Management Science

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

Go West Young Firm: The Impact of Startup Migration on the Performance of Migrants

Jorge Guzman

informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–23
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Go West Young Firm: The Impact of Startup Migration on the Performance of Migrants

Jorge Guzman[a,b]

[a] Columbia Business School Management Division, Columbia University, New York, New York 10027; [b] National Bureau of Economic Research, Cambridge, Massachusetts 02138

Contact: jag2367@gsb.columbia.edu, https://orcid.org/0000-0002-0826-8306 (JG)

Received: November 16, 2020
Revised: May 12, 2022; November 5, 2022
Accepted: January 26, 2023
Published Online in Articles in Advance: September 21, 2023

https://doi.org/10.1287/mnsc.2023.4924

Abstract. This paper studies how regional migration to tech clusters impacts the performance of startups within the United States. Startups that move to Silicon Valley experience a significant improvement in performance. This improvement is higher than migrations to other regions in the United States, many of which report null treatment effects. The startups that benefit the most from migration are those leaving low performing entrepreneurial ecosystems and moving to high performing ecosystems, consistent with an agglomeration mechanism. Within different measures of the ecosystem, the level of local patenting predicts startup improvements more than venture capital or the quality-adjusted number of startups, suggesting the local innovation environment is more important to migrant performance than financing or the presence of other startup peers.

History: Accepted by Toby Suart, entrepreneurship and innovation.

Funding: This work was supported by Ewing Marion Kauffman Foundation, the Jean Hammond (1986) and Michael Krasner (1974) Entrepreneurship Fund at Massachusetts Institute of Technology (MIT), and the Edward B. Roberts (1957) Entrepreneurship Fund at MIT.

Supplemental Material: The data files and online appendix are available at https://doi.org/10.1287/mnsc.2023.4924.

Keywords: entrepreneurship • location • migration

1. Introduction

In 2014, a firm born in Silicon Valley (the San Francisco Bay Area) was 35 times more likely to achieve an initial public offering (IPO) or a high-value acquisition than one born in San Antonio, Texas (Andrews et al. 2020, author calculations). Under most metrics, there are large differences in the outcomes of startups across regions. Substantial research documents that a core driver of such differences is agglomeration—the co-occurrence of a range of regional inputs and institutions that support better startup selection and performance (Stuart and Sorenson 2003b, Glaeser and Kerr 2009, Samila and Sorenson 2011).[1] For example, tech clusters that provide startups access to high-quality financing options, that provide a richer level of local innovative ideas, or that have a large number of other peer startups support the entry of better startups and allow them to both grow faster and be more successful (Delgado et al. 2010, Lerner 2012, Kerr and Robert-Nicoud 2020, among others).

Although the importance of agglomeration on performance is well established, *who* can exploit these regional inputs is less clear. Prior work has by and large studied the extent to which firms born in different ecosystems obtain different outcomes. However, because regional agglomeration influences both the types of firms created and their subsequent growth, the high performance of startups across regions may be driven more by the birth of better startups than the way these startups take advantage of inputs to grow. Small university towns appear to have this feature. They serve as "nursery" locations for innovative startups (Duranton and Puga 2001), but these startups are unable to grow and tend to leave, resulting in high out-migration and very low in-migration (Bryan and Guzman 2021). The importance of differentiating between how agglomeration supports firm selection and firm performance is particularly salient for migrant firms: the benefits of location are only captured by migrants if locations enhance growth after founding. Indeed, even then, the ability of migrants to capture agglomeration benefits is uncertain because doing so may also require local embeddedness and social networks (Michelacci and Silva 2007, Dahl and Sorenson 2012). For any of these reasons, a migrant moving to Silicon Valley may not see a benefit, even when local firms outperform.

In addition, easier access to locally agglomerated resources is not the only way migration could lead to higher performance for migrants. The act of migration in and of itself—through mechanisms such as changing individual ambition, removing secondary options, or increasing entrepreneurial impetus—can change the performance of startups. Tech clusters such as Silicon Valley are also characterized by a distinct culture and institutions (Saxenian 1994), or even desirable weather, which in itself impacts productivity (Dell et al. 2012). Yet, such culture and institutions do not directly represent the closer presence of economic resources and their relationship to better performance.

To take agglomeration seriously, it is necessary to go beyond asking if there is a benefit from migration and consider specific conditions under which startups would benefit from it. In agglomeration, the benefits of migration do not come from migration itself but from the (relative) improvement in the quality of the local entrepreneurial ecosystem experienced by the migrant through its change in location.

Understanding these mechanisms is also important for both strategy and economic policy.

For strategy, observing a positive value of migration would imply location choice is key to formulating startup strategy. Startups then not only need to know how to use their local ecosystem, as is already explained in existing work (Gans et al. 2023). They can choose (and change) their ecosystem, even after founding. For example, a key tradeoff faced by founders may be whether to seek high agglomeration tech clusters or prefer to grow locally due to their personal embeddedness (Dahl and Sorenson 2012). If migration can be valuable, *choosing*, and not simply using, a location would be core to entrepreneurial strategy.

For economic policy, migration is a different lever for regional ecosystem design, with a distinct timeline and potential impact, than the endogenous development of an entrepreneurial ecosystem's local firms—the core of current work (Stam 2007, Feldman et al. 2019). Developing an entrepreneurial ecosystem requires the symbiotic evolution of several stakeholder groups. Migrant high-quality startups may help kickstart, grow, or sustain this process. Some entrepreneurial programs, such as Startup Chile have been built based on this hypothesis. However, Startup Chile is one case study, and existing academic work on this program (Gonzalez-Uribe and Leatherbee 2018) has studied the accelerator's design but not the benefits of migration for startups. This paucity of evidence on how migrant startups engage with local agglomeration limits its consideration in economic policy.[2]

This paper aims to develop a better understanding of startup agglomeration by empirically assessing whether and how migrants benefit from migration. To be sure, not all startups that move do so seeking higher performance. In a companion paper (Bryan and Guzman 2021), we show that the value migrant startups have for U.S. destinations appears to be more related to differences in the amenities offered to founders—such as weather, personal taxes, and cost of living—than startup inputs. Perhaps surprisingly, top startup hubs such as Silicon Valley, Boston, or New York lose more startups than they gain. These facts shape the empirical strategy of this paper.

Before delving into the results, it is useful to overview the core empirical challenges a study of the impact of migration on performance would face and the way these are addressed. First, both defining and observing the migration of a startup is difficult. To measure migration, I build from prior related work (Bryan and Guzman 2021, Conti and Guzman 2023) and take advantage of business registration records to document the timing and relocation of firms across states. Business registration is the formal organization of any corporation, limited liability company, or limited partnership in a state. They are filed both at the state of founding and any other state in which a startup establishes operations. The details in the business registration forms allow documenting the timing of entry into a state and separating the migration of headquarters from expansions of offices across states. I define migration as the geographic relocation of headquarters. I study specifically the migration of firms registered under Delaware jurisdiction (i.e., Delaware is their legal jurisdiction, although not their home location) and their movement across 36 states. Delaware firms are a small share of all firms but account for a large portion of growth-oriented startups (Guzman and Stern 2020).

Second, there is the problem of identification. The choice to move is obviously endogenous and not random. To address this endogeneity, I implement three independent strategies. The first uses machine learning on the cross-sectional version of the data, incorporating a large number of observables at founding (and before moving) to estimate the predicted performance of firms (their estimated 'quality') in their home region and control for differences across firms. Introducing machine learning allows me to use a much larger number of observables to better account for potential omitted variables than prior approaches, such as kitchen-sink regressions and coarsened exact matching. The second identification approach uses variation in the timing of migration within migrants in a panel version of the data, including fixed effects for each migrant, their age, and their time of migration. The empirical comparison here is within movers, comparing the performance of migrants to other migrants with the same destination but who move at different times. The third identification approach uses an instrumental variable—the age of the founder—in cross-sectional data. The key assumption is that older founders are less likely to move conditional on firm characteristics (Molloy et al. 2011) due to either higher personal costs of moving (e.g., spouse wage or relocating a family) or differences in personal preferences for moving across age. To validate the exclusion restriction, I show that estimated startup quality does not vary across founder age conditional on my controls, nor does founder age predict performance for nonmovers. The two-stage least squares (2SLS) estimate compares firms that moved to counterfactual firms that did not move due to higher *personal* founder costs.

Using this approach, the present paper documents four interrelated insights.

First, as in prior work (Guzman and Stern 2015, 2020; Catalini et al. 2019), there are systematic startup characteristics at founding that predict performance for growth-oriented firms in the United States. Relative to that prior work, this paper introduces machine learning as an enhancement of the logit model used previously to develop estimates of startup quality. It includes a higher number of firm characteristics and focuses specifically on estimating the predicted performance of Delaware startups that do not move. That is, it creates estimates of the predicted performance of high-growth companies at home.

The out-of-sample area under the receiver operating characteristic curve (ROC score) of the random forest model used to predict performance is 0.8, implying about 60% of the variation in outcomes is accounted for by the model. Although there is no benchmark for what should be the ROC of the perfect model, it is worth emphasizing that even the perfect model mapping founding characteristics to outcomes would not have a ROC of 1.0 because there are shocks occurring *after* founding that necessarily shape firm performance.

Second, the performance improvements that migrants get when moving to Silicon Valley are substantial. In the cross-sectional estimates, migrants increase the probability of achieving an equity outcome such as IPO or acquisition by 2.77 times after moving. This machine learning estimate accounts for selection better than more traditional approaches. The point estimate is 42% lower than a naïve estimate without any controls, and 23% lower than estimates using coarsened exact matching (Iacus et al. 2012). The precision also increases. Furthermore, when I use an Oster framework on coefficient stability to estimate the potential role of remaining unobservables (Altonji et al. 2005, Oster 2019), the estimate only drops slightly to 2.6, and the size of unobservables required to make the effect zero is implausible. Small acquisitions and sell-offs of firms do not drive these effects: The estimates are similar (and even larger) in regressions dropping all acquisitions with reported values below $100 million and when considering only IPOs. The panel data regressions report a comparable coefficient to the machine learning model, and the instrumental variables result is, if anything, slightly larger, although not statistically different. The panel data also do not show any pre-trends in the outcomes before moving.

The results also appear similar when considering other outcomes related to agglomeration and performance. Movers to Silicon Valley are 3 times more likely to receive venture capital (VC), 1.6 times more likely to file for a patent, and 47% more likely to get a trademark (a proxy for introducing a product). The agglomeration benefits of Silicon Valley are substantial.

Third, other regions also offer benefits to migrants, but they are smaller than Silicon Valley. When I repeat the same methodology on the fourteen most popular destinations in my data outside Silicon Valley, I find that Silicon Valley is an outlier in the benefits to migrants. Only one region, the Denver-Aurora, Colorado, Metropolitan Statistical Area (MSA) has an estimated treatment effect on equity outcomes comparable to Silicon Valley, whereas New York City or Boston have treatment effects that are much lower. The differences in the benefits of migration to get venture capital are even more dramatic, with Silicon Valley being a distinct outlier.

Finally, fourth, I study the mechanisms through which migrants benefit from migration by considering the *relative* increase in ecosystem characteristics for migrants. I focus on three variables that are traditionally used as measures for a strong startup ecosystem—the local level of total entrepreneurship (Marshall 1890), the amount of local idea generation (Forman et al. 2016) (proxied by patents per capita), and the level of local venture capital (Chen et al. 2010). For movers to Silicon Valley, increases in each of these variables predict higher performance when considered independently, but they are highly correlated. Only patenting per capita remains positive once I consider them together. This suggests the innovation environment is more valuable to migrants than financing or the presence of startup peers. The results remain very similar, with comparable coefficients, when I expand the analysis to all regions and include origin state by year and destination city by year fixed effects.

The fact that the estimates relating the relative ecosystem improvements to migrant performance are positive and similar across regions is consistent with an agglomeration mechanism. Startups that move from worse to better regions see the highest boost from migration, and the effect is well determined by how big the improvement in the ecosystem was. It suggests agglomeration as a straightforward explanation for the outsized benefits of Silicon Valley: Because it is a regional outlier in the availability of resources, the relative improvements experienced by migrants to this destination are naturally much larger. These results are also not consistent with other potential mechanisms, such as the culture of the destination or a treatment effect of the act of moving by changing founder incentives.

Together, these results provide a comprehensive assessment of the way in which agglomeration—and particularly the innovation environment—mediate startup performance.

2. Literature Review: Tech Clusters, Migration, and Startup Performance

Understanding how location shapes firm performance has been at the center of social science research for over a century. Initial work by Marshall (1890) and Weber (1929) focused on how industry concentration creates an agglomeration of resources that enables higher firm productivity. Roback (1982) built on this insight to use spatial equilibrium—a concept representing the idea that the marginal mover is indifferent between locations—to characterize the migration of people when firm locations are taken as exogenous. Spatial equilibrium continues to be the workhorse tool of economic geography to understand how differences in the attributes of regions attract individuals (Glaeser and Gottlieb 2009, Albouy 2016, Hsieh and Moretti 2019).

A related, but separate, line of work pushed on the Roback assumption that the location of firms can be taken as exogenous by building on Marshall's insight that firm location is codependent to the location of other valuable resources. Krugman (1991) showed economies of scale both determine and maintain the preference of industry for a location due to proximity to its market. Entrepreneurship and innovation research moved beyond market effects to instead recognize the more specific role of ideas in driving the incidence and performance of regional entrepreneurship (Marshall 1890, Jaffe et al. 1993, Audretsch and Feldman 1996, Zucker et al. 1998). The ability to grow these ideas, in turn, depends on access to early-stage capital to build organizations that commercialize them (Stuart and Sorenson 2003a, Chen et al. 2010). These insights have led to entrepreneurial regions being represented as a tripartite relation of ideas, capital, and entrepreneurship, commonly called "tech clusters" (Stuart and Sorenson 2003b, Kerr and Robert-Nicoud 2020, Moretti 2021). The Boston, Massachusetts, area and Silicon Valley are the quintessential case examples of this phenomenon (Saxenian 1994). Not merely a conceptual object, tech clusters today account for an important share of the regional divergence in economic outcomes experienced by U.S. regions (Andrews et al. 2020, Moretti 2021), creating recent calls to understand and replicate them in both policy[3] and academia (Stam 2007).

Among the possible mechanisms of tech clusters, a particularly rich literature has emerged in understanding how the social structure of regions impacts the incidence and type of firm formation in them (see Sorenson 2018 for a review). Entrepreneurship in a region does not occur in a vacuum. It stems from individual interactions across local stakeholders. Regional financing depends on access to local social networks (Sorenson and Stuart 2001), knowledge spillovers often require people changing jobs across companies to occur (Almeida and Kogut 1999), and eventual entrepreneurs develop an inclination toward starting a company through work and school peers (Nanda and Sørensen 2010, Lerner and Malmendier 2013). More importantly, the choice to start a company depends critically on the local embeddedness of individuals and how it enables them to access resources (Michelacci and Silva 2007, Dahl and Sorenson 2012). Yet, social networks are not everything, leaving an important role for agency. As emphasized by Feldman (2014): "Entrepreneurs benefit from location. But entrepreneurs are also pivotal agents of change that can transform local communities" (see also Feldman and Francis 2003).

The analysis of mechanisms has also deepened our understanding of the process of agglomeration itself and how it applies to tech clusters. Although the work of Marshall emphasized three distinct agglomeration costs mediating the benefits of proximity—the transportation costs of goods, of people, and of ideas—more recent work building on the tradition of regional competitive advantage (Delgado et al. 2014) has instead sought to understand the resource characteristics that drive the emergence of tech clusters. A range of approaches have been created, with perhaps the most widely applied in policy circles being Massachusetts Institute of Technology (MIT)'s Regional Entrepreneurship Acceleration Program (REAP) (Budden and Murray 2019). At a broad level, this cluster-based analysis recognizes the interdependent nature of a few key inputs, such as the local entrepreneurial capacity, the local availability of innovation and ideas, the local supply of risk capital, the nature of regional competitiveness in the global economy, and the rules and regulations of a location (Marx et al. 2009, Assenova and Sorenson 2017), as the key ingredients of a productive regional ecosystem.

Together, the previous arguments recognize the importance of location to the performance of startups, but also predict differing benefits and costs from changing a startup's location and how these relate to its eventual success. For example, although the value of networks emphasizes the nature of being a local actor in an ecosystem, the value of agglomeration instead focuses on the role of physical proximity to local resources, which could, in principle, be at least partially accessed by relocating firms. Fundamentally, understanding whether startups benefit from relocating to economic clusters is an empirical question that speaks clearly to the ways in which tech clusters beget resources for startups and the determinants of performance within these.

Is migration to a tech cluster such as Silicon Valley beneficial to a startup born in a different location? If it is, what type of agglomeration mechanism is responsible for the benefits of relocation? This is the question to which this paper now turns.

3. Data

The data are built from business registration records of 36 U.S. states from 1988 to 2014. These states represent 82% of the U.S. population and 86% of the 50 largest metro areas.[4] These data were retrieved during the first phase of the broader effort of the Startup Cartography Project (Andrews et al. 2020) and then improved to measure the migration of startups across states in Bryan and Guzman (2021).

Business registration records are public records created endogenously when founders register their firm as a corporation, limited partnership, or limited liability company with the Secretary of State (or Commonwealth) of any U.S. state. Their initial filing represents an early moment of a firm when it transitions into a formal entity to undertake a specific business purpose. From a legal perspective, business registration marks the legal birth of the firm.

To focus this study on startups with high growth intention, I select the subset of startups registered under Delaware jurisdiction. These are not firms headquartered in Delaware. They operate in every U.S. state but have chosen to establish under Delaware Corporate Law rather than the regime of their home state. There are some significant benefits to registering in Delaware for startups.[5] These benefits are more useful for startups that will be large or for startups interacting with investors, including venture capitalists. Conversely, registering under Delaware jurisdiction also carries extra costs because it requires maintaining two registrations (one in Delaware and one in the state of operation), imposing extra fees.[6] This creates a natural separating equilibrium, with growth-oriented companies choosing to register in Delaware, but the bulk of firms registering locally. Although Delaware companies represent only about 4% of all firms, they account for more than half of all publicly listed firms and more than 60% of all VC financing (Catalini et al. 2019). In terms of outcomes, Delaware-founded startups are 23 times more likely to achieve an IPO or be acquired than non-Delaware firms (Guzman and Stern 2017, 2020).

The complete data set contains the registration of 405,536 new Delaware firms observed in their home state. I enhance business registration data by using a name-matching algorithm[7] to merge business registrations with four other data sets: (i) three types of intellectual property filings from the U.S. Patent and Trademark Office (trademark applications, patent applications, and patent assignments), (ii) all new IPOs in the United States from the SDC New Issues database, (iii) all U.S. M&A activity reported in the SDC Mergers and Acquisitions database, and (iv) venture capital activity from Thompson Reuters VentureXpert.

3.1. Firm Observables at Founding

I create two binary measures from business registration records indicating whether a firm is a corporation and whether it is an LLC. Building on existing evidence that firm name length predicts performance (Green and Jame 2013), I create 12 measures of firm name length, including a continuous measure of the number of words in the firm name, the square of the number of words, and 10 binary indicators for whether the name is exactly 1 through 10 words long. I create industry measures by using the approach of Guzman and Stern (2015, 2020), which uses a large sample of firms with industry (NAICS) code to create a name-based algorithm that allows categorizing firms into different economic clusters from the U.S. Cluster Mapping Project (Delgado et al. 2014). There are 13 binary measures following this approach, one for each of the following groups: Agriculture and Food, Automotive, Chemicals, Clothing, Consumer Apparel, Distribution and Shipping, Energy, High Technology, Local Industries, Mining, Paper and Plastic, Publishing, and Services. I also create five more measures for names associated with specific high-tech industries that have accounted for a meaningful share of high-growth entrepreneurship in this period: Information Technology (IT), Biotechnology, E-Commerce, Medical Devices, and Semiconductors. Finally, I create six measures from intellectual property filings. Three indicate whether the firm applies for a patent in its first year, has a patent assigned (from a prior inventor) in its first year, or files for a trademark in its first year. The other three indicate whether the firm applies for more than one patent, is assigned more than one patent, or applies for more than one trademark in its first year.

In total, there are 38 measures observable at the time of firm founding, which can be combined in 703 ways in two-way interactions, for a total of 741 observable measures at founding. The goal in this step is to remain as flexible as possible in creating observables to use in the machine learning approach. Then, in a subsequent step, I allow a penalized variable selection procedure to select only those that best predict treatment or outcomes.

3.2. Defining and Measuring Migration

I use the registration of companies across states to track their entry into a state and categorize whether this entry constitutes an expansion through a satellite office or the geographic relocation of the startup's headquarters. Consistent with other work (Bryan and Guzman 2021, Conti and Guzman 2023), I define only change of headquarters as migration.

Companies need to register in the state in which they are founded and in every state in which they rent an office location or any real estate, hire people, or set up a local bank account. The matching across states is easy since (except for minor exceptions) firm names are required to be sufficiently different from each other to avoid customer confusion and must be consistent in their use.[8] Together with considerable research assistant support and coauthors on other projects, I spent meaningful effort over several years validating the quality of these matches and confirming the records used in a state's records office tracked the destination address of firms well. State registrations include three different addresses: the address of the firm's local office within the state, the address of the principal office (headquarters), and the address of the registered agent or lawyer. Only a change in the second one, the headquarters address, would constitute a migration of the startup.

To provide a tangible example, I include in Figure A2 in the online appendix the *California* business registration records for two MIT startups founded in 2010: Ginger.io and Sociometric Solutions (later Humanyze). Both startups were founded at the MIT Media Laboratory by PhD students of Professor Alex (Sandy) Pentland based on work done during their dissertations. Both startups focused on the application of analytics to handheld devices to understand social dynamics. However, Ginger.io decided to move early on to Silicon Valley, whereas Sociometric did not. Accordingly, Ginger.io shows a business registration with a principal executive office in Silicon Valley. We also see the address of the chief executive office (CEO; which is often used as validation in the measurement) is also in Silicon Valley. In contrast, Sociometric Solutions shows a principal executive office in Boston, and a CEO office in Boston. The only address in California is the address of principal office in California, indicating that Sociometric Solutions' role in California is only a satellite office. In this case, Ginger.io would be considered a migration, but Sociometric Solutions would not.

I use the time of initial registration at the destination as the migration date. To guarantee the firm was established in the origin region first, I require that the time elapsed between registration in the origin state and destination is at least three months. Furthermore, I exclude all migrations where the origin state is also part of the destination MSA to avoid cross-state migrations within the same metro area. Finally, I focus only on migrations within the first two years of founding, the early stages of the firm, to allow time to experience outcomes after founding.

Moves to Silicon Valley is a variable is equal to one if a firm moves to Silicon Valley in the first two years after founding, where Silicon Valley is defined as the union of the San Jose-Sunnyvale-Santa Clara, California, MSA and the San Francisco-Oakland-Hayward, California, MSA. Migrations to other destinations are defined in the same way.

Although useful in tracking migration, this approach does come with some limitations.

First, by only observing state-to-state migrations, I am unable to include the migration of startups across cities within the same state, such as Los Angeles, California, to Silicon Valley. Although migrations within state are undoubtedly of theoretical interest, the empirical concerns caused by not having them are minimal since the empirical approach focuses on comparing the performance of migrants to similar startups from the origin region, which we do observe. Out-of-state migrations also represent the modal migration of startups in the United States and better capture the phenomenon of a distant firm approaching a new destination region that motivates this paper.

Second, my approach only focuses on headquarters migrations, creating a binary treatment variable for a phenomenon that is often continuous. In some cases, headquarter migrations would only be partial migrations as a share of the company remains in the origin region. Understanding *what* moves and *how much* of the firm moves are important avenues for future work.

Finally, focusing on business registration records to observe firm birth creates a sample selection of firms that have engaged in a relatively meaningful investment in the origin region but then moved. However, it abstracts away from several other types of movers. Two important ones include the founders and firms that move even before registering in the origin region (e.g., Netscape), and the startups that are never founded because founders do not wish to move but know they would have to. Therefore, although the estimates in this paper are useful to understand the gains of movers, it is clearly a local average treatment effect among the broader population of potential startup movers. Future work may engage in understanding other types of moving more fully to make broader claims about welfare and the allocation of U.S. entrepreneurial activity.

3.3. Measures of Regional Ecosystem Characteristics

Next, I include several regional ecosystem characteristics porting directly from the data built in Bryan and Guzman (2021). The measures are five. *Patenting per Capita* is the number of patents filed per thousand of the population in a year and region. *VC per Capita* is the total number of millions of VC dollars invested per thousand population in a year and region. *Entrepreneurship per Capita* is the quality-adjusted quantity of startups founded in an MSA and year per 10,000 of population (to have all variables at similar scales), using data from the Startup Cartography Project. *Personal Tax Rate* is the personal tax rate faced by individuals in that city at the 95th percentile of income, estimated by Moretti and Wilson (2017). *Sunshine Percentage* is the share of days that have sunshine.

3.4. Measures of Firm Performance

I develop four outcome measures based on the firms' observed performance six years after founding. *Equity Growth*, the key outcome of interest, is equal to one if a firm achieves an IPO or acquisition and zero otherwise. Although rare, equity growth represents a highly desirable outcome for entrepreneurs with high growth intention (the sale of their company) and closely matches the anecdotal incentives sought by many high growth founders. *Gets a Patent* is equal to one if the firm files or acquires (is assigned) one or more patents, excluding the first-year window used for at-founding observables. *Gets a Trademark* is equal to one if the firm files one or more trademarks, excluding the first-year window.[9] *Gets Venture Capital* is equal to one if the firm receives venture capital.

3.5. Measuring the Age of Founders with Ancestry.com

Finally, I also measure the age at founding for a subsample of founders. To do so, I developed a web-scraping algorithm that scrapes Ancestry.com to get the founder's year of birth and estimates the age of the founders at the time they are founding the firm.

Ancestry.com is the industry standard for ancestry research and contains substantial information about U.S. resident individuals that has been reported in public records. This information includes birth and death certificates (of self and kin), marriage certificates, deeds, and others. Using my university library subscription to Ancestry.com, I developed an algorithm that searches Ancestry.com for all instances of the names of the firm directors only in the firm's state of origin and downloads the first 10 matches for each individual. I record the year of birth noted in these 10 matches, which is reported in about two-thirds of the returned records. I filter to only matches where the founder was between 17 and 70 years old at startup founding. Next, I assign each startup the top match based on the "priority" of the founders' title in the director records (CEO, President, and Managing Directors first, then other executive titles, and then general directors). Within the highest priority bracket, I choose the record with the smallest edit (Levenshtein) distance between the founder's name in the corporate record and the public record name (this allows accounting for differences in middle initials and nicknames). Although this scraping was quite slow, averaging 30 seconds per search, and many founder names did not return a match, about four months of scraping allowed me to recover the age at founding for 23,794 firms. The summary statistics of this data set are reported in Table A1 in the online appendix.

There are several limitations to using scraped data from Ancestry.com to measure founder age. The biggest one is that the name of the CEO in the registration record may sometimes spuriously match a record returned by Ancestry.com with the same name but that is not the same person. In these cases, the data will record the wrong age for the founder, leading to noisier estimates. This concern is minimized because age is only used as an instrumental variable in my setting. As long as the age is measured well enough to allow a strong first-stage regression and the mismatches in age are uncorrelated to the startup performance (to maintain the exclusion restriction), this noisy measurement can still be valid. Therefore, although noisy measurement will lead to a noisier estimate, there are no apparent ways these mismatches of names would cause bias.

3.6. Summary Statistics

Table 1 shows summary statistics for some of the key variables. Figure 1 compares the incidence of outcomes for all firms and for movers to different destinations. The average likelihood of an equity growth outcome for firms in my data are 1.6%. However, this measure increases to about 4% when considering migrants and is similar (3.5%) when considering only migrants to four non-Silicon Valley startup hubs—New York City, Boston, Austin, and Seattle. However, the estimate for Silicon Valley migrants is more than twice as large as other regions, at 9%. Similar, although less stark, differences are observed in other variables. Migrants appear to perform better than locals, and Silicon Valley migrants are outliers within this migrant group.

Although the over-representation of migrants in success outcomes is significant, these simple statistics also quickly raise concerns about selection bias. The following section explains the econometric methodology.

4. Empirical Strategy

4.1. Estimating Entrepreneurial Quality

A core requirement of the empirical approach is to measure the expected performance of startups if they had remained in their home region. To do so, I implement the "entrepreneurial quality" approach in Guzman and Stern (2015) and related work (Catalini et al. 2019; Fazio et al. 2019, 2021; Guzman and Kacperczyk 2019; Andrews et al. 2020; Guzman and Stern 2020). This approach combines three interrelated insights. First, as the challenges to reach a growth outcome as a sole proprietorship are formidable, a practical requirement for any entrepreneur to achieve growth is business registration (as a corporation, limited partnership, or limited liability company). This allows for forming a quasi-population of entrepreneurs "at risk" of growth at a similar (and foundational) stage of the entrepreneurial process. Second, it is possible to distinguish among business registrants by observing choices the founders make at or close to the time of registration, informed by their ambitions and expectations for the firm. Examples of these choices include whether the founders name the firm after themselves (eponymy), whether the firm is organized to facilitate equity financing (e.g., registering as a corporation or in Delaware)

Table 1. Summary Statistics

	Count	Mean	Standard deviation	Minimum	Maximum
Startup outcomes					
Equity growth	405,536	0.015	0.123	0	1
Gets venture capital	405,536	0.020	0.141	0	1
Gets a patent	405,536	0.054	0.225	0	1
Gets a trademark	405,536	0.071	0.257	0	1
Migration measures (within 2 years)					
Moves (anywhere)	399,804	0.022	0.148	0	1
Moves to Silicon Valley	268,201	0.002	0.0462	0	1
Regional measures					
Patents per capita	315,959	0.011	0.0156	0.00045	0.078
Venture capital per capita	315,959	0.006	0.0115	0	0.054
(Quality-adjusted) entrepreneurship per capita	304,080	0.000	0.0000621	1.0 e-06	0.00023
Sunshine percentage	315,959	0.614	0.0699	0	0.84
Personal income tax rate (at 95th percentile)	315,959	0.236	0.0242	0.2	0.27
Founding characteristics					
Patent	405,536	0.033	0.179	0	1
Trademark	405,536	0.016	0.127	0	1
Corporation	405,536	0.433	0.496	0	1
Short name	405,536	0.470	0.499	0	1
Eponymous	405,536	0.073	0.261	0	1

Notes. Reports observables for the data. *Moves (Anywhere)* is equal to one only for firms that move within two years of founding and zero otherwise. It excludes all firms that move after two years of founding. *Moves to Silicon Valley* is defined only for firms born outside California.

and whether the firm seeks intellectual property protection (e.g., a patent or trademark). Third, one can leverage the fact that, although rare, it is possible to observe meaningful growth outcomes for some firms (e.g., those that achieve an IPO or high-value acquisition).

Combining these insights, consider a firm fully characterized by many (even infinite) founding observables Z_i. Entrepreneurial quality is the relationship between a specific growth outcomes g_i and founding characteristics. Specifically, for a firm i and a growth outcome g_i quality is

$$\theta_i = P(g_i | Z_i). \tag{1}$$

Given a subset of observed founding characteristics $Z_{i'} \in Z_i$, an (albeit imperfect) empirical measurement of quality can then be estimated as the *predicted* out of sample probability of measured founding characteristics on performance—that is, $\hat{\theta}_i = \hat{P}(g_i | Z_{i'})$. I perform this estimate using variable regularization through double least absolute shrinkage estimator (LASSO) (Belloni et al. 2014) and a random forest (Breiman 2001).

To evaluate the performance of these predictions, I use the area under the ROC curve (i.e., the ROC score). This is the most accepted method for evaluating the predictive fit of binary models. In my setting, the ROC score represents an answer to the following problem: if two random startups, one who achieved growth and one which did not, are fed to the machine learning predictive model from (1), what is the probability that this model will score the growth startup higher than the nongrowth startup? A fully uninformative classifier will have a ROC score of 0.5, whereas a perfect classifier will have a ROC score of 1. *Fit* is defined as the share of the distribution between 0.5 and 1 covered by the ROC score. It can also be interpreted as the share of variation in outcomes accounted for by the predictive model, which will be used later in the coefficient stability tests.[10]

$$Fit = (1 - ROC)/0.5 \tag{2}$$

4.2. Cross-Sectional Approach Using Machine Learning

The first empirical strategy uses machine learning in cross-sectional data. Consider many firms indexed by i, born outside Silicon Valley. The firms are fully characterized by a high-dimensional (even infinite) number of observables Z_i. The firm's performance Y_i is determined by two structural functions of these observables, g_1 and g_0, and two additively separable error terms U_{i1} and U_{i0}.

$$Y_i = \begin{cases} Y_i(1) = g_1(Z_i) + U_{i1} & \text{if located in Silicon Valley} \\ Y_i(0) = g_0(Z_i) + U_{i0} & \text{if located outside Silicon Valley} \end{cases}$$

The goal is to estimate the treatment effect on the treated.

$$\Delta = E\left[\frac{Y_i(1) - Y_i(0)}{\theta} \middle| S_i = 1\right] \tag{3}$$

The econometric challenge is that we do not observe $Y_i(0)$ (nor g_0) for those who move and therefore cannot estimate Δ directly. The goal of index models is to use some set of observables $X_i \in Z_i$ to estimate a function

Figure 1. (Color online) Difference Between Locals and Migrants in Regions

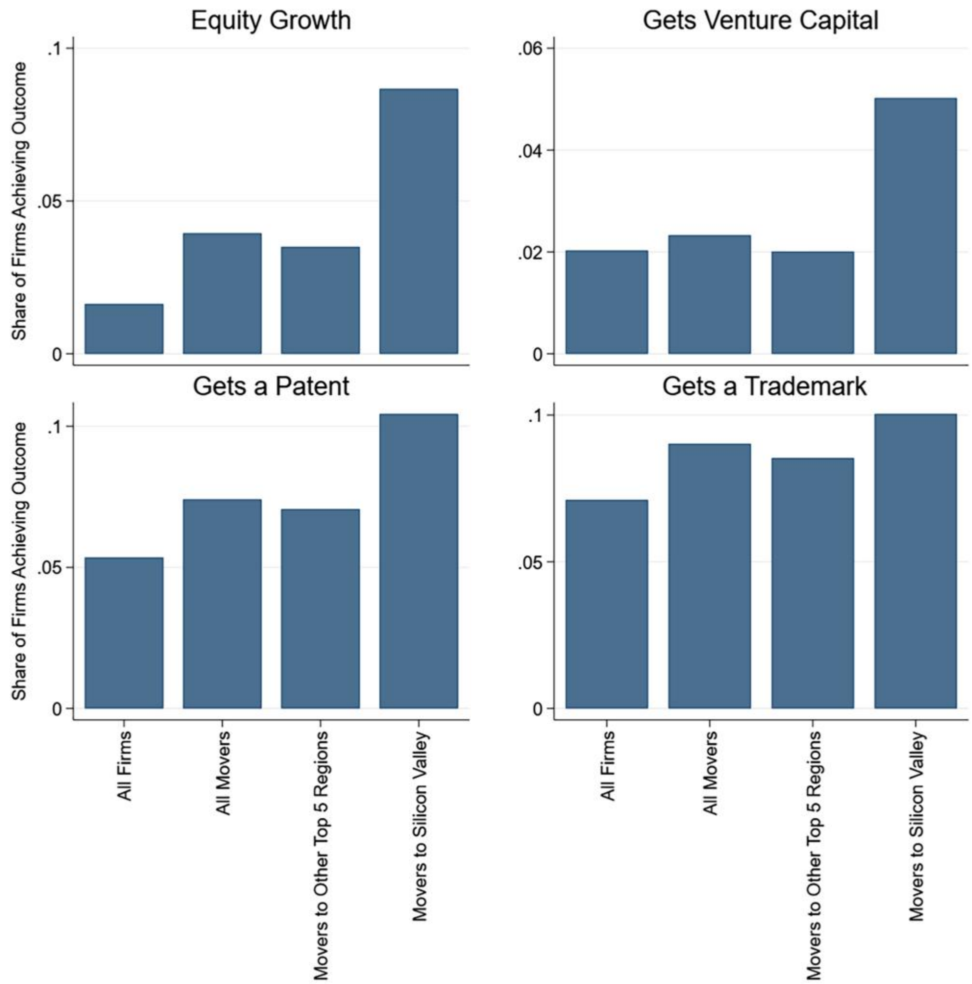

Notes. Plots the share of firms that achieve different startup outcomes across four samples: all firms in the data, all movers, movers to the top five regions in terms of venture capital excluding Silicon Valley, and movers to Silicon Valley. The other top five regions are Boston, New York, Austin, and Seattle.

$\hat{g}_0(X_i)$ such that $\hat{g}_0(X_i) \approx g_0(Z_i)$. If the errors terms can be assumed to be mean-zero (i.e., $E[U_{i1} - U_{i0} | D_i, X_i] = 0$) then Δ is identified. This paper implements the double LASSO approach, which uses a high-dimensional number of observables $Z_{i'}$ to estimate these counterfactuals.

A central assumption in this approach is that the predicted value of performance from observables $\hat{g}_0(Z_{i'})$ is close enough to the true value a firm would see at their origin region. This assumption is testable. To do so, building on Altonji et al. (2005) and Oster (2019), I consider the stability of the coefficients when follow-on information w_i is added. Then, one can back out the importance of remaining unobservables by assuming that unobservables are not more correlated to the outcome than the observables w_i. To do so, I compare the coefficient stability from a regression with and without founding state by founding year fixed effects. Because the mechanisms tests in this paper show the characteristics of the founding state are a key component in determining the benefits of moving, then these fixed effects should be highly correlated with selection into migration.

4.3. Panel Data Regressions

The second empirical approach takes advantage of the data in a panel structure. Because not all migrants move at the same time, I set up a panel that includes age and firm fixed effects and compares the performance of early movers to other movers that have not yet moved. That is, for each migrant firm i, moving at age m, of age t, I run OLS regressions of the form

$$Y_{i,m,t,\tau} = \lambda_i + \gamma_t + \rho_m + \boldsymbol{\beta}'_{\boldsymbol{\tau}} \boldsymbol{M}_{\boldsymbol{i},\boldsymbol{\tau}} + \epsilon_{i,t,m,\tau},$$

where $\boldsymbol{M}_{\boldsymbol{i},\boldsymbol{\tau}}$ is a vector of individual indicators for each value of τ, defined as the difference between the age at migration and t. λ_i is a firm fixed effect, γ_t is an age fixed effect, ρ_m is a fixed effect for age at migration, $Y_{i,m,t,\tau}$ is an outcome measure, and $\epsilon_{i,t,m,\tau}$ is uncorrelated noise. The coefficients of interest are the vector $\boldsymbol{\beta}_\tau$, which represent the differences in the performance of migrants after migration (or before if $\tau < 0$). To take advantage of additional time periods and better observe differences while keeping close to the early stages of the firm, I change the data set to consider the firm's first six years and expand to the quarterly level rather than annually.

4.4. Instrumental Variables Estimates Using Age of Founder

Finally, I perform an instrumental variables estimate using the founder's age as an instrument. The assumption is that the *personal* costs to move change across the founder age, independently of the firm. These costs may include life cycle events independent of the firm, such as having children, a working spouse, or owning a home, or psychological changes in the desire to live in new locations related to age.

Empirically in the United States, there are substantial differences in interstate migration rates by age (Molloy et al. 2011). These migrations are driven both by local productivity and by local amenities (Diamond 2016). In general, young individuals are more mobile as they can more easily move and value productivity more. Middle-aged individuals may find it harder to do so due to children, valuing the cost of living and a larger home, or two-body problems with a spouse or partner. Retirees may have a higher propensity to move but are more focused on maximizing amenity value than productivity. All these differences suggest that, among founders, variation by age would predict the probability of migration independent of the firm, serving as an instrument.

Of course, as documented in Azoulay et al. (2020), the age of the founder can also be related to startup performance, which would violate the exclusion restriction. There are several differences between this prior analysis and my approach that are worth emphasizing.

First, Azoulay et al. (2020) focus mainly on the selection *into* high growth entrepreneurship compared with other labor choices and the average age of founders in this group. The key finding is that, in contrast to popular perceptions, the average age of founders within high-growth firms is in their 40 s, a result replicated in my data. In contrast, my study focuses only on high-growth firms, within which relationship of age to performance is less studied.

Second, and related, this prior work mostly speaks to differences in the quality of startups across age, but the purpose of the approach in this paper is to control for quality. Therefore, the key question is whether quality is predicted by age, *conditional* on the controls and fixed effects I include. Indeed, this leads to two placebo tests probing at the instrument's validity. Conditional on controls, performance should not be related to age for startups that do not move, and estimated quality should not be related to age. I run both validations in the results section, showing both statements are true in my data.

I implement a two-stage least square regressions using the predicted probability of moving by age as the instrument rather than using age directly. This probability is estimated through a probit model that includes founding state fixed effects and founding year fixed effects. Using the predicted value of an instrument allows higher precision in binary treatments (Angrist and Pischke 2008, Wooldridge 2010, Galasso and Schankerman 2015). This method does not allow incorporating founding year by founding state fixed effects in the regressions, only state and year fixed effects because (due also to the smaller sample) the probit would drop a substantial number of observations. As reported in Section 5, the instrument is robust, with a first-stage F statistic of 21.

5. Results

5.1. Entrepreneurial Quality Estimates

I estimate entrepreneurial quality through a random forest (Breiman 2001) and predict the likelihood of equity growth from firm observables at founding. I perform variable regularization using the double LASSO (Belloni et al. 2014). *LASSO Controls* is the set of 91 variables selected. Besides being used in a random forest, these variables can also be used as controls directly in an OLS regression, providing a secondary estimating approach. The random forest model predicts *Equity Growth* from *LASSO Controls* using an out-of-sample 10-fold procedure trained only on nonmigrants. Specifically, I split the data into 10 random equal-size groups. I train the random forest on the first nine and predict on the remaining one to estimate the out-of-sample probably of growth. Then, I repeat the same process using each group as the leave out. The predicted estimates can be interpreted as the out-of-sample probability of success when the firm stays at home.

The ROC plot relating the out-of-sample random forest predictions to outcomes is plotted in Figure 2. It has a value of 0.80, implying a *Fit* of 0.6—60% of the variation in outcomes is accounted for by the model.

5.2. Impact of Migration to Silicon Valley on Migrants

5.2.1. Machine Learning Cross-Sectional Estimates on Equity Outcomes. The empirical results begin by considering the benefits of moving to Silicon Valley for firms that move from other regions. In Figure 3, I report the change in odds of *Equity Growth* when comparing migrants and nonmigrants across several statistical models. Standard errors are double clustered by the state of founding and year of founding.

Models (1) and (2) are naïve regressions that estimate the change in probability of growth using no controls.

Figure 2. (Color online) Out-of-Sample ROC Curve

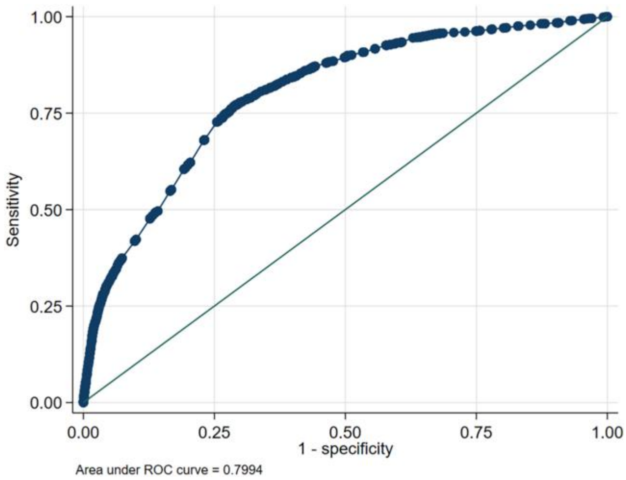

Notes. Reports the ROC curve of the main random forest model, which estimates the out of sample probability that a startup achieves an equity growth outcome if it remains at home. It is trained through a 10-fold cross-validation approach, using only nonmigrants for the training of the model. The ROC value is 0.80, implying 60% of the variation in outcomes is accounted for by the model.

The dependent variable is a binary measure of equity growth divided by the mean value of this measure to allow it to be interpreted as a change in odds. The estimated coefficients are large, 4.8 and 4.61, respectively. According to this estimate, migrants that move to Silicon Valley experience an increase of 461% in the probability of having an equity outcome.

Model (3) presents a regression reflecting a more typical setup for this setting. It includes founding state by founding year fixed effects and controls for whether a firm has a patent, a trademark, or is a corporation. These controls reduce bias, causing the coefficient to lower to a point estimate of 4.19.

Model (4) tries to take the selection issues more seriously by using a coarsened exact matching model (CEM; Iacus et al. 2012). CEM offers the advantage of creating weights for observations based on how likely they are to be in the region of common support. This method makes the coefficient drop further down to 3.7. Relative to the naïve regression, CEM offers a reduction in the treatment effect of 23%.

Models (5) and (6) are the machine learning models proposed in this paper. They incorporate two improvements over the prior approach. First, they account more systematically for the right counterfactual by dividing the outcome by the predicted probability of success at home for the migrants rather than simply the mean of the outcome variable. And second, they control systematically for selection by using a large number of observables and including them as *LASSO Controls*. Model (5) has only the controls. Model (6) is the preferred specification, also including founding state by founding year fixed effects.

The differences are substantial. For Model (6), the coefficient drops to 2.77, a decrease of 25% from the CEM effect and almost half the initial naïve effect. An unreported test using seemingly unrelated regressions

Figure 3. (Color online) Main Effect on Migration to Silicon Valley

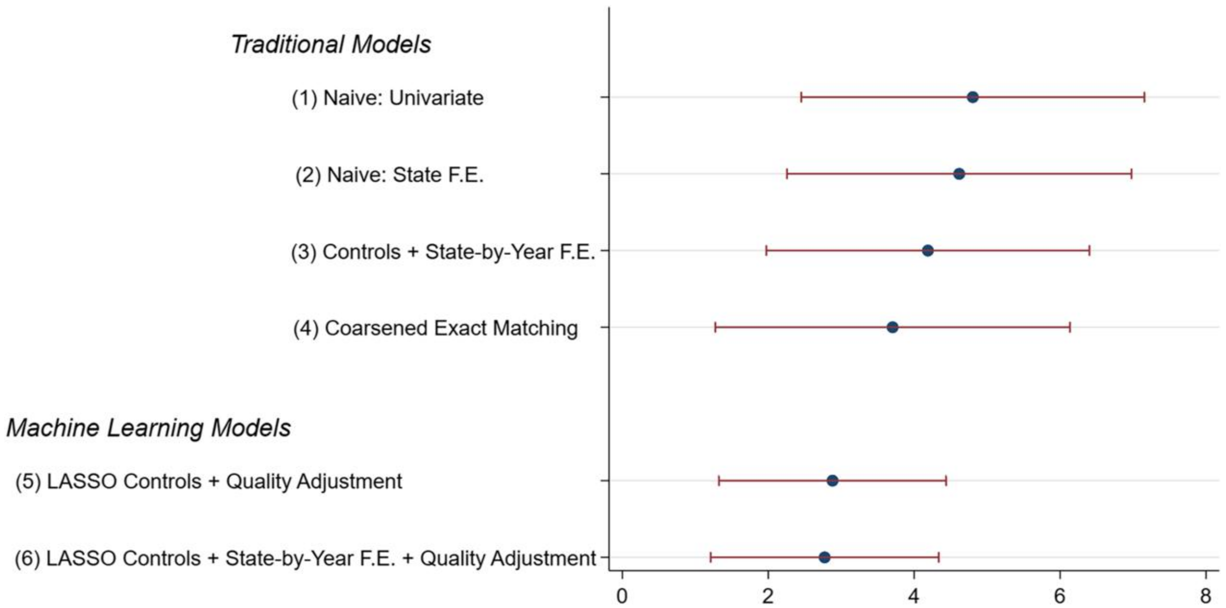

Notes. Reports change in the likelihood of achieving an equity outcome (IPO or acquisition) across several specifications. Models (1) through (4) use observables in traditional ways, including introducing controls and using coarsened exact matching. Models (5) and (6) are the machine learning models proposed in this paper. The machine learning model has lower bias, and a higher level of precision than the traditional models.

Table 2. Oster Tests of Coefficient Stability

	Equity growth	
	(1)	(2)
Moves to Silicon Valley	2.881**	2.773**
	(0.794)	(0.797)
Founding state by founding year fixed effects	No	Yes
N	268,201	268,201
R^2	0.0261	0.0307
Oster estimates		
β^*		2.559
δ (size of unobservables) for $\beta^* = 0$		51.53
δ (size of unobservables) for β^* not significant		14.81

Notes. Reports the tests of coefficient stability as developed in Oster (2019). β^* is the estimated value of the main effect after accounting for unobservables using the usual parameters. δ is the size the unobservables need to be to either be zero or not significant at 95% confidence. Standard errors are clustered by founding year and founding state.
$^*p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

also documents that this coefficient is statistically different from Models (1) through (3).[11] Furthermore, the standard errors of Model (6) are smaller than the traditional models, indicating more precision in the estimate.

Together, these results emphasize the benefits of using machine learning to account flexibly for firm selection compared with more traditional methods. Although there is a large benefit from Silicon Valley, it is smaller than the one suggested by the traditional approach.

5.2.2. Oster Robustness Tests on Equity Growth Estimates. I evaluate the robustness of this estimate through a coefficient stability test (Altonji et al. 2005, Oster 2019) as described in Section 3. The results are reported in Table 2, comparing Models (5) and (6) of Figure 3. The assumption in comparing these two models is that founding state by founding year fixed effects would matter for selection (see Bryan and Guzman (2021) for statistical evidence). After including them, we see that the R^2 increases by 18%, from 0.026 to 0.031. This result validates the value of these fixed effects as useful exemplars of potential omitted variables. The small R^2 values are due to having a binary outcome variable rather than the low fit of the observables. This is evidenced in the high ROC score of the machine learning model (0.80), a measure that evaluates the fit of binary predictions. It is also not a concern in this approach since we are only interested in relative increases in R^2 not its absolute size.

The estimate β^* is a new estimate of the main effect using the standard parameters in Oster (2019) to adjust for unobservables. Its value is 2.56, a point estimate substantially close to the main estimate. Even if unobservables are missing, a reasonable estimate accounting for them is quite similar to the one reported in this paper.

The estimates of δ take unobservables to the limit and ask how large they would have to be to make the effect disappear. I estimate two versions of this parameter: making the point estimate zero and statistically insignificant at the 5% level (assuming the same standard errors as column (2)). Their values are 52 and 15, respectively. Making the coefficient not statistically significant requires unobservables to be 15 times larger than observables used. However, given a ROC of 0.80, this is not realistic. The observables are already accounting for 60% of the variation.

These estimates suggest significant stability of the estimate and that the approach thus far is able to account reasonably well for selection.

5.2.3. More Stringent Definitions of Growth. Next, I consider whether the results could be driven by the incidence of low-value acquisitions rather than larger equity growth outcomes. The alternative hypothesis, in this case, is that given the thick labor and technology markets of Silicon Valley, entrepreneurs may be able to sell off their company simply for the team and not the product (acqui-hire) and investors may be able to force this sale to recoup their original investment even if it does not lead to outsized performance or any gain for the founders.

To study this question, I consider in Table 3 the relationship between startups and more stringent definitions

Table 3. More Stringent Definitions of Growth

	(1) IPO and only acquisitions with any reported value	(2) IPO and only acquisitions with reported value > 50 million USD	(3) IPO and only acquisitions with reported value > 100 million USD	(4) IPO only
Moves to Silicon Valley	4.872**	5.677*	5.778*	5.498*
	(1.773)	(2.860)	(2.910)	(2.866)
Founding state by year fixed effects	Yes	Yes	Yes	Yes
N	254,046	253,386	253,385	256,281
R^2	0.0223	0.0177	0.0177	0.0167

Notes. Reports the change in the odds of different equity outcomes after moving into Silicon Valley. The main specification in the paper includes all acquisitions. This table focuses only on acquisitions with a reported value in SDC Platinum, and columns (2) and (3) impose more stringent hurdles on the valuation. Standard errors are clustered by founding year and founding state.
$^*p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

of equity outcomes by taking advantage of the reported acquisition value in SDC Platinum. Only 30% of the firms have an acquisition value reported, with larger deals being more likely to report. I consider four potential specifications: dropping all acquisitions with no reported value, dropping all acquisitions with no reported value or reported value below $50 million, dropping all acquisitions with no reported value or reported value below $50 million, and keeping only IPOs. The estimates have broader confidence intervals but are, if anything, larger.

These results do not suggest acqui-hires or low-value acquisitions are not more common for movers to Silicon Valley. Rather, the relative increase in performance from migration accrues across the range of outcomes broadly, equally improving low- and high-value ones.

5.2.4. Other Outcomes. In Table 4, I expand the analysis of the benefits of migration to Silicon Valley to other outcomes besides equity growth. I report three models: a naïve model, a model including controls for founding patent, founding trademark, whether the firm is a corporation, and founding year and founding state fixed effects, and a machine learning model (Models (1), (3), and (6) of Figure 3). For each outcome, I rerun the LASSO procedure to match the best observables and rerun the random forest model to estimate the expected performance of migrants at home for this outcome.

Column (1) repeats the main dependent variable reported in Figure 3.

Column (2) focuses on venture capital financing and column (3) on trademarks (as a proxy for product introductions). The effects are dramatic. For venture capital, the naïve approach estimates a 376% increase in the likelihood of raising venture capital on moving, but this number drops to almost half (218%) when using the machine learning approach. The coefficient for trademarks is similarly halved.

The result for filing patents, in column (4), is the only one showing a different pattern. Although it decreases with controls, it increases when using the machine learning approach. Moving to Silicon Valley increases the probability of getting a patent by 162%.

Together, these results report a significant benefit of moving to Silicon Valley across a range of outcomes for migrants.

5.2.5. Panel Estimates Within Movers. In Figure 4, I report a second identification strategy using panel data and the timing of migration, as overviewed in Section 4.3.[12] The figure plots the difference in the dependent variable of this regression for four outcomes from two years before moving to six years after moving. To provide a richer assessment of the changes from migration, I also focus on more continuous variation to incorporate the intensive margin by using continuous outcomes of venture capital financing, number of trademarks, and number of patents. Unreported results with binary outcome measures are similar. Standard errors are clustered by the firm and founding year level.

Three key results are relevant in this figure. First, there is no obvious pretrend in any of the specifications. The level of outcomes for the preperiod is not significantly different from zero, and it is not different from migrations to other destinations.

Second, there are large treatment effects from migration. All variables increase meaningfully initially and

Table 4. Startup Performance on Migration to Silicon Valley: Other Outcomes

	(1) Equity growth	(2) Gets venture capital	(3) Gets trademark	(4) Gets patent
Panel A: Naïve				
Moves to Silicon Valley	4.657***	3.762***	0.989***	1.435***
	(1.163)	(0.487)	(0.126)	(0.197)
Naive p(outcome)	0.0134	0.0150	0.0665	0.0483
Panel B: Controls + fixed effects approach				
Moves to Silicon Valley	4.060***	2.931***	0.592***	0.827**
	(1.095)	(0.393)	(0.0792)	(0.259)
Founding state by founding year fixed effects	Yes	Yes	Yes	Yes
Naive p(outcome)	0.0134	0.0150	0.0665	0.0483
Panel C: Machine learning approach				
Moves to Silicon Valley	2.773***	2.181***	0.472**	1.628***
	(0.797)	(0.471)	(0.218)	(0.310)
Founding state by founding year fixed effects	Yes	Yes	Yes	Yes
N	268,201	256,281	256,281	256,281
Random forest p(outcome)	0.0188	0.0225	0.0790	0.0606

Notes. Reports the OLS coefficient for three different models per dependent variable. The preferred model is Panel C, the machine learning approach. Standard errors clustered at the founding state and founding year level.
*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

Figure 4. (Color online) Panel Data Regressions on the Impact of Migration to Silicon Valley

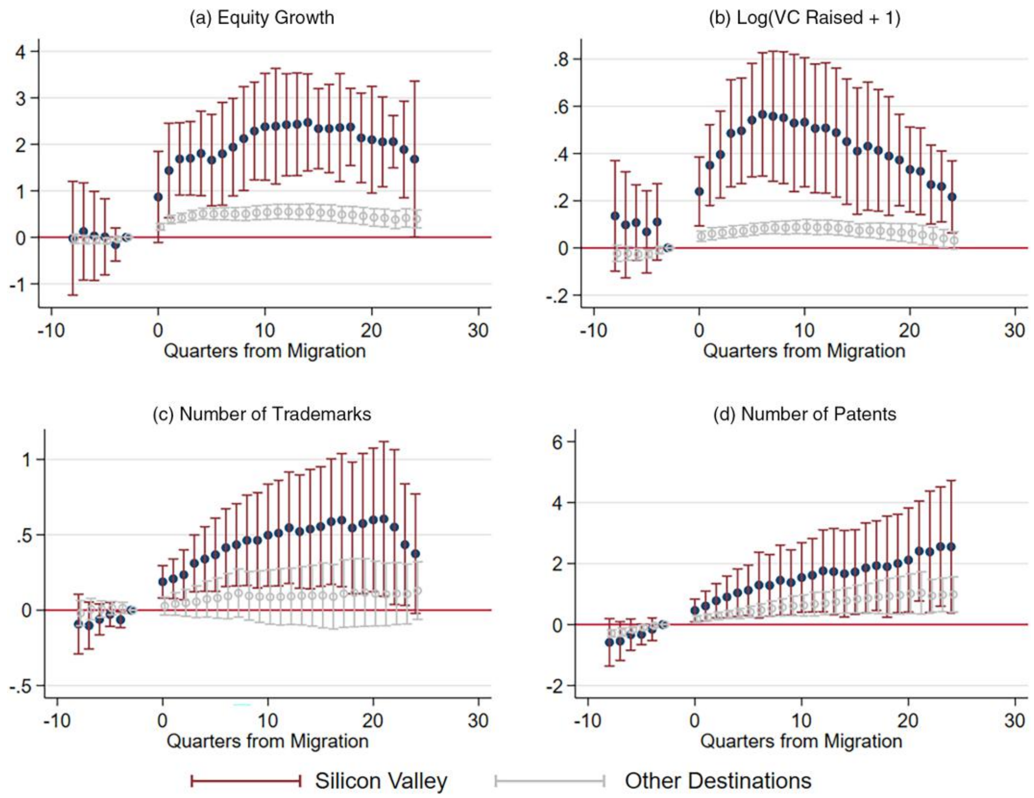

Notes. Reports the coefficients of panel OLS regressions by quarter with fixed effects for startup, quarter of age, quarter of migration, and founding year. The patent regression is too noisy with quarterly fixed effects for age and migration, and I use yearly instead. Bands indicate 95% confidence interval. The baseline period is three quarters before migration, and I exclude the following two quarters to account for noise in the precise timing of migration. Standard errors are clustered by startup . and founding year.

then increase even further as time passes. There are also increases for equity growth, venture capital, and patents when moving to other destinations, but they are substantially smaller and statistically different from the migration to Silicon Valley. This suggests, again, that Silicon Valley is not only quite valuable but distinctly valuable in migrations.

Third, there are interesting differences in the evolution of these effects over time. The benefits of venture capital increase the fastest and then drop after about three years at a destination. The benefits on equity outcomes are slower, but still appear to peak after a few years. Conversely, the benefits on innovation (patents) are slower and do not appear to see their peak within the sample window. This suggests interesting heterogeneity in the time that one needs to be local to absorb agglomeration benefits: migrants appear to more quickly be able to access the VC market and the equity market, than fundamentally shaping their innovation intensity.

5.2.6. Instrumental Variables Estimates Using Age. I next report the instrumental variable regressions using founder age as an instrument.

The distribution of age across founders is plotted in Figure A3 in the online appendix. Consistent with the systematic analysis of founder age and high growth startups in Azoulay et al. (2020), it shows an increasing age distribution from their early 20 s, peaking at about 42 years old.

Figure 5 plots a range of correlations between age and migration. Figure 5(a) is simply the average migration rate by decade of age at founding. We see an intuitive pattern consistent with overall migration patterns in the United States. Youngest individuals are the most likely to move. There is a decreasing trend up to founders in their 40 s and a slight increase after 50, possibly related to migration closer to retirement age (as documented for the United States at large by Chen and Rosenthal (2008)). Figure 5(b) is the same relationship in a binned scatterplot. Figure 5(c) and (d), reports the correlation of the predicted probability of moving based on age and actual migration. This predicted probability of moving is estimated through a probit model (reported in Table A3 in the online appendix) using a cubic age function and founding state and founding year fixed effects. Figure 5(c) is the raw correlation, and Figure 5(d) includes all controls included in the main regression. The correlation

Figure 5. (Color online) Graphical First Stage

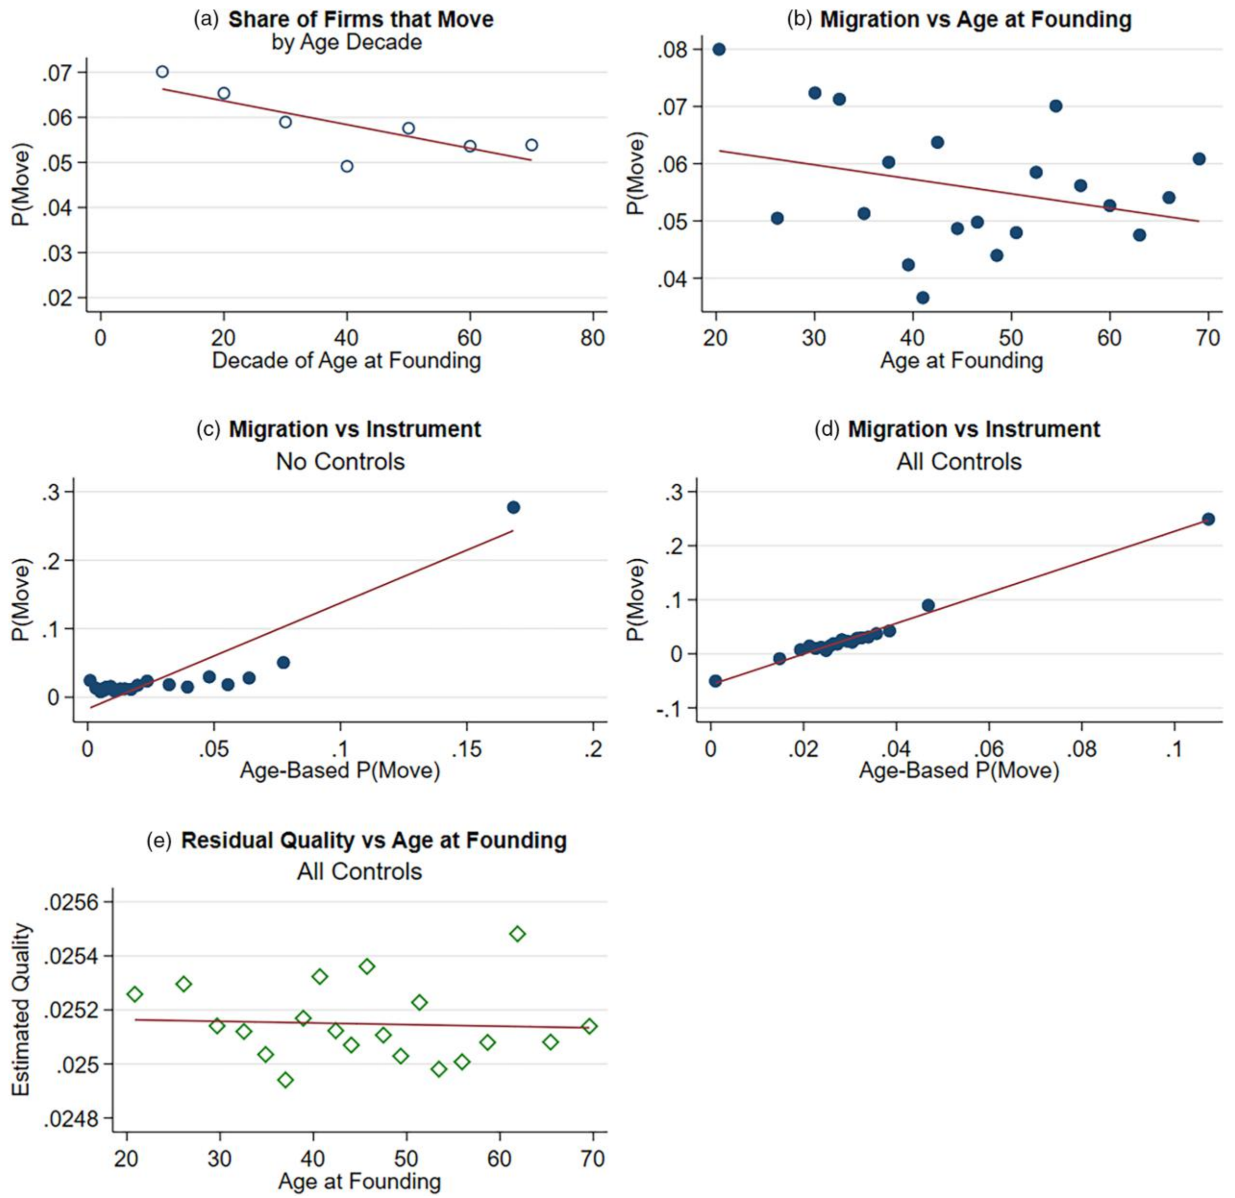

Notes. (a)–(d) Several versions of the first-stage regression in the instrumental variable model. The instrument is the age of the founder at founding, retrieved through a web-scraping algorithm from Ancestry.com. (a) Basic scatterplot by decade. (b) Binned scatterplot. (c) Binned scatterplot with the instrument used—the predicted probability of moving based on age. (d) Series of controls used in the analysis: LASSO controls, founding year fixed effects, and founding state fixed effects. (e) Exclusion restriction by assessing how does age predict quality conditional on controls, finding no relationship.

between the instrument and actual migration appears quite meaningful. Finally, Figure 5(e) is a falsification test considering age and estimated quality at founding for all firms conditional on these controls. Reassuringly, there is no relationship between the two.

The main results are presented in Table 5. Column (1) replicates the main result using the machine learning methodology of Table 4 within the smaller sample used here. The coefficient is 2.99, close to the main estimate.

Column (2) is the first stage regression on moving to Silicon Valley. The instrument predicts moving well. The *F* statistic of this regression (reported in column (4)) is reasonably large, at 21.38.

Column (3) is a preliminary estimate without controls, and column (4) is the main estimate. It has a coefficient of 4.489, which is statistically significant at the 1% level. This point estimate is higher, but not statistically different, from the main estimate of column (1). The most likely source of these differences in estimates may be the conceptual difference between the treatment effect estimated through the whole sample and the local average treatment effect created by the instrument's variation.

Table 5. Instrumental Variables Estimates of the Impact of Migration on Startup Performance

	(1) OLS *Equity Growth*	(2) OLS *Moves to S.V.* (first stage)	(3) 2SLS *Equity Growth*	(4) 2SLS *Equity Growth*	(5) OLS *Equity Growth* (nonmovers only)	(6) 2SLS *Equity Growth*
Moves to Silicon Valley	2.989**		4.234***	4.489***		
	(1.318)		(1.155)	(1.187)		
Age-based p(move)		1.526***			−0.105	
		(0.330)			(2.562)	
Move outside California						0.0728
						(0.869)
Controls	Yes	Yes	No	Yes	Yes	Yes
Founding year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Founding state fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
First-stage F statistic			20.96	21.38		71.74
N	22,534	22,524	22,524	22,524	22,425	22,943
R^2	0.0462	0.166	0.00154	0.0344	0.0436	0.0318

Notes. Reports the instrumental variables regression using the age of founder as the instrument as explained in Section 3. The age of founder is estimated through an algorithm scraping Ancestry.com. Standard errors clustered by founding state and founding year.
*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

Columns (5) is a placebo test considering whether age can be correlated to performance independent of migration (a test of the exclusion restriction). It reports the relationship of the age instrument to performance only for nonmigrant firms. The coefficient is small and not significant. Within my sample of nonmigrants, age does not seem to predict differences in outcomes after introducing all controls.

Column (6) considers migrations to other destinations outside California. The estimate, in this case, is small and not statistically different from zero.

5.3. Migration to Other Destinations

I next proceed to open the analysis beyond Silicon Valley to other destination regions. In Figure 6, I report coefficients replicating the machine learning approach for the other 14 most common destinations in the data and compare them to the coefficient reported for Silicon Valley. The differences are dramatic. They emphasize wide heterogeneity across destinations in the performance of firms after moving, and ample benefits from moving to Silicon Valley compared with other destinations.

Figure 6(a) considers the impact of migration on equity growth. Except for the Denver-Aurora, Colorado, MSA, all other destinations have a meaningfully lower estimated effect than Silicon Valley. Although migrants to Silicon Valley showed a relative increase in equity outcomes of 277%, migrants to Boston or New York City see increases of 143% and 73%, respectively. Denver is slightly higher than Silicon Valley at 308%, although not statistically different than Silicon Valley. After Denver, the next largest effect is the Los Angeles area, with 181%.

Figure 6(b) considers the impact of migration on venture capital financing. The differences are even starker. Migrants to Silicon Valley increase their likelihood of getting financing by 218%. In contrast, migrants to Boston, Massachusetts, the second-largest estimated effect and the second-largest venture capital market, only increased by 89%. Most destination cities have an estimate that is closer to zero. Although there are some migration benefits on venture capital they appear to be dramatically larger in Silicon Valley, with the rest being quite smaller.

Figure 6(c) considers patenting as a measure of innovative outcomes. Once again, the point estimate is higher for Silicon Valley than other locations. The largest estimates, after Silicon Valley, are for places that have large health sciences clusters, such as the Washington, DC, area or the Boston area, as well some more idiosyncratic cases of Denver and Chicago. It appears that Silicon Valley has a meaningfully different effect on innovation outcomes.

Finally, Figure 6(d) studies the incidence of trademarks as a proxy for introducing new products. The coefficients are smaller for all cities than for all other outcomes. Furthermore, Silicon Valley has a large effect, but it does not appear substantially different from other large metro areas. Locations such as Boston, Chicago, or Denver also have large effects and the effect of Washington, DC, is about 30% larger, even if not statistically significant in these differences. Although Silicon Valley's agglomeration benefits appear meaningfully better than other regions in the other outcomes, they are only as good (but not better) in the introduction of products.

Together, these results paint a distinct picture of the heterogeneity in agglomeration benefits across regions and the importance of Silicon Valley within these. It also validates the value of focusing on Silicon Valley for the empirical question of this paper as the quintessential location for startups and the agglomeration benefits they can capture.

Figure 6. (Color online) Other Destination

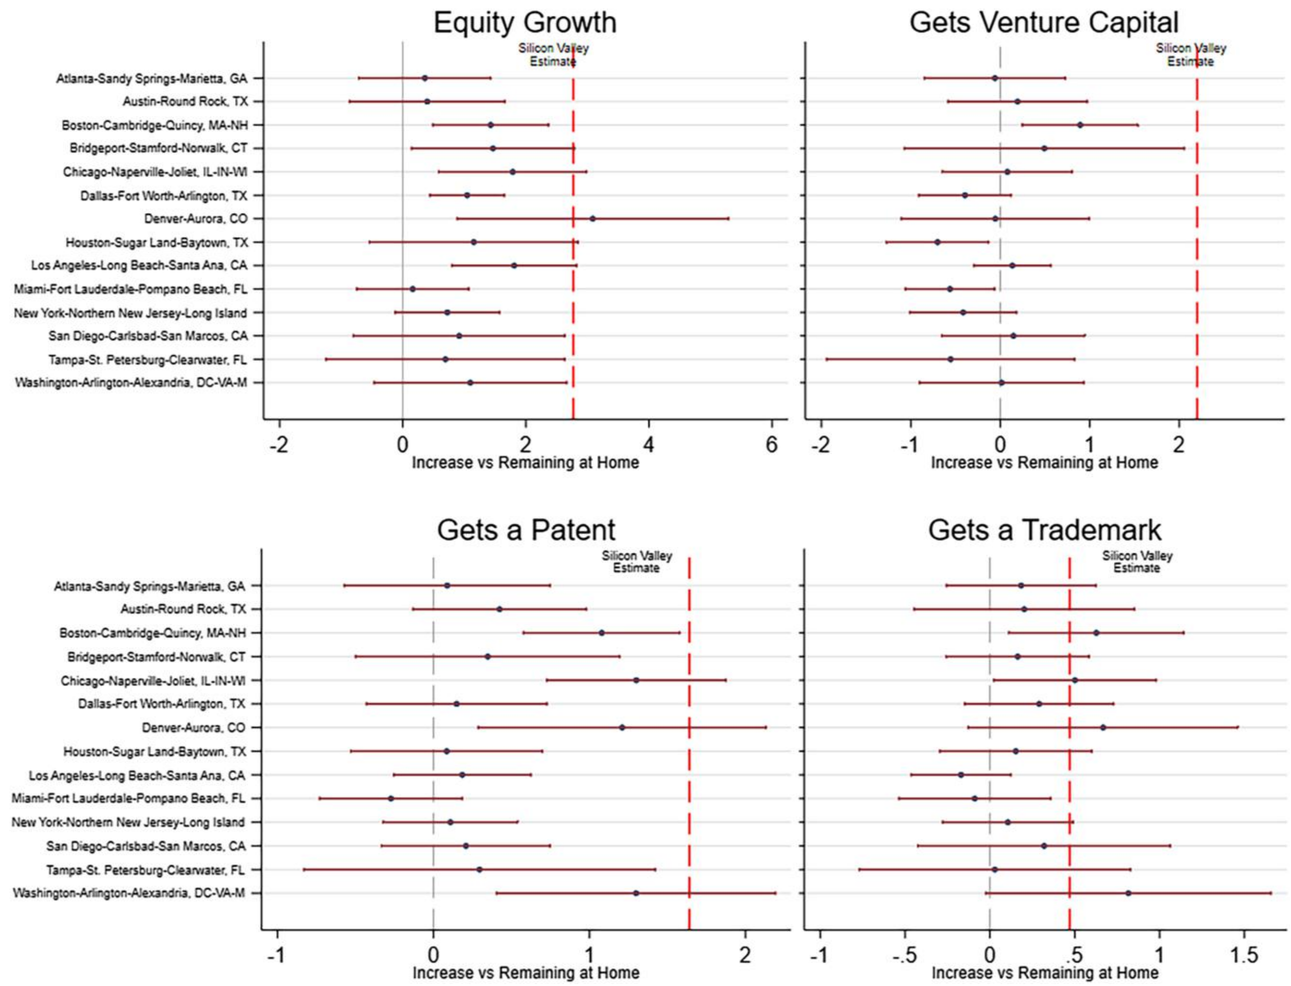

Notes. Reports the coefficient of the main machine learning specification, estimated for each of the top 14 destination metro areas (excluding Silicon Valley), and compares the effect to the coefficient of Silicon Valley. For each metro area and outcome, a new set of LASSO controls is estimated and a different counterfactual of the quality if the firm had stayed at home is used. All regressions include founding state by founding year fixed effects. Band indicates the 95% confidence interval. Standard errors are double clustered at the founding state and founding year level.

5.4. Mechanisms

Finally, I study the mechanisms. The arguments thus far have stated migration leads to a positive performance due to the spatial agglomeration of resources in tech clusters (Kerr and Robert-Nicoud 2020, Moretti 2021), even though agglomeration is not the only way migration could lead to higher performance for migrants. However, agglomeration has a distinct prediction regarding who benefits from migration. It implies that the benefits of migration should vary for startup migrants, not based simply on where they move to, but on the relative improvement in the quality of the local economic cluster they experience in this migration compared with the place they leave.

5.4.1. Relative Gains from Migrants to Silicon Valley. Table 6 studies relative improvements in the entrepreneurial ecosystem for migrants to Silicon Valley. To do so, I introduce new measures of the relative increase in a region's entrepreneurial ecosystem, estimated as the difference in a measure of ecosystem caliber between the startup's founding region and Silicon Valley at the time of its founding. All nonmigrants experience an increase of zero in these measures. I report OLS regressions on equity outcomes, including founding state and founding year fixed effects. These fixed effects allow identification to come from the time-varying nature of the origin ecosystem, not its permanent characteristics nor yearly fluctuations in Silicon Valley's quality or attractiveness. The summary statistics of these measures are included in Table A4 in the online appendix. Their scales are adjusted to have comparable magnitudes. Standard errors are double clustered by founding year and founding state.

Column (1) replicates the main result using founding state and founding year fixed effects rather than founding state by founding year (which cannot be included

Table 6. Ecosystem Changes for Migrants to Silicon Valley

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Moves to Silicon Valley	2.785***	1.311	1.244	1.951*	2.530***	2.401***	1.430
	(0.800)	(0.962)	(0.899)	(1.109)	(0.857)	(0.853)	(0.964)
Main mechanisms							
Δ VC Per Capita		0.0844***					−0.0361
		(0.0188)					(0.0766)
Δ Patents Per Capita			0.0700***				0.114*
			(0.0117)				(0.0628)
Δ Entrepreneurship Ecosystem				0.0707			−0.0505
				(0.0432)			(0.0587)
Placebos							
Δ Personal Tax Rate at 95th Percentile					−0.0204		−0.0643
					(0.0604)		(0.0566)
Δ Sunshine						0.0178	0.0152
						(0.0284)	(0.0284)
LASSO controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Founding state fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Founding year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	268,201	265,073	265,073	264,589	265,073	265,073	264,589
R^2	0.0282	0.0286	0.0287	0.0285	0.0283	0.0283	0.0288

Notes. Reports OLS regressions replicating the main specification, but including additional variables that measure the change in ecosystem quality for migrants (the change for nonmigrants is zero). Venture capital per capita is the total number of dollars per thousand of population. Patents per capita is total patents per thousand population. Entrepreneurship ecosystem is the quality adjusted quantity measure of the Startup Cartography Project, per 10,000 population (done to have all variables in similar scales). Personal tax rate is estimated at the 95th level by Moretti and Wilson (2017). Percent sunshine is the percent of days that have sunshine in this city, on average. Standard errors are clustered by founding year and founding state.

*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

as state-year variation is the identifying variation). The coefficient is quite similar to the main effect in Table 6.

Columns (2) through (4) report different measures of the strength of the entrepreneurial ecosystem and its relationship to performance. Column (2) measures the ecosystem through the venture capital dollars per capita invested in a region in that year. The coefficient is positive, with a value of 0.084 and significant. Startups coming from locations with less venture capital supply see a larger benefit from migration. These differences are also economically meaningful. It predicts that, although a startup facing the median increase in venture capital per capita sees an increase of 215% on its probability of achieving an equity outcome after migration, one experiencing an increase at the 75th percentile of migrants sees this probability rise to 257% and one at the 90th percentile a dramatic 567%. The main effect of moving to Silicon Valley, in contrast, reduces to about half the magnitude and is not significant.

Column (3) performs the same analysis but uses patents per capita rather than venture capital per capita to measure the ecosystem. The results are very similar. The coefficient is 0.07 and significant. Under this measure, migrants experiencing the median gain in their ecosystem's patents per capita are predicted to see an increase of 201% in their probability of growth. This value rises to 236% for those in the 75th percentile and 621% for those in the 90th percentile. The consistency between the two measures is expected. Patenting per capita and venture capital per capita are highly correlated, so either measure is possibly proxying for a similar type of origin region.

Column (4) uses a broader measure of entrepreneurial ecosystem strength, the level of quality-adjusted entrepreneurship per capita taken from the Startup Cartography Project (Andrews et al. 2020). This is a measure of total entrepreneurship in a city but is less focused directly on the inputs of high technology venture-backed startups. The coefficient, in this case, seems noisier, with a p value of 0.11. The change on the main effect of migration is also more muted. Although this main effect's point estimate had dropped from 2.79 to 1.3 and 1.2 when we introduced venture capital or patenting, it drops only to 1.95. Although moving from locations with a low level of quality-adjusted entrepreneurship to a high level of quality-adjusted entrepreneurship is beneficial, these benefits appear more muted than when we use other measures of ecosystem quality.

Columns (5) and (6) are instead placebo tests, introducing regional amenities that are desirable utilities for migrants but not intuitively related to the performance of the firm. These include the level of personal income taxes and the amount of sunshine, both of which are shown in Bryan and Guzman (2021) to predict migration. The effects for amenities are quite different: they are best interpreted as zero.

The way relative differences in ecosystem quality predict startup performance while amenities do not is consistent with an agglomeration mechanism, where proximity to valuable economic inputs is the core to migration benefits.

I analyze this more in-depth in column (7) by evaluating all ecosystem measures simultaneously. Doing so allows performing a "horse race" of sorts between three potential mechanisms and their benefits for startups: venture capital, local idea generation, and the overall ecosystem and presence of peers. Although these are undoubtedly correlated, they also imply meaningfully different ways a regional ecosystem may improve startups' performance and different versions of the key resource that agglomeration in Silicon Valley provides startups.

Only the patent coefficient remains positive and statistically significant in this regression. The coefficient for venture capital and the local ecosystem is smaller, negative, and not significant. These differences suggest that, within these options, the core benefit of migration into Silicon Valley may be more related to local ideas, knowledge flows and innovation inputs, than capital inputs or the presence of many local peers.

5.4.2. Within Analysis with All Movers. Table 7 expands on the ideas in Table 6 but takes a different approach by considering movers to all destinations and their relative changes in performance from migration. Analyzing within all movers allows including origin-by-year and destination-by-year fixed effects. Therefore, this analysis controls for any fixed aspects in both origin and destination regions at a point in time, using only the relative differences between regions as the identifying variation. Because this analysis is within movers, the table does not include the main effect of moving.

The parallels to Table 6 are striking. Columns (1) and (2) report very similar coefficients to those in Table 6 for venture capital and patenting per capita. Column (3), studying the ecosystem effect, shows a larger and more precise effect than for Silicon Valley, significant at the 5% level. An improvement in the local entrepreneurship level also predicts increased benefits from migration on performance. Columns (4) and (5) report the placebo effects for changes in amenities—personal income taxes and the amount of sunshine—which are again noisy and not significant. Overall, the relationship between regional improvements and performance is similar (though more precisely estimated) for all regions compared with only Silicon Valley.

Column (6) brings all variables together. As in the case of Silicon Valley, the coefficients for venture capital and local entrepreneurship are smaller or negative, and noisy. In contrast, the coefficient for patenting per capita strengthens.

The similarity between the results for all migrants and those to Silicon Valley suggests stability in how ecosystem improvements lead to higher performance for U.S. startups during migration to tech clusters. They also indicate no mythology of Silicon Valley. A rather traditional agglomeration story explains the differences in treatment effects: Silicon Valley simply has the highest

Table 7. Gains in Ecosystem and Performance Within Movers

	(1)	(2)	(3)	(4)	(5)	(6)
Main mechanisms						
Δ VC Per Capita	0.0792***					−0.122
	(0.0275)					(0.110)
Δ Patents Per Capita		0.0648***				0.153*
		(0.0204)				(0.0781)
Δ (Quality-adjusted) Entrepreneurship per Capita			0.114**			0.0154
			(0.0432)			(0.0746)
Placebos						
Δ Personal Tax Rate				0.0876		0.0139
				(0.0645)		(0.0858)
Δ Sunshine					−0.0126	−0.0355
					(0.0338)	(0.0435)
Founding state by founding year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Destination CBSA by founding year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	5,825	5,825	5,503	5,825	5,825	5,503
R^2	0.264	0.264	0.257	0.263	0.263	0.258

Notes. Reports OLS regressions within movers, including additional variables that measure the change in ecosystem quality for migrants (the change for nonmigrants is zero) and fixed effects for both origin region by year and destination region by year. Identification comes from the pair-wise differences in regions. Venture capital per capita is the total number of dollars per thousand of population. Patents per capita is total patents per thousand population. Entrepreneurship ecosystem is the quality adjusted quantity measure of the Startup Cartography Project, per 10,000 population (done to have all variables in similar scales). Personal tax rate is estimated at the 95th level by Moretti and Wilson (2017). Percent sunshine is the percent of days that have sunshine in this city, on average. Standard errors are clustered by founding year and founding state.

*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

level ecosystem under any measure—venture capital, patenting, and startup formation—which then translates into a large treatment effect from moving to it.

Although different ecosystem inputs tend to correlate across regions, the analysis suggests the local level of innovation is particularly predictive of the benefits of migration on the performance of firms. When companies move, they benefit by accessing a location ripe with more and better ideas rather than simply more capital. Several agglomeration theories support this perspective.

An important avenue for future work would be to consider this result across multiple theories linking local innovation to entrepreneurship outcomes. For example, do the benefits of local innovation stem from the specialization of innovation and its local returns to scale, such as in Marshall (1890), or is it instead the diversity of innovations creating distinct ideas as in Jacobs (1970)?

Furthermore, the results also do not support other "second-order" mechanisms unrelated to the local economic agglomeration effects. In particular, there is no effect from founders accessing better amenities or having a higher share of personal income migration, nor is there evidence of a distinct fixed effect of Silicon Valley for migrants that can be measured independently from the actual provision of resources. Whether these regional elements develop each other over time is a valuable area for future study.

6. Conclusion

This paper studied how migration affects the performance of startups and presented evidence that migration leads to higher performance by improving a startup's ecosystem and its access to locally agglomerated resources. Within a few potential mechanisms, the role of the local innovation ecosystem appears more critical than local financing or the incidence of local startup peers. One region in the United States, Silicon Valley, is an outlier in the strength of its ecosystem, independent of the measure used. Accordingly, migrations to Silicon Valley lead to a substantially higher treatment effect than migrations to other destinations, particularly for migrants that leave low-performing ecosystems. In contrast, there is no evidence for additional benefits related to taxes, weather, or the act of moving, on the performance of migrants.

These results call for several avenues of future research.

First, empirically, the benefits of migration to many regions are not positive, in which case migration did not improve performance on average for migrants. Possibly, this is because amenities also drive migrations (Bryan and Guzman 2021)—These founders might not have moved to seek better performance after all. This result highlights that startup choices are done by founders and that studying together the startup and the founder in entrepreneurial strategy and policy is important. For example, although an extensive literature since Marshall (1890) argues that founders are well matched to the elements of their regions, the results of this paper suggest other areas may be even better suited for some startups, but they do not move due to founder costs.

Second, moving to mechanisms, the results in this paper focus on how ecosystem improvements predict performance. This allows concluding that it is some version of agglomeration driving the results and that, in a horse race between different measures of the ecosystem, the local innovation level appears to be most important. However, much more can be done to understand how do these pieces of the ecosystem relate to each other in developing the grounds that support startup performance. For example, one possibility from the results in this paper is that venture capital matters more than the results allow but it does so in motivating individuals to create high-level innovations rather than direct financing of existing firms. In general, the equilibrium relationship of the variables studied is likely to be more complex than one can appreciate thus far.

Third, understanding why Silicon Valley is such an outlier as a startup ecosystem, and therefore creates outsized performance for startups that move to it, is key to replicate innovation-driven growth across the United States. An ample line of work has documented the skewness in entrepreneurial inputs across cities (Moretti 2012) and theorized that it drives differences in performance. However, this is the first paper introducing a cleaner counterfactual by considering only firms born outside a city and their performance improvements after moving and providing evidence that differences in the spatial distribution of resources is itself a key mechanism. Follow-up work should take advantage of the results in this paper to ask more substantively how to create the conditions of Silicon Valley elsewhere to create and support other high-performing entrepreneurial ecosystems.

Fourth, bringing it more to the present, there is increasing commentary in the press that Silicon Valley's unique status has reduced over time and other clusters have emerged (Russell 2014). For example, recently, large firms like Hewlett-Packard and Tesla moved to out of the Bay Area to Texas. In general, cities such as Austin and Dallas appear to be some of the most attractive for startup migrants (Bryan and Guzman 2021). However, the results in this paper suggest the value of Silicon Valley has not reduced. It is true that the economic role of some Southern cities, such as Atlanta or those in Texas, has increased over time, and that in consequence also their total amount of innovation and venture capital. However, the United States has also been in a long venture capital boom. When considered in relative terms, the share of venture capital going to Silicon Valley has remained at about 50% for more than a decade (PwC/CB Insights 2022). This paper shows benefits of location are driven by relative differences in the amount of inputs, and there is no evidence thus far that this relative difference has changed in a generalized way.

Finally, a different dynamic that could affect the role of Silicon Valley is the changing importance of location itself on startup performance, powered by remote work technologies. Of particular note is how the ongoing COVID-19 pandemic has disrupted the locational distribution of teams and startups in the United States, creating an increasing number of remote teams (Brynjolfsson et al. 2020). Although venture capital has traditionally been very local, the number of remote investments is increasing because of video conferencing technologies and crowdfunding platforms (Han et al. 2022). One possibility is that this trend will weaken the importance of Silicon Valley over time. Many individuals will now have relatively more access to talent and investors without physically moving to this location, making physical colocation less important. In this scenario, the patterns uncovered in this paper are about to change significantly. An opposite possibility is that remote work strengthens the role of Silicon Valley by being a complement to the success and scaling of startups in this region. In this scenario, startups in Silicon Valley can surpass limits in the local employment market, and investors in more remote areas can also invest in Silicon Valley startups of higher quality than their local ones. If universities such as Stanford and Berkeley have high levels of innovation, marginal startups that were originally not competitive in the local financing and labor market in the Bay Area may be competitive in the global one, being induced to enter. Then, the patterns in this paper will only strengthen in the future. Indeed, this form of geographic polarization was the result of the first wave of crowdfunding (Agrawal et al. 2015): Crowdfunding strengthened the hand of top startup locations as remote investors in smaller cities went for startups in startup hubs more so than Silicon Valley investors seeking to invest remotely. The impact of COVID-19 and remote work on the geography of entrepreneurship is an important and quickly evolving area of research.

Acknowledgments

The author thanks Scott Stern, Pierre Azoulay, Christian Catalini, and Olenka Kacperczyk for advice and support of this project; Ashish Arora, Sharon Belenzon, Kevin Bryan, Bo Cowgill, J. P. Eggers, Cathy Fazio, Joshua Gans, Deepak Hedge, David Hsu, Bill Kerr, Daniel Keum, Ramana Nanda, Rob Seamans, Dan Wang, and Sam Zyontz, as well participants in seminars at New York University, London Business School, Columbia, Kellogg, MIT, Harvard Business School, Duke, Toronto, and University of California San Diego and the Consortium of Cooperation and Competitiveness (CCC) and Roundtable for Engineering Entrepreneurship Research (REER). He also thanks the editor, Toby Stuart, the associate editor, and three anonymous referees.

Endnotes

[1] Some core drivers of agglomeration include culture (Saxenian 1994, Florida 2002); distance and infrastructure (Belenzon and Schankerman 2013, Kerr and Kominers 2015, Agrawal et al. 2017); built space and serendipitous interaction (Roche 2019, Roche et al. 2020); social networks (Sorenson 2018); financial liquidity (Stuart and Sorenson 2003a, Chen et al. 2010); Marshallian agglomerations (Delgado et al. 2010, Chatterji et al. 2014, Andrews et al. 2020); laws, regulations, and political borders (Marx et al. 2009, Singh and Marx 2013); and immigration (Kerr 2008, Balsmeier et al. 2020).

[2] An important exception is Conti and Guzman (2023), who use migration of Israeli firms to perform a comparative analysis between the United States and the Israeli entrepreneurial ecosystems.

[3] For example, the National Science Foundation's Directorate for Technology, Innovation, and Partnerships (TIP) recently launched the Regional Innovation Engines initiative, which invests $1.5 billion dollars over 10 years in creating regional innovation clusters outside of the main tech regions in the United States (Gianchandani 2022).

[4] A map of all states in the data are found in Figure A1 of the online appendix. Although the Startup Cartography Project has data for 49 states and Washington, DC, this analysis omits 12 states in which we could not adequately separate local state addresses from headquarters addresses. Further details are available in Bryan and Guzman (2021).

[5] The Delaware General Corporate Law is the best understood corporate law in the United States, with a long cannon of decisions that are useful in creating predictable contracts even in cases of significant complexity. Delaware also has an advanced institutional setup to deal with corporate arbitration, including its highly reputed Court of the Chancery. Furthermore, Delaware's decisions and legal framework are generally regarded as probusiness.

[6] On average, these fees amount to a few thousand dollars per year.

[7] The matching approach builds on the existing approaches of Kerr and Fu (2008) and Balasubramanian and Sivadasan (2009). Further details are available in the supplementary materials of Guzman and Stern (2015, 2020) and Andrews et al. (2020).

[8] The only exception are cases when a local firm already has the name of the incoming startup migrant. In these cases, the Delaware migrant adds its origin state to its name to differentiate, such as "Microsoft from Washington." This is rare and would not cause my approach to mistakenly observe a migration of a Delaware firm that did not move but, rather, would only underestimate it.

[9] For migrant firms, patent and trademark outcomes are only included if the patent is assigned to a firm in the destination state to avoid potential biases occurring through innovative activity occurring at home.

[10] This statistic holds a close parallel to Somers' D, one of the most common measures of fit for binary predictive models.

[11] It is not possible to apply the test in Model (4) due to differences in the observation weights created by the CEM approach.

[12] The patent regression does not converge when using quarter fixed effects for age. I use annual age fixed effects only in this case.

References

Agrawal A, Catalini C, Goldfarb A (2015) Crowdfunding: Geography, social networks, and the timing of investment decisions. *J. Econom. Management Strategy* 24(2):253–274.

Agrawal A, Galasso A, Oettl A (2017) Roads and innovation. *Rev. Econom. Statist.* 99(3):417–434.

Albouy D (2016) What are cities worth? Land rents, local productivity, and the total value of amenities. *Rev. Econom. Statist.* 98(3):477–487.

Almeida P, Kogut B (1999) Localization of knowledge and the mobility of engineers in regional networks. *Management Sci.* 45(7):905–917.

Altonji JG, Elder TE, Taber CR (2005) Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *J. Political Econom.* 113(1):151–184.

Andrews RJ, Fazio C, Guzman J, Liu Y, Stern S (2020) The startup cartography project: Measuring and mapping entrepreneurial ecosystems. Working paper, Info We Trust.

Angrist JD, Pischke JS (2008) *Mostly Harmless Econometrics* (Princeton University Press, Princeton, NJ).

Assenova VA, Sorenson O (2017) Legitimacy and the benefits of firm formalization. *Organ. Sci.* 28(5):804–818.

Audretsch DB, Feldman MP (1996) R&D spillovers and the geography of innovation and production. *Amer. Econom. Rev.* 86(3): 630–640.

Azoulay P, Jones BF, Kim JD, Miranda J (2020) Age and high-growth entrepreneurship. *Amer. Econom. Rev. Insights* 2(1):65–82.

Balasubramanian N, Sivadasan J (2009) NBER patent data-BR bridge: User guide and technical documentation. Working paper, Syracuse University, Syracuse, NY.

Balsmeier B, Fleming L, Marx M, Shin SR (2020) Skilled human capital and high-growth entrepreneurship: Evidence from inventor inflows. Technical report, National Bureau of Economic Research, Cambridge, MA.

Belenzon S, Schankerman M (2013) Spreading the word: Geography, policy, and knowledge spillovers. *Rev. Econom. Statist.* 95(3):884–903.

Belloni A, Chernozhukov V, Hansen C (2014) High-dimensional methods and inference on structural and treatment effects. *J. Econom. Perspective* 28(2):29–50.

Breiman L (2001) Random forests. *Machine Learn.* 45(1):5–32.

Bryan K, Guzman J (2021) Entrepreneurial migration. Preprint, submitted October 11, http://dx.doi.org/10.2139/ssrn.3914010.

Brynjolfsson E, Horton JJ, Ozimek A, Rock D, Sharma G, TuYe HY (2020) Covid-19 and remote work: An early look at us data. Technical report, National Bureau of Economic Research, Cambridge, MA.

Budden P, Murray F (2019) MIT's stakeholder framework for building & accelerating innovation ecosystems.

Catalini C, Guzman J, Stern S (2019) Hidden in plain sight: Venture growth with and without venture capital. NBER Working Paper No. 26521, National Bureau of Economic Research, Cambridge, MA.

Chatterji A, Glaeser E, Kerr W (2014) Clusters of entrepreneurship and innovation. *Innovation Policy and the Economy*, vol. 14.

Chen H, Gompers P, Kovner A, Lerner J (2010) Buy local? The geography of successful venture capital expansion. *J. Urban Econom.* 67(1):90–102.

Chen Y, Rosenthal SS (2008) Local amenities and life-cycle migration: Do people move for jobs or fun? *J. Urban Econom.* 64(3):519–537.

Conti A, Guzman J (2023) What is the US comparative advantage in entrepreneurship? Evidence from Israeli migration to the United States. *Rev. Econom. Statist.* 105(3):528–544.

Dahl MS, Sorenson O (2012) Home sweet home: Entrepreneurs' location choices and the performance of their ventures. *Management Sci.* 58(6):1059–1071.

Delgado M, Porter ME, Stern S (2010) Clusters and entrepreneurship. *J. Econom. Geography* 10(4):495–518.

Delgado M, Porter ME, Stern S (2014) Clusters, convergence, and economic performance. *Res. Policy* 43(10):1785–1799.

Dell M, Jones BF, Olken BA (2012) Temperature shocks and economic growth: Evidence from the last half century. *Amer. Econom. J. Macroeconom.* 4(3):66–95.

Diamond R (2016) The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *Amer. Econom. Rev.* 106(3):479–524.

Duranton G, Puga D (2001) Nursery cities: Urban diversity, process innovation, and the life cycle of products. *Amer. Econom. Rev.* 91(5):1454–1477.

Fazio C, Guzman J, Stern S (2019) The impact of state-level r&d tax credits on the quantity and quality of entrepreneurship. Technical report, National Bureau of Economic Research, Cambridge, MA.

Fazio CE, Guzman J, Liu Y, Stern S (2021) How is covid changing the geography of entrepreneurship? evidence from the startup cartography project. Technical report, National Bureau of Economic Research, Cambridge, MA.

Feldman M, Siegel DS, Wright M (2019) New developments in innovation and entrepreneurial ecosystems. *Industrial Corporate Change* 28(4):817–826.

Feldman MP (2014) The character of innovative places: Entrepreneurial strategy, economic development, and prosperity. *Small Bus. Econom.* 43(1):9–20.

Feldman MP, Francis JL (2003) Fortune favours the prepared region: The case of entrepreneurship and the capitol region biotechnology cluster. *Eur. Planning Stud.* 11(7):765–788.

Florida R (2002) Bohemia and economic geography. *J. Econom. Geography* 2(1):55–71.

Forman C, Goldfarb A, Greenstein S (2016) Agglomeration of invention in the bay area: Not just ICT. *Amer. Econom. Rev.* 106(5):146–151.

Galasso A, Schankerman M (2015) Patents and cumulative innovation: Causal evidence from the courts. *Quart. J. Econom.* 130(1):317–369.

Gans J, Scott E, Stern S (2023) *Entrepreneurial Strategy*. Forthcoming.

Gianchandani E (2022) NSF Regional Innovation Engines (NSF Engines) Program, Accessed June, 2022, https://beta.nsf.gov/funding/opportunities/nsf-regional-innovation-engines-nsf-engines.

Glaeser EL, Gottlieb JD (2009) The wealth of cities: Agglomeration economies and spatial equilibrium in the United States. *J. Econom. Literature* 47(4):983–1028.

Glaeser EL, Kerr WR (2009) Local industrial conditions and entrepreneurship: How much of the spatial distribution can we explain? *J. Econom. Management Strategy* 18(3):623–663.

Gonzalez-Uribe J, Leatherbee M (2018) The effects of business accelerators on venture performance: Evidence from start-up chile. *Rev. Financial Stud.* 31(4):1566–1603.

Green TC, Jame R (2013) Company name fluency, investor recognition, and firm value. *J. Financial Econom.* 109(3):813–834.

Guzman J, Kacperczyk AO (2019) Gender gap in entrepreneurship. *Res. Policy* 48(7):1666–1680.

Guzman J, Stern S (2015) Where is Silicon Valley? *Science* 347(6222): 606–609.

Guzman J, Stern S (2017) Nowcasting and placecasting entrepreneurial quality and performance. *Proc. NBER/CRIW Measuring Entrepreneurial Businesses: Current Knowledge and Challenges Conf.*

Guzman J, Stern S (2020) The state of American entrepreneurship: New estimates of the quantity and quality of entrepreneurship for 32 us states, 1988–2014. *Amer. Econom. J. Econom. Policy* 12(4):212–243.

Han P, Liu C, Tian X, Wang K (2022) Invest local or remote? The effects of COVID-19 lockdowns on venture capital investment around the world. Preprint, submitted November 14, http://dx.doi.org/10.2139/ssrn.3871317.

Hsieh C-T, Moretti E (2019) Housing constraints and spatial misallocation. *Amer. Econom. J. Macroeconom.* 11(2):1–39.

Iacus SM, King G, Porro G (2012) Causal inference without balance checking: Coarsened exact matching. *Political Anal.* 20(1):1–24.

Jacobs J (1970) *The Economy of Cities* (Vintage, New York).

Jaffe JA, Trajtenberg M, Henderson R (1993) Geographic localization of knowledge spillovers as evidenced by patent citations. *Quart. J. Econom.* 108(3):577–598.

Kerr W (2008) Ethnic scientific communities and international technology diffusion. *Rev. Econom. Statist.* 90(3):518–537.

Kerr W, Fu S (2008) The survey of industrial r and d: Patent database link project. *J. Tech. Transfer* 33(2):173–186.

Kerr W, Kominers SD (2015) Agglomerative forces and cluster shapes. *Rev. Econom. Statist.* 97(4):877–899.

Kerr WR, Robert-Nicoud F (2020) Tech clusters. *J. Econom. Perspective* 34(3):50–76.

Krugman P (1991) Increasing returns and economic geography. *J. Political Econom.* 99(3):483–499.

Lerner J (2012) *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed: And What to Do About It* (Princeton University Press, Princeton, NJ).

Lerner J, Malmendier U (2013) With a little help from my (random) friends: Success and failure in post-business school entrepreneurship. *Rev. Financial Stud.* 26(10):2411–2452.

Marshall A (1890) *Principles of Economics* (MacMillan and Co., New York).

Marx M, Strumsky D, Fleming L (2009) Mobility, skills, and the Michigan non-compete experiment. *Management Sci.* 55(6):875–889.

Michelacci C, Silva O (2007) Why so many local entrepreneurs? *Rev. Econom. Statist.* 89(November):615–633.

Molloy R, Smith CL, Wozniak A (2011) Internal migration in the united states. *J. Econom. Perspective* 25(3):173–196.

Moretti E (2012) *The New Geography of Jobs* (Mariner Books, Boston).

Moretti E (2021) The effect of high-tech clusters on the productivity of top inventors. *Amer. Econom. Rev.* 111(10):3328–3375.

Moretti E, Wilson DJ (2017) The effect of state taxes on the geographical location of top earners: Evidence from star scientists. *Amer. Econom. Rev.* 107(7):1858–1903.

Nanda R, Sorensen JB (2010) Workplace peers and entrepreneurship. *Management Sci.* 56(7):1116–1126.

Oster E (2019) Unobservable selection and coefficient stability: Theory and evidence. *J. Bus. Econom. Statist.* 37(2):187–204.

PwC/CB Insights (2022) PwC/CB insights Moneytree explorer. Accessed May 10, 2022, https://www.pwc.com/us/en/industries/tmt/technology/moneytree/explorer.html.

Roback J (1982) Wages, rents, and the quality of life. *J. Political Econom.* 90(6):1257–1278.

Roche M, Oettl A, Catalini C (2020) Entrepreneurs (co-) working in close proximity: Impacts on technology adoption and startup performance outcomes. Working paper.

Roche MP (2019) Taking innovation to the streets: Microgeography, physical structure and innovation. *Rev. Econom. Statist.* 1–47.

Russell J (2014) It's settled: Silicon valley is dying. So what's next? *Wired*, https://www.smartcitiesdive.com/ex/sustainablecitiescollective/it-s-settled-silicon-valley-dying-so-what-s-next/222886/.

Samila S, Sorenson O (2011) Venture capital, entrepreneurship, and economic growth. *Rev. Econom. Statist.* 93(1):338–349.

Saxenian AL (1994) *Regional Advantage: Culture and Competition in Silicon Valley and Route 128* (Harvard University Press, Cambridge, MA).

Singh J, Marx M (2013) Geographic constraints on knowledge spillovers: Political borders vs. spatial proximity. *Management Sci.* 59(9):2056–2078.

Sorenson O (2018) Innovation policy in a networked world. *Innovative Policy Econom.* 18(1):53–77.

Sorenson O, Stuart TE (2001) Syndication networks and the spatial distribution of venture capital investments. *Amer. J. Sociol.* 106(6): 1546–1588.

Stam E (2007) Why butterflies don't leave: Locational behavior of entrepreneurial firms. *Econom. Geography* 83(1):27–50.

Stuart T, Sorenson O (2003a) The geography of opportunity: Spatial heterogeneity in founding rates and the performance of biotechnology firms. *Res. Policy* 32(2):229–253.

Stuart T, Sorenson O (2003b) Liquidity events and the geographic distribution of entrepreneurial activity. *Admin. Sci. Quart.* 48(2): 175–201.

Weber A (1929) *Uber den Standort der Industrie (Theory of the Location of Industries)* (University of Chicago Press, Chicago).

Wooldridge JM (2010) *Econometric Analysis of Cross Section and Panel Data* (MIT Press, Cambridge, MA).

Zucker LG, Darby MR, Brewer MB (1998) Intellectual human capital and the birth of US biotechnology enterprises. *Amer. Econom. Rev.* 88(1):290.