

# The Impact of Venture Capital Monitoring: Evidence from a Natural Experiment\*

Shai Bernstein  
Stanford University  
Graduate School of Business

Xavier Giroud  
Massachusetts Institute of Technology  
Sloan School of Management

Richard Townsend  
Dartmouth College  
Tuck School of Business

February 23, 2014

## Abstract

Do VCs contribute to the innovation and success of their portfolio companies, or do they merely select companies that are already poised to innovate and succeed even absent their involvement? To address this question, we exploit exogenous variation in VC involvement stemming from the introduction of new airline routes that reduce the travel time between VCs and their *existing* portfolio companies, thereby holding company selection fixed. We find that, within an existing VC-company relationship, reductions in travel time lead to increased innovation and a higher likelihood of an IPO. These effects concentrate in routes that connect portfolio companies with their lead VC, as opposed to other investors. Overall, the results indicate that VC involvement is an important determinant of innovation and success.

**JEL Classification:** D81, G24, L26, M13, O31, O32

**Keywords:** Venture Capital, Monitoring, Innovation, IPO

---

\*We are grateful to Yael Hochberg, Dirk Jenter, Anna Kovner, Josh Lerner, Holger Mueller, Francisco Pérez-González, David Robinson, Antoinette Schoar, as well as seminar participants at the NBER Entrepreneurship Meetings, Stanford (GSB), and Dartmouth (Tuck), for helpful comments and suggestions.

# 1 Introduction

It is often argued that venture capital (VC) plays an important role in promoting innovation and growth. Consistent with this belief, governments around the world have pursued a number of policies aimed at fostering VC activity (Lerner, 2009). However, there remains scarce evidence that the activities of venture capitalists actually play a causal role in stimulating the creation of innovative and successful companies. Indeed, VCs may simply select companies that are poised to innovate and succeed, even absent their involvement. In this paper, we examine whether the activities of VCs do affect portfolio company outcomes.

Answering this question is difficult for two reasons. First, since VCs back companies that are small and young as well as privately held, data are usually not available for observably similar non-VC-backed companies. As a result, there is often no control group available to estimate a counterfactual. Fortunately, recent work has helped to overcome this limitation by using survey data (Hellmann and Puri, 2000) as well as data from the U.S. Census Bureau (Chemmanur et al., 2011; Puri and Zarutskie, 2012). However, even after overcoming data limitations, there inevitably remains a second issue, which is the endogeneity of VC funding. In particular, it is entirely plausible that even among companies that are similar along the coarse dimensions we can observe, VCs tend to invest in those with greater potential.

Naturally, an ideal experiment would be to randomly provide certain companies with VC funding and others not. Such an experiment would randomize away the selection of companies (“screening”), thus allowing us to estimate the effect of VC involvement (“monitoring”).<sup>1</sup> Unfortunately, it is quite difficult to find a setting that convincingly approximates this experiment. That being said, another useful experiment would be to instead randomly vary VC involvement *after* initial investments are made. This would allow us to identify the effect

---

<sup>1</sup>Kaplan and Strömberg (2001) review the screening and monitoring roles of VCs, and emphasize the difficulty of disentangling them.

of VC involvement, holding company selection fixed. In particular, if differences in outcomes for VC-backed companies are driven purely by selection, post-investment involvement of the VCs should have no effect. In this paper we attempt to approximate this second experiment.

The source of exogenous variation in VC involvement that we exploit is the introduction of new airline routes that reduce the travel time between VC firms and their existing portfolio companies. Previous work suggests that travel time reductions lower monitoring costs for firms with headquarters that are geographically separated from their production facilities (Giroud, 2013). Similarly, there is evidence that proximity lowers monitoring and information acquisition costs in the context of mutual funds (Coval and Moskowitz, 1999, 2001) and banks (Petersen and Rajan, 2002). If VC activities do matter, reductions in the cost of monitoring should translate into better portfolio company performance by allowing VCs to engage in more of these activities. For example, VCs may be able to spend more time advising and shaping senior management, providing access to key resources, and aiding in company professionalization in myriad other ways (e.g., Lerner, 1995; Hellmann and Puri, 2000, 2002; Kaplan and Strömberg, 2004; Bottazzi et al., 2008).

There is ample anecdotal evidence that venture capitalists are sensitive to distance and travel time. For example, in response to a new United Airlines flight between Raleigh-Durham and San Francisco in 2012, the president of the Durham Chamber of Commerce stated that the new route would be valuable to “venture capitalists who like to be a direct flight away from any company they’re going to invest in” (*News & Observer*, August 12, 2012). Similarly, the lack of direct flights to Indianapolis is seen as an impediment to venture capital in the area: “Layovers and complicated connections are aggravating [...] That’s an important consideration because most venture capitalists want to keep close tabs on the companies they invest in, which requires frequent in-person visits” (*Indianapolis Star*,

October 8, 2000).<sup>2</sup> Consistent with the anecdotal evidence, the academic literature shows that VC activity is sensitive to distance and travel time. For example, Lerner (1995) finds that VCs are more likely to sit on boards of geographically proximate companies. In addition, Chen et al. (2010) find that VCs are more likely to invest in a distant region if they already visit one portfolio company in the same region.

We begin by documenting that there is significant VC activity outside of the three main regions of Northern California, New England, and New York. Indeed, approximately 50% of both VC-backed companies and VC investment firms are located outside of these three regions, consistent with the findings of Chen et al. (2010). Moreover, we show that it is fairly common for VCs to invest in distant portfolio companies. Given these patterns, we explore how the introduction of new airline routes that reduce the travel time between VCs and their portfolio companies affect company-level outcomes. The primary outcomes we examine are the quantity and quality of innovation (as measured by the patent count and citations per patent, respectively), as well as success (as measured by exit via IPO or acquisition). Using a difference-in-differences estimation framework, we find that the introduction of a new airline route leads to a 3.1% increase in the number of patents the portfolio company produces and a 5.8% increase in the number of citations per patent it receives. Furthermore, the treatment increases the probability of going public by 1.0%, and of having a successful exit (via IPO or acquisition) by 2.5%. These results indicate that VC involvement is an important determinant of innovation and success.

A natural concern is that local shocks, in the region of either the VC or the portfolio

---

<sup>2</sup>Relatedly, the president of England & Company, an investment bank with major activities in venture capital, argues that limited air service to Madison tends to discourage VC investment there: “Many potential venture capital investors on the East and West coasts aren’t willing to travel anywhere that isn’t serviced by a direct flight, and early-stage investors often like to play an active role in a company’s development, which is difficult to do from afar” (*The Wisconsin State Journal*, November 18, 2004). In an interview, Jim Nichols, director at the North Carolina Department of Commerce was asked: “So something seemingly as marginal as a direct flight to the money and technology mecca of Silicon Valley is absolutely crucial for [North Carolina]’s growth?” He answered: “You wouldn’t think it would be that important, but it really is a factor for these companies” (*Tribune Business News*, November 27, 2000).

company, could be driving the results. For example, a booming local economy may lead to both increased innovation and the introduction of a new airline route. In this case, we may estimate a spurious positive effect of travel time reductions on innovation. However, since our treatment is defined at the VC-company pair level, we can control for such local shocks. Specifically, we include two full sets of MSA (Metropolitan Statistical Area) by year fixed effects for the MSAs of both the VC and the portfolio company.

One remaining concern that is not addressed by controlling for local shocks is the possibility that if a portfolio company is performing very well, a new airline route may be introduced in response. While we do not believe this to be likely, it would bias our estimates. To ensure that such pre-existing trends are not driving our results, we examine the dynamics of how company outcomes change in the years surrounding the treatment. We find the bulk of the effect coming 1 to 2 years after the treatment, with no “effect” prior to the treatment. Moreover, we show that our results are robust to considering only new airline routes that are the outcome of a merger between two airlines or the opening of a new hub. Such treatments are likely to be even more exogenous to any given VC-company pair.

Next, we investigate the channel through which these effects operate. Our main hypothesis is that a reduction in travel time increases VC involvement, which in turn improves portfolio company performance. Unfortunately, we cannot directly observe whether VC involvement actually increases when monitoring costs decline. However, we take advantage of the fact that certain VCs should be more sensitive to changes in monitoring costs than others. Specifically, VCs often syndicate their investments, and when this occurs, one typically takes the role of the lead investor. The lead investor is generally more actively involved in the monitoring of the portfolio company, while others act more as passive providers of capital (Gorman and Sahlman, 1989). Given that lead VCs play a greater role in monitoring, their monitoring effort should be more sensitive to reductions in monitoring costs, as should

portfolio company performance. Indeed, we find that our results are driven primarily by reductions in travel time for lead VCs rather than other members of the investment syndicate. Finally, reductions in monitoring costs should be greater the greater the reduction in travel time. Consistent with this argument, we find larger effects associated with larger travel time reductions.

In the last part of the paper, we extend our analysis by studying whether the introduction of new airline routes affects aggregate venture capital flows between MSAs. Accordingly, we move from doing analysis at the VC-company-pair level (“relationship analysis”), to analysis at the MSA-pair level (“regional analysis”). For the regional analysis we again use a difference-in-differences estimation framework, controlling for local shocks in both source and target MSAs. We find that the introduction of a new airline route between two MSAs leads to a 4.6% increase in total VC investments as well as a 2.5% increase in the likelihood of VC activity between the two MSAs. These results indicate that better airline connections foster VC flows between regions. In addition, these results have relevant policy implications. Governments around the world employ a wide range of policies to promote VC activity, ranging from providing tax subsidies to getting directly involved in the investment process through government-run venture capital programs (Lerner, 2009). Our evidence suggests that policies encouraging new routes to VC hubs are an effective tool for stimulating VC activity.

Our paper contributes to a growing literature that studies the effect of VCs on portfolio company outcomes. As previously mentioned, much of this literature tries to disentangle VC monitoring from screening by comparing outcomes of VC-backed and non-VC-backed companies (e.g., Hellmann and Puri, 2000, 2002; Chemmanur et al., 2011; Puri and Zarutskie, 2012). These papers are valuable given the scarcity of data on young companies that are not affiliated with a VC. However, even if both groups of companies are matched on the

basis of observables, it is quite plausible that VCs select companies with higher potential *ex ante*—an inherently unobservable characteristic. In contrast, our setting allows us to identify the effect of VC monitoring holding selection fixed, because we exploit exogenous reductions in monitoring costs after initial investments are made. Other papers rely on structural modeling. In particular, Sorensen (2007) models the two-sided matching process of VCs and entrepreneurs to structurally estimate the relative importance of VC monitoring and screening as explanations for why companies backed by more experienced VCs outperform. Relatedly, Kortum and Lerner (2000) structurally estimate industry-level patent production functions with corporate R&D and venture capital as inputs in order to compare their relative potency. Our paper differs from these in that it does not require any structural assumptions for identification.

The remainder of this paper is organized as follows. Section 2 discusses the data and key variables. Section 3 discusses our empirical strategy. Section 4 presents the results of the relationship analysis. Section 5 presents the results of the regional analysis, and Section 6 concludes.

## 2 Data

### 2.1 Sources and Sample Selection

We obtain data on venture-backed companies from the Thomson Reuters VentureXpert database (formerly called Venture Economics). VentureXpert, along with Dow Jones' VentureSource (formerly VentureOne), are the two primary venture capital data sources available. Both have been validated by previous researchers against known financing rounds (Kaplan et al., 2002). We choose to use VentureXpert because VentureSource starts later and is less comprehensive in earlier years, when many new airline routes were introduced.

VentureXpert began compiling data in 1977. It contains detailed information about the dates of venture financing rounds, the investors and portfolio companies involved, the estimated amounts invested by each party, and the ultimate portfolio company outcome. The database also contains detailed information on the location of each VC firm and portfolio company. It should be noted that one shortcoming of these data for our purposes is that VentureXpert only associates a VC firm with a single location (its main office). However, some of the larger VC firms operate out of multiple offices. While ideally we would observe all of these offices, this should not present a systematic source of bias.<sup>3</sup> We limit the sample to U.S. based portfolio companies coded as being in a venture stage (seed, early, expansion, or later stage) in their first observed financing round. For our baseline analysis, we further limit the sample to only VC-company pairs involving the lead investor, which will be defined in Section 2.2.3.

To measure the innovative output of portfolio companies, we combine VentureXpert with data from the NBER Patent Data Project (Hall et al., 2001). The NBER data cover all utility patents granted by the U.S. Patent and Trademark Office (USPTO) from 1976 to 2006.<sup>4</sup> Among other things, the data provide information on the date a patent was applied for and ultimately granted as well as its detailed technology class. If a patent was assigned to one or more companies (“assignees”), the data also provide information on assignee name(s)/location(s). We match the NBER data with VentureXpert using standardized company and location names along with the company’s founding date and the date of the assignee’s first patent application. The details of the matching procedure are provided in Appendix A. Finally, we also supplement the NBER data with citation data from Google

---

<sup>3</sup>If the monitoring is done out of local offices, not accounting for them would merely go against us finding any effect.

<sup>4</sup>In addition to utility patents, there are three other minor patent categories: design, reissue, and plant patents. Following the literature we focus only on utility patents, which represent approximately 99% of all awards (Jaffe and Trajtenberg, 2002).



patents in some cases so that we can observe citations in a three-year window following the grant date for all patents, including those at the end of the NBER sample in 2006.

Data on airline routes are obtained from the T-100 Domestic Segment Database (for the period 1990 to 2006) and ER-586 Service Segment Data (for the period 1977 to 1989), which are compiled from Form 41 of the U.S. Department of Transportation (DOT). All airlines operating flights in the U.S. are required by law to file Form 41 with the DOT and are subject to fines for misreporting. Strictly speaking, the T-100 and ER-586 are not samples: they include all flights that have taken place between any two airports in the U.S. The T-100 and ER-586 contain monthly data for each airline and route (segment). The data include, for example, the origin and destination airports, flight duration (ramp-to-ramp time), scheduled departures, performed departures, enplaned passengers, and aircraft type.

After combining these three data sources we are left with a sample of venture-backed companies that were active between 1977 (the beginning of the airline data) and 2006 (the end of the patent data). In total, we observe 22,986 companies, receiving funding from 3,158 lead VC firms. Table 1 shows the composition of the sample. Panel A shows the company region distribution broken down according to whether the company was ever treated or not (i.e. experienced a reduction in travel time to its lead VC). Similarly, Panel C shows the VC region distribution broken down according to whether the venture firm was ever part of a treatment or not. Perhaps the most striking finding from these tables is that, contrary to common perception, a significant amount of venture capital activity takes place outside of Northern California, New England, and New York. Indeed, approximately 50% of venture-backed companies and VC firms are located outside of these three regions. This is consistent with the findings of Chen et al. (2010). Overall, treated and untreated companies are distributed similarly across regions; however, as one might expect, treated companies are less likely to be located in Northern California. Similarly, Panel C shows that VCs that are

part of a treatment are also less likely to be located in Northern California. Finally, Panel B shows that treated and untreated companies are also distributed similarly across industries, although treated companies are somewhat less likely to be in the Internet sector.

While Table 1 shows that both portfolio companies and VC firms are fairly dispersed geographically, it does not directly show whether it is common for VCs to invest in distant portfolio companies. If, to a first approximation, all VCs invested locally, we would not have sufficient power to identify an effect, since there would be few reductions in travel time due to new airline routes. Figure 1 provides some perspective on the distance between VCs and portfolio companies graphically. First, it shows the distribution of portfolio companies across states, depicting states with more companies in darker shades. More interestingly, the height of the bar over each state indicates the percentage of companies located in that state, which are funded by a lead VC from the same state. As can be seen, many states have a relatively low percentage of locally funded portfolio companies. Thus, airline routes could potentially be an important determinant of monitoring costs for many companies. To examine this issue more directly still, we plot the cumulative density function of the VC-company distance distribution in Figure 2. Consistent with what one might expect, we find that a large fraction of VC investments are local, with around 30% being located close to zero miles from their lead VC. However, the median distance between a portfolio company and its lead VC is approximately 200 miles and the 60th percentile is approximately 500 miles. Thus, around 40% of portfolio companies are located more than 500 miles from their lead VC. This both suggests that we will likely have enough power to identify an effect if one is present, and that the long-distance pairs that we use for identification are not particularly unusual.

## 2.2 Definitions of Variables

### 2.2.1 Treatment

To estimate the effect of reductions in travel time on portfolio company outcomes, we define a treatment indicator variable equal to one if a new airline route is introduced that reduces the travel time between the VC firm and the portfolio company. Travel time is estimated as the time it would take to travel from the VC’s ZIP code to the company’s ZIP code using the optimal itinerary and means of transportation (car or airplane). The details of the algorithm used to compute optimal itineraries and travel times are described in Appendix B. During our sample period (1977-2006), there are 1,131 treated VC-company pairs. The average travel time reduction is 126 minutes round-trip. Note, however, that this estimated reduction in travel time is likely a lower bound as it does not take into account the compounding probability of delays and cancellations when taking indirect flights (see also Section 4.3.1).

### 2.2.2 Innovation

We use patent-based measures of the scale and quality of a company’s innovation (Jaffe and Trajtenberg, 2002; Lanjouw et al., 1998). These measures have been widely adopted over the past two decades.<sup>5</sup> Our primary measure of the scale of a company’s innovation during a year is the number of (eventually granted) patents it applied for. Our primary measure of the quality of a company’s innovation during a year is the number of citations it received per patent. Patent citations are important in patent filings since they serve as “property markers” delineating the scope of the granted claims. Hall et al. (2005) illustrate that citations are a good measure of innovation quality and economic importance. Specifically, they find that an extra citation per patent boosts a firm’s market value by 3%. Moreover, Kogan et al. (2012) show that the stock market reaction to patent approvals is a strong predictor of the number

---

<sup>5</sup>Recent examples include Lerner et al. (2011); Aghion et al. (2013); Seru (2014).

of future citations a patent receives.

One challenge in measuring patent citations is that patents granted at the end of the sample period have less time to garner citations than those granted at the beginning. To address this issue, we only consider citations that occur during a three-year window following the date a patent is granted. In addition, we check that our results are robust to correcting for truncation using the estimated shape of the citation-lag distribution as in Hall et al. (2001). An additional consideration is that citation rates vary over time and across technologies. To ensure this does not affect our results, we also explore scaling each patent’s citation count by the average citation count for patents granted in the same year and technology class. Finally, we take logs and add one to both the patent count and citation variables.

### **2.2.3 Other Variables**

In addition to innovation, we also measure success annually. We define company success in two ways. The first is an indicator variable equal to one if the company went public during a given year. The second is an indicator variable equal to one if the company went public or was acquired. The issue with the second definition is that it likely captures some acquisitions that were not positive outcomes. Specifically, an acquisition may be a sell-off that was not very profitable for the company’s investors or founders. Unfortunately, it is difficult to distinguish these cases in the data because the amount paid by the acquirer is frequently undisclosed. Nonetheless, given the increasing importance of acquisitions as a means of exit for successful venture-backed companies, this broader measure may better capture company success. To be conservative we focus primarily on the IPO indicator in our analysis; however, we show that including acquisitions as well yields even stronger results.

Finally, as previously mentioned, in our baseline analysis, we limit the sample to only VC-company pairs involving the lead investor. We focus on the lead investor because it is

likely to be the one most involved in monitoring. Following Gompers (1996), we define the lead investor as the one that has invested in the company the longest.<sup>6</sup> This is also consistent with Gorman and Sahlman’s (1989) finding that the venture firm originating the investment is usually the firm that acquires a board seat first and has the most input into the decisions of the company, even though it might not end up ultimately owning the largest equity stake. Our results are also robust to other commonly used definitions of the lead investor, such as the investor that invested the most in a given round.

### 3 Empirical Strategy

#### 3.1 Relationship Analysis

The introduction of new airline routes that reduce the travel time between VC firms and their portfolio companies make it easier for VCs to spend time at their portfolio companies. If VC activities do matter, such reduction in travel time should translate into better portfolio company performance by allowing VCs to engage in more of these activities. To estimate the effect of the introduction of new airline routes (“treatments”) on company outcomes, we adopt a difference-in-differences methodology similar to Giroud (2013). Specifically, we estimate the following regression:

$$y_{ijt} = \beta \times treatment_{ijt} + \boldsymbol{\gamma}'\mathbf{X}_{ijt} + \alpha_{ij} + \alpha_{MSA(i)} \times \alpha_t + \alpha_{MSA(j)} \times \alpha_t + \epsilon_{ijt}, \quad (1)$$

where  $i$  indexes portfolio companies,  $j$  indexes VC firms,  $t$  indexes years,  $MSA(i)$  indexes the Metropolitan Statistical Area (MSA) in which portfolio company  $i$  is located, and  $MSA(j)$  indexes the MSA in which VC  $j$  is located;  $y$  is the dependent variable of interest (e.g.,

---

<sup>6</sup>We break ties by selecting the firm that invested the most. If there are still ties, we classify all of the tied VC firms as lead investors.

number of patents, citations per patent, IPO), treatment is an indicator variable (“treatment indicator”) that equals one if a new airline route that reduces the travel time between company  $i$ ’s ZIP code and VC  $j$ ’s ZIP code has been introduced by year  $t$ ;  $\mathbf{X}$  is the vector of control variables, which includes company age and a set of indicator variables for the stage of VC financing;  $\alpha_t$  and  $\alpha_{ij}$  are year and VC-company pair fixed effects, respectively;  $\alpha_{MSA(i)} \times \alpha_t$  and  $\alpha_{MSA(j)} \times \alpha_t$  are MSA by year fixed effects with respect to company  $i$ ’s MSA and VC  $j$ ’s MSA, respectively;  $\epsilon$  is the error term. This methodology fully controls for fixed differences between treated and non-treated VC-company pairs via the inclusion of pair fixed effects. The inclusion of MSA by year fixed effects further accounts for local shocks that may correlate with the introduction of new airline routes. To allow for serial dependence of the error terms, we cluster standard errors at the portfolio company level. The coefficient of interest is  $\beta$  which measures the effect of the introduction of new airline routes on  $y$ .<sup>7</sup>

Our identification strategy can be illustrated with a simple example. From 1986 to 1994, Anesta Corporation, a biopharmaceutical company located in Salt Lake City, UT, was receiving VC funding from Flagship Ventures, a VC firm in Cambridge, MA. Until 1988, the fastest way to travel between Boston Logan Airport (BOS) and Salt Lake City International Airport (SLC) was an indirect flight operated by Delta Airlines with one stopover at Chicago O’Hare (ORD). In 1988, Delta introduced a direct flight between BOS and SLC, which substantially reduced the travel time between the two locations. To measure how this “treatment” affects, for example, the number of patents filed by Anesta, one could compute the difference in the number of patents before and after 1988. However, other events may have occurred around 1988, which may also have affected patenting. To account for this

---

<sup>7</sup>Note that the use of private jets is not widespread in the VC industry, and was not widespread in general for much of our sample period. Moreover, if anything, the use of private jets would merely go against us finding any effect.

possibility, we use a control group that consists of all VC-company pairs that have not been treated by 1988. We then compare the difference in the number of patents at Anesta before and after 1988 with the difference in the number of patents at the control companies before and after 1988. The difference between these two differences is the estimated effect of the treatment on patenting at Anesta.

### 3.1.1 Local Shocks

Including a control group accounts for the possibility of economy-wide shocks that are contemporaneous with the introduction of the new airline routes. However, since a treatment is defined at the VC-company level, we can tighten the identification by also controlling for local shocks in the portfolio company’s MSA, thereby separating out the effect of the new airline routes from the effect of contemporaneous local shocks. For example, Systemed Inc. is another biopharmaceutical company located in Salt Lake City. Around 1988, Systemed was receiving VC funding from Summit Capital Associates, a New York-based VC. (Direct flights between New York’s John F. Kennedy Airport and SLC were offered in each year during our sample.) If patenting at Systemed also increases around 1988, then an increase in patenting at Anesta might not be due to the new airline route between BOS and SLC, but rather due to a contemporaneous local shock that affects patenting in the Salt Lake City MSA. In Equation (1), we control for such local shocks by including the full set of MSA fixed effects (pertaining to the portfolio company’s location) interacted with year fixed effects ( $\alpha_{MSA(i)} \times \alpha_t$ ).

In addition, since a treatment is defined at the VC-company level, we can make the identification even tighter by also controlling for shocks at the location of the VC firm. In the above example, suppose there is a local shock that affects patenting in Boston in 1988. This local shock may affect Flagship Ventures, the Cambridge VC financing Anesta, and in

turn Anesta’s ability to innovate. In this case, however, patenting should also increase in the Boston area. In Equation (1), we control for such local shocks by including MSA fixed effects (pertaining to the VC’s location) interacted with year fixed effects ( $\alpha_{MSA(j)} \times \alpha_t$ ).<sup>89</sup>

### 3.1.2 Pair-Specific Shocks

One potential concern that is not addressed by controlling for local shocks, is the possibility that a pair-specific shock (i.e., a shock that is specific to a VC-company pair, but not to the MSA of the company, or the MSA of the VC) is driving both company-level outcomes (e.g., patenting) and the introduction of the new airline route. For example, it could be that a portfolio company that is successful in patenting becomes more salient to its VC. In response, the VC may want to spend more time at that company and hence may lobby for better airline connections to the company’s location. Nevertheless, such alternative stories are unlikely for several reasons. First, portfolio companies and VC firms are relatively small business entities. Hence, it seems unlikely that a VC-company pair is sufficiently powerful to successfully lobby for better airline connections (or that an airline would introduce a new route in response to a shock to that pair). To further rule out this concern, we have verified that our results also hold if we restrict our sample to portfolio companies and VC firms whose size is below the median in our sample, i.e. those companies and VCs that are even less able to successfully lobby for a new airline route. Second, we examine the dynamic effects of the treatment. Arguably, if the new airline routes are introduced in response to pair-specific shocks, one may already observe an “effect” of the new airline routes before they are even introduced. However, when we examine the dynamics of the treatment, we find

---

<sup>8</sup>In practice, it is computationally difficult to estimate a regression that has so many layers of fixed effects. Fortunately, recent algorithms have been developed that can handle such high-dimensional fixed effect regressions. In our analysis, we use the iterative algorithm of Guimarães and Portugal (2010). See Gormley and Matsa (2013) for details.

<sup>9</sup>In robustness checks, we further show that our results are similar if we allow local shocks to be industry specific, i.e. instead of including MSA by year fixed effects, we include the full set of MSA by industry by year fixed effects (see Section 4.4.6).



no such evidence: most of the effects we observe occur between 12 and 24 months after the introduction of the new airline routes. Third, in robustness checks, we show that our results also hold if we consider new airline routes that are introduced as part of the opening of a new hub or a merger between two airlines. Arguably, it is unlikely that a shock that is specific to a VC-company pair is sufficiently large to lead to a hub opening or an airline merger.

### **3.1.3 Differences between Treated and Non-Treated Pairs**

In order to be treated, a VC-company pair needs to be sufficiently far apart so that air travel is the optimal means of transportation between the two. Thus, by construction, treated pairs are farther apart than the average VC-company pair in the U.S. This is confirmed by looking at the summary statistics in Table 2. On average, treated pairs are located approximately 500 miles farther away than non-treated pairs. The other characteristics shown in the table further indicate that, for treated pairs, portfolio companies receive less funding, are less innovative, and tend to receive funding from VCs that are more experienced and more diversified.

While these differences may be intuitive, they do raise the concern of whether our control group is an appropriate one. Nevertheless, this concern is minimized for several reasons. First, in all our regressions, we include VC-company pair fixed effects, which fully controls for any fixed differences between treated and non-treated VC-company pairs. Since the main difference—the distance between VC and portfolio company—is a fixed characteristic, it seems likely that most of the relevant differences between the two groups are absorbed away. Second, because of the staggered introduction of the new airline routes over time, the eventually treated pairs are both control and treatment pairs (i.e., they remain in the control group until they become treated). Third, we show that our results are robust if we restrict the control group to those control pairs whose average distance matches the average

distance in the treatment group (i.e., we exclude short-distance control pairs so that the average distance is the same in both groups).<sup>10</sup> Fourth, we show that our results also hold if we allow pairs that differ on the basis of the characteristics in Table 2 to be on different time trends. More precisely, this test is conducted by including as additional controls the characteristics in Table 2 interacted with a full set of year fixed effects (see Bertrand and Mullainathan (2003) for a similar robustness check).

Finally, another helpful robustness check proposed by Bertrand and Mullainathan (2003) consists of estimating the difference-in-differences specification using only observations of the eventually treated pairs—essentially, due to the staggered introduction of the new airline routes, Equation (1) can be estimated using only this subsample (in this case, the control group consists exclusively of pairs that are subsequently treated). In our context, a caveat of this test is that the number of observations drops to 7,978 pair-year observations, which makes it infeasible to control for MSA by year fixed effects. Nevertheless, we show that our results are robust if we perform this test (dropping the MSA by year fixed effects from the regressions).

### 3.2 Regional Analysis

The difference-in-differences specification in Equation (1) can be extended to study whether proximity fosters VC activity at the regional level. To conduct this analysis, we aggregate our data from the VC-company level to the MSA-pair level. We then estimate the following regression:

$$y_{mnt} = \beta \times treatment_{mnt} + \alpha_t + \alpha_{mn} + \alpha_m \times \alpha_t + \alpha_n \times \alpha_t + \epsilon_{mnt}, \quad (2)$$

---

<sup>10</sup>In conducting this test, we lose about half of the control sample. The fact that the other half qualifies as a “distance-matched” control group confirms that long-distance VC-company relationships are fairly common, as discussed in Section 2 (see also Figures 1 and 2).

where  $m$  indexes MSAs from which VC funding is coming (i.e., MSAs of the VC firms),  $n$  indexes MSAs to which VC funding is going (i.e., MSAs of the portfolio companies), and  $t$  indexes years;  $y$  is the dependent variable of interest (e.g., the total amount of VC funding provided by VCs in MSA  $m$  to portfolio companies in MSA  $n$ ); treatment is the treatment indicator at the MSA-pair level;  $\alpha_t$  and  $\alpha_{mn}$  are year and MSA-pair fixed effects, respectively;  $\alpha_m \times \alpha_t$  and  $\alpha_n \times \alpha_t$  are the two sets of MSA by year fixed effects;  $\epsilon$  is the error term. Standard errors are clustered at the MSA-pair level. The identification strategy is analogous to that at the VC-company pair level. In particular, we are able to include MSA-pair fixed effects as well as the two sets of MSA by year fixed effects, thus controlling for local shocks that may be correlated with airlines' decisions to introduce new airline routes.

There are two main differences compared to the relationship analysis. First, “treatments” are coded in a different way. At the relationship level, a treatment is the introduction of a new airline route that reduces the travel time between the VC’s ZIP code and the company’s ZIP code, taking into account the optimal itinerary and means of transportation (see Appendix B). Since an MSA covers several ZIP codes, there is no notion of an “optimal itinerary” at the MSA-pair level. Instead, we code as a treatment the first time a direct flight is introduced between any two locations in the two MSAs. Second, there are a large number of MSA pairs between which no VC activity ever occurred during our sample period. For these pairs, any dependent variable would be set to zero in all years, and thus be absorbed by the inclusion of MSA-pair fixed effects. In the regressions, we drop these MSA-pairs from the sample. This follows common practice in the trade literature in which a similar issue arises when measuring trade flows between country pairs (e.g. Feyrer, 2009).

## 4 Relationship Analysis

### 4.1 Main Results

We estimate variants of Equation (1) to examine whether the introduction of new airline routes that reduce the travel time between VC firms and their portfolio companies affect portfolio companies' innovation and success. The results are presented in Table 3. In Columns (1)-(3), the dependent variable is the number of patents (in logs). The regression in Column (1) includes VC-company pair and year fixed effects. In Column (2), we also control for company age and a set of indicators for the stage of VC financing. In Column (3), we further control for local shocks by including the two sets of MSA by year fixed effects. The coefficient on the treatment indicator is very stable across all specifications. It lies between 0.031 and 0.037, which implies that the number of patents increases by 3.1% to 3.7% after the treatment. In Columns (4)-(6), we re-estimate these specifications using citations per patent (in logs) as the dependent variable. The coefficient on the treatment indicator varies between 0.058 and 0.074, corresponding to an increase in citations per patent of 5.8% to 7.4%. In Columns (7)-(9) the dependent variable is an indicator variable equal to one if the company goes public (IPO) during the year. We find that the introduction of new airline routes leads to an increase in the likelihood of going public by approximately 1.0%. Overall, our findings indicate that a reduction in VC monitoring costs leads to significant increases in innovation and the likelihood of an IPO.

### 4.2 Dynamic Effects of the Treatment

In Table 4, we study the dynamic effects of the introduction of new airline routes. Specifically, we replace the treatment indicator in Equation (1) with a set of four indicator variables representing the years around the treatment. For example, the indicator "Treatment (-1)" equals

one if the VC-company pair observation is recorded in the year preceding the treatment. The other indicator variables are defined analogously with respect to the year of the treatment (0), the first year after the treatment (1), and two or more years after the treatment (2+). The underlying specification is the conservative specification used in Columns (3), (6), and (9) of Table 3, i.e. the specification that includes control variables, VC-company pair fixed effects, year fixed effects, as well as the two sets of MSA by year fixed effects (henceforth, the “baseline specification”). We observe a very similar pattern for all three dependent variables. In particular, we always find that the coefficient on Treatment (−1), which measures the “effect” of the new airline routes before their introduction, is small and insignificant, suggesting that there are no pre-existing trends in the data.<sup>11</sup> The effect is positive but small in the year of the treatment (year 0). It is only one year after the treatment (year 1) that the effect becomes large and significant. Finally, the effect is persistent in the longer run (years 2+). In sum, the dynamic pattern suggests that it takes about 12 to 24 months until the reduction in travel time materializes into greater innovation and a higher likelihood of going public.

## 4.3 Discussion

### 4.3.1 Do VCs Respond to the Treatment?

Our results indicate that the introduction of new airline routes between VCs and their existing portfolio companies leads to increased innovation and a higher likelihood of going public. Arguably, the reduction in travel time increases VC involvement, which in turn improves innovation and success. While we do not directly observe the travel behavior of VCs, looking at survey evidence helps to put these results in context. Gorman and

---

<sup>11</sup>We cannot identify the coefficient of Treatment (−1) in the regression whose dependent variable is the IPO indicator. Since companies exit the sample upon going public, companies that go public before the treatment cannot be in the treatment group by construction.

Sahlman (1989) survey VC firms that account for roughly 40 percent of industry capital under management. They find that, on average, VCs invest in 9 companies at a time and sit on 5 boards. The lead investor visits the company site roughly 20 times per year and spends approximately 5 hours per visit, and 100 hours annually. On average, a treatment saves roughly 2 hours per trip, which at 20 trips per year is 40 hours per year of a VC’s time. Moreover, as mentioned before, the estimated 2 hours saved per trip is likely a lower bound, as it does not take into account the compounding probability of delays and cancellations when taking indirect flights.<sup>12</sup> Accordingly, our treatments do correspond to fairly large reductions in monitoring costs. Given the high opportunity costs of a VC’s time—e.g., Kaplan and Strömberg (2001) comment that “the scarcest commodity a VC has is time, not capital” (p. 428) and further refer to the anecdotal accounts of Gladstone (1988) and Quindlen (2000)—it seems reasonable that VCs are indeed fairly responsive to the treatment. This is in line with the ample anecdotal evidence that VCs are sensitive to travel time and flight connections (see Section 1).

### 4.3.2 Lead versus Non-Lead VCs

To further ensure that our results are driven by increased VC monitoring following the treatment, we take advantage of the fact that, ex-ante, certain VCs are expected to be more sensitive to changes in monitoring costs than others. In particular, VC investments are often syndicated with one VC taking the role of the lead investor. The lead investor typically is the one primarily in charge of monitoring, while other investors are more passive providers of capital. Indeed, Gorman and Sahlman (1989) find that a VC acting as lead investor spends ten times the number of hours on a company than he or she would otherwise. Accordingly, we expect the treatment effect to be concentrated in routes that connect portfolio companies

---

<sup>12</sup>In addition, in calculating travel time reductions, we conservatively assumed that the layover time is one hour (see Giroud, 2013).

with their lead VC, as opposed to other syndicate members.

To investigate this hypothesis, we re-estimate our baseline specification in the sample of VC-company pairs involving a non-lead investor located in a different MSA than the lead investor. We now set the treatment indicator to one if a new airline route is introduced that reduces the travel time between a portfolio company and a non-lead investor. The results are shown in Table 5. We find that, for all dependent variables, the estimated treatment effect is statistically insignificant. Moreover, the sample size in this analysis is comparable to that from the baseline analysis and the point estimates are close to zero, suggesting these are well-estimated zero effects. These results are consistent with the argument that VC involvement increases following the treatment—travel time reductions appear to matter primarily for active investors.<sup>13</sup>

### 4.3.3 Small versus Large Reductions in Travel Time

Finally, if travel time indeed matters, we expect to find a stronger treatment effect for larger reductions in travel time. In our baseline analysis, any new airline route that reduces the travel time between a VC firm and its portfolio company was coded as a treatment, regardless of the magnitude of the travel time reduction. We now interact the treatment indicator with two dummy variables indicating whether the reduction in travel time is “large” or “small”. We consider a travel time reduction to be large if it is more than one hour. The results are reported in Table 6. For travel time reductions of less than one hour, the treatment effect is small and insignificant. In contrast, the treatment effect is strongest and highly significant for travel time reductions of more than one hour.

---

<sup>13</sup>In addition, these results reinforce our identification as they can be viewed as a placebo test. Company outcomes do not always improve with the introduction of a new airline route—they only improve when that airline route connects the company to an active investor.

## 4.4 Robustness

### 4.4.1 Hub Openings and Airline Mergers

As explained in Section 3.1.2, one potential concern that is not addressed by controlling for local shocks is the possibility that a VC-company pair-specific shock is driving both company outcomes and the introduction of a new airline route (e.g., through lobbying). Given the relatively small size of portfolio companies and VC firms, such alternative stories seem unlikely. Moreover, we have verified that our results are robust if we restrict our sample to portfolio companies and VC firms whose size is below the median; that is, companies and VCs that are even less able to successfully lobby for a new airline route. In addition, if a new airline route is introduced in response to a pair-specific shock, one may already observe an “effect” of the new airline route before it is even introduced. However, when we looked at the dynamics of the treatment effect, we found no evidence for such pre-existing trends.

Another way to rule out this concern is by considering new airline routes that are introduced as part of a hub opening or a merger between airlines. Arguably, it is unlikely that a pair-specific shock could induce the opening of a new hub or the merger of two airlines. Thus, new airline routes of this kind are more likely to be exogenous. The data on hub openings and airline mergers are obtained from Giroud (2013). Hub and merger treatments account for about 15% of the treatments in our sample. In Panel A of Table 7, we replace the treatment indicator in our baseline specification with two dummy variables indicating hub/merger treatments (“Hub or Merger”) and other treatments (“Other”), respectively. As can be seen, our results are robust when considering hub and merger treatments, which alleviates concerns that our results may be driven by unobservable pair-specific shocks.<sup>14</sup>

---

<sup>14</sup>The treatment effect is larger for hub and merger treatments compared to other treatments. This likely reflects the fact that new airline routes that are introduced as part of a hub opening or airline merger are mostly long-distance routes, which tend to be associated with larger travel time reductions.



#### 4.4.2 Distance-Matched Control Group

As discussed in Section 3.1.3, in order to be treated, a VC-company pair needs to be sufficiently far apart so that air travel is the optimal means of transportation between the two. Thus, by construction, treated pairs are farther away than control pairs. This difference raises the concern of whether our control group is an appropriate one. While the inclusion of VC-company pair fixed effects accounts for any time-invariant differences between pairs (such as differences in distance), a remaining concern is that long-distance VC-company pairs may be on a different trend. To mitigate this concern, we re-estimate our baseline specification after restricting the control group to those control pairs whose average distance matches the average distance in the treatment group. More precisely, we exclude short-distance control pairs (in increasing distance) until the average distance is the same in both groups. The results are presented in Panel B of Table 7. As is shown, our results are robust to using this “distance-matched” control group.

#### 4.4.3 Heterogeneous Time Trends

Another way to address the possibility that control and treated pairs may be on different trends is to explicitly control for such heterogeneous time trends. This can be done by interacting the cross-sectional characteristic of interest (e.g., distance) with the full set of year fixed effects (see Bertrand and Mullainathan, 2003). Specifically, we interact all characteristics from Table 2 with year fixed effects and re-estimate our baseline specification with these additional controls.<sup>15</sup> The results are reported in Panel C of Table 7. The estimated treatment effects are very similar to before.

---

<sup>15</sup>Time-varying characteristics are measured in the first year of the pair (baseline year), see Bertrand and Mullainathan (2003).

#### 4.4.4 Eventually Treated Pairs

Another helpful robustness check proposed by Bertrand and Mullainathan (2003) consists of estimating the difference-in-differences specification using only observations of the eventually treated pairs—essentially, due to the staggered introduction of the new airline routes, Equation (1) can be estimated using only this subsample (in this case, the control group consists exclusively of pairs that are subsequently treated). This further helps to alleviate concerns about the comparability of the control group. In our context, a caveat of this test is that the number of observations drops to 7,978 pair-year observations, which makes it infeasible to control for MSA by year fixed effects. The results without these fixed effects are reported in Panel D of Table 7. Again, they remain similar.

#### 4.4.5 Alternative Dependent Variables

In Panel E of Table 7 we explore whether our results are robust to alternative definitions of our main dependent variables. As discussed in Section 2.2.2, we only consider citations during a three-year window following a patent grant, so that all patents in our sample have the same amount of time to garner citations. Hall et al. (2001) propose an alternative adjustment method that uses the estimated shape of the citation-lag distribution. In Column (1), we re-estimate our baseline specification, adjusting for truncation in this manner. The coefficient on the treatment indicator is similar to that in our baseline specification in Table 3. Another common practice in the literature is the use of citation-weighted patent counts (Trajtenberg, 1990). Column (2) shows that using this weighting leads to qualitatively similar results. Citation intensity also varies considerably across time and industries. In Column (3) we normalize each patent’s (three-year) citation count by the mean citation count for patents granted in the same year and in the same technology class. This again yields similar results. Finally, as discussed in Section 2.2.3, our IPO indicator variable may be too narrow

a measure of success. Therefore, we define a broader success indicator variable equal to one in the case of an IPO or acquisition. Column (4) shows that the treatment effect is again similar, albeit larger in magnitude.

#### **4.4.6 Industry-Specific Local Shocks**

Finally, we refine our baseline specification by allowing local shocks to be industry specific. Doing so accounts for the possibility that the new airline routes are introduced in response to local shocks that are specific to a particular industry. In terms of the regression specification, instead of including MSA by year fixed effects in Equation (1), we now include MSA by industry by year fixed effects (for both the MSAs of the portfolio company and the VC). We partition industries according to the six industry major groups of VentureXpert. The results are presented in Panel F of Table 7. As is shown, the estimates are very similar to our baseline coefficients in Table 3. However, the significance of the treatment effect is lower for all dependent variables (the treatment effect is even marginally insignificant for the IPO indicator). This is not surprising given that the additional layer of industry fixed effects reduces the power of our tests.

## **5 Regional Analysis**

In this section, we extend our analysis to study whether proximity fosters VC flows between regions—to the extent that travel time affects performance outcomes, it likely also affects VCs’ investment decisions. To address this question, we examine whether the introduction of new airline routes lead to increased VC investments between MSAs. The results are presented in Table 8. They are obtained by estimating variants of Equation (2). Observations are at the MSA-pair by year level and all regressions include MSA-pair and year fixed effects. Column (1) of Panel A shows the effect of the introduction of new airline routes between

pairs of MSAs on total VC investment (in logs).<sup>16</sup> The coefficient on the treatment indicator is 0.114, and statistically highly significant. This implies that total investment increases by 11.4% following the treatment. In Column (2), we account for the possibility of local shocks by including the two sets of MSA by year fixed effects. As can be seen, local shocks are an important determinant of VC investments across MSAs and hence accounting for them leads to a smaller treatment effect: the coefficient is now 0.046, corresponding to a 4.6% increase in total VC investment. Importantly, even after controlling for local shocks, the treatment effect remains highly significant and economically important. This finding indicates that better airline connections foster flows of VC investments between MSAs.

In Columns (3)-(6) of Panel A, we decompose total investments into initial investments (extensive margin) and follow-up investments (intensive margin). After controlling for local shocks, the treatment effect is 2.2% and 4.0%, respectively. Both coefficients are significant. Thus, better airline connections lead to higher VC investment along both the extensive and intensive margin. The increase in investment at the intensive margin suggests that proximity not only facilitates the screening of portfolio companies, but also their monitoring after the initial investment—arguably, VCs are more likely to expand their investment in companies they can monitor more easily. Accordingly, the latter finding complements our analysis of performance outcomes at the relationship level. A reduction in monitoring costs (holding selection fixed) not only leads to increased innovation and a higher likelihood of going public, but also to higher follow-up investments.

In Panel B, we explore alternative dependent variables that capture the intensity of VC activity following the treatment. In Columns (1) and (2), the dependent variable is the number of deals (in logs), in Columns (3) and (4), it is an indicator variable equal to one if any VC investment occurs between the two MSAs. After accounting for local shocks, we find

---

<sup>16</sup>Total VC investment is obtained by aggregating VC investment (i.e., VC funding) at the VC-company level into the MSA-pair level.

that the number of deals increases by 3.2%, and the likelihood of any VC activity increases by 2.5%.

In Table 9 we examine the dynamic effects of the treatment. As in the relationship analysis, we do so by replacing the treatment indicator with a set of four indicator variables representing the years around the treatment. We observe a very similar pattern for all dependent variables. The effect is small and insignificant in the year preceding the treatment (year  $-1$ ), which suggests that there are no pre-existing trends in the data. In the year of the treatment (year  $0$ ), we find that the treatment effect is positive, but relatively small and insignificant. It is only in the first year after the treatment (year  $1$ ) that the effect becomes large and significant. It remains somewhat stable thereafter (year  $2+$ ). This pattern suggests that it takes about 1 to 2 years for the new airline routes to translate into higher flows of VC investment between MSAs.

## 6 Conclusion

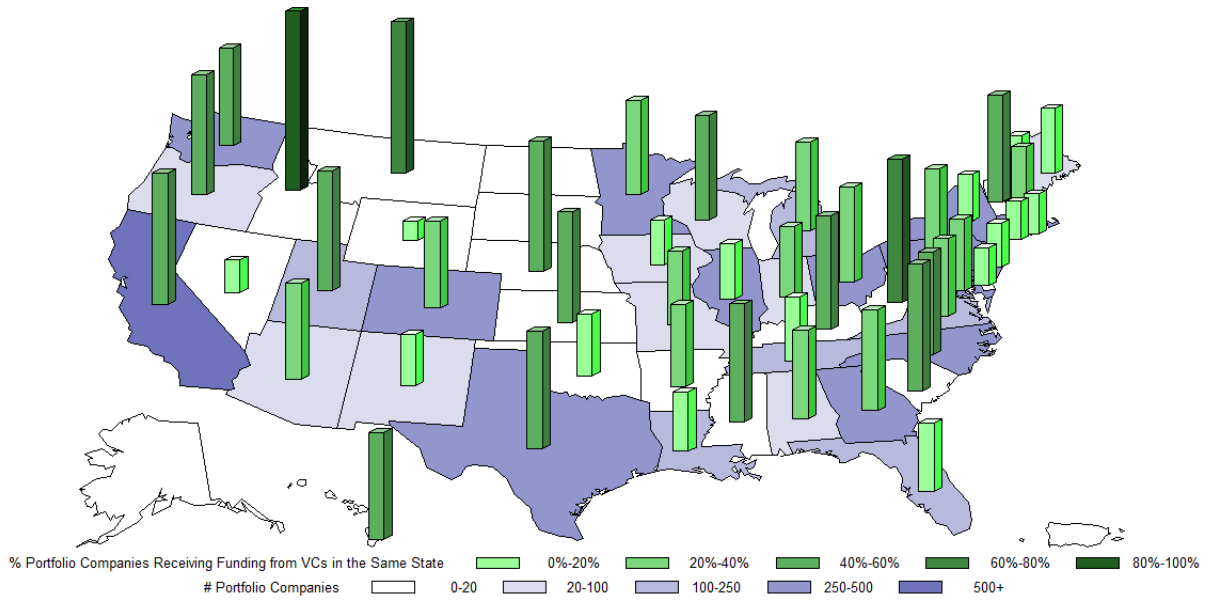
Do VCs contribute to the innovation and success of their portfolio companies, or do they simply identify and invest in companies that are already poised to innovate and succeed even absent their involvement? Our results suggest that VC involvement does matter. Specifically, we exploit exogenous reductions in monitoring costs stemming from the introduction of new airline routes that reduce the travel time between VCs and their existing portfolio companies, thereby holding company selection fixed. If differences in outcomes for portfolio companies are driven only by selection, reductions in monitoring costs subsequent to selection should have no effect. On the other hand, if VC activities do matter, reductions in monitoring costs should translate into better portfolio company performance by allowing VCs to engage in more of these activities.

We find that, within an existing relationship, reductions in travel time lead to an increase in the number of patents and number of citations per patent of the portfolio company, as well as an increase in the likelihood of an IPO or acquisition. These results are robust to controlling for local shocks that could potentially drive the introduction of the new airline routes. We further document that the effect is concentrated in routes that connect lead VCs (as opposed to other investors) with portfolio companies. In addition, we find that the effect is stronger for larger reductions in travel time.

Finally, we also study whether travel time affects aggregate venture capital flows between regions. We find that the introduction of new airline routes between two MSAs leads to an increase in VC investment along both the intensive and extensive margin. This suggest that policies encouraging new airline routes to VC hubs may be an effective tool for stimulating VC activity.


**Figure 1**  
**VC-Company Pairs**

This figure shows the distribution of portfolio companies across states graphically, where darker states are those with more portfolio companies. The height of the bars indicates the percentage of companies funded by a lead VC in the same state.



**Figure 2**  
**CDF of Distance Distribution**

This figure plots the cumulative density function (CDF) of the VC-company distance distribution.





**Table 1**  
**Sample Composition**

This table shows the composition of portfolio companies and VC firms in the sample. Portfolio companies are categorized as “Never Treated” if they never experienced a reduction in travel time to their lead VC investor, and “Ever Treated” otherwise. Similarly, VC firms are categorized as “Never Treated” if they never experienced a reduction in travel time to any of the companies in their portfolio (for which they were a lead investor), and “Ever Treated” otherwise. Panel A shows the company region distribution. Panel B shows the company industry distribution. Panel C shows the VC region distribution.

**Panel A: Company Region**

	Never Treated		Ever Treated		All	
	Freq	Percent	Freq	Percent	Freq	Percent
Alaska/Hawaii	22	0.10	1	0.09	23	0.10
Great Lakes	1054	4.81	51	4.66	1105	4.81
Great Plains	738	3.37	44	4.02	782	3.40
Mid-Atlantic	1178	5.38	59	5.39	1237	5.38
N. California	5464	24.96	146	13.35	5610	24.41
New England	2529	11.55	115	10.51	2644	11.50
New York Tri - State	2355	10.76	90	8.23	2445	10.64
Northwest	854	3.90	48	4.39	902	3.92
Ohio Valley	1169	5.34	59	5.39	1228	5.34
Rocky Mountains	875	4.00	44	4.02	919	4.00
S. California	1980	9.04	120	10.97	2100	9.14
South	432	1.97	67	6.12	499	2.17
Southeast	1475	6.74	121	11.06	1596	6.94
Southwest	1740	7.95	129	11.79	1869	8.13
US Territories	27	0.12	0	0	27	0.12
Total	21892	100.00	1094	100.00	22986	100.00

**Panel B: Company Industry**

	Never Treated		Ever Treated		All	
	Freq	Percent	Freq	Percent	Freq	Percent
Biotechnology	1221	5.58	70	6.40	1291	5.62
Communications and Media	2243	10.25	109	9.96	2352	10.23
Computer Hardware	1307	5.97	75	6.86	1382	6.01
Computer Software and Services	4526	20.67	192	17.55	4718	20.53
Consumer Related	1428	6.52	91	8.32	1519	6.61
Industrial/Energy	1222	5.58	77	7.04	1299	5.65
Internet Specific	4137	18.90	135	12.34	4272	18.59
Medical/Health	2329	10.64	144	13.16	2473	10.76
Other Products	1955	8.93	124	11.33	2079	9.04
Semiconductors/Other Elect.	1524	6.96	77	7.04	1601	6.97
Total	21892	100.00	1094	100.00	22986	100.00

**Table 1**  
**(Continued)**

**Panel C: VC Region**

	Never Treated		Ever Treated		All	
	Freq	Percent	Freq	Percent	Freq	Percent
Alaska/Hawaii	4	0.15	0	0	4	0.13
Great Lakes	174	6.65	38	7.04	212	6.71
Great Plains	90	3.44	29	5.37	119	3.77
Mid-Atlantic	126	4.81	34	6.30	160	5.07
N. California	502	19.17	60	11.11	562	17.80
New England	210	8.02	84	15.56	294	9.31
New York Tri - State	615	23.49	129	23.89	744	23.56
Northwest	67	2.56	9	1.67	76	2.41
Ohio Valley	143	5.46	34	6.30	177	5.60
Rocky Mountains	82	3.13	13	2.41	95	3.01
S. California	204	7.79	27	5.00	231	7.31
South	58	2.22	20	3.70	78	2.47
Southeast	145	5.54	26	4.81	171	5.41
Southwest	196	7.49	37	6.85	233	7.38
US Territories	2	0.08	0	0	2	0.06
Total	2618	100.00	540	100.00	3158	100.00

**Table 2**  
**Summary Statistics**

This table shows summary statistics for our main variables. Observations are shown at the level at which variables vary and are broken down by those that are “Never Treated” and those that are “Ever Treated,” as defined in Table 1. Great circle distance is the distance (in miles) between the VC’s ZIP code and the company’s ZIP code. Travel time is the amount of time (in minutes) it takes to travel from the VC’s ZIP code to the company’s ZIP code (round trip) based on the optimal itinerary and means of transportation (see Appendix B). Change in travel time is the reduction in travel time that occurs due to the treatment. Patents is the raw patent count, citations per patent is the number of citations garnered per patent in the three years after being granted, investment is the funding the portfolio company receives from all VCs in a given year. VC firm experience is measured as the number of years since firm founding, the number of companies invested in to date, and the number of investments that have gone public to date.

	Never Treated			Ever Treated		
	Obs	Mean	Std Dev	Obs	Mean	Std Dev
<i>Company-VC Pair Level:</i>						
Great Circle Distance (Miles)	30373	735.89	931.84	1131	1236.13	845.38
Travel Time (Minutes)	30373	470.22	551.17	1131	719.82	252.37
Change in Travel Time (Minutes)	—	—	—	1131	126.18	87.57
<i>Company-Year Level:</i>						
Patents	111959	0.44	6.37	9293	0.28	1.28
Citations Per Patent	111959	1.43	7.89	9293	1.03	6.09
Investment (Millions)	111959	3.28	10.86	9293	1.70	7.14
<i>VC-Year Level:</i>						
Experience (Years)	17404	11.00	13.43	8554	14.98	12.16
Experience (Companies)	17404	16.18	27.28	8554	53.85	74.36
Experience (IPOs)	17404	1.94	5.21	8554	8.26	15.21



**Table 4****Relationship Analysis: Dynamics**

This table shows the dynamics of the treatment effects in the relationship analysis. All variables are defined as in Table 3. The variable Treatment(-1) is an indicator variable equal to one if the observation is recorded in the year preceding the treatment. Treatment(0), Treatment(1), and Treatment(2+) are defined analogously with respect to the year of the treatment, the first year after the treatment, and two or more years after the treatment, respectively. Standard errors, clustered by portfolio company, are shown in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment(-1)	0.00639 (0.0147)	0.0170 (0.0285)	
Treatment(0)	0.0165 (0.0155)	0.0244 (0.0283)	0.00682 (0.00502)
Treatment(1)	0.0391** (0.0182)	0.0690** (0.0333)	0.00805 (0.00644)
Treatment(2+)	0.0494*** (0.0182)	0.106*** (0.0326)	0.0158** (0.00655)
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
MSA(VC) $\times$ Year FE	Yes	Yes	Yes
MSA(Company) $\times$ Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.668	0.576	0.494
Observations	130169	130169	130169

**Table 5****Relationship Analysis: Non-Lead VCs**

This table repeats the analysis of Table 3, but limiting the sample to company-VC pairs that do not involve a lead investor. Non-lead VCs located in the same MSA as the lead VC are also excluded from the sample. Standard errors, clustered by portfolio company, are shown in parentheses. \* , \*\* , and \*\*\* denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment	-0.0128 (0.0203)	-0.0205 (0.0368)	0.00761 (0.00691)
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
MSA(VC) $\times$ Year FE	Yes	Yes	Yes
MSA(Company) $\times$ Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.758	0.688	0.673
Observations	90609	90609	90609

**Table 6****Relationship Analysis: Intensity of the Treatment**

This table repeats the analysis of Table 3, but separating the treatment indicator into two variables. Treatment  $\times$  Large is an indicator variable equal to one if the treatment is associated with a travel time reduction of at least 60 minutes. Treatment  $\times$  Small is an indicator variable equal to one if the treatment is associated with a travel time reduction of less than 60 minutes. Standard errors, clustered by portfolio company, are shown in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment $\times$ Large	0.0336** (0.0143)	0.0684*** (0.0248)	0.0115** (0.00524)
Treatment $\times$ Small	0.0259 (0.0173)	0.0359 (0.0333)	0.00822 (0.00683)
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
MSA(VC) $\times$ Year FE	Yes	Yes	Yes
MSA(Company) $\times$ Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.668	0.576	0.494
Observations	130169	130169	130169

**Table 7****Relationship Analysis: Robustness**

Panel A of this table repeats the analysis of Table 3, but separating the treatment indicator into two variables. Treatment (Hub or Merger) is an indicator variable equal to one if the treatment is due to the opening of a new airline hub, or the merger of two airlines. Treatment (Other) is an indicator variable equal to one if the treatment is not due to a hub opening or merger. Panel B restricts the control group to those control pairs whose average distance matches the average distance in the treatment group. That is, we exclude short-distance control pairs so that the average distance is the same in both groups. Panel C controls for heterogeneous time trends by interacting baseline characteristics (funding, patents, experience, and distance) with year fixed effects. Panel D restricts the sample to the eventually treated pairs. Panel E uses alternative definitions of the dependent variables. HJT CPP adjusts for truncation in citations per patent by using the estimated shape of the citation-lag distribution following Hall et al. (2001). HJT WPC represents citation-weighted patent counts (Trajtenberg, 1990), again using the HJT method to adjust for citation truncation. Relative CPP normalizes 3-year citations per patent by the mean citations per patent for other patents granted in the same year and technology class. Success is an indicator variable equal to one if the company has an IPO or is acquired that year. Panel F controls for MSA-industry-year fixed effects. Industries are partitioned according to VentureXpert's major industry groups. Standard errors, clustered by portfolio company, are shown in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% level, respectively.

**Panel A: Hub Openings and Airline Mergers**

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment (Hub or Merger)	0.0540** (0.0255)	0.116** (0.0508)	0.0237* (0.0142)
Treatment (Other)	0.0273** (0.0126)	0.0475** (0.0219)	0.00842* (0.00433)
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
MSA(VC) $\times$ Year FE	Yes	Yes	Yes
MSA(Company) $\times$ Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.668	0.576	0.494
Observations	130169	130169	130169

**Panel B: Distance-Matched Control Sample**

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment	0.0382*** (0.0126)	0.0660*** (0.0226)	0.00923** (0.00455)
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
MSA(VC) $\times$ Year FE	Yes	Yes	Yes
MSA(Company) $\times$ Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.687	0.595	0.542
Observations	77129	77129	77129



**Table 7**  
(Continued)

**Panel C: Heterogeneous Time Trends**

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment	0.0322*** (0.0117)	0.0593*** (0.0209)	0.0102** (0.00437)
Baseline Characteristics × Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
MSA(VC) × Year FE	Yes	Yes	Yes
MSA(Company) × Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.668	0.577	0.497
Observations	130169	130169	130169

**Panel D: Eventually Treated Pairs**

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment	0.0314*** (0.0107)	0.0354* (0.0207)	0.0250*** (0.00414)
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.582	0.440	0.218
Observations	7978	7978	7978

**Panel E: Alternative Dependent Variables**

	(1) HJT CPP	(2) HJT WPC	(3) Relative CPP	(4) Success
Treatment	0.0860*** (0.0268)	0.107*** (0.0325)	0.0295*** (0.00922)	0.0252*** (0.00840)
Controls	Yes	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
MSA(VC) × Year FE	Yes	Yes	Yes	Yes
MSA(Company) × Year FE	Yes	Yes	Yes	Yes
R <sup>2</sup>	0.589	0.640	0.567	0.424
Observations	130169	130169	130169	130169

**Table 7**  
(Continued)

**Panel F: Industry-Specific Local Shocks**

	(1) Patents	(2) Citations/Patent	(3) IPO
Treatment	0.0339** (0.0133)	0.0551** (0.0261)	0.00738 (0.00484)
Controls	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
MSA(VC) × Industry × Year FE	Yes	Yes	Yes
MSA(Company) × Industry × Year FE	Yes	Yes	Yes
R <sup>2</sup>	0.743	0.653	0.611
Observations	130169	130169	130169

**Table 8****Regional Analysis: Main Regressions**

This table shows the main results of the regional analysis. Observations are at the MSA-pair by year level. Only MSA pairs that ever have venture capital flows between them are included in the sample. Treatment is an indicator variable equal to one if a direct flight has been introduced between the two MSAs. Total investment is the log of (one plus) the total amount invested by VCs in the source MSA to companies in the target MSA. Initial investment represents investment in new companies. Follow-up investment represents investment in existing companies. Number of deals represents the number of rounds of funding closed between VCs in the source MSA and companies in the target MSA. VC activity is an indicator variable equal to one if any VC from the source MSA invested in a company in the target MSA that year. Standard errors, clustered by MSA-pair, are shown in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% level, respectively.

**Panel A: Investment**

	Total Investment		Initial Investment		Follow-up Investment	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.114*** (0.0215)	0.0455*** (0.0171)	0.0486*** (0.0116)	0.0215** (0.0103)	0.0981*** (0.0208)	0.0398** (0.0167)
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
MSA(VC) $\times$ Year FE	No	Yes	No	Yes	No	Yes
MSA(Company) $\times$ Year FE	No	Yes	No	Yes	No	Yes
R <sup>2</sup>	0.499	0.618	0.378	0.468	0.477	0.602
Observations	182970	182970	182970	182970	182970	182970

**Panel B: Deals**

	Number of Deals		VC Activity	
	(1)	(2)	(3)	(4)
Treatment	0.0827*** (0.0134)	0.0318*** (0.0122)	0.0618*** (0.00728)	0.0248*** (0.00742)
Pair FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
MSA(VC) $\times$ Year FE	No	Yes	No	Yes
MSA(Company) $\times$ Year FE	No	Yes	No	Yes
R <sup>2</sup>	0.612	0.693	0.363	0.463
Observations	182970	182970	182970	182970

**Table 9****Regional Analysis: Dynamics**

This table shows the dynamics of the treatment effects in the regional analysis. All variables are defined as in Table 8. The variable Treatment(-1) is an indicator variable equal to one if the MSA-pair observation is recorded in the year preceding the treatment. Treatment(0), Treatment(1), and Treatment(2+) are defined analogously with respect to the year of the treatment, the first year after the treatment, and two or more years after the treatment, respectively. Standard errors, clustered by MSA-pair, are shown in parentheses. \*, \*\*, and \*\*\* denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)
	Total Inv	Initial Inv	Follow-up Inv	Num Deals	VC Activity
Treatment(-1)	0.00461 (0.0180)	0.00535 (0.0115)	0.00898 (0.0171)	0.00641 (0.0139)	-0.00376 (0.0106)
Treatment(0)	0.0218 (0.0197)	0.0132 (0.0125)	0.0246 (0.0185)	0.0226 (0.0148)	0.0162 (0.0116)
Treatment(1)	0.0524*** (0.0196)	0.0234* (0.0137)	0.0477** (0.0187)	0.0414*** (0.0149)	0.0196* (0.0117)
Treatment(2+)	0.0484** (0.0200)	0.0232* (0.0121)	0.0423** (0.0195)	0.0328** (0.0142)	0.0258*** (0.00863)
Pair FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
MSA(VC) × Year FE	Yes	Yes	Yes	Yes	Yes
MSA(Company) × Year FE	Yes	Yes	Yes	Yes	Yes
R <sup>2</sup>	0.618	0.468	0.602	0.693	0.463
Observations	182970	182970	182970	182970	182970

## References

- Aghion, Philippe, John Van Reenen, and Luigi Zingales, 2013, Innovation and institutional ownership, *American Economic Review*, 103, 277–304.
- Bertrand, Marianne, and Sendhil Mullainathan, 2003, Enjoying the quiet life? corporate governance and managerial preferences, *Journal of Political Economy*, 111, 1043–1075.
- Bottazzi, Laura, Marco Da Rin, and Thomas Hellmann, 2008, Who are the active investors? evidence from venture capital, *Journal of Financial Economics*, 89, 488–512.
- Chemmanur, Thomas J., Karthik Krishnan, and Debarshi K. Nandy, 2011, How does venture capital financing improve efficiency in private firms? a look beneath the surface, *Review of Financial Studies*, 24, 4037–4090.
- Chen, Henry, Paul Gompers, Anna Kovner, and Josh Lerner, 2010, Buy local? the geography of venture capital, *Journal of Urban Economics*, 67, 90–102.
- Coval, Joshua D., and Tobias J. Moskowitz, 1999, Home bias at home: Local equity preference in domestic portfolios, *Journal of Finance*, 54, 2045–2073.
- Coval, Joshua D., and Tobias J. Moskowitz, 2001, The geography of investment: Informed trading and asset prices, *Journal of Political Economy*, 109, 811–841.
- Feyrer, James, 2009, Trade and income—exploiting time series in geography, NBER Working Paper 14910.
- Giroud, Xavier, 2013, Proximity and investment: Evidence from plant-level data, *Quarterly Journal of Economics*, 128, 861–915.
- Gladstone, David, 1988, *Venture Capital Handbook* (Englewood Cliffs, NJ: Prentice Hall).
- Gompers, Paul A., 1996, Grandstanding in the venture capital industry, *Journal of Financial Economics*, 42, 133–156.
- Gorman, Michael, and William A. Sahlman, 1989, What do venture capitalists do?, *Journal of Business Venturing*, 4, 231–248.
- Gormley, Todd A., and David A. Matsa, 2013, Common errors: How to (and not to) control for unobserved heterogeneity, *Review of Financial Studies*, forthcoming.
- Guimarães, Paulo, and Pedro Portugal, 2010, A simple feasible procedure to fit models with high-dimensional fixed effects, *Stata Journal*, 10, 628–649.
- Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg, 2001, The NBER patent citations data file: Lessons, insights and methodological tools, NBER Working Paper 8498.
- Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg, 2005, Market value and patent citations, *RAND Journal of Economics*, 36, 16–38.

- Hellmann, Thomas, and Manju Puri, 2000, The interaction between product market and financing strategy: The role of venture capital, *Review of Financial Studies*, 13, 959–984.
- Hellmann, Thomas, and Manju Puri, 2002, Venture capital and the professionalization of start-up firms: Empirical evidence, *Journal of Finance*, 57, 169–197.
- Jaffe, Adam B., and Manuel Trajtenberg, 2002, *Patents, Citations, and Innovations: A Window on the Knowledge Economy* (Cambridge and London: MIT Press).
- Kaplan, Steven N., and Per Strömberg, 2001, Venture capitalists as principals: Contracting, screening, and monitoring, *American Economic Review*, 91, 426–430.
- Kaplan, Steven N., and Per Strömberg, 2004, Characteristics, contracts, and actions: Evidence from venture capitalist analyses, *Journal of Finance*, 59, 2177–2210.
- Kaplan, Steven N., Per Strömberg, and Berk A. Sensoy, 2002, How well do venture capital databases reflect actual investments?, Working Paper, University of Chicago.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2012, Technological innovation, resource allocation, and growth, NBER Working Paper 17769.
- Kortum, Samuel, and Josh Lerner, 2000, Assessing the contribution of venture capital to innovation, *RAND Journal of Economics*, 31, 674–692.
- Lanjouw, Jean O., Ariel Pakes, and Jonathan Putnam, 1998, How to count patents and value intellectual property: The uses of patent renewal and application data, *Journal of Industrial Economics*, 46, 405–432.
- Lerner, Josh, 1995, Venture capitalists and the oversight of private firms, *Journal of Finance*, 50, 301–318.
- Lerner, Josh, 2009, *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed—and What to Do About It* (Kauffman Foundation Series on Innovation and Entrepreneurship. Princeton and Oxford: Princeton University Press).
- Lerner, Josh, Morten Sorensen, and Per Strömberg, 2011, Private equity and long-run investment: The case of innovation, *Journal of Finance*, 66, 445–477.
- Petersen, Mitchell A., and Raghuram G. Rajan, 2002, Does distance still matter? the information revolution in small business lending, *Journal of Finance*, 57, 2533–2570.
- Puri, Manju, and Rebecca Zarutskie, 2012, On the life cycle dynamics of venture-capital- and non-venture-capital-financed firms, *Journal of Finance*, 67, 2247–2293.
- Quindlen, Ruthann, 2000, *Confessions of a Venture Capitalist* (New York, NY: Warner).
- Seru, Amit, 2014, Firm boundaries matter: Evidence from conglomerates and R&D activity, *Journal of Financial Economics*, 111, 381–405.

Sorensen, Morten, 2007, How smart is smart money? a two-sided matching model of venture capital, *Journal of Finance*, 62, 2725–2762.

Trajtenberg, Manuel, 1990, A penny for your quotes: Patent citations and the value of innovations, *RAND Journal of Economics*, 21, 172–187.

# Appendix

## A Matching VentureXpert with NBER Patent Data

### A.1 Name Standardization

In order to match VentureXpert with data from the NBER Patent Project, we begin by standardizing the company names in both, using the name standardization routines developed by the NBER Patent Data Project to create a bridge file to COMPUSTAT.<sup>17</sup> These routines standardize common company prefixes and suffixes building on a list created by Derwent World Patent Index (Thomson-Reuters); they also identify a company's stem name excluding these prefixes and suffixes. Similarly, we standardize the location names from both datasets. This is done to correct for spelling errors as well as other types of errors that commonly occur, particularly in the patent data. For example, in some cases a neighborhood name is used rather than the name of a city. In other cases country codes are listed as state codes, e.g. a patent assignee from Germany (DE) may be coded as being from Delaware (DE). The city name standardization is done by running all location names through the Google Maps API, which automatically corrects close, but inaccurate text representations of location names and returns a standardized name broken down into its component parts (city, state, country), along with latitude and longitude information.

### A.2 The Matching Procedure

With these standardized company and city names we then use the following matching procedure:

1. Each standardized name associated with a company in VentureXpert is matched with

---

<sup>17</sup><https://sites.google.com/site/patentdatapoint/>



standardized names from the NBER data.<sup>18</sup> If an exact match is found, this is taken to be the same company and hence it is removed from the set of names that needs to be matched.

2. For the remaining companies in VentureXpert, each stem name associated with a company is matched with stem names from the NBER data. If an exact match is found and enough other identifying information matches as well, this is taken to be the same company and it is removed from the set of names that need to be matched. If an exact match is found, but not enough other identifying information matches as well, the match is added to a list of borderline matches to be checked manually.

- (a) For a stem match to be considered definite, the standardized city/state combination also has to match, or the state has to match along with the time period (first patent application was after the company founding year).

3. For the remaining companies in VentureXpert, each stem name associated with a company is matched with up to 10 close stem names from the NBER data using a padded bi-gram comparator. Fuzzy matches with match quality between 1.5 and 2 that also had a city/state match were kept for review, as were fuzzy matches with quality above 2 with only a state match.
4. The borderline matches identified using the above procedure were reviewed by hand, now also using other qualitative information from both data sources, including full patent abstracts, and paragraph-long company descriptions.

---

<sup>18</sup>Many companies have multiple names listed in VentureXpert, reflecting the fact that young companies often change their name as they mature.

## B Measuring Travel Time

The procedure to compute travel times between VC firms and portfolio companies is the same as in Giroud (2013). The core of the algorithm is done using Visual Basic in the MS Mappoint software. Importantly, the results are not sensitive to the various assumptions listed below. The algorithm goes as follows:

1. Using MS Mappoint, we first compute the travel time by car (in minutes) between the two ZIP codes. This travel time is used as a benchmark and is compared to the travel time by air based on the fastest airline route. Whenever traveling by car is faster, air transportation is ruled out by optimality, and the relevant travel time is the driving time by car.
2. To determine the fastest airline route between any two ZIP codes, we use the itinerary information from the T-100 and ER-586 data. The fastest airline route minimizes the total travel time between the VC and the company. The total travel time consists of three components: (1) the travel time by car between the VC and the origin airport; (2) the duration of the flight, including the time spent at airports and, for indirect flights, the layover time; and (3) the travel time by car between the destination airport and the company. The travel time by car to and from airports is obtained from MS Mappoint. Flight duration per segment is obtained from the T-100 and ER-586 data, which include the average ramp-to-ramp time of all flights performed between any two airports in the United States. The only unobservable quantities are the time spent at airports and the layover time. We assume that one hour is spent at the origin and destination airports combined and that each layover takes one hour.
3. Additional assumptions we made are as follows:

- (a) If the distance between the two ZIP codes is less than 100 miles, driving is always optimal.
- (b) A new route dominates a previous one if the time saving is more than 15 minutes one-way (i.e., 30 minutes round-trip).
- (c) In the data, we also “smoothed” the optimal itinerary by keeping the previously optimal route if a new route is introduced but does not dominate the current route (e.g., a new flight from LGA instead of JFK with a saving of merely 5 minutes).