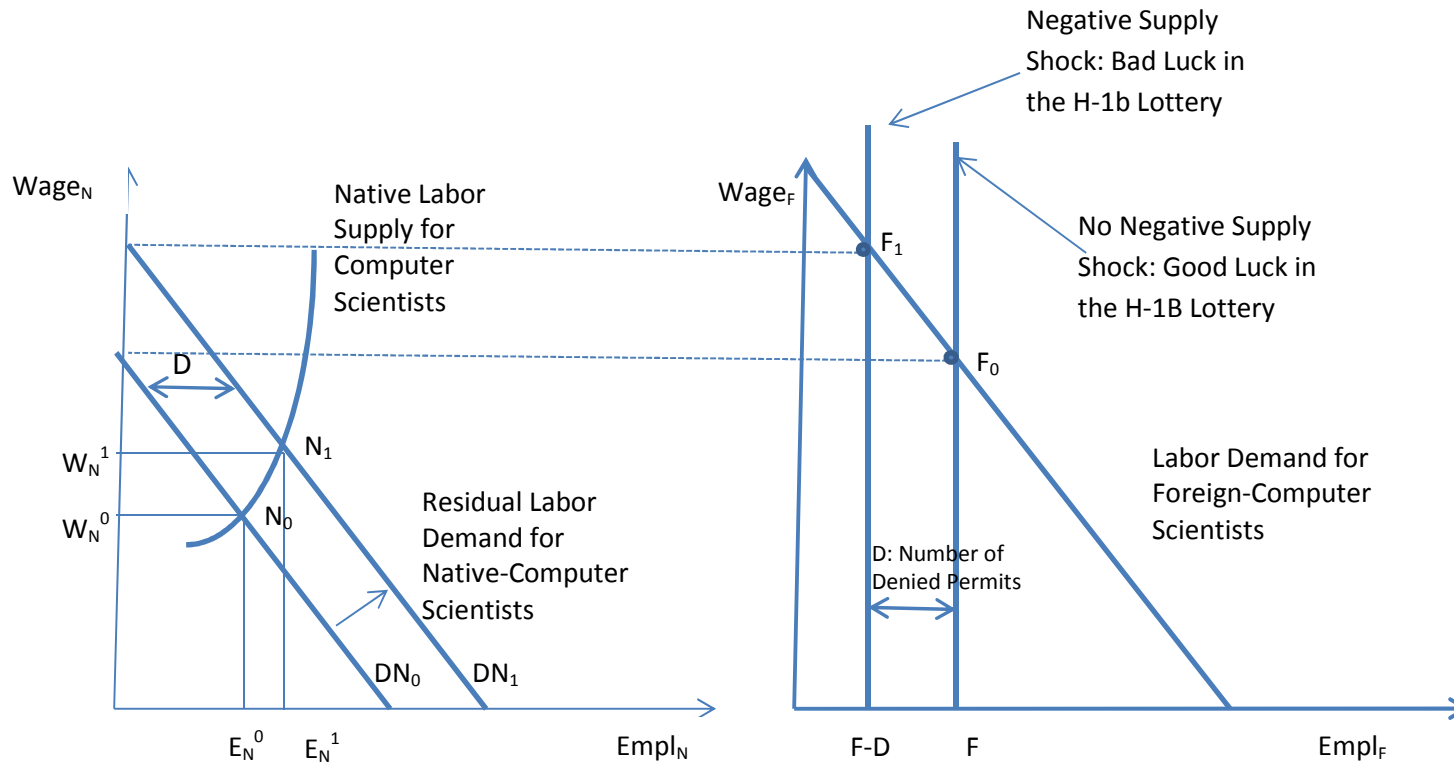


Figure 4

Effect of a negative H-1B supply shock when native and foreign-born computer scientists are complements or have positive productivity effects

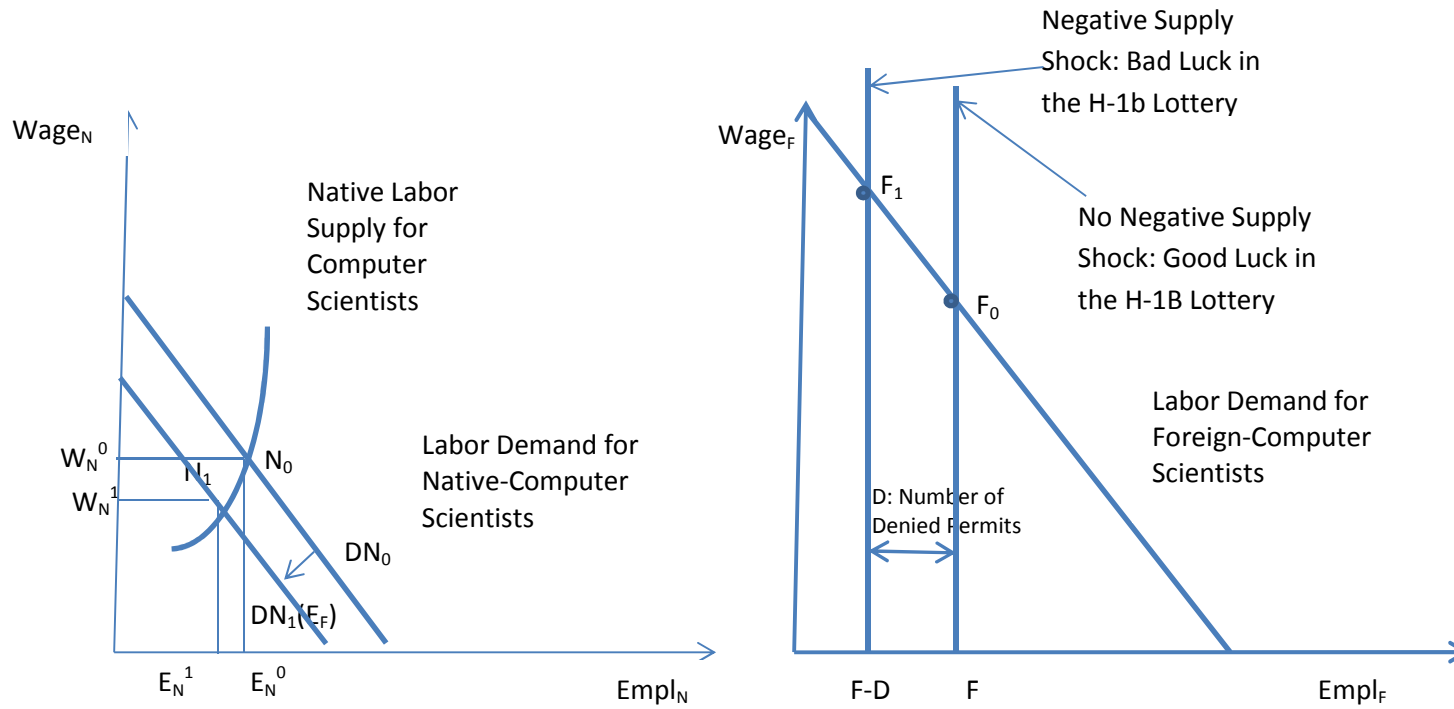


Figure 5

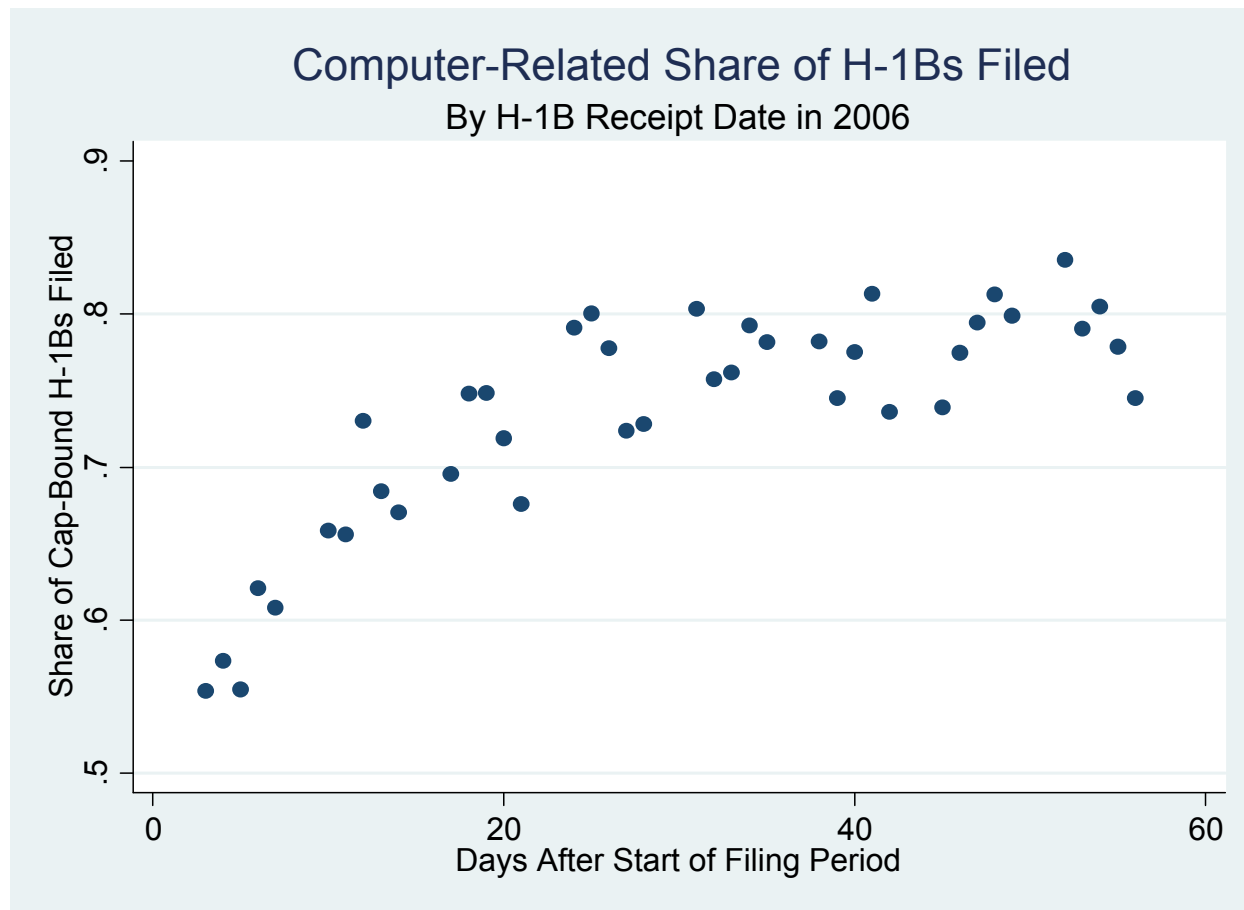


Figure 6

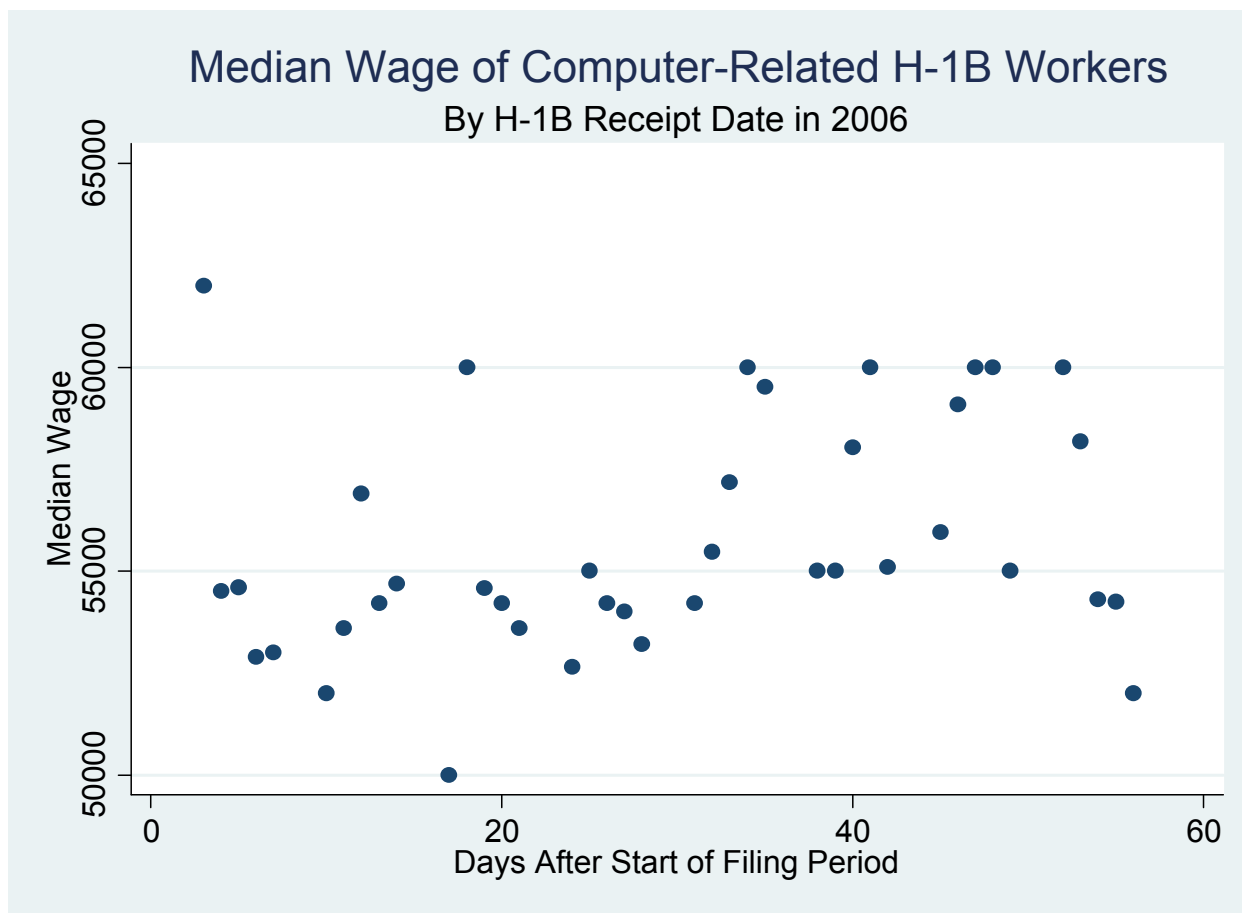


Figure 7

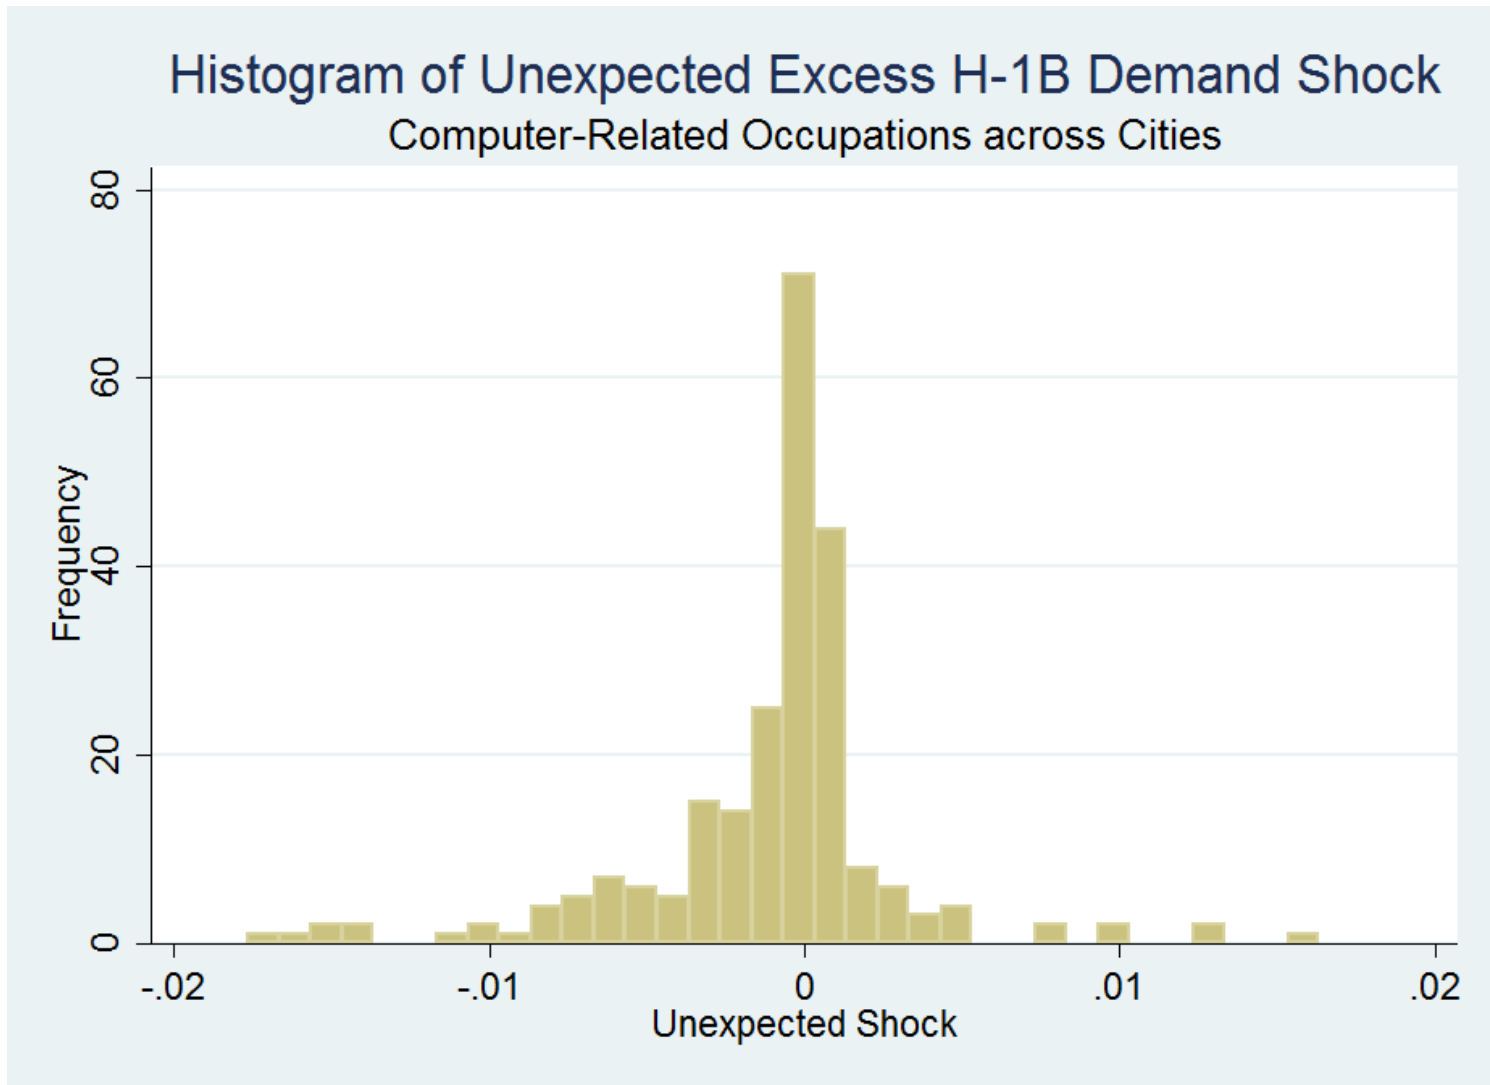


Table A1
The H-1B supply shock and pre-lottery native employment growth

		Unweighted	Weighted by Dependence	Weighted by ln(H-1B Dependence)
<i>Model:</i>				
Normalized Excess Demand (Table 3)				
Growth Rate between 2000 and 2005-06	Total Rationing	-1.021** (0.427)	-0.023 (0.198)	-0.612* (0.319)
Decomposed Shock (Table 5)				
Growth Rate between 2005-06 and 2008-09	Average Rationing	-1.000** (0.394)	-0.094 (0.141)	-0.618** (0.279)
	Unexpected Shock	5.188 (3.414)	6.005*** (2.228)	5.158* (2.752)
Decomposed Shock, Four Outliers Omitted				
Growth Rate between 2005-06 and 2009-10	Average Rationing	-1.214** (0.471)	-0.236* (0.132)	-0.788** (0.321)
	Unexpected Shock	-0.455 (7.457)	-4.504 (3.667)	-2.313 (5.665)
Highly Rationed Cities (Table 6)				
Growth Rate between 2005-06 and 2010-11	Average Rationing	-0.018 (0.164)	0.068 (0.101)	-0.006 (0.148)
	Unexpected Shock	-9.039 (6.248)	-5.053 (3.980)	-8.255 (5.774)
Top H-1B Cities (Table 7)				
Growth Rate between 2005-06 and 2010-11	Average Rationing	-0.605** (0.272)	-0.210 (0.140)	-0.511* (0.248)
	Unexpected Shock	5.522 (4.806)	1.709 (4.113)	4.445 (4.682)

Note: The models above test the relationship between the explanatory variable and pre-lottery native-born, college-educated, computer-related employment growth. Each entry in the table is the coefficient of the explanatory variable d_c or its components as described in the text. Panels correspond to the models (and tables) indicated. The observations are 238 metropolitan areas in the US. The method of estimation is Least squares, using different weights (as described in each column).