

The Impact of Venture Capital Monitoring

SHAI BERNSTEIN, XAVIER GIROUD, and RICHARD R. TOWNSEND*

ABSTRACT

We show that venture capitalists' (VCs) on-site involvement with their portfolio companies leads to an increase in (1) innovation and (2) the likelihood of a successful exit. We rule out selection effects by exploiting an exogenous source of variation in VC involvement: the introduction of new airline routes that reduce VCs' travel times to their *existing* portfolio companies. We confirm the importance of this channel by conducting a large-scale survey of VCs, of whom almost 90% indicate that direct flights increase their interaction with their portfolio companies and management, and help them better understand companies' activities.

*Shai Bernstein is with Stanford University, Graduate School of Business; Xavier Giroud is with the MIT Sloan School of Management, NBER, and CEPR; Richard Townsend is with the Tuck School of Business at Dartmouth College. We are grateful to Michael Roberts (the editor), two anonymous referees, Bo Becker, José-Miguel Gaspar, Yael Hochberg, Dirk Jenter, Anna Kovner, Josh Lerner, Laura Lindsey, Holger Mueller, Francisco Pérez-González, David Robinson, Antoinette Schoar, Rob Seamans, Scott Stern, as well as seminar participants at Columbia (GSB), Dartmouth (Tuck), Duisenberg School of Finance and Tinbergen Institute, EFA (Lugano), EPFL and University of Lausanne, Economics of Entrepreneurship and Innovation Conference (Queen's), HKUST, Harvard (HBS), LBS Private Equity Findings Symposium, MIT (Sloan), MIT Micro@Sloan Conference, NBER Entrepreneurship Meetings, Nanyang Technological University, Northeastern, Northwestern (Kellogg), Norwegian School of Economics, SFS Finance Cavalcade (Georgetown), Singapore Management University, Stanford (GSB), UCSD (Rady), UNC (Kenan-Flagler), University of Bergen, University of Illinois at Chicago, University of Maryland (Smith), and University of Miami School of Business for helpful comments and suggestions.

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.

An ideal experiment to establish the impact of VCs would be to randomly provide certain companies with VC funding and others not. Such an experiment would eliminate the selection of companies (“screening”), thus allowing us to estimate the effect of VC involvement (“monitoring”).¹ 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 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). 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.

To obtain direct evidence on whether VC involvement increases following reductions in travel time, we conduct a large-scale survey of VC investors. Almost 90% of the 306 survey participants agreed that they would visit a portfolio company more frequently following the introduction of a direct flight. Survey participants also agreed that the introduction of a direct flight would help them

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

establish better relationships with management teams, better understand the state of their companies, and generally add more value. This qualitative evidence supports our underlying assumption that VC involvement is responsive to the introduction of direct flights, and is consistent with the academic literature that shows that VC activity is sensitive to geographic proximity.²

We then 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 1.4%. 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 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. Moreover, we find that pre-existing trends are not driving our results, and the 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.

We provide further evidence on the underlying channel through which these effects operate by

²For example, Lerner (1995) finds that VCs are more likely to sit on boards of geographically proximate companies. 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.

taking 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.

Our paper contributes to a growing literature that studies the effect of VCs on portfolio company outcomes. 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; Chemmanur et al., 2011a; Hellmann and Puri, 2002; 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.

Our paper also contributes to a large, mostly theoretical literature that explores how financial contracts shape the interaction between entrepreneurs and VC firms, alleviating moral hazard and agency problems. For example, several papers consider the optimal contractual arrangement that leads both entrepreneurs and VCs to contribute effort to promote a venture’s success in a double

moral hazard setting (e.g., Schmidt, 2003; Casamatta, 2003; Inderst and Mueller, 2004; Hellmann, 2006). Other theoretical work highlights the importance of contractual arrangements on the VC refinancing versus termination decision (e.g., Repullo and Suarez, 2004; Dessi, 2005; Cornelli and Yosha, 2003), as well as the effective allocation of control (e.g., Berglof, 1994; Cestone, 2014). Kaplan and Strömberg (2003; 2004) provide empirical evidence on such contractual arrangements.³ Our paper complements this literature by highlighting the role of geographic proximity, in addition to contracts, in shaping the interactions between entrepreneurs and VC firms.

The remainder of this paper is organized as follows. Section I discusses the data and key variables. Section II discusses our empirical strategy. Section III presents the survey. Section IV presents the results, and Section V concludes.

I. Data

A. *Data 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

³For a comprehensive review of this literature, see Da Rin et al. (2013).

offices, this should not present a systematic source of bias.⁴ Similarly, we only observe one head-quarter location per portfolio company, despite the fact that portfolio companies can potentially move. The location we observe represents the company’s latest known address. Again, this should not present a systematic source of bias. 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 restrict the sample to only VC-company pairs involving the lead investor, which will be defined in Section I.B.3. In subsequent analysis, we examine whether the results hold for non-lead investors as well.

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 2006.⁵ 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 Section II of the Internet Appendix. Finally, we also supplement the NBER data with citation data from Google 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.

⁴If the monitoring is done out of local offices, not accounting for them would merely go against us finding any effect.

⁵In 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).

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 I 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 I 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.

B. Definitions of Variables

B.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 Section III of the Internet Appendix. During our sample period (1977 to 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. Moreover, a 126-minute travel time reduction could mean

the difference between being able to fly back on the same day versus having to stay overnight.⁶

B.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.⁷ 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 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.

⁶In addition, indirect flights may induce other types of disutility, e.g., anxiety about missing a connection or fatigue due to longer time in transit (e.g., Boeh and Beamish, 2011, 2012)

⁷Recent examples include Lerner et al. (2011); Aghion et al. (2013); Seru (2014).

B.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 may capture 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. Indeed, Metrick and Yasuda (2011) find that first-round (second-round) investors do not even recover their invested capital in 38% (46%) of acquisitions. Unfortunately, due to data limitations, we cannot calculate gross value multiples analogously. Therefore, to ensure that we only count significant acquisitions as positive outcomes, we obtain acquisition values from SDC Platinum and CapitalIQ.⁸ We then include only acquisitions at values over \$25M (in 2000 dollars) in our success measure.

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.⁹ 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.

⁸These two databases are merged with VentureXpert using standardized target names. We further require that the acquisition date be within 30 days of the date reported in VentureXpert. Using this methodology, we are able to match 66% of companies classified as acquired in VentureXpert. We assume that any acquisition that is not in SDC/CapitalIQ, or whose acquisition value is unknown in those two databases, is not significant.

⁹We 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.

II. Methodology

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.¹⁰ 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:

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

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., 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 (the number of days since the first round of financing) and company stage of development (a set of indicator variables for the 8-point stage classification used by VentureXpert); α_{ij} are VC-company pair fixed effects; $\alpha_{MSA(i)} \times \alpha_t$ 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; ϵ 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 β which measures the effect of the introduction of new airline routes on y .

¹⁰Note 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.

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 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.

A. 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 ($\alpha_{MSA(j)} \times \alpha_t$).¹¹¹²

B. 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, that is, those companies and VCs that are even less

¹¹In 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 (2014) for details.

¹²In robustness checks, we further show that our results are similar if we allow local shocks to be industry specific, that is, instead of including MSA by year fixed effects, we include the full set of MSA by industry by year fixed effects (see Section IV.E.3).

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 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.

C. 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 II. 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. Fourth, we show that our results also hold if we allow pairs that differ on the basis of the characteristics in Table II to be on different time trends. More precisely, this test is conducted by including as additional controls the characteristics in Table II 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). Again, we show that our results are robust if we perform this test.

III. Survey of VCs

The key assumption underlying our empirical strategy is that VCs are responsive to the treatment, i.e., VC involvement increases following a reduction in travel time. Since VC involvement is not observable, we cannot directly test this assumption. Instead, to assess the plausibility of this assumption, we conduct a large-scale survey of VCs.

Surveying VC investors is difficult because these investors are time constrained and also notoriously reluctant to provide data on their operations. In order to increase the likelihood of participation, we limited our survey population to alumni from our respective academic institutions (Stanford, MIT, and Dartmouth). In total, we identified 2,109 alumni with current or past VC experience. We distributed the survey electronically to these alumni and obtained 306 responses (corresponding to a response rate of 14.5%).¹³

As is typically the case with VC surveys, our sample is unlikely to be perfectly representative of the VC universe. Nonetheless, we see no reason to believe that the sample should be biased toward

¹³By way of comparison, Gorman and Sahlman (1989) obtained responses from 49 venture capitalists. More recently, Gompers et al. (2014) obtained responses from 79 buyout investors.

VCS whose monitoring is more sensitive to travel time reductions. The average assets under management (AUM) of the VC firms of survey participants was \$1.2 billion (median \$448 million).¹⁴ For comparison, the average assets under management of VC firms as reported by Thomson Reuters is \$213 million. This difference reflects the large representation of top VC firms in our survey. Panel A of Table III illustrates the geographical distribution of the survey participants. Approximately 20% were located outside the U.S. The most strongly represented international locations were Germany, Brazil, and China. Within the U.S., there was a clear bias toward California, which accounted for 70% of the participants. As expected, other prominent locations were Massachusetts and New York.

Panel B of Table III shows that 77% of the survey participants were partners at their VC firms. The average portfolio size among the respondents was 6.58 companies (median 5), and an average of 3.79 of these companies were local (median 3). A local company was defined as a company within 50 miles of the investor. The average number of visits to a given portfolio company was 9.93 per year (median 6). Interestingly, the survey participants reported that they spend 48% of their time monitoring and assisting portfolio companies. Moreover, 71% reported that they tend to visit local companies more than non-local companies, suggesting that proximity affects their level of involvement with a company. We explored this hypothesis more directly in the remainder of the survey.

A common issue in survey design is the possibility of social desirability bias (SDB). This refers to the tendency of research participants to present themselves in a positive or socially acceptable way (Maccoby and Maccoby, 1954). In the context of our survey, an important concern is that SDB may lead participants to be hesitant in revealing the effect of direct flights on their level of monitoring. Indeed, admitting that direct flights matter could suggest that respondents do not provide sufficient monitoring and assistance to distant companies. Rather, respondents might want to portray themselves as “always doing whatever is necessary to help their portfolio companies regardless of other factors.”

VCS may want to portray themselves in this way to maintain a positive image in their own mind,

¹⁴We define assets under management as the aggregate size of all non-liquidated funds.

or in the minds of others they imagine may get access to the survey results. Indeed, several VCs that we consulted when designing the survey independently brought up the concern that respondents may not be truthful for this reason. To help elicit truthful responses, we administered the survey in an anonymous manner and informed the participants that their identity could not be linked with their response. However, past research shows that anonymizing surveys has somewhat limited effectiveness in reducing SDB (Dillman et al., 1996; Singer et al., 1995). In our case, participants may worry that the online survey platform that we utilized does not completely strip identifying information (e.g., IP address) from their response. Therefore, in addition to anonymization, we also used the well-known technique of “indirect questioning” to further mitigate the possibility of SDB (Haire, 1950; Calder and Burnkrant, 1977; Anderson, 1978).

Specifically, in our first set of key questions, rather than asking research participants about their own behavior, we asked about their beliefs about general VC behavior. While this approach has been shown to mitigate SDB, one may be concerned that VCs incorrectly perceive the sensitivity of others to reductions in travel time. Therefore, we also asked a second set of key questions regarding VCs’ own behavior, recognizing that responses in this case may be more affected by SDB. For the second set of questions, we described a situation in which an indirect flight (Seattle to Raleigh-Durham via Chicago) is replaced by a non-stop flight.¹⁵

Finally, according to survey design conventions, key questions are generally asked in a variety of different but closely related ways. Among other things, this helps ensure that the results are not driven by participants misunderstanding a single question. This is also helpful to better understand the mechanism through which direct flights improve VC involvement with portfolio companies. Therefore, for both our general and specific questions, we asked several variations related to different dimensions along which direct flights may matter. That being said, our main interest is in whether VCs report that they are likely to spend more time at a company in person if a direct flight is introduced.

¹⁵We did not use the Boston to Salt Lake City example given in Section II, because our pre-testing indicated that participants found it confusing to think about a scenario in which there is no direct flight between Boston and Salt Lake City (since there are currently direct flights between them). In contrast, at the time of the survey, there was no direct flight between Seattle and Raleigh-Durham.

On all questions, participants were asked to state their degree of agreement with various statements about the effect of direct flights on VCs. We used a standard 6-point Likert scale, where potential responses ranged from *Strongly Disagree* ($= 1$) to *Strongly Agree* ($= 6$). We note that it is not clear what percentage of respondents would need to agree with these statements for our survey results to be consistent with our main empirical results. For example, if 30% of VCs become more involved when a direct flight is introduced, that may be enough for us to find a statistically significant effect on average. However, this point is largely irrelevant, as we find that a large majority of respondents agreed with all of the statements presented.

Panel C of Table III summarizes the responses to our general questions. The precise wording of the questions is shown in Section VI of the Internet Appendix. In terms of the main question, 86% of respondents agreed that direct flights between VCs and portfolio companies increase the time VCs spend at companies in person, with a mean response of 4.43 out of 6. The mean response was also statistically different from the neutral mid-point response of 3.5 at the 1% significance levels. We find that over 80% agreed with all of the remaining questions of the general variety and that the mean response was significantly non-neutral. For example, 83% agreed that direct flights allow VCs to more effectively advise companies, and 80% agreed that direct flights allow VCs to better understand key challenges and issues that portfolio companies are facing. The full distribution of responses to the indirect questions is depicted graphically in the first four rows of Figure 3.

Panel D of Table III summarizes the responses to our specific questions. Again, the precise wording of the questions is shown in Section VI of the Internet Appendix. In this case, 83% agreed that the introduction of a direct flight would increase the frequency with which they visit a portfolio company and 89% agreed that a direct flight would increase flexibility to visit a portfolio company when most useful. Interestingly, there was the least agreement (72%) that a direct flight would improve communications with the company. This is likely due to recent advances in communication technology. The remaining questions yielded agreement ranging from 75-81%, as survey participants agreed that the introduction of direct flights would help them establish better relationships with management teams, better understand the state of companies, and generally add more value. In all

cases, the mean response was again significantly non-neutral at the 1% level. The full distribution of responses to the indirect questions is depicted graphically in the last six rows of Figure 3.

Finally, in unreported tests, we also explored whether the sensitivity to travel time differs across different types of participants. Because the survey was anonymous, we can only partition the sample based on answers to the preliminary questions. We find greater agreement with all 10 statements for VC partners, those currently working at a VC firm, as well as those who manage a greater number of companies. However, in most cases, the differences are not statistically significant. Nonetheless, these results provide some comfort that our overall findings do not primarily reflect the views of low level associates or those that have not worked in the industry in many years.

To summarize, the survey results indicate that VCs are likely to spend more time at their portfolio companies following a reduction in travel time. This qualitative evidence supports our underlying assumption that VC involvement is responsive to the treatment.¹⁶

IV. Results

A. Main Results

Next, we estimate variants of Equation (1) to examine whether the introduction of new airline routes that reduce the travel time between lead VC firms and their portfolio companies affects portfolio companies' innovation and success. The results are presented in Table IV. In Columns (1)-(3) of Panel A, 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

¹⁶In Internet Appendix Table IA.I, we provide additional supporting evidence based on aggregate travel patterns. Specifically, we use data from the Airline Origin and Destination Survey (DB1B)—a 10% sample of airline tickets from reporting carriers collected by the U.S. Bureau of Transportation Statistics—to study whether general passenger flows between cities increase following the treatment. As is shown in the table, the treatment leads to a 14.5% to 15.5% increase in passenger flows. Thus, it does appear that general passengers are sensitive to reductions in travel time. This suggests that VCs may be sensitive as well.

that the number of patents increases by 3.1% to 3.7% after the treatment. In Columns (4)-(6) of Panel A, 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 (1)-(3) of Panel B, the dependent variable is an indicator 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%. Finally, in Columns (4)-(6) of Panel B, the dependent variable is an indicator equal to one if the company goes public or is acquired during the year (“Success”). As is shown, the success likelihood increases by 1.1% to 1.4% following the treatment. Overall, our findings indicate that a reduction in VC monitoring costs leads to significant increases in innovation and the likelihood of a successful exit.¹⁷

B. Dynamic Effects of the Treatment

In Table V, 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) and (6) of both panels of Table IV, 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 four dependent variables. In particular, we always find that the coefficient of Treatment (−1), which measures the “effect” of the new airline routes before their introduction,

¹⁷Increased innovation and higher exit likelihood are desirable outcomes. As such, they should translate into higher payoffs to the VC firm. Since these payoffs are not observable, assessing their magnitude requires a set of simplifying assumptions. Mindful of this caveat, in Section IV of the Internet Appendix, we conduct a simple back of the envelope calculation. This calculation suggests that there is no “money left on the table,” i.e., the increase in expected payoffs to the VC is unlikely to be large enough to justify taking costly measures to maintain higher levels of involvement prior to the treatment (e.g., hiring additional partners).

is small and insignificant, suggesting that there are no pre-existing trends in the data.¹⁸ 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 higher likelihood of a successful exit.

C. Lead versus Non-Lead VCs

The results thus far 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 or being acquired. Our interpretation is that reduced travel time increases VC involvement, which in turn improves portfolio company outcomes. Still, because we do not observe VC involvement, we cannot definitively show that VC monitoring increases following a reduction in travel time. 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 10 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 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

¹⁸We cannot identify the coefficient of Treatment (−1) in the regressions where the dependent variable is the IPO indicator or the success indicator. Since they would exit the sample, companies that go public or are acquired before the treatment cannot be in the treatment group by construction.

time between a portfolio company and a non-lead investor. The results are shown in Panel A of Table VI. 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. One potential concern with this analysis, however, is that reducing travel time to the MSA of a non-lead VC may have less of an impact, not because those VCs are less involved, but because they are located in different types of cities. For example, they may be located in cities with fewer resources for start-ups. A priori, there is no reason to expect this to be the case, particularly because a typical VC acts in both a lead and non-lead capacity on different deals. Nonetheless, we explore the possibility that non-lead treatments connect companies to different types of MSAs in Panel B. As can be seen, non-lead treatments and lead treatments connect companies to MSAs that are similar in terms of population, income, as well as geography. Thus, the results overall are consistent with the argument that VC involvement increases following the treatment—travel time reductions appear to matter primarily for active investors.¹⁹

D. Small versus Large Reductions in Travel Time

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 VII. For travel time reductions of

¹⁹The results presented in Table VI 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. In Internet Appendix Table IA.II, we conduct a more formal placebo test. Specifically, we replace each company’s real VC with a random VC that made investments in the same year the company was initially funded. We require that the placebo VC be located in a different MSA than any of the company’s real VCs. In addition, to strengthen the test, we require that the placebo VC be located in the San Francisco, San Jose, Boston, or New York MSAs, as these are generally regarded as the major innovation hubs in the U.S. We then reconstruct our *treatment* indicator as before based on these placebo VC relationships. As is shown in Panel A, these placebo treatments are not associated with improvements in company outcomes. This is despite the fact that, as Panel B shows, these placebo treatments connect companies to richer and more populous MSAs than real treatments.

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. We note, however, that the difference between the two coefficients is not statistically significant at conventional levels. Given the limited number of treatments we observe in the data, we may not have sufficient power to identify cross-sectional differences, even if they are present.

E. Robustness

E.1. Hub Openings and Airline Mergers

As explained in Section II.B, 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 VIII, 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.²⁰

E.2. Eventually Treated Pairs

As discussed in Section II.C, 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 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 (for a similar robustness check, see Bertrand and Mullainathan, 2003). In this case, the control group consists exclusively of pairs that are subsequently treated, thus alleviating 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 B of Table VIII. They are similar to our baseline estimates.

E.3. Miscellaneous Robustness Checks

This section presents additional robustness checks. For brevity’s sake, the results are tabulated in the Internet Appendix.

Distance-matched control group. To further mitigate the concern that control and treated pairs may be on different trends, we re-estimate our baseline specification after restricting the control

²⁰The treatment effect is larger for hub and merger treatments compared to other treatments, although the difference is not statistically significant. The larger point estimates likely reflect 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.

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 A of Internet Appendix Table IA.III. As is shown, our results are robust to using this “distance-matched” control group.

Heterogenous 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 characteristics 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 II with year fixed effects and re-estimate our baseline specification with these additional controls. The results are reported in Panel B of Internet Appendix Table IA.III. The estimated treatment effects are very similar to before.

Alternative dependent variables. In Panel C of Internet Appendix Table IA.III, we explore whether our results are robust to alternative definitions of our main dependent variables. As discussed in Section I.B.2, in our baseline analysis 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 before. 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.

Industry-specific local shocks. Next, we refine our baseline specification by allowing local shocks to be industry specific, that is, 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 major industry groups of VentureXpert. The results are presented in Panel D of Internet Appendix Table IA.III. As is shown, the estimates are very similar to our baseline coefficients in Table IV. However, the significance of the treatment effect is lower for all dependent variables (the treatment effect is even marginally insignificant for the IPO and Success indicators). This is not surprising given that the additional layer of industry fixed effects reduces the power of our tests.

Two-way clustering. In Panel E of Internet Appendix Table IA.III, we re-estimate our baseline specification, clustering standard errors at both the portfolio company level and the VC firm level. As is shown, this changes our standard errors little and all results continue to be statistically significant.

Access to non-VC resources. Finally, one potential concern is that a portfolio company might improve after the treatment not because of increased VC involvement, but because the portfolio company gains access to other resources at the VC’s location (e.g., universities, technology centers, trade shows). This concern is mitigated by the inclusion of MSA by year fixed effects, since all companies at the same location would benefit from more direct access to these resources (regardless of the location of their VC). Nevertheless, to further rule out this alternative explanation, we examine whether the treatment leads to an increase in citations made to patents of non-VC individuals/organizations at the VC’s location. More precisely, for each patent a portfolio company is granted, we calculate the percentage of citations that the patent makes to firms or inventors located in the MSA of the portfolio company’s VC.²¹ We then examine whether this percentage increases following a reduction in travel time. The results are reported in Internet Appendix Table IA.IV. As can be seen, the estimated coefficients are statistically indistinguishable from zero and the magnitude of the point estimates is small as well. Thus, there is no evidence that the innovative activity of the treated companies is influenced by increased access to non-VC resources at the VC’s location.

²¹Inventor location data are obtained from the USPTO’s raw XML files. Location names are standardized using the same procedure as the one described in Section II of the Internet Appendix.

F. Extensions

F.1. Cross-Sectional Heterogeneity

In Internet Appendix Table IA.V, we explore whether the treatment effect differs based on cross-sectional characteristics (e.g., early versus late stage companies). To conduct this analysis, we interact the *treatment* indicator with the characteristics of interest. This analysis is subject to two caveats. First, while the treatment is arguably exogenous, the variables interacted with the treatment may not be—that is, they may correlate with unobservable characteristics that affect the extent to which VC-company pairs react to the treatment. Second, given the limited number of treatments we observe in the data, we may not have sufficient power to identify heterogeneous effects, even if they are present. Despite these caveats, we do find suggestive results that go in the direction one might expect.

In Panel A, we interact the treatment variable with an “Early Stage” indicator equal to one if the company is classified as “Seed” or “Early Stage” in the year under observation. We find evidence that the treatment effect is larger for early stage companies, suggesting that VC involvement may matter more early on. For patents (citations), the coefficient on the interaction term is positive and statistically significant. In terms of magnitudes, the effect of a reduction in travel time on patents is estimated to be about 4.5 (3.4) times larger for early stage companies. For IPOs (success), the point estimates suggest that the effect is about 2.6 (1.5) times larger; however, in this case, the difference is not statistically significant at conventional levels, potentially due to lack of power.

Panel B interacts the treatment variable with an “Other VC Close” indicator equal to one if a non-lead VC (that is part of the investment syndicate) is located in the same MSA as the portfolio company. The point estimates of the interaction term coefficient are negative, suggesting that the treatment effect is smaller when a non-lead VC is located nearby. However, the difference in the effect is not statistically significant. This may reflect the fact that non-lead VCs are less actively involved in monitoring as discussed in Section IV.C. Finally, Panel C interacts the treatment variable with a “Syndicated” indicator equal to one if more than one VC invested in the company.

The point estimates are again negative but not statistically significant.

F.2. Regional Analysis

Lastly, a natural extension of our analysis is to study whether proximity fosters VC flows between regions—to the extent that travel time affects performance outcomes within existing VC-company relationships, it likely also affects VCs’ investment decisions. This analysis is provided in Section V of the Internet Appendix. In a nutshell, 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 do indeed foster VC flows between regions.

V. 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 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 also document that the effect is concentrated in routes that connect lead VCs (as opposed to other investors) with portfolio companies. Overall, our results

indicate that VCs' on-site involvement with their portfolio companies is an important determinant of innovation and success.

We confirm the importance of this channel by conducting a large-scale survey of VC investors. We find that almost 90% of the respondents agreed that they would visit a portfolio company more frequently if an indirect flight were replaced by a direct flight. Moreover, survey participants also agreed that the introduction of a direct flight would help them establish better relationships with management teams, better understand the state of their companies, and generally add more value.

REFERENCES

- Aghion, Philippe, John Van Reenen, and Luigi Zingales, 2013, Innovation and institutional ownership, *American Economic Review*, 103, 277–304.
- Anderson, James C., 1978, The validity of haire’s shopping list projective technique, *Journal of Marketing Research*, 15, 644–649.
- Berglof, Erik, 1994, A control theory of venture capital finance, *Journal of Law, Economics, and Organization*, 10, 247–267.
- Bertrand, Marianne, and Sendhil Mullainathan, 2003, Enjoying the quiet life? Corporate governance and managerial preferences, *Journal of Political Economy*, 111, 1043–1075.
- Boeh, Kevin K., and Paul W. Beamish, 2011, Connecting flights: The time sink that kills profits, *Harvard Business Review*, 89, 30.
- Boeh, Kevin K., and Paul W. Beamish, 2012, Travel time and the liability of distance in foreign direct investment: Location choice and entry mode, *Journal of International Business Studies*, 43, 525–535.
- Calder, Bobby J., and Robert E. Burnkrant, 1977, Interpersonal influence on consumer behavior: An attribution theory approach, *Journal of Consumer Research*, 4, 29–38.
- Casamatta, Catherine, 2003, Financing and advising: Optimal financial contracts with venture capitalists, *Journal of Finance*, 58, 2059–2085.
- Cestone, Giacinta, 2014, Venture capital meets contract theory: Risky claims or formal control?, *Review of Finance*, 18, 1097–1137.
- Chemmanur, Thomas J., Karthik Krishnan, and Debarshi K. Nandy, 2011a, 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.
- Cornelli, Francesca, and Oved Yosha, 2003, Stage financing and the role of convertible securities, *Review of Economic Studies*, 70, 1–32.
- Da Rin, Marco, Thomas Hellmann, and Manju Puri, 2013, A survey of venture capital research, in Milton Harris and Rene M. Stulz George M. Constantinides, ed., *Handbook of the Economics of Finance*, volume Volume 2, Part A, 573–648 (Elsevier).
- Dessi, Roberta, 2005, Start-up finance, monitoring, and collusion, *RAND Journal of Economics*, 36, 255–274.
- Dillman, Don A., Eleanor Singer, Jon R. Clark, and James B. Treat, 1996, Effects of benefits appeals, mandatory appeals, and variations in statements of confidentiality on completion rates for census questionnaires, *Public Opinion Quarterly*, 60, 376–389.
- Giroud, Xavier, 2013, Proximity and investment: Evidence from plant-level data, *Quarterly Journal of Economics*, 128, 861–915.

- Gompers, Paul A., 1996, Grandstanding in the venture capital industry, *Journal of Financial Economics*, 42, 133–156.
- Gompers, Paul A., Steven N. Kaplan, and Vladimir Mukharlyamov, 2014, What do private equity firms (say they) do?, *Working Paper*.
- 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, 2014, Common errors: How to (and not to) control for unobserved heterogeneity, *Review of Financial Studies*, 27, 617–661.
- Guimarães, Paulo, and Pedro Portugal, 2010, A simple feasible procedure to fit models with high-dimensional fixed effects, *Stata Journal*, 10, 628–649.
- Haire, Mason, 1950, Projective techniques in marketing research, *Journal of Marketing*, 14, 649–656.
- Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg, 2001, The NBER patent citations data file: Lessons, insights and methodological tools, *Working Paper*.
- Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg, 2005, Market value and patent citations, *RAND Journal of Economics*, 36, 16–38.
- Hellmann, Thomas, 2006, IPOs, acquisitions, and the use of convertible securities in venture capital, *Journal of Financial Economics*, 81, 649–679.
- 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.
- Inderst, Roman, and Holger Mueller, 2004, The effect of capital market characteristics on the value of start-up firms, *Journal of Financial Economics*, 72, 319–356.
- 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, 2003, Financial contracting theory meets the real world: An empirical analysis of venture capital contracts, *Review of Economic Studies*, 70, 281–315.
- 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*.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2012, Technological innovation, resource allocation, and growth, *Working Paper*.
- 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.
- Maccoby, Eleanor E., and Nathan Maccoby, 1954, The interview: A tool of social science, in Gardiner Lindzey, ed., *Handbook of Social Psychology*, volume 1, 449–487 (Addison-Wesley, Cambridge, MA).
- Metrick, Andrew, and Ayako Yasuda, 2011, *Venture capital and the finance of innovation*, second edition (Wiley, New York).
- 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.
- Repullo, Rafael, and Javier Suarez, 2004, Venture capital finance: A security design approach, *Review of Finance*, 8, 75–108.
- Schmidt, Klaus M., 2003, Convertible securities and venture capital finance, *Journal of Finance*, 58, 1139–1165.
- Seru, Amit, 2014, Firm boundaries matter: evidence from conglomerates and R&D activity, *Journal of Financial Economics*, 111, 381–405.
- Singer, Eleanor, Dawn R. Von Thurn, and Esther R. Miller, 1995, Confidentiality assurances and response a quantitative review of the experimental literature, *Public Opinion Quarterly*, 59, 66–77.
- 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.

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.

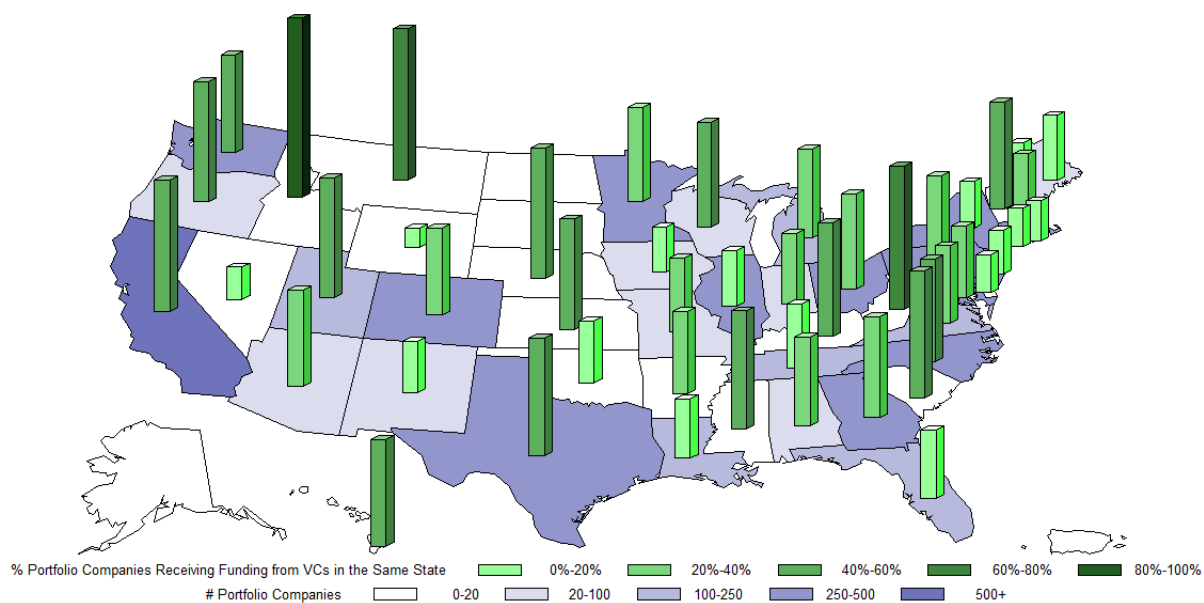


Figure 2
CDF of Distance Distribution

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




Figure 3
Survey Responses

This figure shows the distributions of responses to the survey questions provided in Section VI of the Internet Appendix. The first four rows show the responses to the general questions, the last six rows show the responses to the specific questions. On the horizontal axis, positive (negative) percentages refer to “agree” (“disagree”) responses.




Table I**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 I
(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 II
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 I. 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. 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 III**Survey Evidence**

This table shows the results of a survey of VC investors. Panel A shows the distribution of respondents across countries and U.S. states. Panel B summarizes the responses to the preliminary questions. Panels C and D summarize the responses to the key questions shown in Section VI of the Internet Appendix. The questions in Panel C regard general VC behavior, whereas the questions in Panel D regard the behavior of the respondents. On all questions, a standard 6-point likert scale is used, where potential responses range from *Strongly Disagree* (= 1) to *Strongly Agree* (= 6). The % Agree column represents the percent of respondents that somewhat agreed, agreed, or strongly agreed with the statement. A *t*-test is done to determine if the mean response is statistically different from the neutral mid-point of 3.5. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: Geographical Distribution

Nation	Percent	State	Percent
Brazil	2.26	California	69.52
Canada	0.75	Connecticut	0.95
China	2.75	Illinois	0.95
Germany	3.76	Louisiana	0.95
Hong Kong	1.5	Maryland	0.95
Israel	2.26	Massachusetts	13.33
Japan	0.75	New Hampshire	0.95
Poland	0.75	New Mexico	0.95
Portugal	0.75	New York	4.76
Russia	0.75	Pennsylvania	0.95
Singapore	0.75	Texas	2.86
South Africa	1.5	Utah	0.95
South Korea	1.5	Virginia	0.95
Sweden	0.75	Washington	0.95
Switzerland	0.75		
United States	78.95		

Panel B: Preliminary Questions

	N	Mean	Median	St. Dev.
Partner (0/1)	306	0.77	1	0.42
Currently a VC (0/1)	306	0.61	1	0.49
Years Since Last Worked in VC	113	8.65	6	7.82
Number of Companies	306	6.58	5	3.98
Number of Local Companies	304	3.79	3	3.52
Number of Visits Per Year	303	9.93	6	12.6
Visit Local Companies More Than Non-Local (0/1)	303	0.71	1	0.46
Percent of Time Spent Monitoring	306	0.48	0.50	0.18

Table III
(Continued)

Panel C: General Questions				
Direct flights allow VCs to...	N	% Agree	Mean	St. Dev.
Spend more time assisting/monitoring in person	306	0.86	4.43***	1.18
More effectively advise companies	306	0.83	4.31***	1.20
Add more value to companies	306	0.80	4.20***	1.27
Better understand key challenges/issues facing companies	306	0.80	4.32***	1.31
Panel D: Specific Questions				
The introduction of a direct flight will...	N	% Agree	Mean	St. Dev.
Increase frequency of travel	306	0.83	4.28***	1.14
Increase flexibility to travel when useful	306	0.89	4.74***	1.13
Help communicate more effectively	306	0.72	3.90***	1.22
Help establish better relationships	306	0.81	4.31***	1.20
Help add more value	306	0.75	3.93***	1.20
Help understand state of company	306	0.76	4.16***	1.22

Table IV

Main Regressions

This table shows the main results. Observations are at the company-VC pair by year level. Only pairs involving a lead investor are included in the sample. Treatment is an indicator variable equal to one if a new airline route that reduces the travel time between the VC and the portfolio company has been introduced. Patents is equal to the log of (one plus) the number of patents the portfolio company applied for during the year. Citations/patent is equal to the log of (one plus) the number of citations those patents received (in the three years following their grant date) divided by the number of patents. IPO is an indicator variable equal to one if the company went public that year. Success is an indicator variable equal to one if the company went public or was acquired (for over 25 million in 2000 dollars) that year. Standard errors, clustered by portfolio company, are shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: Innovation

	Patents			Citations/Patent		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0371*** (0.00975)	0.0352*** (0.00971)	0.0310*** (0.0113)	0.0744*** (0.0178)	0.0698*** (0.0178)	0.0575*** (0.0203)
Controls	No	Yes	Yes	No	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	No	Yes	Yes	No
MSA(VC) \times Year FE	No	No	Yes	No	No	Yes
MSA(Company) \times Year FE	No	No	Yes	No	No	Yes
R ²	0.638	0.640	0.668	0.546	0.547	0.576
Observations	130169	130169	130169	130169	130169	130169

Panel B: Exits

	IPO			Success		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0103*** (0.00378)	0.00994*** (0.00373)	0.0104** (0.00429)	0.0113** (0.00507)	0.0112** (0.00493)	0.0135** (0.00577)
Controls	No	Yes	Yes	No	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	No	Yes	Yes	No
MSA(VC) \times Year FE	No	No	Yes	No	No	Yes
MSA(Company) \times Year FE	No	No	Yes	No	No	Yes
R ²	0.435	0.440	0.494	0.399	0.405	0.453
Observations	130169	130169	130169	130169	130169	130169

Table V
Dynamics

This table shows the dynamics of the treatment effects. All variables are defined as in Table IV. 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	(4) Success
Treatment(-1)	0.00639 (0.0147)	0.0170 (0.0285)		
Treatment(0)	0.0165 (0.0155)	0.0244 (0.0283)	0.00682 (0.00502)	0.0114 (0.00710)
Treatment(1)	0.0391** (0.0182)	0.0690** (0.0333)	0.00805 (0.00644)	0.0110 (0.00842)
Treatment(2+)	0.0494*** (0.0182)	0.106*** (0.0326)	0.0158** (0.00655)	0.0172** (0.00831)
Controls	Yes	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes
MSA(VC) \times Year FE	Yes	Yes	Yes	Yes
MSA(Company) \times Year FE	Yes	Yes	Yes	Yes
R ²	0.668	0.576	0.494	0.453
Observations	130169	130169	130169	130169

Table VI
Non-Lead VCs

Panel A repeats the analysis of Table IV, but restricting the sample to company-VC pairs that do not involve a lead investor. Panel B compares mean VC MSA characteristics (in the treatment year) for treatments involving lead and non-lead investors. Non-lead VCs located in the same MSA as the lead VC are excluded from the sample in both panels. Standard errors, clustered by portfolio company, are shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: Effect of Non-Lead Treatment				
	(1) Patents	(2) Citations/Patent	(3) IPO	(4) Success
Treatment	-0.0128 (0.0203)	-0.0205 (0.0368)	0.00761 (0.00691)	0.0139 (0.00972)
Controls	Yes	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes
MSA(VC) \times Year FE	Yes	Yes	Yes	Yes
MSA(Company) \times Year FE	Yes	Yes	Yes	Yes
R ²	0.758	0.688	0.673	0.627
Observations	90609	90609	90609	90609

Panel B: Lead vs Non-Lead Treatment Characteristics						
	Lead Treat		Non-Lead Treat		Difference	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Err
VC MSA Income (Billions)	161.1	201.4	150.6	190.1	10.5	8.53
VC MSA Population (Millions)	4.85	5.56	4.49	5.18	0.36	0.23
VC MSA Income Per Capita (Thousands)	33.1	12.1	33.5	13.1	-0.41	0.55
VC in Northern California	0.087	0.28	0.10	0.30	-0.014	0.012
VC in New York Tri-State	0.22	0.42	0.20	0.40	0.018	0.017
VC in New England	0.19	0.39	0.19	0.40	-0.0037	0.017
Observations	1131		1068		2199	

Table VII
Intensity of the Treatment

This table repeats the analysis of Table IV, 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	(4) Success
Treatment \times Large	0.0336** (0.0143)	0.0684*** (0.0248)	0.0115** (0.00524)	0.0138** (0.00701)
Treatment \times Small	0.0259 (0.0173)	0.0359 (0.0333)	0.00822 (0.00683)	0.0129 (0.00948)
Controls	Yes	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes
MSA(VC) \times Year FE	Yes	Yes	Yes	Yes
MSA(Company) \times Year FE	Yes	Yes	Yes	Yes
R ²	0.668	0.576	0.494	0.453
Observations	130169	130169	130169	130169

Table VIII**Robustness**

All regressions presented in this table are variants of the baseline specification in Table IV. Panel A separates 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 sample to the eventually treated pairs. 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	(4) Success
Treatment (Hub or Merger)	0.0540** (0.0255)	0.116** (0.0508)	0.0237* (0.0142)	0.0325* (0.0176)
Treatment (Other)	0.0273** (0.0126)	0.0475** (0.0219)	0.00842* (0.00433)	0.0105* (0.00593)
Controls	Yes	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes
MSA(VC) \times Year FE	Yes	Yes	Yes	Yes
MSA(Company) \times Year FE	Yes	Yes	Yes	Yes
R ²	0.668	0.576	0.494	0.453
Observations	130169	130169	130169	130169

Panel B: Eventually Treated Pairs

	(1) Patents	(2) Citations/Patent	(3) IPO	(4) Success
Treatment	0.0314*** (0.0107)	0.0354* (0.0207)	0.0250*** (0.00414)	0.0376*** (0.00517)
Controls	Yes	Yes	Yes	Yes
Pair FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
R ²	0.582	0.440	0.218	0.211
Observations	7978	7978	7978	7978