The Causal Effects of Proximity on Investment: Evidence from Flight Introductions

Jesse Ellis
North Carolina State University
jaellis5@ncsu.edu

Leonardo Madureira
Case Western Reserve University
leonardo.madureira@case.edu

Shane Underwood
Baylor University
shane_underwood@baylor.edu

November 17, 2016

Abstract

We address fundamental endogeneity issues inherent in the local bias literature by using direct flight introductions as an exogenous shock to the travel time between mutual funds and firms, thereby allowing us to understand the causal effects of proximity on fund investment decisions and performance. We find that a fund invests significantly more in firms that become more proximate due to the introduction of direct flights that reduces the fund's cost of traveling to the firm, and that these more proximate investments exhibit superior performance. Our findings are robust to the inclusion of a variety of fixed effects and specifications that control for potential confounders such as firm-level shocks, fund-level shocks, and time trends. Collectively, our results indicate that proximity enhances investors' ability to acquire value-relevant information about firms.

Keywords: Local Bias, Mutual Funds, Portfolio Choice

JEL classification: G11, G23

1. Introduction

There is a large body of evidence indicating that investors exhibit a preference for investing in firms that are located nearby.[1] This local bias could arise from many factors. For one, it could be a symptom of a behavioral bias whereby investors may be more comfortable making risky decisions in contexts where they consider themselves knowledgeable or more familiar (Heath and Tversky, 1991). Investors may be more familiar with local companies, and favor them simply because of their increased salience. However, there is also a rational story for local bias: increased proximity could enhance an investor's ability to acquire information about the firm, thereby leading to better investment decisions (Coval and Moskowitz, 2001). Understanding the causes and consequences of local bias is important, because it can help us understand how investors become informed as well as the extent to which they succumb to behavioral biases.

Naturally, much of the local bias literature studies investors and firms that are geographically proximate to one another. However, there are important methodological issues that limit our ability to make causal inferences from studies of geographic proximity. The primary endogeneity issue is that firms and investors choose where to locate, and thus there could be unobservable factors driving certain investors and firms to locate in the same place. For instance, it could be that mutual fund managers that have a unique talent or preference for investing in technology stocks choose to locate in California to be closer to their preferred investment habitat. If that were the case, the fund's local investment choices and performance may be the result of their initial location decision, rather than their proximity to local firms.

[1] For example, there is evidence that institutional investors such as mutual funds invest disproportionately in the stocks of firms that are headquartered or operate nearby (e.g., Coval and Moskowitz, 1999 and 2001; Bernile, Kumar, and Sulaeman, 2015a and 2015b; and Bernile, Kogan, and Sulaeman, 2015). Also, individual investors appear to exhibit a similar preference for local stocks (e.g., Ivkovic and Weisbenner, 2005; Huberman, 2001; Massa and Simonov, 2006; Bodnaruk, 2009). There is even evidence that fund of fund investors prefer local hedge funds (Sialm, Sun, and Zheng, 2013).

Even if we set aside the problem of endogenous location decisions, we are limited in what we can learn from the extant literature as to the key mechanisms underlying local bias. For example, even if we assume that proximity (and not location decisions) causes funds to favor local companies, a further issue is identifying whether the effects of proximity stem from enhanced information acquisition or merely a preference for the familiar. Going a step further, even if we assume that local bias is at least partly driven by information advantages, we still know little about how investors become informed about local stocks. One explanation is that funds are skilled in their local investments because they have a lower cost of meeting with firm managers in order to obtain information. Another explanation is that firm information gets disseminated throughout the local community, and funds pick up on locally generated rumors without actually engaging with the firm. Or perhaps funds are simply herding locally— following the trading behavior of other local investors without actually possessing superior information themselves (Pool, Stoffman, and Yonker, 2015). In other words, by being located in the same community, there are several potential confounding effects that obscure our ability to understand why there is a local bias and what it implies for capital market participants.

One way to overcome the endogeneity issues inherent in much of the local bias literature is to examine proximity through the lens of travel time, rather than geographic distance. Although distance is fixed, travel time varies as a function of the availability of airline flights between fund and firm headquarter locations. Thus, funds and firms can effectively become more proximate to one another when new airline routes are introduced that significantly reduce the travel time between them. Importantly, flight introductions are plausibly exogenous to the investment patterns between funds and firms, meaning we can use them as a means of identifying the causal effects of proximity on fund investment. Giroud (2013) was the first to make use of exogenous changes in travel time in a study of how firms allocate capital within their internal capital markets. In this paper, we follow Giroud's methodology and identify the causal effect of proximity by exploiting exogenous variation in travel time caused by the initiation of new airline routes between funds and firms.

Specifically, we use a difference-in-difference style approach to examine the changes in holdings and performance of a fund's investments in Metropolitan Statistical Areas (MSAs) that have become "closer" to the fund following the introduction of a direct flight between the MSAs of the fund and of the firm. We consider the introduction of a new direct flight between a fund-MSA pair to be an exogenous proximity treatment. For example, consider the Janus Twenty fund, which is located in the Denver, CO MSA, and the Fayetteville-Rogers-Springdale (Arkansas) MSA. This MSA is home to several publicly traded corporations, including Wal-Mart. Prior to 2004, the fastest route between Denver and this MSA took a minimum of 3 hours 50 minutes, as it required a connecting flight through another city such as Dallas. However, in 2004, a non-stop flight was introduced between Denver and the Northwest Arkansas regional airport that dropped the travel time to 1 hour 50 minutes.[2] Thus, we consider Janus Twenty and its investments in the northwest Arkansas area to be treated in 2004, and we examine how that treatment affected the fund's holdings and investments in this area relative to its investments in other MSAs over the same period. The advantage of this methodology is that it allows us to isolate the channel through which proximity affects investment. Because treated fund-MSA pairs are not geographically proximate, the changes we observe following flight introductions cannot be attributed to factors like location-based familiarity or community affinity effects. Rather, they can only be attributed

[2] It is important to note that the reduction in travel time of about 2 hours in this example is a lower bound on the true reduction in perceived costs, since adding a layover also increases the risk of delays, cancellations, etc.

to changes in the cost of travel between the two places, suggesting the effects we uncover are based on reductions in information acquisition costs.

We find evidence that when a new direct flight is introduced between a fund and an MSA, the fund invests significantly more in firms headquartered in that MSA. However, as Giroud (2013) points out, the major concern with this type of analysis is that the change in travel time could be related to unobservable local shocks, thus leading to a spurious effect between holdings and travel time. For example, if the MSA experiences a wave of new firm growth, it may also experience both an increase in flights and an increase in fund investment activity. Thus, we follow Bernstein, Giroud, and Townsend (2016) and control for local shocks at both the firm and the fund's headquarters by including both fund-year and MSA-year fixed effects. MSA-year fixed effects control for the possibility that the MSA has experienced a shock that has led to increased investment in firms from that region. Fund-year fixed effects further reduce the variation to only the changes within a fund's portfolio in a point in time, to control for the possibility the fund has grown or changed its overall investment strategy. Finally, we include fund-MSA pair fixed effects to account for any unobservable fixed characteristics between pairs that are treated vs. those that are not treated. Including all these fixed effects, we find that the introduction of a new airline route leads to a 27% increase in fund holdings in the treated region.

Next, we examine the effect of reduced travel time on a fund's stock selection skill. If reduced travel costs improves a fund's ability to acquire information, then we should observe an improvement in the performance of the fund's holdings in treated regions. Using a similar difference in difference methodology (including a full complement of fixed effects), we find evidence that is consistent with this prediction. Specifically, we find that the risk-adjusted performance of stocks held in treated regions is 1.74% per year higher than the performance of stocks in untreated regions. This supports the notion that increased proximity reduces a fund's cost of acquiring value-relevant information.

Finally, we subject our primary results to a battery of robustness tests to ensure our findings reflect a causal relationship between travel time, investment and performance. First, we find similar results when we focus on the sub-sample of eventually treated fund-MSA pairs only, which addresses the possibility that the treated pairs are somehow fundamentally different from the non-treated pairs. Second, we control for time trends by interacting the treatment indicator with a set of indicators for years around the treatment, and find no effect in the years leading up to the treatment. This is an important result because it indicates that the introduction of the flight is not being caused by a pre-existing trend of increased fund investment. Finally, to ensure our results are not spurious, we run two placebo tests. In our first placebo test, we randomize the treatment date between fund-MSA pairs and our results go away. In our second placebo test, we replace our set of actively managed funds with a sample of index funds. Index funds are passive by definition, and thus should not be changing their investment allocations or holdings performance based on changes in information costs. When we re-run our results using only index funds, the results go away as well. This helps to confirm that our results are not symptomatic of some unobservable trends between fund-firm location pairs.

This paper contributes to the extensive literature on local bias and the means by which investors acquire value relevant information. There is a plethora of evidence suggesting that investors invest disproportionately in firms that operate or are headquartered nearby, however the evidence is more mixed as to whether investors benefit from their local preference. Some studies find evidence that both professional and individual investors earn superior returns on their local investments, suggesting that a local bias may be rational and driven by an investor's enhanced ability to obtain value-relevant information from local information sources (Coval and Moskowitz, 2001); Ivkovic and Weisbenner, 2005; Bernile, Kumar, and Sulaeman, 2015a). However, other studies (e.g., Seasholes and Zhu, 2010; and Sulaeman, 2014) do not find evidence of superior performance in local investments, thereby calling into question the rationality of overweighting local stocks. Our paper contributes to this debate by addressing the endogeneity issues inherent to studies that rely on geographic proximity to understand local bias. By using exogenous shocks to travel time, we are able to cleanly isolate the effects of decreased information acquisition costs on investor outcomes apart from potentially confounding familiarity or location selection effects. Moreover, because our evidence is based on flights between geographically distant locales, it is more likely that we are detecting evidence of investors acquiring information through direct contact with firms, as opposed to secondhand exposure to information diffused through a community-wide social network.

In addition, we are able to contribute to the growing literature examining the impact of decreased travel time on investment outcomes. Giroud (2013) was the first to come up with the strategy of using exogenous changes in travel time in a study of how firms allocate resources within their internal capital markets. However, Giroud (2013) notes that the nature of his study made it difficult to disentangle whether the effects came from enhanced information acquisition, enhanced monitoring, or both. By focusing on proximity shocks to pre-existing venture capital relationships, Bernstein, Giroud, and Townsend (2016) argue that their evidence suggests proximity can enhance investor monitoring, though they do not rule out a role for information acquisition. Our study uses a setting where the proximity effects are unlikely to come from monitoring, given that mutual funds are not typically active monitors, and typically exhibit short holding periods and small investment proportions in our sample of proximity treatments. Thus, our evidence that proximity affects investment decisions and performance suggests that proximity can also enhance investors' ability to acquire value-relevant information.

We proceed as follows. Section 2 describes the data, defines the basic variables of the study, and provides some preliminary analysis of the relationship between proximity and mutual fund holdings. Section 3 presents the main results on how holdings and performance of funds are affected by the introduction of direct flights. Section 4 presents more analyses to verify the robustness of the results. Section 5 concludes.

2. Data

2.1 Data Sources

In order to track the introduction of direct flights through time, we use the T-100 data from the Bureau of Transportation Statistics. The T-100 dataset begins in 1990 and tracks all domestic flights on a monthly basis, and includes origin, destination, total ramp-to-ramp travel times, and total number of flights, as well as a host of other variables. We consider only flights that are listed as Class "F" or scheduled passenger service. This eliminates cargo-only flights as well as a number of other miscellaneous categories of air service. From this data we then construct several important variables for our analysis.

We first define for each pair of MSAs and each month a dummy variable that is equal to one if there is a direct flight between the MSAs in that month. We also record the number of direct flights that took place between the two MSAs in that month. If an MSA has more than one commercial airport, we consider all airports located in the MSA in constructing these variables. For the purpose of recognizing the gain in travel time when a direct flight is introduced, we also construct measures of the number of legs required to fly between two MSAs, as well as the total time needed to fly between the MSAs. The construction of these variables requires us to find the shortest flight path between two MSAs. To find the optimal route, we use the NETFLOW procedure in SAS and assume that each layover adds 60 minutes to the total travel time. This is consistent with Giroud (2013) and is likely to be a conservative estimate of the cost of a layover, given the increased likelihood of flight delays and cancellations as legs are added to a route. In determining the flight time between a pair of airports with a direct flight, we take the average across all flights in all months during the sample. This reduces noise in the data due, for example, to excessive delays in any particular month. In constructing the travel time between each MSA pair, we make a number of assumptions. First, we assume that driving is optimal for any two MSAs within 100 miles of each other. If the MSAs are more than 1000 miles apart, we assume flying is optimal and consider the total travel time to be the flight time between the two MSAs. For MSA pairs between 100 and 1000 miles apart, we compare the shortest flight time to the driving time between the two cities and take the minimum of these two. Drive times are calculated using Microsoft MapPoint.

For firm location, we use the headquarters zip code from Compustat. For mutual funds, we use the zip code from the CRSP mutual fund database. This variable is available beginning in 2000, and we backfill to earlier years for funds in the period prior to the year 2000. We also supplement with location data pulled from NSAR filings of the mutual funds.[3] In some cases, the NSAR filings show multiple zip codes for a single fund due to sub-advising relationships. In these cases we exclude the fund from our analysis.

We collect mutual fund holdings data from the Thomson S12 holdings database. We consider holdings by all U.S. equity mutual funds over the 1990-2015 period. We then supplement

[3] We thank Chris Clifford for providing this data.

this with active share data obtained from Martijn Cremer's website. For the bulk of our analysis, we follow Cremers and Petajisto (2009) and we restrict the sample to only include funds with active share above 50%, so as to exclude index funds and closet indexers who are less likely to be making active investment decisions.

We measure holdings at the fund-MSA pair level. That is, for each fund we calculate their total holdings in dollars, and as a percentage of their total portfolio, in each MSA. We also calculate a measure of fund performance in stocks at the MSA level. For each MSA in which the fund held stocks in a given year, we calculate the value-weighted risk-adjusted returns of the fund's holdings over the next 12 months. Stock returns are risk-adjusted according to the procedure in Daniel, Grinblatt, Titman, and Wermers (1997), so henceforth we refer to them as DGTW-adjusted returns.

2.2 Defining Treatments

We define treatments at the level of pairs of MSAs. More specifically, a pair of MSAs is treated during our sample period at time *t* if at that time a direct flight is initiated between the two MSAs. We also require that the initiation occurs with at least 10 flights in the month of initiation and that there be no direct flights between the two MSAs at any time prior to this initiation.[4] Moreover, this analysis ignores direct flights introduced between MSAs located less than 100 miles apart (as we discussed above, we consider driving to be optimal for such situation).

[4] We use 10 flight per month cutoff to 1) ensure that there is a sufficient number of flights such that the fund could have reliable access to the flight (i.e., the flight introduction represents a real improvement in the travel time choice set) and 2) to ensure the flight introduction is exogenous and not related to pre-existing time trends. If a flight is introduced sporadically at first and then slowly builds to 10 flights per month, it is much more likely to represent endogenous swings in demand. On the other hand, a discrete flight introduction that begins by offering 10 or more flights per month is much more likely to represent an exogenous shock to travel time between regions. Later we discuss robustness tests where we consider different thresholds for defining a treatment. Our results are robust to a variety of cutoffs for the number of flights in the initiation month.

Figures 1 through 3 provide summary evidence on the distribution of treatments through time and across the country during our sample. There are a total of 558 pairs of MSAs that are treated during our sample period. This means there were 558 new direct flights introduced between MSAs in which there are funds and companies during our sample period.

<Insert Figure 1 here>

Figure 1 shows the geographical dispersion of MSAs where mutual funds are located (henceforth fund MSA) and the number of treatments per fund MSA. That is, we summarize the number of times that each fund MSA sees a direct flight introduced to an MSA that is the home of firms which are potential investments. Cities with the most treatments include Atlanta (38), Dallas (30), and Cincinnati (28), but there is a good deal of geographic variation. In all, there are 67 fund MSAs that receive at least one treatment over the 1990-2014 period.

<Insert Figure 2 here>

Figure 2 repeats the analysis with respect to MSAs containing companies that can be considered investments to mutual funds (henceforth company MSAs). There is a greater level of geographic dispersion of company MSAs, reflecting the broader spread of corporations across the U.S. relative to mutual funds. In all, there are 231 company MSAs with at least one treatment over the 1990-2014 period, and 29 with at least 5 treatments. The most frequently treated MSAs are Fayetteville-Springdale-Rogers (Arkansas) and Trenton, NJ, each with 13 treatments.

<Insert Figure 3 here>

Finally, Figure 3 shows the number of treatments per year. While treatments are concentrated in the early years of the sample, there are only five years with less than 10 treatment events. The figure shows that the number of treatments in each year ranged from a high of 67 (1990) to a low of 6 (2015). Given that we only allow a pair of MSAs to be treated once, we would expect a gradual decline in the number of treatments over time as fewer MSA pairs have the possibility of experiencing a new direct flight. We can also observe in Figure 3 some general macroeconomic effects, as there are fewer treatments in the early 2000's recession as well as in the financial crisis later in the decade.

2.3 Proximity and Mutual Fund Holdings

In Table 1, we provide initial evidence on the relation between travel time and mutual fund holdings. The observational units in our analysis are yearly fund-MSA pairs. For each fund-MSA pair we analyze the fund's yearly holding of stocks located in that MSA. We initially sort fund-MSA pairs into eleven groups based on the distance between the fund's MSA and the MSA of the stocks. Within each of these 250-mile distance groups, we sort the fund-MSA pairs into quartiles based on the travel time between the two locations. We then measure the average fund holdings in each bin. Panel A presents results for the number of stocks held in a fund's portfolio, and Panel B presents results for the weight (in percent) of these stocks in the fund's portfolio. The results consistently suggest that travel time is an important determinant of a fund's decision to hold a particular stock, even after controlling for the distance between the fund and the firm. For example, consider the row representing fund-MSA pairs that are 2,001-2,250 miles apart in Panel A. Even though fund-MSA pairs within this group are all very similar in terms of the physical distance between each other, we see that pairs with a shorter travel time have significantly higher holdings than pairs with longer travel times. Fund-MSA pairs in the shortest travel time quartile have average holdings of 3.5 stocks, while fund-MSA pairs in the longest travel time quartile have average holdings of just over 2 stocks. The difference between these two is significant at the 1% level. For all distance groups and for both measures of fund holdings, we see similar results. Taken together, these results suggest that travel time is an important metric that may affect a fund's decision to invest in a firm.

<Insert Table 1 here>

Table 2 provides summary statistics on yearly fund-MSA holdings depending on treatment status and the relative distance between fund's location and the MSA. We discuss results aggregating all yearly fund-MSA holdings in the sample. Corroborating the local bias extensively discussed in the literature, the highest average holdings occur for "local" fund-MSA pairs that is, fund's holdings in stocks located in the same MSA as the fund's MSA. On the other hand, the smallest average holdings happen for the "never treated" fund-MSA pairs— the ones whose locations are never served by direct flights during our sample period.

<Insert Table 2 here>

As a first view of the potential effects of introduction of direct flights, the results show that holdings are substantially higher in "treated" fund-MSA pairs compared to pairs that are "later treated" in the sample. Holdings are even higher for fund-MSA pairs whose locations are "always served" by direct flights— that is, the pair is not treated in our sample period because a direct flight was already available at the beginning of the period. Finally, funds that are less than 100 miles away from an MSA have average holdings between those of "later treated" and those of already "treated" fund-MSA pairs. Table 2 also shows summary stats for three different years in the sample to highlight that the inferences above are not depending on specific eras of our sample period.

In summary, the results provide evidence that treated fund-MSA pairs are associated with increased holdings. However, the summary statistics also highlight some concerns about the sample. For example, there is a pattern of increased holdings over time. Also, different groups of fund-MSA pairs are materially different— for example, "never treated" fund-MSA pairs have substantially lower holdings compared to fund-MSA pairs that are treated in our sample. We next turn to regression analyses to verify the robustness of the univariate patterns in view of these concerns.

3. Results

3.1 Holdings

We examine whether a fund's holdings of stocks in an MSA relate to the introduction of a direct flight between the fund's location (defined by the MSA of the fund's headquarters) and the stocks' MSA. More specifically, we estimate the following regression model to explain fund holdings in a specific MSA:

$$Y_{f,m,t}=\beta*Treatment_{f,m,t}+\alpha_f*\alpha_m+\alpha_f*\alpha_t+\alpha_m*\alpha_t+\varepsilon_{f,m,t} \qquad (1)$$

A data point in this regression is a tuple (fund *f*, MSA *m*, year *t*) where *f* identifies a fund, *m* represents an MSA where a fund might hold stocks, and *t* is the year of the observation. $Y_{f,m,t}$ is the measure of holdings of the fund in that MSA in that year.[5] Treatment is equal to 1 if at some point prior to that year a direct flight was initiated between the fund's location and the MSA. The α components refer to fixed effects terms: $\alpha_f*\alpha_m$ are fund-MSA pair fixed effects, $\alpha_f*\alpha_t$ are fund by year fixed effects, and $\alpha_m*\alpha_t$ are MSA by year fixed effects. For this and all other regression specifications in the paper, standard errors are clustered at the fund-MSA pair level.

<Insert Table 3 here>

Table 3 presents the results from running variations of the model explaining fund holdings in an MSA. Columns (1) through (5) have (the logarithm of) holdings as the dependent variable.

---

[5] The sample includes for each fund and each period all MSAs from which the fund could potentially hold stocks. These are the MSAs in that period that have at least one company whose stock is held by a fund in the sample. For an MSA from which a fund holds no stocks, holdings are set to zero. Zero holdings represent the fund's decision not to hold stocks in that MSA

We start in column (1) using only year fixed effects, then progressively extend the model with the addition of other fixed effects. The results in column (1) appear solely for illustrative purposes: the coefficient on the treatment dummy is equal to 2.12 (and significantly different from zero), but most of its magnitude is later subsumed when accounting for additional fixed effects.

In column (2) we control for fund-MSA pair fixed effects. Using fund-MSA pair fixed effects addresses two concerns. First, treatments may correlate with fund holdings given the characteristics of the treated pair of locations. For example treatments may be more likely to occur between economically robust locations— with more funds in one location and/or more companies in the other location. Second, there is the possibility that fund-MSA characteristics could drive the treatment as funds could lobby for a direct flight to locations more relevant in terms of its— current or prospective— portfolio of companies. When controlling for fund-MSA pairs, the coefficient on the treatment dummy drops from 2.12 to 0.55.

We next add fixed effects intended to control for local shocks that may correlate with the inclusion of direct flights. First, local shocks may appear at the location of companies that can be held by mutual funds— say an economic shock that both lifts the fortunes of the local companies and increases the attractiveness of the location for new airline flights. Thus, an increase in holdings of companies located in a treated MSA may result from this economic shock rather than the new direct flight. To account for this possibility, we include MSA by year fixed effects. Results in column (4) show that such inclusion, when compared to the specification adopting only fixed effects by year and fund-MSA pairs, leads to a further reduction in the coefficient on the treatment dummy, from 0.55 to 0.22.

Second, an economic shock at the location of a fund may increase the likelihood of that location being the starting point for new airline flights. If investors of the funds in that location are also geographically proximate, thus benefitting from the economic shock, we could see a correlation between treatments and increase in holdings from these funds. To control for this possibility, we employ fund by year fixed effects. However, results in column (3) show no further change in the treatment coefficient when fund by year fixed effects are added to the specification using fixed effects by year and fund-MSA pairs.

Finally, column (5) presents the results from the full-fledged model in equation (1), thus including all four types of fixed effects. The coefficient on the treatment dummy is still highly significant— with *t*=3.80. Given the logarithmic nature of the specification, the coefficient estimate of 0.24 implies a 27% increase in the fund's holdings of stocks in an MSA after the introduction of a direct flight between the fund location and that MSA.

In columns (6) through (10) we repeat the analysis with an alternative proxy for holdings. The dependent variable for an observation (fund, MSA, year) is an indicator equal to one if the fund has any holdings of stocks located in the MSA in that year. Inferences are similar to the ones obtained through the analysis of dollar amount of holdings: A treated pair fund-MSA is associated with an increased likelihood that the fund holds stocks in that MSA. The coefficients on the treatment dummy are positive and significant in all specifications, and become smaller as we add different types of fixed effects. In the model employing all four fixed effects, the coefficient on treatment equals 0.013 (*t*=3.14), suggesting that, other things equal, the introduction of a direct flight leads to an increase of 1.3% in the likelihood of the fund holding stocks in the treated MSA.

One limitation of the results in Table 3 is that it does not address the potential for a trend affecting a specific fund-MSA pair. For example, consider a shock that makes the companies in some specific MSA more attractive to a certain fund, so that the fund increases holdings in these companies after the shock (but other funds do not). This shock may also be the catalyst for the initiation of a direct flight between the fund location and that specific MSA— say, because the fund lobbies the airline for a direct flight to that MSA given its increased relevance. Fund-MSA pair fixed effects cannot subsume such a shock as, by definition, the shock implies a change in the characteristics of a fund-MSA pair. Given that the shock takes place for some specific fund-MSA pair, it does not affect other holdings of the same fund, so the adoption of fund by year fixed effects cannot control for such shock. Finally, MSA by year fixed effects are of no help either, as the shock may not affect holdings of the companies in that MSA by other funds.

If the introduction of the direct flight is the result of fund-MSA pair shocks, the increase in holdings that also results from these shocks may occur even before the flight introduction. In our next examination, we thus look at possible effects taking place even before our formal recognition of a treatment. For this test, we replace the single dummy for before- vs. after-treatment status with multiple dummies to have a more detailed view of the timing of the observation with respect to the initiation of a direct flight. More specifically, the new model employs three dummies: Treatment (-1) equals 1 to indicate the observation is from the year preceding the initiation of the direct flight; Treatment (0) equals 1 to indicate the observation is from the actual year of such initiation; and Treatment (1+) equals 1 to indicate the observations is from one year or more after the initiation.

<Insert Table 4 here>

Table 4 shows the results. As in Table 3, we show multiple specifications, where we vary the dependent variable and the usage of different types of fixed effects. Results for the full-fledged models in columns (5) and (10) show that the coefficients measuring the effect before the introduction (Treatment (-1)) or at the year of the introduction (Treatment (0)) of the flight are not significantly different from zero. On the other hand, Treatment (1+) is significant and similar in magnitude to the corresponding treatment coefficients in Table 3. In sum, this is evidence that the effect of the introduction of a new direct flight happens strictly after such introduction, and that the introduction of the flight is not being caused by some pre-existing trend of increased fund investment in the treated MSA.

3.2 Performance

We next examine the effect of introduction of direct flights on a fund's stock selection skill. If proximity derived from reduced travel time increases the fund's ability to acquire information, and therefore improve its investment decisions, we should observe an improvement in the performance of the fund's holdings in the treated MSA. In this section we test whether the performance of a fund's holdings in an MSA relates to the introduction of direct flight between the fund's location and the MSA.

Notice that that the analysis of performance in treated MSAs is not conditional on the results of the previous analysis of holdings. The hypothesis here is that reduction in travel time increases the fund's ability to *select* its portfolio of stocks— that may include increasing holdings of stocks with good prospects and decreasing holdings of stocks with bad prospects. Thus, in theory it would be possible for a fund not to increase its holdings in a treated MSA— say, if the average prospect of firms in that MSA is poor— and still attain better performance for the portfolio of stocks held from that MSA.

We measure performance of a fund at the level of each MSA and each point in time. More specifically, for a tuple (fund f, MSA m, year t) we define performance as the one-year ahead risk-adjusted return of the portfolio of stocks from MSA m that are held by the fund f in year t. Naturally, in some MSAs a fund might not hold any stocks; in these cases we define the fund performance as zero. This reflects the fact that not holding stocks at all from an MSA is a choice of the fund— with the fund consciously deciding not to pursue any return from that MSA. In summary, a fund has two options regarding an MSA: hold some stocks in the MSA, in which case its return is the return on those stocks, or not hold any stocks, which by definition leads to a zero return. The model explaining fund performance in a specific MSA thus becomes:

$$Perf_{f,m,t} = \beta * Held_{f,m,t} + \gamma * Treatment_{f,m,t} + \delta * Held * Treatment_{f,m,t} + \alpha_f * \alpha_m + \alpha_f * \alpha_t + \alpha_m * \alpha_t + \varepsilon_{f,m,t} \quad (2)$$

where *f* identifies a fund, *m* an MSA where a fund might hold stocks, and *t* a year. The dependent variable $Perf_{f,m,t}$ is the performance of the fund *f*'s holding in the MSA *m* in the year *t*. We measure performance as the 12-months ahead weighted average DGTW-adjusted returns over the stocks the fund holds within the MSA, where the weights refer to the fund's holdings. For the MSA for which the fund carries no stocks, the fund performance is set to zero. As in equation (1), the α components refer to fixed effects terms.

The effect of a flight introduction is captured by the dummies Held and Treatment. Held is a dummy equal to 1 if the fund has any positive holdings in stocks from that MSA. Treatment is equal to 1 if at some point prior to that year a direct flight was initiated between the fund's location and the MSA. The variable of interest in this analysis becomes the interaction term *Held*Treatment*: it captures the difference-in-difference of the excess performance of the fund's portfolios with positive holdings vs. the zero holdings in treated MSAs compared to the same excess performance in non-treated MSAs.

<Insert Table 5 here>

Table 5 presents the results from running equation (2). As before, we show variations of the model depending on which fixed effects are being utilized. Since most inferences are similar across the specifications, we focus the discussion on the full-fledged model, employing all fixed effects, in column (5). Recall that the baseline performance refers to a fund and MSA such that the fund has no holdings in the MSA (Held=0) and such that the pair has not yet been treated (Treatment=0). The coefficient on Held thus captures the DGTW-adjusted performance of a fund's holdings when the MSA is not treated. The negative coefficient suggests that the average abnormal performance of fund holdings is negative. The coefficient on the Treatment variable is not significantly different from zero: it confirms that simply treating an MSA does not change the performance unless the fund actually has holdings there. Finally, the coefficient on the interaction term Held*Treatment is positive and significant in all specifications. It suggest an improvement in the fund's ability to choose stocks in an MSA after the initiation of a direct flight between the fund's location and the MSA. More specifically, in column (5) the coefficient is equal to 0.0174 (t=3.765), suggesting an improvement of 1.74% in yearly adjusted-return for the performance of funds' portfolios with positive holdings vs. the zero holdings in treated MSAs compared to the same gap in non-treated MSAs.

<Insert Table 6 here>

We next examine the dynamic effects of introduction of flights on performance. As with the analysis of holdings, we replace the Treatment coefficient with three dummies that better distinguish the timing of the observation with respect to the introduction of the direct flight. Table 6 shows the results. Among the interaction terms of Held and the treatment dummies, only the Held*Treatment(1+) shows up significantly, suggesting that the performance effect appears only after the treatment takes place.

3.3 Monitoring vs. Information Acquisition

While we have documented strong results concerning the holdings and performance of mutual funds after the introduction of a nonstop flight, we now attempt to shed some light on the source of the outperformance that we observe. It could be the case that mutual funds who are long-term investors are better monitors of the firm, thus leading to improved performance in their investments without necessarily acquiring new information about the firm. However, we would only expect to see such monitoring on relatively large investments in the firm's portfolio. Thus we consider a series of regressions where we limit our sample to only pairs of fund and firm such that the fund's holdings in the firm are always below some threshold.

<Insert Table 7 here>

The results are shown in Table 7 for holdings (Panel A) and performance (Panel B). The baseline scenario (Column 1) is our original regression with all fixed effects and all levels of holdings included. In Columns 2-4 we restrict the sample to only observations where the fund always holds less than 0.5%, 1%, or 5% of its portfolio in the stock. For both holdings and performance regressions, the coefficient estimates for the relevant treatment variables are all significant at the 1% level. This suggests that the treatment effect we observe is present even for relatively small holdings. We interpret this as evidence that the performance benefit is more likely to be a result of information acquisition through contact with the firm, as opposed to benefits of increased monitoring of the firm.

4. Robustness

In this section we run a battery of tests to strengthen the interpretation of a causal relationship between reduction in travel time, fund holdings, and performance.

4.1 Number of Flights

Recall that the definition of a treatment between locations *A* and *B* at time *t* refers to the initiation of a direct flight between the two locations such that: the initiation has at least 10 flights per month and there were no direct flights linking *A* to *B* prior to time *t*. A natural question is whether there is something special about the threshold of 10 flights per month. To address this concern we repeat the analysis of holdings and performance for different cutoff values for number of flights in the definition of a treatment.

<Insert Table 8 here>

Results appear in Table 8. Each column in the output refers to a different cutoff: a treatment for a model with *# of initial flights*=*n* denotes the initiation of a direct flight having least *n* occurrences per month, and such that there were no direct flights prior to that period. Results are robust to the different cutoffs, with the only noticeable shift in the magnitude of the coefficients happening when going from *n*=1 to *n*=2. [6]

[6] An alternative way to define a treatment comes from relaxing the restriction that there should be no direct flight prior to the treatment period. For example, one could define a treatment when a minimum of 10 direct flights per month occurs for the first time between locations *A* and *B*. In this case, an initiation of a direct flight with 9 occurrences between *A* and *B* would not be recognized as a treatment. Moreover, an initiation of a direct flight with 10 occurrences would be classified as a treatment, though this treatment could be preceded by months with direct flights already operating between locations *A* and *B*. In unreported results, we repeat the analysis of holdings and performance with this alternative definition. All results qualitatively mimic the results discussed here, with exception of the analysis of dynamic effects. Under this alternative definition, we do see increases in holdings prior to the formal recognition of the treatment—not unexpectedly, though, given that we allow some direct flights to operate well before a treatment is formally recognized.

4.2 Eventually Treated Pairs

Fund-MSA pairs that experience the initiation of a direct flight are different from the average fund-MSA pair. This is not unexpected, at least because direct flights tend to link more economically robust locations— locations that thus have more funds and more companies, so that average holdings for these eventually treated MSA-pairs are higher. This raises concerns regarding the comparability between the treated observations and the rest of the sample that works as control group (see also the analysis in Table 2).

<Insert Table 9 here>

To address this concern, we employ the robustness test suggested by Bertrand and Mullainathan (2003), in which the analysis includes only observations from fund-MSA pairs that are eventually treated. Thus, the same fund-MSA pair can function both as control units (for the periods before the treatment) and treatment units (for the periods after the treatment takes place). This design is possible given the staggered introduction of direct flights (see Figure 3).

Table 9 shows the results of estimating the regressions for the analysis of holdings (column (1) of Panel A) and performance (column (1) of Panel B) for the subsample of eventually treated pairs. The results are qualitatively similar to those for the overall sample.

4.3 Placebo Tests

We finish our robustness analysis with two placebo tests. In the first one, we create placebo treatments, in which each quarter we randomly select pairs of MSAs as treated locations. The pool of candidates each quarter includes all pairs of MSAs such that: (1) the pair has not had any direct flight prior to that point; and (2) the pair has at least one fund in on location and one company in the other location. Requirement (1) attests to the definition of a treated pair of locations and requirement (2) avoids placebo treatments that are irrelevant with respect to fund holdings. Further, we force the number of placebo treatments each quarter to match the actual number of treated pairs of locations in the real sample.

We then repeat the analysis of fund holdings and performance under the new definition of treatment. Results appear in column (2) of Panels A and B of Table 9. All the inferences regarding increases in holdings and improvements in performance completely disappear.

Our second placebo analysis keeps the actual sample of treatments but creates a placebo sample of fund holdings. For this, we replace the set of actively managed funds with index funds. Index funds are passive by definition. They are unable to capitalize on any potential improvement in information acquisition when travel time is reduced to update holdings in the treated MSA. Therefore, we should not expect any effect from flight initiations on index funds' portfolio holdings and on the performance of these holdings. When we rerun the holdings and performance models for the sample of index funds (results appear in Table 9), all the inferences regarding increase in holdings and improvement in performance completely disappear.

5. Conclusion

In this paper we revisit the question of whether local bias affects the holdings decision and investment performance of mutual fund managers. Using the introduction of direct flights as a plausibly exogenous shock to the "proximity" between a fund and a firm, we find that fund managers increase their holdings when they become "closer" to a firm. Furthermore, we find that a fund's returns in holdings from a particular MSA improve when a new flight facilitates easier access to the firm for the manager(s) of the fund. Our results are robust to a number of different ways of defining treatments and are not present for a sample of index funds, suggesting the introduction of flights does have a meaningful effect on the behavior of active mutual fund managers.

Our results make an important contribution to the literature on geographic proximity and investment performance. Some evidence in the extant literature suggests that there is an informational advantage that comes from being located close to a firm. However, existing studies are unable to determine whether the improvement in investment performance comes from the acquisition of information directly from the firm or from simply being in the community near the firm (and thus benefitting from word-of-mouth interactions with other informed investors). Since our results involve investors located far away from the firm in question, we interpret our results as strong evidence that managers benefit from the decreased costs of visiting the firms themselves.

References

Bernile, G., Kogan, S. and Sulaeman, J., 2015. Information Environment and the Geography of Firms and Investors. Working Paper, Southern Methodist University.

Bernile, G., Kumar, A. and Sulaeman, J., 2015. Home away from home: Geography of information and local investors. *Review of Financial Studies* 28, 2009-2049.

Bernile, G., Kumar, A. and Sulaeman, J., 2015. Leveling the Playing Field: Financial Regulation and Disappearing Local Bias of Institutional Investors, Working Paper, Southern Methodist University.

Bernstein, S., X. Giroud, and R. Townsend, 2016. The Impact of Venture Capital Monitoring. *The Journal of Finance*, 71, 1591-1622.

Bertrand, M. and S. Mullainathan, 2003. Enjoying the Quiet Life? Corporate Governance and Managerial Preferences. *Journal of Political Economy* 111, 1043-1075.

Bodnaruk, A., 2009. Proximity Always Matters: Local Bias When the Set of Local Companies Changes. *Review of Finance* 13, 629-656.

Coval, J., and T. Moskowitz, 1999. Home Bias at Home: Local Equity Preference in Domestic Portfolios. *The Journal of Finance* 54, 2045-2073.

Coval, J., and T. Moskowitz, 2001. The Geography of Investment: Informed Trading and Asset Prices. *Journal of Political Economy* 109, 811-841

Cremers, M. and A. Petajisto, 2009. How Active Is Your Fund Manager? A New Measure That Predicts Performance. *Review of Financial Studies* 22, 3329-3365.

Daniel, K., M. Grinblatt, S. Titman, and R. Wermers, 1997. Measuring Mutual Fund Performance with Characteristic-based Benchmarks. *Journal of Finance*, 52, 1035–1058.

Giroud, X., 2013. Proximity and Investment: Evidence from Plant-Level Data, *Quarterly Journal of Economics*. 128, 861-915.

Heath. C., and A. Tversky, 1991. Preference and Beliefs: Ambiguity and Competence in Choice Under Uncertainty. *Journal of Risk and Uncertainty* 4, 5-28.

Huberman, G., 2001. Familiarity Breeds Investment. *Review of financial Studies*, *14*(3), pp.659-680.

Ivković, Z., and S. Weisbenner, 2005. Local Does as Local Is: Information Content of the Geography of Individual Investors' Common Stock Investments. *The Journal of Finance* 60, 267–306

Massa, M. and Simonov, A., 2006. Hedging, Familiarity and Portfolio Choice. *Review of Financial Studies* 19, 633-685.

Pool, V., N. Stoffman, and S. Yonker, 2015. The People in Your Neighborhood: Social Interactions and Mutual Fund Portfolios. *Journal of Finance* 70, 2679–2732.

Seasholes, M.S. and Zhu, N., 2010. Individual Investors and Local Bias. *The Journal of Finance* 65, 1987-2010.

Sialm, C., Sun, Z. and Zheng, L., 2013. Home Bias and Local Contagion: Evidence from Funds of Hedge Funds, *National Bureau of Economic Research* Working Paper.

Sulaeman, J., 2014. Do Local Investors Know More? Evidence from Mutual Fund Location and Investments. *The Quarterly Journal of Finance*, *4*(02), p.1450010.

Figure 1. Number of treatments per Fund MSA

This figure plots the geographical distribution of MSAs with funds in our sample, and the number of treatments per MSA. A treatment means that at some point during our sample period there was the initiation of a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) having that MSA as one of the end-points.

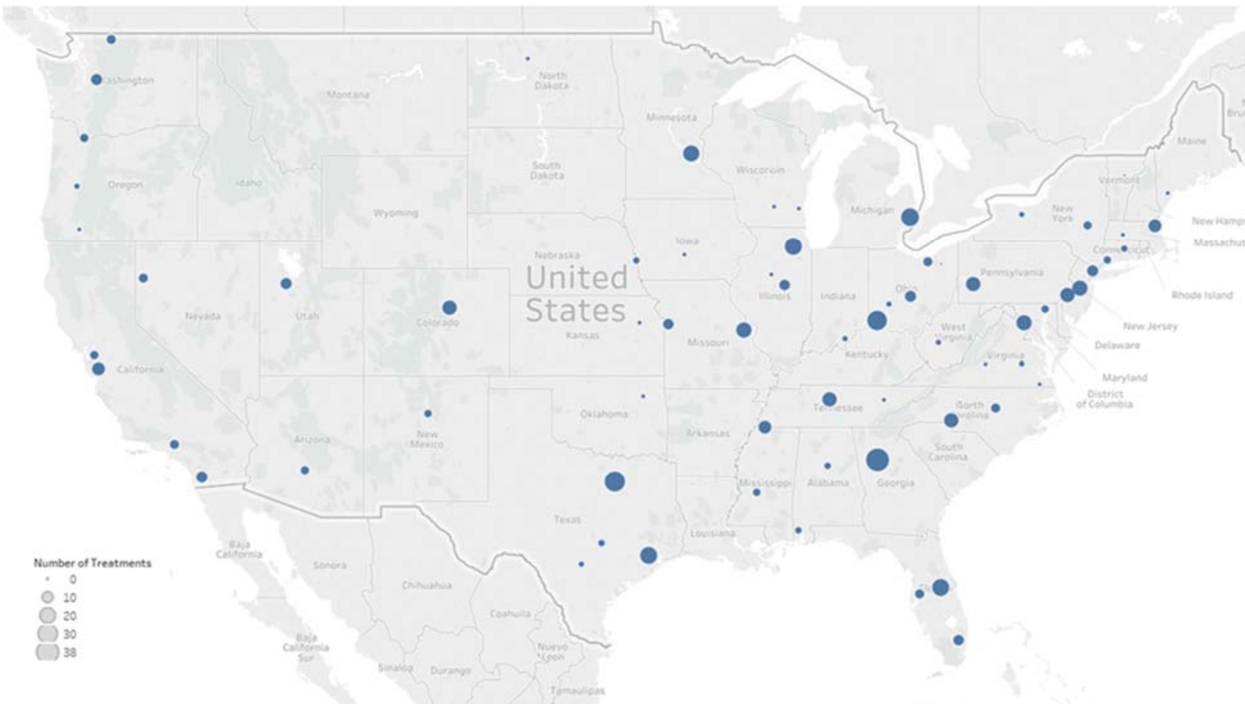

Figure 2. Number of treatments per Company MSA

This figure plots the geographical distribution of MSAs with companies (whose stocks are held by mutual funds) in our sample, and the number of treatments per MSA. A treatment means that at some point during our sample period there was the initiation of a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) having that MSA as one of the end-points.

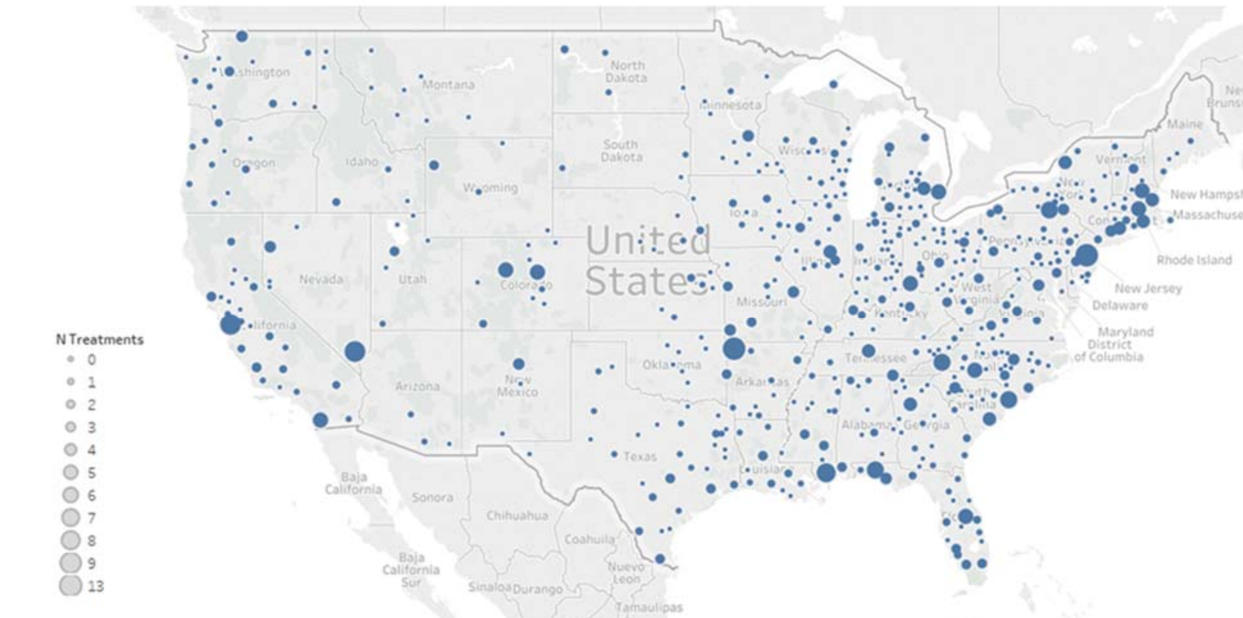

Figure 3. Number of Treated Pairs of MSAs per Year

This figure presents the yearly distribution of treated pairs of MSAs. A treated pair is a pair of MSAs such that at some point during our sample period a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) was initiated between the two MSAs.

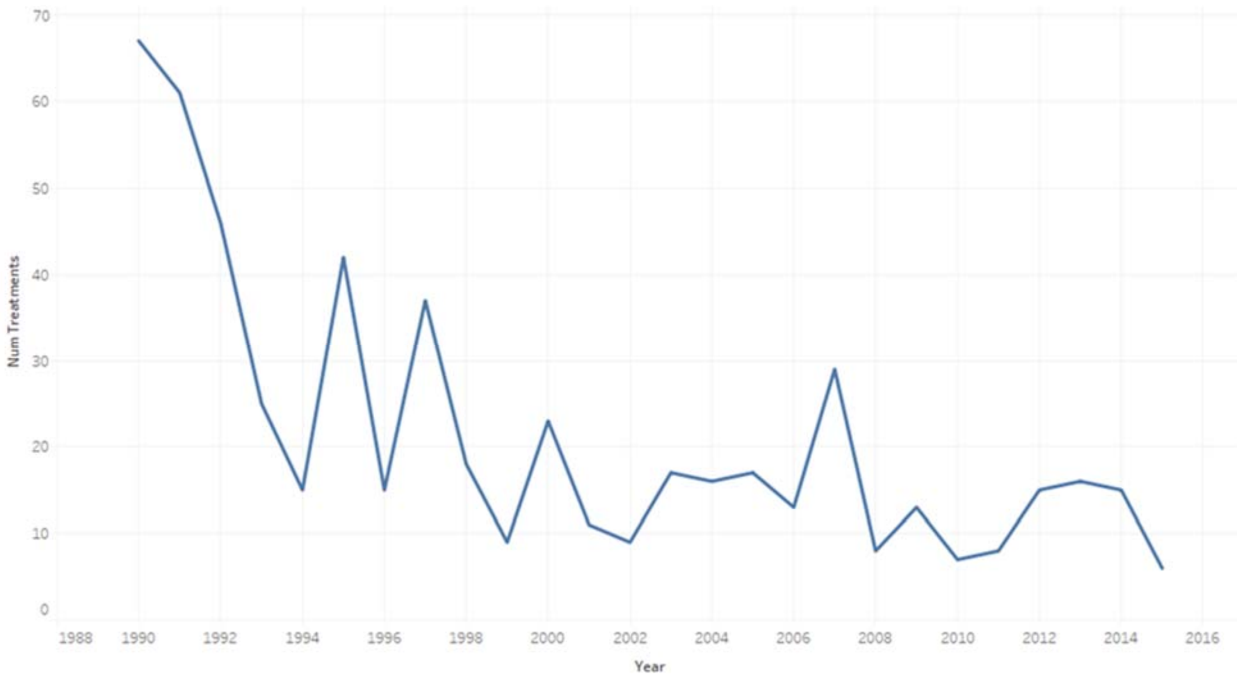

Table 1. Holdings by Travel Time and Distance

This table presents the average number of stocks held in mutual fund portfolios (Panel A) and the average weight of those stocks in a fund's portfolio (Panel B) based on distance between the fund's MSA and the company's MSA, as well as the travel time between the two MSA's. *** indicates significance at the 1% level.

Panel A: Number of stocks held in portfolio

	Quartiles of Travel Time				
Distance (miles)	1 (shortest)	2	3	4 (longest)	Q1-Q4
0-250	4.5	3.3	5.9	1.9	2.59***
251-500	3.3	3	2.6	1.9	1.36***
501-750	2.4	3	3.8	1.5	0.92***
751-1000	3.2	3	2.8	1.6	1.61***
1001-1250	3.4	3.2	2.3	1.6	1.80***
1251-1500	3.8	5.1	3	2	1.83***
1501-1750	4.1	5.8	3.7	1.7	2.34***
1751-2000	3.9	3.8	2.3	2.2	1.75***
2001-2250	3.5	3.3	2.2	2	1.42***
2251-2500	5.2	4.1	3.2	2.9	2.27***
>2500	7.5	6.2	4.5	2.4	5.08***

Panel B: Weight in Portfolio

	Quartiles of Travel Time				
Distance (miles)	1 (shortest)	2	3	4 (longest)	Q1-Q4
0-250	4.3	3	6.9	1.8	2.53***
251-500	3	2.8	2.6	1.7	1.26***
501-750	2.7	3.1	4.2	1.4	1.22***
751-1000	3.5	3.2	2.7	1.4	2.07***
1001-1250	3.6	3.3	2.1	1.3	2.29***
1251-1500	3.3	5.4	2.7	1.7	1.56***
1501-1750	4.2	5.7	3.1	1.5	2.71***
1751-2000	4	3.4	2.1	1.7	2.25***
2001-2250	3.1	2.9	2	1.6	1.46***
2251-2500	5	3.4	3.1	2.7	2.27***
>2500	8.5	5.1	4.3	2.1	6.40***

Table 2. Holdings by Treatment Status

This table presents summary statistics on yearly fund holdings in MSAs. The observations are tuples (fund *f*, MSA *m*, year *t*), defined as the dollar value of holdings by fund *f* of stocks located in MSA *m* in year *t.* Results are summarized depending on the classification of the observation. A "later treated" observation denotes that there is no direct flight between fund *f*'s location and MSA *m* in year *t*, but a direct flight is later established between the two locations. Treated means a direct flight has been initiated between fund *f*'s location and MSA *m* during our sample period, but prior to year *t*. Never treated denotes that there has never been a direct flight between fund *f*'s location and MSA *m* during our sample period. Always served denotes+ that a direct flight between fund *f*'s location and MSA *m* was already available in the beginning of our sample period. Local means the fund is located in the MSA *m*. Close denotes that fund *f* is not located in MSA *m*, but fund *f*'s location and MSA *m* are less than 100 miles apart.

		Later treated	Treated	Never treated	Always served	Local	Close
	Mean	1.61	15.34	0.31	21.17	57.45	4.12
All	95th percentile	1.26	35.09	0.00	70.82	261.10	8.91
	% with zero holdings	92%	77%	97%	53%	18%	84%
	Mean	2.28	3.15	0.21	8.40	24.06	1.62
Year=1991	95th percentile	6.73	14.46	0.00	30.87	97.24	5.50
	% with zero holdings	85%	74%	97%	57%	20%	84%
	Mean	1.16	11.44	0.26	19.36	51.11	2.61
Year=2001	95th percentile	0.19	21.99	0.00	58.90	294.16	5.47
	% with zero holdings	94%	78%	98%	56%	19%	87%
	Mean	4.67	23.95	0.50	27.38	73.39	7.39
Year=2011	95th percentile	1.63	61.12	0.00	102.49	309.05	22.58
	% with zero holdings	87%	74%	97%	51%	18%	82%

Table 3. Holdings Regressions

This table presents the results of regressions explaining yearly fund holdings in MSAs. The observations are tuples (fund *f*, MSA *m*, year *t*). The dependent variable is either the logarithm of 1+holdings of fund *f* in MSA *m* in year *t*, or a dummy for whether these holdings are positive. Treatment is equal to 1 if at some point prior to year *t* a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) was initiated between the fund *f*'s MSA and MSA *m*. There are 5 regression models per dependent variable, depending on which set of fixed effects is being used. All regressions use clustered standard errors per fund-MSA pair. ***, **, * denote statistical significance at the 1%, 5%, 10% levels, respectively.

	Ln(1+Holdings)					Holdings>0				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	2.1156***	0.5465***	0.5619***	0.2237***	0.2363***	0.1332***	0.0309***	0.0317***	0.0121***	0.0127***
	[134.343]	[8.336]	[8.662]	[3.543]	[3.804]	[128.213]	[7.478]	[7.781]	[2.919]	[3.139]
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund x MSA FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Fund x Year FE	No	No	Yes	No	Yes	No	No	Yes	No	Yes
MSA x Year FE	No	No	No	Yes	Yes	No	No	No	Yes	Yes
Adj-R^2	0.00545	0.598	0.604	0.606	0.612	0.00492	0.573	0.580	0.580	0.586
Observations	4,640,695	4,586,922	4,586,922	4,586,922	4,586,922	4,640,695	4,586,922	4,586,922	4,586,922	4,586,922

Table 4. Holdings: Dynamic Effects

This table presents the results of regressions analyzing the dynamics of treatment effects on yearly fund holdings in MSAs. The observations are tuples (fund *f*, MSA *m*, year *t*). The dependent variable is either the logarithm of 1+holdings of fund *f* in MSA *m* in year *t*, or a dummy for whether these holdings are positive. Treatment is equal to 1 if at some point prior to year *t* a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) was initiated between fund *f*'s MSA and MSA *m*. A treatment refers to the initiation of a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) linking fund *f*'s location and MSA *m*. Treatment (-1) is equal to 1 for the year preceding such initiation. Treatment (0) is equal to 1 for the year of the initiation. Treatment (+1) is equal to 1 for any year after the initiation of the direct flight. There are 5 regression models per dependent variable, depending on which set of fixed effects is being used. All regressions use clustered standard errors per fund-MSA pair. ***, **, * denote statistical significance at the 1%, 5%, 10% levels, respectively.

	Ln(1+Holdings)					Holdings>0				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment (-1)	0.3453***	0.1213*	0.1538**	0.0289	0.0640	0.0240***	0.0090*	0.0109**	0.0029	0.0050
	[4.551]	[1.785]	[2.270]	[0.428]	[0.949]	[4.793]	[1.955]	[2.396]	[0.617]	[1.082]
Treatment (0)	0.1899**	-0.0119	0.0179	-0.0757	-0.0442	0.0133**	0.0020	0.0037	-0.0042	-0.0024
	[2.266]	[-0.142]	[0.215]	[-0.936]	[-0.549]	[2.399]	[0.360]	[0.678]	[-0.767]	[-0.436]
Treatment (1+)	2.1852***	0.6644***	0.6871***	0.2798***	0.2999***	0.1376***	0.0378***	0.0390***	0.0155***	0.0166***
	[136.375]	[8.652]	[9.054]	[3.787]	[4.133]	[130.109]	[7.832]	[8.223]	[3.211]	[3.521]
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund x MSA FE	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Fund x Year FE	No	No	Yes	No	Yes	No	No	Yes	No	Yes
MSA x Year FE	No	No	No	Yes	Yes	No	No	No	Yes	Yes
Adj-R^2	0.006	0.643	0.650	0.651	0.658	0.005	0.622	0.629	0.628	0.635
Observations	4,640,695	4,586,922	4,586,922	4,586,922	4,586,922	4,640,695	4,586,922	4,586,922	4,586,922	4,586,922

Table 5. Performance Regressions

This table presents the results of regressions explaining yearly performance of fund holdings in MSAs. The observations are tuples (fund *f*, MSA *m*, year *t*). The dependent variable is the performance of fund *f*'s holding in MSA *m* in year *t*. We measure performance as the 12-months ahead weighted average DGTW-adjusted returns over the stocks the fund holds in the MSA, where the weights refer to the fund's holdings. For an MSA for which the fund holds no stocks, the fund performance is set to zero. Held is a dummy equal to 1 if the fund *f* has any positive holdings in stocks from MSA *m* in year *t*. Treatment is equal to 1 if at some point prior to year *t* a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) was initiated between the fund *f*'s MSA and MSA *m*. There are 5 regression models, depending on which set of fixed effects is being used. All regressions use clustered standard errors per fund-MSA pair. ***, **, * denote statistical significance at the 1%, 5%, 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Held	0.0042***	-0.0079***	-0.0082***	-0.0057***	-0.0059***
	[24.897]	[-10.389]	[-10.676]	[-7.476]	[-7.710]
Treatment	-0.0001	0.0015	0.0017	-0.0015	-0.0014
	[-0.244]	[1.169]	[1.271]	[-1.029]	[-0.929]
Held*Treatment	0.0281***	0.0152***	0.0154***	0.0172***	0.0174***
	[31.968]	[3.290]	[3.343]	[3.707]	[3.765]
Year FE	Yes	Yes	Yes	Yes	Yes
Fund x MSA FE	No	Yes	Yes	Yes	Yes
Fund x Year FE	No	No	Yes	No	Yes
MSA x Year FE	No	No	No	Yes	Yes
Adj-R^2	0.000874	-0.0190	-0.0173	0.0255	0.0273
Observations	4,254,909	4,199,345	4,199,345	4,199,343	4,199,343

Table 6. Performance: Dynamic Effects

This table presents the results of regressions analyzing the dynamics of treatment effects on yearly performance of fund holdings in MSAs. The observations are tuples (fund *f*, MSA *m*, year *t*). The dependent variable is the performance of fund *f*'s holding in MSA *m* in year *t*. We measure performance as the 12-months ahead weighted average DGTW-adjusted returns over the stocks the fund holds in the MSA, where the weights refer to the fund's holdings. For an MSA for which the fund holds no stocks, the fund performance is set to zero. Held is a dummy equal to 1 if the fund *f* has any positive holdings in stocks from MSA *m* in year *t*. A treatment refers to the initiation of a direct flight (with at least 10 occurrences per month, and such that there were no direct flights prior to that period) linking fund *f*'s location and MSA *m*. Treatment (-1) is equal to 1 for the year preceding such initiation. Treatment (0) is equal to 1 for the year of the initiation. Treatment (+1) is equal to 1 for any year after the initiation of the direct flight. There are 5 regression models, depending on which set of fixed effects is being used. All regressions use clustered standard errors per fund-MSA pair. ***, **, * denote statistical significance at the 1%, 5%, 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Held	0.0042***	-0.0079***	-0.0081***	-0.0057***	-0.0059***
	[24.961]	[-10.343]	[-10.626]	[-7.487]	[-7.717]
Treatment (-1)	0.0002	0.0012	0.0017	0.0031**	0.0036***
	[0.093]	[1.155]	[1.596]	[2.542]	[2.939]
Treatment (0)	-0.0000	-0.0004	-0.0004	-0.0063***	-0.0064***
	[-0.009]	[-0.307]	[-0.304]	[-3.015]	[-2.989]
Treatment (1+)	-0.0001	0.0016	0.0018	0.0004	0.0007
	[-0.247]	[1.023]	[1.185]	[0.273]	[0.449]
Held*Treatment (-1)	-0.0131**	-0.0136	-0.0152	0.0079	0.0063
	[-2.408]	[-0.844]	[-0.942]	[0.486]	[0.389]
Held*Treatment (0)	0.0413***	0.0367*	0.0366*	0.0244	0.0244
	[6.633]	[1.900]	[1.901]	[1.286]	[1.286]
Held*Treatment (1+)	0.0278***	0.0141***	0.0142***	0.0171***	0.0172***
	[31.360]	[3.029]	[3.071]	[3.658]	[3.706]
Year FE	Yes	Yes	Yes	Yes	Yes
Fund x MSA FE	No	Yes	Yes	Yes	Yes
Fund x Year FE	No	No	Yes	No	Yes
MSA x Year FE	No	No	No	Yes	Yes
Adj-R^2	0.000876	-0.0190	-0.0173	0.0255	0.0273
Observations	4,254,909	4,199,345	4,199,345	4,199,343	4,199,343

Table 7. Subsampling based on Size of Holdings

This table presents the results of regressions explaining yearly fund holdings in MSAs (Panel A) and yearly performance of fund holdings in MSAs (Panel B). The observations and variables are as defined in table 3 (for Panel A) and table 5 (for Panel B). Further, the sample in Model (2)[3]{4} is based only on pairs fund-firm such that the fund's holdings in the firm are always below a threshold of 5%[1%]{0.5%} of the outstanding shares of the firm. All regressions use clustered standard errors per fund-MSA pair. ***, **, * denote statistical significance at the 1%, 5%, 10% levels, respectively.

Panel A: Dependent variable=Ln(1+Holdings)

	(1) All	(2) Threshold=5%	(3) Threshold=1%	(4) Threshold=0.5%
Treatment	0.2363*** [3.804]	0.2301*** [3.665]	0.2061*** [3.365]	0.1844*** [3.177]
Year FE	Yes	Yes	Yes	Yes
Fund x MSA FE	Yes	Yes	Yes	Yes
Fund x Year FE	Yes	Yes	Yes	Yes
MSA x Year FE	Yes	Yes	Yes	Yes
Adj-R^2	0.612	0.608	0.595	0.587
Observations	4,586,922	4,552,653	4,390,217	4,260,412

Panel B: Dependent variable=Performance

	(1) All	(2) Threshold=5%	(3) Threshold=1%	(4) Threshold=0.5%
Held	-0.0059*** [-7.710]	-0.0059*** [-7.695]	-0.0058*** [-7.656]	-0.0058*** [-7.580]
Treatment	-0.0014 [-0.929]	-0.0014 [-0.926]	-0.0013 [-0.878]	-0.0015 [-0.987]
Held*Treatment	0.0174*** [3.765]	0.0174*** [3.756]	0.0173*** [3.740]	0.0174*** [3.769]
Year FE	Yes	Yes	Yes	Yes
Fund x MSA FE	Yes	Yes	Yes	Yes
Fund x Year FE	Yes	Yes	Yes	Yes
MSA x Year FE	Yes	Yes	Yes	Yes
Adj-R^2	0.0273	0.0273	0.0263	0.0260
Observations	4,199,343	4,175,963	4,054,141	3,946,546

Table 8. Redefining Treatment based on Different Numbers of Flights

This table presents the results of regressions explaining yearly fund holdings in MSAs (Panel A) and yearly performance of fund holdings in MSAs (Panel B). The observations and variables are as defined in table 3 (for Panel A) and table 5 (for Panel B). Each model is based on the *number of flights=n* variable, such that a treatment in that model is equal to 1 if at some point prior to that year a direct flight (with at least *n* occurrences per month, and such that there were no direct flights prior to that period) was initiated between the fund's MSA and the MSA. All regressions use clustered standard errors per fund-MSA pair. ***, **, * denote statistical significance at the 1%, 5%, 10% levels, respectively.

Panel A: Dependent variable=Ln(1+Holdings)

# of initial flights:	1	2	3	4	5	6	7	8	9	10
Treatment	0.0911***	0.2505***	0.2814***	0.2338***	0.2363***	0.2383***	0.2452***	0.2459***	0.2368***	0.2363***
	[4.628]	[4.961]	[5.002]	[4.123]	[4.081]	[4.080]	[4.059]	[4.034]	[3.867]	[3.804]
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund x MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund x Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA x Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj-R^2	0.612	0.612	0.612	0.612	0.612	0.612	0.612	0.612	0.612	0.612
Observations	4,586,922	4,586,922	4,586,922	4,586,922	4,586,922	4,586,922	4,586,922	4,586,922	4,586,922	4,586,922

Table 8. (*Continued*)

Panel B: Dependent variable=Performance

# of initial flights:	1	2	3	4	5	6	7	8	9	10
Held	-0.0067***	-0.0062***	-0.0061***	-0.0058***	-0.0059***	-0.0059***	-0.0059***	-0.0058***	-0.0059***	-0.0059***
	[-7.885]	[-8.076]	[-7.961]	[-7.665]	[-7.698]	[-7.684]	[-7.693]	[-7.677]	[-7.709]	[-7.710]
Treatment	-0.0006	0.0003	-0.0019	-0.0015	-0.0017	-0.0017	-0.0017	-0.0016	-0.0014	-0.0014
	[-1.515]	[0.220]	[-1.415]	[-1.100]	[-1.207]	[-1.190]	[-1.112]	[-1.090]	[-0.934]	[-0.929]
Held* Treatment	0.0059***	0.0153***	0.0203***	0.0155***	0.0167***	0.0166***	0.0169***	0.0166***	0.0174***	0.0174***
	[3.440]	[4.723]	[4.691]	[3.399]	[3.614]	[3.629]	[3.678]	[3.586]	[3.760]	[3.765]
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund x MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund x Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA x Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj-R^2	0.0273	0.0274	0.0274	0.0273	0.0273	0.0273	0.0273	0.0273	0.0273	0.0273
Observations	4,199,343	4,199,343	4,199,343	4,199,343	4,199,343	4,199,343	4,199,343	4,199,343	4,199,343	4,199,343

Table 9. Robustness Analysis

This table presents the results of regressions explaining yearly fund holdings in MSAs (Panel A) and yearly performance of fund holdings in MSAs (Panel B). The observations and variables are as defined in table 3 (for Panel A) and table 5 (for Panel B). Further, the sample in Model (1) is restricted to only fund-MSA pairs that are eventually treated. In model (2) treatments are randomly assigned: each quarter we randomly select pairs of MSAs as treated locations. The pool of candidates each quarter includes all pairs of MSAs such that: (1) the pair has not had any direct flight prior to that point; and (2) the pair has at least one fund in one location and one company in the other location. We force the number of placebo treatments each quarter to match the actual number of treated pairs of locations in the real sample. The sample in Model (3) is restricted to index funds only. All regressions use clustered standard errors per fund-MSA pair. ***, **, * denote statistical significance at the 1%, 5%, 10% levels, respectively.

Panel A: Dependent variable=Ln(1+Holdings)

	(1) Eventually treated pairs	(2) Placebo treatments	(3) Index funds
Treatment	0.8113***	0.0759	-0.1662
	[4.858]	[1.078]	[-0.986]
Year FE	Yes	Yes	Yes
Fund x MSA FE	Yes	Yes	Yes
Fund x Year FE	Yes	Yes	Yes
MSA x Year FE	Yes	Yes	Yes
Adj-R^2	0.701	0.612	0.861
Observations	99,447	4,586,922	659,835

Panel B: Dependent variable=Performance

	(1) Eventually treated pairs	(2) Placebo treatments	(3) Index funds
Held	-0.0074	-0.0054***	-0.0101***
	[-0.779]	[-7.088]	[-5.581]
Treatment	-0.0041	0.0015	-0.0015
	[-0.522]	[1.425]	[-0.321]
Held*Treatment	0.0219**	-0.0003	0.0072
	[2.042]	[-0.029]	[0.935]
Year FE	Yes	Yes	Yes
Fund x MSA FE	Yes	Yes	Yes
Fund x Year FE	Yes	Yes	Yes
MSA x Year FE	Yes	Yes	Yes
Adj-R^2	0.0578	0.0273	0.207
Observations	94,190	4,199,343	598,238