Buying Lottery Tickets for Foreign Workers: Lost Quota Rents Induced by H-1B Policy*

Rishi Sharma (Colgate University) Chad Sparber (Colgate University)

November 25, 2023

Abstract

The H-1B program allows firms in the United States to temporarily hire high-skilled foreign citizens. The government restricts inflows of new H-1B workers and therefore creates potential rents typical of a quota. Importantly, however, the US allocates H-1B status by random lottery. We develop a theoretical model demonstrating that this lottery creates a negative externality that destroys quota rents by incentivizing firms to search for more workers than can actually be hired. Some firms specialize in hiring foreign labor and contracting out those workers' services to third-party sites. These outsourcing firms exacerbate the search externality. Numerical exercises suggest that these processes result in an annual economic loss exceeding $10,000 per new H-1B worker hired relative to what would occur in the absence of lottery allocation.

Key Words: Skilled Workers, H-1B, Quota Rents, Outsourcing.

JEL Codes: J61, J68, F22

*Address: Rishi Sharma, Department of Economics, Colgate University, rsharma1@colgate.edu; Chad Sparber, Department of Economics, Colgate University, csparber@colgate.edu. The authors thank Isla Globus-Harris and Erik Lillethun for their input. All views contained herein are the authors' own.

1 Introduction

The H-1B program provides an important pathway for college-educated foreign citizens to secure temporary legal employment in the United States. Many economists credit the program and high-skilled immigrants for significantly enhancing US innovation, technology, productivity, wages, and GDP growth.[1] Less sanguine views often focus on two concerns. The first is a fear that infusions of skilled foreign-born workers reduce wages paid to subsets of native-born labor.[2] The second is that firms engage in domestic outsourcing by hiring H-1B workers and contracting out those workers' services to third-party clients.[3]

Our paper develops a theoretical model arguing that the US government's response to this first concern has created costs exacerbated by the second: Restrictions on the number of new H-1B entrants creates the potential for quota rents. Random lottery allocation of H-1B status generates search cost externalities that would not exist under alternative allocation methods. Firms engage in competitive rent seeking behavior that destroys quota rents. Domestic outsourcing specialists amplify negative search cost externalities.

To understand these effects more fully, consider relevant characteristics of the H-1B program in turn. First, policy-makers seeking to reduce potential labor market competition with skilled immigrants face options including a tax on H-1B employment or a minimum wage for foreign-born employees. Instead, the US limits the number of new H-1B workers to just 85,000 per year. As with any quota, this cap generates potential rents.

Second, interest in the H-1B program far exceeds the quota. During the first week of application eligibility for fiscal year (FY) 2017, for example, United States Citizenship and Immigration Services (USCIS) received 2.8 times more petitions for new H-1B status than they could legally approve for the entire year. Though the US could use market mechanisms (e.g., willingness to pay) to allocate quota-restricted H-1Bs, current policy instead uses a lottery. Our theoretical model demonstrates that this random allocation incentivizes firms to search for (and extend offers to) far more workers than can actually be hired. Any offer extended by one firm creates a negative externality on competing firms by reducing the chances that they will win the lottery and secure permission to hire the individuals they would like to employ. The negative externality and competitive rent seeking combine to destroy quota rents.

Third, negative externalities are compounded by firms that specialize in renting their

[1] See Hunt and Gauthier-Loiselle (2010), Kerr and Lincoln (2010), Hunt (2011), Moser et al. (2014), Kerr, Kerr, and Lincoln (2015), Peri, Shih, and Sparber (2015), National Academies of Sciences, Engineering, and Medicine (2017), Bernstein et al. (2018), and Gunadi (2019).

[2] See Borjas and Doran (2012), Bound et al. (2015), Ma (2020), Turner (2020), and Doran et al. (2022).

[3] Costa and Hira (2020) provide definitions and criticisms of this form of outsourcing activity.

workers' services to other firms. Such specialization might carry economic benefits in a competitive free market, but in the context of quota restrictions and lottery allocation, outsourcing specialists further increase the number of petitions for H-1B status. Thus, while the lottery is responsible for creating the negative search cost externality, outsourcing firms amplify those costs.

More formally, firms in our model generate output according to a standard Constant Elasticity of Substitution (CES) production function and face monopolistically competitive demand for their products. They search for foreign-employees but can only hire them if they randomly win the right to do so through a lottery. They can, however, contract with other firms to buy or sell the services provided by foreign workers in a secondary market. An important implication of this environment is that changes in the H-1B quota do not generate proportional changes in the probability of winning the lottery since the quota also affects the number of H-1B petitions that firms submit. We provide a closed form solution demonstrating that the elasticity of the win probability with respect to the quota is positively related to the share of firm expenditures spent hiring H-1B workers in the economy and the premium paid to outsourced workers. The larger this elasticity, the more the quota leads to higher search externalities and increased H-1B concentration.

Our model distinguishes between two firm types: While most "regular" firms experience increasing marginal costs associated with labor searches, "outsourcing" firms experience constant marginal costs, possibly because they have a large pool of potential H-1B workers already employed at overseas headquarters or affiliates.[4] The existence of outsourcing specialists is not necessary for the strategic behavior described above; lottery allocation incentivizes all firms to petition to hire more workers than they actually intend to employ. However, outsourcing firms do increase the size of the search externality.

We provide numerical exercises suggesting that the negative job search externality cost US firms more than $1 billion during the late 2010s when firms submitted an extraordinarily high number of H-1B petitions on behalf of prospective employees. In other words, the government forfeited between $12,000 and $22,000 per new H-1B worker hired that could have been retained as tax revenue or labor welfare under alternative allocation methods.

This paper connects to several branches of the economics literature. Our theoretical approach has parallels to models examining raffles, stock externalities, and all-pay auctions.[5] A firm's decision to bid for a finite resource imposes a negative externality on other firms by

[4] Note, for example, that critics often list Indian-based companies as outsourcing firms. Indian-born workers accounted for more than half of H-1Bs granted for initial employment and nearly three quarters of H-1Bs overall in FY 2019.

[5] See Morgan (2000), Landry et al. (2006), Lange et al. (2007), Gordon (1954), Smith (1968), Brito and Intriligator (1987), and Baye, Kovenock, and de Vries (1996).

reducing their chances of winning. While a firm with high demand can increase its chance of winning by submitting a large number of bids, it cannot guarantee success.

Our focus on competitive rent seeking and the destruction of quota rents is rooted in the international trade literature. Model implications resemble those of Krueger (1974). In Krueger's political economy setting, firms spend resources to obtain import licenses (e.g. through bribery or lobbying). Our analysis shows that the H-1B lottery can cause the same type of rent destruction even in the absence of political economy or lobbying processes.

We also add to at least three strands of related work on lottery allocation, outsourcing, and employment concentration within the context of the H-1B program. First, while past studies including Kerr et al. (2010), Peri et al. (2015a), Sparber (2018), and Mayda et al. (2023) have argued that H-1B restrictions and the H-1B lottery harm aggregate productivity, wages, the selection of foreign labor, and firm outcomes, our paper is unique in arguing that the lottery allocation mechanism and outsourcing behavior combine to generate substantial quota rent losses associated with job search externalities. This inefficiency is not the simple deadweight loss that would typically arise due to an intervention into a well-functioning market. Rather, it is an additional cost specific to the lottery that would not arise through other allocation schemes. Second, our modeling of domestic outsourcing behavior (i.e., sending employees to third-party worksites within the United States) complements work by Glennon (2020), who finds that firms respond to H-1B restrictions by moving production to foreign affiliates. Third, our model argues that the quota, lottery allocation, and outsourcing lead to an increased concentration of H-1B employment among a small group of employers, thereby building a theoretical foundation to explain empirical evidence documented in Kim and Pei (2023) and Mayda et al. (2018).

Finally our model provides key economic insights for ongoing policy debates. Between 2018 and 2021, the US Department of Homeland Security (DHS), Department of Labor (DOL), and USCIS issued several memoranda and rule changes attempting to limit domestic outsourcing and to replace the lottery by allocating H-1Bs to applicants receiving the highest wage offers. Though courts overturned those policies on procedural or technical grounds, they did not rule against their adoption in the future (through legislation, for example).[6] Our paper identifies distinct costs generated by lottery allocation and exacerbated by outsourcing that should be of interest to policy-makers considering H-1B program modifications.

[6] A working paper version of this article provides more detail about legal developments. Also see National Law Review (2020a, 2020b), USCIS (2020b), Anderson (2020, 2021), and Chamber of Commerce of the United States of America v. United States Department of Homeland Security (2021).

2 H-1B Program Background

This section summarizes the process of applying for H-1B status, important program changes across time, and data trends. Three facts are especially pertinent for our model in Section 3: 1) Firms demand more new H-1B workers than the US government allows them to hire. 2) USCIS currently allocates H-1B status by random lottery. And 3) Some firms specialize in hiring H-1B workers and contracting their services to third-party sites.

2.1 H-1B Application Process

A firm wanting to hire an H-1B worker begins by filing a Labor Condition Application (LCA) with the US DOL attesting that it will abide by specific work conditions and wage requirements. There is no limit to the number of LCAs a firm can file, a single LCA can be used for multiple open job positions, and there is no LCA filing fee. An approved LCA serves as *de facto* government permission for a firm to include foreign workers in its job search.

When a firm offers a job to a qualified foreign-born worker, it must then file an I-129 petition with USCIS applying for H-1B status on the individual's behalf. These petitions are unlike LCAs in three important ways. First, I-129 forms are linked with specific individuals. Thus, they include information about both the employer and the prospective employee, and they provide a cleaner indication that the firm has identified a foreign worker it intends to hire. Second, I-129 submissions require basic filing fees plus additional charges for fraud detection, "Education and Training," and optional premium processing – all of which vary across firms. The Consolidated Appropriations Act of 2016 created an additional $4,000 fee for "H-1B Dependent Employers" – firms employing 50 or more employees with more than half of those workers on H-1B status. In total, a single submission cost a firm between $1,710 and $7,900 in FY 2020 (and up to $8,960 in FY 2023) before including any attorney fees. Third, the government has always imposed a statutory limit on the number of new H-1Bs issued per fiscal year.[7] USCIS must therefore reject an otherwise-valid I-129 petition for processing if it exceeds the quota.

2.2 Changes to the H-1B Quota and Allocation Procedure

The H-1B quota has changed over time. The 1990 American Competitiveness and Workforce Improvement Act (ACWIA) created the H-1B program and initially set it at 65,000. Congress increased the cap to 115,000 for fiscal years 1999 and 2000, then raised it again

[7] All employees of colleges, universities, and non-profit institutions are exempt from the cap, as are workers applying to renew their H-1B status.

to 195,000 for FY 2001-2003. The quota reverted back to 65,000 per year in FY 2004. It has held at 85,000 since FY 2006 when the H-1B Visa Reform Act added 20,000 H-1Bs for people who have obtained a masters degree or more education from a US institution.

In principle, USCIS has always allocated H-1B status on a first-come, first-served basis. In practice, allocation has depended upon fluctuations in both the statutory quota and foreign labor demand. Economists often exploit differences in firm and individual behavior across three allocation regimes to evaluate effects of H-1B restrictions. Though our model is concerned with the current policy environment, it is useful to understand each period.

The first lasted from the inception of the program in 1990 through 2003. The quota was effectively non-binding during these years. The market allocated H-1B status, with foreign workers going to whichever firms were willing to pay for their services. It experienced neither a deadweight loss nor quota rents.

The possibility for quota rents emerged in FY 2004 when the cap declined and became binding. USCIS initially continued its practice of allocating H-1B status to individuals on a first-come first-served basis. Firms were able to file petitions beginning on April 1 for the fiscal year starting six months later. USCIS approved early applications meeting all legal requirements. The comparably few petitions received on the last date of receipt were subjected to a random lottery for processing. This allocation regime was relatively short, operating from 2004-2007 and 2010-2013, and it taught firms that they needed to act quickly to hire the workers they desired. The filing season lasted 323 days for FY 2004, but just 56 days three years later. Although we do not formally model this era, rents would presumably have accrued to early-movers. Firms with specialties and expertise with the H-1B program started to employ a larger share of these workers during this era.

These lessons, coupled with increased demand for H-1B workers, gave rise to the third allocation regime (2008, 2009, and 2014-2024) – the current era of acute restrictions and severe rationing (as described in Mayda et al. (2018, 2023)). Though the quota has remained at 85,000 since FY 2006, demand has risen such that the number of cap-subject I-129 petitions submitted during the start of the application season now exceeds the total H-1B cap for the entire year. The first and last dates of receipt occur contemporaneously. USCIS's preferred first-come first-served model is no longer feasible, and it allocates all cap-subject H-1Bs with a random lottery. This paper's theoretical model seeks to understand the economic implications of H-1B policy during these lottery years.

2.3 H-1B Demand and Application Trends

Figure 1 illustrates H-1B filing trends since 2002. USCIS received roughly 89,000 cap-subject I-129 petitions when the quota exceeded the number of applications in FYs 2002 and 2003. That figure remained roughly constant in the early years of the reduced and binding cap. In FYs 2008 and 2009, however, USCIS received 123,000 and 163,000 petitions (respectively) in the first week of the application period. This rise in early H-1B submissions forced USCIS to abandon its first-come, first-served allocation policy. Applications surpassed 236,000 for FY 2017 before dipping to around 200,000 in FYs 2018-2020.

Although Figure 1 helps to form a foundation for understanding H-1B trends and our numerical exercises in Section 5, we caution that firms' H-1B submissions equal actual demand for new H-1B workers only for FY 2002 and 2003. This is due to two data limitations. First, firms with an interest in hiring an H-1B worker after national limits have been exceeded will never petition USCIS for H-1B approval. USCIS cannot measure this demand. However, Mayda et al. (2018) estimate that in the absence of the reduced H-1B cap, demand would have been 16-26% higher than the 85,000 cap in pre-lottery years (FYs 2004-2007) and 22-33% higher in the early lottery years (FYs 2008-2009). This would imply a maximum of new H-1B demand between 103,700 and 113,050 in each of these years, and yet firms applied for 163,000 new H-1B workers in FY 2009. This points to the second data limitation using I-129 petitions as a measure of H-1B demand. Section 3 will show that petitions rise simply because firms compete against each other for a labor source that is constrained by artificial limits and allocated by lottery. The number of jobs offered exceeds the number of employees actually demanded. This negative externality reduces welfare and destroys quota rents.

Figure 1 ends in FY 2020 due to a change in registration procedure beginning in FY 2021 that required firms filing cap-subject petitions to first electronically register with USCIS during March preceding the fiscal year. USCIS (2023) reported receiving 274,237 unique registrations for FY 2021; 308,613 for FY 2022; 483,927 for FY 2023; and a staggering 780,884 for FY 2024. Given this explosion in filings, the numerical analysis in this paper relies upon data before this procedural change and therefore characterizes 2017 as the year of peak demand. The quota, lottery, and firm behavior modeled by the theory are still in effect, however, so the model's implications still apply to the current policy environment.

2.4 Outsourcing Firms and H-1B Concentration Trends

Policy papers and the popular media commonly criticize the H-1B program for providing a vehicle for domestic outsourcing activities. Costa and Hira (2021) argue that just 30 firms accounted for nearly 40% of H-1B issuances in FY 2020, more than half of which used an outsourcing model. Park (2015) reported that "In recent years, global outsourcing companies have dominated the program, winning tens of thousands of visas and squeezing out many American companies, including smaller start-ups. 13 outsourcing companies took nearly one-third of all H-1B visas in 2014." North (2020) similarly argues that "outsourcing firms... control of [*sic*] 36 percent of the H-1B visas." Loten (2020), Torres (2017), and Fernandez Campbell (2016) provide further anecdotes describing how although some companies hire H-1B workers directly, many do so through business services companies including Accenture, Cognizant, and Infosys.

More systematic evidence is not immediately available since administrative datasets on the H-1B program do not definitively identify outsourcing firms. Moreover, other firms might provide outsourcing services even if that is not their primary business function. Instead, we use data on approved individual I-129 petitions from FYs 2002-2011 acquired through a Freedom of Information Act (FOIA) request to USCIS to illustrate a related firm behavior: Approved, cap-subject, H-1B workers are increasingly concentrated among fewer firms. For example, when H-1B issuances were unconstrained by the quota in FYs 2002 and 2003, 69% of cap-dependent firms seeking new H-1B workers filed only a single petition. These unique firm names accounted for 40% of all new cap-subject H-1Bs issued during that period. Contrast this outcome with FYs 2008 and 2009 when the quota had decreased to 85,000 and all new H-1Bs were allocated by lottery. Unique firm names petitioning for a single cap-subject H-1B worker accounted for 55% of firms but just 18% of new issuances. New H-1B employment was far more concentrated in a fewer set of firms during these lottery years.

Figure 2 provides further evidence of increased H-1B concentration by separately considering cap-subject firms versus colleges, universities, and non-profit research institutions. The latter employers were exempt from H-1B limits throughout this time period and therefore serve as a rudimentary control group.[8] The top panel displays the share of new H-1Bs awarded to employees of firms hiring five or fewer total H-1B workers in a given year. This value was above 50% in FYs 2002 and 2003 for cap-subject firms, but it declined precipitously after the quota became binding in 2004 and has fluctuated between 25% and 30% since then. The share among cap-exempt employers, however, remained fairly static at 20%.

Figure 2 also augments our FOIA evidence from 2002-2011 with aggregate statistics provided by the USCIS H-1B Employer Data Hub from 2009-2019. One limitation of this data is that it does not provide an identifier for whether a firm is subject to H-1B caps.

[8]Given the timing of H-1B petitions, we record all I-129 petitions beginning in April and moving forward as belong to H-1B caps for the fiscal year beginning in the subsequent October. Mayda et al. (2018) follow this strategy as well. The appendix describes an automated routine that we implement to correct potential name misspellings. Also, appendix Figure A1 displays trends in LCA filings that support observed I-129 behavior in Figure 2.

Instead, we proxy for cap-exempt employers as those with "UNIV"(except for those with "UNIVERSAL"), "COLLEGE", or "INSTITUTE" in their names. The two datasets reasonably resemble each other for the three overlapping fiscal years. The Data Hub data suggests a continuation of previous years' trends: Very few H-1B employees have worked for firms hiring five or fewer total H-1B workers in recent years.

The bottom panel displays the shares of new H-1Bs awarded to workers at firms hiring 250 or more total H-1B employees in a given year. H-1B employment concentration has perhaps declined among cap-exempt employers over time, but it rose dramatically from about 10% in FY 2002 and 2003 to over 35% in 2011 for cap-subject firms. Data Hub statistics suggest that this share has usually been above 40% in recent years.[9]

H-1B hiring has become a specialized endeavor. Our paper contends that increasing H-1B concentration – unique among cap-subject firms reflects firms' rational responses to restrictive limits on H-1B hiring. Though this concentration occurred before all cap-subject H-1Bs were allocated by lottery, we interpret this as evidence that H-1B outsourcing firms began to play a more dominant role in the program as soon as it became clear that new H-1B issuances would be restricted. Palagashvili and O'Connor (2021, p. 6) offer qualitative evidence supporting this view. They summarize a survey conducted among roughly 400 executives of small US technology firms by noting that "The responses from small startups indicate that changes to the H-1B visa program to make it more restrictive and costly will not necessarily lead small businesses to hire more Americans. Instead, small startups will continue to settle for hiring overseas contractors rather than hiring what they would prefer – foreign workers as employees." The following section develops a theory to model domestic outsourcing behavior and its cost implications.

3 Theoretical Model

This section builds a theoretical model to understand how H-1B caps affect firm behavior and welfare in an environment with lottery allocation and H-1B outsourcing activity. Before turning to the mathematics, it is helpful to understand the intuition that will drive the main results. The section closes with simple graphs illustrating predicted welfare losses relative to what would occur under alternative allocation mechanisms.

Imagine a lottery contest. Suppose the government decreases the number of winners and therefore the probability of any single entry (or bid) from winning. A competitor with a strong interest in winning might rationally respond by increasing its number of entries.

[9] Roughly 40% of H-1Bs are awarded to workers at firms employing more than 5 but fewer than 250 H-1B workers each year. This regularity has been steady throughout our time period.

This further reduces the probability of any single bid from winning and therefore imposes an externality on other competitors. If all competitors behave this way then it will create a multiplier effect that amplifies large costs. In the context of our paper, the competitors are firms that increase their number of lottery entries by offering an increasing number of jobs to foreign-born workers, many of whom will lose the H-1B lottery.

Entering the lottery is not a costless proposition. Firms wanting H-1B workers will need to conduct job searches, make job offers, complete government forms, and submit processing fees. Firms are heterogenous – some will find the procedure more cumbersome than others. Those finding the process to be particularly onerous might not participate at all.

Firms failing to hire their desired number of H-1B workers directly are not completely shut off from the foreign labor market. They can indirectly hire such workers by contracting with firms that employ surplus H-1B labor to provide desired services. These service providers include both domestic outsourcing specialists and firms that just happened to win the right to hire more H-1B workers than they had intended.[10]

With that intuition established, we begin the formal modeling by considering a firm (i) that faces CES demand for its product and can produce x_i units using a combination of native and foreign labor. In stage 1, the firm chooses the number of H-1B applications, F_{iA}, to maximize expected profits. In stage 2, the lottery outcome is realized and the firm learns its number of successful bids, $F_{iS} \leq F_{iA}$. The firm then chooses the number of immigrant workers (F_i) and native workers (N_i) to use. Thus, the firm – regardless of whether it specializes in providing outsourcing services or not – may interact with other firms to either rent out or contract for immigrant workers in the second stage. We can solve the firm's problem backwards starting from stage 2.

3.1 Stage 2

3.1.1 Setup

After the lottery outcome is realized, the firm learns its number of successful H-1B bids F_{iS}. It can hire up to this number of workers at immigrant labor wage w^F, a perfectly elastic reservation wage representing the payment actually received by H-1B workers. It can choose to either rent out or contract for immigrant labor at a wage rate of w^O, which in equilibrium will be greater than w^F. The firm also employs native-born workers at wage w^N.

[10]Our paper builds a theoretical model of outsourcing labor to third-party sites within the United States. Complementary but distinct empirical analysis in Glennon (2020) finds that restrictive H-1B policies cause firms to offshore jobs to their foreign affiliates, while Olney and Pozzoli (2021) estimate a negative relationship between immigration and firm-level offshoring in a Danish context.

3.1.2 Demand and Revenue

Firm i produces a differentiated product and faces CES demand derived from an underlying quasi-linear preference structure provided in Appendix B.1. Demand for good x_i is:

$$x_i = p_i^{-\sigma} \alpha_i \frac{E}{P^{1-\sigma}}, \tag{1}$$

where E is the total expenditure on all varieties and is constant; P is the price index for all varieties and is constant from the firm's perspective;[11] α_i is a taste parameter; and σ is the elasticity of substitution across products.

The firm's revenue generated by sales of its good (i.e., excluding any potential income from renting out excess labor) is:

$$R_i = \left(\alpha_i \frac{E}{P^{1-\sigma}} \right)^{1/\sigma} x_i^{\frac{\sigma-1}{\sigma}} \tag{2}$$

3.1.3 Production Function and Cost

Firms produce x_i according to a CES production function that combines native (N_i) and foreign (F_i) labor with an elasticity of substitution of ϵ, augmented by a firm-specific productivity parameter (A_i):

$$x_i = A_i \left(N_i^{\frac{\epsilon-1}{\epsilon}} + F_i^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \tag{3}$$

The cost of production accounting for the net cost of outsourcing is:

$$C_i = w^N N_i + w^F F_i + \left(w^O - w^F \right) (F_i - F_{iS}) \tag{4}$$

Note that $F_i - F_{iS}$ is the difference between the actual number of foreign workers the firm uses in the production process and the number of successful H-1B bids it has received. When positive, it therefore represents labor services supplied by third-party providers at wage rate w^O. These providers are likely to be outsourcing specialists, but they may also be regular firms that simply won the right to employ more workers than it intended to hire. Note that C_i can be negative since it includes revenues firms receive from renting out their H-1B workers to firms desiring more labor.

[11] The constant E setup is useful for expositional simplicity but could be relaxed without affecting the thrust of our analysis.

3.1.4 Profit and Factor Expenditure Functions

Subsection 3.2 will describe search costs incurred in the first stage. These costs are sunk in the firm's second stage decision. The firm's relevant profit function is to maximize the difference between the revenues in (2) and costs net of outsourcing in (4). Given the realized number of successful H-1B bids, the firm chooses optimal native and foreign employment fulfilling two first-order wage conditions:

$$\frac{\partial R_i}{\partial x_i}\frac{\partial x_i}{\partial N_i} = w^N \tag{5}$$

$$\frac{\partial R_i}{\partial x_i}\frac{\partial x_i}{\partial F_i} = w^O \tag{6}$$

We see from (6) that the marginal revenue product of immigrant labor is equal to the outsourcing wage at the optimum. This may be surprising given that the firm will generally use at least some labor from its own successful bids, in which case it would pay w^F instead of w^O. However, since H-1B workers can always be rented out to other firms, the opportunity cost of using a worker in production is still given by w^O for all firms regardless of whether they insource or outsource foreign labor.

The production function in (3) and wage conditions in (5) and (6) combine to deliver a standard marginal rate of technical substitution condition:

$$\frac{F_i}{N_i} = \left(\frac{w^N}{w^O}\right)^{\epsilon} \tag{7}$$

We next insert (7) into (5) and (6) to solve for the optimal expenditure on native and foreign-born labor:

$$w^N N_i = \frac{\left(w^N\right)^{1-\epsilon}}{\left(w^O\right)^{1-\epsilon} + \left(w^N\right)^{1-\epsilon}}\frac{\sigma-1}{\sigma}R_i \tag{8}$$

$$w^O F_i = \frac{\left(w^O\right)^{1-\epsilon}}{\left(w^O\right)^{1-\epsilon} + \left(w^N\right)^{1-\epsilon}}\frac{\sigma-1}{\sigma}R_i \tag{9}$$

These are standard CES factor expenditure functions, except with the opportunity cost of hiring foreign-born labor, w^O, representing the relevant wage for immigrants.

3.1.5 Profits

We can plug expenditure functions (8) and (9) into the cost expression (4), then state the firm's profits as:

$$\Pi_i = \frac{1}{\sigma} R_i + \left(w^O - w^F\right) F_{iS} \tag{10}$$

The above expression is similar to the standard CES profit maximizing choice except that the second term captures the benefit of having more successful H-1B submissions.

To calculate the optimal level of output in terms of the parameters of the model, first substitute the cost expressions (8) and (9) into the production function (3). Then insert the revenue function from (2). The resulting output function becomes:

$$x_i = \left(\frac{\sigma - 1}{\sigma} A_i\right)^{\sigma} \left(\alpha_i \frac{E}{P^{1-\sigma}}\right) \left[\left(w^N\right)^{1-\epsilon} + \left(w^O\right)^{1-\epsilon}\right]^{\frac{\sigma}{\epsilon - 1}} \tag{11}$$

Equation (11) gives us the firm's optimal choice as a function of wages as well as demand and productivity parameters from the firm and industry. This finally allows us to write profits as a function of parameters and factor prices:

$$\Pi_i = K_i \left[\left(w^N\right)^{1-\epsilon} + \left(w^O\right)^{1-\epsilon}\right]^{\frac{\sigma-1}{\epsilon-1}} + \left(w^O - w^F\right) F_{iS}, \tag{12}$$

where $K_i \equiv \left(\frac{\sigma-1}{\sigma} A_i\right)^{\sigma-1} \frac{1}{\sigma} \left(\alpha_i \frac{E}{P^{1-\sigma}}\right)$ is a constant from the standpoint of the firm. We can now use (12) in the firm's first stage problem.

3.2 Stage 1 Problem

3.2.1 Probability Distribution and Density Functions

In stage 1, the firm chooses the optimal number of H-1B applications, F_{iA}, so as to maximize expected profits, taking as given the optimal choices it will make in stage 2. In order to write down the firm's expected profits when accounting for search costs, we need to specify the stochastic process that determines the successful number of bids, F_{iS}. This poses a potential problem because while we can follow standard practice in ignoring any integer constraint on F_{iA}, we would still somehow have to account for the fact that the distribution of F_{iS} will be binomial with a constant probability (ρ) of winning the lottery. We deal with this challenge by using a continuous version of the binomial distribution (Ilienko, 2013), the cumulative density function (CDF) of which is:

$$G\left(F_{iS}; F_{iA}\right) = I_{1-p}\left(F_{iA} - F_{iS}, 1 + F_{iS}\right), \tag{13}$$

where $I_{1-\rho}(.)$ is the regularized incomplete beta function. The associated probability density function, $\frac{dG(F_{iS};F_{iA})}{dF_{iS}}$, exists but does not have a simple closed form. This continuous binomial distribution is a generalization of the regular binomial distribution that is also defined for non-integer values of F_{iA} and F_{iS} through the use of gamma functions in place of factorial-based expressions.

3.2.2 Expected Profits

Given the probability distribution function, we can write the firm's expected profits ($E[\pi_i]$) when accounting for search costs as:

$$E\left[\pi_i\right] = \int_0^{F_{iA}} \Pi_i\left(F_{iS}\right) dG\left(F_{iS}; F_{iA}\right) - cF_{iA} - c_i \frac{1}{\lambda} F_{iA}^{\lambda} \tag{14}$$

The general cost of filing H-1B petitions is represented by cF_{iA}. The search costs associated with finding F_{iA} potential H-1B workers is $c_i \frac{1}{\lambda} F_{iA}^{\lambda}$, with c_i serving as a firm-specific search cost parameter. We assume that that firms face either constant ($\lambda = 1$) or increasing ($\lambda > 1$) marginal costs of finding H-1B workers.[12]

By substituting the second-stage maximum profit condition (12) into the first-stage expected profit function (14), we get:

$$E\left[\pi_i\right] = K_i \left[\left(w^N\right)^{1-\epsilon} + \left(w^O\right)^{1-\epsilon}\right]^{\frac{\sigma-1}{\epsilon-1}} + \left(w^O - w^F\right)\rho F_{iA} - cF_{iA} - c_i \frac{1}{\lambda} F_{iA}^{\lambda}, \tag{15}$$

where ρF_{iA} is the mean of the binomial distribution. The marginal expected profit from filing an H-1B petition is then:

$$\frac{\partial E[\pi_i]}{\partial F_{iA}} = \left(w^O - w^F\right)\rho - c - c_i F_{iA}^{\lambda-1} \tag{16}$$

The first term represents the marginal benefit of a bid, which equals the premium an employer can earn by renting a worker to another firm $\left(w^O - w^F\right)$ multiplied by the probability that the bid will be successful (ρ). The latter two terms represent the marginal costs associated with searching for a prospective worker and filing a petition.

For firms facing increasing marginal search costs ($\lambda > 0$), we can set this expression equal to zero and solve for its optimal number of H-1B petitions:

$$F_{iA} = \left[\frac{\left(w^O - w^F\right)\rho - c}{c_i}\right]^{\frac{1}{\lambda-1}} \tag{17}$$

Intuitively, this value is increasing in the expected premium, $\left(w^O - w^F\right)\rho$, and decreasing in both the application cost c and the search cost parameter c_i. Perhaps surprisingly, the optimal number of bids is entirely independent of the firm's actual needs, which would

[12] We choose this specific iso-elastic functional form for these costs for simplicity and ease of interpretation. Qualitatively, any increasing marginal cost function would give us similar results.

depend on its production costs and demand for its output. This is because the firm's ability to rent out or contract for labor fixes the benefit of a successful bid to equal $w^O - w^F$ at the firm's optimum.

3.2.3 Regular and Outsourcing Firms

Although any firm can rent its labor to other firms, we exogenously separate firms into two types: Regular firms (with each firm i filing F_{iA} petitions) and outsourcing firms that specialize in providing labor to third-party sites (with each firm j filing F^O_{jA} petitions).

For simplicity, we assume that these firm types differ in that regular firms face increasing marginal search costs ($\lambda > 1$) and an optimal number of H-1B petitions governed by (17), whereas outsourcing firms face a constant marginal search cost ($\lambda = 1$) equal to c_X. On the surface, this dichotomy might appear to constitute a *deus ex machina* assumption, but we believe it accurately reflects the reality noted in Section 1 that many outsourcing specialists are headquartered overseas (particularly in India) and are likely to have a large pool of potential H-1B workers already on their payroll abroad. Thus, they would not face the same challenges as domestic firms when looking for more people to work in the US. Appendix B.2 demonstrates that this constant marginal cost assumption is a limiting case of a more general model where the outsourcing firms simply have a higher marginal search cost than regular firms.

With constant marginal costs and substituting corresponding F^O_{jA}, c_X, and $\lambda = 1$ values into (16), the marginal profits for outsourcing firms become constant and equal to:

$$\frac{\partial E[\pi_i]}{\partial F^O_{jA}} = \left(w^O - w^F\right)\rho - c - c_X \tag{18}$$

Outsourcers will never file an H-1B petition – and will therefore cease to exist – if marginal profits and the expression in (18) are negative. Since this outcome is inconsistent with real world observation, we largely dismiss it as a curiosity except to note that in the equilibrium identified in the next subsection, the sum of H-1B bids filed by outsourcing firms would also necessarily equal zero in this scenario (i.e., $\sum_j F^O_{jA} = 0$).[13] Conversely, the expression seemingly suggests that outsourcing firms will submit an infinite number of petitions if $\left(w^O - w^F\right)\rho > c + c_X$. Note, however, that this explosion in applications would drive down the probability that any one petition will be selected in a lottery with a fixed number of available H-1B positions. That is, ρ would collapse until (18) is no longer positive. Thus, the only stable equilibrium in which outsourcers exist occurs where $\frac{\partial E[\pi_i]}{\partial F_{iA}} = \frac{\partial E[\pi_i]}{\partial F^O_{jA}} = 0$ for regular and outsourcing firms alike. Whereas regular firms file H-1B petitions according to

[13] Appendix C.2 provides a lengthier discussion about implications in the absence of outsourcing firms.

(17) above, outsourcers do so with perfect elasticity such that the premium for renting out an H-1B worker is characterized by:

$$\left(w^O - w^F\right) = \frac{c + c_X}{\rho} \tag{19}$$

3.3 Equilibrium

3.3.1 Equilibrium Concept

The above analysis assumes that a firm faces uncertainty about the number of successful bids it receives (F_{iS}) but not about other values such as the probability of success (ρ) and the wage it will have to pay to H-1B workers, whether employed directly or through its decisions to contract for or rent out labor. This requires a rational expectations equilibrium concept where firms' choices regarding the number of submitted H-1B petitions are optimal given the realized wages and ρ in the second stage. We also assume that there is no aggregate uncertainty. Ultimately, the number of petitions filed, the probability of success, and wages paid are endogenous variables determined by the size of the H-1B quota and the model's cost and elasticity parameters.

3.3.2 Equilibrium Conditions

Assume that the total supply of immigrant labor would be perfectly elastic in a free market but is fixed by the H-1B quota, Ω. The supply of native workers is assumed to be constant and equal to $\bar{N}$. We can write the outsourcing wage in terms of the native wage by combining the optimal expenditure functions in (8) and (9) and then summing over all firms. This delivers a marginal rate of technical substitution relation that is the aggregate analogue to the firm-level equivalent expressed in (7):

$$\frac{\Omega}{\bar{N}} = \left(\frac{w^N}{w^O}\right)^{\epsilon} \tag{20}$$

We can also sum over firms' foreign labor expenditures in (9) and simplify to get:

$$w^O \Omega = \frac{\left(w^O\right)^{1-\epsilon}}{\left(w^O\right)^{1-\epsilon} + \left(w^N\right)^{1-\epsilon}} \frac{\sigma - 1}{\sigma} E \tag{21}$$

The model is closed by setting the successful number of bids equal to the actual H-1B quota:

$$\rho\left(\sum_i F_{iA} + \sum_j F_{jA}^O\right) = \Omega \tag{22}$$

Here, we separate the number of H-1B petitions submitted by regular firms i and those filed by outsourcing firms j. We can now use equations (17) and (19)-(22) to determine the endogenous variables F_{iA}, $\sum F_{jA}^{O}$, w^{O}, w^{N}, and ρ. Note that the distribution of H-1B petitions is indeterminate because firms make no economic profits from outsourcing activities and are therefore willing to provide any number of workers in equilibrium.

3.3.3 Solving for the Equilibrium

To solve for the equilibrium values of our endogenous variables, we first recall that the perfectly elastic bid-submissions of outsourcing firms identified in (19) determines the outsourcing premium – a net price for renting out H-1B workers. One can rearrange (19) to deliver the equilibrium probability of a successful bid: $\rho = \frac{c+c_X}{w^O - w^F}$. This value rises in proportion to outsourcing firms' marginal search costs and H-1B filing fees. The more difficult it is to find foreign workers, the fewer applications will be submitted, and hence the probability of winning will be larger. If the outsourcing premium $(w^O - w^F)$ is large, then firms will be incentivized to submit more H-1B applications, thus reducing the probability of winning any one individual bid.

We can then combine the outsourcing premium condition in (19) with the optimal H-1B petition condition in (17) to solve for the number of bids submitted by a regular firm. The resulting Equation (23) intuitively demonstrates that a firm submits fewer H-1B applications as its own marginal costs of searching for workers rise. It will submit more applications if outsourcing specialists' marginal search costs are high since these costs would be passed along to regular firms in the form of higher labor expenses.

$$F_{iA} = \left(\frac{c_X}{c_i}\right)^{\frac{1}{\lambda-1}} \tag{23}$$

Next combine (19), (20), and (21) to determine the outsourcing wage as a function of the H-1B quota, native employment, and the elasticity parameters of the model.

$$w^O = \frac{\sigma-1}{\sigma} E\left[\Omega + \left(\bar{N}\right)^{\frac{\epsilon-1}{\epsilon}} (\Omega)^{\frac{1}{\epsilon}}\right]^{-1} \tag{24}$$

The outsourcing wage is increasing in the total expenditure on goods (E) and decreasing in the quota-constrained supply of H-1B workers (Ω). It is also decreasing in the supply of native labor ($\bar{N}$) as long as native and immigrant workers are more substitutable than what would be implied by Cobb-Douglas production (i.e., if $\epsilon > 1$).

Finally, the simple supply and demand framework in (22) determines the total number of H-1B applications filed by outsourcing firms ($\sum F_{jA}^{O}$). Consistent with the discussion above, this sum total in (25) increases if the outsourcing premium is high. It decreases if marginal search costs, application fees, or the sum of regular firms' filings are high.

$$\sum_j F^O_{jA} = \frac{\Omega}{c + c_X} \left(w^O - w^F \right) - \sum_i F_{iA} \tag{25}$$

The model developed above assumes the existence of an H-1B quota, lottery allocation of H-1B status, and outsourcing firms. Section 4 will examine how the quota affects search costs and other firm behaviors in this environment. We will show that implications crucially depend upon the extent to which the quota induces firms to submit excess H-1B petitions. First, however, we turn to a graphical illustration of the job search externality and demonstrate that it exists only with lottery allocation.

3.4 Graphical Illustration of Alternative Allocation Mechanisms

A binding quota will generate potential quota rents and inefficiency, but lottery allocation entails costs on top of the more conventional deadweight loss resulting from quantitative restrictions. Figure 3 provides an illustration of these additional costs by comparing the lottery to alternative allocation mechanisms. All counterfactuals that we consider assume a stylized case such that marginal search costs for regular firms are so high that all hiring is performed by outsourcing firms. We impose this to simplify the graphical illustration.[14]

The left panel in Figure 3 illustrates a free market equilibrium for H-1B workers. Assume that foreign workers supply their labor with perfect elasticity at wage w^F_0. Since all hiring is performed by outsourcers in this diagram, firms face constant marginal search cost $c + c_X$. The intersection between firms' demand for H-1B workers and the marginal cost of *hiring* them (the wage plus the search cost) determines equilibrium H-1B employment at F^*.

The middle panel supposes that the government has an interest in limiting the number of new H-1B hires to Ω. Any intervention for achieving this outcome will generate a readily-identifiable deadweight loss inefficiency represented by the triangle with base $F^* - \Omega$ and height $(w^F_1 + c + c_X) - (w^F_0 + c + c_X)$. Box C (which is definitionally the same size as Box A) represents total search costs. Boxes A and B together represent quota rents: The rents people could earn if they had the exclusive right to buy at a cheaper external price and resell at the more expensive domestic price. These quota rents are simply transfers of surplus from firms to some other entity.

Consider how different policies would affect the distribution of these quota rents. For

[14]This section focuses on search costs arising from the lottery and outsourcing in the context of a quota. We do not illustrate how the quota leads to a rising concentration of H-1Bs at outsourcing firms. We address that question with a numerical exercise in Section 5.4.

example, the government could impose a tax (τ) on firms for each H-1B worker hired such that $\tau = w_1^F - w_0^F$. Quota rents A and B would then be transferred from firms to the government in the form of tax revenue. Alternatively, the government could institute an H-1B minimum wage equal to w_1^F. Under this policy, the quota rents would be awarded to H-1B workers in the form of higher wages. The most direct method of achieving the government's H-1B employment target would be to enact a quota equal to Ω representing the total number of new H-1B workers that can be hired. If labor is allocated according to firms' willingness to pay, then the quota rents accrue to H-1B workers just as under a minimum wage. This would be feasible if USCIS collected the I-129 petitions submitted at the beginning of the filing period, sorted them by wage offer, and approved H-1B status to workers at the top of the list.

Importantly, these scenarios represent counterfactuals to current US policy. In each case, the marginal *search* cost of a hired worker is the same as it would be without government intervention (i.e., the free market panel). The intervention simply raises the marginal cost of hiring a worker to $w_1^F + c + c_X$. Foreign workers earn this amount, less the marginal cost of searching.

Contrast these outcomes with current US policy illustrated in the third panel. In this scenario, the government imposes an H-1B quota but allocates workers according to a random lottery. Note that the deadweight loss triangle induced by the quota is identical to that from the alternative scenarios presented in the middle panel. The unique feature of lottery allocation is that it breaks the one-for-one link between the number of workers with job offers and the number who are employed. Firms extend offers to Ω/ρ individuals for just Ω available slots. Thus, the marginal search cost for a successful hire rises to $\frac{c+c_X}{\rho}$.

Firms pay H-1B workers the free market wage of w_0^F. Box C therefore continues to represent the total search cost for *hired* H-1B workers, but it no longer represents total search costs overall. That value must incorporate search costs for everyone who received a job offer. Since the market-clearing outsourcing wage must equal w^O, we know that the vertical distance $w^O - w_0^F$ equals $\frac{c+c_X}{\rho}$, just as predicted in Equation (19). In other words, the lottery creates new losses associated with a quota. Quota rents – firms' free market surplus in boxes A and B – are now entirely wasted on unnecessary search costs when they would have been distributed as welfare to foreign workers or as tax revenues to the government under alternative allocation systems.[15]

[15] In our setup, the outsourcing sector is competitive and earns zero equilibrium profits. This is analogous to competitive rent-seeking in the trade quota literature, which also leads to the complete destruction of quota rents (Krueger, 1974). If there were imperfect competition among outsourcing firms, this distortion would partly counteract the lottery externality and preserve a portion of quota rents as pure profits. This would similarly be the case in a model such as the one in Appendix B.2 where outsourcing firms also have

This waste – not the deadweight loss inefficiency – is the negative search cost externality that we seek to identify in this paper. Our numerical exercises in Section 5 estimate these losses. Parameters like $\frac{c+c_X}{\rho}$, w^F, and w^O are unobservable, however, so we will instead rely upon other methods to assess the search costs associated with excess demand for H-1B workers. The quota will be useful in backing out parameters necessary for that exercise.

4 Results: Theoretical Implications

Whereas Section 3 described equilibrium H-1B outcomes in the presence of a quota, lottery, and domestic outsourcing behavior, this section demonstrates that domestic outsourcing exacerbates negative externalities associated with the lottery. It then shows how a tightening of the quota would affect both concentration and search externalities in this environment.

4.1 Search Cost Implications of Outsourcing Firms

As long as the H-1B quota is binding, outsourcing firms cannot affect the total number of workers actually hired. The deadweight loss inefficiency is common to all quota allocation systems discussed in Section 3.4. We therefore focus instead on how outsourcing firms affect lost quota rents in the form of a search cost externality. One way to think about this issue is to consider possible policy actions and responses.

The government could increase outsourcing firms' costs by imposing restrictions or taxes on their hiring activity. Since this is specific to outsourcing specialists, we can characterize this as increasing c_X. The Consolidated Appropriations Act of 2016 enacted such a policy by charging H-1B dependent firms higher I-129 submission fees. Note, however, that these increased expenses also reduce regular firms' use of the service provided by outsourcing firms. This service has value since it allows firms to find immigrant labor at a lower cost than would otherwise be possible, but excess bids generate negative externalities. The net effect of such a restriction is therefore not obvious, *a priori*. By focusing on a continuous metric such as c_X to evaluate the effect of such restrictions, we can evaluate externality implications using simple differentiation. Our key result is summarized in Proposition 1.

Proposition 1 *Economy-wide H-1B search and bid costs are decreasing in outsourcing specialits' marginal search costs (c_X).*

To derive Proposition 1, first write out the economy-wide H-1B search and bid cost, C:

an increasing marginal search cost and earn some positive economic profit from their outsourcing activities.

$$C = \sum_i \left[cF_{iA} + c_i \frac{1}{\lambda} F_{iA}^{\lambda} \right] + \sum_j (c + c_X) F_{jA}^{O} \tag{26}$$

Note that the first summation in (26) captures the costs incurred by regular firms and the second reflects those incurred by outsourcing specialists. After substituting the total application behavior described in (23) and (25) into this aggregate cost function, we can differentiate with respect to c_X to find the key implication identified in (27): An increase in the marginal search costs of outsourcing firms reduces aggregate search costs:

$$\frac{dC}{dc_X} = -\sum_i F_{iA} < 0 \tag{27}$$

Appendix C.1 provides a complete derivation of this result. Its intuition begins by recognizing that higher marginal costs cause a relatively large decline in bids from outsourcing specialists. Regular firms increase their bids and fill some of this void (as predicted by (23)), but since those firms face increasing marginal search costs ($\lambda > 1$) whereas outsourcers experience constant costs ($\lambda = 1$), the magnitude of this response is much smaller. Even when combined with the rise in marginal costs for continued filings, these effects are less than the declining costs associated with the exodus among outsourcers. Government attempts to impose restrictions or taxes on outsourcing firms therefore decrease aggregate spending on searches.

A second possible policy action would be for the government to shut down outsourcing firms altogether. This need not involve a ban or prohibition, but could instead take the form of fees that increase the marginal costs of hiring for H-1B dependent employers so high that outsourcing firms choose to completely shut down on their own accord. Our modeling structure does not produce a closed form solution for what that limiting value of c_X would equal. However, Appendix C.2 provides numerical and graphical illustrations of the consequences of this policy. The most important insight is that the search cost externality still exists; the government cannot remove it simply by eliminating outsourcing firms. As long as the probability of randomly winning a bid is less than one, all firms face an incentive to search for more workers than they can hire. More search activity will therefore take place under lottery allocation than what a social planner would choose. Outsourcing specialists worsen, but do not create, this negative externality by submitting a higher number of bids. Moreover, the presence of outsourcing firms is harmful specifically because of the lottery. Under alternative rationing systems in which firms can hire all the workers they search for (such as those illustrated in the middle panel of Figure 3), outsourcing specialists are beneficial because they can find workers at a low marginal search cost. With a lottery, however, outsourcing firms increase the number of searches per hired worker by enough to increase total costs, undermining the potential benefits of these firms.

Since the incentive to submit petitions for more workers than a firm expects to hire is common among all employers in the presence of a quota with lottery allocation, a third policy counterfactual worth considering is a prohibition on all outsourcing activity, not just the practices of outsourcing specialists. Under this policy, firms can only use the services of their own directly-employed workers. They would not be permitted to rent their employees' services to third-party sites, nor could they contract to use other firms' workers. Appendix C.3 details this scenario. In contrast to the effect of shutting down outsourcing firms, the effect of shutting down outsourcing activity on bidding behavior is generally ambiguous. This follows from two conflicting mechanisms. On the one hand, there would be no benefit to having excess H-1B workers – a factor that reduces the marginal benefit of bidding. On the other hand, a firm winning fewer H-1B bids than expected would have no access to a secondary market to cover the shortfall – a factor that increases the marginal benefit of bidding. This version of the model is analytically intractable for reasons related to the continuous binomial probability distribution.

While this counterfactual is useful for thinking about the model's mechanisms, we also think the policy is unrealistic, notwithstanding recent USCIS, DOL, and DHS memoranda attempting to make it more difficult for H-1B labor to work at third-party sites. Foreign workers are scarce inputs in the production process. It is unlikely that a firm would decline to hire an H-1B worker simply because more of its potential hires won the lottery than it had anticipated. Trade in tasks models recognize that, in practice, there is little distinction between a firm selling its services to another company and renting out the services provided by its employees. Even if the government prohibited H-1B workers from physically working at third-party sites, it might not have any real effect if firms trade their goods and services after producing them at their own facilities. By comparison, the imposition of additional fees on outsourcing firms would be a more realistic policy action, the consequence of which is clearly identified by Proposition 1 and Equation (27).

4.2 Implications from Tightening The H-1B Quota

Section 2.2 and Figure 1 highlight changes to the H-1B quota that have occurred over time. Here, we describe consequences of tightening it in an environment characterized by lottery allocation. Although our primary interest is in lost rents, the intuition is best illustrated by beginning with a discussion of wage and H-1B employment concentration predictions.

First, differentiate the outsourcing wage in (24) with respect to the quota. Intuitively, increasing (i.e. loosening) the quota would decrease the outsourcing wage, w^O.

$$\frac{dw^O}{d\Omega} = -\frac{\sigma - 1}{\sigma} E \frac{1 + \frac{1}{\epsilon}\left(\bar{N}/\Omega\right)^{\frac{\epsilon-1}{\epsilon}}}{\left[\Omega + \left(\bar{N}\right)^{\frac{\epsilon-1}{\epsilon}} \left(\Omega\right)^{\frac{1}{\epsilon}}\right]^2} < 0 \tag{28}$$

Next, note that Equation (23) implies that the total number of H-1B petitions filed by regular firms is entirely independent of the quota. The hiring behavior of outsourcing firms identified in (25), however, is not. This latter equation is therefore key in understanding how government decisions to tighten H-1B restrictions will affect the number and concentration of H-1B petitions filed by outsourcing firms. By differentiating the aggregate number of outsourcers' bids in (25) and multiplying by -1, we get:

$$-\frac{d\sum_j F^O_{jA}}{d\Omega} = -\frac{1}{\rho} - \frac{\Omega}{c_X + c}\frac{dw^O}{d\Omega} \tag{29}$$

Two competing effects in (29) imply that a tightening of the H-1B quota could either increase or decrease the number of bids by outsourcing firms and therefore either increase or decrease the concentration of H-1B employment. The first is the direct effect of the quota; the second operates through the probability of securing a successful bid.

Consider the first channel captured by $-\frac{1}{\rho} < 0$. There is a mechanical relationship such that the win probability equals the quota divided by the total number of bids. Holding the probability constant, any decline in the quota would be offset by a decline in H-1B applications equal to $-\frac{1}{\rho}$. A lower quota means that firms are restricted to hiring fewer H-1B workers. Thus, there is a reduced need to rely on the services of outsourcing firms. Conversely, if Ω was large, then at a constant ρ the industry would hire a larger number of workers and would rely on outsourcing firms to avoid otherwise higher search costs.

The second channel, represented by $-\frac{\Omega}{c+c_X}\frac{dw^O}{d\Omega} > 0$, recognizes that the probability of success will not remain constant. Thus, there is an indirect effect that operates through this probability. Since a lower Ω will reduce ρ, it means that to hire a fixed aggregate number of H-1B workers, firms need to search for more labor and submit more bids. Increasing search costs faced by regular firms raise demand for services provided by outsourcing firms.

This probability channel is a lottery externality. To see why, consider a rather artificial but nonetheless useful counterfactual. Suppose that instead of individual firms choosing their optimal values of F_{iA} in a decentralized manner, a social planner chose the bids for each firm. In this scenario, the social planner is still subject to the lottery system but makes the decisions on behalf of firms. The planner would have no reason to submit bids in excess of available slots since any petition filed above Ω entails search and bid social costs that offer no corresponding social benefit. The externality arises from the fact that when a firm submits a bid, it internalizes the potential benefit to itself but not the cost to other firms in the form of a reduced probability of success. Any probability of success less than one reflects an inefficient use of aggregate resources that arises only in a lottery environment.

There is no *a priori* reason why one of the channels discussed above must dominate the other. However, we are able to sign the effect on the basis of a key elasticity, as summarized in Proposition 2.

Proposition 2 *A decrease (increase) in the H-1B quota increases (decreases) the number and share of applications submitted by outsourcing firms if the elasticity of the win probability with respect to the quota is greater than one.*

To derive Proposition 2, rearrange the supply and demand condition in (22), differentiate with respect to Ω, and multiply by -1 to get:

$$-\frac{d\sum_j F^O_{jA}}{d\Omega} = \frac{1}{\rho}\left(\frac{d\rho}{d\Omega}\frac{\Omega}{\rho} - 1\right) \tag{30}$$

The right hand side of this equation includes a term representing the elasticity of the win probability with respect to the H-1B quota, $\frac{d\rho}{d\Omega}\frac{\Omega}{\rho}$. This elasticity, which we henceforth represent with ψ, is the key parameter that drives the main implications of our model. The left hand side of (30) is positive so long as $\psi > 1$. The random allocation of H-1B status implies that a lower H-1B cap will cause an increasing concentration of H-1B employment among outsourcing firms when the elasticity of the win probability with respect to the quota is greater than one. The intuition of this result is straightforward. An elasticity of one would imply that the total number of bids remains constant when the quota is decreased so that the probability of success also decreases proportionately.

These insights now let us consider how the quota affects the costs of bidding, with the key result summarized in Proposition 3.

Proposition 3 *A decrease (increase) in the H-1B quota increases (decreases) total search and bidding costs if the elasticity of the win probability with respect to the quota is greater than one.*

Recall that Equation (26) identified total search and bidding costs and that its first summation term (which reflects hiring costs of regular firms) is independent of the quota – The quota affects search costs only through outsourcers' behaviors. By combining these implications with the equilibrium outsourcing premium in (19), we see that the H-1B quota affects the total costs of H-1B searches according to (31):

$$-\frac{dC}{d\ln(\Omega)} = \left(w^O - w^F\right)\Omega\left(\psi - 1\right) \tag{31}$$

Lottery allocation generates a search externality and lost rents associated with the quota. These losses would not occur under alternative allocation methods. They only arise when firms increase their bids in an effort to win the lottery, implying that the elasticity of the win probability with respect to the quota (ψ) is above one. Outsourcing specialists can alter cost magnitudes through ψ, but it is the lottery that creates the general incentive structure causing search costs to rise in response to quota restrictions.[16]

The next section will numerically show that the total number of bids likely increases due to the quota in reality, thus implying $\psi > 1$. This leads to a greater use of outsourcing firms, an increased concentration of bids, and a larger search externality beyond the losses that would occur under alternative allocation mechanisms.

5 Results: Numerical Exercises

The elasticity of the win probability with respect to the quota is the key parameter for determining many of the model's economic predictions. Its size is an empirical question, but data for performing regression analysis do not exist. In what follows, we propose different conceptual approaches to measuring this elasticity with available data. We then use alternative values of ψ to quantify how the H-1B quota affects search costs and H-1B employment concentration in the context of lottery allocation and domestic outsourcing.

5.1 Win Elasticity Equation

We gain insight into the elasticity of the win probability with respect to the quota by re-expressing it as a ratio of log differentials, $\psi = \frac{d\ln(\rho)}{d\ln(\Omega)}$, and comparing two different eras (t): One ($t = 0$) with a high and non-binding quota in which $\rho_0 = 1$, and another ($t = 1$) with constrained H-1B inflows allocated by lottery such that $\rho_1 < 1$. This comparison simplifies the numerator to $d\ln(\rho) = \ln(\rho_1) - \ln(1) = \ln(\rho_1)$ so that the percentage change in the probability of winning solely reflects the win rate under the lottery scenario.

The denominator is more complicated. While the probability of winning the lottery initially increases with the quota, ψ becomes perfectly inelastic in the limit and the probability of winning the lottery equals one for a large range of non-constraining Ω values. Thus, although we observe an H-1B cap of $\Omega = 85,000$ during the lottery era that far exceeds the high, non-binding, H-1B cap of $\Omega = 195,000$ from a period of H-1B certainty, it

[16] Appendix C.2 demonstrates that Equation (31) holds even in the absence of outsourcing specialists.

would be inappropriate to simply calculate the log quota differential between these periods as $d\ln(\Omega) = \ln(85,000) - \ln(195,000)$ since the win probability may have equalled one at much lower quota values during unconstrained years.

As a solution, we re-express the elasticity as a function of the change in log H-1B filings. First, let $\bar{F} = \sum_i F_{iA} + \sum_j F^O_{jA}$ represent the aggregate number of cap-subject H-1B petitions submitted by regular and outsourcing firms. Next, note that since the win probability is defined as $\rho = \frac{\Omega}{\bar{F}}$, we can rearrange, log-linearize, and partially differentiate to find $d\ln(\Omega) = d\ln(\rho) + d\ln(\bar{F})$. We can then rewrite our elasticity of interest as a function of the win probability when the quota is binding and the change in log H-1B petitions filed between the binding and non-binding periods:

$$\psi = \frac{d\ln(\rho)}{d\ln(\Omega)} = \frac{(\ln(\rho_1) - \ln(1))}{(\ln(\rho_1) - \ln(1)) + d\ln(\bar{F})} = \frac{\ln(\rho_1)}{\ln(\rho_1) + d\ln(\bar{F})} \tag{32}$$

Conceptually, $d\ln(\bar{F})$ represents the percentage-change in H-1B petitions attributable to the imposition of a binding quota – a supply shock – allocated by lottery. In calculating this value, one must therefore net out demand-driven changes in submissions. This task is complicated by the latent nature of demand described in Section 2. The following discussion considers three approaches to measuring $d\ln(\bar{F})$. All three are counter-factual exercises that reduce ψ to a function of two parameters: The win probability during the lottery year (ρ_1) and excess H-1B filings as a share (or percentage) of total filings ($\phi = \frac{\bar{F}-D}{\bar{F}}$).[17] Section 5.2 then computes win probability elasticities for values of ρ and ϕ suggested by the data.

5.1.1 Win Elasticity: Method 1

Our first method for calculating the percentage-change in filings due to a quota-induced supply shock compares excess filings in lottery years relative to the binding quota itself. This approach effectively assumes that pre-lottery filings are equal to the lottery era quota (an assumption supportable by the data in Figure 1). Note that by using the quota as its reference point, this method implies that $d\ln(\bar{F})$ reflects the growth in total filings (net of demand trends) relative to what the social planner would choose if confronted with a cap – a pollyannaish view of the lottery in which firms exhibit no strategic behavior after the quota is reached and make no attempt to submit additional H-1B petitions in order to secure a better chance of hiring the number of workers they actually desire. Mathematically:

$$d\ln(\bar{F}) = \ln(\phi + \rho_1) - \ln(\rho_1) \tag{33}$$

[17] The Appendix provides more detailed derivations of equations in this subsection. The three methods here are not exhaustive, but each delivers similar implications.

After substituting (33) into equation (32), we can express the win elasticity with respect to the quota as a simple function of the lottery-period win probability and excess H-1B filings as a share of total filings:

$$\psi = \frac{\ln(\rho_1)}{\ln(\rho_1 + \phi)} \tag{34}$$

5.1.2 Win Elasticity: Method 2

A second approach to comparing outcomes across eras is to assume that actual demand is constant so that all lottery-era filings exceeding this amount arise solely as a response to lottery allocation. The growth in submissions and the win elasticity become:

$$d\ln\left(\bar{F}\right) = -\ln(1-\phi) \tag{35}$$

$$\psi = \frac{\ln(\rho_1)}{\ln(\rho_1) - \ln(1-\phi)} \tag{36}$$

Interestingly, this is the same counter-factual one would use for comparing observed filings to a scenario in which the government set the quota equal to lottery-era demand but allowed non-random market mechanisms (e.g., willingness to pay) to allocate H-1B status. Thus, whereas the Method 1 computation of $d\ln(\bar{F})$ represents petition growth due to the lottery when the cap restricts H-1B availability below actual demand, Method 2 instead reflects petition growth due to lottery allocation when the cap is binding only at the margin.

5.1.3 Win Elasticity: Method 3

The final approach to measuring supply-induced changes in H-1B filings is much like the second, but instead represents the discrete-form percentage change in filings, holding demand constant. This implies that $d\ln(\bar{F})$ and ψ equal:

$$d\ln\left(\bar{F}\right) = \frac{d\bar{F}\big|_{D=\bar{D}}}{\bar{F}} = \phi \tag{37}$$

$$\psi = \frac{\ln(\rho_1)}{\ln(\rho_1) + \phi} \tag{38}$$

Importantly, note that $\phi \approx -\ln(1-\phi)$ if excess filings as a share of total filings are small. In that event, Methods 2 and 3 will yield both comparable petition growth and win elasticity calculations, which is fitting given the similarity in the setup of those approaches.

5.2 Win Elasticity Calculations

5.2.1 Win Elasticities Predicted by Petition Data

We use observed data, past evidence, and Equations (34), (36), and (38) to perform calibration exercises that calculate win probability with respect to the quota elasticities (ψ) for a range of win probabilities (ρ) and excess filings as a share of total filings (ϕ). The columns in Table 1 provide calculations for three years. The first two – 2008 and 2009 – were the first fiscal years in which all H-1Bs were allocated by lottery. Though only 85,000 new H-1Bs were available in those years, USCIS received more than 120,000 petitions during the first week of the application period for 2008 and roughly 163,000 for the following year. The third covers 2017 when firms submitted over 236,000 H-1B petitions – the peak year for H-1B applications prior to the DHS pre-registration policy for FY 2021 and beyond. In each of these years, H-1B applications almost certainly exceeded actual demand for new H-1B labor due to lottery allocation of quota-restricted H-1B status.

Each exercise in Table 1 requires specific values of ρ and ϕ. The former is a straightforward ratio of the cap (85,000) relative to the number of H-1B petitions filed in a given year. An individual applicant had about a two in three chance of winning the lottery in 2008, but just a one in three chance in 2017. To understand ϕ, we turn to evidence in Mayda et al. (2018) that H-1B employment was 22%-33% lower in FY 2008 and 2009 than what would have occurred in the absence of a quota. If we take these values to be reflective of actual demand, and if 85,000 represents actual H-1B employment, then firms demanded between 103,700 and 113,050 real H-1B workers in those years – figures much lower than the number of H-1B petitions filed and an indication that firms are conducting H-1B searches beyond their actual demand for foreign labor. The first three columns of Table 1 use the upper-bound demand estimate for 2008, 2009, and 2017; Columns (4)-(6) use the lower-bound figure. The gap between actual filings and these demand figures, normalized by total filings, then provides positive values of ϕ. These values are much lower when assumed H-1B demand is high.[18]

Panels are distinguished by the three subtly different counterfactual comparisons and methods for computing the change in log H-1B filings described in Section 5.1, which in turn affect the assumed value of the percentage change in the quota between the pre-binding and binding periods since $d\ln(\Omega) = d\ln(\rho) + d\ln(\bar{F})$. The top panel adopts Method 1 and computes filings relative to the 85,000 H-1B quota. It compares filings (net of demand) relative to what a social planner would do if confronted with the current quota. Methods

[18] We caution that 2017 is beyond the sample period included in the Mayda et al. (2018) analysis. Trends in employment demand might imply that even a 33% increase in real demand above the quota could be a lower-bound estimate for that year. Thus, our resulting calculation of excess filings might be larger than reality for 2017.

2 (middle panel) and 3 (bottom panel) compare actual filings relative to what the social planner would do if the cap was only marginally-binding. Method 2 does so by measuring filings relative to estimated contemporary real H-1B demand, whereas Method 3 equates the growth rate of H-1B submissions with excess filings as a share of total filings (ϕ).

In the absence of a lottery, $\psi = 1$ and a 1% decrease in the quota would generate a 1% decrease in H-1B submissions. But lottery allocation gives rise to strategic behavior implying that $\psi > 1$. Table 1 illustrates this by reporting a range of ψ calculations. At a minimum, we find a value of 1.29 in 2008 when excess H-1B filings were comparatively modest and we assume that real demand for foreign labor was high. The elasticity is always smallest when using Method 3, which delivers quantitatively similar results to using Method 2. Elasticities are highest when calculating filings relative to the quota (Method 1). Not surprising, the largest elasticity values arise in 2017 when the number of H-1B petitions was particularly high. Most importantly, however, all values are comfortably above one. This result implies that lottery allocation causes H-1B restrictions to generate lost quota rents associated with negative job search externalities (as predicted by Equation (31)), and it increases H-1B employment concentration among outsourcing firms (as predicted by Equation (30)). These effects would not occur under a quota alone and are entirely distinct from the deadweight loss inefficiency common among alternative forms of government intervention into the labor market.

Next, we turn to algebraic and graphical analysis for understanding the interplay between ϕ, ρ, and ψ. This is somewhat complicated in the sense that the parameter values are endogenously determined by the model – this is not a typical comparative statics exercise in which one simply changes an exogenous variable in order to examine the consequences of a shock. Nonetheless, we believe the graphs are informative about how the endogenous variables are related to each other and are helpful in identifying elasticity values consistent with real world data.

It is straightforward to show that the elasticity ψ is increasing in both ϕ and ρ, conditional upon $\phi \in [0, 1-\rho]$ and $\rho \in [0, 1]$, where the former range restriction arises from an assumption that real H-1B demand in lottery years equals the size of the quota at a minimum. Consider the intuition behind these results in turn. The first is almost tautological. It implies that if excess filings are low, then a change to the quota will lead to an almost proportional change in the win probability. High values of ψ, in contrast, arise when firms file a large number of H-1B petitions in excess of demand, thus implying that the win probability is very responsive to quota restrictions. The importance of this result becomes more apparent when recognizing that ψ is guaranteed to be greater than one so long as any firm submits more H-1B petitions in excess of their actual demand for H-1B workers. Moreover, any elasticity above one is indicative of a negative externality that is sufficiently large to induce an increase in the total number of bids. Even a small proportion of H-1B petitions filed in excess of actual H-1B demand would be enough to generate such costs.

The positive relationship between ρ and ψ implies that when firms are confident that they can win the lottery, they respond to quota restrictions by conducting an excess number of searches and extending multiple offers in an attempt to secure their desired number of workers. Collectively, this causes the actual win probability to drop precipitously.

Figure 4 plots elasticity values as a function of the percentage of excess H-1B applications that arise from Methods 1-3 described above. Each graph displays contour plots for different values of ρ increasing in increments of 0.05 from a minimum of $\rho = 0.30$ (light colored) to a maximum of $\rho = 0.75$ (dark). Each panel also includes point values arising from observed win probabilities and the number of excess H-1B filings implied by the given methodology. The red line corresponds to the assumption that real H-1B demand is 22% above the quota; the blue line assumes demand is 33% higher than the quota. Graphs share the same vertical axis scale and are displayed side-by-side for comparison purposes.

Not surprisingly, all three methods demonstrate that the win elasticity with respect to the quota is high when firms submit a large number of petitions in an effort to win the lottery. Methods 2 and 3 produce quantitatively similar elasticities, though Method 3 has the advantage of producing values that are the least sensitive to differences in ϕ and ρ values. Using real world data, we find a minimum elasticity of 1.29 and a maximum of about 2.25. Despite this variation, Section 5.3 will demonstrate that the three methods produce similar values for search cost externalities resulting from the quota and lottery.

5.2.2 Win Elasticities Predicted by Wage Data

Table 1 provides a range of reasonable win probability elasticities based upon H-1B filing behavior, but we can also gain insight about ψ by appealing to wage data. We begin by noting that if native and foreign labor is relatively substitutable, the win probability elasticity can be approximated by Equation (39).[19]

$$\psi \approx \left(\frac{w^O \Omega}{w^O \Omega + w^N \bar{N}} \right) \left(\frac{w^O / w^F}{w^O / w^F - 1} \right) \tag{39}$$

The first term is the share of wages paid to H-1B workers in the sectors of the economy that employ them. The higher this value, the larger the elasticity: As H-1B workers become more important to the production process, firms become more responsive to changes in the

[19]See the Appendix for a full derivation.

H-1B quota and are more likely to aggressively change their number of excess H-1B filings. The second term is a transformation of the outsourcing premium $\frac{w^O}{w^F}$. The win probability elasticity decreases as this wage gap rises and is governed solely by H-1B workers' share of income at the limit.

If we could measure H-1B workers' share of income and the $\frac{w^O}{w^F}$ relative wage, we could back out an estimate for our key probability elasticity. Unfortunately – and like the measures described above – they are unknowable. We cannot observe the wages of all current H-1B workers in representative datasets, and given the small size of the H-1B program relative to the US economy, we believe it would be a mistake to insert a literal measure of total H-1B expenditures as a share of the national wage bill into Equation (39). Similarly, the ratio $\frac{w^O}{w^F}$ is conceptual and reflects the gap between what employers earn from renting out their workers and what they actually pay their H-1B workers. This too is unobservable.

Despite these limitations, we can turn to proxies in the data for insight. First, note that H-1B workers are highly concentrated among computer-related occupations and 91% of new H-1B recipients are under age 40.[20] Motivated by these stylized facts, we turn to the 2008 and 2009 American Community Surveys (ACSs) and calculate that foreign-born labor accounted for 40% of wages paid to college-educated employees under age 40 in computer-related occupations. We then substitute this value for our first term in (39).

Second, let the conceptual wage gap discussed above be similar to the gap in wages paid by regular firms and those paid by low-wage H-1B dependent employers. The H-1B Visa Reform Act of 2004 set a *de facto* minimum wage of \$60,000 for the latter group, and wage offers exhibit a well-populated mode at this value. We set $w^F = \$60,000$ accordingly. For w^O, we instead use the average wage offer to new H-1B workers receiving more than \$60,000 in fiscal years 2008 and 2009. After removing extreme outliers, this amounts to \$83,500 and implies a wage gap of $w^O - w^F = \$23,500$ and an outsourcing premium of $\frac{w^O}{w^F} = 1.39$.[21] After substituting this ratio into the second term of (39), the probability elasticity with respect to the quota becomes $\psi = 1.43$. This is consistent with the lower range of elasticities in Table 1 calculated by exploiting differences in firm H-1B filing behavior over time. Most importantly, it still delivers an elasticity greater than one. The condition required for search cost externalities and increased H-1B concentration as a result of the lottery is therefore easily satisfied through evidence from both H-1B petition and wage data.

[20] See USCIS (2020a) and Mayda et al. (2018).

[21] Note also that in his February 2016 testimony to the US Senate, Hira (2016) argued that firms heavily relying upon H-1B workers generate profit margins of 20-25%, figures broadly consistent with this premium.

5.3 Marginal Cost and Search Cost Externality Calculations

5.3.1 Marginal Bid and Search Costs

The wage data outlined in Section 5.2.2 is also useful for evaluating the implied bid and search costs for an outsourcing firm. Equation (19) expresses the relationship between the outsourcing premium, marginal costs, and the win probability. If we assume a wage gap of $w^O - w^F = \$23,500$ (as above) and lottery success rate of 0.50 (comparable to the lottery in 2009), then it implies a marginal search cost of $c + c_X = \$11,750$. If we take a literal interpretation of application cost and we assume a \$2,000 charge near the lower range of the H-1B filing fees discussed in Section 2, then an outsourcing firm's marginal cost of searching for and extending an offer to a worker, c_X, equals roughly \$9,750.

5.3.2 Search Cost Externality

Equation (31) describes how lottery allocation of quota-restricted H-1B status generates lost rents in the form of a search cost externality. We can rearrange this equation to express cost implications as a function of increased filings.

$$dC = \left(w^O - w^F\right)\ \Omega\ d\ln\left(\bar{F}\right) \tag{40}$$

Given the wage gap and quota described above, calibrated total costs depend only upon the method used to compute filings growth and ψ outlined in Section 5.1. Interestingly, Equation (40) plus the filings calculations in (33), (35), and (37) imply that the value of the win probability itself plays no role in identifying these costs when using Methods 2 and 3. Not surprisingly, all three methods find that excess filings increase search externalities.

Figure 5 displays search costs in millions of dollars as a function of the percentage of excess filings (ϕ). Panels are distinguished by their use of Methods 1 through 3 for calculating the elasticity ψ. Graphs also include point values implied by observed data and Mayda et al. (2018) estimates of H-1B demand for each fiscal year in which a lottery allocated all cap-subject H-1Bs. (Table 1 provides a subset of these calculations.) Red dots assume that actual H-1B demand was 22% above the quota; blue dots assume 33%.

As with our computations of ψ, smaller search cost externalities arise when real demand is assumed to be large, when using Method 3 for calculating ψ, and when the win probability is high (69% in 2008). They are largest when assuming low levels of real demand, Method 1's calculation of ψ, and low win probabilities. Nonetheless, results are similar across methods. Our most conservative calculations imply that even when excess filings were a modest 8.4% of all H-1B petitions submitted for fiscal year 2008, firms' strategic behavior resulted in at least a $169 million search cost externality, or about $2,000 per new H-1B worker hired.[22] Costs were substantially higher in 2009 when nearly 31% of firms' H-1B submissions were filed in excess of demand: Quota rents lost via this externality range between $612 million and $1.057 billion for that year – roughly $7,000 to $12,000 per hire. If real H-1B demand remained at a constant 33% above the quota for years in the late 2010s, then annual search costs exceed $1 billion in that period. In 2017 when losses were maximized, for example, this exercise implies a search cost externality of $12,000 to $22,000 per new H-1B worker. Given that firms must actually search for these workers and file paperwork with the government on their behalf, we believe this range of values is a reasonable reflection of reality.

These costs arise purely because the lottery incentivizes firms to search for an excess number of H-1B workers to meet their hiring target; they are not costs associated with losing workers that firms actually wish to hire. Note the absence of an externality when $\phi = 0$. If firms submit H-1B petitions only for workers they can actually hire, then costs associated with attempts to win the lottery do not exist. The data, however, strongly show that firms engage in strategic behavior, resulting in large costs to society. Specifically, Table 1 suggests that the US government forfeited $1.042 billion to $1.879 billion in 2017 that could have been retained as tax revenue or distributed as welfare to foreign workers under alternative methods of restricting H-1B flows. Instead, the lottery destroyed those rents.

5.4 Concentration Implications

In addition to generating unnecessary search cost externalities, the H-1B lottery also leads to an increased concentration of H-1B employment among a small set of firms. To quantify this, let $\theta = \frac{\sum_j F^O_{jA}}{\bar{F}}$ represent the share of H-1B approvals awarded to outsourcing firms. By combining the response in outsourcers' filings identified in (30) with the identity $d\ln(\Omega) = d\ln(\rho) + d\ln(\bar{F})$, we can determine how the H-1B quota restriction affects the proportion of H-1B approvals granted to H-1B outsourcing specialists:

$$-\frac{d\theta}{d\ln(\Omega)} = (\psi - 1)(1 - \theta) \tag{41}$$

This expression states that H-1B concentration among outsourcing firms increases linearly with the win elasticity ψ and is positive so long as $\psi > 1$, which – as we have seen above – is likely. The marginal change in this share decreases as the share itself approaches one.

To get a sense of the magnitude of the prediction in (41), first consider possible values of θ from the data. Individuals and organizations generally opposed to the H-1B program

[22] Since currency figures are derived from nominal wage premia from 2008 and 2009, it would be reasonable to interpret these outcomes as reflecting losses in real 2008 or 2009 dollars.

argue that at least a third of cap-subject H-1B workers are hired by outsourcing firms.[23] Thus, $\theta = 0.33$ serves as a useful lower bound. An upper bound of $\theta = 0.60$ would reflect the percentage of firms that pay H-1B workers low wages according to Costa and Hira (2020).[24]

Table 1 provides calculations of the change in the outsourcing employment share for 2008, 2009, and 2017 at these values of θ. Focusing on the first three columns and the middle panel, we infer that the quota caused the proportion of H-1Bs awarded to outsourcing firms to rise by 4 to 6 percentage points in 2008, depending upon the assumed starting value of θ. This effect is much larger and sometimes exceeds 50 percentage points in 2009 and 2017 when the probability of winning the lottery was much lower and the elasticity of the win probability with respect to the quota was much higher. Our third method (bottom panel) for calculating the change in log H-1B filings produces the lowest elasticity values. It therefore also produces the most muted response in the share of H-1Bs awarded to outsourcing firms. At its maximum, this method predicts that outsourcers controlled 24 percentage points more of the H-1B market in 2009, and 37 percentage points more of the market in 2017, than they would have in the absence of a lottery.

Further insight into these magnitudes comes from considering that the top 1% of H-1B employers accounted for 24.3% of new H-1B employment in 2002 and 2003 when the H-1B cap was high and non-binding, and 36.7% of H-1Bs in the 2008 and 2009 lottery years. If we think of these firms as outsourcers, then this implies a 12.4 percentage point increase in θ. If we broaden our definition to include the top 5% of H-1B employers, then θ rose by nearly 20 percentage points (from 38.0% to 57.7%) between these two periods. Thus, these descriptive statistics fall within the range of values predicted by our theoretical model.

6 Conclusion

The H-1B program allows highly-educated foreign-born individuals to temporarily work in the United States. The government limits the number of new H-1Bs that can be distributed each year. Demand for the program is high, and the US allocates H-1B status through a random lottery. H-1B petitions peaked at over 236,000 in 2017 when only about a third of petitioners won the lottery.

Past research has estimated several losses induced by the H-1B quota and lottery distribution. By restricting the skilled labor force, the quota reduces US productivity and wages paid to American-born workers. Sales and profits have declined. H-1B employment

[23] See North (2020) and Park (2015).

[24] Namely, the authors find that "Sixty percent of H-1B positions certified by the US Department of Labor are assigned wage levels well below the local median wage for the occupation" and that "three-fifths of all H-1B jobs were certified at the two lowest prevailing wage levels in 2019."

is increasingly concentrated. Since firms receive permission to hire only those who win the lottery, lottery allocation prevents employers from selecting the foreign workers whom they most desire. Some firms have responded by moving operations overseas.

This paper adds to the existing evidence by modeling heretofore undocumented costs induced by the quota and lottery: Firms waste resources and destroy quota rents by searching for workers who cannot legally be hired. Indeed, firms search for far more workers (and submit H-1B petitions on their behalf) than they truly intend to hire. They do this knowingly in an attempt to better their chances of securing enough workers who will win the lottery and satisfy actual H-1B demand. Moreover, the quota and lottery have concentrated H-1B employment among outsourcing firms that specialize in hiring H-1B labor for the purpose of renting their services to other firms. This exacerbates negative search cost externalities.

We estimate that lottery allocation and outsourcing behavior cost American firms $169 million in search costs in 2008, more than $600 million in 2009, and more than $1 billion in the late 2010s when the probability of winning the lottery was very low. The loss was over $10,000 per H-1B worker hired in 2017. Importantly, these are not the typical deadweight losses that exist under other methods of reducing H-1B labor inflows. Rather, they are lost quota rents due to the strategic behavior of firms responding to lottery allocation.

References

[1] Anderson, Stuart (2020). "USCIS Rescinds Policy Memos That Increased H-1B Visa Denials." *Forbes.* June 22, 2020. Retrieved at https://www.forbes.com/sites/stuartanderson/2020/06/22/uscis-rescinds-policy-memos-that-increased-h-1b-visa-denials/#669df4d827eb

[2] Anderson, Stuart (2021). "Judge Kills The Last Trump H-1B Visa Rule Left Standing." *Forbes.* September 17, 2021. Retrieved at https://www.forbes.com/sites/stuartanderson/2021/09/17/judge-kills-the-last-trump-h-1b-visa-rule-left-standing/?sh=141203647340

[3] Baye, Michael R., Dan Kovenock, and Casper G. de Vries (1996). "The All-Pay Auction with Complete Information." *Economic Theory*, Vol. 8: 291-305.

[4] Bernstein, Shai, Rebecca Diamond, Timothy James McQuade, and Beatriz Pousada (2012). "The Contribution of High-Skilled Immigrants to Innovation in the United States." Stanford Graduate School of Business Working Paper No. 3748.

[5] 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*, Vol. 127(3): 1143-1203.

[6] Borjas, George J., Jeffrey Grogger, and Gordon H. Hanson (2012). "Comment: On Estimating Elasticities of Substitution." *Journal of the European Economic Association*, Vol. 10(1): 198-210.

[7] Bound, John, Breno Braga, Joseph M. Golden, and Gaurav Khanna (2015). "Recruitment of Foreigners in the Market for Computer Scientists in the United States." *Journal of Labor Economics.* Vol. 33(S1): S187-223.

[8] Brito, Dagobert L. and Michael D. Intriligator (1987). "Stock Externalities, Pigovian Taxation and Dynamic Stability." *Journal of Public Economics*, Vol. 33: 59-72.

[9] Costa, Daniel and Ron Hira (2020). "H-1B Visas and Prevailing Wage Levels." Economic Policy Institute Report. May 4, 2020. Retrieved at https://files.epi.org/pdf/186895.pdf

[10] Costa, Daniel and Ron Hira (2021). "The H-1B Visa Program Remains the 'Outsourcing Visa."' Economic Policy Institute Report. March 31, 2021. Retrieved at https://www.epi.org/blog/the-h-1b-visa-program-remains-the-outsourcing-visa-more-than-half-of-the-top-30-h-1b-employers-were-outsourcing-firms/

[11] Chamber of Commerce of the United States of America, et al. v. United States Department of Homeland Security, et al. (2021). United States District Court, Northern District of California. Case No. 20-cv-07331-JSW, Document 157

[12] DHS (2011). "Registration Requirement for Petitioners Seeking to File H-1B Petitions on Behalf of Aliens Subject to the Numerical Limitations." *Federal Register*, Vol. 76(42): 11686-11699. March 3, 2011.

[13] DHS (2018). "Registration Requirement for Petitioners Seeking To File H–1B Petitions on Behalf of Cap-Subject Aliens." *Federal Register*, Vol. 83(232): 62406-62447. December 3, 2018.

[14] DHS (2021). "Modification of Registration Requirement for Petitioners Seeking To File Cap-Subject H-1B Petitions." *Federal Register*, Vol. 86(5): 1676-1735. January 8, 2021.

[15] Depew, Briggs, Peter Norlander, and Todd A. Sorensen (2017). "Inter-firm Mobility and Return Migration Patterns of Skilled Guest Workers." *Journal of Population Economics*. Vol 30: 681-721.

[16] Douglas, Genevieve (2021). "Debate Over H-1B Lottery Revamp Heats Up as Cases Go Forward." *Bloomberg Law*. June 1, 2021. Retrieved at https://news.bloomberglaw.com/daily-labor-report/debate-over-h-1b-lottery-revamp-heats-up-as-cases-go-forward

[17] Dustmann, Christian and Ian Preston (2012). "Comment: Estimating the Effect of Immigration on Wages." *Journal of the European Economic Association*, Vol. 10(1): 216-223.

[18] Fernandez Campbell, Alexia. (2016). "There's a Clear Way to Fix the H-1B Visa Program." *The Atlantic*, 6 December 2016.

[19] Doran, Kirk, Alexander Gelber, and Adam Isen (2022). "The Effect of High-Skill Immigration Policy on Firms: Evidence from H-1B Visa Lotteries." *The Journal of Political Economy*, Vol. 130(10): 2501-2533.

[20] Glennon, Britta (2020). "How do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program." Mimeo. University of Pennsylvania. http://dx.doi.org/10.2139/ssrn.3547655

[21] Gordon, H. Scott (1954). "The Economic Theory of a Common-Property Resource: The Fishery." *The Journal of Political Economy*, Vol. 62(2): 124-142.

[22] Gunadi, Christian (2019). "An Inquiry on the Impact of Highly-Skilled STEM Immigration on the U.S. Economy." *Labour Economics*. Vol. 61.

[23] Hira, Ronil (2016). Testimony on the impact of high-skilled immigration on US workers. US Senate Subcommittee on Immigration and the National Interest . February 25, 2016

[24] Hunt, Jennifer. (2011). "Which Immigrants are Most Innovative and Entrepreneurial? Distinctions by Entry Visa." *Journal of Labor Economics*. Vol. 29(3): 417-457.

[25] Hunt, Jennifer and Marjolaine Gauthier-Loiselle. (2010). "How Much Does Immigration Boost Innovation?" *American Economic Journal: Macroeconomics*. Vol. 2(2): 31-56.

[26] Hunt, Jennifer and Bin Xie (2019). "How Restricted is the Job Mobility of Skilled Temporary Work Visa Holders?" *Journal of Policy Analysis and Management*. Vol. 38(1): 41-64.

[27] Ilienko, Andrii. (2013). "Continuous Counterparts of Poisson and Binomial Distributions and their Properties." *arXiv preprint arXiv:1303.5990.*

[28] Kerr, Sari Pekkala, William Kerr, and William F. Lincoln (2015). "Skilled Immigration and the Employment Structures of US Firms." *Journal of Labor Economics*. Vol. 33(S1): S147-S186.

[29] Kerr, William and William F. Lincoln. (2010) "The Supply Side of Innovation: H-1B permit Reforms and U.S. Ethnic Invention." *Journal of Labor Economics* 28(3): 2010, pp. 473-508.

[30] Kim, Seohee and Alison Pei (20123). "Monopsony in the High-Skilled Migrant Labor Market – Evidence from H-1B Petition Data." Working Paper Available at SSRN: http://dx.doi.org/10.2139/ssrn.4010152.

[31] Lange, Andreas, John A. List, and Michael K. Price. (2007) "Using lotteries to finance public goods: Theory and experimental evidence. " *International Economic Review*. Vol. 48(3): 901-927.

[32] Landry, Craig E., Andreas Lange, John A. List, Michael K. Price, and Nicholas G. Rupp (2006). "Toward an understanding of the economics of charity: Evidence from a field experiment. " *Quarterly Journal of Economics*. Vol. 121(2): 747-782.

[33] Loten, Angus. (2020). "H-1B Overhaul Pressures Outsourcing Firms and Their U.S. Clients." *The Wall Street Journal*, 8 October 2020.

[34] Ma, Jie. (2020). "High Skilled Immigration and the Market for Skilled Labor: The Role of Occupational Choice." *Labour Economics*. Vol. 63(C).

[35] Manacorda, Marco, Alan Manning, and Jonathan Wadsworth (2012). "The Impact of Immigration on the Structure of Wages: Theory and Evidence from Britain." *Journal of the European Economic Association*, Vol. 10(1): 120-51.

[36] Mayda, Anna Maria, Francesc Ortega, Giovanni Peri, Kevin Shih, and Chad Sparber (2018). "The Effect of the H-1B Quota on Employment and Selection of Foreign-Born Labor." *European Economic Review*, Vol. 108(C): 105-128.

[37] Mayda, Anna Maria, Francesc Ortega, Giovanni Peri, Kevin Shih, and Chad Sparber (2023). "Coping with H-1B Shortages: Firm Performance and Mitigation Strategies." *ILR Review*, 76(5), 919-943. https://doi.org/10.1177/00197939231165770.

[38] Morgan, John. (2000). "Financing public goods by means of lotteries." *The Review of Economic Studies*. Vol. 67(4): 761-784.

[39] Moser, Petra, Alessandra Voena, and Fabian Waldinger. (2014). "German Jewish Émigrés and US Invention." *American Economic Review*, Vol. 104(10): 3222-55.

[40] National Academies of Sciences, Engineering, and Medicine 2017. The Economic and Fiscal Consequences of Immigration. Washington, DC: The National Academies Press. https://doi.org/10.17226/23550.

[41] National Foundation for American Policy (2020). "H-1B Denial Rates and Numerical Limits as Indicators of Current Restrictions." NFAP Policy Brief, May 2020. Retrieved at https://nfap.com/wp-content/uploads/2020/05/H-1B-Denial-Rates-and-Numerical-Limits-as-Indicators-of-Current-Restrictions.NFAP-Policy-Brief.May-2020-1.pdf

[42] National Law Review (2020a). "Rescission of USCIS Policy Memoranda in Favor of H-1B Petitioners in the IT Third Party Placement Industry." Vol. X (188). June 22, 2020.

[43] National Law Review (2020b). "Recent USCIS Settlement Offers Substantial Relief to H-1B Employers." Vol. X (188). July 6, 2020.

[44] North, David (2020). "The Government Announces Stricter H-1B Provisions for Outsourcing Firms."Press Release. Center for Immigration Studies. October 6, 2020.

[45] Olney, William W. and Dario Pozzoli (2021). "The Impact of Immigration on Firm-Level Offshoring." *The Review of Economics and Statistics*, Vol. 103(1): 177-195.

[46] Ottaviano, Gianmarco I. P. and Giovanni Peri (2012). "Rethinking the Effect of Immigration on Wages." *Journal of the European Economic Association*, Vol. 10(1): 152–97.

[47] Office of the Federal Register (2011). *Federal Register* Vol. 76, No. 42: 11667-11936.

[48] Palagashvili, Liya and Patrick O'Connor (2021). "Unintended Consequences of Restrictions on H-1B Visas." Policy Brief, Mercatus Center, George Mason University. January 2021.

[49] Park, Haeyoun (2015). "How Outsourcing Companies Are Gaming the Visa System." *New York Times*, November 10, 2015.

[50] Peri Giovanni, Kevin Shih and Chad Sparber (2015a). "STEM Workers, H-1B permits and Productivity in US Cities." *Journal of Labor Economics*, Vol. 33, no. S1 Part 2 (2015): S225-S255.

[51] Peri Giovanni, Kevin Shih and Chad Sparber (2015b). "Foreign and Native Skilled Workers: What Can We Learn from H-1B Lotteries?" NBER Working Paper 21175.

[52] Reif, Julian (2010). "STRGROUP: Stata Module to Match Strings Based on Their Levenshtein Edit Distance." Statistical Software Components S457151, Boston College Department of Economics, revised 27 Sep 2015.

[53] Smith, Vernon L. (1968). "Economics of Production from Natural Resources." *The American Economic Review*, Vol. 58(3): 409-431.

[54] Sparber, Chad (2018). "Choosing Skilled Foreign-Born Workers: Evaluating Alternative Methods for Allocating H-1B Work Permits." *Industrial Relations: A Journal of Economy and Society*, Vol. 57 (1): 3-34.

[55] Sparber, Chad (2019). "Substitution between Groups of Highly-Educated, Foreign-Born, H-1B Workers." *Labour Economics*, Vol. 61.

[56] Torres, Nicole. (2017). "The H-1B Visa Debate, Explained." *Harvard Business Review*, 4 May 2017.

[57] Turner, Patrick, (2020). "High-Skilled Immigration and the Labor Market: Evidence from the H-1B Visa Program." Mimeo. University of Notre Dame.

[58] USCIS (2007). "USCIS Runs Random Selection Process for H-1Bs." USCIS Update issued April 13, 2007.

[59] USCIS (2018a). "Contracts and Itineraries Requirements for H-1B Petitions Involving Third-Party Worksites." Policy Memorandum PM-602-0157. Issued February 22, 2018.

[60] USCIS (2018b). "USCIS Completes the H-1B Cap Random Selection Process for FY 2019." Retrieved at https://www.uscis.gov/news/alerts/uscis-completes-h-1b-cap-random-selection-process-fy-2019

[61] USCIS (2019). "USCIS Completes the H-1B Cap Random Selection Process for FY 2020 and Reaches the Advanced Degree Exemption Cap." Retrieved at https://www.uscis.gov/news/alerts/uscis-completes-h-1b-cap-random-selection-process-fy-2020-and-reaches-advanced-degree-exemption-cap

[62] USCIS (2020a). "Characteristics of H-1B Specialty Occupation Workers; Fiscal Year 2019 Annual Report to Congress." Retrieved at https://www.uscis.gov/sites/default/files/reports-studies/Characteristics_of_Specialty_Occupation_Workers_H-1B_Fiscal_Year_2019.pdf

[63] USCIS (2020b). "Rescission of Policy Memoranda." Policy Memorandum PM-602-0114. Issued June 17, 2020.

[64] USCIS (2020c). "FY 2021 H-1B Cap Petitions May be Filed as of April 1." Release Date: April 1, 2020. Retrieved at https://www.uscis.gov/news/news-releases/fy-2021-h-1b-cap-petitions-may-be-filed-april-1.

[65] USCIS (2021). "DHS Delays Effective Date of H-1B Selection Final Rule." Release Date: February 4, 2021. Retrieved at https://www.uscis.gov/news/alerts/dhs-delays-effective-date-of-h-1b-selection-final-rule.

[66] USCIS (2022). "H-1B Electronic Registration Process." Release Date: April 25, 2022. Retrieved at https://www.uscis.gov/working-in-the-united-states/temporary-workers/h-1b-specialty-occupations-and-fashion-models/h-1b-electronic-registration-process.

[67] USCIS (2023). "H-1B Electronic Registration Process." Release Date: April 24, 2023. Retrieved at https://www.uscis.gov/working-in-the-united-states/temporary-workers/h-1b-specialty-occupations-and-fashion-models/h-1b-electronic-registration-process.

A H-1B Employment Concentration Data

A.1 Firm Names

Raw summary statistics on the uniqueness of firm names may be misleading in the sense that a single firm in the dataset can be represented by different name variants due to typos, misspellings, or alternative naming conventions. It is possible that firms took greater care to ensure consistent name spellings after the quota became binding, thus leading to a spurious change in the measured concentration of new H-1B hires across time. Figure 2 limits this possibility by employing an automated routine to correct misspellings based upon the Levenshtein edit distance between firm names (Reif (2010)). This method compares the number of characters in a firm name that would need to change for two firms to share a name. We choose a relatively high threshold value of 25% for this procedure: If 25% or fewer characters would have to change, we assume that the two names belong to the same firm, and we group workers accordingly.[25]

A.2 LCA Data

Although I-129 petitions represent the best available information on firms' H-1B hiring behaviors, Figure A1 complements our Figure 2 illustration of H-1B employment concentration using LCA data. These filings provide a crude measure of a firm's intent to hire foreign labor, but the number of job openings with LCA approval will far exceed the number of job offers that firms eventually extend to foreign workers.[26] Specified work start dates are merely projections at best and might not correspond to actual start dates at all.[27] LCAs provide no indication of whether the employer is cap exempt. Moreover, LCAs are not linked to employees, so they will not indicate whether the worker will be seeking new (cap-subject) or continued (cap-exempt) H-1B status.

Despite these limitations, the documented trends for LCA filings in Figure A1 (after employing the automated name correction routine and cap-exempt employer proxy described above) mirror those of H-1B petitions in Figure 2. We graph trends across calendar years of stated work start dates. Among cap-subject firms, the top panel of Figure A1 shows a sustained decline in the proportion of LCAs filed by firms submitting five or fewer total

[25]We also employ a 10% threshold and no correction at all. These differences do not affect our qualitative conclusions.

[26]Sparber (2019) describes this in more detail.

[27]Firms need to have approved LCAs in March if they intend to file an I-129 petition in April, but they can only submit an LCA six months in advance of the stated work start date. Peri et al. (2015b) document pre-dating behavior in which firms file LCAs with stated start dates in September for work that is more accurately scheduled to begin in October.

LCAs per year, falling from over 17% in 2002 and 2003, and stabilizing at about 5.5 to 7.5% from 2012 through 2019. The figure hovers around 6.5 to 8.5% for cap-exempt employers throughout the period. The bottom panel shows that since 2012, roughly 75% or more of LCAs filed by cap-subject firms have originated with firms seeking 250 or more H-1B workers. This compares to a figure below 50% in years when statutory H-1B limits were high. The share among cap-exempt employers declines over this period, however.

B Additional Theory Details

B.1 Partial Equilibrium Setup

This section describes underlying assumptions on preferences that rationalize the partial equilibrium model in the main text. We assume that the representative consumer has quasi-linear utility over the aggregate H-1B sector good (X) and a numeraire good representing all other goods (m), with constant utility parameter γ:

$$U = m + \gamma \times log\left(X\right) \tag{42}$$

The H-1B sector good is produced according to:

$$X = \left[\sum_i \left(\alpha_i\right)^{\frac{1}{\sigma}} \left(x_i\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{43}$$

The household is subject to a budget constraint:

$$m + PX = I + \Pi + w^N \bar{N} + w^F \Omega \tag{44}$$

Total income excluding the H-1B sector (I) is exogenous given the partial equilibrium setting; Π represents the total profits of the H-1B sector; and P is the price index for the aggregate H-1B sector output:

$$P = \left[\sum_i \alpha_i \left(p_i\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{45}$$

Note that while we include both native and immigrant H-1B wages in (44) – taking the representative household to include immigrant workers – the analysis would be basically identical if we only focused on native households. Maximizing (42) gives us:

$$PX = \gamma \tag{46}$$

This rationalizes our constant E assumption since $PX = E$. The constant E here depends on the assumption of log substitutability. This constant E setup is helpful for expositional clarity and simplicity, but could easily be relaxed without changing the thrust of our overall analysis. Finally, we could find the optimal choice of each variety x_i within (43), and this would give us the standard CES demand in (1).

B.2 Generalizing Marginal Costs for Outsourcing Firms

We can consider a generalization of the model where instead of a constant marginal search cost, outsourcing firms have an increasing marginal cost function with parameter value λ_O. Our model in the text corresponds to the limiting case of this general model when $\lambda_O \to 1$. In this version of the model, (17) for outsourcing firms is instead:

$$F_{jA}^{O} = \left[\frac{\left(w^O - w^F\right)\rho - c}{c_X} \right]^{\frac{1}{\lambda_O - 1}}$$

As $\lambda_O \to 1$, F_{jA}^{O} can be finite (and non-zero) only if: $\frac{\left(w^O - w^F\right)\rho - c}{c_X} = 1$, i.e. $\left(w^O - w^F\right)\rho - c - c_X = 0$. This is the same as (19) in the baseline model.

In this generalized version of the model, the equilibrium number of filings by a regular firm – defined by (23) in the limiting case – becomes:

$$F_{iA} = \left[\frac{\left(F_{jA}^{O}\right)^{\lambda_O - 1} c_X}{c_i} \right]^{\frac{1}{\lambda - 1}}$$

In contrast to what we find in the baseline model, F_{iA} now depends on the value of F_{jA}^{O} and, therefore through (22), also on the value of Ω. As λ_O gets closer to 1, F_{iA} becomes less affected by F_{jA}^{O} and Ω. The limiting case again corresponds exactly to the baseline model in this respect, where a regular firm's filings (F_{iA}) are independent of the quota (Ω).

C Outsourcing Firms and Behavior

This appendix provides more information about our model's assumptions and predictions regarding outsourcing firms and behavior.

C.1 Marginal Search Costs of Outsourcing Firms

Section 4.1 argued that an increase in the marginal search cost of outsourcing firms will reduce the equilibrium number of bids submitted by those firms, raise regular firms' bids, and lower the total search costs of H-1B hiring. This section provides proofs of those arguments.

C.1.1 Number of Bids

Outsourcing firms' bids unambiguously decrease in their marginal search costs. This can be found by differentiating their equilibrium number of filings in (25) with respect to c_X. In doing so, note that Equation (24) implies that w^O is independent of c_X. This is because w^O is determined by the demand and supply for H-1B workers in the second stage, neither of which depend on c_X. The perfectly elastic supply of foreign-workers implies that w^F is also constant. The derivative therefore equals:

$$\frac{d\sum_j F_{jA}^O}{dc_X} = \frac{-\Omega\left(w^O - w^F\right)}{\left(c + c_X\right)^2} - \sum_i \frac{dF_{iA}}{dc_X} < 0$$

The sign of this response becomes clearer when recognizing that the equilibrium condition in (23) implies that each regular firm's filing behavior is positively related to outsourcing specialists' search costs:

$$\frac{dF_{iA}}{dc_X} = \left(\frac{1}{\lambda - 1}\right)\left(\frac{c_x^{2-\lambda}}{c_i}\right)^{\frac{1}{\lambda-1}} > 0$$

Together, these responses demonstrate that an increase in outsourcing firms' marginal search costs will shift H-1B hiring activity away from outsourcing specialists and toward regular firms. If we further define $\theta = \sum_j F_{jA}^O / (\sum_i F_{iA} + \sum_j F_{jA}^O)$ as the share of petitions submitted by outsourcing firms, we know that this too must decline in c_X.

C.1.2 Total Search Costs

We now provide a proof of Proposition 1. To understand the effect of outsourcing firms' marginal search costs on aggregate costs, first substitute the equilibrium total number of outsourcers' bids (25) into the total search cost function (26). Regular firms' total search costs equal $\sum_i \left(cF_{iA} + c_i \frac{1}{\lambda} F_{iA}^\lambda\right)$; outsourcing firms' costs are represented by $\Omega\left(w^O - w^F\right) - (c + c_X)\sum_i F_{iA}$:

$$C = \sum_i \left(cF_{iA} + c_i \frac{1}{\lambda} F_{iA}^\lambda\right) + \Omega\left(w^O - w^F\right) - (c + c_X)\sum_i F_{iA}$$

Total costs then simplify to:

$$C = \sum_i c_i \frac{1}{\lambda} F_{iA}^\lambda + \Omega\left(w^O - w^F\right) - c_X \sum_i F_{iA}$$

The total costs incurred by outsourcing firms under a quota and lottery allocation unambiguously decrease in c_X as high marginal costs greatly reduce those firms' submissions.[28] There is, however, a potential counteracting effect noted above: The shift away from hiring by outsourcing firms toward regular firms could lead to an increase in search costs. The net effect can be calculated by differentiating total costs with respect to c_X:

$$\frac{dC}{dc_X} = \sum_i c_i F_{iA}^{\lambda-1} \frac{dF_{iA}}{dc_X} - c_X \sum_i \frac{dF_{iA}}{dc_X} - \sum_i F_{iA}$$

Substituting the equilibrium number of bids filed by a regular firm (23) delivers:

$$\frac{dC}{dc_X} = \sum_i \left(c_i \left(\frac{c_X}{c_i} \right)^{\frac{\lambda-1}{\lambda-1}} \frac{dF_{iA}}{dc_X} \right) - c_X \sum_i \frac{dF_{iA}}{dc_X} - \sum_i F_{iA}$$

This further simplifies to Equation (27) from Section 4.1.

$$\frac{dC}{dc_X} = (c_X - c_X) \sum_i \frac{dF_{iA}}{dc_X} - \sum_i F_{iA} = -\sum_i F_{iA} < 0$$

Hence, this derivation demonstrates that the response of regular firms is insufficient to counteract the decline in outsourcing firms' bidding costs. Altogether, if the government acts to increase the marginal search costs for outsourcing firms, it would simultaneously lower economy-wide search costs while reducing the share of H-1B workers employed by (and petitions filed by) outsourcing specialists. In other words, outsourcing firms exacerbate search costs created by the lottery.

C.2 Shut Down Outsourcing Firms

C.2.1 Search Costs at Shutdown

Appendix C.1 proves that outsourcing specialists worsen the search cost externality. This section examines a related question about whether the externality would exist at all in the absence of outsourcing firms. This is a continuation of the exercise above. If the government escalates restrictions on outsourcing firms, then c_X will eventually grow so large that all such firms will cease operation. A social planner operating under a quota would never file bids in excess of the number of available positions. Therefore, any value of $\rho < 1$ at this shutdown value of c_X would be sufficient for demonstrating that the lottery creates an inefficient search cost externality even if outsourcing firms are prohibited.

[28]Also note that unlike in many other economic contexts, there is no welfare loss from the quantity reduction here because real economic activity is in any case fixed by the quota.

Unfortunately, two limitations make it impossible to solve definitively for this shutdown value of c_X. First, its calculation would require knowledge about the full microstructure of the search process for regular firms, including the value of λ and the distribution of search costs across regular firms (c_i). Second, it is not possible to obtain a closed-form solution for this value. One can see both of these limitations by substituting (23) into (25), setting the expression equal to zero, and isolating the c_X terms to one side:

$$c_X^{\frac{1}{\lambda-1}}\left(c+c_X\right)=\frac{\Omega\left(w^O-w^F\right)}{\sum_i c_i^{\frac{1}{1-\lambda}}}$$

Given these limitations, we turn to graphical analysis to illustrate that – given the lottery setup – ρ is regularly less than one and the externality exists in the absence of outsourcing firms. We stress that this analysis is best understood as providing an informal numerical example, not a more formal calibration evaluating the counterfactual equilibrium that would occur without any outsourcing firms. As noted, that task would require more information about the microstructure of the search process. Though we can appeal to past evidence for some parameter values, we lack information for others.

Let $\Omega = 85,000$ to match the current quota. The main text argues for an outsourcing premium of $\left(w^O - w^F\right) = \$23,500$ and a common submission cost of $c = \$2,000$. For convenience and with less support from the data or literature, we assume $\lambda = 2$ and $c_i = \$2,700$ for each of 40,000 firms.

Figure C1 performs the first exercise. The left panel displays the share of petitions filed by outsourcing firms (θ) as a function of those firms' marginal search costs. This can be calculated by substituting the assumed parameter values and resulting c_X calculation into equilibrium conditions (23) and (25). The range is specified to display θ values from 0 to 50%. As c_X rises, outsourcing firms shut down.

The middle panel computes values of the win probability implied by dividing the quota by the total number of filings. Recall that as long as the probability of winning a bid is less than one, there is more search taking place than what a social planner would choose. Here we see that $\rho = 0.54$ at the shut-down cost, indicating that the externality persists in the absence of outsourcing firms. The government would not eliminate the H-1B search externality by prohibiting outsourcing.

The final panel displays total search costs, C, normalizing the initial value to 100 when H-1B employment is evenly split between regular and outsourcing firms. Consistent with our analytical results in the main text, increasing marginal search costs for outsourcing firms cause total search costs to fall. Given our parameters, the search costs drop by almost 30%

when outsourcing firms are shut down relative to when they account for half of bids. Note that since the probability remains less than one at shut down, the lottery inefficiency is still present; a social planner would optimally search for exactly the number of workers made available by the quota.

In the Figure C1 example, ρ remains less than one when outsourcing firms shut down. It is natural to ask whether this will generally be the case. As long as the search costs for regular firms is not too large, this will indeed be the case. If the search costs for regular firms are large enough, the quota will not bind and $\rho = 1$. We see this from Figure C2, which shows that our assumed values of c, c_i, and λ were not necessary for $\rho < 1$ at shut down. The graphs in this figure display the win probability (ρ) occurring at the shutdown value of c_X as a function of the application cost component common to outsourcing and regular firms (c). The left panel continues to assume that $\lambda = 2$. It displays three curves distinguished by different assumed values for a single marginal search cost $\left(c_i F_{iA}^{\lambda-1}\right)$ shared by each of 40,000 regular firms: $c_i = \{2000, 4000, 8000\}$. The right panel displays the same, but instead assumes $\lambda = 1.1$ – a value much closer to that experienced by outsourcing firms.

In all cases, we see that increasing filing costs (c) and marginal search costs for regular firms (c_i) are both associated with increasing win probabilities for regular firms at the shut-down point of c_X for outsourcing specialists. At particularly high search costs, regular firms mimic the social planner's solution, file petitions equal to the number of available positions, and win the lottery with certainty. Altogether, the search externality is a consequence of lottery allocation. Outsourcing firms can worsen the lottery's externality, but they are not required for the externality to exist.

C.2.2 Search Externality without Outsourcing Firms

Equation (31) shows the effect of a quota reduction on total search costs in the presence of outsourcing firms. The same equation still holds even in the absence of outsourcing firms. Outsourcing firms therefore exacerbate externalities by amplifying the win probability elasticity with respect to the quota (ψ), not by changing the externality function. The derivation of this condition, however, is much different without those firms.

To see this, begin with the economy-wide total cost of bidding function in (26) but with $F_{jA}^{O} = 0$. Then differentiate with respect to the quota:

$$\frac{dC}{d\Omega} = \sum_i \left[c + c_i F_{iA}^{\lambda-1}\right] \frac{dF_{iA}}{d\Omega}$$

In the absence of outsourcing firms, equilibrium conditions (19) and (23) cease to hold. Regular firms' petitions are governed by (17), and F_{iA} is no longer independent of the quota

(Ω). Substituting (17) into the above, we find:

$$\frac{dC}{d\Omega} = \left(w^O - w^F\right)\rho \sum_i \frac{dF_{iA}}{d\Omega}$$

Since the number of successful bids must equal the quota, (22) now becomes $\sum F_{iA} = \Omega/\rho$. Then converting to a semi-elasticity, substituting for $\frac{dF_{iA}}{d\Omega}$, and simplifying, we get:

$$-\frac{dC}{d\ln(\Omega)} = \left(w^O - w^F\right)\Omega\left(\frac{d\rho}{d\Omega}\frac{\Omega}{\rho} - 1\right)$$

Though the derivation differs, this expression is equivalent to the quota-induced search cost externality in (31). Lottery allocation ensures this expression is positive so long as firms respond to tightening quotas by submitting more petitions (i.e., if $\psi > 1$). As shown above, regular firms continue to submit bids in excess of available positions even in the absence of outsourcing specialists.

C.3 Shut Down Outsourcing

This section considers how the model is altered if the government outlaws the practice of outsourcing altogether so that firms can only use the workers they receive through the H-1B lottery. Denote the firm's unconstrained optimum (i.e. if it had enough H-1B workers at w^O) as F_i^*. The number of workers used by a firm will then be $F_i = \min\{F_i^*, F_{iS}\}$; the firm will choose the unconstrained optimum if it can and otherwise will use the maximum number of available workers. Firms that have too many successful H-1B bids will simply not hire the excess workers. We can write the expected profits as:

$$E[\pi_i] = \int_0^{F_i^*} \Pi_{iL}(F_{iS})\,dG(F_{iS};F_{iA}) + (1 - G(F_i^*;F_{iA}))\,\Pi_{iH} - cF_{iA} - c_i\frac{1}{\lambda}F_{iA}^{\lambda}$$

Before accounting for search costs, the optimal profit when the firm is constrained equals Π_{iL}, and Π_{iH} is the unconstrained optimal profit. This latter outcome's profit does not depend on F_{iS}. The endogeneous equilibrium variables now are just w_N and ρ, which are determined by equations analogous to (22) and (20).

$$\rho = \frac{\sum_i F_{iA}}{\Omega}$$

$$\left(\frac{w^N}{w^O}\right)^{\epsilon} = \frac{\sum_i \left\{\int_0^{F_i^*} \Pi_{iL}(F_{iS})\,dG(F_{iS};F_{iA}) + (1 - G(F^*;F_{iA}))\,F_i^*\right\}}{N}$$

Note that the first term in the numerator of the ρ calculation is the number of workers actually used, which is no longer Ω and no longer includes outsourcing specialists.

While the firms' problem here is still differentiable, we cannot analytically find the optimal F_{iA} under these assumptions. Owing to the two possible cases (i.e. unconstrained and constrained), the optimum will now depend more specifically on derivatives of the probability distribution and density functions. These functions are complex for the continuous binomial distribution. This also means that we are not able to analytically characterize the resulting equilibrium and compare it to the corresponding equilibrium without outsourcing firms.

In general, the prohibition of outsourcing behavior would have an ambiguous effect on F_{iA} and ρ due to conflicting mechanisms discussed in Section 4.1. Firms would not gain from winning too many H-1B workers, a factor that reduces the marginal benefit of bidding. But firms with too few workers would have no access to a secondary H-1B market, a factor that increases the marginal benefit of bidding. We can verify this ambiguity with numerical exercises, available upon request.

D Calculating the Win Probability Elasticity

This appendix provides more complete detail on computing key equations from Section 5. The first three subsections provide non-exhaustive methods for calculating changes in pre- versus post-lottery filings behavior and the elasticity of the win probability with respect to the quota (ψ) in Equation (32). The next illustrates how ψ relates to wages as in Equation (39). The final subsection derives the externality arising from the quota in Equation (40).

D.1 Win Elasticity: Method 1

The first method of calculating ψ begins by assuming that the number of filings, real H-1B demand, and the eventually-binding quota were equivalent in pre-lottery years so that $\bar{F}_0 = D_0 = \Omega_1 = \rho_1 \bar{F}_1$. Data helps to justify this assumption. Figure 1 indicates that firms submitted approximately 85,000 petitions both when the quota was non-binding (2002 and 2003), and when it was constraining but most H-1Bs were allocated on a first-come first-served basis (2004-2007 and 2010-2013).

Next, we calculate filings during the current lottery era $\left(\bar{F}_1\right)$ after removing the increase in actual demand ($\Delta D = D_1 - \Omega_1 = D_1 - \rho_1 \bar{F}_1$) that may have occurred over time. Since excess filings can be expressed as $\bar{F}_1 - D_1 = \phi \bar{F}_1$, we can record lottery-era filings net of demand as $\bar{F}_1 - \Delta D = \left(\phi + \rho_1\right) \bar{F}_1$. The difference in relevant log-filings between eras is then assumed to equal:

$$d\ln\left(\bar{F}\right) = \ln\left(\bar{F}_1 - \Delta D\right) - \ln\left(\bar{F}_0\right)$$

This then simplifies to Equation (33) from the main text:

$$d\ln\left(\bar{F}\right) = \ln\left(\phi + \rho_1\right) - \ln\left(\rho_1\right)$$

Since pre-lottery filings are assumed to equal the lottery era quota, this definition of $d\ln\left(\bar{F}\right)$ represents the growth in total filings (net of demand trends) relative to what the social planner would choose if confronted with a cap – a submission total equalling the number of available H-1Bs so that $\rho = 1$. That is, it compares excess filings above the quota to a world in which firms never file any H-1B petitions above the cap. Note that filings and demand equal the binding quota under the tax, minimum wage, and willingness to pay mechanisms illustrated in Figure 3. Thus, these alternative policies can serve as a baseline reference for how the imposition of a quota under a lottery system affects filing behavior. In the absence of a lottery, firms exhibit no strategic behavior after the quota is reached, and they make no attempt to submit additional H-1B petitions in order to secure a better chance of hiring the number of workers they actually desire. Firms do not even bother demanding any real workers beyond the quota since they know they cannot be hired. After substituting this value of $d\ln\left(\bar{F}\right)$ into (32), we get the formula for ψ in (34).

D.2 Win Elasticity: Method 2

The second approach compares lottery-era H-1B petitions to actual (but latent) demand:

$$d\ln\left(\bar{F}\right) = \ln\left(\bar{F}_1\right) - \ln\left(D_1\right)$$

We again take advantage of the excess filings identity and rearrange so that $\left(\frac{\bar{F}_1}{D_1}\right) = \frac{1}{1-\phi}$. Filings growth then simplifies to Equation (35) from the main text.

$$d\ln\left(\bar{F}\right) = -\ln\left(1 - \phi\right)$$

This setup yields multiple interpretations of what filings growth represents. First, (35) reflects a scenario in which real H-1B demand is constant ($D_1 = D_0$). Any lottery-era filings $\left(\bar{F}_1\right)$ exceeding this amount therefore reflect a response to supply restrictions, which drive all H-1B petition growth. Second, it could be that firms are able to hire all of the H-1B workers they desire ($\rho = 1$) so long as the quota is large enough to satisfy real demand: $\Omega \geq D_1$. Filings explode once the quota falls below that threshold, however. A comparison of lottery-era filings to lottery-era demand would then represent an attempt to compute ψ only from the portion of Equation (32) that is not perfectly inelastic. A third and related interpretation supposes that the government set the quota at lottery-era demand so that $\Omega = D_1$, but it chose a non-random allocation mechanism (such as willingness to pay). The gap between lottery-era H-1B petitions and actual demand would then reflect filing growth attributable to the lottery alone.

D.3 Win Elasticity: Method 3

The third approach begins by recognizing that our relevant growth calculation, $d\ln\left(\bar{F}\right)$, represents a discrete change in filings ($\bar{F}$) at a constant level of demand ($\bar{D}$). In elasticity form, this means that we must calculate the difference in both H-1B filings ($\Delta\bar{F}$) and demand (ΔD) between the two periods, normalized by total filings.

$$d\ln\left(\bar{F}\right) = \frac{d\bar{F}\big|_{D=\bar{D}}}{\bar{F}} = \frac{\Delta\bar{F} - \Delta D}{\bar{F}}$$

We again note that in the initial period with a non-binding quota, filings and demand are equivalent ($\bar{F}_0 = D_0$). The numerator then simplifies to $\bar{F}_1 - D_1$. Taking advantage of of the excess filings identity, this setup then delivers Equation (37) from the text:

$$d\ln\left(\bar{F}\right) = \phi$$

D.4 Win Elasticities from Wage Data

To estimate ψ from wage data, first rewrite it in terms of income shares and the elasticity of substitution between native and foreign labor. Using (28), we get:

$$\psi = \frac{d\rho}{d\Omega}\frac{\Omega}{\rho} = \frac{w^O\Omega}{w^O\Omega + w^N\bar{N}}\frac{w^O/w^F}{w^O/w^F - 1}\left[1 + \frac{1}{\epsilon}\left(w^O/w^N\right)^{\epsilon-1}\right] \tag{47}$$

Next, rearrange the marginal rate of technical substitution equilibrium from (20) and insert the relative wage into the bracketed expression. If we then assume that native and foreign labor is relatively substitutable so that ϵ is large,[29] Equation (47) simplifies to Equation (39) from the text.

[29] Dustmann and Preston (2012), Manacorda, Manning, and Wadsworth (2012), Ottaviano and Peri (2012), and Borjas, Grogger, and Hanson (2012) provide extensive discussion about this parameter and challenges in measuring it. The lowest estimates – and hence greatest complementarity – in the literature that we are aware of come from Manacorda, Manning, and Wadsworth (2012). They find a minimum value of $\epsilon = 4.6$ and a baseline value of $\epsilon = 7.8$. Ottaviano and Peri (2012) prefer a value of $\epsilon = 20$, whereas some Borjas, Grogger, and Hanson (2012) values suggest perfect substitutability ($\epsilon = \infty$).

D.5 Search Cost Externalities

Equation (31) describes how the H-1B quota, in the presence of a lottery, generates lost quota rents in the form of a search cost externality that would not arise under alternative allocation methods. This expression can be rearranged to reflect cost implications as a function of the win elasticity and quota change:

$$dC = -\left(w^O - w^F\right)\ \Omega\ \left(\psi - 1\right)\ d\ln\left(\Omega\right)$$

Recalling that the elasticity is simply $\psi = \frac{d\ln(\rho)}{d\ln(\Omega)}$, this expression can be rewritten as:

$$dC = -\left(w^O - w^F\right)\ \Omega\ \left(d\ln\left(\rho\right) - d\ln\left(\Omega\right)\right)$$

The final term in this equation is identical to $-d\ln\left(\bar{F}\right)$. Making this substitution, we see that externality costs can be expressed as a function of excess filings as in Equation (40). The methods for calculating $d\ln\left(\bar{F}\right)$ described above can then be used to numerically evaluate these costs.

Figure 1

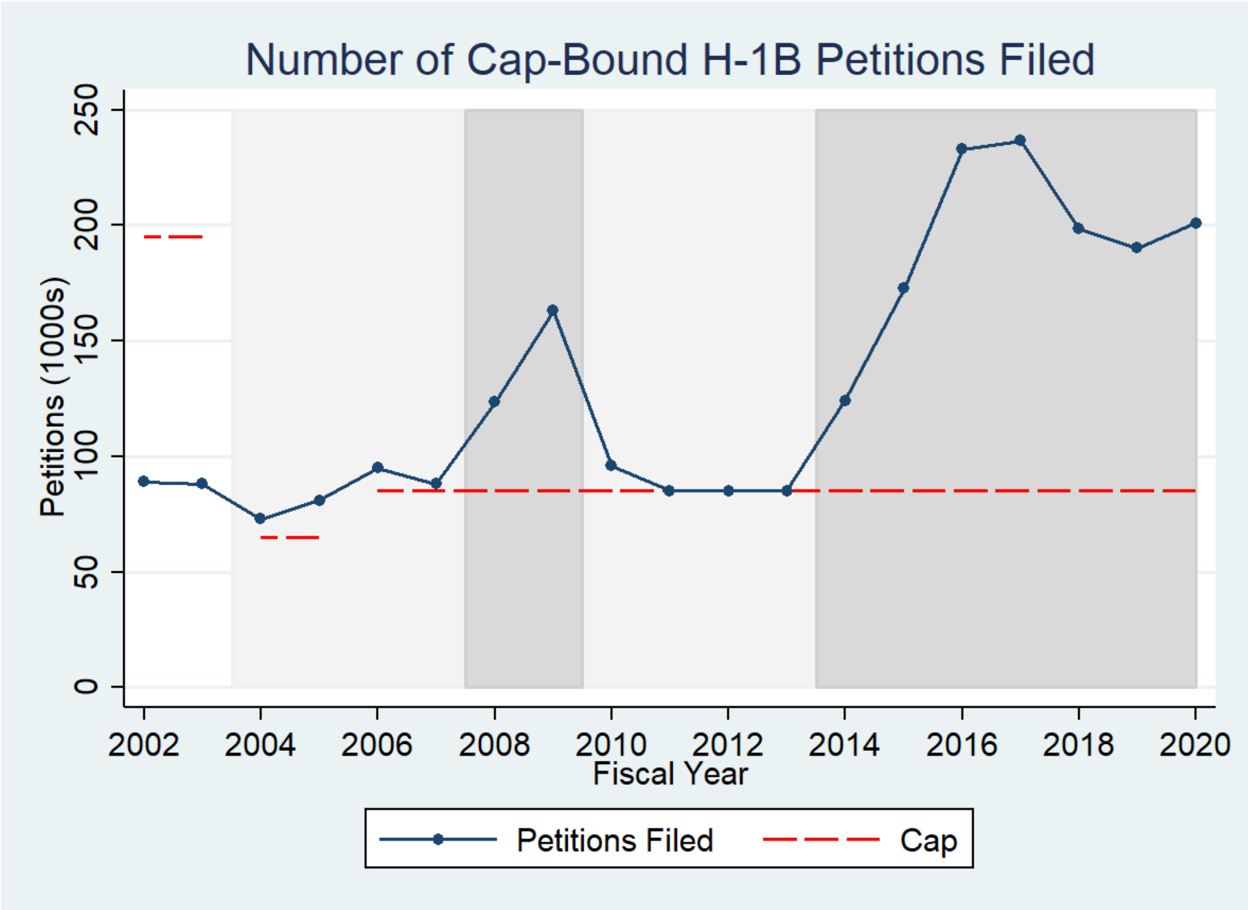

This figure displays the H-1B cap (discontinuous red dashed line) and the number of I-129 petitions filed (solid blue line) for cap-subject H-1B workers in fiscal years 2002-2020. Years with a binding H-1B cap lightly shaded. Years in which USCIS allocated all cap-bound H-1Bs by random lottery are darkly shaded. Sources include USCIS (2007, 2018b, 2019), DHS (2011, 2018), and Office of the Federal Register (2011).

Figure 2

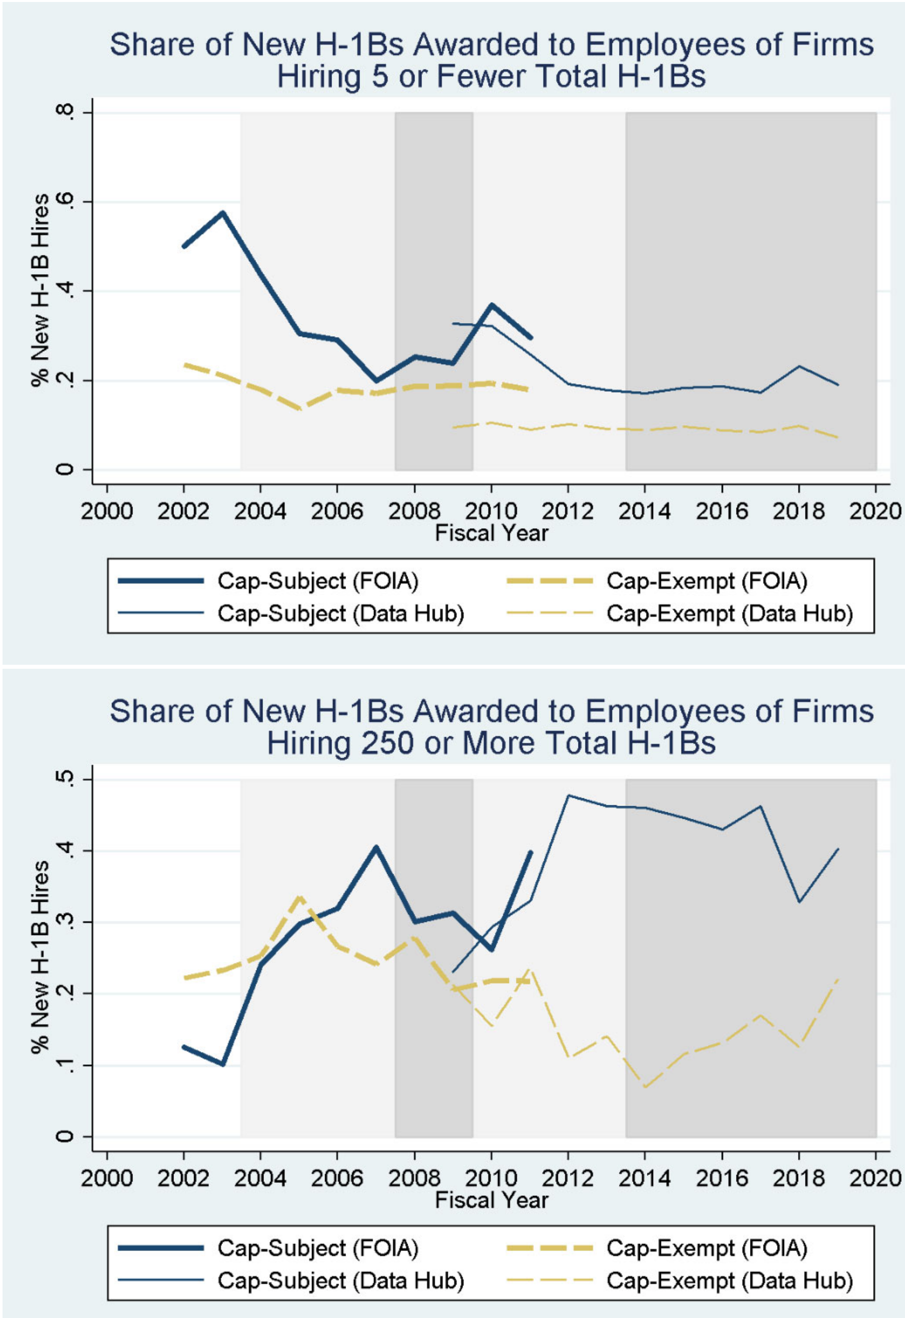

This figure displays H-1B hiring trends at cap-subject and cap-exempt employers. The top panel displays the percentage of new H-1B workers hired by firms employing five or fewer total H-1B workers for a given fiscal year; the bottom panel displays the percentage hired at firms employing 250 or more total H-1B workers. Years with a binding H-1B cap lightly shaded. Years in which USCIS allocated all cap-bound H-1Bs by random lottery are darkly shaded. Data is based upon I-129 petitions: Data used in bold curves come from a FOIA request to USCIS; data used in thinner curves come from the USCIS Data Hub.

Figure 3

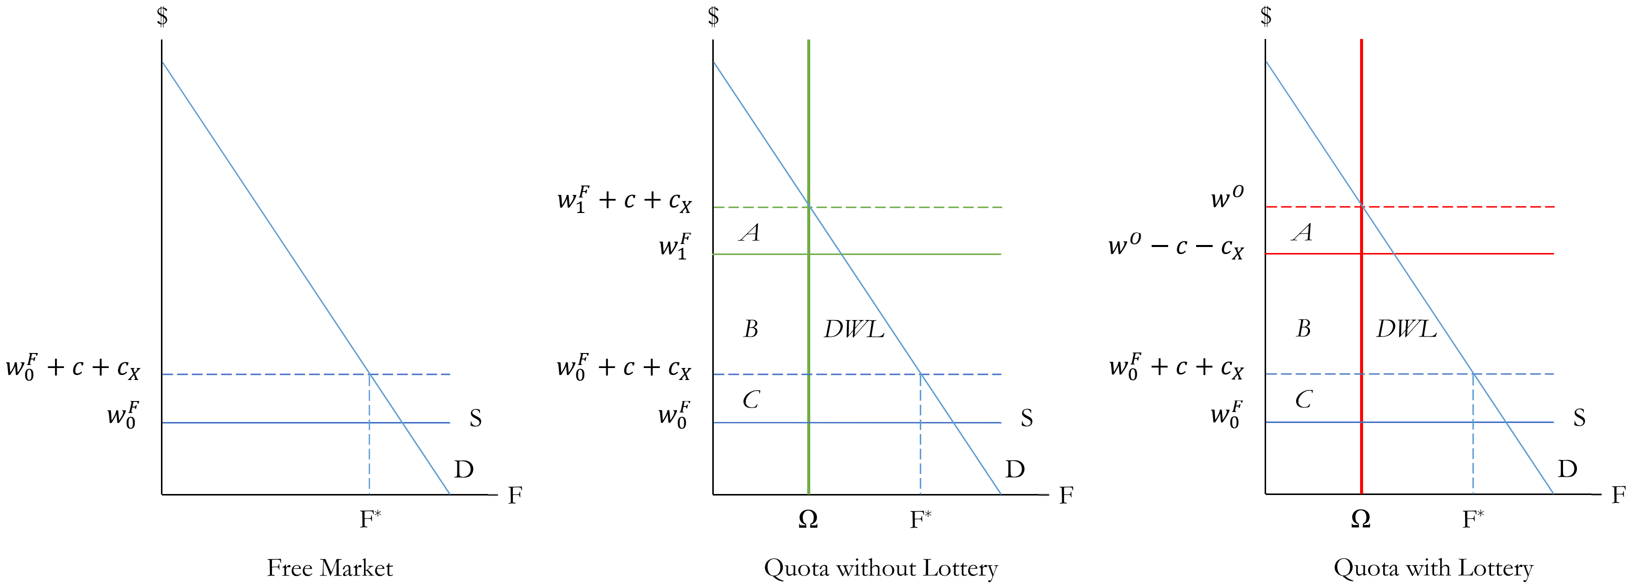

Each graph displays the Supply (*S*) and Demand (*D*) of foreign-born H-1B labor (*F*) and identifies equilibrium wages (*w*) given outsourcing firms' marginal job search and filing costs ($c + c_X$) and a possible H-1B quota (Ω). The left panel illustrates a free market equilibrium. The center displays outcomes with a quota and willingness to pay H-1B allocation. The right shows outcomes with a quota and lottery allocation. The latter two graphs yield predictable deadweight loss triangles (*DWL*). Without a lottery, the quota transfers surplus (*A* and *B*) to H-1B workers. With a lottery, that surplus is completely absorbed by additional search costs. See the text for more details.

Figure 4

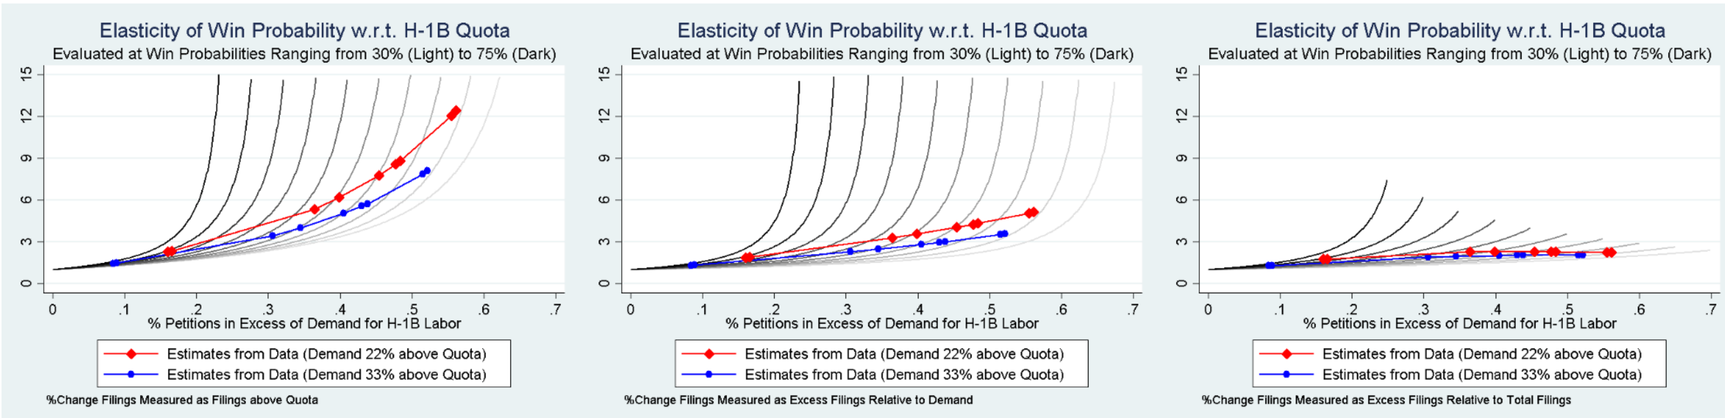

This figure displays estimates for the elasticity of the H-1B win probability with respect to the quota (ψ) as a function of the win probability (ρ) and petitions filed in excess of demand for H-1B labor (ϕ). Plausible values of ϕ range from zero to $1 - \rho$. Panels differ in the methods used to compute the percentage change in filings due to the quota.

Figure 5

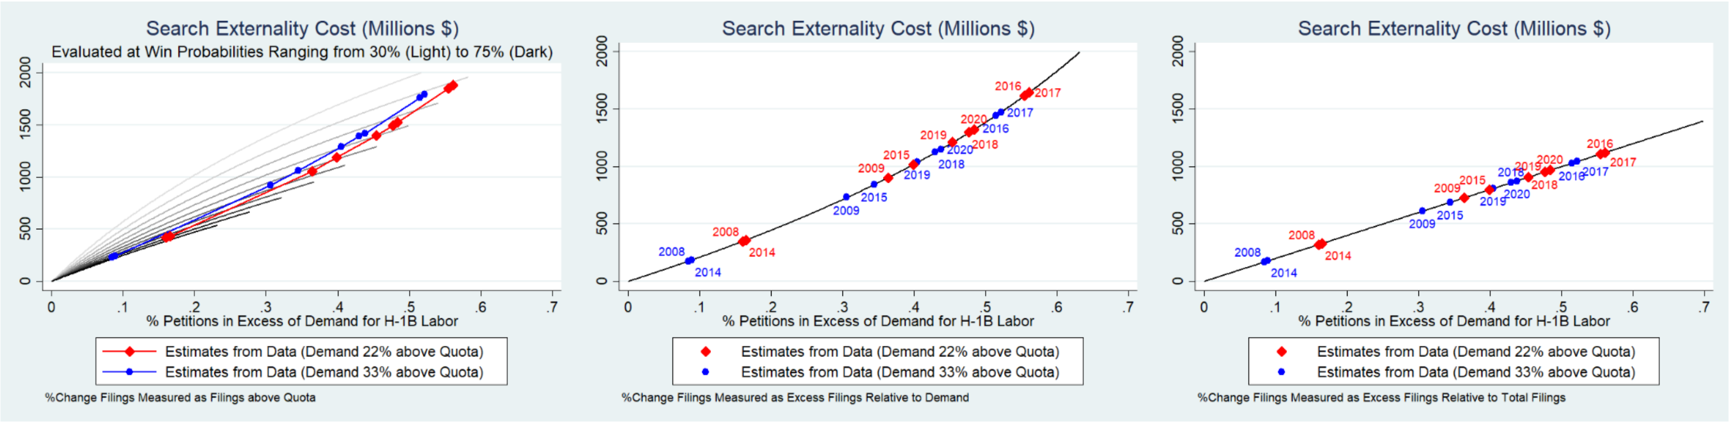

This figure displays estimates for the search cost externality generated by the H-1B quota as a function of the H-1B win probability (ρ, middle panel only) and petitions filed in excess of demand for H-1B labor (ϕ). Plausible values of ϕ range from zero to 1 – ρ. Panels differ in the methods used to compute the percentage change in filings due to the quota. Year labels suppressed in left panel to preserve clarity.

Table 1

Implied Demand Shock		33%			22%	
Assumed Real Demand		113,050			103,700	
Year of Observation	**2008**	**2009**	**2017**	**2008**	**2009**	**2017**
Applications	123,480	163,000	236,444	123,480	163,000	236,444
Excess Filings	10,430	49,950	123,394	19,780	59,300	132,744
% Excess Filings (ϕ)	0.084	0.306	0.522	0.160	0.364	0.561
Win Probability (ρ)	68.8%	52.1%	35.9%	68.8%	52.1%	35.9%
	Δln(Filings) Measured as Filings above Quota					
Δln(Filings)	0.12	0.46	0.90	0.21	0.53	0.94
Δln(Ω)	-0.26	-0.19	-0.13	-0.16	-0.12	-0.08
Elasticity (ψ)	1.45	3.45	8.10	2.27	5.34	12.42
Search Cost Externality ($M)	231	923	1791	418	1057	1879
Externality per Hire ($)	2,720	10,863	21,074	4,916	12,437	22,106
$\Delta\theta$, at θ=1/3	0.08	0.31	0.60	0.14	0.35	0.63
$\Delta\theta$, at θ=3/5	0.05	0.18	0.36	0.08	0.21	0.38
	Δln(Filings) Measured as Excess Filings Relative to Demand					
Δln(Filings)	0.09	0.37	0.74	0.17	0.45	0.82
Δln(Ω)	-0.29	-0.29	-0.29	-0.20	-0.20	-0.20
Elasticity (ψ)	1.31	2.28	3.59	1.88	3.27	5.14
Search Cost Externality ($M)	176	731	1474	349	903	1646
Externality per Hire ($)	2,074	8,599	17,340	4,103	10,628	19,369
$\Delta\theta$, at θ=1/3	0.06	0.24	0.49	0.12	0.30	0.55
$\Delta\theta$, at θ=3/5	0.04	0.15	0.30	0.07	0.18	0.33
	Δln(Filings) Measured as Excess Filings Relative to Total					
Δln(Filings)	0.08	0.31	0.52	0.16	0.36	0.56
Δln(Ω)	-0.29	-0.34	-0.50	-0.21	-0.29	-0.46
Elasticity (ψ)	1.29	1.89	2.04	1.75	2.27	2.22
Search Cost Externality ($M)	169	612	1042	320	727	1121
Externality per Hire ($)	1,985	7,201	12,264	3,764	8,549	13,193
$\Delta\theta$, at θ=1/3	0.06	0.20	0.35	0.11	0.24	0.37
$\Delta\theta$, at θ=3/5	0.03	0.12	0.21	0.06	0.15	0.22

The first rows of the table reflect estimates of real H-1B demand from Mayda et al. (2018). The second set of rows display the observed number of H-1B filings according to USCIS, plus the implied win probabilities and excess filings implied by the data and estimated real demand. The remaining sets of rows show the search cost externality (measured in millions of dollars), externality per H-1B hire (in dollars), the change in H-1B concentration at outsourcing firms (θ), and component measures caused by the quota, lottery allocation, and outsourcing according to the three methodologies outlined in the main text.

Figure A1

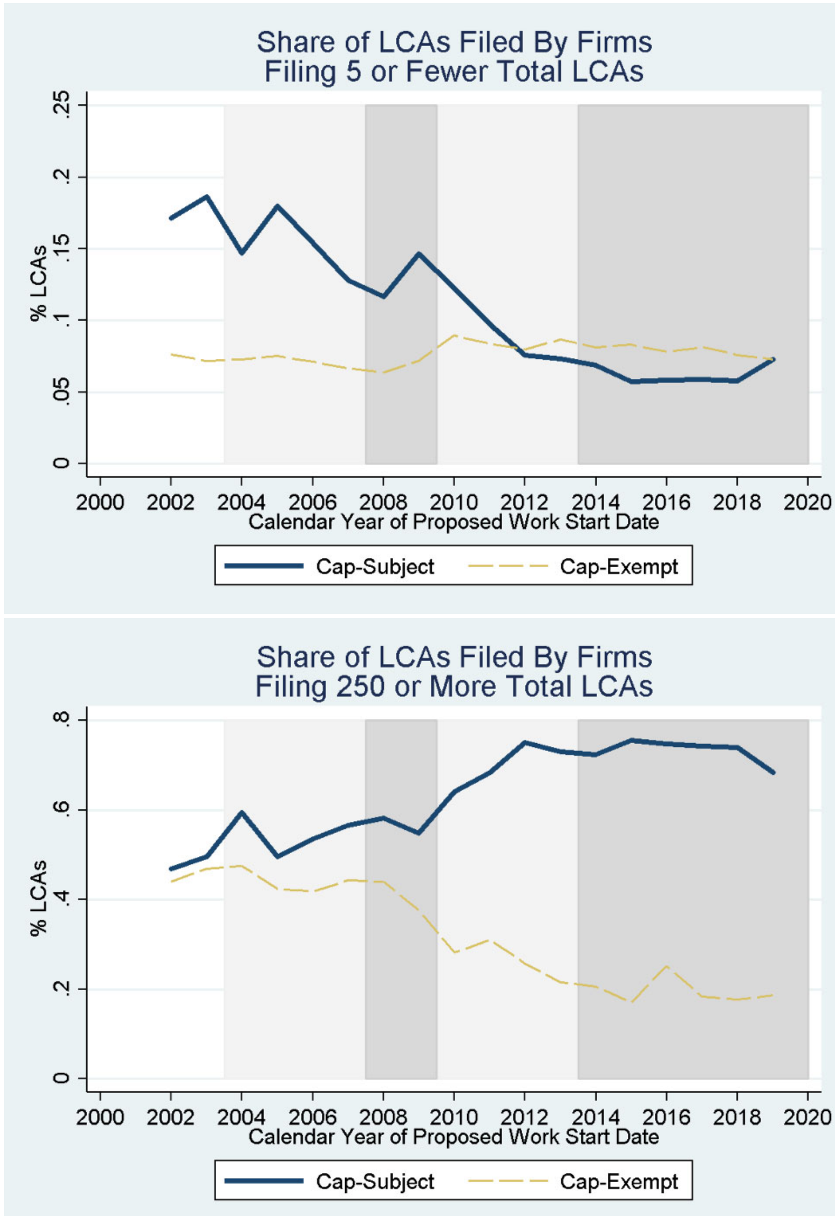

This figure displays the percentage of LCAs filed by firms filing five or fewer LCAs (top panel) and 250 or more LCAs (bottom panel) for a given fiscal year. Years with a binding H-1B cap lightly shaded. Years in which USCIS allocated all cap-bound H-1Bs by random lottery are darkly shaded. Data is based upon LCA filings.

Figure C1

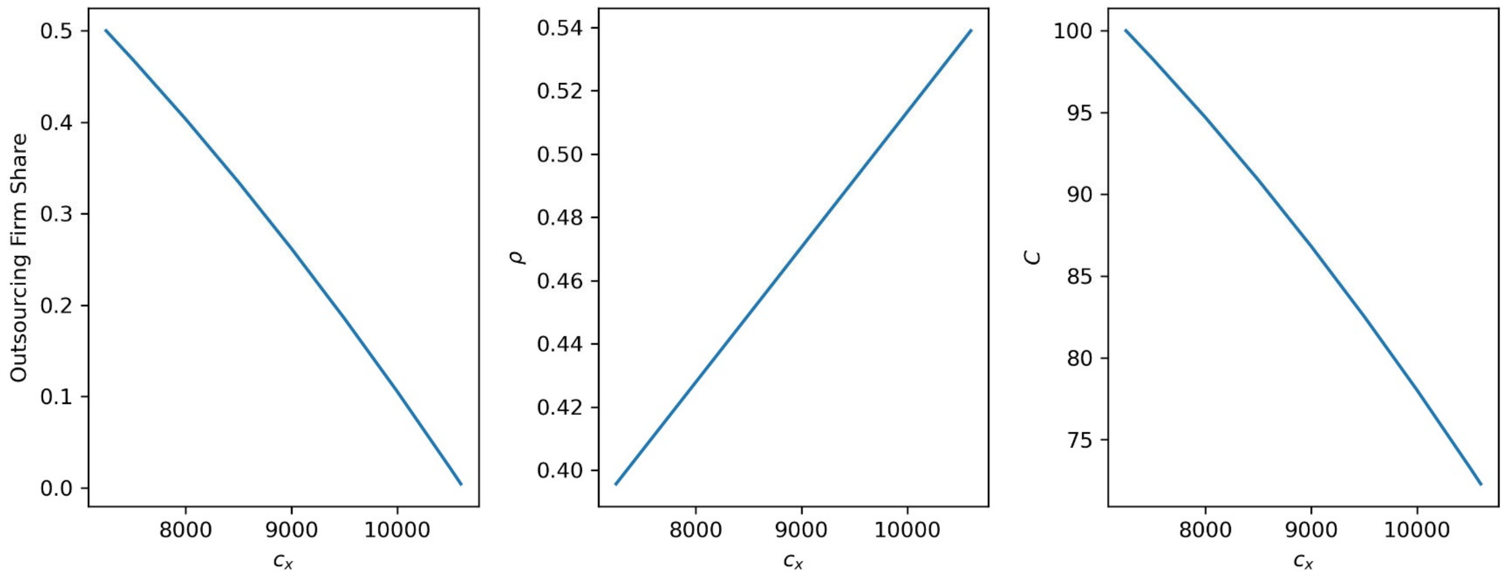

Graphs display the consequences of increasing the marginal search costs of outsourcing firms (c_X) until they shut down. Assumed parameter values are in the text. Left panel displays the share of H-1Bs submitted (and hired) by outsourcing firms; middle panel displays the probability (ρ) that a submission wins the lottery; right panel displays economy-wide search coasts (C) normalized to equal 100 when H-1Bs are evenly split between regular and outsourcing firms.

Figure C2

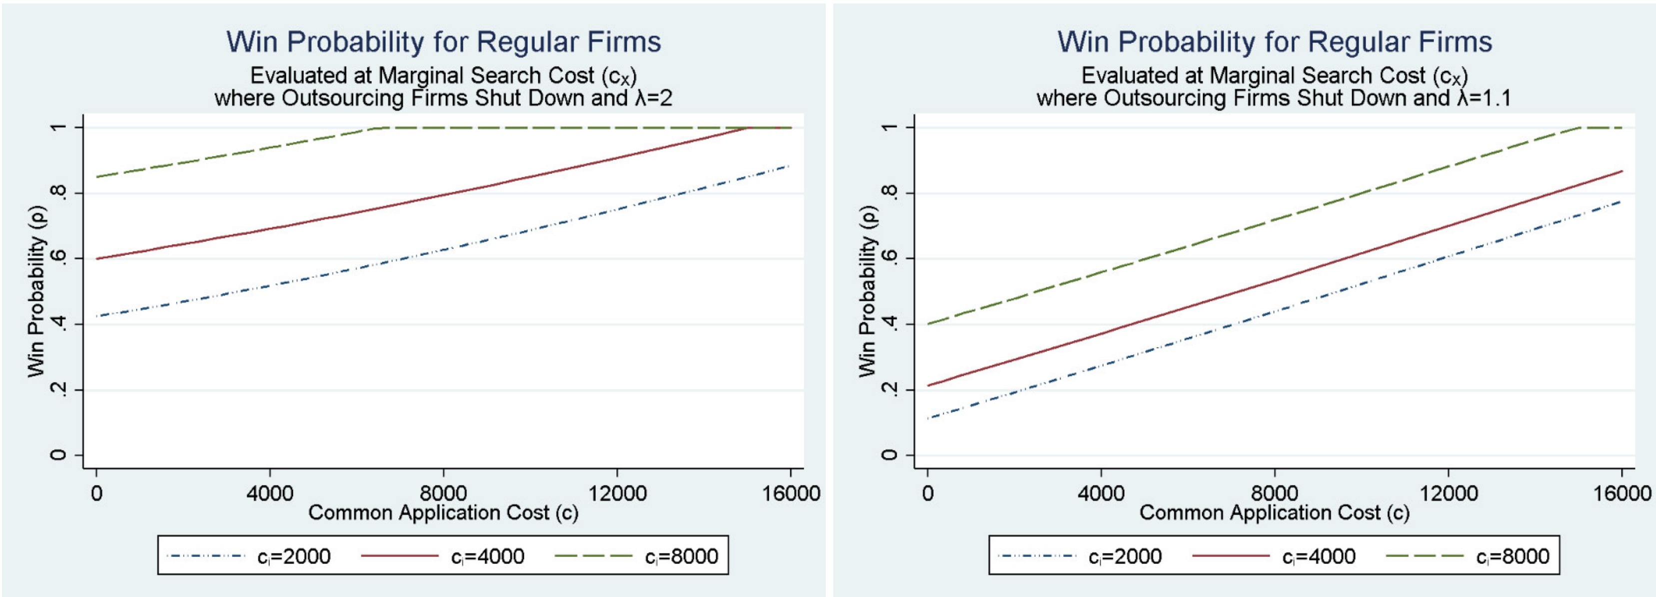

Graphs display the probability (ρ) that an H-1B submission wins the lottery when outsourcing firms' marginal search costs (c_X) are so high that they shut down and all petitions are filed by regular firms (i). This probability is a function of both the common application cost (c) and the marginal search cost of regular firms $\left(c_i F_{iA}^{\lambda-1}\right)$, which is assumed to be common across all regular firms for this exercise. Outsourcing firms are governed λ=1, implying that they face constant marginal search costs. Left panel assumes λ=2 for regular firms; right panel assumes λ=1.1.