

PROXIMITY AND INVESTMENT: EVIDENCE FROM PLANT-LEVEL DATA*

XAVIER GIROUD

Proximity to plants makes it easier for headquarters to monitor and acquire information about plants. In this article, I estimate the effects of headquarters' proximity to plants on plant-level investment and productivity. Using the introduction of new airline routes as a source of exogenous variation in proximity, I find that new airline routes that reduce the travel time between headquarters and plants lead to an increase in plant-level investment of 8% to 9% and an increase in plants' total factor productivity of 1.3% to 1.4%. The results are robust when I control for local and firm-level shocks that could potentially drive the introduction of new airline routes, when I consider only new airline routes that are the outcome of a merger between two airlines or the opening of a new hub, and when I consider only indirect flights where either the last leg of the flight (involving the plant's home airport) or the first leg of the flight (involving headquarters' home airport) remains unchanged. Moreover, the results are stronger in the earlier years of the sample period and for firms whose headquarters is more time-constrained. In addition, they also hold at the extensive margin, that is, when I consider plant openings and closures. *JEL* Codes: D24, G31.

I. INTRODUCTION

Proximity facilitates monitoring and access to information. For instance, venture capitalists are more likely to serve on the boards of local firms, where monitoring is easier (Lerner 1995). Likewise, mutual fund managers are more likely to hold shares of local firms—and they earn significant abnormal returns from these investments—suggesting “improved monitoring

*This article is based on my dissertation submitted to New York University. I thank my advisor, Holger Mueller; the editors, Jeremy Stein, Andrei Shleifer, and Larry Katz; three anonymous referees; as well as Viral Acharya, Ashwini Agrawal, Allan Collard-Wexler, Carola Frydman, Xavier Gabaix, Kose John, Marcin Kacperczyk, Andrew Karolyi, Leonid Kogan, Anthony Lynch, Javier Miranda, Adair Morse, Dimitris Papanikolaou, Adriano Rampini, Michael Roberts, Alexi Savov, Philipp Schnabl, Antoinette Schoar, Amit Seru, Daniel Wolfenzon; and seminar participants at MIT, Chicago, Stanford, NYU, Wharton, Kellogg, UCLA, Yale, Duke, Ohio State, Cornell, and USC for valuable comments and suggestions. The research in this article was conducted while the author was a Special Sworn Status researcher of the U.S. Census Bureau at the New York and Boston Census Research Data Centers. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

© The Author(s) 2012. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2013), 861–915. doi:10.1093/qje/qjs073.

Advance Access publication on December 19, 2012.

capabilities or access to private information of geographically proximate firms” (Coval and Moskowitz 1999, 2001, 812). Finally, banks located closer to their borrowers are more likely to lend to informationally difficult borrowers, for example, borrowers without any financial records (Petersen and Rajan 2002; Mian 2006; Sufi 2007).

All of these examples come from arm’s-length transactions. Much less, if anything, is known about the role of proximity *within* firms. For instance, is it true that—in analogy to the findings in the mutual funds and banking literatures—headquarters is more likely to invest in plants that are located closer to headquarters? Does proximity to headquarters improve plant-level productivity? Understanding plant-level investment and productivity is important, not least because they affect economic growth.¹ One difficulty in answering these questions is that they require data on the locations of plants and headquarters. Another, more serious issue is that the locations of plants and headquarters are choice variables. Accordingly, commonly used proxies for proximity—such as the physical distance between plants and headquarters—are likely to be endogenous, making it difficult to establish causality.

I attempt to address both of these issues. For the first issue, I use plant-level data provided by the U.S. Census Bureau for the manufacturing sector for the period 1977 to 2005, which include the locations of plants and headquarters. For the second issue, I notice that the main reason empirical studies are interested in (geographical) proximity is because it proxies for the ease of monitoring and acquiring information. I argue that a more direct proxy is travel time. For instance, a plant may be located far away from headquarters, yet monitoring may be easy, because there exists a short, direct flight. Conversely, a plant may be

1. Anecdotal evidence suggests that proximity to headquarters is a potentially important determinant of plant-level investment. For instance, when Tesla Motors decided on the location of a manufacturing plant to produce its electric Tesla roadster, it announced that the plant would be located “as close to our headquarters as possible,” citing as a reason “to keep better control over production” (*Silicon Valley/San Jose Business Journal*, June 30, 2008). As for the effects of proximity on productivity, Ray Kroc, the founder of McDonald’s, writes in his autobiography: “One thing I liked about that house was that it was perched on a hill looking down on a McDonald’s store on the main thoroughfare. I could pick up a pair of binoculars and watch business in that store from my living room window. It drove the manager crazy when I told him about it. But he sure had one hell of a hard-working crew!” (Kroc 1992, 141).

located in the same state as headquarters, yet monitoring may be costly, because it involves a long and tedious road trip. Of course, in the cross-section, geographical proximity and travel time are highly correlated. However, the advantage of using travel time is that it entails plausibly exogenous variation, allowing me to address the endogeneity issue.

Precisely, I combine the census plant-level data with airline data from the U.S. Department of Transportation, which contain information about all flights that have taken place between any two airports in the United States. The specific source of exogenous variation that I exploit is the introduction of new airline routes that reduce the travel time between headquarters and plants. Using a difference-in-differences approach, I find that the introduction of new airline routes leads to an increase in plant-level investment of 8% to 9%, corresponding to an increase in capital expenditures of \$158,000 to \$177,000 (in 1997 dollars). Moreover, I find that plants' total factor productivity increases by 1.3% to 1.4%, corresponding to an increase in plant-level profits of \$275,000 to \$296,000 (in 1997 dollars).

My identification strategy can be illustrated with a simple example. Consider a company headquartered in Boston with a plant in Memphis. In 1985, the fastest way to travel from Boston to Memphis was an indirect flight with one stopover in Atlanta. In 1986, Northwest Airlines opened a new hub in Memphis and started operating direct flights between Boston and Memphis. The introduction of this new airline route substantially reduced the travel time between the Boston headquarters and the Memphis plant and is therefore coded as a "treatment" of the Memphis plant.²

An important concern is that local shocks in the plant's vicinity could be driving both the introduction of new airline routes and plant-level investment. For instance, suppose the local economy in Memphis is booming. As a result, the company headquartered in Boston may find it more attractive to increase investment at its Memphis plant. At the same time, airlines may find it more attractive to introduce new flights to Memphis (e.g., because of lobbying by local plants). In this case, finding a positive treatment effect would be a spurious outcome of an omitted shock in the

2. Overall, there are 10,533 plants in my sample that experience a reduction in the travel time to headquarters due to the introduction of new airline routes.

Memphis area. Fortunately, because a treatment is uniquely defined by *two* (airport) locations—the plant’s and headquarters’ home airports—I can control for such local shocks, making the identification even tighter. Specifically, I include metropolitan statistical area (MSA)–year controls in all my regressions. Similarly, to account for omitted shocks at the firm level, I include firm-year controls in all my regressions. Both types of controls are identified here, because not all local plants have their headquarters in the same city or region and not all plants of a given firm are affected by the introduction of a new airline route, respectively.

Although the inclusion of MSA- and firm-year controls accounts for omitted local and firm-level shocks, it remains the possibility of an omitted shock that is specific to a single plant—that is, the shock does not affect other plants within the same region or firm. I address this issue in three ways. First, I consider the dynamic effects of the introduction of new airline routes. If a new airline route is the (endogenous) outcome of a preexisting plant-specific shock, then I should find an “effect” of the treatment already before the new airline route is introduced. However, I find no such effect. Second, I show that my results are robust when I consider only new airline routes that are the outcome of a merger between two airlines or the opening of a new hub. Arguably, it is less likely that a shock to a single plant would trigger an airline merger or a hub opening. Third, I show that my results are robust when I consider only indirect flights where either the last leg of the flight (involving the plant’s home airport) or the first leg of the flight (involving headquarters’ home airport) remains unchanged. Arguably, it is less likely that a single plant or headquarters can successfully lobby for the introduction of a new flight elsewhere, that is, a flight that does not involve its respective home airport.

A limitation of my study is that by design, it relies on exogenous variation in travel time, not variation in monitoring or access to information. With this caveat in mind, I think it is plausible that a travel time reduction leads to an improvement in monitoring and information acquisition and, as a result, to an increase in plant-level investment and productivity. For instance, monitoring by headquarters may induce higher effort by plant managers and workers (see the McDonald’s example in note 1), thus improving the plant’s productivity and, along with it, the plant’s marginal return on investment. Moreover, just like

mutual fund managers may face less uncertainty with respect to local stocks, headquarters may face less uncertainty with respect to local projects. Consequently, headquarters may be in a better position to evaluate local projects and, as a result, assign a larger investment budget to them.

That being said, there are alternative stories that I cannot rule out. For instance, visiting the plant more often may allow the CEO to give better advice, thereby improving the plant's productivity and, by implication, its marginal return on investment. Or the plant managers may simply get a better sense of the company's needs. Or it may improve the plant managers' morale, who may think they have a better chance of getting promoted if their actions become more visible to headquarters. Or headquarters may devote more of its "limited attention," for example, in the budget allocation process, to plants that have become more salient to it.³ Although all of these alternative stories can be broadly categorized under the notion of "information transmission," they arguably have a different flavor. Importantly, however, all of them have to do with proximity, which is the main hypothesis explored in this article. Moreover, as the new airline routes are commercial (not cargo), they also all have to do with personal travel. Thus, the reason for plant-level investment increases is not, for example, that it becomes cheaper to ship equipment to the plant.

In the final part of this article, I provide auxiliary evidence that a reduction in travel time facilitates monitoring and information acquisition. For instance, I show that my results are stronger for plants whose headquarters is more "time-constrained," based on the notion that time constraints limit the ability to monitor and gather information about plants. Likewise, I show that my results are stronger in the earlier years of the sample period, where other, nonpersonal means of information transmission (e.g., Internet, corporate intranet, video conferencing) were either unavailable or less developed. I also examine what happens at the "extensive margin," that is, with respect to plant openings and closures. Specifically, I show that companies whose headquarters is more time-constrained are more likely to open new plants—and are less likely to close

3. For models of limited attention, see Gabaix et al. (2006), Gennaioli and Shleifer (2010), and Bordalo, Gennaioli, and Shleifer (2012). See also DellaVigna (2009, Section 4.2) for a survey of the literature.

existing plants—in close proximity to headquarters. Likewise, I show that over time—as innovations in information technology have reduced the need to personally travel to plants—firms have become more likely to open new plants at remote locations and close existing plants at proximate locations. The latter result is consistent with Petersen and Rajan's (2002) finding that, owing to innovations in information technology, the distance between small business borrowers and their lenders has increased over time.

The results in this article have important policy and welfare implications. From a policy perspective, they point to an important externality of transportation infrastructure: the facilitation of monitoring and information flows within firms. That being said, this externality seems to have become less important in the later years of the sample period, where innovations in information technology have facilitated information flows both within and across company units. Likewise, the results suggest that state or regional competition for firms to set up new plants may not be a zero-sum game. Specifically, if distant states or regions prevail in the competition—for example, because they can offer better tax breaks or other incentives—plant-level investment and productivity may be lower than if relatively proximate states or regions had won. Finally, a possible welfare implication of my results is that firms may be induced to locate plants suboptimally—that is, closer to headquarters than would be optimal in a frictionless world with perfect and symmetric information—for the sole purpose of facilitating monitoring and information transmission.

The rest of this article is organized as follows. Section II describes the data and empirical methodology. Section III presents the main results. Section IV contains robustness checks. Section V considers heterogeneity in the treatment effect. Section VI considers plant openings and closures. Section VII concludes. Appendix A and Appendix B provide information regarding the construction and measurement of variables.

II. DATA

II.A. Data Sources and Sample Selection

Plant-Level Data. The data on manufacturing plants are obtained from three different data sets provided by the U.S. Census

Bureau. The first data set is the Census of Manufactures (CMF). The CMF covers all U.S. manufacturing plants with at least one paid employee. The CMF is conducted every five years in years ending with 2 and 7 (“Census years”). The second data set is the Annual Survey of Manufactures (ASM). The ASM is conducted in all non-Census years and covers a subset of the plants covered by the CMF: plants with more than 250 employees are included in every ASM year, whereas plants with fewer employees are randomly selected every five years, where the probability of being selected is higher for larger plants. Although the ASM is referred to as a “survey,” reporting is mandatory, and fines are levied for misreporting. The CMF and ASM cover approximately 350,000 and 50,000 plants per year, respectively, and contain information about key plant variables, such as capital expenditures, total assets, value of shipments, material inputs, employment, industry sector, and location. The third data set is the Longitudinal Business Database (LBD), which is compiled from the Business Register. The LBD is available annually and covers all U.S. business establishments with at least one paid employee.⁴ The LBD contains longitudinal establishment identifiers along with data on employment, payroll, industry sector, location, and corporate affiliation. I use the longitudinal establishment identifiers to construct longitudinal linkages between the CMF and ASM.

Given that the LBD covers the entire U.S. economy, it also contains information about non-manufacturing establishments of companies that have plants in either the CMF or the ASM. I use this information to construct firm-level variables, such as the total number of employees and the number of establishments per firm. For my analysis, the most important firm-level variable is the ZIP code of the company’s headquarters. At the firm level, the Census Bureau distinguishes between single- and multi-unit firms. Single-unit firms consist of a single establishment, which means headquarters and the plant are located in the same unit. Multi-unit firms consist of two or more LBD establishments, with one establishment being the company’s headquarters.

To determine the location of headquarters, I supplement the LBD with data from two other data sets provided by the Census Bureau: the Auxiliary Establishment Survey (AES) and the

4. An establishment is a “single physical location where business is conducted” (Jarmin and Miranda 2003, 15). Establishments are the economic units used in the Census data sets.

Standard Statistical Establishment List (SSEL). The AES contains information on non-production (“auxiliary”) establishments, including information on headquarters. The SSEL contains the names and addresses of all U.S. business establishments. Appendix A outlines the procedure used to obtain the location of headquarters from these data sets. The main source of information about headquarters, the AES, is available every five years between 1977 and 2002. To fill in the missing years, I use the information from the latest available AES. Given that the Census years are deterministic, this measurement error is unlikely to introduce any bias. It merely introduces noise into the regression, which makes it harder for me to find any significant results.

My sample covers the period from 1977 to 2005. (The first available AES year is 1977; 2005 is the last available ASM year.) To be included in my sample, I require that a plant have a minimum of two consecutive years of data. Following common practice in the literature (e.g., Foster, Haltiwanger, and Syverson 2008), I exclude plants whose information is imputed from administrative records rather than directly collected. I also exclude plant-year observations for which employment is either zero or missing. Finally, to ensure that the physical distance between plants and headquarters is comparable across years, I exclude firms that change the location of their headquarters during the sample period. The results are virtually identical if these firms are included.

The foregoing selection criteria leave me with 1,332,824 plant-year observations. In my regressions, I use a 10-year window around the treatment date, meaning treated plants are included from 5 years before the treatment to 5 years after the treatment. Using a 10-year treatment window reduces my sample only slightly, leaving me with a final sample of 1,291,280 plant-year observations. That being said, the length of the treatment window is immaterial for my results. All results are similar if I use a different treatment window or no treatment window at all, meaning all plant-year observations of treated plants are included either before or after the treatment.

Airline Data. The data on airline routes are obtained from the T-100 Domestic Segment Database (for the period 1990 to 2005) 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).⁵ All airlines operating flights in the United States are required by law to file Form 41 with the DOT and are subject to fines for misreporting. Strictly speaking, the T-100 and ER-586 are not samples: They include *all* flights that have taken place between any two airports in the United States.

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.

II.B. Empirical Methodology

The introduction of new airline routes that reduce the travel time between headquarters and plants makes it easier for headquarters to monitor and acquire information about plants. To examine the effects on plant-level investment and productivity, I use a difference-in-differences approach. I estimate:

$$(1) \quad y_{ijlt} = \alpha_i + \alpha_t + \beta \times \text{treatment}_{it} + \gamma' \mathbf{X}_{ijlt} + \varepsilon_{ijlt},$$

where i indexes plants, j indexes firms, l indexes plant location, t indexes years, y_{ijlt} is the dependent variable of interest (plant investment or productivity), α_i and α_t are plant and year fixed effects, treatment is a dummy variable that equals 1 if a new airline route that reduces the travel time between plant i and its headquarters has been introduced by time t , \mathbf{X} is a vector of control variables, and ε is the error term. Location is defined at the MSA level.⁶ The main coefficient of interest is β , which measures the effect of the introduction of new airline routes.

5. The T-100 Domestic Segment Database is provided by the Bureau of Transportation Statistics. The annual files of the ER-586 Service Segment Data are maintained in the form of magnetic tapes at the U.S. National Archives and Records Administration (NARA). I obtained a copy of these tapes from NARA.

6. As defined by the Office of Management and Budget, an MSA consists of a core area that contains a substantial population nucleus together with adjacent communities that have a high degree of social and economic integration with that core. MSAs include one or more counties, and some MSAs contain counties from several states. For instance, the New York MSA includes counties from four states: New York, New Jersey, Connecticut, and Pennsylvania. Because MSAs represent economically integrated areas, they are likely to be affected by the same local shocks. By definition, the MSA classification is only available for urban areas. For rural areas, I consider the rural part of each state as a separate region. There are 366 MSAs in the United States and 50 rural areas based on state boundaries.

If the relationship between plants and headquarters is governed by symmetric information and no agency problems, then the introduction of new airline routes might not matter. In all other cases, it might matter. For instance, headquarters may invest more in plants that are easier to monitor and are less likely to have private information.⁷ Likewise, better monitoring may improve plant managers' incentives, and learning about a plant may allow headquarters to improve the plant's productivity. On the other hand, if headquarters becomes "too well informed" or "monitors too much," this may impair plant managers' incentives to create new investment opportunities (Aghion and Tirole 1997) or work hard (Cr mer 1995).

My identification strategy can be illustrated with a simple example. Suppose a company headquartered in Boston has a plant located in Memphis. In 1985, no direct flight was offered between Boston Logan International Airport (BOS) and Memphis International Airport (MEM). The fastest way to connect both airports was an indirect flight operated by Delta Airlines with a stopover in Atlanta. In 1986, Northwest Airlines opened a new hub in MEM. As part of this expansion, Northwest started operating direct flights between BOS and MEM as of October 1986. The introduction of this new airline route reduced the travel time between BOS and MEM and is consequently coded as a "treatment" of the Memphis plant in 1986.

To measure the effect of this treatment on, for example, plant-level investment, one could simply compare investment at the Memphis plant before and after 1986. However, other events in 1986 might have also affected investment at the Memphis plant. For instance, there might have been a nationwide surge in investment due to favorable economic conditions or low interest rates. To account for this possibility, I include a control group that consists of all plants that have not (yet) been treated. Due to the staggered nature of the introduction of new airline routes,

(The District of Columbia has no rural area.) For expositional simplicity, I refer to these 416 geographical units as MSAs.

7. A standard result in the capital budgeting literature with asymmetric information is that there is likely to be underinvestment under the optimal mechanism (e.g., Harris and Raviv 1996; Malenko 2011). See also Seru (forthcoming), who provides empirical evidence consistent with the idea that headquarters is less likely to invest in projects that rely on division managers' private information. Likewise, moral hazard, which can be alleviated through monitoring, typically leads to underinvestment in equilibrium (e.g., Tirole 2006, chapters 3 and 4).

this implies a plant remains in the control group until it is treated (which, for some plants, may be never). I then compare the difference in investment at the Memphis plant before and after 1986 with the difference in investment at the control plants before and after 1986. The difference between the two differences is the estimated effect of the introduction of the new airline route between BOS and MEM on investment at the Memphis plant.

Airlines' decisions to introduce new routes depend on several factors, including economic and strategic considerations as well as lobbying. As long as these factors are unrelated to plant-level investment or productivity, this is not a concern. However, if there are (omitted) factors that are driving both the introduction of new airline routes and plant-level investment or productivity, then any relationship between the two could be spurious. I now discuss how my identification strategy can account for such omitted factors at the local, firm, and plant level.

Local Shocks. To continue with the example, suppose the Memphis area is booming. As a consequence, the company headquartered in Boston may find it more attractive to increase investment at the Memphis plant. At the same time, airlines may find it more attractive to introduce new flights to Memphis (e.g., due to lobbying by local plants). Fortunately, because a treatment is uniquely defined by *two* (airport) locations—the plant's and headquarters' home airports—I can control for such local shocks, thus separating out the effects of the new airline routes from the effects of contemporaneous local shocks.

Suppose, for instance, that another plant, also located in Memphis, has its headquarters in Chicago. (The travel time between Chicago and Memphis was not affected by the introduction of new airline routes during 1985 and 1986.) If investment at this other Memphis plant also increases in 1986, then an increase in investment at the first Memphis plant (with headquarters in Boston) might not be due to the newly introduced airline route between MEM and BOS but due to a contemporaneous local shock in the Memphis area. In principle, I could control for such local shocks by including a full set of MSA fixed effects interacted with year fixed effects. However, doing so would require the inclusion of $416 \text{ MSAs} \times 29 \text{ years} = 12,064$ additional fixed effects. Unfortunately, the computing resources at the Census Research Data Center were insufficient to handle this task. I therefore

adopt the methodology used in Bertrand and Mullainathan (2003) and account for local shocks by including MSA-year controls, which are computed as the mean of the dependent variable (e.g., plant-level investment) in the plant's MSA in a given year, excluding the plant itself.⁸

An alternative way to account for local shocks is to focus only on new airline routes whose introduction is unlikely to be driven by such shocks. Specifically, in a subset of cases, a new indirect flight replaces a previously optimal indirect flight, but the last leg of the flight—that is, the leg involving the plant's home airport—remains unchanged. For instance, suppose a (different) company headquartered in Boston has a plant in Little Rock. In 1985, the fastest way to connect BOS and Little Rock National Airport (LIT) was an indirect flight with stopovers in Atlanta (ATL) and MEM. In 1986, Northwest Airlines started operating direct flights between BOS and MEM, with the effect that the previously optimal indirect flight BOS-ATL-MEM-LIT is replaced with a new, faster indirect flight BOS-MEM-LIT. Importantly, the last leg of the flight—that between MEM and LIT—remains unchanged. Arguably, it is unlikely that a local shock in the Little Rock area would be responsible for the introduction of a new airline route between BOS and MEM. As I show in robustness checks, I obtain similar results when I consider only new airline routes where the last leg of the flight remains unchanged.

Firm-Level Shocks. I am also able to control for firm-level shocks. For instance, suppose the company headquartered in Boston has another plant in Queens in New York City. (The travel time between Queens and Memphis was not affected by the introduction of new airline routes during 1985 and 1986.) If investment at the Queens plant also increases in 1986, then an increase in investment at the Memphis plant might not be due to the newly introduced airline route between MEM and BOS but to a contemporaneous shock at the firm level. Analogous to the construction of the MSA-year controls, I can account for firm-level shocks by including firm-year controls, which are

8. An alternative approach would be to use a coarser definition of location, such as the nine census regions. This would only require the inclusion of $9 \text{ regions} \times 29 \text{ years} = 261$ additional fixed effects. I have done this, and all my results remain similar.

computed as the mean of the dependent variable across all of the firm's plants in a given year, excluding the plant itself.

Similar to the case of local shocks, an alternative way to account for firm-level shocks is to focus only on new airline routes whose introduction is unlikely to be driven by such shocks. Specifically, in a subset of cases, a new indirect flight replaces a previously optimal indirect flight, but the first leg of the flight—that is, the leg involving headquarters' home airport—remains unchanged. As I show in robustness checks, I obtain similar results when I focus only on this subset of new airline routes.

Plant-Specific Shocks. There is one remaining possibility: shocks that are specific to a *single* plant.⁹ Because such shocks do not affect other plants in the same region, they cannot be accounted for by the inclusion of MSA-year controls. Likewise, as the shocks do not affect other plants of the same company, they cannot be accounted for by the inclusion of firm-year controls. I address this possibility in three different ways.

First, I consider the dynamic effects of the introduction of new airline routes. If a new airline route is the (endogenous) outcome of a preexisting plant-specific shock, then I should find an "effect" of the treatment already before the new airline route is introduced. However, I find no such effect. On the contrary, I find that plant-level investment (productivity) increases only with a lag of 6–12 (12–18) months after the introduction of the new airline route, implying there is no effect either before or immediately after.

Second, it could be that a new airline route is introduced in anticipation of a future plant-specific shock. Or it could be that the shock leads first to the introduction of a new airline route and only later to an increase in plant-level investment (productivity). Both interpretations are consistent with the absence of preexisting trends. To address this issue, I show that my results are

9. For expositional simplicity, I refer to such shocks as "plant-specific shocks." Strictly speaking, this category encompasses any shock whose dimension is at the plant-headquarters level, that is, any shock that is collinear with the treatment. This includes, for example, "pair-specific" shocks to trade between the plant's and headquarters' locations, in the sense that only plants in location A with headquarters in location B—but not plants in location A in general—are affected by the shock. Naturally, a shock that is specific to a single plant is also a pair-specific shock, given that each plant is associated with a unique headquarters' location.

robust when I consider only new airline routes that are the outcome of a merger between two airlines or the opening of a new hub. Arguably, it is less likely that a shock to a single plant would be responsible for an airline merger or a hub opening.

Third, I show that my results are robust when I consider only indirect flights where either the last or first leg of the flight remains unchanged. This addresses not only the possibility of local and firm-level shocks but also the possibility of shocks that are specific to a single plant.

Finally, I should mention that I obtain similar results when I consider only small plants or plants of small firms. Arguably, it is less likely that small plants or firms can successfully lobby for the introduction of a new airline route. By the same token, airlines are less likely to respond to shocks affecting small plants or plants of small firms. I should also point out that there are 10,533 plants overall in my sample that experience a reduction in travel time due to the introduction of new airline routes. Even if it were true that, in some cases, the new airline route was the outcome of lobbying by individual plants (or firms), this would still imply that the treatment is exogenous for the remaining 10,000+ plants.

Miscellaneous Methodological Issues. In addition to accounting for the possibility of local shocks, firm-level shocks, and plant-specific shocks, my empirical design can address several other concerns.

- (1) The time variation in travel time used to construct the treatment dummy comes entirely from the introduction of new airline routes. In reality, travel time can also vary for other reasons, such as the introduction of new roads, changes in speed limits, and the expansion of railroad networks. Unfortunately, lack of comprehensive data makes it difficult to account for these sources of travel time variation. Nevertheless, their omission is unlikely to affect my results. First, I show in robustness checks that my results are only significant for large reductions in travel time (at least two hours round-trip), which almost always come from long-distance trips where air travel is the optimal means of transportation. Second, plants whose travel

time to headquarters is reduced through the expansion of roads and railroad networks are part of the control group. Thus, to the extent that these sources of travel time reduction lead to an increase in plant-level investment (productivity), their omission would imply that my results *understate* the true effects of travel time reductions.¹⁰

- (2) I do not consider the termination of existing airline routes, only the introduction of new airline routes. Terminations are much less frequent than introductions. Moreover, as routes that are discontinued are mostly minor regional routes, the resulting increase in travel time is generally modest. That being said, I show in robustness checks that my results are unchanged if I account for the termination of existing airline routes. Precisely, I augment the specification in equation (1) by adding a second treatment dummy that equals 1 whenever the termination of an existing airline route leads to an increase in travel time between a plant and its headquarters. Including this second treatment dummy has no effect on the main treatment dummy (see Table VII later).
- (3) Some companies may own private jets. However, if companies use private jets to fly to plants, then the introduction of new airline routes should not matter. Although this is unlikely to introduce any systematic bias, it introduces noise into the regressions, making it only harder for me to find any significant results.
- (4) My sample spans 29 years of data (from 1977 to 2005). In my regressions, I use a 10-year treatment window that begins 5 years before the treatment and ends 5

10. I should note that large reductions in travel time through the expansion of roads and railroad networks are less likely during my sample period, given that most of today's road and railroad infrastructure was already in place before the beginning of my sample in 1977. Most of the railroad network was built prior to World War I. The latest major extension of the road network was the completion of the Interstate Highway System. Construction began in 1956 after the enactment of the National Interstate and Defense Highways Act. By 1975, the system was mostly complete (Michaels 2008). In contrast, the airline industry was deregulated early during my sample period (Airline Deregulation Act of 1978), which triggered an expansion of airline routes in the following decades. Hence, most of the time series variation in travel time during my sample period is due to changes in airline routes, not due to the expansion of roads and railroad networks.

years after the treatment. However, my results are similar if I use different treatment windows (6, 8, 12, 14 years) or no treatment window at all, meaning all plant-year observations of treated plants are included either before or after the treatment.

- (5) An important concern—especially with regard to difference-in-differences estimations—is that serial correlation of the error term can lead to understated standard errors. In my regressions, I cluster standard errors at the MSA level. This clustering not only accounts for the presence of serial correlation within the same plant, it also accounts for any arbitrary correlation of the error terms across plants in the same MSA in any given year as well as over time. My results are similar if I cluster standard errors at the firm level or at both the MSA and firm level. I also obtain similar results if I collapse the data into two periods, before and after the introduction of a new airline route, using the residual aggregation method described in Bertrand, Duflo, and Mullainathan (2004).

II.C. Definition of Variables

Measuring Investment. Investment is total capital expenditures divided by capital stock. Both the numerator and denominator are expressed in 1997 dollars.¹¹ Investment is industry-adjusted by subtracting the industry median in a given three-digit SIC industry and year.¹² To mitigate the effect of outliers, I winsorize investment at the 2.5th and 97.5th percentiles of its empirical distribution.

Measuring Productivity. My main measure of plant productivity is total factor productivity (TFP). TFP is the difference between actual and predicted output. Predicted output is the

11. Capital expenditures are deflated by the four-digit SIC investment deflator from the NBER-CES Manufacturing Industry Database. Appendix B describes how real capital stock is constructed.

12. Instead of industry-adjusting investment, I could alternatively include industry-year controls, which are computed analogously to the MSA- and firm-year controls. My results would be unchanged.

amount of output a plant is expected to produce for given levels of inputs. To compute predicted output, I use a log-linear Cobb-Douglas production function (e.g., Lichtenberg 1992; Schoar 2002; Bertrand and Mullainathan 2003; Syverson 2004; Foster, Haltiwanger, and Syverson 2008). Specifically, TFP of plant i in year t is the estimated residual from the regression

$$(2) \quad y_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \varepsilon_{it},$$

where y is the logarithm of output and k , l , and m are the logarithms of capital, labor, and material inputs, respectively. To allow for different factor intensities across industries and over time, I estimate equation (2) separately for each industry and year. Accordingly, TFP can be interpreted as the relative productivity of a plant within its industry. Industries are classified using three-digit SIC codes. (The results are qualitatively similar if I use two- or four-digit SIC codes).¹³ To match the variables of the production function as closely as possible, I use data from the longitudinal linkage of the CMF and ASM. Appendix B describes how these variables are constructed and how inflation and depreciation are accounted for.

In my main analysis, I estimate equation (2) by ordinary least squares (OLS). Though this approach is common in the literature (see Syverson 2011 for a survey), it is not uncontroversial. Research in industrial organization has argued that two econometric issues arise when production functions are estimated by OLS (see Akerberg et al. 2007 for a review). To illustrate these issues, it is helpful to decompose the error term in equation (2) into two components: $\varepsilon_{it} = \omega_{it} + \eta_{it}$. Although both components are unobservable to the econometrician, only η_{it} is unobservable to the plant. The other component, ω_{it} , represents productivity shocks that are observed or predictable by the plant at the time when it makes its input decisions. Intuitively, ω_{it} may represent

13. SIC codes were the basis for all Census Bureau publications until 1996. In 1997, the Census Bureau switched to the North American Industry Classification System (NAICS). SIC codes were not discontinued until the 2002 census, however. From 2002 to 2005, SIC codes are obtained as follows. For plants “born” before 2002, I use the latest available SIC code. For plants born between 2002 and 2005, I convert NAICS codes into SIC codes using the concordance table of the Census Bureau. This concordance is not always one-to-one, however. Whenever a NAICS code corresponds to multiple SIC codes, I use the SIC code with the largest shipment share within the NAICS industry. Shipment shares are obtained from the 1997 CMF, which reports both NAICS and SIC codes.

variables such as the expected downtime due to machine breakdowns or temporary productivity losses due to the integration of newly hired workers. A classic endogeneity problem arises now because the plant's optimal choices of inputs k_{it} , l_{it} , and m_{it} will generally be correlated with the observed or predictable productivity shock ω_{it} . As a result, OLS estimates of the coefficients in equation (2) may be biased and inconsistent. This endogeneity problem is often referred to as a "simultaneity problem."

The second endogeneity issue, the "selection problem," arises when a plant whose observed or predictable productivity shock ω_{it} is below a certain threshold is shut down. Since plants have knowledge of ω_{it} prior to the shutdown decision, surviving plants will have ω_{it} drawn from a selected sample. The selection criteria may depend on the production inputs. For instance, plants with larger capital stock may afford to survive longer at lower productivity levels, inducing a negative correlation between ω_{it} and k_{it} in the sample of surviving plants. This correlation, in turn, may render the OLS estimates biased and inconsistent.

A variety of techniques have been suggested to address the simultaneity and selection problems. In robustness checks, I employ the structural techniques of Olley and Pakes (OP; 1996) and Levinsohn and Petrin (LP; 2003).¹⁴ OP and LP address the simultaneity problem by using investment and intermediate inputs, respectively, to proxy for the productivity shock ω_{it} . The selection problem is addressed by estimating plant survival propensity scores. Regardless of which method I use, I find that my results are similar to those obtained by estimating TFP by OLS (see Table A.5 in the Online Appendix).

TFP measures rely on structural assumptions (e.g., Cobb-Douglas production function). In robustness checks, I use two alternative measures of plant-level productivity that are free of such assumptions: operating margin (OM) and return on capital (ROC). OM is shipments minus labor and material costs, all divided by shipments.¹⁵ ROC is defined analogously, except that the denominator is capital stock (instead of shipments).

14. A description of how OP's and LP's techniques can be implemented using plant-level data is available from the author on request.

15. All dollar values are expressed in 1997 dollars. Deflators for shipments and material costs are available at the four-digit SIC level from the NBER-CES Manufacturing Industry Database. Deflators for labor costs are available at the two-digit SIC level from the Bureau of Economic Analysis.

OM and ROC are industry-adjusted by subtracting the industry median in a given three-digit SIC industry and year. Regardless of which measure I use, I find that my results are similar to my baseline results (see Table A.5 in the Online Appendix).

All productivity measures are subject to extreme values. To avoid outliers driving my results, I winsorize all productivity measures at the 2.5th and 97.5th percentiles of their respective empirical distributions.

Measuring Travel Time Reductions. The itinerary between headquarters and plants is constructed to reflect as closely as possible the decision making of managers. I assume that managers make optimal decisions. Accordingly, they choose the route and means of transportation (e.g., car, plane) that minimizes the travel time between headquarters and plants.

To identify the location of headquarters and plants, I use five-digit ZIP codes from the LBD. (Precisely, I use the latitude and longitude corresponding to the centroid of the area spanned by the ZIP code.) The travel time between any two ZIP codes is computed as follows. Using MS Mappoint, I first compute the travel time by car (in minutes) between the two ZIP codes. This travel time is used as a benchmark and is compared to the travel time by air based on the fastest airline route. Whenever traveling by car is faster, air transportation is ruled out by optimality, and the relevant travel time is the driving time by car.

To determine the fastest airline route between any two ZIP codes, I use the itinerary information from the T-100 and ER-586 data. The fastest airline route minimizes the total travel time between the plant and headquarters. The total travel time consists of three components: (1) the travel time by car between headquarters and the origin airport, (2) the duration of the flight, including the time spent at airports and, for indirect flights, the layover time, and (3) the travel time by car between the destination airport and the plant. The travel time by car to and from airports is obtained from MS Mappoint. Flight duration per segment is obtained from the T-100 and ER-586 data, which include the average ramp-to-ramp time of all flights performed between any two airports in the United States. The only unobservable quantities are the time spent at airports and the layover time. I assume that one hour is spent at the origin and destination airports combined and that each layover takes one hour.

Although these assumptions reflect what I believe are sensible estimates, none of my results depend on them. I obtain virtually identical results when making different assumptions.¹⁶

I sometimes refer to the physical distance between headquarters and plants. The physical distance in miles (“mileage”) is computed using the great-circle distance formula used in physics and navigation. The great-circle distance is the shortest distance between any two points on the surface of a sphere and is obtained from the formula

$$r \times \arcsin(\sin \lambda_P \sin \lambda_{HQ} + \cos \lambda_P \cos \lambda_{HQ} \cos[\phi_P - \phi_{HQ}]),$$

where λ_P (λ_{HQ}) and ϕ_P (ϕ_{HQ}) are the latitude and longitude, respectively, of the ZIP code of the plant (headquarters), and where r is the approximate radius of the Earth (3,959 miles).

II.D. Summary Statistics

Table I provides summary statistics for all 1,291,280 plant-year observations (column (1)) and separately for plants that are treated during the sample period (column (2)) and plants that are never treated during the sample period (column (3)). For each plant characteristic, the table reports the mean and standard deviation (in parentheses).¹⁷ All dollar values are expressed in 1997 dollars.

As shown, the group of eventually treated plants accounts for a relatively small fraction of the total plant-year observations. This is not a concern, however. Reliable identification of the treatment dummy requires only that this group be sufficiently large in absolute terms. A sample of 70,467 plant-year observations is a sufficiently large sample. The summary statistics also show that eventually treated plants are larger and are located farther away from headquarters. These differences make sense. To be treated, a plant needs to be sufficiently far away from headquarters, such that air travel is the optimal means of transportation. Besides,

16. To obtain an estimate of the average layover time, I randomly selected 100 indirect flights from the most recent year of my sample and used the airlines' current websites to obtain estimates of the layover time. The average layover time based on these calculations is approximately one hour. The time spent at the origin and destination airports is immaterial as it cancels out when comparing old and new flights.

17. Due to the Census Bureau's disclosure policy, I cannot report median or other quantile values.

TABLE I
SUMMARY STATISTICS: PLANTS

	(1)	(2)	(3)
	All plants	Eventually new airline route	No new airline route
Total value of shipments	50,196 (360,930)	75,752 (222,685)	48,721 (367,270)
Capital stock	20,710 (106,473)	33,615 (118,024)	19,965 (105,719)
Employees	213 (568)	300 (638)	208 (564)
Distance to headquarters (miles)	312 (563)	854 (616)	281 (544)
Travel time (minutes)	126 (170)	362 (135)	113 (162)
Number of observations	1,291,280	70,467	1,220,813

Notes. “All plants” refers to all plants in the sample. “Eventually new airline route” refers to plants that are treated during the sample period, that is, plants whose travel time to headquarters is reduced through the introduction of a new airline route. “No new airline route” refers to plants that are not treated during the sample period. Total value of shipments and capital stock are expressed in 1997 dollars (in 1,000s) using four-digit SIC deflators from the NBER-CES Manufacturing Industry Database. Capital stock is constructed using the perpetual inventory method described in Appendix B. Employees is the number of employees of the plant. Distance to headquarters is the great-circle distance between the plant’s ZIP code and the ZIP code of headquarters (in miles). Travel time is the total travel time based on the fastest route and means of transportation (car or plane) between the plant’s ZIP code and the ZIP code of headquarters (in minutes). All figures are sample means based on unadjusted (i.e., nonwinsorized) distributions. Standard deviations are in parentheses. The sample period is from 1977 to 2005.

plants that are located farther away from headquarters typically belong to larger companies that own larger plants. In robustness checks, I show that my results are similar if I restrict the sample to the 70,467 plant-year observations of eventually treated plants (see Table VII later).¹⁸ Also, the difference between eventually treated plants and nontreated plants comes largely from the fact that the latter include single-unit firms, that is, firms with a single plant. Naturally, these plants are relatively small. Importantly, they cannot be possibly affected by the introduction of new airline routes—as headquarters and the plant are located in the same unit—which implies they are in the control group. In robustness checks, I show that my results are virtually

18. Due to the staggered nature of the introduction of new airline routes, eventually treated plants are first in the control group and only later—when they are treated—in the treatment group. Also, I control for plant size and age in all my regressions, and I obtain identical results if I allow time shocks to differentially affect plants of different size by interacting plant size with a full set of year dummies (see columns (3) and (6) of Table III).

unchanged if I exclude single-unit firms from the sample (see Table VII later).

The 70,467 plant-year observations in column (2) of Table I correspond to 10,533 treated plants.¹⁹ In Table II, I provide auxiliary information about the nature of the treatments. New airline routes can be classified into four categories: (1) “Direct to Direct”: a new direct flight using a different route replaces a previously optimal direct flight, for example, the new flight involves an airport that is closer to either headquarters or the plant; (2) “Indirect to Indirect”: a new indirect flight using a different route replaces a previously optimal indirect flight, for example, the new indirect flight has only one stopover, while the previously optimal indirect flight had two stopovers; (3) “Indirect to Direct”: a new direct flight replaces a previously optimal indirect flight, for example, as in the Boston-Memphis example; (4) “Road to Flight”: a new direct or indirect flight replaces car travel as the previously optimal means of transportation.²⁰

For all treated plants (column (1)) and separately also for each of the above four categories (columns (2)–(5)), Table II reports the average distance (in miles) between headquarters and plants, the average travel time before and after the introduction of the new airline routes, and the average reduction in travel time, both in absolute and relative terms. As column (1) shows, the average travel time reduction across all treated plants is 1 hour, 43 minutes for a one-way trip, which amounts to a travel time reduction of 25%. The breakdown in columns (2) to (5) shows that the category “Indirect to Indirect” accounts for the largest reduction in travel time (2 hours, 26 minutes), followed by the category “Indirect to Direct” (2 hours, 7 minutes) and the category “Direct to Direct” (1 hour, 12 minutes). Also, as one would expect, larger travel time reductions are associated with

19. Thus, on average, I have about seven years of data for each treated plant. I have verified that my results are robust if I only include plants for which I have data for the entire 10-year treatment window.

20. To give an example of a “Direct to Direct” treatment, suppose a firm headquartered in Atlanta has a plant in Naples, FL. In 1982, the fastest way to travel from Atlanta to Naples was to take a direct flight from ATL to FLL (Fort Lauderdale International Airport) and then to drive to Naples. In 1983, Delta Airlines started operating direct flights between ATL and RSW (Southwest Florida International Airport), which is located right next to Naples. Thus, in this case, a previously optimal direct flight (ATL to FLL) is replaced by a new, and faster, direct flight (ATL to RSW).

TABLE II
SUMMARY STATISTICS: TRAVEL TIME REDUCTIONS

	(1)	(2)	(3)	(4)	(5)
	All	Indirect to indirect	Indirect to direct	Direct to direct	Road to flight
Distance to headquarters (miles)	854	1,211	942	726	191
Travel time before (minutes)	417	566	466	338	253
Travel time after (minutes)	314	420	339	266	206
Δ travel time (minutes)	-103	-146	-127	-72	-47
Δ travel time (%)	-25	-26	-27	-21	-19
Number of observations	10,533	1,911	3,469	4,544	609

Notes. “All” refers to all treated plants, that is, plants whose travel time to headquarters is reduced through the introduction of a new airline route during the sample period. “Indirect to indirect” refers to the subset of treated plants for which a new indirect flight using a different route replaces a previously optimal indirect flight. “Indirect to direct” refers to the subset of treated plants for which a new direct flight replaces a previously optimal indirect flight. “Direct to direct” refers to the subset of treated plants for which a new direct flight using a different route replaces a previously optimal direct flight. “Road to flight” refers to the subset of treated plants for which a new direct or indirect flight replaces car travel as the previously optimal means of transportation. Distance to headquarters is the great-circle distance between the plant’s ZIP code and the ZIP code of headquarters (in miles). Travel time before is the total travel time between the plant’s ZIP code and the ZIP code of headquarters based on the fastest route and means of transportation (car or plane) before the introduction of the new airline route (in minutes). Travel time after is defined accordingly. Δ travel time is the difference between travel time after and travel time before, expressed either in minutes or as a percentage of travel time before. All figures are sample means across all plants. The sample period is from 1977 to 2005.

longer physical distances. Finally, the “Road to Flight” category applies only to a small subset of treated plants (609 plants) whose location is relatively close to headquarters (191 miles), which explains why for these plants travel by car was previously the optimal means of transportation. Not surprisingly, the average travel time reduction is rather small for this category (47 minutes).

III. RESULTS

III.A. Main Results

Table III contains the main results. All regressions include plant and year fixed effects. Column (1) shows the effect of the introduction of new airline routes on plant-level investment. Investment is defined as capital expenditures divided by capital stock and is industry-adjusted at the three-digit SIC level. As is shown, the coefficient on the treatment dummy is 0.008, which implies that investment increases by 0.8 percentage points on

TABLE III
THE EFFECT OF NEW AIRLINE ROUTES ON PLANT INVESTMENT AND PRODUCTIVITY

Dependent variable	Investment			TFP		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.008*** (0.001)	0.009*** (0.001)	0.010*** (0.001)	0.014*** (0.003)	0.013*** (0.003)	0.013*** (0.003)
MSA-year		0.153*** (0.022)	0.148*** (0.022)		0.080*** (0.012)	0.080*** (0.012)
Firm-year		0.205*** (0.006)	0.205*** (0.006)		0.186*** (0.005)	0.186*** (0.005)
Age		-0.060*** (0.002)	-0.061*** (0.002)		0.015*** (0.002)	0.018*** (0.003)
Size		0.029*** (0.001)			0.012*** (0.002)	
Plant fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Size × year fixed effects	No	No	Yes	No	No	Yes
R-squared	0.39	0.41	0.41	0.60	0.61	0.61
Number of observations	1,291,280	1,291,280	1,291,280	1,291,280	1,291,280	1,291,280

Notes. Investment is the ratio of capital expenditures to capital stock, which is industry-adjusted by subtracting the industry median across all plants in a given three-digit SIC industry and year. Total factor productivity (TFP) is the residual from estimating a log-linear Cobb-Douglas production function by ordinary least squares for each three-digit SIC industry and year at the plant level. Investment and TFP are winsorized at the 2.5th and 97.5th percentiles of their empirical distribution. Treatment is a dummy variable that equals 1 if a new airline route that reduces the travel time between a plant and its headquarters has been introduced. MSA-year and firm-year indicate the mean of the dependent variable in the plant's MSA and firm, respectively, excluding the plant itself. Age is the natural logarithm of one plus the number of years since the plant has been in the LBD. Size is the natural logarithm of the number of employees of the plant. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

average. The coefficient is statistically highly significant. It is also economically significant. Given that the sample mean of plant-level investment is 0.10, an increase of 0.8 percentage points implies that investment increases by 8%, corresponding to an increase in capital expenditures of \$158,000 (in 1997 dollars).²¹

In column (2), I account for the possibility of local shocks (by including MSA-year controls) and shocks at the firm level (by including firm-year controls). The MSA- and firm-year controls are described in Section II.B. I also control for plant age and size. Age is the logarithm of one plus the number of years since the plant has been covered in the LBD. Size is the logarithm of the number of employees. As shown, the results are not sensitive to the inclusion of controls. If anything, the coefficient on the treatment dummy becomes slightly larger: the coefficient is now 0.009, which implies that plant-level investment increases by 9%, corresponding to an increase in capital expenditures of \$177,000 (in 1997 dollars). In column (3), I allow time shocks to differentially affect plants of different size by interacting plant size with a full set of year dummies. Again, this has little impact on my results.

In columns (4) to (6), I reestimate the specifications in columns (1) to (3) with TFP as the dependent variable. TFP is defined in Section II.C. Recall that TFP measures the relative productivity of a plant within an industry. The coefficient on the treatment dummy is between 0.013 and 0.014, which implies TFP increases by 1.3% to 1.4%, respectively, corresponding to an increase in profits (in 1997 dollars) of \$275,000 and \$296,000, respectively.²²

21. A coefficient of 0.008 implies an increase in capital expenditures equal to 0.8% of capital stock. Given that the average capital stock of treated plants is \$19.7 million—based on the winsorized distribution of capital stock, consistent with the way investment is constructed (see Section II.C)—this implies an increase in plant-level capital expenditures of $0.008 \times \$19.7 \text{ million} = \$158,000$.

22. A 1.3% increase in TFP implies, by definition, that the plant produces a 1.3% higher value of shipments with exactly the same inputs. Accordingly, the increase in plant-level profits can be approximated by multiplying the pretreatment average value of shipments of treated plants based on the winsorized distribution (consistent with the way TFP is constructed, see Section II.C) by 0.013. This pretreatment average is \$32.5 million, implying an increase in plant-level profits of $0.65 \times 0.013 \times \$32.5 \text{ million} = \$275,000$ (using a corporate tax rate of 35%). This figure is likely an overstatement, however, as it only accounts for expenses incurred by the plant but not for those incurred by headquarters (e.g., travel costs and other headquarters-related expenses that may arise following the treatment).

In the remainder of this article, I use the specification in columns (2) and (5)—which includes MSA- and firm-year controls, plant age, and plant size—as my baseline specification. All my results are similar if I exclude these four controls, if I include only a subset, or if I additionally control for firm age and firm size.

III.B. Dynamic Effects of New Airline Routes

As discussed in Section II.B, an important concern is that omitted plant-specific shocks could be driving both the introduction of new airline routes and plant-level investment or productivity. Because these shocks do not affect other plants in the same region, they cannot be accounted for by the inclusion of MSA-year controls. Likewise, as they do not affect other plants of the same firm, they cannot be accounted for by the inclusion of firm-year controls.

If a new airline route is the (endogenous) outcome of a pre-existing plant-specific shock, then I should find an “effect” of the treatment already before the new airline route is introduced. To investigate this issue, I study in detail the dynamic effects of the introduction of new airline routes. Given that annual records in the CMF and ASM are measured in calendar years, the last month of each plant-year observation is December. Because the T-100 and ER-586 segment data are at monthly frequency, this means I know precisely in which month a new airline route is introduced. Accordingly, I am able to reconstruct how many months before or after the introduction of a new airline route a given plant-year observation is recorded. For instance, consider again the Boston-Memphis example. In this example, the 1985 plant-year observation of the Memphis plant is recorded 9 months before the treatment; the 1986 plant-year observation is recorded 3 months after the treatment; the 1987 plant-year observation is recorded 15 months after the treatment, and so on.

By exploiting the detailed knowledge of the months in which new airline routes are introduced, I can replace the treatment dummy in equation (1) with a set of dummies indicating the time interval between a given plant-year observation and the treatment. I use eight dummies. The first dummy, “Treatment (−12 m, −6 m),” equals 1 if the plant-year observation is recorded between 12 and 6 months before the treatment. The other dummies are defined accordingly with respect to the

intervals (−6 m, 0 m), (0 m, 6 m), (6 m, 12 m), (12 m, 18 m), (18 m, 24 m), (24 m, 30 m), and 30 months and beyond (“30 m +”).

Table IV shows the results. In column (1), the dependent variable is plant-level investment. The main variables of interest are Treatment (−12 m, −6 m) and Treatment (−6 m, 0 m), which measure the “effect” of the new airline routes before their introduction. As is shown, the coefficients on both variables are small and insignificant, which suggests that there are no pre-existing trends in the data. Interestingly, the coefficient on Treatment (0 m, 6 m), which captures the effect of the new airline routes within the first six months after their introduction, is also insignificant. Moreover, although the effect becomes significant after six months, it remains initially small in economic terms. It is only after 12 months that the effect becomes large and highly significant. Precisely, the coefficients on Treatment (12 m, 18 m), Treatment (18 m, 24 m), and Treatment (24 m, 30 m) are between 0.013 and 0.014, which implies that plant-level investment increases by 13% to 14%. In the longer run—that is 30 months and beyond—the magnitude of the coefficient reverts to a slightly lower level. In column (2), the dependent variable is TFP. The pattern is similar to above, except that the increase in TFP occurs six months after the increase in investment. Accordingly, the effect on TFP becomes significant only after 12 months, and it becomes economically large only after 18 months.

III.C. Discussion

Money Left on the Table? My results suggest that plant-level profits increase by \$275,000 to \$296,000 on average (see Section III.A). This raises the question of whether there is or, precisely, was “money left on the table.” In particular, if the increase in profits is so large, why did the CEO (or some other senior executive) not already fly more often to the plants before? One answer is that the CEO is so time-constrained that, absent a reduction in travel time, he was simply unable to travel more often. Consistent with this argument, I document in Section V.A that the treatment effect is stronger for firms whose managers are more time-constrained.

Another way of looking at the magnitudes of my estimates is from an agency perspective.²³ Indeed, the CEO (or some other

23. I am very grateful to the editor, Andrei Shleifer, for suggesting this argument.

TABLE IV
DYNAMIC EFFECTS OF NEW AIRLINE ROUTES

Dependent variable	(1) Investment	(2) TFP
Treatment (-12 m, -6 m)	-0.000 (0.003)	-0.001 (0.005)
Treatment (-6 m, 0 m)	-0.001 (0.002)	-0.001 (0.004)
Treatment (0 m, 6 m)	0.003 (0.003)	0.001 (0.005)
Treatment (6 m, 12 m)	0.005** (0.002)	0.006 (0.005)
Treatment (12 m, 18 m)	0.013*** (0.003)	0.012** (0.005)
Treatment (18 m, 24 m)	0.014*** (0.002)	0.020*** (0.004)
Treatment (24 m, 30 m)	0.014*** (0.003)	0.020*** (0.005)
Treatment (30 m +)	0.009*** (0.002)	0.013*** (0.004)
MSA-year	0.153*** (0.022)	0.080*** (0.012)
Firm-year	0.205*** (0.006)	0.186*** (0.005)
Age	-0.060*** (0.002)	0.015*** (0.002)
Size	0.029*** (0.001)	0.012*** (0.002)
Plant fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
R-squared	0.41	0.61
Number of observations	1,291,280	1,291,280

Notes. Treatment (-12 m, -6 m) is a dummy variable that equals 1 if the plant-year observation is recorded between 6 and 12 months before the introduction of the new airline route. Treatment (-6 m, 0 m), treatment (0 m, 6 m), treatment (6 m, 12 m), treatment (12 m, 18 m), treatment (18 m, 24 m), treatment (24 m, 30 m), and treatment (30 m +) are defined analogously. All other variables are defined in Table III. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

senior executive who must travel to the plants) pays the time and inconvenience cost out of his own pocket, so to speak, but receives only a fraction of the incremental profits. This personal profit—which the CEO forgoes by not flying more often—is what needs to be compared with his personal cost of traveling. Thus, the relevant question may not be whether there is money left on the table

for the firm, but rather whether there is money left on the table for the CEO himself.

To estimate the CEO's personal profits, I merge my sample with ExecuComp, which contains detailed information on executive compensation, including CEO ownership. Alas, merging my sample with ExecuComp significantly reduces the number of observations, first, because ExecuComp includes only large, publicly traded firms and, second, because ExecuComp begins only in 1992. I compute CEO ownership in two ways. First, I compute the ratio of shares held by the CEO to the total number of shares outstanding. This measure ignores stock options and may therefore underestimate CEO ownership. To mitigate this concern, I also compute the ratio of shares and option deltas held by the CEO to the total number of shares and option deltas outstanding. Stock option deltas are computed using the methodology in Core and Guay (2002). Depending on which measure I use, I find that the average pretreatment CEO ownership is between 1.3% and 1.4%.

I then reestimate my baseline TFP regression using the ExecuComp subsample. I find that the coefficient on the treatment dummy is 0.010, which corresponds to an increase in plant-level profits of \$312,000. (This is slightly higher than the profit increase of \$275,000 to \$296,000 reported in Section III.A, because the plants of companies included in ExecuComp are on average larger than those in my main sample.) Given that CEO ownership is between 1.3% and 1.4%, this implies an increase in the CEO's personal profits of \$4,056 to \$4,368.²⁴ Although this is of course speculative, I find it plausible to think of this personal profit as being relatively small from a (wealthy) CEO's perspective—possibly small enough that it does not compensate him for his personal time and inconvenience cost.

Maybe the CEO would have been more likely to travel if the firm had a corporate jet. Though entry-level jets are available for a couple of million dollars (with top-of-the-line models costing \$30 million and more), the annual operating costs are substantial. For instance, Edgerton (2012, 2188) estimates that “annual operating

24. This is likely an overstatement of a CEO's personal profits. As plant-level profits do not include expenses incurred by headquarters, for example, overhead and travel expenses, the increase in firm-level profits is likely to be less than \$312,000, which implies that the increase in the CEO's personal profits is less than \$4,056 to \$4,368.

costs can be as high as \$5 million per jet, with \$1 million being quite typical.” Also, fractional ownership of jets, an innovation that significantly reduced the costs of using private jets, has become available only relatively late during my sample period.²⁵ Leaving cost considerations aside, there may also be other reasons firms do not use corporate jets. One such reason may be corporate governance, given that private jets may be excessively used for the pursuit of private benefits. Consistent with this picture, Edgerton (2012) finds that firms undergoing leveraged buyouts significantly reduce their fleet of corporate jets. Also, Yermack (2006) finds that the use of corporate jets is significantly related to CEOs’ long-distance golf club memberships, and that companies allowing personal aircraft use by their CEO underperform market benchmarks by about 4% per year.²⁶

Maybe it would pay the firm to hire “delegated monitors” and deploy them at the plant? Again, it is unclear whether a profit increase of \$275,000 to \$296,000 is enough to justify the costs. More important, it is unclear whether delegated monitors would be an effective substitute for the CEO or some other senior manager visiting the plant. Indeed, the very definition of soft information argues that it “cannot be credibly transmitted” (Stein 2002, 1891) and “cannot be directly verified by anyone other than the agent who produces it” (1892). Consequently, to the extent that capital budgeting decisions are (partly) based on soft information, there is no effective substitute for the CEO (or whoever makes the decision) *personally* visiting the plant.

Spillover Effects. That plant-level TFP increases does not automatically imply that the company is better off. Maybe the

25. In 1986, NetJets pioneered the concept of fractional ownership of private jets, but it has become widely popular only since the mid-1990s (National Business Aviation Administration 2004). Nevertheless, the increased use of private jets since the mid-1990s may be an alternative explanation—in addition to the one given in Section V.B—for why the treatment effect has become significantly weaker during the later years of my sample period (see Table IX for details).

26. See also “Corporate Jet Set: Leisure vs. Business” (*Wall Street Journal*, June 16, 2011), reporting that, for some companies, more than half of all flights go to resort areas such as Palm Beach, Aspen, or the Bahamas.

CEO or other managers—who can now travel more easily to the plant—do not really add any monitoring value at all, but they think they do. If so, they might draw comfort from their added involvement and invest more in the plant. This might make the plant more productive (i.e., TFP increases), especially if the plant was below the first-best level of investment before, for example, because the firm is financially constrained. However, under this alternative scenario, the rest of the firm would probably suffer, because all that is going on is an inefficient reallocation of capital based on a mistake. Likewise, if the reduction in travel time makes it easier for plant-level managers to lobby headquarters, investment and productivity at the treated plant might increase (at the expense of other plants). However, if all that is going on is an inefficient reallocation based on lobbying, overall productivity (at the firm level) should not increase. In fact, it should probably decline.

In follow-up work, I take a closer look at such spillover effects (Giroud and Mueller 2012). The research is motivated by the fact that theory on internal capital markets predicts that—provided firms are financially constrained—positive shocks to investment at one plant should lead to a decline in investment at other plants (e.g., Stein 1997). However, if the reallocation is overall efficient, firm-level productivity should rise. The main results are as follows. On average, there are no spillovers to other plants. Thus, on average, the increase in investment and productivity at the treated plant documented here is indeed a “net effect,” which translates into higher investment and productivity for the firm as a whole. That being said, when we look separately at financially constrained and unconstrained firms, we find evidence of negative spillovers among financially constrained firms—that is, investment at other plants declines. In contrast, we find no evidence of spillovers at financially unconstrained firms. However, even in situations where spillovers exist, overall firm-level productivity improves, suggesting that the resource reallocation is beneficial for the firm as a whole.

IV. ROBUSTNESS

Not finding a significant treatment effect either before or immediately after the introduction of new airline routes (see Table IV) mitigates concerns that my results are driven by

preexisting plant-specific shocks. However, it could still be that a new airline route is introduced in anticipation of a future plant-specific shock. Or it could be that the shock leads first to the introduction of a new airline route and only later to an increase in plant-level investment (productivity). Both interpretations are consistent with the absence of preexisting trends. More generally, any shock whose dimension is at the plant-headquarters level is a potential remaining concern. As such shocks are collinear with the treatment, they cannot be filtered out by the inclusion of MSA- or firm-year controls. This includes, for example, the possibility of pair-specific shocks to trade between the plant's and headquarters' locations, in the sense that only plants in location A with headquarters in location B—but not plants in location A in general—are affected by the shock. The robustness tests in this section are meant to address this remaining concern.

IV.A. Hub Openings and Airline Mergers

In this section, I show that my results are robust when I consider only new airline routes that are the outcome of a merger between two airlines or the opening of a new hub.²⁷ Arguably, it is rather unlikely that a shock to a single plant—or a pair-specific shock that affects only plants in location A with headquarters in location B—would be responsible for an airline merger or a hub opening.

Table A.1 in the Online Appendix provides a list of airline hubs that were opened during the sample period. The list is compiled from two sources: newspaper reports and airlines' annual reports. The newspaper reports are obtained from various newspaper databases (ProQuest, Factiva, and Newsbank America's Newspapers). Precisely, I ran a search for articles that contain the airline name, the airport name, and the word "hub." These articles are supplemented with information about hub openings that airlines self-report in their annual reports. As can be seen, most of the hub openings date back to the 1980s. In the years following the Airline Deregulation Act of October 1978, airlines started competing for strategic hub locations, and as a result, the 1980s witnessed a substantial number of new hub openings (Ivy 1993).

27. I thank Adair Morse for suggesting the idea to look at hub openings.

Table A.2 in the Online Appendix provides a list of airline mergers that were completed during the sample period.²⁸ The list is compiled from the same sources as the list of hub openings and is supplemented with merger information from Thompson's Securities Data Corporation (SDC) database. Although many airline mergers were completed during the sample period, I consider only mergers that account for at least one treatment in my sample. Mergers of small commuter airlines servicing few locations often do not satisfy this criterion.²⁹ As is shown in Table A.2, the pattern of airline mergers mirrors that of new hub openings. The increase in competition induced by the Airline Deregulation Act of 1978 forced many airlines to file for bankruptcy or merge with another airline. By 1990, this consolidation phase was largely completed. As a result, industry-wide concentration had increased sharply, with the nine largest airlines representing a total market share of over 90% of domestic revenue passenger miles (Goetz and Sutton 1997).

28. Airline mergers can lead to both the introduction of new airline routes and the termination of existing routes. New airline routes are typically introduced as the acquirer airline takes over the gates of the target airline at airports that were previously not serviced by the acquirer. For instance, in 1986, American Airlines acquired Air California (AirCal), a regional carrier operating in California. AirCal had previously serviced regional airports such as Sacramento, Palm Springs, and Oakland. After taking over AirCal's gates at these airports, American Airlines introduced several new airline routes, for example, from Chicago to Sacramento or from Nashville to Oakland. Route terminations are examined separately in Section IV.C.

29. I apply three additional criteria when compiling the list of airline mergers. First, I consider only mergers that resulted in an actual merger of the airlines' operations. For example, Southwest Airlines acquired Muse Air in 1985 and operated it as a fully owned subsidiary until its liquidation in 1987. Since an integration of the Muse Air routes into the Southwest network never occurred, I do not code this event as a merger. Second, the year of the merger in Table A.2 is the year in which the airlines actually merged their operations, not the year in which the merger was consummated. For example, Delta Airlines acquired Western Airlines on December 16, 1986. For a few months, Western was operated as a fully owned subsidiary. It is only several months later, on April 1, 1987, that Western's operations were merged into the Delta network. Hence, in Table A.2, the relevant merger year is 1987. Third, in two cases, the term "Acquirer Airline" refers to the name of the merged entity, not the actual acquirer. In the 1997 merger of AirTran Airways and ValueJet Airlines, the acquirer was actually ValueJet. However, the merged carrier retained the AirTran name, brand, and identity. Likewise, in the 1982 merger of Continental Airlines and Texas International Airlines, the acquirer was Texas Air (the owner of Texas International Airlines). The merged airline retained the Continental name, however.

Based on the list of hub openings and airline mergers, I divide the 10,533 treated plants into three categories: “hub treatments,” “merger treatments,” and “other treatments.” Hub treatments involve new airline routes that are introduced by airlines in the same year as they open a new hub. Merger treatments are defined analogously with respect to airline mergers.³⁰ In total, my sample includes 1,761 hub treatments and 535 merger treatments, which together account for 22% of all treated plants. This high percentage indicates that hub openings and airline mergers are significant events in the lives of airlines.

Figure I provides additional statistics. The diamond dots mark the number of newly treated plants (“treatments”) per year where the treatment involves a new airline route that is introduced by an airline that opens a new hub in year zero (event year). All years (i.e., -2 , -1 , *etc.*) are measured relative to the event year. As shown, the number of new treatments involving airlines that open a new hub is roughly constant in the years before and after the hub opening. However, in the year of the hub opening, the number of new treatments is about three times higher. I obtain a similar pattern when I consider new treatments involving airlines that merge in year zero (marked by square dots in Figure I).³¹ In either case, the spike in the event year confirms that airlines substantially expand their route networks when opening a new hub or integrating other airlines’ routes into their own operations.

In Table V, I replace the treatment dummy in equation (1) with a set of three dummies indicating whether the treatment is a hub treatment, merger treatment, or other treatment. As shown, the coefficients on all three dummies are statistically significant and economically large. The coefficient is largest for hub treatments, slightly smaller for merger treatments, and smallest for the other treatments. The differences among the coefficients are likely reflective of 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. As I show in robustness checks

30. If a merger treatment coincides with a hub treatment, I classify the event as a hub treatment. For instance, in 1987, Delta Airlines merged the operations of Western Airlines into their network and opened a new hub in Salt Lake City on the basis of the former Western hub.

31. In the years preceding the merger, the number of new treatments per year includes treatments associated with both the acquirer and target airlines.




FIGURE I

Number of New Treatments around Hub Openings and Airline Mergers

The vertical axis indicates the number of newly treated plants per year (“Treatments”) by airlines that open a new hub (“Hubs”) or merge with another airline (“Mergers”) in year 0. Years are expressed in event time with respect to year 0 (event year). In the years preceding the airline mergers, the number of newly treated plants per year includes treatments by both the acquirer and target airlines.

(see Section IV.D), larger travel time reductions are associated with stronger treatment effects. Importantly, however, that all three coefficients—especially those associated with hub and merger treatments—are large and significant mitigates concerns that my results are driven by plant-specific shocks.

IV.B. New Airline Routes with Same Last Leg or Same First Leg

Another way to account for the possibility of plant-specific shocks—or likewise, pair-specific shocks that are collinear with the treatment—is to consider only new airline routes whose introduction is unlikely to be driven by such shocks. Precisely, in a subset of cases, a new indirect flight replaces a previously optimal indirect flight, but either the last leg or the first leg of the flight—that is, the leg involving either the plant’s or headquarters’ home airport—remains unchanged. I show now that my results are robust when I consider only such new airline

TABLE V
HUB OPENINGS AND AIRLINE MERGERS

Dependent variable	(1) Investment	(2) TFP
Treatment (hub)	0.017*** (0.002)	0.019*** (0.005)
Treatment (merger)	0.014** (0.006)	0.018** (0.009)
Treatment (other)	0.008*** (0.001)	0.011*** (0.003)
MSA-year	0.153*** (0.022)	0.080*** (0.012)
Firm-year	0.205*** (0.006)	0.186*** (0.005)
Age	-0.060*** (0.002)	0.015*** (0.002)
Size	0.029*** (0.001)	0.012*** (0.002)
Plant fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
R-squared	0.41	0.61
Number of observations	1,291,280	1,291,280

Notes. Treatment (hub) and treatment (merger) are dummy variables that equal 1 if the treatment dummy equals 1 and the new airline route is introduced by an airline in the same year as it opens a new hub or merges its operations with another airline, respectively. Treatment (other) is a dummy variable that equals 1 if the treatment dummy equals 1 and the treatment is not a hub or merger treatment as defined above. All other variables are defined in Table III. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

routes.³² This not only addresses the possibility of plant or pair-specific shocks, it also addresses the possibility of local and firm-specific shocks to the extent that these are not already being fully accounted for by the inclusion of MSA- and firm-year controls.

As Table II shows, there are 10,533 treatments in total, of which 1,911 (18%) are due to a new indirect flight replacing a previously optimal indirect flight (“indirect to indirect”). In 977 of these cases, the new indirect flight operates the same last leg as the previously optimal indirect flight. For instance, a previously optimal indirect flight with two stopovers (three legs) might be replaced by a new indirect flight with only one stopover (two legs),

32. I thank Leonid Kogan and Dimitris Papanikolaou for suggesting this robustness check.

but the last leg of the flight—that is, the leg connecting the plant’s home airport—remains unchanged (see the BOS-ATL-MEM-LIT example in Section II.B). Because the last leg of the flight is unchanged, it is rather unlikely that this new airline route was triggered by a plant-specific shock or local shock in the plant’s vicinity. In the remaining 934 cases, the new indirect flight operates the same first leg as the previously optimal indirect flight. (Hence, there exists no “indirect to indirect” treatment where either both the last and first leg have changed or where both legs remain unchanged.) By the same token as above, given that the first leg of the flight—the leg connecting headquarters’ home airport—remains unchanged, it is rather unlikely that this new airline route was triggered by a shock at the firm level.

In Table VI, I replace the treatment dummy in equation (1) with a set of three dummies indicating whether the treatment is due to a new indirect flight operating the same last leg (“same last leg”), a new indirect flight operating the same first leg (“same first leg”), or any other new flight (“other”). As shown, the coefficients on all three dummies are statistically significant and economically large. The coefficient is largest for the “same first leg” and “same last leg” treatments, which is reflective of the fact that “indirect to indirect” treatments, are associated with larger travel time reductions (see Table II). Importantly, however, the fact that all three coefficients—especially those associated with the “same first leg” and “same last leg” treatments—are large and significant alleviates concerns that my findings are driven by plant- or pair-specific shocks, local shocks, or shocks at the firm level.

IV.C. Alternative Control Groups

In my baseline specification, the control group consists of all plants that have not (yet) been treated. Due to the staggered nature of the introduction of new airline routes, this includes plants that are never treated during the sample period as well as plants that will be treated at some future time. In this section, I examine the robustness of my results to using alternative control groups. The results are presented in Table VII.

Multiunit Firms. In columns (1) and (2), I exclude single-unit firms from the sample, which means the sample consists exclusively of multiunit firms. As explained in Section II.D, single-unit

TABLE VI
NEW AIRLINE ROUTES WITH SAME LAST LEG OR SAME FIRST LEG

Dependent variable	(1) Investment	(2) TFP
Treatment (same last leg)	0.012*** (0.004)	0.014* (0.008)
Treatment (same first leg)	0.013*** (0.004)	0.015** (0.008)
Treatment (other)	0.009*** (0.001)	0.012*** (0.003)
MSA-year	0.153*** (0.022)	0.080*** (0.012)
Firm-year	0.205*** (0.006)	0.186*** (0.005)
Age	-0.060*** (0.002)	0.015*** (0.002)
Size	0.029*** (0.001)	0.012*** (0.002)
Plant fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
R-squared	0.41	0.61
Number of observations	1,291,280	1,291,280

Notes. Treatment (same last leg) and treatment (same first leg) are dummy variables that equal 1 if the treatment dummy equals 1 and the new airline route operates the same last leg or the same first leg, respectively, as the previously optimal airline route. Treatment (other) is a dummy variable that equals 1 if the treatment dummy equals 1 and the new airline route operates neither the same last leg nor the same first leg as the previously optimal airline route. All other variables are defined in Table III. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

firms cannot be possibly affected by the introduction of new airline routes, which implies they are necessarily in the control group. As shown, my results are unchanged if single-unit firms are excluded.

Eventually Treated Plants. As discussed in Section II.D, eventually treated plants are larger than plants that are never treated during the sample period. In columns (3) and (4), I exclude the latter plants from the sample, which means the sample consists exclusively of eventually treated plants (see Bertrand and Mullainathan 2003 for a similar robustness check). This is possible, because—due to staggered nature of the introduction of new airline routes—eventually treated plants are first in the control group and only later, when they are treated, in the treatment

TABLE VII
ALTERNATIVE CONTROL GROUPS

Dependent variable	Multiunit firms		Eventually treated plants		Increase in travel time	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	Investment 0.009*** (0.001)	TFP 0.012*** (0.003)	Investment 0.011*** (0.002)	TFP 0.010*** (0.003)	Investment 0.009*** (0.001)	TFP 0.013*** (0.003)
Increase in travel time					-0.005** (0.002)	-0.008* (0.005)
MSA-year	0.133*** (0.019)	0.090*** (0.015)	0.084** (0.042)	0.084** (0.041)	0.153*** (0.022)	0.083*** (0.012)
Firm-year	0.207*** (0.006)	0.186*** (0.005)	0.257*** (0.016)	0.293*** (0.016)	0.205*** (0.006)	0.185*** (0.005)
Age	-0.072*** (0.002)	0.013*** (0.003)	-0.047*** (0.004)	0.038*** (0.009)	-0.060*** (0.002)	0.015*** (0.002)
Size	0.026*** (0.001)	0.020*** (0.002)	0.027*** (0.002)	0.025*** (0.006)	0.029*** (0.001)	0.012*** (0.002)
Plant fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.37	0.61	0.32	0.65	0.41	0.61
Number of observations	825,097	825,097	70,467	70,467	1,282,228	1,282,228

Notes. In columns (1) and (2), the sample is restricted to plants that belong to multiunit firms consisting of more than one establishment. In columns (3) and (4), the sample is restricted to plants that are eventually treated—that is, plants whose travel time to headquarters is reduced through the introduction of a new airline route—during the sample period. Increase in travel time is a dummy variable that equals 1 if the travel time to headquarters increases during the sample period due to the termination of an existing airline route. All other variables are defined in Table III. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

group. As is shown, my results are similar if nontreated plants are excluded.

Increases in Travel Time. In my main analysis, I consider only the introduction of new airline routes and not the termination of existing routes. Terminations are much less frequent than introductions. Moreover, as routes that are discontinued are mostly minor regional routes, the resulting increase in travel time (and thus the treatment effect) is likely to be modest. In columns (5) and (6), I add a second dummy that equals 1 whenever the termination of an existing airline route leads to an increase in travel time between a plant and its headquarters. As shown, the coefficient on this “increase in travel time” dummy is of the opposite sign as the coefficient on the main treatment dummy, which is what one might expect. Importantly, however, the coefficient on the main treatment dummy remains unchanged (see, columns (2) and (5) of Table III), which implies my results are unaffected if I additionally account for the termination of existing airline routes.

IV.D. *Miscellaneous Robustness Checks*

Small versus Large Reductions in Travel Time. In the analysis so far, any new airline route that reduces the travel time between a plant and its headquarters is coded as a treatment. Arguably, the treatment effect may be stronger for larger reductions in travel time. To see whether this is true, I interact the treatment dummy in equation (1) with a set of five dummies indicating the magnitude of the travel time reduction: ($\Delta t \leq 30$ min), ($\Delta t > 30$ min and $\Delta t \leq 1$ hr), ($\Delta t > 1$ hr and $\Delta t \leq 1$ hr 30 min), ($\Delta t > 1$ hr 30 min and $\Delta t \leq 2$ hr), and ($\Delta t > 2$ hr), where Δt is the reduction in travel time based on a *one-way* trip. As Table A.3 in the Online Appendix shows, regardless of whether I use investment or TFP as the dependent variable, the treatment effect is monotonic in the magnitude of the travel time reduction, is small and insignificant when the travel time reduction is less than one hour (one way), and is strongest when the travel time reduction exceeds two hours.

ASM Sample Weights. As described in Section II.A, the ASM includes all manufacturing plants with more than 250 employees,

whereas smaller plants are randomly sampled every five years. As a result, the average plant in the sample is larger than the average U.S. manufacturing plant. Fortunately, the ASM contains plants' sample weights—that is, the inverse of their sampling probabilities—which can be used to create a representative sample in terms of plant size (see, e.g., Greenstone, List, and Syverson 2011). Specifically, I reestimate my baseline regressions by weighting observations either by their respective ASM sample weight (for plant-weighting) or by the product of their respective ASM sample weight and their deflated shipments (for dollar-weighting). The results are presented in Table A.4 in the Online Appendix. Regardless of which weighting procedure I use, and regardless of whether I use investment or TFP as the dependent variable, the results are similar to those in Table III.

Alternative Measures of Productive Efficiency. My main measure of plant-level productivity is TFP. A potential drawback of TFP is its reliance on structural assumptions (e.g., Cobb-Douglas production function). To assess the robustness of the TFP results, I consider two margin-based measures of plant-level productivity that are free of such assumptions: return on capital (ROC) and operating margin (OM). ROC is shipments minus labor and material costs, all divided by capital stock. OM is defined similarly, except that the denominator is shipments. Both measures are industry-adjusted by subtracting the industry median in a given three-digit SIC industry and year. As is shown in columns (1) and (2) of Table A.5 in the Online Appendix, the results are similar to my baseline results. This is not surprising, given that TFP is highly correlated with both ROC (60%) and OM (50%).

Another potential drawback is that in my main analysis, TFP is estimated by OLS. As discussed in Section II.C, this approach—though common in the literature (e.g., Schoar 2002; Bertrand and Mullainathan 2003)—has been criticized on grounds that it gives rise to simultaneity and selection problems. In columns (3) and (4), I employ the structural techniques of Olley and Pakes (OP, 1996) and Levinsohn and Petrin (LP, 2003), respectively, which have been designed to address these problems. As is shown, the results are similar to my baseline results. This is again not surprising, given that the correlation between TFP estimated by OLS and TFP estimated using OP's (LP's) technique is 81% (84%).

V. HETEROGENEITY IN THE TREATMENT EFFECT

V.A. *Headquarters' Time Constraints*

Monitoring requires that managers travel to plants. The same is true for collecting “soft” information, that is, information that “cannot be credibly transmitted” (Stein 2002, 1891) and “cannot be directly verified by anyone other than the agent who produces it” (1892).³³ Given that both activities are time-consuming, I would expect the treatment effect to be stronger for plants whose headquarters is more time-constrained. To examine this hypothesis, I construct two measures of headquarters’ time constraints. The first measure is the number of managers employed at headquarters divided by the number of plants of the firm (Managers/Plants). The second measure is the number of managers employed at headquarters divided by the total distance (in miles) between headquarters and all of the firm’s plants (Managers/Total Distance). The lower the ratio of managers to plants—or the greater the average distance the managers must travel—the more time-constrained the headquarters.

A caveat is in order. When constructing these measures, I (somewhat generously) treated all of headquarters’ employees as “managers.” Although it is true that all of headquarters’ employees are white-collar employees, not all of them are managers. They may also include secretaries and clerical employees. Unfortunately, the Census data do not allow me to distinguish between managers and other white-collar employees. Nevertheless, as long as the number of other white-collar employees is roughly proportional to the number of managers—which seems like a reasonable assumption—this measurement error is unlikely to affect my results, for it merely implies that the number of managers is scaled by a constant. Bearing this caveat in mind, I sort treated plants into two categories—“high time constraints” and “low time constraints”—based on whether the measure of time constraints is above or below the median of all treated plants in the year prior to the treatment. Using pretreatment values to sort plants mitigates concerns that the categorization is affected by the treatment itself.

33. To provide further evidence that proximity to plants makes it easier for headquarters to monitor and collect (soft) information about plants, I show in Table A.6 in the Online Appendix that the treatment effect is stronger for larger firms and conglomerates, where one could plausibly argue that the loss of soft information inside the firm hierarchy is larger.

An important limitation of my analysis is that the two measures of headquarters' time constraints—although computed in the year prior to the treatment—are endogenous, for example, with respect to plant location or employment. Accordingly, the evidence presented next may not warrant a causal interpretation.³⁴

To examine whether the treatment effect is stronger for plants whose headquarters is more time-constrained, I interact the treatment dummy in equation (1) with two dummies indicating whether time constraints are low and high, respectively. The results are presented in Table VIII. In columns (1) and (3), the dependent variable is plant-level investment. As shown, the treatment effect is between two and two and a half times stronger when time constraints are high. Specifically, in column (1), where time constraints are measured by "Managers/Plants," the coefficient on the interacted treatment dummy is 0.012 when time constraints are high but only 0.006 when time constraints are low. The difference is significant at the 5% level ($p=.028$). The difference becomes even more pronounced when the measure of headquarters' time constraints takes into account the geographic dispersion of plants. In column (3), where time constraints are measured by "Managers/Total Distance," the coefficient on the interacted treatment dummy is 0.013 when time constraints are high but only 0.005 when time constraints are low. The difference is now significant at the 1% level ($p=.002$). The results when TFP is the dependent variable (columns (2) and (4)) mirror those for investment.

V.B. Innovations in Information Technology

The sample period from 1977 to 2005 witnessed major innovations in information technology. These innovations (e.g., Internet, corporate intranet, video conferencing) facilitated information flows both within and across company units, reducing the need to personally travel to plants. Accordingly, I would expect the treatment effect to be stronger in the earlier years of my sample period, where other, nonpersonal means of exchanging

34. In the absence of a concrete theory of optimal headquarters size and plant location, it is difficult to speculate what the direction or magnitude of a potential bias might be.

TABLE VIII
HEADQUARTERS' TIME CONSTRAINTS

Dependent variable	Managers/plants		Managers/total distance	
	(1) Investment	(2) TFP	(3) Investment	(4) TFP
Treatment \times high time constraints	0.012*** (0.002)	0.015*** (0.004)	0.013*** (0.002)	0.015*** (0.003)
Treatment \times low time constraints	0.006*** (0.002)	0.010** (0.004)	0.005** (0.002)	0.009* (0.005)
MSA-year	0.153*** (0.022)	0.080*** (0.012)	0.153*** (0.022)	0.080*** (0.012)
Firm-year	0.205*** (0.006)	0.186*** (0.005)	0.205*** (0.006)	0.186*** (0.005)
Age	-0.060*** (0.002)	0.015*** (0.002)	-0.060*** (0.002)	0.015*** (0.002)
Size	0.029*** (0.001)	0.012*** (0.002)	0.029*** (0.001)	0.012*** (0.002)
Plant fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
R-squared	0.41	0.61	0.41	0.61
Number of observations	1,291,280	1,291,280	1,291,280	1,291,280

Notes. In columns (1) and (2), headquarters' time constraints are measured as the number of (white-collar) employees at headquarters divided by the total number of plants of the company ("managers/plants"). In columns (3) and (4), headquarters' time constraints are measured as the number of (white-collar) employees at headquarters divided by the total distance (in miles) between headquarters and all of the company's plants ("managers/total distance"). High time constraints is a dummy variable that equals 1 if the measure of headquarters' time constraint lies above the median value across all treated plants in the year prior to the treatment. Low time constraints is defined analogously. All other variables are defined in Table III. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

information were either unavailable or less developed. To examine this hypothesis, I interact the treatment dummy in equation (1) with three dummies indicating different time periods: before 1986 (9 years), between 1986 and 1995 (10 years), and after 1995 (10 years). The results are presented in Table IX. As is shown, the treatment effect is indeed stronger in the earlier years of the sample period. When the dependent variable is plant investment, the coefficient on the interacted treatment dummy is 0.013 in the pre-1986 period, 0.010 in the period between 1986 and 1995, and 0.005 in the post-1995 period. The difference between the pre-1986 and post-1995 coefficients is significant at the 5% level ($p=.012$). The results are similar when TFP is the dependent variable.

TABLE IX
INNOVATIONS IN INFORMATION TECHNOLOGY

Dependent variable	(1) Investment	(2) TFP
Treatment × pre 1986	0.013*** (0.002)	0.019*** (0.004)
Treatment × between 1986 and 1995	0.010*** (0.002)	0.012*** (0.004)
Treatment × post 1995	0.005** (0.002)	0.009* (0.005)
MSA-year	0.153*** (0.022)	0.080*** (0.012)
Firm-year	0.205*** (0.006)	0.186*** (0.005)
Age	-0.060*** (0.002)	0.015*** (0.002)
Size	0.029*** (0.001)	0.012*** (0.002)
Plant fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
R-squared	0.41	0.61
Number of observations	1,291,280	1,291,280

Notes. Pre 1986, between 1986 and 1995, and post 1995 are dummy variables that equal 1 if the plant-year observation lies within the specified time interval. All other variables are defined in Table III. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

VI. PLANT OPENINGS AND CLOSURES

If proximity facilitates monitoring and information gathering, one might expect that it also matters at the “extensive margin,” that is, for plant openings and closures. My sample includes 14,691 plant openings and 22,281 plant closures.

The first hypothesis that I test is whether—in analogy to the hypothesis tested in Section V.A—firms whose headquarters is more time-constrained are more likely to open new plants in close proximity to headquarters. Likewise, such firms may be more likely to close plants that are located far away from headquarters. Using the subsample of plant openings and closures, respectively, I regress the logarithm of distance—either the geographical distance or the travel time between headquarters and the plant—on the logarithm of headquarters’ time constraints (using either of the two measures introduced in Section V.A), controls (plant size, firm size, plant age, and firm age, all in logarithms), and year

fixed effects.³⁵ To mitigate concerns that plant openings or closures might affect headquarters' time constraints, I lag both measures of time constraints by one year.

Importantly, although both measures of time constraints are lagged, the same caveat that applied in Section V.A also applies here: The measures are endogenous. Specifically, the concern is that there may be a persistent component in headquarters' time constraints that correlates with omitted variables that in turn affect the locations of plant openings or closures. Hence, the evidence presented below may not warrant a causal interpretation.

Table X contains the results. As is shown in columns (1) to (4), companies with more time-constrained headquarters are indeed more likely to open new plants in close proximity to headquarters. Specifically, a one standard deviation reduction in headquarters' time constraints—that is, an increase in the measures “Managers/Plants” or “Managers/Total Distance”—is associated with a significant increase in distance of 29% to 39%. Likewise, as is shown in columns (5) to (8), companies with more time-constrained headquarters are more likely to close plants that are located far away from headquarters.

The second hypothesis that I test is whether—in analogy to the hypothesis tested in Section V.B—over time, as innovations in information technology have reduced the need to travel to plants, the distance between headquarters and plants has increased. Specifically, I test whether firms have become more likely to open new plants in distant locations and to close existing plants in proximate locations. Using again the subsample of plant openings and closures, respectively, I regress the logarithm of distance—either the geographical distance or the travel time between headquarters and the plant—on dummies indicating the time periods 1986 to 1995 and post-1995, respectively, and controls (plant size, firm size, plant age, and firm age, all in logarithms). The pre-1986 period constitutes the base group.

The results are presented in Table XI. As is shown in columns (1) and (2), the distance between headquarters and newly opened plants has increased by 18.5%–19% (29.7%–30.5%) from the pre-1986 period to the 1986–1995 (post-1995) period. The results for plant closures in columns (3) and (4) mirror those for plant

35. Naturally, plant age is only included as a control in the plant closure regressions.

TABLE X
HEADQUARTERS' TIME CONSTRAINTS AND PLANT LOCATION

Sample	Plant openings			Plant closures				
	Geographical distance	Travel time	Geographical distance	Travel time	Geographical distance	Travel time		
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Managers/plants	0.148*** (0.027)		0.121*** (0.022)		-0.209*** (0.019)		-0.192*** (0.015)	
Managers/total distance		0.424*** (0.010)		0.313*** (0.009)		-0.428*** (0.009)		-0.350*** (0.008)
Firm size	0.557*** (0.014)	0.314*** (0.014)	0.439*** (0.011)	0.259*** (0.011)	0.647*** (0.014)	0.362*** (0.015)	0.543*** (0.010)	0.314*** (0.010)
Firm age	0.523*** (0.041)	0.384*** (0.036)	0.544*** (0.037)	0.442*** (0.033)	0.728*** (0.100)	0.496*** (0.089)	0.701*** (0.089)	0.510*** (0.078)
Size	-0.279*** (0.018)	-0.173*** (0.016)	-0.246*** (0.014)	-0.169*** (0.013)	-0.282*** (0.013)	-0.179*** (0.012)	-0.251*** (0.010)	-0.170*** (0.009)
Age					-0.304*** (0.034)	-0.257*** (0.031)	-0.284*** (0.025)	-0.246*** (0.023)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.27	0.39	0.28	0.38	0.38	0.48	0.40	0.49
Number of observations	14,691	14,691	14,691	14,691	22,281	22,281	22,281	22,281

Notes: The sample of plant openings (closures) in columns (1)-(4) (5)-(8) includes all sample plants that are coded as "birth" ("death") in the LBD. Geographical distance is the natural logarithm of one plus the great-circle distance between the plant's ZIP code and the ZIP code of headquarters. Travel time is the natural logarithm of one plus the total travel time based on the fastest route and means of transportation (car or plane) between the plant's ZIP code and the ZIP code of headquarters. Firm size is the natural logarithm of the total number of employees of the parent company to which the plant belongs. Firm age is the natural logarithm of one plus the number of years the parent company has been in the LBD. All other variables are defined in Table III. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

TABLE XI
 INNOVATIONS IN INFORMATION TECHNOLOGY AND PLANT LOCATION

Sample	Plant openings		Plant closures	
	(1) Geographical distance	(2) Travel time	(3) Geographical distance	(4) Travel time
Between 1986 and 1995	0.185*** (0.063)	0.190*** (0.049)	-0.126** (0.057)	-0.104** (0.047)
Post 1995	0.297*** (0.078)	0.305*** (0.063)	-0.197** (0.077)	-0.168*** (0.064)
Firm size	0.580*** (0.014)	0.461*** (0.011)	0.686*** (0.014)	0.580*** (0.010)
Firm age	0.451*** (0.035)	0.459*** (0.029)	0.338*** (0.061)	0.341*** (0.051)
Size	-0.319*** (0.017)	-0.283*** (0.014)	-0.317*** (0.013)	-0.282*** (0.010)
Age			-0.314*** (0.033)	-0.294*** (0.025)
<i>R</i> -squared	0.26	0.27	0.37	0.39
Number of observations	14,691	14,691	22,281	22,281

Notes. Between 1986 and 1995 and post 1995 are dummy variables that equal 1 if the plant is opened (columns (1)–(2)) or closed (columns (3)–(4)) within the specified time interval. All other variables are defined in Tables III and X. Standard errors are clustered at the MSA level. The sample period is from 1977 to 2005. Standard errors are in parentheses. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

openings. Overall, these results suggest that the distance between headquarters and plants has increased over time. This is consistent with Petersen and Rajan's (2002) finding that, for related reasons, the distance between small business borrowers and their lenders has increased.

VII. CONCLUSION

"Proximity breeds investment." Empirical evidence supporting this hypothesis has been found in many contexts (see Section I). However, all of this evidence comes from arm's-length transactions. In contrast, little is known about the role of proximity *within* firms. For instance, is it true that—in analogy to the findings in the mutual funds and banking literatures—headquarters is more likely to invest in plants that are located closer to headquarters? And does proximity to headquarters improve plant-level productivity? In this study, I attempt to address

these questions using plant-level data from the U.S. Census Bureau. My contribution is to provide plausibly exogenous variation in the proximity between headquarters and plants. Specifically, I notice that the main reason empirical studies are interested in (geographical) proximity is because it proxies for the ease of monitoring and acquiring information. I argue that a more direct proxy is travel time. Using the introduction of new airline routes as a source of exogenous variation in plants' proximity to headquarters, I estimate the effects on plant-level investment and productivity. I find that new airline routes that reduce the travel time between headquarters and plants lead to an increase in plant-level investment of 8% to 9% and to an increase in plants' total factor productivity of 1.3% to 1.4%, corresponding to an increase in plant-level profits of \$275,000 to \$296,000 (in 1997 dollars).

Although these magnitudes represent the average treatment effect, there is substantial heterogeneity. For instance, the treatment effect is stronger for plants whose headquarters is more time-constrained, consistent with the notion that time constraints limit the ability of headquarters to monitor and acquire information about plants. Also, I show that the results are stronger in the earlier years of the sample period, where other, non-personal means of exchanging information (e.g., Internet, corporate intranet, video conferencing) were either unavailable or less developed. Finally, I provide auxiliary evidence suggesting that monitoring and information acquisition are potentially important determinants of plant location. Specifically, I show that companies whose headquarters is more time-constrained are more likely to open plants in close proximity to headquarters and close plants that are located far away from headquarters. Similarly, over time—as innovations in information technology have reduced the need to travel to plants—firms have become more likely to open new plants in distant locations and close existing plants in proximate locations, implying the distance between headquarters and plants has increased.

APPENDIX A: LOCATION OF HEADQUARTERS

The primary source of headquarters data is the AES, which contains information on auxiliary establishments every five years from 1977 to 2002. An auxiliary is any establishment whose

principal function is to “manage, administer, service, or support the activities of the company’s other establishments” (U.S. Census Bureau 1996, 133). Auxiliary establishments include headquarters, warehouses, garages, and other facilities primarily engaged in servicing a company’s operating establishments.

To distinguish between headquarters and other auxiliary establishments, I use the selection criteria employed in Aarland et al. (2007). Specifically, in 1997 and 2002, headquarters is identified by the six-digit NAICS industry code 551114.³⁶ In prior years (1977, 1982, 1987, and 1992)—that is, before the introduction of NAICS codes by the Census Bureau—headquarters is identified as an establishment for which the joint category of management, administrative, and clerical employees dominates each of the other employment categories.

These criteria do not differentiate between a company’s main headquarters and regional or divisional administrative offices. As a result, they may yield more than one “headquarters” per company. In my manufacturing sample, 20% of the multiunit companies have multiple headquarters in the AES. To identify the main headquarters, I supplement the AES with information from the SSEL. The SSEL contains the names and addresses of all U.S. business establishments. This information typically includes a brief description of the establishment. Accordingly, I search for keywords that explicitly point to the main headquarters (such as “corporate headquarters” or “company headquarters”). This procedure identifies the main headquarters for 24% of the companies with multiple headquarters. For the remaining companies, I supplement the AES with payroll information from the LBD. The main headquarters is then identified as the headquarters with the highest payroll. The intuition behind this criterion is twofold. First, the main company headquarters is likely to be substantially larger than either regional or divisional administrative offices. Second, the main headquarters employs the CEO and most senior executives of the company, whose salaries are likely to translate into relatively higher payroll figures.

36. The NAICS Industry 551114 comprises “establishments (except government establishments) primarily engaged in administering, overseeing, and managing other establishments of the company or enterprise. These establishments normally undertake the strategic or organizational planning and decision-making role of the company or enterprise. Establishments in this industry may hold the securities of the company or enterprise” (U.S. Census Bureau 2000, Appendix B).

Not all multiunit companies have headquarters data in the AES. Because by definition, auxiliary establishments are physically separated from production facilities, the AES covers only *stand-alone* headquarters. For example, headquarters that are integrated into manufacturing plants are classified as manufacturing establishments and appear in the CMF. To determine the headquarters' location of companies without stand-alone headquarters, I apply similar criteria as above. Specifically, all LBD establishments of the company are matched to the SSEL. Whenever the name and address provided in the SSEL is not sufficient to determine the corporate headquarters, I select the establishment with the highest payroll from the LBD (or the highest white-collar payroll from the CMF if all establishments are manufacturing plants). Arguably, the latter criterion is subject to misclassification if, for example, headquarters is located in the smallest plant of the company. Fortunately, the impact of such misclassification is likely to be small. In my sample of manufacturing firms, companies without stand-alone headquarters are mainly small companies with only a few plants. These plants are typically located in the same MSA or county, which makes air travel an unlikely means of transportation between headquarters and the plants. My results are unaffected if I exclude these plants from the sample.

To assess the accuracy of the headquarters location obtained from the Census micro data, I merge my data set with Compustat using the Compustat-SSEL bridge maintained by the Census Bureau. Compustat contains firm-level information on large publicly traded U.S. companies, including the ZIP code of the company's headquarters. A drawback is that Compustat's ZIP codes are available only for the latest available year of the database and may therefore be an incorrect benchmark for companies whose headquarters has moved since the last AES year. Nevertheless, this inaccuracy will merely understate the actual match between headquarters locations from Compustat and the Census micro data. The merged sample consists of 4,045 companies corresponding to 312,774 plant-year observations. The headquarters location is the same for 84% of the companies, which account for 91% of the plant-year observations. Though this match may be considered satisfactory, I have verified that my results are similar if I restrict the sample to the publicly traded companies listed in Compustat and use the headquarters ZIP codes from Compustat instead.

APPENDIX B: VARIABLES OF THE PRODUCTION FUNCTION

This appendix describes how the variables of the production function are constructed. Unless otherwise specified, all variables are measured at the plant level and are obtained from the longitudinal linkage of the CMF and ASM.

Output is the total value of shipments plus changes in the value of inventories for finished goods and work in process, divided by the four-digit SIC shipment deflator from the NBER-CES Manufacturing Industry Database. *Material* is the sum of cost of materials and parts, cost of fuels, cost of purchased electricity, cost of resales, and cost of contract work, divided by the four-digit SIC material deflator from the NBER-CES Manufacturing Industry Database. *Labor* is measured in “production worker-equivalent hours” using the procedure described in Lichtenberg (1992). Specifically, labor is calculated as production worker hours times the ratio of total wages (including supplemental labor costs) to wages of production workers. This procedure assumes that the ratio of production to nonproduction wage rates is equal to the ratio of their marginal products.

Following Lichtenberg (1992), *capital* is calculated using the perpetual inventory method. This method requires an initial value of real capital stock. For each plant, I select the earliest available book value of capital. To account for depreciation, I multiply this value by the two-digit SIC adjustment factor from the Bureau of Economic Analysis (BEA). This adjustment factor is the ratio of industry net capital stock in current dollars to industry gross capital stock in historical dollars. The adjusted book value of capital is then divided by the four-digit SIC investment deflator from the NBER-CES Manufacturing Industry Database. If the earliest available book value of capital corresponds to the year in which the plant was “born” (as identified by the “birth” flag in the LBD), no adjustment for depreciation is needed. In this case, the book value is simply divided by the four-digit SIC investment deflator.

The initial value of real capital stock is then written forward using the recursive perpetual inventory formula

$$K_{it} = K_{it-1} \times (1 - \delta_{it}) + I_{it},$$

where i indexes plants, t indexes years, K is the value of real capital stock, δ is the two-digit SIC depreciation rate from the BEA, and I is capital expenditures divided by the four-digit SIC

investment deflator. Until the 1997 Census, all necessary variables are available separately for buildings and machinery. Accordingly, I calculate the capital stock for each asset category and add them together to obtain the final measure of capital stock. As of 1997, only aggregate capital stock variables are available. Another limitation is that in 1986 as well as after the 1987 Census, no questions are asked about assets in the ASM. However, questions about capital expenditures are asked every year. Accordingly, for plants that enter in 1986 or between Census years after 1987, I use asset information from the CMF in the following Census year and iterate the recursive perpetual inventory formula backward to construct capital stock in the preceding ASM years (e.g., if a plant enters in 1989, I use asset information from the 1992 CMF and iterate the formula backward to obtain capital stock for 1989–1991). Plants whose birth and death occur between two Census years post 1987 are dropped from the sample because capital stock cannot be constructed. For a more detailed discussion of these issues, see the data appendix in Abraham and White (2006).

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournals.org).

REFERENCES

- Aarland, Kristin, James C. Davis, J. Vernon Henderson, and Yukako Ono, "Spatial Organization of Firms: The Decision to Split Production and Administration," *RAND Journal of Economics*, 38 (2007), 480–494.
- Abraham, Arpad, and Kirk White, "The Dynamics of Plant-Level Productivity in U.S. Manufacturing," CES Working Paper 06-20, 2006.
- Ackerberg, Daniel A., C. Lanier Benkard, Steven Berry, and Ariel Pakes, "Econometric Tools for Analyzing Market Outcomes," in *Handbook of Econometrics*, James J. Heckman, and Edward E. Leamer, eds. (Amsterdam: Elsevier/North-Holland, 2007).
- Aghion, Philippe, and Jean Tirole, "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105 (1997), 1–29.
- Bertrand, Marianne, and Sendhil Mullainathan, "Enjoying the Quiet Life? Corporate Governance and Managerial Preferences," *Journal of Political Economy*, 111 (2003), 1043–1075.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119 (2004), 249–275.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience Theory of Choice under Risk," *Quarterly Journal of Economics*, 127 (2012), 1243–1285.

- Core, John, and Wayne Guay, "Estimating the Value of Employee Stock Option Portfolios and their Sensitivities to Price and Volatility," *Journal of Accounting Research*, 40 (2002), 613–630.
- Coval, Joshua D., and Tobias J. Moskowitz, "Home Bias at Home: Local Equity Preference in Domestic Portfolios," *Journal of Finance*, 54 (1999), 2045–2074.
- , "The Geography of Investment: Informed Trading and Asset Prices," *Journal of Political Economy*, 109 (2001), 811–841.
- Crémer, Jacques, "Arm's Length Relationships," *Quarterly Journal of Economics*, 110 (1995), 275–295.
- DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47 (2009), 315–372.
- Edgerton, Jesse, "Agency Problems in Public Firms: Evidence from Corporate Jets in Leveraged Buyouts," *Journal of Finance*, 67 (2012), 2187–2214.
- Foster, Lucia, John C. Haltiwanger, and Chad Syverson, "Reallocation Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98 (2008), 394–425.
- Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg, "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," *American Economic Review*, 96 (2006), 1043–1068.
- Gennaioli, Nicola, and Andrei Shleifer, "What Comes to Mind," *Quarterly Journal of Economics*, 125 (2010), 1399–1433.
- Giroud, Xavier, and Holger M. Mueller, "Capital and Labor Reallocation Inside Firms," Mimeo, Massachusetts Institute of Technology, 2012.
- Goetz, Andrew R., and Christopher J. Sutton, "The Geography of Deregulation in the U.S. Airline Industry," *Annals of the Association of American Geographers*, 87 (1997), 238–263.
- Greenstone, Michael, John A. List, and Chad Syverson, "The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing," CES Working Paper 11-03, 2011.
- Harris, Milton, and Artur Raviv, "The Capital Budgeting Process: Incentives and Information," *Journal of Finance*, 51 (1996), 1139–1174.
- Ivy, Russell L., "Variations in Hub Service in the US Domestic Air Transportation Network," *Journal of Transport Geography*, 1 (1993), 211–218.
- Jarmin, Ron S., and Javier Miranda, "The Longitudinal Business Database," CES Working Paper 02-17, 2003.
- Kroc, Ray. *Grinding it Out: The Making of McDonald's* (New York: St. Martin's Paperbacks, 1992).
- Lerner, Josh, "Venture Capitalists and the Oversight of Private Firms," *Journal of Finance*, 50 (1995), 301–318.
- Levinsohn, James, and Amil Petrin, "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies*, 70 (2003), 317–341.
- Lichtenberg, Frank R. *Corporate Takeovers and Productivity* (Cambridge, MA: MIT Press, 1992).
- Malenko, Andrey, "Optimal Design of Internal Capital Markets," Mimeo, Stanford University, 2011.
- Mian, Atif, "Distance Constraints: The Limits of Foreign Lending in Poor Economies," *Journal of Finance*, 61 (2006), 1465–1505.
- Michaels, Guy, "The Effect of Trade on the Demand for Skill — Evidence from the Interstate Highway System," *Review of Economics and Statistics*, 90 (2008), 683–701.
- National Business Aviation Administration, *NBAA Business Aviation Fact Book 2004* (Washington, DC: National Business Aviation Administration, 2004).
- Olley, Steven G., and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64 (1996), 1263–1297.
- Petersen, Mitchell A., and Raghuram G. Rajan, "Does Distance Still Matter? The Information Revolution in Small Business Lending," *Journal of Finance*, 57 (2002), 2533–2570.
- Schoar, Antoinette, "The Effect of Diversification on Firm Productivity," *Journal of Finance*, 62 (2002), 2379–2403.

- Seru, Amit, "Firm Boundaries Matter: Evidence from Conglomerates and R&D Activity," *Journal of Financial Economics* (forthcoming).
- Stein, Jeremy C., "Internal Capital Markets and the Competition for Corporate Resources," *Journal of Finance*, 52 (1997), 111–133.
- , "Information Production and Capital Allocation: Decentralized vs. Hierarchical Firms," *Journal of Finance*, 57 (2002), 1891–1921.
- Sufi, Amir, "Information Asymmetry and Financing Arrangements: Evidence from Syndicated Loans," *Journal of Finance*, 62 (2007), 629–668.
- Syverson, Chad, "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 112 (2004), 1181–1222.
- , "What Determines Productivity?" *Journal of Economic Literature*, 49 (2011), 326–365.
- Tirole, Jean, *The Theory of Corporate Finance* (Princeton, NJ: Princeton University Press, 2006).
- U.S. Census Bureau, *History of the 1992 Economic Census* (Washington, DC: U.S. Government Printing Office, 1996).
- , *1997 Economic Census: Management of Companies and Enterprises* (Washington, DC: U.S. Government Printing Office, 2000).
- Yermack, David, "Flights of Fancy: Corporate Jets, CEO Perquisites, and Inferior Shareholder Returns," *Journal of Financial Economics*, 80 (2006), 211–242.