

Testimony of
Dr. Britta Glennon, PhD
Assistant Professor of Management
at
The Wharton School of the University of Pennsylvania
before
The U.S. Senate Committee on the Budget

Wednesday September 13, 2023

Hearing: “Unlocking America’s Potential: How Immigration Fuels Economic Growth and Our
Competitive Advantage”

The existing academic research strongly indicates that immigration leads to large gains in productivity, innovation, and business and economic growth. As a result, the current, outdated U.S. immigration system is putting the U.S. at a competitive disadvantage. When we restrict immigration, we lose...and other countries gain instead.

Our restrictive immigration policies have already motivated companies to move jobs and investment out of the U.S. For example, my workⁱ shows that the 2004 reduction in the H-1B visa cap led U.S. multinational companies to move tens of thousands of jobs to the subsidiaries that they operate abroad – both at existing subsidiaries and by opening brand new ones! In other words, the U.S.’s restrictive immigration policies sent jobs and investment to other countries—especially China, India, and Canada—instead. On the other hand, evidence also shows that welcoming migrants can increase domestic investment. For instance, William Olney and Dario Pozzoli found that an influx of refugees led to firm-level *onshoring* of investmentⁱⁱ. In short, policymakers should consider that restrictive immigration policies implemented to protect U.S. jobs are actually likely to have the unintended consequence of pushing U.S. business investment abroad.

But our failure to reform outdated immigration laws is not only causing us to lose workers. We are also losing start-ups. As one example, my work with Saerom Leeⁱⁱⁱ has documented that immigrant would-be entrepreneurs in the U.S. are going to Canada instead to form their start-ups because there is no start-up visa in the U.S. and thus no clear legal path for immigrants to found a business here^{iv}. The loss of immigrant entrepreneurship to other countries is especially significant since recent work^{vi} has found that immigrants—of all skill levels—were 80% more likely than U.S.-born citizens to start a business in the U.S. Moreover, immigrants were more likely to start a business at every business size, from small mom-and-pop shops all the way to high-growth start-ups. And these immigrant-founded companies created 42% more jobs than U.S. citizen-founded companies. So, if immigrants originally destined for the U.S. instead choose to go to Canada—for example—they end up creating new jobs for Canadians, rather than for Americans.

Immigrants also disproportionately positively contribute to innovation, which has historically given the U.S. a competitive edge relative to countries that do not attract or permit immigrants. A recent study^{vii} found that immigrants, who account for 14% of the U.S. population, are responsible for 36% of aggregate innovation in the U.S. Not only do immigrants patent more than U.S.-born workers, but they also increase U.S.-born workers’ productivity by 13 percentage points due to increased collaboration between U.S.-born and immigrant workers and the recombination of diverse knowledge and ideas. Another recent study^{viii} finds that inflows of immigrants since 1965 contributed to an additional 8% growth in per capita innovation and 5% growth in local wages and per capita output in the U.S. Indeed, a recent review in the *Journal of Economic Perspectives*^{ix} of the toolkit available to policymakers for boosting innovation points to immigration policy as one of the most effective and highest net benefit tools available. Yet another study^x turns to history for a clear example of what happens to American innovation when immigration flows are restricted: immigration quotas in the 1920s caused a 68% decline in patenting, in large part because Americans were less innovative without immigrants around. We risk repeating that history.

Finally, my own work^{xi} with Francisco Morales, Seth Carnahan, and Exequiel Hernandez shows that businesses based in countries with more restrictive immigration laws are at a competitive disadvantage. They perform worse on average than firms in countries with more open immigration systems, both because they get less talent and because of the lost collaboration between U.S.-born and immigrant workers. But I am not the only one to show that access to immigrants affects business performance. For instance, other research focused on start-ups^{xix} has shown that winning the H-1B lottery improves the likelihood of receiving venture capital funding and successfully going public. Similarly, Michael Clemens

and Ethan Lewis find that the ability to obtain seasonal immigrant workers increased firm revenue and investment with no decline in U.S.-born worker employment^{xiv}. In short, the ability to hire immigrants of all skill levels matters a lot for firm performance and innovation.

The implications of the empirical evidence are clear. When we choose to not reform our outdated and restrictive immigration system, we put our companies at a competitive disadvantage relative to companies in other countries. Start-ups in particular are more likely to fail when they cannot hire immigrants. The U.S.'s outdated immigration laws hurt our competitiveness and economic dynamism, sending jobs, start-ups, investment, and innovation to countries that do recognize the economic benefits that immigrants bring.

ⁱ Glennon, Britta. (2023) "How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program." *Management Science*. <https://pubsonline.informs.org/doi/full/10.1287/mnsc.2023.4715>

ⁱⁱ Olney, W. W., & Pozzoli, D. (2021). The impact of immigration on firm-level offshoring. *Review of Economics and Statistics*, 103(1), 177-195. <https://direct.mit.edu/rest/article-abstract/103/1/177/97762>

ⁱⁱⁱ Lee, Saerom, and Britta Glennon. (2023) "The Effect of Immigration Policy on Founding Location Choice: Evidence from Canada's Start-up Visa Program." No. w31634. National Bureau of Economic Research. <https://www.nber.org/papers/w31634>

^{iv} Agarwal, R., Ganco, M., & Raffiee, J. (2022). "Immigrant entrepreneurship: the effect of early career immigration constraints on new venture formation." *Organization Science*, 33(4), 1372-1395. <https://dl.acm.org/doi/abs/10.1287/orsc.2021.1485>

^v Kerr, S. P., & Kerr, W. R. (2022). "Immigration Policy Levers for US Innovation and Start-Ups." *Innovation and Public Policy*, 85.

^{vi} Azoulay, Pierre, Benjamin F. Jones, J. Daniel Kim, and Javier Miranda. (2022) "Immigration and Entrepreneurship in the United States." *American Economic Review: Insights*, 4 (1): 71-88.

<https://www.aeaweb.org/articles?id=10.1257/aeri.20200588>

^{vii} Bernstein, S., Diamond, R., Jiranaphawiboon, A., McQuade, T., & Pousada, B. (2022). "The contribution of high-skilled immigrants to innovation in the United States." No. w30797. National Bureau of Economic Research.

<https://www.nber.org/papers/w30797>

^{viii} Burchardi, K. B., Chaney, T., Hassan, T. A., Tarquinio, L., & Terry, S. J. (2020). "Immigration, Innovation, and Growth" No. w27075. National Bureau of Economic Research. <https://www.nber.org/papers/w27075>

^{ix} Bloom, Nicholas, John Van Reenen, and Heidi Williams. (2019). "A Toolkit of Policies to Promote Innovation." *Journal of Economic Perspectives*, 33 (3): 163-84. <https://www.aeaweb.org/articles?id=10.1257/jep.33.3.163>

^x Moser, Petra and San, Shmuel. (2020) "Immigration, Science, and Invention. Lessons from the Quota Acts." <https://ssrn.com/abstract=3558718>

^{xi} Glennon, B., Morales, F., Carnahan, S., & Hernandez, E. (2021). "Does employing skilled immigrants enhance competitive performance? Evidence from European football clubs." *Management Science*, accepted.

<https://www.nber.org/papers/w29446>

^{xii} Chen, J., Hsieh, S., & Zhang, F. (2021). "The role of high-skilled foreign labor in startup performance: Evidence from two natural experiments." *Journal of Financial Economics*, 142(1), 430-452.

<https://www.sciencedirect.com/science/article/abs/pii/S0304405X2100235X?dgcid>

^{xiii} Dimmock, S. G., Huang, J., & Weisbenner, S. J. (2022). "Give me your tired, your poor, your high-skilled labor: H-1b lottery outcomes and entrepreneurial success." *Management Science*, 68(9), 6950-6970.

<https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2021.4152>

^{xiv} Clemens, M. A., & Lewis, E. G. (2022). "The effect of low-skill immigration restrictions on US firms and workers: Evidence from a randomized lottery" No. w30589. National Bureau of Economic Research.

<https://www.nber.org/papers/w30589>



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

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

Britta Glennon

To cite this article:

Britta Glennon (2023) How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program. Management Science

Published online in Articles in Advance 15 Mar 2023

. <https://doi.org/10.1287/mnsc.2023.4715>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2023 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

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

Britta Glennon^a

^a Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Contact: bglennon@wharton.upenn.edu,  <https://orcid.org/0000-0003-4702-8599> (BG)

Received: February 4, 2021

Revised: September 2, 2021; May 17, 2022;
June 17, 2022

Accepted: June 27, 2022

Published Online in Articles in Advance:
March 15, 2023

<https://doi.org/10.1287/mnsc.2023.4715>

Copyright: © 2023 The Author(s)

Abstract. Highly skilled workers are not only a crucial and relatively scarce input into firms' productive and innovative processes, but are also a critical resource determining competitive advantage. An increasingly high proportion of these workers in the United States were born abroad and permitted to work on skilled worker visas. How do multinational firms respond when artificial constraints, namely, policies restricting skilled immigration, are placed on their ability to hire scarce human capital? This paper combines visa microdata and comprehensive data on U.S. multinational firm activity to demonstrate that firms respond to restrictions on H-1B immigration by increasing foreign affiliate employment at the intensive and extensive margins, particularly in China, India, and Canada. The most impacted jobs were R&D-intensive ones, but there is some evidence that non-R&D employment was also affected. The paper highlights a means by which firms can circumvent constraining policies and mitigate country-level risk, and it also suggests that, for the average multinational company (MNC), this means is imperfect; for every visa rejection, they hire 0.4 employees abroad. The most globalized MNCs are the most likely to respond to these restrictions by offshoring, highlighting that firm capabilities—in the form of prior internationalization—shape the decision and ability to offshore in response to skilled immigration restrictions; indeed, these firms hire 0.9 employees abroad for every visa rejection. More broadly, the paper provides evidence of a push factor for internationalizing knowledge activity: artificial constraints on resources result in firms circumventing restrictive policies in ways that may not be anticipated by policy makers.

History: Accepted by Alfonso Gambardella, business strategy.



Open Access Statement: This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. You are free to download this work and share with others, but cannot change in any way or use commercially without permission, and you must attribute this work as "Management Science. Copyright © 2023 The Author(s). <https://doi.org/10.1287/mnsc.2023.4715>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by-nc-nd/4.0/>."

Funding: This work was supported by the Mack Institute for Innovation Management.

Supplemental Material: The online appendix and data are available at <https://doi.org/10.1287/mnsc.2023.4715>.

Keywords: immigration • offshoring • multinational companies • human capital • global strategy

1. Introduction

Access to talent is a critical resource determining firm competitive advantage (e.g., Hatch and Dyer 2004, Campbell et al. 2012, Mayer et al. 2012), and that talent is not constrained by country borders; indeed, the relevant talent pool is global. In the United States, firms tend to meet demand for top talent by importing foreign labor using skilled worker visas. However, government-imposed restrictions on skilled immigration constrain the ability of firms to import such foreign talent.¹ This paper examines the extent to which U.S. firms circumvent such constraining policies and mitigate country-level risk through increasing foreign affiliate employment.

Although there is an extensive body of literature documenting the impact of skilled immigration on native

labor market outcomes, there is surprisingly little focus on how skilled immigration shapes firm decision making and outcomes (Kerr et al. 2015a). Fortunately, in recent years, a series of careful empirical papers began to address this gap, examining the impact of skilled immigration on a range of firm outcomes, such as innovation (Ashraf and Ray 2017, Wu 2017, Laursen et al. 2020), firm structure and employment (Kerr et al. 2015b, Mayda et al. 2018, Doran et al. 2022), stock market valuation (Bahar et al. 2020), venture capital funding (Li 2020, Dimmock et al. 2022), and performance (Ghosh et al. 2015, Hernandez and Kulchina 2020, Mayda et al. 2020, Glennon et al. 2021). Although these studies significantly contribute to our understanding of the impact of hiring skilled foreign labor on some firm outcomes, we

know very little about how restrictions on firm ability to hire skilled foreign labor in the home country shape the strategic global expansion decisions of multinational firms.

Knowledge—and the human capital in whom knowledge resides—is a key source of sustained competitive advantage for firms (Kogut and Zander 1992). Multinational companies (MNCs) in particular are posited to exist because they have an advantage over external mechanisms in creating and transferring knowledge across borders, and knowledge is a critical source of advantage that MNCs seek to exploit or obtain when expanding abroad (Kogut and Zander 1993). However, although MNCs are able to transfer knowledge across borders better than other actors, there are strong reasons to concentrate knowledge activities and knowledge workers in the home country. First, because knowledge contains both explicit and tacit components (Polanyi 1958) and distance creates challenges for the transfer of tacit knowledge (Szulanski 1996, Gupta and Govindarajan 2000), transferring knowledge across borders is costly even within an MNC (Teece 1977). Second, R&D activities are subject to economies of scale and scope, making it less effective to have R&D in multiple locations. Finally, because of the strategic importance of knowledge as a core capability of the firm, firms may wish to keep it close to the core of the company to control strategic information and minimize external technology diffusion. For these reasons, MNCs typically prefer to keep their most important knowledge workers and R&D activity near headquarters—according to the Bureau of Economic Analysis (BEA), 85% of global R&D expenditures by U.S. MNCs are in the United States²—unless other factors push or pull them to locate such activities abroad.

Prior work considers many such pull factors. Firms may choose to conduct R&D abroad in order to support local manufacturing operations and adapt products to large local markets (Mansfield et al. 1979, Hirschey and Caves 1981) or because of cheaper input costs (Kumar 2001, Athukorala and Kohpaiboon 2010) or fiscal and tax incentives (Hines 1995, Hall and Van Reenen 2000). Some MNCs conduct R&D abroad in order to gain competitive advantage by tapping into, harnessing, and recombining different sources of expertise and knowledge around the world (Cantwell 1995, Florida 1997, Kuemmerle 1997, Chung and Alcácer 2002).

In this paper, I introduce an important but overlooked push factor that could drive MNCs to expand their knowledge activity and employment abroad: domestic restrictions on the ability to bring skilled knowledge workers from abroad. As described above, skilled workers are crucial and relatively scarce inputs into firms' productive and innovative processes and are a critical resource determining competitive advantage. Because an increasingly large proportion of these workers are

born abroad, multinational firms need to obtain visas for them to work at headquarters. When home country immigration regimes become more restrictive and obtaining visas is no longer a consistent option, it is plausible that MNCs respond to such restrictions on skilled immigration by offshoring employment to their foreign affiliates. Examining that possibility is the focus of this paper.

Determining the causal impact of skilled visa restrictions on offshoring is a difficult exercise as unobserved factors may simultaneously affect a firm's choice to hire skilled immigrants and to expand its foreign activity. Two exogenous shocks to high-skilled immigration make it possible to empirically control for such confounding factors: (1) a 70% drop in permitted H-1B visa issuance in the United States and (2) random rationing of visas in the 2007 and 2008 H-1B visa lotteries. Although the two shocks are different in nature, the results are consistent, suggesting that there is an underlying empirical regularity in firms' responses.

The empirical regularity is identified by linking three different data sets that make it possible to measure each MNC's constraints and reactions to constraints on foreign hiring. First, they provide information on each firm's requested visas (demand) and issued visas (realized supply). Second, they measure each firm's response in terms of the expansion or contraction of foreign affiliate employment in each country. The demand and supply data are derived from H-1B visa microdata obtained from Labor Condition Applications (LCA) and by Freedom of Information Act (FOIA) request, respectively.³ Detailed microdata on the financial and operating characteristics of both the U.S. parent companies and their foreign affiliates are obtained from the BEA annual surveys on the U.S. Direct Investment Abroad (USDIA) database.

I find that foreign labor supply restrictions caused increases in foreign affiliate activity at both the intensive margin (U.S. multinationals employed more people at their existing foreign affiliates) and the extensive margin (U.S. multinationals opened foreign affiliates in new countries). On the intensive margin, the effect is concentrated on R&D foreign affiliate employment, but when the intensive and extensive margins are considered jointly, both R&D and non-R&D employment increase in response to foreign labor supply restrictions, suggesting that complementarities between production and R&D may magnify the effect of the immigration restrictions and lead to the offshoring of both skilled and unskilled labor and their respective activities. The expansion of foreign affiliate employment is especially concentrated in three countries—China, India, and Canada—suggesting that the expansion of foreign affiliate activity might operate through both direct (direct access to raw human capital) and indirect (indirect access to skilled immigrants in a less restrictive environment) channels. Finally, more globalized MNCs are the most likely to

respond to these restrictions, and their growth is the least constrained by the restrictions, highlighting that firm capabilities—in the form of prior internationalization—shape the decision and ability to offshore in response to skilled immigration restrictions. These results imply that immigration restrictions serve as an important push factor for internationalizing knowledge workers and activity; firms prefer to have key knowledge workers at headquarters, but skilled immigration restrictions push them to locate said workers abroad instead.

Five aspects of this paper are novel. First, it provides causal empirical evidence that restrictions on high-skilled immigration cause the offshoring of skilled jobs, highlighting a new and important but previously overlooked explanation for the global expansion of MNCs and the geographic location of their foreign affiliates. Whereas high-profile anecdotal cases⁴ suggest that restricting skilled immigration flows could lead to the offshoring of jobs, the degree to which these anecdotal examples can be viewed as a systematic response to skilled immigration restrictions—or foreign supply shocks of any kind—was previously unknown. Second, it documents which job types are most affected by this push factor: whereas R&D foreign affiliate jobs are the most affected, non-R&D foreign affiliate jobs are also impacted. Third, it documents that prior internationalization is an important moderator of the impact of skilled immigration restrictions on offshoring employment and firm-specific capabilities influence the rate of such offshoring. Fourth, to the best of my knowledge, it is the first paper to use a matched firm-level data set of H-1B visas and multinational firm activity and to show how such a data set can be used to answer previously unanswered questions about the link between immigration and multinational firm activity. And, fifth, it contributes to a growing but still small body of evidence showing that immigration influences firm behavior as well as providing insights into the connections between the fields of immigration and (global) strategy.

2. Multinational Firms and the Search for Global Talent

The resource-based view of the firm explains sustained differences in firm performance through access to valuable and rare resources (e.g., Barney 1986, Conner 1991, Mahoney and Pandian 1992, Peteraf 1993). In particular, access to talent is a critical resource determining firm advantage (e.g., Hatch and Dyer 2004, Campbell et al. 2012, Mayer et al. 2012). But companies face growing global competition for the scarcity of talented people in the world (e.g., Collings et al. 2019). Bartlett and Ghoshal (2002) argue that the scarcity of talented people, in fact, is the key constraining resource companies face.

The relevant talent pool is no longer a local one, but instead spans the globe. The global distribution of science and engineering university degrees provides some

suggestive evidence of the global nature of the relevant talent pool; 10% of these degrees went to students in the United States, whereas 25%, 22%, and 12%, respectively, went to students in India, China, and the European Union (National Science Foundation 2020). Kerr et al. (2017) describe a global distribution of talent, whereas Arora and Gambardella (2005) document the abundant supply of engineering and technology graduates in emerging economies. Branstetter et al. (2021) argue that multinational firms that are able to plug into the large and growing pool of scientists and engineers in the developing world may have a productivity advantage.

U.S. firms—and especially those in knowledge-based industries such as information technology (IT)—have relied heavily on the H-1B program to access foreign talent, particularly foreign talent educated at U.S. universities. Bound et al. (2015) document the increase in high-skilled foreign-born IT workers in parallel to the rising importance of IT in the United States and note that foreign-born IT workers are more likely to hold advanced degrees than their U.S. counterparts. According to the 2015 Silicon Valley Competitiveness and Innovation Report, 56% of STEM workers and 70% of software engineers in Silicon Valley were foreign-born. The preference of U.S. firms to import foreign talent to meet demand for top talent at their U.S. offices is well-documented and explained as a consequence of agglomeration economies by Kerr et al. (2017).

However, government-imposed restrictions on skilled immigration visas restrict the ability of U.S. firms to import such foreign talent. Furthermore, these restrictions differentially constrain the ability of firms located in the restrictive policy-affected country to access top talent relative to firms located in other parts of the world. The inability to access the top talent, therefore, constrains the affected firms' ability to sustain their competitive advantage and provides an advantage to firms located in other parts of the world unless they can find another solution. In this paper, I suggest that one such solution is to hire such talent at their foreign affiliates instead although my results suggest that this solution may be an imperfect one.

3. Firm Capabilities and the Relationship Between Offshoring and Immigration

Although some recent research has begun to incorporate immigration and offshoring as interconnected decisions in the same general equilibrium framework (Olney 2012, Ottaviano et al. 2013, Peters 2017, Morales 2019, Mehra and Shen 2022), there is limited empirical evidence about how they are interconnected, particularly when the focus is on high-skilled immigration. The expected sign and magnitude of the effects, if any, of high-skilled immigration restrictions on MNC foreign affiliate employment are theoretically ambiguous, and furthermore, there is

very little causal evidence of the effect of skilled migration on firm offshoring decisions.

There are four broad reasons to expect complementarities between immigrants and foreign direct investment (FDI) and trade. The first is through consumer preferences/tastes; diasporas may increase demand for import of “nostalgia” goods (Bronnenberg et al. 2012, Atkin 2013). A second is through an information channel; immigrants may have superior knowledge about the host country that can be leveraged by the firm to overcome the distances that give rise to liability of foreignness (Zaheer 1995, Saxenian 2002, Marino et al. 2020). A third is through the provision of social capital to enforce contracts when courts do not work well (Greif 1993, Besley and Coate 1995). Finally, migrants may make the destination and origin more “similar” in terms of skill endowments. The implication is that a reduction in immigration could actually reduce foreign affiliate employment and activity because setting up new foreign affiliates would be more difficult without the diaspora networks. There is some empirical support for this view; Caliendo et al. (2021), Gould (1994), Hiller (2013), Iranzo and Peri (2009), and Murat and Pistori (2009) establish a positive link between immigration and trade, whereas Buch et al. (2006), Burchardi et al. (2019), Cuadros et al. (2018), Foley and Kerr (2013), Hernandez (2014), Javorcik et al. (2011), Kugler and Rapoport (2011), and Morales (2019) find a positive link between FDI and immigration.

However, a few recent papers and books find some evidence of substitutability between immigrants and offshore workers. Lewin et al. (2009) and Branstetter et al. (2018) argue that shortages of technical talent is an important driver of offshoring innovation using survey data and descriptive evidence. Ottaviano et al. (2018) find that immigrants may substitute for imported intermediate inputs in the UK services sector. And Olney and Pozzoli (2021) find that an influx of refugees into Denmark reduced offshoring as measured by import data. The latter two papers, although they take a causal approach, are unable to measure the direct substitution of immigrants and foreign affiliate labor because of data constraints, instead capturing offshoring with import data.

Much of this literature focuses on unskilled immigration and offshoring with a few exceptions (Lewin et al. 2009, Ottaviano et al. 2018, Mehra and Shen 2022), but the discussion is more nuanced when the focus is on skilled immigration and the decision to offshore innovative activity. There are strong potential benefits to conducting R&D or other skilled activities abroad as well as significant potential costs. Different firms are likely to weigh this trade-off differently; only some firms have the capabilities to—and want to—overcome such costs to leverage the potential opportunities of global R&D. In the following paragraphs, I lay out the determinants of R&D offshoring and discuss how firm capabilities might affect the weighing of the inherent trade-offs.

The first potential benefit to conducting R&D or other skilled activities abroad comes from a market access motive: firms may wish to support local manufacturing operations and tap into growth opportunities through R&D activities related to adapting products or technology to host country conditions (Mansfield et al. 1979, Hirschey and Caves 1981). MNCs may want to seize the opportunity to arbitrage on international cost differences (Athukorala and Kohpaiboon 2010), fiscal and tax incentives (Hines 1995, Hall and Van Reenen 2000), or other different policy environments (e.g., Zhao 2006).

A more recent but important motivator is the desire to gain competitive advantage by tapping into, harnessing, and recombining different sources of expertise and knowledge around the world. Because knowledge spillovers do not pass easily across national borders (Jaffe et al. 1993, Audretsch and Feldman 1996), countries have different knowledge profiles (Furman et al. 2002, Alcácer and Chung 2007). Conducting R&D abroad, thus, provides an opportunity to access and recombine diverse knowledge to create new innovation. A related reason, of course, is the search for global talent discussed in more detail in Section 2. Global R&D networks then provide an opportunity for economies of specialization and an international division of innovative labor within the firm. Firms might leverage specialized skills and human capital from different parts of their network to engage in larger projects; any single location will eventually be constrained by limits in resources or specialized knowledge. Alternatively, the firm might reshape the organization of R&D such that some locations specialize in less technologically intensive R&D activities, whereas others specialize in more technologically intensive R&D activities. Indeed, these types of motivations could have significant implications for firm growth; those firms able to leverage this kind of international division of innovative labor within the firm might see substantial benefits from global R&D (Branstetter et al. 2021).

Hence, pursuing R&D opportunities abroad can be an important source of competitive advantage but not for every firm. Unlike production activities, knowledge—and skilled human capital—is a critical source of competitive advantage for the firm (Kogut and Zander 1992, Alcácer 2006), and as a result, there are strong reasons to keep it close to home. Because knowledge is strategically important, firms may want to keep knowledge and the human capital in which it resides close to the core of the company to maintain control and reduce knowledge spillovers. In addition, because tacit knowledge is not easily transferred across borders, even within an MNC (Polanyi 1966, Teece 1977, Szulanski 1996, Gupta and Govindarajan 2000), it may be more efficient to concentrate knowledge activities and workers in one location. Finally, skilled workers are not easily substitutable. In short, the costs of transferring knowledge internally are nontrivial and vary substantially across firms.

Only some firms can overcome—and are interested in overcoming—these challenges. In particular, one might expect that firms with more prior international experience would be best positioned to respond to skilled immigration restrictions by offshoring and the least likely to have their growth constrained by such restrictions. These firms have already had time to develop and learn the optimal strategies⁵ for overcoming many of the challenges inherent in operating skilled activities. They also have already invested in the fixed costs involved in opening up skilled activity in new locations. Less internationalized MNCs may be at a disadvantage because they have had less time to develop the skills and capabilities for conducting R&D abroad and, thus, may be less responsive to immigrant restrictions or, if forced to respond, may find it more harmful.

However, drawing on insights around multinationality and operational flexibility, one might also expect more internationalized MNCs to be less likely to respond to immigration restrictions by hiring more abroad. Multinationality gives firms real options and the managerial capabilities to respond flexibly to external shocks, which less globalized firms lack (Kogut and Kulatilaka 1994, Lee and Makhija 2009, Chang et al. 2016). As a result, more global firms may have complementary resources that allow them to respond along dimensions besides offshoring, whereas offshoring may be one of the few options available to less internationalized MNCs. For instance, perhaps more global MNCs can rely on relationships with other firms via alliances or outsourcing, or perhaps they have the flexibility to adjust their geographic division of tasks rather than individuals. In short, more globalized MNCs might respond to immigration restrictions by using other organization choices to substitute for skilled labor rather than by offshoring. These choices, of course, imply quite constrained firm growth regardless of the previous internationalization of the firm. In short, *ex ante*, it is unclear how one might expect firm capabilities in the form of prior internationalization to shape the substitutability of skilled immigrants and foreign affiliate skilled workers within the firm.

Finally, this discussion assumes that restrictions on skilled immigration only affect skilled labor and activities abroad. And yet there is evidence in the literature of a synergistic relationship between the production and R&D functions of the firm (Cohen and Zysman 1987, Naghavi and Ottaviano 2009, Pisano and Shih 2012, Fort et al. 2020); the process of developing an idea from a concept into a marketable product or service may require continuous collaboration and knowledge transfer between R&D and production personnel (Schumpeter 1939). But, because the complementarities between R&D and production are likely to depend on the nature of the R&D being undertaken and the goods being manufactured and because it is unclear whether skilled immigration restrictions affect both the intensive and extensive

margin, it is not clear *ex ante* whether non-R&D foreign affiliate employment is also affected by skilled immigration restrictions.

4. Research Design and Data

4.1. Overview of the H-1B Program

There are multiple ways in which firms can hire foreign high-skilled workers: the H-1B, L-1, O, OPT, and TN visas are just a few examples.⁶ The first is the most widely used and is the focus of this paper, but I provide some discussion of the alternatives (especially the L-1, which permits overseas branches or subsidiaries of MNCs to transfer foreign workers within the company) in the online appendix. Although the H-1B visa is the focus of this paper—only painting a partial picture of immigration restrictions—it is worth noting that multinational firms are probably less constrained than other firms because they have the L-1 option. As such, one might imagine that the estimates in this paper are lower bounds.

The H-1B visa is a nonimmigrant visa that enables firms to hire foreign workers in the United States for a three-year period, renewable once for a total of six years. They are called “nonimmigrant” visas because they allow those with H-1Bs to stay in the United States only temporarily. However, they are also “dual intent” visas, which means that workers can reside in the United States with a nonimmigrant status while simultaneously applying for permanent residency. H-1B visas make up about 50% of temporary work visas and are used to employ foreign workers in “specialty occupations.”⁷ Firm interviews conducted with the author suggest that U.S. firms typically use H-1B visas to hire international students at domestic universities.

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

Second, the H-1B application is a two-stage process. In the first stage, firms must file an LCA with the Department of Labor (DOL) Employment and Training Administration. This first stage measures demand. There is no limit—beyond cost—on the number of LCAs that a firm can file, so demand is measured independent of whether an H-1B is ultimately issued or not. In the second stage, after LCA approval, the firm must file an I-129 petition with U.S. Citizenship and Immigration Services (USCIS), which makes the ultimate determination about the visa application. H-1B cap constraints are imposed in the second stage, in which the final decision is made, so this

stage measures realized supply. The two-stage structure of the H-1B application process allows measurement of exactly how constrained each firm was as the cap grew more restrictive over time by comparing LCA requests (demand) and issued H-1B visas (realized supply).

Third, variations in the cap on H-1B visa supply provide a source of exogenous variation. The number of new H-1B visas that can be issued to private sector businesses has been subject to a cap since their inception in the Immigration Act of 1990. Online Figure A8-1 plots the cap on the number of H-1B visas by fiscal year. This cap is set by Congress and the President and only applies to new H-1B visas issued to private sector businesses.⁸

There are three discrete phases of interest in terms of hiring constraints over time. The first phase is one in which the hiring constraint was not binding: throughout most of the 1990s, the cap was set at 65,000 visas, and applications rarely outstripped supply.⁹ Phase 2 began in 1998–2000, when the cap was increased to 195,000 by the American Competitiveness and Workforce Improvement Act of 1998 and the American Competitiveness in the Twenty-First Century Act of 2000. During this period, the cap limits were never reached. Phase 3 occurred when trends in increasing visa availability reversed in 2004, and the cap reverted to 65,000 visas although 20,000 additional visas were granted to applicants with a graduate degree in 2006 (for a total of 85,000). Since then, the cap has not changed, and it has been (increasingly) binding in every year since 2004. Because of data constraints, this paper focuses on Phases 2 and 3.

The fourth relevant characteristic of the H-1B program is the random variation introduced by the process by which H-1B visas are distributed. H-1B petitions are distributed in a first-come, first-served fashion or by lottery in especially high demand years. The process is

illustrated in Online Figure A8-2. On the first business day of each April, USCIS begins accepting H-1B applications from firms seeking permits that count toward the following fiscal year. Because the H-1B visa program operates on a first-come, first-served basis, petitions are accepted until the cap hits, at which point no more petitions are processed. The end of the application period is demarcated by the “final receipt date,” which is the date on which they receive enough applications to fill the remaining available permits under the cap. Any cap-subject petitions submitted after the final receipt date are automatically rejected. This date is announced by USCIS in a press release, and it varies every year as shown in Table 1. On the date(s) that the available permits are exhausted, a computer-generated random selection process selects the petitions that will be processed. The dates of the lottery are not announced in advance and are, in fact, unknown in advance; they are determined by the number of applications received on different dates. These dates are only made known to firms after the cap is reached. In April 2007 and 2008 (as well as several other years, demarcated by a star in Table 1), USCIS received so many petitions within the first week that all cap-subject petitions were distributed by lottery for fiscal years 2008 and 2009. The lottery generated a random negative shock in the supply of foreign-born skilled workers to firms; the second empirical approach exploits the lottery-generated randomized variation from the H-1B visa lotteries in those two years, which allows for a causal interpretation of the effect of constrained foreign-born skilled worker supply on the offshoring of skilled jobs.

The final relevant characteristic is significant for measuring demand. The timing of petitions can be used to reveal whether an application is for a cap-subject H-1B

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

Fiscal year	Final receipt date	Days in filing period	Number of lottery-subject H-1B petitions received during the filing period
2004	February 17, 2004	323	
2005	October 1, 2004	184	
2006	August 10, 2005	132	
2007	May 26, 2006	56	
2008*	April 3, 2007	3	150,000
2009*	April 7, 2008	7	163,000
2010	December 21, 2009	265	
2011	January 26, 2011	301	
2012	November 22, 2011	236	
2013	June 11, 2012	72	
2014*	April 7, 2013	7	124,000
2015*	April 7, 2014	7	172,500
2016*	April 7, 2015	7	233,000
2017*	April 7, 2016	7	236,000
2018*	April 7, 2017	5	199,000
2019*	April 6, 2018	5	190,098

Note. Years that are demarcated with an asterisk (*) indicate years in which USCIS received so many petitions within the first week that all cap-subject petitions were distributed by lottery.

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

However, the rising demand for H-1B visas and the first-come, first-served nature of the distribution process changed firm behavior. Firms that needed cap-subject H-1B visas wanted to submit their petitions as early as possible—and certainly before April—to ensure the submission would be before the final receipt date. Online Figure A8-3 illustrates the change in the timing of LCA applications; as demand for cap-dependent H-1B visas increased, LCA applications were filed earlier. I infer that LCA applications submitted in the first quarter of the calendar year are for cap-dependent H-1B petitions.

4.2. Data

I use three sources of data to generate a unique data set that permits the analysis of the link between MNC hiring decisions and their response to high-skilled immigration constraints. The first data set provides information about multinational activity, including employment and R&D expenditures. The second and third are particularly useful because they provide information about both the demand and the realized supply for foreign workers.

4.2.1. Multinational Activity Data. The data used to examine multinational activity are confidential firm-level data from the Bureau of Economic Analysis's annual surveys on U.S. Direct Investment Abroad. BEA is under a congressional mandate¹⁰ to track investment into and out of the United States, and as such, their data comprise the most comprehensive available data on U.S. multinational activity abroad. Of particular importance is that the data includes foreign affiliate employment, which is the primary variable of interest for this paper. In benchmark years (every five years), foreign affiliate employment is further categorized by whether it is R&D employment or not. In parts of the analysis, R&D foreign affiliate employment is used as an imperfect proxy for skilled foreign affiliate employment. An additional—related—variable of interest is R&D expenditures, which is used for distinguishing between foreign affiliates that perform R&D and those that do not. Detailed definitions of R&D employment and expenditures are in the online appendix.

I constructed a panel data set of this activity from 1994 through 2014.¹¹ Each firm may report on a consolidated

basis for multiple affiliates in the same country under certain conditions,¹² so rather than conducting analysis at the affiliate level, I aggregate all foreign affiliate activity up to the host country level for a given parent firm in a given year. The panel contains 2,263 U.S.-based firms with multinational activity. Importantly, the U.S. MNCs in the BEA data form the basis of the analytical sample, and as such, the sample does not include U.S. domestic firms or foreign multinational firms.

4.2.2. H-1B Data. The next two data sets allow measurement of firm-level hiring patterns of foreign-born workers. The information sources are worker-level application records from the U.S. DOL and worker-level approved H-1B petition data from USCIS.

Measures of firm-level demand for H-1B visas come from the DOL LCA data. Before a firm can file a petition with USCIS, it must file an LCA with the DOL. These applications have been made publicly available by the DOL since 2001 and contain information on the employer's name and address, the occupation code of and wage offered to the worker, and the geographic location of the position to be filled by the visa recipient. There is no limit (other than financial constraints) on the number of LCAs that a firm can file. The primary purpose of the LCA is for employers to attest to the employment details of H-1B applicants and affirm that the worker will be employed in accordance with U.S. law.¹³ This data set comprises 6.4 million records between 2001 and 2016, which I aggregate by employer-year and then link to the BEA data.

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

Measures of realized H-1B labor supply come from I-129 H-1B visa applications, obtained by FOIA request. These data are used in the second identification strategy. The original data set contains I-129 petitions from fiscal years 2004–2014, consisting of about 3.3 million petitions, with information on the final decision regarding each petition, the type of visa being requested, the beneficiary's country of birth, the employer name and location, the job code, and other administrative details. An I-129 form is needed for many types of visas, but for the purpose of this paper, the most relevant is the H-1B visa.

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

An examination of the heterogeneity of H-1B petition filings by firm and country—shown in Online Appendix A5—provides some intuition regarding where the expansion of foreign affiliate activity might be expected to take place: 85% of H-1B petition filings were for workers from India or China, suggesting one might expect to see large increases in foreign affiliate employment in India and China. Finally, there is significant skewness in H-1B visas across firms; a nonlinear approach may be appropriate.

4.2.3. Final Data Set. The final data set is at the firm–country–year level and contains 2,263 multinationals.

Of the 2,263 multinationals, 28% filed at least one LCA in 2001, 29% applied for at least one U.S. Patent and Trademark Office (USPTO) patent, and 15% had both at least one LCA application in 2001 and at least one USPTO patent at some point. Summary statistics are shown in Table 2.

5. Empirical Approach and Results

In this section, I empirically examine whether restrictive high-skilled immigration policies caused U.S. multinational companies to hire more foreign labor at their foreign affiliates. My empirical analysis has five parts. In part one, I estimate the impact of constrained skilled immigration supply on the foreign affiliate employment intensive margin (i.e., did foreign affiliate employment at existing foreign affiliates change). A strong identification strategy is necessary to answer this question as many unobserved factors may simultaneously affect skilled immigration and foreign activity expansion, so I use two that are different in nature. The first exploits the 70% drop in permitted H-1B visa issuance in the United States in 2004, whereas the second exploits randomized variation from the H-1B visa lotteries in high demand years. In part two, I estimate the impact on the extensive margin (did multinational firms open more foreign affiliates). In part three, I take a closer look at which jobs and foreign affiliates are most and least affected by restrictions on high-skilled immigration. In part four, I examine geographic location choice, particularly focusing

Table 2. Summary Statistics

Summary statistics of MNCs in 2001					
	N	Mean	Standard deviation	10th percentile	90th percentile
Reporter employment	2,263	7,785.2	29,915.7		
Ratio of LCA applications to U.S. employment in percentage point units	2,263	0.176	0.812	0	0.2835
Number of countries active in	2,263	2.831	5.646		
Number of LCA applications	2,263	11.87	79.85		
Summary statistics of MNCs in 2013					
	N	Mean	Standard deviation		
Reporter employment	2,263	6,300.6	35,868.7		
Number of countries active in	2,263	2.422	5.992		
Summary statistics of existing foreign affiliates in 2001					
	N	Mean	Standard deviation		
FA employment	6,407	1,151.0	3,964.3		
FA R&D expenditure	6,407	2,765.6	26,276.7		
Summary statistics of still-existing foreign affiliates in 2013					
	N	Mean	Standard deviation		
FA employment	5,482	1,559.4	8,307.1		
FA R&D expenditure	5,482	5,741.0	37,289.6		

Note. The 10th and 90th percentiles are based on an average of firms centered on that percentile.

on differences between countries that possess the necessary raw human capital (such as India or China, the primary source of H-1B applicants) versus countries with less restrictive skilled immigration policies (such as Canada).¹⁴ Finally, in part five, I examine whether firm capabilities—in the form of prior internationalization—shape the decision to offshore in response to skilled immigration restrictions.

5.1. Part One: The Intensive Margin

I begin by estimating the impact of constrained skilled immigration supply on the foreign affiliate employment intensive margin. In this part, the dependent variable is total foreign affiliate employment; I break down the effects on R&D and non-R&D foreign affiliate employment in part three. In addition, this part concentrates on country-level foreign affiliate activity that existed prior to the policy changes. The effects measured in this part do not capture new foreign affiliates in new countries that appear as a result of immigration restrictions; I examine the effect on new foreign affiliates in part two. It does, however, capture new entries into any country in which the MNC has already entered. Finally, I use two identification strategies in this part. The first exploits the 2004 drop in the H-1B visa cap, whereas the second exploits randomized variation from the H-1B visa lotteries in high demand years. The robustness of the results to both shocks—which are quite different in nature—is reassuring and suggests that there is an underlying empirical regularity in firms’ responses.

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

My empirical specification can be interpreted as a difference-in-difference estimator—similar to that used in Ashraf and Ray (2017), Ghosh et al. (2015), Kerr and Lincoln (2010), and Xu (2016)—in which the treatment and control groups are categories of firms with different levels of H-1B dependency in 2001 (when the cap was not binding) and the treatment is the reduction in the annual H-1B cap in 2004. In other words, the difference-in-differences approach relies on preexisting variation in demand for foreign-born skilled workers to identify how exogenous constraints in supply affected foreign affiliate employment. Accordingly, the regression compares the change in foreign affiliate employment before and after the policy change across multinationals within the same firm and country that were more dependent on H-1B visas prior to the policy change (the “treatment” group) relative to less dependent on H-1B visas prior to


the policy change (the “control” group). Figure 1 provides a graphic version of the strategy and shows that, whereas foreign affiliate employment growth for non-H-1B dependent firms remained fairly flat after the policy change, extremely H-1B dependent firms experienced rapid growth in foreign affiliate employment after the policy change. Furthermore, the trajectory of foreign affiliate employment growth of both types of firms remained parallel and quite flat prior to the policy change. The regression results confirm the associations in the raw data presented in Figure 1.¹⁵

In the baseline specifications, $H1BDep_i$ is defined as the total LCA applications for a given multinational divided by that multinational’s U.S. employment in 2001 as in Kerr and Lincoln (2010) and Xu (2016). The dependency measure is calculated in a prepolicy year to help address the problem of reverse causality. The dependency measure is my preferred metric because it measures demand for H-1B visas, and it is measured independently of whether an H-1B visa is ultimately issued or not. Furthermore, because of the high cost of application, the dependency measure can be seen as reflecting real measured demand. Finally, the dependency measure closely mirrors the DOL’s own measure of H-1B dependency, namely, “The determination as to whether an employer is H-1B dependent is a function of the number of H-1B nonimmigrants employed as a proportion of the total number of full-time equivalent employees employed in the United States.”¹⁶ The regression specification is as follows:

$$\ln(FAemp_{ict}) = \alpha_c + \alpha_i + \alpha_t + \beta_1 H1BDep_i * policy_t + \epsilon_{ict} \quad (1)$$

where i indexes the firm, c indexes the country, and t indexes the year. Country, firm, and year fixed effects

Figure 1. (Color online) Average Foreign Affiliate Employment by H-1B Dependency



Notes. Non-H-1B dependent firms had zero LCA applications in 2001, whereas very H-1B dependent firms were in the top category of H-1B dependency in 2001. The solid vertical line demarcates the year of the 2004 policy change.

are captured by α_c , α_i , and α_t , respectively. $FAemp$ is foreign affiliate employment. $Policy$ is a dummy variable equal to one for all years after and encompassing 2004 and zero otherwise. This is interacted with $H1BDep$, which is defined in one of two ways. The first is the continuous version of H-1B dependency described earlier, in which the number of LCA applications (the measure of H-1B demand) by firm i in prepolicy year 2001 is normalized by a multinational's employment in the United States in prepolicy year 2001, whereas the second is a new variable equal to one if the firm was in the top group of H-1B dependency in 2001 and zero if the firm had zero LCA applications in 2001. All standard errors are clustered at the firm level, but results are also robust to clustering at other levels, shown in the online appendix. I expect β_1 to be positive; after the 2004 reduction in the H-1B cap, firms that were ex ante more dependent on H-1B visas should be more affected by the policy change and, therefore, more likely to expand their

foreign affiliate activity. Column (1) of Table 3 shows the results of the continuous version, whereas column (2) shows the results of the binary version. Both are statistically and economically significant. The coefficient in the binary version tells us that the 2004 policy change caused highly H-1B dependent firms to increase their foreign affiliate employment by 27% more than a non-H-1B dependent firm. An additional advantage of this specification is that, unlike the long-differences version (shown next), it ensures a consistent sample of firms throughout.

A modification of this specification that interacts year fixed effects with the treatment allows observation of time-varying effects and answers the question: was the firm response to the cap change a one-off effect, or did the effect grow over time as firms felt more constrained? The results are shown in column (3) and provide evidence that firms accelerated offshoring over time as constraints tightened.

Table 3. Effect on Foreign Affiliate Employment for All Firms, 1994–2014

	(1) ln(<i>FA emp</i>)	(2) ln(<i>FA emp</i>)	(3) ln(<i>FA emp</i>)
<i>Continuous treatment</i>	0.0615*** (0.0176)		
<i>Binary treatment</i>		0.274*** (0.0906)	
<i>Year = 2004 # continuous treatment</i>			0.0249 (0.0164)
<i>Year = 2005 # continuous treatment</i>			0.0432** (0.0216)
<i>Year = 2006 # continuous treatment</i>			0.0408* (0.0230)
<i>Year = 2007 # continuous treatment</i>			0.0477* (0.0245)
<i>Year = 2008 # continuous treatment</i>			0.0562*** (0.0171)
<i>Year = 2009 # continuous treatment</i>			0.0720*** (0.0208)
<i>Year = 2010 # continuous treatment</i>			0.0522*** (0.0192)
<i>Year = 2011 # continuous treatment</i>			0.0713*** (0.0179)
<i>Year = 2012 # continuous treatment</i>			0.0848*** (0.0172)
<i>Year = 2013 # continuous treatment</i>			0.103*** (0.0179)
<i>Year = 2014 # continuous treatment</i>			0.132*** (0.0222)
Country fixed effects	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Observations	137,459	57,984	137,459
R^2	0.429	0.429	0.429

Notes. Treatment is constructed as the interaction between a dummy variable equal to one in 2004 onward and zero before and a measure of H-1B dependency. In the continuous case, H-1B dependency is measured as the total number of LCA applications for a given multinational divided by that multinational's U.S. employment in 2001. In the binary case, it is a new variable equal to one if the firm was in the top group of H-1B dependency in 2001 and zero if the firm had zero LCA applications in 2001. Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

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

$$\Delta \ln(FA_emp_{ic,t-2001}) = \alpha_j + \alpha_c + \beta_1 \left(\frac{LCAapps}{US_emp} \right)_i^{2001} + \Delta \epsilon_{ic,t-2001} \quad (2)$$

where i indexes the firm, j indexes the industry, c indexes the country, and t is a postpolicy year. As before, FA_emp is foreign affiliate employment in country c by firm i , $LCAapps$ is the measure of demand (the number of LCA applications) by firm i in prepolicy year 2001, US_emp is a multinational's employment in the United States in prepolicy year 2001, and α_j and α_c capture industry (North American Industry Classification System (NAICS) four-digit level) and country time trends, respectively. The dependent variable is the logged differenced foreign affiliate employment of firm i in country c between year t and a prepolicy year (2001). I expect β_1 to be positive in postpolicy change years (2005–2014) and null in prepolicy change years (2002–2003). After the 2004 reduction in the H-1B cap, firms that were more dependent on H-1B visas should be more affected by the policy change and, therefore, more likely to expand their foreign affiliate activity. Before the 2004 cap change, any preexisting variation in demand for foreign-born skilled workers, as measured by H-1B dependency, should not be correlated with foreign affiliate employment growth. The main advantage of this specification is that I can control for industry time trends or other trends among sets of firms.

The results are shown in Table 4, in which each column represents a long difference between 2001 and a later year. The results provide evidence that there were no existing pretrends in the differences in foreign affiliate employment growth that correspond with the measure of H-1B dependency; β_1 is not statistically significant until 2005, after the policy change.¹⁷ The results also show that firms that were one percentage point more H-1B dependent than average saw a 3%–8% larger increase in foreign affiliate employment than average as a result of increased immigration restrictions resulting from the 2004 cap drop. Because the average foreign affiliate in 2001 had 1,151 employees, the estimated effect at the mean is, thus, an increase of 35–90 employees at every foreign affiliate.

As with the differences-in-differences specification, the approach requires that pretreatment trends in foreign affiliate employment were the same for the treatment and control groups (the parallel trends assumption). If one looks simply at how the level of foreign affiliate activity differed between our treatment and control groups, it is apparent that firms that filed more LCAs in 2001 were not identical to firms with fewer LCAs. For example, firms that filed large numbers of LCAs tend to do more

Table 4. Effect on Growth in Foreign Affiliate Employment for All Firms, 2001 Base Year

	(1) 2002	(2) 2003	(3) 2004	(4) 2005	(5) 2006	(6) 2007	(7) 2008	(8) 2009	(9) 2010	(10) 2011	(11) 2012	(12) 2013	(13) 2014
<i>H-1B dependency</i>	0.0098 (0.0103)	0.0053 (0.0122)	0.0113 (0.0153)	0.0287* (0.0172)	0.0375** (0.0185)	0.0492* (0.0260)	0.0466** (0.0203)	0.0443* (0.0235)	0.0329* (0.0200)	0.0602*** (0.0229)	0.0679*** (0.0259)	0.0724*** (0.0258)	0.0547 (0.0419)
Industry fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,619	5,277	4,657	4,418	4,343	4,122	3,798	4,160	3,687	3,449	3,237	3,130	3,277
R ²	0.052	0.062	0.085	0.085	0.091	0.134	0.143	0.147	0.162	0.181	0.196	0.208	0.198

Notes. Each column represents a long difference in logged foreign affiliate employment between 2001 and a later year. H-1B dependency is defined as the total number of LCA applications for a given multinational divided by that multinational's U.S. employment in 2001. The differences in observations in each year are driven by the fact that the dependent variable is logged; hence, any observations equal to zero are dropped. Thus, this regression only measures the effect on the intensive margin. Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

R&D abroad but have fewer employees abroad. However, the difference in levels is not in itself problematic for my specification; the threat to identification is if my measure of H-1B dependency were correlated with pretreatment changes in foreign affiliate employment.

To test for this possibility, I include the 1994–1999 pretrend in the baseline specification. These results are shown in column (2) of Table 5 with the baseline results shown in column (1) for the purpose of comparison. The coefficient of interest gets slightly smaller but remains positive and statistically significant, hence lessening any concerns about endogeneity. Note that I only report the results for the 2013–2001 long difference; they are robust to choosing any post-2004 year, but because of space constraints and constraints from the use of confidential data, I only report the 2013 results here.

Another concern might be that results are driven by systematic growth rate differences across firms of different size, internationalization, or innovativeness to begin with. If, for instance high-patenting firms naturally expand their activity abroad more quickly than nonpatenting firms, then my results could reflect that correlation rather than the effect of the policy change. I, therefore, test whether the coefficient of interest changes when including controls for the size or type of firm in terms of their 2001 employment, 2001 sales, 2001 R&D spending, 2001 total patenting, and number of foreign affiliates in 2001 quantiles. These controls are included as

group fixed effects. The identification in these regressions, once various controls are added, is now based on the comparison of trajectories within the same industry and the same category of employment, sales, or R&D size or patenting amount or degree of internationalization, depending on the control. Column (3) of Table 5 presents the results of a specification containing indicators for each firm's 2001 patenting quantile. The other results are in the online appendix and show that the main results remain qualitatively unaffected. The robustness of the results to a variety of firm characteristics is reassuring; results do not appear to be driven by systematic growth rate differences across firm characteristics other than H-1B dependency.

Of course, the relationship between growth in multinational foreign affiliate employment and the share of H-1B workers might be nonlinear as suggested by the skewness in applications. I use a nonparametric approach to examine this possibility: I divide multinationals into groups according to their H-1B dependence in 2001. I create seven categories of multinationals, in which the base category is all multinationals with zero LCA applications, and the remainder are divided into five quantiles with the top category divided into two groups. Again, I expect positive coefficients with especially large coefficients on high-dependency multinationals, and column (4) of Table 5 shows exactly that pattern. In particular, I find large, positive, and statistically significant

Table 5. Effect on Growth in Foreign Affiliate Employment for All Firms, 2001–2013

	(1) Baseline	(2) Controlling for pretrend	(3) Base year patent control	(4) Nonlinear specification
<i>H-1B dependency</i>	0.0724*** (0.0258)	0.0687** (0.0283)	0.0876*** (0.0294)	
<i>Change in ln(FA emp), 1994–1999</i>		–0.0822** (0.0326)		
<i>Second quantile of H-1B dependency</i>				0.0509 (0.108)
<i>Third quantile of H-1B dependency</i>				–0.00229 (0.118)
<i>Fourth quantile of H-1B dependency</i>				–0.0518 (0.0990)
<i>Fifth quantile of H-1B dependency</i>				–0.0792 (0.135)
<i>Sixth quantile of H-1B dependency</i>				–0.129 (0.154)
<i>Top quantile of H-1B dependency</i>				0.462** (0.209)
Industry fixed effects	Yes	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes	Yes
Patent group fixed effects			Yes	
Observations	3,130	1,913	3,130	3,130
R ²	0.208	0.247	0.211	0.209

Notes. The dependent variable in every column is the long difference in logged foreign affiliate employment between 2001 and 2013. H-1B dependency is defined as the total number of LCA applications for a given MNC divided by that MNC's U.S. employment in 2001. Results are robust to the choice of any post-2004 year, but I only show 2013 results here because of space constraints and because of the confidentiality of the data. Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

coefficients for the top bracket (with an LCA application—U.S. employment ratio of at least 0.0158 in 2001). In other brackets, there is no statistical significance. These estimates suggest that the positive effect of H-1B restrictions on foreign affiliate employment is being driven by the heaviest users of H-1B visas. This result may raise concerns that the effects are driven by just a handful of firms, but results are robust to dropping the top 50 greatest users of H-1B visas.

I implement a number of additional robustness checks. First, I use 2002 as the base year instead of 2001 to ensure that any oddities about the year 2001 are not driving the results. Second, instead of normalizing applications by U.S. employment as the H-1B dependency measure, I use the count of applications. A third robustness check addresses the concern that LCAs are an imperfect measure of demand for skilled H-1B visas. In particular, we know that some firms continuously file LCA applications and only utilize some of them. In addition, some of these applications are for continuing and transfer H-1B visa applications. A robustness check that measures H-1B dependency in 2001—our treatment measure—using H-1B petitions to USCIS in 2001 rather than LCA applications to DOL in 2001 can be found in the online appendix (Table A4-2, column (1)) and confirms that the results are not driven by the imperfectness of LCAs as a measure of demand. The online appendix (Table A4-2, column (2)) also includes a version of the baseline differences-in-differences that combines the extensive and intensive margins by using the inverse hyperbolic sine of foreign affiliate employment as the dependent variable rather than the natural log. One might be concerned that different countries experience different and nonlinear trends in foreign affiliate growth; India and China both experienced a sharp increase in FDI during this time period. Hence, I also test robustness to country-year interacted fixed effects (column (1) of Online Table A4-3). One might also be concerned that changes in foreign trade barriers are driving the results; if they occur at the same time as the immigration restrictions, they could be the omitted variable driving FDI. The results are robust to including a control for tariff rates (column (2) of Online Table A4-3). I test robustness to standard errors clustered at a variety of different levels shown in Online Table A4-3. Finally, I show robustness to using the pseudo-Poisson maximum likelihood estimator approach suggested by Santos Silva and Tenreiro (2006) in column (8) of Online Table A4-3. The results are robust to all of these tests.

The main threat to identification comes from any shocks correlated with both the timing of the H-1B policy and its effects across firms. In particular, the tech bubble in the late 1990s and early 2000s may have been correlated with increases in the cap. After the bubble burst, the H-1B visa cap was higher than average, and the economy experienced a downturn. To the extent that

the recession particularly affected H-1B dependent firms, the estimates could be biased. The direction, however, is unclear. They may have been more likely to increase foreign affiliate activity to escape the recession in the United States, which would lead to upward bias, or they may have been more likely to shrink their firms, which would lead to a downward bias. However, the robustness of the results to the inclusion of industry time trends in all regressions suggests that this is not problematic; any unobserved demand shocks for highly skilled workers would need to vary across firms within the same industry for there to be any bias.

To address remaining concerns about omitted variable bias that may not have been addressed by the robustness checks, I also implement a matching approach, also shown in the online appendix (Table A4-4). I cannot implement a standard propensity score matching difference-in-difference specification (Heckman et al. 1997) because the treatment in question is continuous rather than binary, so I instead take two variants on this approach. In the first, I follow the generalized propensity score method introduced by Hirano and Imbens (2005). In the second, I flexibly control for observables—with indicator variables for each quantile of the number of foreign affiliates, employees, sales, and R&D spending in 2001—following the classic Rosenbaum and Rubin (1983) proof that shows propensity score matching is equivalent to including flexible controls.

A separate concern surrounds the parallel trends assumption and any anticipation of the cap decline, namely, did firms behave differently leading up to the policy change in anticipation of soon experiencing immigration constraints, and did these behavioral differences correlate with my constructed measures of H-1B dependency? Figure 1 does not show any clear change in trends prior to 2004, and more rigorous tests, shown in Tables 3–5 and in the online appendix, further support the case that firms did not change behavior in advance. These concerns are further ameliorated by the results of a second strategy, which does not suffer from the same sources of potential bias and yet produces consistent results.

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

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

causal interpretation of the effect of constrained foreign-born skilled worker supply on the offshoring of jobs.

The identification in this strategy exploits random variation in the allocation of H-1B workers across U.S. multinational firms resulting from the H-1B lotteries of 2007 and 2008. In both of those years, the number of cap-dependent H-1B visa petitions submitted within the first month of the filing period far exceeded the annual limit of available permits in those years as shown in Table 1. In those years, all petitions received by the final receipt date (April 3 and 8, respectively) were put through a computer-generated random selection process that selected which petitions would be processed. This produced a random shock to the supply of H-1B workers; some firms were successful in the lottery, whereas others were not. The data support the random nature of the lottery; the mean win rates for the multinational firms in the sample are 0.46 and 0.53 in the 2007 and 2008 lotteries, respectively, with standard deviations of 0.37 and 0.38. My approach exploits this random H-1B variation.

The dependent variable in this approach is the same as the differences-in-differences approach: the change in foreign affiliate employment. However, instead of regressing the change in foreign affiliate employment on a measure of the firm's pre-policy change H-1B dependency, I regress it on a measure of excess demand for foreign labor that is driven by exogenous supply shocks. Following Peri et al. (2015), I calculate excess demand as the difference between the firm-level demand for new H-1B workers (LCA applications that were filed early) and the firm-level capped supply of H-1B workers (the lottery allocation of permits). I scale this absolute measure of excess demand by the firm's U.S.-based employment in 2007. There are two mechanisms at work in this approach, both generating variation in normalized excess demand across firms. The first is the same mechanism at work in the differences-in-differences approach: firms that are more H-1B-dependent feel the effects of H-1B

supply constraints more acutely than those that do not hire many H-1B workers. The second is the unexpected supply shock coming from the lottery.

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

$$\Delta \ln(\text{FA_emp}_{ic,t-2005}) = \alpha_j + \alpha_c + \beta_1 \left(\frac{\text{ExcessDemand}_i^{2007+2008}}{\text{US_emp}_i^{2007}} \right) + \beta_2 \text{LCA}_i^{07+08} + \Delta \epsilon_{ic,t-2005} \quad (3)$$

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

Table 6 shows that β_1 is indeed significantly positive. The columns display foreign affiliate employment growth one, two, three, four, and five years after the lottery and illustrate a persistent positive effect. The results show that a random negative shock to H-1B supply equal

Table 6. Effect on Growth in Foreign Affiliate Employment for All Firms, 2005 Base Year

	(1) 2010	(2) 2011	(3) 2012	(4) 2013	(5) 2014
Adjusted excess demand in 2007–2008	0.0484 (0.0334)	0.137*** (0.0496)	0.121** (0.0499)	0.161** (0.0681)	0.142* (0.0836)
Number of LCAs in 2007–2008	–0.0000989 (0.000113)	–0.000342** (0.000139)	–0.000205 (0.000157)	–0.000358* (0.000185)	–0.000295 (0.000216)
Industry fixed effects	Yes	Yes	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	4,308	4,029	3,776	3,685	3,834
R ²	0.140	0.162	0.172	0.190	0.179

Notes. Each column represents a long difference in logged foreign affiliate employment between 2005 and a postlottery year (2010–2014). Adjusted excess demand is defined as the total number of cap-subject LCA applications minus cap-subject H-1B petitions issued for a given multinational in 2007 and 2008, divided by that multinational's U.S. employment in 2007. This number is multiplied by 100 for purposes of interpretation. Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

to one percentage point of initial employment caused an increase in the foreign affiliate growth rate of between 12% and 16%.

Translating these coefficients into the number of jobs offshored, I find that about 0.42 foreign affiliate jobs were created for every unfilled H-1B position. Additional details of this calculation can be found in Online Appendix A6. This 0.42 estimate of substitution could be considered a lower bound for several reasons. First, it relies on calculations on the intensive margin and does not consider increased foreign affiliate employment on the extensive margin. Second, there are likely at least some firms that did not submit their petition(s) in time to be considered for the lottery, which means that their behavior is not captured in the analysis. Third, this estimate only captures how existing firms modified their strategy; new firms born into the visa-restrictive world may have incorporated offshoring into their strategy from the beginning. Finally, the estimate does not include foreign multinational company behavior, and one expects that foreign MNCs would be even more likely to substitute foreign-born U.S.-based skilled workers for foreign-born skilled workers at headquarters or other foreign affiliates. In short, it is possible that the true effect is larger than 0.42, but it is still unlikely to be close to one.

If the true effect is indeed less than one, it implies a constraint on firm growth. A simple example illustrates; imagine an MNC applies for 10 H-1B visas with the goal of hiring 10 skilled immigrants domestically, but only five of those applications wins the lottery. The 0.42 estimate of substitution then suggests that the MNC will expand foreign affiliate employment by two, whereas the null results on U.S. employment (Online Appendix A3) suggests zero change in domestic native employment. Thus, the MNC grows by only seven employees rather than the original planned 10 employees. In short, this back-of-the-envelope calculation suggests that the MNC does not grow employment as much as it might have preferred although it does not provide a precise measure of the change in firm growth. This is perhaps not surprising; if a global firm's first best strategy were to hire abroad rather than applying for H-1B visas, they would not appear in the sample. Much more can and should be done in future research to understand the impacts of this and similar constraining policies on firm growth.

One might be concerned that some firms anticipated the lottery and increased their submitted applications to improve their chances of winning the lottery. The two lottery years in the empirical analysis were selected precisely to address this concern; these were the first two years in which a lottery was held to distribute all H-1B visas. To further address this concern, Online Table A4-5 duplicates the same analysis but with only the 2007 lottery and finds the same qualitative result: that a random negative shock to H-1B supply caused an increase in the

foreign affiliate growth rate. One might also be concerned that the error term is correlated with a firm's U.S.-based employment and the number of LCA applications. Online Table A4-6 duplicates the analysis but without the LCA control and without normalizing the independent variable and also finds the same qualitative result.

Overall, the positive effect from the lottery approach is consistent with the results from the differences-in-differences approach¹⁸ and provides further, reassuring, support for the hypothesis that restrictions on high-skilled immigration cause increased offshoring of high-skilled jobs.

5.2. Part Two: The Extensive Margin

The results, thus, far show the effect of immigration restrictions on foreign affiliate employment at the intensive margin; because foreign affiliate employment was logged in all regressions to this point, regressions so far have only captured the change in employment within countries in which a firm is active in both 2001 and the later postpolicy period. It does not capture the effect of any foreign affiliates that were opened in a new country after the policy change in response to the policy. In short, it captures expansion within a country but not expansion into new countries. Both effects are of interest here, and in fact, any impact on the extensive margin represents a much stronger strategic response by firms—and a larger long-term impact—than an impact on the intensive margin.

To measure the extensive margin effect, I use the long-differenced regression specification, but I change the dependent variable to a binary variable equal to one if the firm has a foreign affiliate in a given country by the postpolicy year and zero otherwise. I use a linear probability model so that I can include time trends, but the results are robust to a logit model. The results are shown in Table 7 and illustrate a statistically significant and consistent response on the extensive margin in all years after—not before—the H-1B cap change in 2004.

An important caveat to the results must be made here: as noted in the data section, each firm may report on a consolidated basis for multiple affiliates in the same country under certain conditions,¹⁹ so I aggregate all foreign affiliate activity up to the host country level for a given parent firm for a given year. This implies that effects on the extensive margin are only observed when a firm enters a new country; the extensive margin results do not, for example, reflect a firm opening a foreign affiliate in Vancouver if it already has a foreign affiliate in Toronto. The paper is, thus, estimating the impact of immigration restrictions on the country-level extensive margin of foreign direct investment. The coefficients as a result are somewhat small, but in the next section, I show that they become larger when restricting for R&D-performing foreign affiliates.

Table 7. Effect of H-1B Policy Restriction on the Likelihood of a U.S. MNC Entering a New Country

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014
<i>H-1B dependency</i>	0.000132 (0.000175)	0.000330 (0.000222)	0.00130*** (0.000428)	0.00175*** (0.000404)	0.00225*** (0.000426)	0.00340*** (0.000445)	0.00326*** (0.000455)	0.00443*** (0.000512)	0.00314*** (0.000478)	0.00307*** (0.000478)	0.00336*** (0.000480)	0.00327*** (0.000484)	0.00347*** (0.000529)
<i>Constant</i>	0.00217*** (0.000145)	0.00350*** (0.000185)	0.0132*** (0.000357)	0.0115*** (0.000336)	0.0129*** (0.000355)	0.0139*** (0.000371)	0.0146*** (0.000379)	0.0185*** (0.000426)	0.0162*** (0.000397)	0.0162*** (0.000398)	0.0163*** (0.000400)	0.0166*** (0.000403)	0.0200*** (0.000440)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	108,623	108,623	108,623	108,623	108,623	108,623	108,623	108,623	108,623	108,623	108,623	108,623	108,623
R ²	0.001	0.002	0.005	0.004	0.005	0.006	0.006	0.008	0.008	0.008	0.008	0.008	0.010

Notes. The dependent variable here is a binary variable equal to one if the firm had opened operations in a new country by a postpolicy year (that they did not have in 2001) and equal to zero otherwise. The model is a linear probability model. H-1B dependency is defined as the total number of LCA applications for a given multinational divided by that multinational's U.S. employment in 2001. Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

5.3. Part Three: Were R&D-Intensive Jobs and Foreign Affiliates Disproportionately Impacted?

The results thus far are estimated using aggregate employment, and yet one would expect skilled (here, proxied as R&D) employment to be disproportionately affected because the restrictions target skilled migrants. It is less clear whether non-R&D employment is affected; existing research documents complementarities between R&D activities and production (Naghavi and Ottaviano 2009, Pisano and Shih 2012). I, therefore, estimate the baseline specification from Equation (1), in which the dependent variable is one of two variables: foreign affiliate R&D employment or foreign affiliate non-R&D employment, and foreign affiliate non-R&D employment is the difference between total foreign affiliate employment and foreign affiliate R&D employment. An important caveat is that aggregate foreign affiliate employment is not broken into more disaggregated pieces in most years, so foreign affiliate R&D employment is only reported in benchmark years, which are every five years. In this sample, the benchmark surveys are 1994, 1999, 2004, 2009, and 2014. Therefore, these specifications are also restricted to that sample.

Table 8 shows the results, in which the dependent variable is logged in columns (1) and (2) and it is transformed using the inverse hyperbolic sine transformation in columns (3) and (4). Columns (1) and (2) show the effect of skilled immigration restrictions on foreign affiliate R&D and non-R&D employment at the intensive margin because the dependent variable is logged; in other words, they only capture the change in R&D and non-R&D employment in countries where there were R&D and non-R&D employees from the beginning. The sample in column (2) is restricted to be the same sample as column (1) to make coefficients comparable. The results show that, on the intensive margin, the substitution is entirely driven by R&D employment.

Columns (3) and (4), by using the inverse hyperbolic sine transformation, show the effect on both the intensive and extensive margin, capturing both within-country employment expansion and entry into new countries. Interestingly, there is evidence of complementarities between R&D and non-R&D employment in these results; both foreign R&D and non-R&D employment increase. The results suggest that when high-skilled immigration restrictions are implemented, any new foreign affiliates in new countries that open in response to said restrictions contain both types of employees. Complementarities between production and R&D may magnify the effect of the immigration restrictions and lead to the offshoring of both skilled and unskilled labor and their respective activities.

Next, Table 9 drills slightly deeper into the extensive margin results using 2013 with an eye toward heterogeneity along R&D-intensity. Column (1) shows that that

Table 8. Effect on Foreign Affiliate Employment for All Firms by Type of Employment, Benchmark Years in 1994–2014

	(1) ln(R&D emp)	(2) ln(non-R&D emp)	(3) ihs(R&D emp)	(4) ihs(non-R&D emp)
<i>Continuous treatment</i>	0.0536*** (0.0149)	−0.000521 (0.0199)	0.0953*** (0.0324)	0.145** (0.0624)
<i>Constant</i>	2.756*** (0.00629)	6.280*** (0.00840)	0.225*** (0.00532)	2.503*** (0.0102)
Country fixed effects	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	6,507	6,507	97,570	97,570
R ²	0.405	0.547	0.193	0.271

Notes. Continuous treatment is constructed as the interaction between a dummy variable equal to one in 2004 onward and zero before and a measure of H-1B dependency in 2001. The sample is benchmark years only: 1994, 1999, 2004, 2009, and 2014. The coefficient on continuous treatment is statistically different between columns (1) and (2) ($p = 0.002$) but not between columns (3) and (4) ($p = 0.34$). Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

the 2004 cap decline made a one percentage point more H-1B dependent firm 0.3% more likely to open a foreign affiliate in a new country than the average firm by 2013. Column (2) examines the likelihood of a multinational initiating R&D activity abroad. Here, the coefficient is slightly larger (although not statistically significantly larger) as expected because one would expect skilled immigration restrictions to affect the opening of an R&D laboratory but not necessarily to affect the opening of a new manufacturing facility.

The coefficients in columns (1) and (2) are fairly small even if they are statistically significant. This is in large part because of the measurement issue described earlier as well as the preponderance of zeros in the data; the average U.S. multinational firm in the BEA data are active in about two countries, whereas the data set

includes 48 countries.²⁰ To counteract this noise, column (3) shows the same regression but for a subset of the 25 countries with the most activity in the data and with the binary version of the independent variable. As expected, the coefficients become substantially larger. The results show that the 2004 cap decline made highly H-1B dependent firms 5% more likely to open a foreign affiliate in a new country and 10% more likely to start conducting R&D in a new country than a non-H-1B-dependent firm.

5.4. Part Four: Geographic Location Choice

In this section, I begin by examining host country heterogeneity and then move into firm heterogeneity. The largest countries of origin for H-1B visa holders are China and India (85% of H-1B petition filings), whereas many of the prominent examples of companies opening

Table 9. Effect on the Extensive Margin, Exploring Heterogeneity

	(1) DV: New FA by 2013 Sample of 48 countries	(2) DV: New R&D by 2013 Sample of 48 countries	(3) DV: New FA by 2013 Subsample of 25 countries	(4) DV: New R&D by 2013 Subsample of 25 countries
<i>Continuous form of H-1B dependency</i>	0.00327*** (0.000484)	0.00421*** (0.000352)		
<i>Binary form of H-1B dependency</i>			0.0532*** (0.0121)	0.0967*** (0.00915)
<i>Constant</i>	0.0166*** (0.000403)	0.00826*** (0.000292)	0.0762*** (0.00305)	0.0366*** (0.00230)
Industry fixed effects	Yes	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes	Yes
Observations	108,623	108,623	8,305	8,305
R ²	0.008	0.010	0.042	0.034

Notes. The dependent variable in columns (1) and (3) is a binary variable equal to one if the firm had opened new operations in a given country by 2013 (that they did not have in 2001) and equal to zero otherwise. In columns (2) and (4), it is a binary variable equal to one if the firm has initiated R&D activity abroad in a given country by 2013 (that they did not have in 2001) and equal to zero otherwise. I use a linear probability model, but the results are robust to a logit model. The continuous form of H-1B dependency is defined as the total number of LCA applications for a given multinational divided by that multinational's U.S. employment in 2001. The binary version of H-1B dependency compares the top quantile to the bottom quantile. Results are robust to the choice of any post-2004 year, but I only show 2013 results here because of space constraints and because of the confidentiality of the data. The coefficient on H-1B dependency in columns (1) and (2) are not statistically different ($p = 0.51$), but they are statistically different between columns (3) and (4) ($p = 0.04$). Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

foreign affiliates abroad in response to H-1B restrictions are concentrated in Canada.²¹ Canada is a special case for U.S. firms; Canada is geographically close to the United States, and in addition, it has much less restrictive high-skilled immigration policies than the United States. These characteristics mean that the fixed costs of offshoring are relatively low. These facts suggest that the expansion of foreign affiliate activity could operate through two channels: (1) a direct channel, whereby multinationals expand foreign affiliate activity in countries where the human capital they need is located (e.g., India or China), or (2) an indirect channel, whereby multinationals expand foreign affiliate activity in countries such as Canada, where it is easy to open foreign affiliates housing immigrants from other countries. These facts also suggest that, whereas U.S. firms are likely using offshoring to hire Chinese and Indian talent in China and India, they are—for the most part—not using offshoring to hire Canadian talent in Canada. Instead, they may often be hiring the same skilled (non-Canadian) immigrants in Canada that they had originally sought to hire in the United States under the H-1B system.

To test the relative effects, I construct two samples—one of foreign affiliates in China and India (the “raw human capital” countries) and one of foreign affiliates in Canada—and run the baseline differences-in-differences regression model on these different samples, shown in columns (2) and (3) of Table 10. This approach is an imperfect way of measuring direct and indirect flows, but it does provide some sense of the possible operational channels. The effects are statistically significant, positive, and larger than the base sample for both subsets, suggesting that a lot of the expansion was concentrated in these three countries. However, a fourth regression—shown in column (4)—that excludes Canada, India, and China shows that these three countries are not the only ones

impacted. The coefficient is smaller when the three countries are excluded, but it is still economically and statistically significant. Hence, whereas expansion in response to immigration restrictions was concentrated in Canada, China, and India, they were not the only countries affected.

These results suggest that host country immigration policy regimes may be an important factor influencing MNC geographic location choice. Although this analysis does not explicitly examine host country immigration policies as a factor influencing geographic location choice, the concentration of the offshoring to Canada in response to restrictive skilled immigration policies is suggestive of its role.

5.5. Part Five: Firm Heterogeneity

Next, I turn to firm heterogeneity: what firm capabilities might serve as a moderator of the impact of skilled immigration restrictions on offshoring. As described in Section 3, one might expect firm capabilities in the form of international experience to impact the decision of whether to offshore workers in response to skilled immigration restrictions. To operationalize a firm’s international presence, for each MNC, I compute the number of countries in which they are active in 2001 prior to the policy change. I then interact this count with the treatment variable. The specification utilizes the inverse hyperbolic sine transformation of foreign affiliate employment as the dependent variable in order to capture both intensive and extensive margin expansion. Column (5) of Table 10 displays the results and shows that firms with activity in more countries are more likely to offshore when a restrictive immigration policy is in place than firms with activity in fewer countries. In other words, the more global a firm is when the policy hits (if it is H-1B dependent), the more likely it is to respond by offshoring.

Table 10. Exploring Country- and Firm-Level Heterogeneity in Growth in Foreign Affiliate Employment for All Firms

	(1)	(2)	(3)	(4)	(5)
	Baseline	Canada	China and India	Excluding Canada, China, and India	<i>ihs(FA emp)</i>
<i>Continuous treatment</i>	0.0615*** (0.0176)	0.105** (0.0424)	0.223*** (0.0319)	0.0483*** (0.0173)	−0.0209 (0.0554)
<i>Continuous treatment</i> × <i>Number of countries in 2001</i>					0.00540** (0.00269)
Country fixed effects	Yes	No	No	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	137,459	12,220	5,704	119,207	409,794
R ²	0.429	0.840	0.676	0.420	0.285

Notes. The first four columns show the same regression as column (1) of Table 3 but with different country subsamples, whereas before, the dependent variable is logged foreign affiliate employment and continuous treatment is constructed by multiplying H-1B dependency in 2001 by a post-2004 indicator variable. The coefficient on continuous treatment is weakly statistically different between columns (2) and (4) ($p = 0.14$) and strongly statistically different between columns (3) and (4) ($p = 0.00$). The fifth column has a change in the dependent variable (inverse hyperbolic sine of foreign affiliate employment instead of logged) and adds an interaction with the number of countries in which the MNC is active in 2001. Firm-clustered standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Indeed, a recalculation of the substitution between hiring H-1B workers and hiring workers abroad for just the most internationalized firms (those present in at least 15 countries—the 90th percentile) indicates not only that these firms are more likely to offshore, but also that their growth is much less constrained by the policy. Their substitution ratio is considerably larger than for the entire sample; for the most internationalized firms, an H-1B rejection increases foreign affiliate employment by 0.93.²² This suggests that, whereas both less and more global firms are at least somewhat constrained in their employment growth, less global firms are especially constrained. In addition to being less likely to offshore, they are also less likely to have the types of complementary resources and operational flexibility discussed in Section 3 for responding to immigration restrictions, so they are in a particularly difficult position. In contrast, the most global firms are the best positioned to respond to these restrictions by offshoring; they have already invested the fixed costs in creating many foreign affiliates abroad, and they have developed the necessary managerial capabilities and strategies for managing knowledge workers abroad. Indeed, they appear to be able to almost perfectly substitute foreign workers for H-1B workers.

6. Discussion and Conclusion

This study finds causal evidence that U.S. MNCs respond to restrictive skilled immigration policies by expanding foreign affiliate employment at existing foreign affiliates and by adding new foreign affiliates. It also finds some evidence that this solution is imperfect from the average firm's perspective: for every visa rejection, firms hire 0.42 employees abroad. However, firm capabilities in the form of previous internationalization are shown to be an important moderator of the impact of skilled immigration on offshoring: for every visa rejection, the most internationalized firms hire 0.93 employees abroad.

The study also finds that expansion of foreign affiliate employment has been significantly—but not exclusively—concentrated in three countries: China, India, and Canada. These location choices suggest that the expansion of foreign affiliate activity could operate through two channels: (1) a direct channel, in which raw human capital is the motive (e.g., India or China), or (2) an indirect channel, in which the ability to hire immigrants from elsewhere under a more relaxed regime is the motive (e.g. Canada).

The results illustrate the importance of skilled human capital to firms, the global nature of the relevant labor pool, and the lengths to which firms will go to obtain the human capital they need. At the same time, they show that, whereas multinational firms can leverage their cross-border presence to obtain such human capital, they are still somewhat constrained by country-level

policies, and these policies can be costly. Indeed, firms based in countries without such restrictions may have a competitive advantage over those based in countries with immigration restrictions.

6.1. Welfare/Policy Implications

The results from this paper may help to explain the existing mixed empirical evidence on the impact of skilled immigration on the labor market outcomes of the native-born. Whereas some papers find negligible effects on the wages and employment of natives (Mithas and Lucas 2010, Ottaviano et al. 2013, Kerr et al. 2015b, Peri et al. 2015, Mayda et al. 2018), others find some evidence of substitution or downward pressure on wages (Borjas 2005, Bound et al. 2017, Turner 2022, Doran et al. 2022). The data used in this paper permit an examination of a third avenue available to firms: hiring the skilled employees they need at their foreign affiliates instead of at home. Whereas the estimation of the impact on native-born outcomes is not the focus of this paper, this paper's results do suggest that the existing immigration literature's focus on the substitution between native-born and foreign workers within the country may be incomplete. Indeed, policymakers should consider that restrictive immigration policies implemented to protect native jobs are likely to have the unintended consequence of offshoring. More broadly, this paper provides new evidence on how firms, particularly multinational firms, respond to policy changes in ways that may run completely counter to policymakers' intentions. The results also underline the fact that firms, managers, and policymakers live in a complex and globalized world that limits the policy autonomy of nations. Indeed, policymakers—and researchers studying inference within domestic borders—should recognize that responses to constraining domestic policies may have impacts that extend beyond their own borders.

The finding that skilled foreign-born workers are hired at foreign affiliates rather than in the United States also may have important implications for the innovative capacity of the United States. If skilled foreign-born workers are at a U.S. firm's foreign affiliate instead of in the United States, the innovative spillovers that they generate will go to another country instead. Furthermore, the finding that immigrants often are not equally innovative outside the United States (Kahn and Macgarvie 2016) has even wider welfare implications. In short, restrictive H-1B policies could not only be exporting more jobs and businesses to countries such as Canada, but they also could be causing the United States's innovative capacity to fall behind. Indeed, restrictive immigration policies may be doing more to enhance Canadian competitiveness than American competitiveness. Establishing the impact on the innovative capacity of the United States—and countries such as Canada who have

likely benefited from these policies—is a worthy topic for future research.

6.2. Management Implications and Future Work

The results have a number of potential implications for managers and for management research, some of which reinforce observations from previous work and some of which raise questions for future research. First, the results in this paper reinforce existing work on the synergies between R&D and production. Although the offshoring described in this paper is concentrated in R&D foreign affiliate employment, it is not exclusive to R&D employment, suggesting that complementarities between production and R&D may have magnified the effect of the H-1B visa restrictions, pushing not just skilled, but also unskilled employment abroad. These results resonate with the literature studying the innovative implications of the geographic separation of R&D and production (Cohen and Zysman 1987, Fuchs and Kirchain 2010, Pisano and Shih 2012, Fort et al. 2020, Branstetter et al. 2021), some—but not all—of which argues that the synergistic relationship between the manufacturing and R&D functions of the firm means that geographic separation could undermine the innovative capacity of the firm. From a firm capabilities and organizational design perspective (e.g., Teece 1997, Eisenhardt and Martin 2000, DeSanctis et al. 2002, Winter 2003, Alexy et al. 2021), this synergy has profound implications for how multinational firms structure their activities in different geographies; firms may not be able to separate knowledge and production capabilities fully to yield value from their R&D spending.

Second, the results underline the arbitraging of international differences as a strategic choice for MNCs under substantial constraints (Ghemawat 2003). Indeed, the fact that Canada appears to be an important destination for offshoring in this study suggests that host country immigration policy regimes may be an important factor influencing MNC geographic location choice. In the face of immigration restrictions at home, MNCs can arbitrage international differences in immigration policy to access human capital that domestic firms cannot easily access. More generally, skilled immigration restrictions serve as an important, counterintuitive, and previously overlooked push factor for internationalizing knowledge activity. Skilled immigration restrictions create constraints on the ability to hire from a global talent pool, but multinational firms can—at least partially—circumvent such restrictions and mitigate country-level risk by hiring foreign talent at their foreign affiliates instead. The arbitrage described in this study is likely not limited to immigration; multinational firms are likely to be able to respond to artificial institutional constraints on many types of resources by leveraging their cross-border presence.

Third, this paper represents a first step toward better understanding the ways in which firms can strategically respond to government-imposed constraints on scarce resources and the extent to which they are able to avoid costs entirely or continue to be at least partially constrained. This paper documents one area—restrictions on foreign human capital in the form of visa restrictions—but the boundary conditions are as yet unclear. The result that the average firm hires 0.42 employees abroad for every visa rejection suggests that there may be a portfolio of responses available to firms beyond the response commonly described in the immigration literature (i.e., substitution for domestic workers) and the response highlighted in this paper (offshoring). Future work should expand the set of responses, the trade-offs involved in each, and which firms are best positioned to act on each.

In addition to examining the range of responses, future work should push further to understand better the costs of circumventing such constraints and how those costs vary across firms and responses. The business decision of where to place knowledge activities and workers is a deliberate and strategic choice, weighing the centripetal forces compelling R&D to stay near headquarters, such as economies of scale and reduced transfer and control costs, against the centrifugal forces that pull R&D to specific host countries, such as cost considerations, market access, and location-specific talents or knowledge (Hirschey and Caves 1981, Pearce 1999). If, under an open immigration system, the firm decides that the optimal location for employing key knowledge workers is at headquarters, then policies that block the firm's ability to do so come at a cost. The partial (0.42), rather than 1:1, substitution identified in this paper reflects that cost. However, the paper also finds that some firms are able to come much closer to fully substituting foreign workers for lost H-1B visas; for every visa rejection, the most internationalized MNCs hire 0.93 employees abroad. This result speaks to the organizational flexibility that multinationality brings (Kogut and Kulatilaka 1994) and suggests that managers of very global firms can respond to immigration restrictions by offshoring without much cost to their growth, whereas managers of less global firms may struggle to fully substitute foreign affiliate workers for skilled immigrants at home. Future research could examine the implications for the concentration of skilled labor; it appears more international firms are better able to respond to these restrictions with skilled labor in other parts of the world, putting them at an advantage relative to smaller, less internationalized firms.

Whereas, in this paper, I speculate that the difference in responses is due to firm capabilities related to internationalization, there are other potential firm-specific capabilities that might also serve as moderators. For example,

prior research shows the importance of within-firm diasporas in facilitating FDI (e.g., Foley and Kerr 2013); perhaps MNCs with greater diasporas are better able to efficiently offshore in response to immigration restrictions. More generally, much more can and should be done in future research to understand the ways in which firms can strategically respond to government-imposed constraints on scarce resources, the costs they incur in responding to such constraints, and the extent to which they are able to avoid costs entirely or continue to be at least partially constrained.

Finally, I see an exciting opportunity for researchers to examine how policies affecting the mobility of human capital shape the global geography of innovation. My findings suggest that host country immigration policy regimes may be an important factor influencing MNC geographic location choice. Although this analysis does not explicitly examine host country immigration policies as a factor influencing geographic location choice, the concentration of the offshoring to countries such as Canada with more open immigration policies in response to restrictive skilled immigration policies is suggestive of its role. Future work could more deeply analyze to what degree host country immigration policy regimes influence MNC geographic location choice for skilled and unskilled activities. Indeed, recent work by Bahar et al. (2022) starts down this path by investigating how host country immigration policy regimes affect MNC patenting through inventor mobility.

Whereas this paper does not explicitly examine the impact of immigration restrictions on the location or direction of innovation, a large literature links the geographic location of skilled immigrants to innovative spillovers (e.g., Moser et al. 2014, Hunt et al. 2017), suggesting that employing skilled foreign workers elsewhere is likely to correspond with doing innovation elsewhere. In addition, because innovative spillovers are geographically localized (Marshall 1890, Jaffe 1986, Jaffe et al. 1993), such impact is likely to extend beyond the innovation done by the MNC, potentially shaping host country innovative ecosystems. The degree to which this occurs is likely to correspond—at least partially—with how exactly MNCs rearrange their global innovative activities. Future work should examine the implications of this paper's findings for the organization of innovative work within MNCs as well as for how resultant spillovers affect the host country.

Acknowledgments

The author is grateful to Alfonso Gambardella and the editorial team at *Management Science*. The author thanks Lee Branstetter, Brian Kovak, J. Bradford Jensen, Giovanni Peri, Ina Ganguli, Elena Kulchina, Erica Fuchs, Katherine Klein, Exequiel Hernandez, Vit Henisz, Claudine Gartenberg, Seth Carnahan, Julia Lane, and Miguel Garza Casado as well as seminar participants at the National Bureau of Economic

Research Summer Institute, Brigham Young University Winter Conference, American Economics Association Annual Meeting, National Bureau of Economic Research Productivity Seminar, University of California Los Angeles, IESE Business School, University of Southern California, Bocconi, Wharton, Duke, Massachusetts Institute of Technology, Federal Reserve Board, Georgetown University, George Washington University, Notre Dame, Georgia Tech, and Harvard Business School for their helpful comments. The author thanks Bill Zeile, Jim Fetzer, and Ray Mataloni for discussions on the Bureau of Economic Analysis (BEA) data. All errors and omissions remain the author's own responsibility. The views expressed in this paper are those of the author and do not necessarily represent the views of the U.S. Bureau of Economic Analysis or the U.S. Department of Commerce. BEA has reviewed this paper for unauthorized disclosure of confidential information and has approved the disclosure avoidance practices applied to this release. The statistical analysis of firm-level data were conducted at the Bureau of Economic Analysis, U.S. Department of Commerce, under arrangements that maintain legal confidentiality requirements. (BEA approval numbers: FY22-E0006-R0001, FY22-E0006-R0002, BEA-FY21-E0006-R0002, BEA-FY21-E0006-R0001, BEA-FY22-E0006-0003, BEA-FY22-E0006-R0004, and additional unnumbered approvals).

Endnotes

¹ Restrictive immigration policies have become increasingly common in recent years around the world with two notable recent examples being the Trump Administration's "Buy American, Hire American" policies and Brexit. Leaders such as Turkey's Recep Tayyip Erdogan, Hungary's Viktor Orbán, French National Front leader Marine Le Pen, and Brazil's Jair Bolsonaro also rose to prominence using anti-immigration rhetoric.

² This number is based on a ratio of R&D performed by U.S. parents divided by R&D performed by U.S. parents and their majority-owned affiliates in 2018, calculated using BEA USDIA public data: <https://www.bea.gov/international/di1usdop>.

³ LCAs are the first step toward H-1B visas for skilled foreign-born workers in the United States. They are described in more detail later in the paper.

⁴ For one example, see "Microsoft opens Canada center in response to US immigration problems," <http://workpermit.com/news/microsoft-opens-canada-center-response-us-immigration-problems-20070710>.

⁵ Such mechanisms might include—but are not limited to—modularity (Zhao 2006), personnel mobility (e.g., Almeida and Kogut 1999, Rosenkopf and Almeida 2003), alliances (e.g., Zaheer and Hernandez 2011), manufacturing integration (Berry 2014), immigrant employees or managers (e.g., Foley and Kerr 2013, Hernandez and Kulchina 2020), and others.

⁶ I provide some discussion of other high-skilled visa alternatives in the online appendix. Yeaple (2019) provides a comprehensive discussion of the differences between L-1 and H-1B visas.

⁷ According to USCIS, "to qualify as a specialty occupation, the position must meet one of the following requirements: (1) a bachelor's or higher degree or its equivalent is normally the minimum entry requirement for the position; (2) the degree requirement is common to the industry in parallel positions among similar organizations or, in the alternative, the position is so complex or unique that it can be performed only by an individual with a degree; (3) the employer normally requires a degree or its equivalent for the

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

⁸ There is no cap for the following categories: (1) those for nonprofit firms, universities, and research labs; (2) those that are an extension of an existing H-1B visa; (3) those that have an existing H-1B visa and are changing jobs during the period of the existing visa; and (4) citizens of countries with whom the United States has a relevant free-trade agreement.

⁹ Fiscal years 1997 and 1998 were the lone instances when the cap was reached.

¹⁰ This is by the International Investment and Trade in Services Survey Act. The act specifies that the survey data may only be used for statistical and analytical purposes.

¹¹ The most extensive data are collected in benchmark years: 1994, 1999, 2004, 2009, and 2014. The reporting requirement threshold varies by year, size of the affiliate, and the parent's ownership stake. The BEA estimates values of some variables of some affiliates in nonbenchmark years in order to estimate a consistent universe across years. I only use the reported data in this paper.

¹² These conditions are that the affiliates operate in the same country and same industry classification or are integral parts of the same business operation.

¹³ There are four main labor conditions that they are required to meet: (1) recipients of the visa must receive the same or better wages and benefits as other similar company employees and as similar employees in the geographic area; (2) working conditions must be similar for all employees; (3) there must not be a "strike, lockout, or work stoppage" at the employment location when the LCA is signed and submitted; and (4) any employee bargaining representatives must be notified of every application submitted.

¹⁴ Much has been made anecdotally of Canada as a destination for firms struggling with immigration constraints in the United States. See, for example, the Envoy 2019 Immigration Trends Report, in which 38% of surveyed firms were thinking about expanding to Canada because their immigration policy is more favorable, and 21% already had at least one office there.

¹⁵ The apparent drop in the growth of average foreign affiliate employment among the top quintile in 2014 is due at least in part to the improved coverage in the 2014 USDIA Benchmark Survey, which increased the number of firms for that survey. More detail can be found at <https://www.bea.gov/sites/default/files/2018-04/MNE%2014-R%20Improvements%20in%20Coverage.pdf>.

¹⁶ See the Labor Condition Application for Nonimmigrant Workers ETA Form 9045CP – General Instructions for the 9035 and 9035E, U.S. Department of Labor: https://icert.doleta.gov/library/ETA_Form_9035CP_2009_Revised_03.18.09.pdf.

¹⁷ The number of observations is different in each column because effects are measured at the intensive margin; the foreign affiliate must have existed in both 2001 and the later year to appear in the sample. These results are robust to dropping firms that went out of business entirely during this time frame as well as to using a balanced panel. Table 3 is also run on a balanced panel.

¹⁸ The coefficients are not directly comparable. The two approaches have different samples (the first is much bigger because it includes multinationals that have never applied for an H-1B, whereas the second only includes the subset of multinationals applying for LCAs in one of those two years). Furthermore, the key regressor is measured differently.

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

²⁰ These 48 countries account for more than 99.5% of foreign R&D and 95% of foreign sales and employment.

²¹ See, for example, <http://www.talenteconomy.io/2017/06/19/tighter-immigration-policy-pushes-firms-open-foreign-satellite-offices/>.

²² This is calculated using the same method described in Online Appendix A6.

References

- Alcácer J (2006) Location choices across the value chain: How activity and capability influence collocation. *Management Sci.* 52(10):1457–1471.
- Alcácer J, Chung W (2007) Location strategies and knowledge spillovers. *Management Sci.* 53(5):760–776.
- Alexy O, et al. (2021) Adaptation or persistence? Emergence and revision of organization designs in new ventures. *Organ. Sci.* 32(6):1439–1472.
- Almeida P, Kogut B (1999) Localization of knowledge and the mobility of engineers in regional networks. *Management Sci.* 45(7):905–917.
- Arora A, Gambardella A, eds. (2005) *From Underdogs to Tigers: The Rise and Growth of the Software Industry in Brazil, China, India, Ireland, and Israel* (Oxford University Press).
- Ashraf R, Ray R (2017) Human capital, skilled immigrants, and innovation. Working paper, Georgia State University.
- Athukorala PC, Kohpaiboon A (2010) Globalization of R&D by US-based multinational enterprises. *Res. Policy* 39(10):1335–1347.
- Atkin D (2013) Trade, tastes, and nutrition in India. *Amer. Econom. Rev.* 103(5):1629–1663.
- Audretsch DB, Feldman MP (1996) R&D spillovers and the geography of innovation and production. *Amer. Econom. Rev.* 86(3):630–640.
- Bahar D, Choudhury P, Glennon B (2020) An executive order worth \$100 billion: The impact of an immigration ban's announcement on Fortune 500 firms' valuation. Preprint, submitted October 20, <https://dx.doi.org/10.2139/ssrn.3713453>.
- Bahar D, Choudhury P, Sappenfield J, Signorelli S (2022) Human mobility and the globalization of knowledge production: Causal evidence from multinational enterprises. Working paper, Harvard Business School Technology & Operations Mgt. Unit, 22–047.
- Barney JB (1986) Strategic factor markets: Expectations, luck, and business strategy. *Management Sci.* 32(10):1231–1241.
- Bartlett CA, Ghoshal S (2002) Building competitive advantage through people. *MIT Sloan Management Rev.* 43(2).
- Berry H (2014) Global integration and innovation: Multicountry knowledge generation within MNCs. *Strategic Management J.* 35:869–890.
- Besley T, Coate S (1995) Group lending, repayment incentives and social collateral. *J. Development Econom.* 46(1):1–18.
- Borjas GJ (2005) The labor-market impact of high-skill immigration. *Amer. Econom. Rev.* 95(2):56–60.
- Bound J, Demirci M, Khanna G, Turner S (2015) Finishing degrees and finding jobs: U.S. higher education and the flow of foreign IT workers. *Innovation Policy and the Economy*, vol. 15, 27–72.
- Bound J, Khanna G, Morales N (2017) Understanding the economic impact of the H-1B program on the United States. High-Skilled Migration to the United States and Its Economic Consequences. University of Chicago Press, 109–175.
- Branstetter L, Glennon B, Jensen JB (2018) The IT revolution and the globalization of R&D. Lerner J, Stern S, eds. *Innovation Policy and the Economy*, vol. 19 (University of Chicago Press).
- Branstetter LG, Glennon B, Jensen JB (2021) The new global invention machine: A look inside the R&D networks of U.S. multinationals. *Multinational Corporations in a Changing Global Economy*.
- Bronnenberg BJ, Dubé J-PH, Gentzkow M (2012) The evolution of brand preferences: Evidence from consumer migration. *Amer. Econom. Rev.* 102(6):2472–2508.
- Buch CM, Kleinert J, Toubal F (2006) Where enterprises lead, people follow? Links between migration and FDI in Germany. *Eur. Econom. Rev.* 50(8):2017–2036.

- Burchardi KB, Chaney T, Hassan TA (2019) Migrants, ancestors, and foreign investments. *Rev. Econom. Stud.* 86(4):1448–1486.
- Caliendo L, Oromolla LD, Parro F, Sforza A (2021) Goods and factor market integration: A quantitative assessment of the EU enlargement. *J. Political Econom.* 129(12): 3491–3545.
- Campbell BA, Coff RW, Kryscynski D (2012) Rethinking sustained competitive advantage from human capital. *Acad. Management Rev.* 37(3):376–395.
- Cantwell J (1995) The globalisation of technology: What remains of the product cycle model? *Cambridge J. Econom.* 19(1):155–174.
- Chang S, Kogut B, Yang JS (2016) Global diversification discount and its discontents: A bit of self-selection makes a world of difference. *Strategic Management J.* 37(11):2254–2274.
- Chung W, Alcácer J (2002) Knowledge seeking and location choice of foreign direct investment in the United States. *Management Sci.* 48(12):1534–1554.
- Cohen SS, Zysman J (1987) Why manufacturing matters: The myth of the post-industrial economy. *California Management Rev.* 29(3):9–26.
- Collings DG, Mellahi K, Cascio WF (2019) Global talent management and performance in multinational enterprises: A multilevel perspective. *J. Management* 45(2):540–566.
- Conner KR (1991) A historical comparison of resource-based theory and five schools of thought within industrial organization economics: Do we have a new theory of the firm? *J. Management* 17(1):121–154.
- Cuadros A, Martín-Montaner J, Paniagua J (2018) Migration and FDI: The role of job skills. *Internat. Rev. Econom. Finance* 59:318–332.
- DeSanctis G, Glass JT, Ensing IM (2002) Organizational designs for R&D. *Acad. Management Perspect.* 16(3):55–66.
- Dimmock SG, Huang J, Weisbenner SJ (2022) Give me your tired, your poor, your high-skilled labor: H-1B lottery outcomes and entrepreneurial success. *Management Sci.* 68(9): 6950–6970.
- Doran K, Gelber A, Isen A (2022) The effects of high-skilled immigration policy on firms: Evidence from visa lotteries. *J. Political Econom.* 130(10):2501–2533.
- Eisenhardt KM, Martin JA (2000) Dynamic capabilities: What are they? *Strategic Management J.* 21(10–11):1105–1121.
- Florida R (1997) The Globalization of R&D: Results of a survey of foreign-affiliated R&D laboratories in the USA. *Res. Policy* 26(1):85–103.
- Foley CF, Kerr WR (2013) Ethnic Innovation and U.S. multinational firm activity. *Management Sci.* 59(7):1529–1544.
- Fort TC, Keller W, Schott PK, Yeaple S, Zolas NJ (2020) Co-location of production and innovation: Evidence from the United States. Working paper, Dartmouth College.
- Fuchs E, Kirchain R (2010) Design for location? The impact of manufacturing offshore on technology competitiveness in the optoelectronics industry. *Management sci.* 56(12), 2323–2349.
- Furman JL, Porter ME, Stern S (2002) The determinants of national innovative capacity. *Res. Policy* 31(6):899–933.
- Ghemawat P (2003) Semiglobalization and international business strategy. *J. Internat. Bus. Stud.* 34(2):138–152.
- Ghosh A, Mayda AM, Ortega F (2015) The impact of skilled foreign workers on firms: An investigation of publicly traded U.S. firms. Working paper, Georgetown University.
- Glennon B, Morales F, Carnahan S, Hernandez E (2021) Does employing skilled immigrants enhance competitive performance? Evidence from European football clubs. NBER Working Paper No. 29446, National Bureau of Economic Research, Cambridge, MA.
- Gould DM (1994) Immigrant links to the home country: Empirical implications for U.S. bilateral trade flows. *Rev. Econom. Statist.* 76(2):302–316.
- Greif A (1993) Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition. *Amer. Econom. Rev.* 83(3):525–548.
- Gupta AK, Govindarajan V (2000) Knowledge flows within multinational corporations. *Strategic Management J.* 21(4):473–496.
- Hall BH, Van Reenen J (2000) How effective are fiscal incentives for R&D? A new review of the evidence. *Res. Policy* 29(4–5):449–469.
- Hatch NW, Dyer JH (2004) Human capital and learning as a source of sustainable competitive advantage. *Strategic Management J.* 25(12):1155–1178.
- Heckman JJ, Ichimura H, Todd PE (1997) Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *Rev. Econom. Stud.* 64(4):605–654.
- Hernandez E (2014) Finding a home away from home: Effects of immigrants on firms' foreign location choice and performance. *Admin. Sci. Quart.* 59(1):73–108.
- Hernandez E, Kulchina E (2020) Immigrants and foreign firm performance. *Organ. Sci.* 31(4):797–820.
- Hiller S (2013) Does immigrant employment matter for export sales? Evidence from Denmark. *Rev. World Econom.* 149(2):369–394.
- Hines JR (1995) Taxes, technology transfer, and the R&D activities of multinational firms. *The Effects of Taxation on Multinational Corporations*, 225–252.
- Hirano K, Imbens GW (2005) The propensity score with continuous treatments. Gelman A, Meng X-L, eds. *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives* (John Wiley & Sons, Ltd.), 73–84.
- Hirschey RC, Caves RE (1981) Research and transfer of technology by multinational enterprises. *Oxford Bull. Econom. Statist.* 43(2): 115–130.
- Hunt J, Gauthier-Loiselle M (2010) How much does immigration boost innovation? *Amer. Econom. J. Macroeconomics* 2(2):31–56.
- Iranzo S, Peri G (2009) Migration and trade: Theory with an application to the Eastern-Western European integration. *J. Internat. Econom.* 79(1):1–19.
- Jaffe AB (1986) Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits and market value.
- Jaffe AB, Trajtenberg M, Henderson R (1993) Geographic localization of knowledge spillovers as evidenced by patent citations. *Quart. J. Econom.* 108(3):577–598.
- Javorcik BS, Özden Ç, Spatareanu M, Neagu C (2011) Migrant networks and foreign direct investment. *J. Development Econom.* 94(2):231–241.
- Kahn S, MacGarvie MJ (2016) How important is US location for research in science? *Rev. Econom. Statist.* 98(2): 397–414.
- Kerr SP, Kerr WR, Lincoln WF (2015a) Firms and the economics of skilled immigration. Kerr WR, Stern S, Lerner J, eds. *Innovation Policy and the Economy*, vol. 15.
- Kerr SP, Kerr WR, Lincoln WF (2015b) Skilled immigration and the employment structures of US firms. *J. Labor Econom.* 33(S1): S147–S186.
- Kerr SP, Kerr W, Özden Ç, Parsons C (2017) High-skilled migration and agglomeration. *Annual Rev. Econom.* 9(1):201–234.
- Kerr WR, Lincoln WF (2010) The supply side of innovation: H-1B visa reforms and U.S. ethnic invention. *J. Labor Econom.* 28(3): 473–508.
- Kogut B, Kulatilaka N (1994) Operating flexibility, global manufacturing, and the option value of a multinational network. *Management Sci.* 40(1):123–139.
- Kogut B, Zander U (1992) Knowledge of the firm, combinative capabilities, and the replication of technology. *Organ. Sci.* 3(3): 383–397.
- Kogut B, Zander U (1993) Knowledge of the firm and the evolutionary theory of the multinational corporation. *J. Internat. Bus. Stud.* 24:625–645.
- Kuemmerle W (1997) Building effective R&D capabilities abroad. *Harvard Bus. Rev.*
- Kugler M, Rapoport H (2011) Migration, FDI, and the Margins of Trade. Working paper, George Mason University.
- Kumar N (2001) Determinants of location of overseas R&D activity of multinational enterprises: The case of US and Japanese corporations. *Res. Policy* 30(1):159–174.

- Laursen K, Leten B, Nguyen NH, Vancauteran M (2020) Mounting corporate innovation performance: The effects of high-skilled migrant hires and integration capacity. *Res. Policy* 49(9):104034.
- Lee SH, Makhija M (2009) Flexibility in internationalization: Is it valuable during an economic crisis? *Strategic Management J.* 30(5):537–555.
- Lewin AY, Massini S, Peeters C (2009) Why are companies offshoring innovation? The emerging global race for talent. *J. Internat. Bus. Stud.* 40:901–925.
- Li A (2020) High-skilled immigrant workers and U.S. firms' access to foreign venture capital. Working paper, Georgia Tech.
- Mahoney JT, Pandian JR (1992) The resource-based view within the conversation of strategic management. *Strategic Management J.* 13(5):363–380.
- Mansfield BE, Teece D, Romeo A (1979) Overseas research and development by US-based firms. *Economica* 46(182):187–196.
- Marino A, Mudambi R, Perri A, Scalera VG (2020) Ties that bind: Ethnic inventors in multinational enterprises' knowledge integration and exploitation. *Res. Policy* 49(9):103956.
- Marshall A (1890) *Principles of Economics, 8th edn* (1920) (McMillan, London).
- Mayda AM, Ortega F, Peri G, Shih K, Sparber C (2018) The effect of the H-1B quota on employment and selection of foreign-born labor. *Eur. Econom. Rev.* 108:105–128.
- Mayda AM, Ortega F, Peri G, Shih K, Sparber C (2020) Coping with H-1B shortages: Firm performance and mitigation strategies. Working paper, Georgetown University.
- Mayer KJ, Somaya D, Williamson IO (2012) Firm-specific, industry-specific, and occupational human capital and the sourcing of knowledge work. *Organ. Sci.* 23(5):1311–1329.
- Mehra M, Shen H (2022) Skilled immigration, firms, and policy. *J. Internat. Econom.* 139:103662.
- Mithas S, Lucas HC, Jr. (2010) Are foreign IT workers cheaper? U.S. visa policies and compensation of information technology professionals. *Management Sci.* 56(5), 745–765.
- Morales N (2019) High-skill migration, multinational companies, and the location of economic activity. *Federal Reserve Bank of Richmond Working Papers*, vol. 19, 1–59.
- Moser P, Voena A, Waldinger F (2014) German Jewish émigrés and US invention. *Amer. Econom. Rev.* 104(10):3222–3255.
- Murat M, Pistoiesi B (2009) Migrant networks: Empirical implications for the Italian bilateral trade. *Internat. Econom. J.* 23(3): 371–390.
- Naghavi A, Ottaviano G (2009) Offshoring and product innovation. *Econom. Theory* 38(3):517–532.
- National Science Foundation (2020) The state of U.S. science and engineering. <https://nces.nsf.gov/indicators>.
- Olney WW (2012) Offshoring, immigration, and the native wage distribution. *Canadian J. Econom.* 45(3):830–856.
- Olney WW, Pozzoli D (2021) The impact of immigration on firm-level offshoring. *Rev. Econom. Statist.* 103(1):177–195.
- Ottaviano GIP, Peri G, Wright GC (2013) Immigration, offshoring, and American jobs. *Amer. Econom. Rev.* 103(5):1925–1959.
- Ottaviano GIP, Peri G, Wright GC (2018) Immigration, trade and productivity in services: Evidence from U.K. firms. *J. Internat. Econom.* 112:88–108.
- Pearce RD (1999) Decentralised R&D and strategic competitiveness: Globalised approaches to generation and use of technology in multinational enterprises (MNEs). *Res. Policy* 28(2–3): 157–178.
- Peri G, Shih K, Sparber C (2015) Foreign and native skilled workers: What can we learn from H-1B lotteries? NBER Working Paper No. 21175, National Bureau of Economic Research, Cambridge, MA.
- Peteraf M (1993) The cornerstones of competitive advantage: A resource-based view. *Strategic Management J.* 14(3):179–191.
- Peters M (2017) *Trading Barriers* (Princeton University Press).
- Pisano GP, Shih WC (2012) *Producing Prosperity: Why America Needs a Manufacturing Renaissance* (Harvard Business Review Press).
- Polanyi M (1958) *Personal Knowledge, Toward a Post Critical Epistemology* (The University of Chicago Press, Chicago).
- Polanyi M (1966) *The Tacit Dimension* (The University of Chicago Press, Chicago).
- Rosenbaum PR, Rubin DB (1983) The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1):41–55.
- Rosenkopf L, Almeida P (2003) Overcoming local search through alliances and mobility. *Management Sci.* 49(6):751–766.
- Santos Silva JMC, Tenreiro S (2006) The log of gravity. *Rev. Econom. Statist.* 88(4):641–658.
- Saxenian AL (2002) *Local and Global Networks of Immigrant Professionals in Silicon Valley* (San Francisco).
- Schumpeter JA (1939) *Business Cycles* (Porcupine Press, Philadelphia).
- Szulanski G (1996) Exploring internal stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management J.* 17(S2):27–43.
- Teece DJ (1977) Technology transfer by multinational firms: The resource cost of transferring technological know-how. *Econom. J.* 87(346):242–261.
- Turner PS (2022) High-skilled immigration and the labor market: Evidence from the H-1B visa program. *J. Pol. Anal. Manage.* 41:92–130.
- Winter SG (2003) Understanding dynamic capabilities. *Strategic Management J.* 24(10):991–995.
- Wu A (2017) Skilled immigration and firm-level innovation: The US H-1B lottery. Working paper, HBS.
- Xu S-J (2016) Skilled labor supply and corporate investment: Evidence from the H-1B visa program. Preprint, submitted December 1, <https://dx.doi.org/10.2139/ssrn.2877241>.
- Yeaple S (2019) The innovation activities of multinational enterprises and the demand for skilled worker, non-immigrant visas. Hanson GH, Kerr WR, Turner S, eds. *High-Skilled Migration to the United States and Its Economic Consequences* (University of Chicago Press), 41–70.
- Zaheer A, Hernandez E (2011) The geographic scope of the MNC and its alliance portfolio: Resolving the paradox of distance. *Global Strategy J.* 1(1–2):109–126.
- Zaheer S (1995) Overcoming the liability of foreignness. *Acad. Management J.* 38(2):341–363.
- Zhao M (2006) Conducting R&D in countries with weak intellectual property rights protection. *Management Sci.* 52(8):1185–1199.