

**INFORMATION AND INVESTMENT:
EVIDENCE FROM PLANT-LEVEL DATA**

by

**Xavier Giroud *
NYU Stern School of Business**

CES 10-38

October, 2010

The research program of the Center for Economic Studies (CES) produces a wide range of economic analyses to improve the statistical programs of the U.S. Census Bureau. Many of these analyses take the form of CES research papers. The papers have not undergone the review accorded Census Bureau publications and no endorsement should be inferred. Any opinions and conclusions expressed herein are those of the author(s) 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. Republication in whole or part must be cleared with the authors.

To obtain information about the series, see www.ces.census.gov or contact Cheryl Grim, Editor, Discussion Papers, U.S. Census Bureau, Center for Economic Studies 2K130B, 4600 Silver Hill Road, Washington, DC 20233, CES.Papers.List@census.gov.

Abstract

A reduction in travel time between headquarters and plants makes it easier for headquarters to monitor plants and gather “soft” information - i.e., information that cannot be transmitted through non-personal means. Using a differences-in-differences methodology, I find that the introduction of new airline routes that reduce the travel time between headquarters and plants leads to an increase in plant-level investment of 8% to 9%. This increase in investment is accompanied by an increase in plants’ total factor productivity of 1.3% to 1.4%. Consistent with the argument that the reduction in travel time makes it easier for headquarters to monitor plants and gather soft information, I find that my results are stronger: i) for plants whose headquarters are more time constrained; ii) for plants operating in soft-information industries; iii) during the earlier years of my sample period; iv) for plants where the information uncertainty is likely to be greater, such as smaller plants, peripheral plants operating in industries that are not the firm’s main industry, and plants operating in industries with more volatile sales or wages.

* I am indebted to my advisor, Holger Mueller, for his guidance and unconditional support. I also thank the other members of my dissertation committee, Kose John, Philipp Schnabl, and Daniel Wolfenzon for valuable comments and suggestions. Moreover, I would like to acknowledge helpful comments from Allan Collard-Wexler, Javier Miranda, and Adair Morse. The research in this paper was conducted while the author was a Special Sworn Status researcher of the U.S. Census Bureau at the New York Census Research Data Center. 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.

1 Introduction

Does higher information uncertainty reduce investment? This question has spurred a large theoretical literature, which makes ambiguous predictions about the relationship between information uncertainty and investment. One class of models assumes perfect competition, constant returns to scale, and symmetric adjustment costs of investment (e.g., Hartman, 1972; Abel, 1983). Together, these assumptions imply that an increase in information uncertainty raises the marginal product of capital, thereby increasing investment. Another class of models emphasizes the role of irreversibility in investment decisions (e.g., Arrow, 1968; Bernanke, 1983; McDonald and Siegel, 1986; Pindyck, 1988; Dixit and Pindyck, 1994). If investment is irreversible, the “option to wait” is valuable, which means the option value constitutes an opportunity cost of current investment. Since the option value increases with information uncertainty, higher information uncertainty reduces current investment.¹

Establishing the sign of the relationship between information uncertainty and investment is also difficult empirically. Earlier studies focus on aggregate data and typically find evidence consistent with models of irreversible investment (e.g., Pindyck and Solimano, 1993; Alesina and Perotti, 1996; Caballero and Pindyck, 1996). However, as Guiso and Parigi (1999) point out, there are important concerns with using aggregate data. First, theories of information uncertainty and investment have been developed at the micro-level; aggregate implications have not been fully worked out. Second, much of the information uncertainty comes from (idiosyncratic) shocks at the firm level. In the aggregate, these shocks are likely to cancel out, implying that aggregate measures of information uncertainty are unlikely to reflect them. Later studies using firm-level data find evidence consistent with models of irreversible investment. Leahy and Whited (1996) and Bloom, Bond, and Van Reenen (2007) measure information uncertainty using the variance of the firm’s daily stock returns in a sample of Compustat and U.K. manufacturing firms, respectively. Guiso and Parigi (1999) use Italian survey data. Their measure

¹Depending on the assumptions, the predictions from either class of models can be reversed, which only adds to the ambiguity of the relationship between information uncertainty and investment. For example, Caballero (1991) shows that the convexity of the marginal product of capital, which is at the core of the first class of models, can be weakened and even overturned by allowing for a sufficient degree of either decreasing returns to scale or imperfect competition. The predictions of the real options literature can be reversed as well. For example, Abel et al. (1996) show that considering the option to expand in addition to the option to wait renders the (net) effect of information uncertainty on investment ambiguous, as it now depends on the relative value of the two options.

of information uncertainty is based on managers' subjective assessment of the distribution of future demand for the firm's products, as reported by the survey.

Empirical studies using firm-level data are potentially subject to the same critique as empirical studies using "aggregate" data. This is because investment takes place at the plant level, not at the firm level. Accordingly, firm-level data are ultimately also aggregate data, and measures of information uncertainty at the firm level are unlikely to reflect (idiosyncratic) shocks at the plant level. Another important concern with existing studies is that omitted variables may influence both investment and the measure of information uncertainty, making it difficult to establish causality.

This paper is an attempt to address both of these concerns. First, I use plant-level data provided by the U.S. Census Bureau for the manufacturing sector for the period 1977 to 2005. Second, I use plausibly exogenous variation in information uncertainty. 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 U.S. The source of exogenous variation which I exploit in this paper is the introduction of new airline routes that reduce the travel time between headquarters and plants.² A reduction in travel time makes it easier for headquarters to monitor plants and gather "soft" information—i.e., information that cannot be transmitted through non-personal means.³ Using a differences-in-differences methodology, I find that the introduction of new airline routes leads to an increase in plant-level investment of 0.8 to 0.9 percentage points. Given that the sample mean of investment is 0.10 (treated and non-treated plants alike), this implies an increase in plant-level investment of 8% to 9%. The effect is stronger for larger reductions in travel time, and it is only significant for travel time reductions of at least one hour.

²Some papers use geographic distance as a proxy for information uncertainty (e.g., Lerner, 1995; Coval and Moskowitz, 1999, 2001; Petersen and Rajan, 2002; Garmaise and Moskowitz, 2004). An advantage of using travel time is that it entails plausibly exogenous variation. In contrast, geographic distance is typically the outcome of (endogenous) locational choices.

³Stein (2002) defines soft information as information that "cannot be credibly transmitted" (p. 1891) and that "cannot be directly verified by anyone other than the agent who produces it" (p. 1892). Similarly, Petersen (2004) notes that "[w]ith soft information the context under which it is collected and the collector of the information are part of the information. It is not possible to separate the two. This is why soft information is collected in person and historically the decision maker was the same person as the information collector" (pp. 7-8). In contrast, with hard information "[t]he collection method need not be personal" (p. 6).

While a reduction in travel time between headquarters and plants may help reduce information uncertainty, it may also help reduce information *asymmetry*. Moreover, in the case of monitoring, it may help improve plant-level managers' incentives, e.g., to increase plant-level productivity or develop new investment opportunities. In either case, the theoretical literature makes ambiguous predictions as to how this may affect investment. For instance, in the capital budgeting literature, private information by plant (or division) managers may lead to either over- or underinvestment by headquarters (e.g., Harris and Raviv, 1996, 1998; Bernardo, Cai and Luo, 2001, 2004). Similarly, while most agency models ascribe a positive role to monitoring, some argue that (too much) monitoring may reduce managerial initiative (e.g., Crémer, 1995; Aghion and Tirole, 1997; Burkart, Gromb, and Panunzi, 1997).

An important question is whether investment is efficient. While neoclassical models of investment generally presume efficiency, a class of models in the organizational and capital budgeting literature argues that investment decisions within firms are distorted by lobbying and rent-seeking activities (e.g., Milgrom, 1988; Meyer, Milgrom, and Roberts, 1992; Rajan, Servaes, and Zingales, 2000; Scharfstein and Stein, 2000). For instance, it could be the case that plant-level investment does not increase because information uncertainty (or asymmetry) is reduced, but rather because the introduction of new airline routes makes it easier for plant-level managers to lobby headquarters for bigger capital budgets. In this case, the increase in investment should be *inefficient*, a result that has been emphasized by the rent-seeking literature. My results do not support this prediction. On the contrary, I find that plants' total factor productivity increases by 1.3% to 1.4%. Likewise, I find that plants' return on capital and operating margin increase by 1.3 and 0.9 percentage points, respectively.

In this paper, a treatment is the introduction of a new airline route that reduces the travel time between headquarters and plants. Thus, a treatment is defined at the plant-headquarters level, meaning it is uniquely defined by *two* airport locations. This provides me with a tight identification. For instance, a concern is that the increase in investment might not be due to the introduction of new airline routes, but rather due to contemporaneous economy-wide shocks, local shocks in the plant's vicinity, or shocks at the firm level. While the differences-in-differences methodology accounts for the possibility of economy-wide shocks (via the control group), I can additionally account for local and firm-level shocks by including MSA-year and firm-year controls, respectively (see Bertrand and Mullainathan (2003) for a similar approach).

Both of these shocks are identified here, because not all local plants have their headquarters in the same city or region, and because not all plants of a company are affected by the introduction of a new airline route, respectively.

Another related concern is that the introduction of new airline routes might be the outcome of lobbying by individual plants or firms experiencing positive investment shocks (“reverse causality”). While the inclusion of MSA- and firm-year controls already mitigates this concern, I provide additional evidence suggesting that my results are unlikely to be driven by reverse causation. First, if a plant experiences a positive investment shock prior to the introduction of a new airline route, then I should find an “effect” already before the new route is introduced. However, I find that investment increases only with a lag of six to twelve months, meaning there is no effect either before the introduction of a new airline route or immediately after. 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 single plant is so powerful that it can successfully lobby for an airline merger or the opening of a new hub.

Even if we take as given the result that the introduction of new airline routes leads to an increase in plant-level investment, there might be alternative stories that have little to do with monitoring or information gathering by headquarters. For instance, suppose a company with headquarters in Boston has a plant in Memphis. In 1986, Northwest Airlines opened a new hub in Memphis and started operating direct flights between Boston and Memphis. (No direct flights between the two cities had been offered previously.) It is possible that the new direct route between Memphis and Boston leads to an increase in commerce between the two cities, with the effect that plants located in Memphis have now access to cheaper input goods from Boston suppliers. Or it might trigger an increase in tourism with more visitors now coming from Boston, stimulating local growth in the Memphis area that benefits, perhaps in some indirect way, local plants. In either case, investment at the Memphis plant might increase for reasons that have nothing to do with monitoring or information gathering by headquarters.

While I cannot completely rule out such alternative stories, they are unlikely to explain my results. First, if investment also increases at other Memphis plants, then this is already accounted for by the inclusion of MSA-year controls. Second, I show that my results are stronger for plants whose headquarters are more time constrained, consistent with the notion that a reduction in

travel time between headquarters and plants relaxes the time constraints of managers (who must regularly travel to plants to monitor them, gather (soft) information, etc.). Third, I show that my results are stronger for plants operating in soft-information industries. Fourth, I show that my results are stronger during the earlier years of the sample period, where other, non-personal, means of monitoring or transmitting information (internet, corporate intranet, video conferencing) have been either unavailable or less developed. Fifth, I show that my results are stronger for plants where the information uncertainty is likely to be greater, such as smaller plants, “peripheral plants” operating in industries that are not the firm’s main industry (and where headquarters is likely less knowledgeable), and plants operating in industries with more volatile sales or wages.

The rest of this paper is organized as follows. Section 2 describes the data and empirical methodology. Section 3 contains the main results. Section 4 presents robustness checks. Section 5 examines interaction effects with plant, firm, and industry characteristics. Section 6 concludes. Appendices A to C contain information on the construction and measurement of variables.

2 Data

2.1 Data Sources and Sample Selection

A. 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, while plants with fewer employees are randomly selected every five years, where the probability of being selected is higher for relatively larger plants. Although the ASM is technically 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-level variables that I use in my analysis, 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 on 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, the number of establishments, and the number of industry segments per firm. For my analysis, the most important firm-level information 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. For these firms, the ZIP code of the establishment coincides with the ZIP code of headquarters by construction. 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 Standard Statistical Establishment List (SSEL). The AES contains information on non-production ("auxiliary") establishments, including information on headquarters. The SSEL contains the name and address of all U.S. business establishments. Appendix A outlines the procedure used to obtain the location of headquarters from these data sets. A potential concern with this procedure is that the main source of information on headquarters, the AES, is available only every five years between 1977 and 2002. To fill in the missing years, I always use the information from the latest available AES. As the Census years are deterministic, this measurement error is unlikely to introduce any systematic bias. It merely introduces noise into the regression, which only makes it harder for me to find any significant results.

My sample covers the period from 1977 to 2005. (1977 is the first available AES year; 2005 is the latest available ASM year.) To be included in my sample, I require that a plant has a

⁴An establishment is a "single physical location where business is conducted" (Jarmin and Miranda, 2003, p. 15). Establishments are the economic units used in the Census data sets. In the manufacturing sector, the term "plant" is commonly used as a synonym for establishment.

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 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 headquarters during the sample period (7% of the firms in my sample). The results are virtually identical if these firms are included.

The above selection criteria leave me with 1,332,824 plant-year observations. In the difference-in-difference estimation, I use a 10-year window around the treatment date, implying that treated plants are included from five years before the treatment to five 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 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 that all plant-year observations of treated plants are included either before or after the treatment (as in, e.g., Bertrand and Mullainathan, 2003).

B. 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 which operate flights in the U.S. are required by law to file Form 41 with the DOT and are subject to sanctions for misreporting. Strictly speaking, the T-100 and ER-586 are thus not samples; they include *all* flights that have taken place between any two airports in the U.S.

The T-100 and ER-586 contain monthly data for each airline and route (“segment”). The data include origin and destination airports, flight duration (“ramp-to-ramp time”), scheduled departures, departures performed, passengers enplaned, and aircraft type.

⁵The 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.

2.2 Empirical Methodology

The introduction of new airline routes that reduce the travel time between plants and headquarters makes it easier for headquarters to monitor plants or gather “soft” information—i.e., information that cannot be transmitted through non-personal means. To examine the implications for plant-level investment and productivity, I use a differences-in-differences approach. Specifically, I estimate:

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

where i indexes plants, j indexes firms, l indexes plant location, t indexes years, y_{ijlt} is the dependent variable of interest (investment or productivity), α_i and α_t are plant and year fixed effects, respectively, treatment is a dummy variable that equals one 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 Metropolitan Statistical Area (MSA) level.⁶ The main coefficient of interest is β , which indicates the effect of the introduction of new airline routes.

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 shortest 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 route reduced the travel time between BOS and MEM and is coded as a “treatment” of the Memphis plant in 1986.⁷

⁶As 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. Since 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 U.S. plus 50 rural areas based on state boundaries (the District of Columbia has no rural area). For expositional simplicity, I refer to these 416 geographic units as “MSAs.”

⁷Section 2.3.C describes how the reduction in travel time in this specific example is computed.

To measure the effect on investment at the Memphis plant, one could simply compare investment after 1986 with investment before 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 by U.S. plants—e.g., due to favorable economic conditions or low interest rates. In that case, an increase in investment at the Memphis plant might not be due to the new airline route between MEM and BOS. To account for this possibility, I include a control group that consists of all plants in the U.S. that have not (yet) been treated. (Due to the staggering introduction of new airline routes, a plant remains in the control group until it is treated (possibly never); see, e.g., Bertrand and Mullainathan, 2003.) 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 these two differences is the estimated effect of the introduction of the new airline route between BOS and MEM on investment at the Memphis plant.

Including a control group accounts for the possibility of economy-wide shocks that are contemporaneous with the introduction of the new airline routes. However, since a treatment is defined at the plant-headquarters level, I can make the identification even tighter by also controlling for local shocks in the plant’s vicinity, thereby separating out the effect of the new airline routes from the effect of contemporaneous local shocks. Suppose, for instance, another plant, which is also located in the Memphis area, has its headquarters in Chicago. (Direct flights between Chicago O’Hare Airport and MEM were offered in each year during my sample.) If investment at the other Memphis plant also increases in 1986, then an increase in investment at the original Memphis plant (with headquarters in Boston) might not be due to the new airline route between MEM and BOS but rather 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. Unfortunately, computational difficulties make it infeasible to estimate a specification with so many fixed effects. Instead, I follow the approach 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 in the plant’s MSA in a given year, excluding the plant itself.

Since a treatment is defined at the plant-headquarters level—meaning each treatment is uniquely defined by *two* airport locations—I can make the identification even tighter by also

controlling for shocks at the firm level. In the above example, suppose the company headquartered in Boston has also other plants that are not located in Memphis. If investment at these other plants also increases in 1986, then an increase in investment at the Memphis plant might not be due to the new airline route between MEM and BOS but rather due to a contemporaneous shock at the firm level. Analogous to the construction of the MSA-year controls, I account for firm-level shocks by including “firm-year” controls, which are computed as the mean of the dependent variable across all of the firm’s plants, excluding the plant itself.⁸

In addition to accounting for economy-wide shocks, local shocks, and shocks at the firm level, my framework can address several concerns.

1. An important concern is that the introduction of new airline routes might be the outcome of lobbying by individual plants or firms experiencing positive investment shocks (“reverse causality”). This concern is very much minimized here. First, if the investment shock also affects other plants in the same MSA—e.g., the local economy experiences a boom, which leads local plants to lobby for a new airline route—then it is already accounted for by the inclusion of MSA-year controls. Similarly, if the investment shock is at the firm level and headquarters lobbies for the introduction of a new airline route to a specific plant, then it is already accounted for by the inclusion of firm-year controls. It remains the possibility that the investment shock is specific to a single plant who lobbies for the introduction of a new airline route to its headquarters. In that case, neither MSA-year controls nor firm-year controls would account for this possibility. However, I provide auxiliary evidence suggesting that this possibility is rather unlikely. First, I show that investment increases only with a lag of six to twelve months, meaning there is no effect either before the introduction of a new airline route or immediately after. 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 single plant is so powerful that it can successfully lobby for an airline merger or the opening of a new hub.
2. Even if we take as given the result that the introduction of new airline routes leads to an increase in plant-level investment, there are alternative stories that have little to do with

⁸Firm-year controls are set to zero for single-unit firms (i.e., firms with a single plant). Since for single-unit firms the ZIP code of the plant and the ZIP code of headquarters are the same by definition, all treated plants necessarily belong to multi-unit firms.

monitoring or information gathering by headquarters. For instance, in the above example, suppose the new airline route between MEM and BOS leads to an increase in commerce between the two cities, with the effect that plants located in Memphis have now access to cheaper input goods from Boston-based suppliers. Or suppose it triggers an increase in tourism with more visitors now coming from Boston, stimulating local growth in the Memphis area that benefits, in some indirect way, local plants. In either case, investment at the Memphis plant might increase for reasons that have nothing to do with monitoring or information gathering by headquarters.

While I cannot completely rule out such alternative stories, they are unlikely to explain my results. First, if investment also increases at other Memphis plants, then this possibility is already accounted for by the inclusion of MSA-year controls. Second, I show that my results are stronger for plants whose headquarters are more time constrained, consistent with the notion that the reduction in travel time relaxes the time constraints of managers (who must regularly travel to plants to monitor them, gather (soft) information, etc.). Third, I show that my results are stronger for plants operating in soft-information industries. Likewise, my results are stronger during the earlier years of the sample period, where other, non-personal, means of monitoring or transmitting information (internet, corporate intranet, video conferencing) have been either unavailable or less developed (see Petersen and Rajan (2002) for a similar research question). Fourth, I show that my results are stronger for plants where information uncertainty is likely to be greater, such as smaller plants, “peripheral plants” operating in industries that are not the firm’s main industry (and where headquarters is likely less knowledgeable), and plants operating in industries with more volatile sales or wages.

3. 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 such sources of travel time variation. Nevertheless, their omission is unlikely to affect my results. First, I show that my results are only significant for large reductions in travel time (at least one hour), 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 assigned to the control group. Thus, to the extent that such alternative sources of travel time reduction lead to an increase in plant-level investment, their omission would imply that my results *understate* the true effects of a reduction in travel time.⁹

4. I do not consider the termination of existing airline routes, but only the introduction of new airline routes. Terminations are much less frequent than introductions. Moreover, since routes that are discontinued are mainly minor regional routes, the resulting increase in travel time is likely to be modest. In robustness checks, I show that my results are unchanged if I additionally account for the termination of existing airline routes. Precisely, I augment the specification in equation (1) by adding a second treatment dummy that equals one whenever the termination of an existing airline route leads to an increase in travel time between plants and headquarters. I find that including this second treatment dummy has no effect on the coefficient on the original treatment dummy.
5. My focus on commercial airlines may be too narrow if companies own private jets (or if they contract with jet charter services). However, if managers use private jets to fly from headquarters to plants, then the introduction of new airline routes should not matter. Accordingly, if the effect documented here comes only from a subset of the treatment group—namely, plants whose headquarters actually makes use of commercial airlines—while the rest of the treatment group does not respond to the treatment (because private jets are used), then this merely implies that the treatment dummy will be noisy and its coefficient will be biased toward zero (“attenuation bias”), which only goes against my finding any significant results.¹⁰

⁹Large 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 WWI. 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 (see 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 likely due to changes in airline routes, not due to the expansion of roads and railroad networks.

¹⁰A story in which the unobserved use of private jets might explain an increase in investment after the intro-

6. My sample spans 29 years of data (from 1977 to 2005). Arguably, the effect of a new airline route on investment may not be permanent, and whatever happens five or more years later may not be due to the treatment. To account for this possibility, I use a 10-year treatment window that begins five years before the treatment and ends five 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 that all plant-year observations of treated plants are included either before or after the treatment. I also provide direct evidence on the dynamic effects of the introduction of new airline routes by considering successive 6-month intervals before and after the treatment. As expected, I find that the effect becomes weaker approximately 30 months after the treatment.

7. An important concern with the differences-in-differences methodology is that serial correlation of the error term can lead to understated standard errors (Bertrand, Duflo, and Mullainathan, 2004). In all 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, but also for any arbitrary correlation of the error terms across plants in the same MSA in any given year as well as over time (see Petersen, 2009). My results are similar if I cluster standard errors at the firm level or at both the MSA and firm level. Finally, 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 in Bertrand, Duflo, and Mullainathan (2004).

2.3 Definition of Variables

A. 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 amount of output a plant is expected to produce for given levels of inputs. To compute predicted output, I use a log-

duction of a new airline route goes as follows. Suppose managers use corporate jets to travel from headquarters to plants before the introduction of a new airline route. Once the new airline route is introduced, however, the private jet is no longer needed. As a consequence, the company sells the jet and uses the proceeds to finance new investment. However, in this case, investment should increase for *all* plants of the company—or at least for other plants besides the treated plant—which is already accounted for by the inclusion of firm-year controls.

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

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

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 3-digit SIC codes. (The results are qualitatively similar if I use 2- or 4-digit SIC codes).¹¹ 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). While this approach is common in the literature (e.g., Schoar, 2002; Bertrand and Mullainathan, 2003), it is not uncontroversial. Recent research in industrial organization argues that two econometric issues arise when production functions are estimated by OLS (for a review, see Akerberg et al., 2007). To illustrate these arguments, it is helpful to decompose the error term in equation (2) into two components: $\varepsilon_{it} = \omega_{it} + \eta_{it}$. While 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 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 here since 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

¹¹Until 1996, SIC codes were the basis for all Census Bureau publications. Since 1997, the Census Bureau has been using the North American Industry Classification System (NAICS). SIC codes were not discontinued until the 2002 Census, though. 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. The 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.

ω_{it} . As a consequence, OLS estimates of the coefficients in equation (2) will be biased and inconsistent. This endogeneity problem is often referred to as “simultaneity problem.”

The second endogeneity issue, the “selection problem,” arises if a plant whose observed or predictable productivity shock ω_{it} is below a certain threshold decides to shut down. Since plants have knowledge of ω_{it} prior to this decision, surviving plants will have ω_{it} drawn from a selected sample. The selection criteria may be determined by 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, renders the OLS estimates biased and inconsistent.

A variety of techniques have been proposed to address the simultaneity and selection problems. In robustness checks, I use the structural techniques suggested by Olley and Pakes (OP, 1996) and Levinsohn and Petrin (LP, 2003). OP and LP address the simultaneity problem by using investment and intermediate inputs, respectively, to proxy for the productivity shock ω_{it} . The intuition is that profit-maximizing behavior induces plants expecting positive productivity shocks to increase investment and use more intermediate inputs, respectively.¹² The selection problem is addressed by estimating plant survival propensity scores. Appendix C describes how OP’s and LP’s techniques are implemented using the plant-level data in my study. Regardless of whether I use the techniques of OP or LP, however, I find that my results are virtually identical to those when I estimate TFP by OLS.

TFP measures rely on structural assumptions (e.g., Cobb-Douglas production function). In robustness checks, I use two alternative measures of plant performance that are independent of structural assumptions: operating margin (OM) and return on capital (ROC). The numerator of OM is the value of shipments minus labor and material costs. This numerator is then divided by the value of shipments. All dollar values are expressed in 1997 dollars.¹³ ROC is defined similarly, except that the numerator is divided by the value of capital. OM and ROC are

¹²OP’s technique requires that investment levels be always strictly positive. This requirement comes at a cost, since discarding plant-year observations with zero investment results in an efficiency loss. The primary motivation of LP’s procedure is to mitigate this efficiency loss by focusing on plants’ choices of intermediary inputs. Unlike investment, intermediary inputs are rarely zero.

¹³Deflators for shipments and material costs are available at the 4-digit SIC level from the NBER-CES Manufacturing Industry Database. Deflators for labor costs are available at the 2-digit SIC level from the Bureau of Economic Analysis.

industry-adjusted by subtracting the industry median in a given 3-digit SIC industry and year.

A limitation of the productivity measures used in my study is that they do not include headquarters and other firm-level costs that are not reported at the plant level (e.g., overhead costs, research and development expenditures). As a consequence, the measures may overestimate the true productivity of plants of companies with high firm-level costs. However, by including the firm-year control in regression (1), I account for any time-varying firm-level component of productivity. Hence, the omission of firm-level costs at the plant level is unlikely to affect any of the regression coefficients.¹⁴

Finally, all of the productivity measures are subject to extreme values. To avoid that outliers drive my results, I winsorize all productivity measures at the 2.5th and 97.5th percentiles of their empirical distribution.

B. Measuring Investment

Investment is defined as total capital expenditures divided by capital stock. Both the numerator and denominator are expressed in 1997 dollars.¹⁵ Investment is industry-adjusted by subtracting the industry median in a given 3-digit SIC industry and year. To mitigate the effect of outliers, I winsorize investment at the 2.5th and 97.5th percentiles of its empirical distribution.

C. 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 are time constrained and that they make optimal decisions. Accordingly, they choose the route and means of transportation (car, plane) that minimizes the travel time between headquarters and plants.

¹⁴Likewise, the productivity measures used in my study may misrepresent a plant's true productivity if transfer prices between vertically related segments of the same firm are strategically distorted. To address this concern, I have constructed a matrix of vertical relatedness from the Input-Output tables of the Bureau of Economic Analysis. Based on this matrix, I identify upstream and downstream segments using the procedure described in the Appendix of Ozbas and Scharfstein (2010) and construct dummy variables indicating whether a plant belongs to a downstream or upstream segment. I obtain virtually identical results when these dummies are included in my regressions.

¹⁵Capital expenditures are deflated by the 4-digit SIC investment deflator from the NBER-CES Manufacturing Industry Database. Appendix B describes how real capital stock is constructed.

To identify the location of headquarters and plants, I use 5-digit ZIP codes from the LBD. (More 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 then used as a benchmark for comparison with the travel time by air based on the shortest airline route. Whenever traveling by car is faster, air transportation is ruled out by optimality, and the relevant travel time is the driving time.

To determine the shortest airline route between any two ZIP codes, I use the itinerary information from the T-100 and ER-586 data. The shortest airline route minimizes the total travel time between 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 (average ramp-to-ramp time of all flights performed between any two airports). 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 together and that each stopover takes another hour.¹⁶

None of my results are sensitive to the assumptions made about the time spent at airports and the layover time. I obtain virtually identical results when making different assumptions. Hence, while the assumptions reflect what I believe are sensible estimates (see footnote 16), they do not impact any of my results. Also, note that while the assumptions affect the measurement of total travel time, they have no impact on the construction of the treatment dummy. A

¹⁶The rationale behind these assumptions is as follows. First, most airlines require passengers to check in at least 30 minutes before departure. This cutoff is a lower bound for the time spent at the origin airport. To the extent that managers are time-constrained, they minimize the time spent at the airport. Accordingly, I assume that managers arrive at the origin airport about 10 to 20 minutes before the 30-minute cutoff and only need another 10 to 20 minutes to exit the destination airport. In total, these timing conventions correspond to an average time of approximately one hour at the origin and destination airports together. Second, most airlines have a minimum connection time of 30 minutes. In many cases, however, the actual connection time is longer. 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.

“treatment,” as defined in my study, is the introduction of a new airline route that reduces the travel time between headquarters and plants. Thus, the identification comes entirely from *changes* in travel time. When comparing different airline routes, the time spent at the origin and destination airports therefore “cancels out,” as it is the same for all routes. Likewise, the layover time only affects the magnitude of the change in travel time, but not whether a change in travel time is coded as a treatment. For example, if a direct flight replaces an indirect flight, the resulting total airborne time is shorter regardless of the assumed duration of the layover.

As an illustration of how the travel time is computed and how the treatment dummy is constructed, consider the example from Section 2.2. In this example, I consider a company headquartered in Boston with a plant located in Memphis. In 1985, no direct flight was offered between BOS and MEM. The shortest way to connect both airports was an indirect flight operated by Delta Airlines with a stopover in Atlanta. The total travel time, including travel to and from the airports by car, was 420 minutes.¹⁷ 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 travel time associated with this new route is 290 minutes.¹⁸ Hence, the introduction of the new direct route from BOS to MEM reduced the travel time from headquarters to the Memphis plant by 130 minutes (32% of the original travel time) and is therefore coded as a “treatment” of the Memphis plant in 1986.

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

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

¹⁷In this example, I assume that headquarters are located in the financial district of Boston (ZIP code 02110) and that the plant is located in an industrial area in Memphis (ZIP code 38115). The total travel time (in minutes) is obtained as follows: 10 (driving time from headquarters to BOS) + 180 (flight time from BOS to ATL) + 95 (flight time from ATL to MEM) + 15 (driving time from MEM to the plant) + 60 (time spent at BOS and MEM airports) + 60 (layover at ATL) = 420 minutes. Note that, in this example, the driving time between headquarters and plant if the entire trip is done by car is 1,175 minutes. Hence, air transportation is clearly the optimal means of transportation.

¹⁸The travel is computed as follows: 10 (driving time from headquarters to BOS) + 205 (flight time from BOS to MEM) + 15 (driving time from MEM to the plant) + 60 (time spent at BOS and MEM airports) = 290 minutes.

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

2.4 Summary Statistics

Table 1 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], “Eventually New Airline Route”) and plants that are never treated during the sample period (column [3], “No New Airline Route”). For each plant characteristic, the table reports the mean and standard deviation (in parentheses).¹⁹ All dollar values are expressed in 1997 dollars.

Column [1] shows that the average plant has shipments of about \$50 million and employs 213 workers. It is located 312 miles away from headquarters, which corresponds to a travel time of slightly over 2 hours (based on the optimal route and means of transportation). Furthermore, the average plant belongs to a firm that employs a total of 9,203 workers, owns 111 establishments, and operates in eight different 3-digit SIC industries.²⁰

The sample of eventually treated plants accounts for a relatively small fraction of the total plant-year observations (about 5.5%). That the group of eventually treated plants is relatively small is not a concern, however. Reliable identification of the treatment dummy only requires that this group be sufficiently large in absolute terms. A sample of 70,467 plant-year observations is a sufficiently large sample.²¹

The summary statistics also indicate that eventually treated plants differ from non-treated

¹⁹Due to the Census Bureau’s disclosure policy, I cannot report median or other quantile values.

²⁰The number of establishments (“plants”) reported here differs from the number in Bertrand and Mullainathan (2003, p. 1054), who report an average of 45 plants per firm. The reason is that I use the LBD to obtain the number of establishments, which includes not only all manufacturing plants but *all* (i.e., manufacturing and non-manufacturing) plants. In contrast, Bertrand and Mullainathan only report the number of manufacturing plants per firm.

²¹The treatment group is based on 10,533 treated plants. Thus, on average, I have about seven years of data for each treated plant. Since I consider a symmetric window of five years before and after the introduction of a new airline route, the maximum possible number of years is ten by construction. Missing years can occur for three reasons. First, as stated in Section 2.1, the ASM is a rotating probability sample. While plants with more than 250 employees are included in each ASM year, plants with fewer than 250 employees are randomly selected every five years. Second, plants can be acquired by other firms after the treatment. Third, a small number of plant-year observations is lost due to missing values of the variables used in the regressions. I have verified that my results are robust if I include only plants for which I have data for the entire 10-year treatment window.

plants in important ways. Specifically, eventually treated plants are larger, belong to larger firms, and are located farther away from headquarters. All of these differences make sense intuitively. In order to be treated, a plant needs to be sufficiently far away from headquarters, so that air travel is the optimal means of transportation. (The average distance to headquarters is 854 miles in column [2], while it is only 281 miles in column [3].) However, plants that are located farther away from headquarters are often part of larger companies that own many, and bigger, plants, and that operate in many different industries. As the firm-level statistics in Table 1 show, eventually treated plants belong to parent companies that are 2.6 times bigger (measured by employees), have 2.8 times more plants, and operate in more than twice as many industries as the parent companies of plants that are never treated.

The differences between eventually treated plants and non-treated plants are not the result of the introduction of new airline routes, as they exist even before the introduction of these new routes. That said, they raise the important question of whether the control group is an appropriate one. There are several reasons why this should not be a serious concern. First, I control for plant size and age in all my regressions. (I could additionally control for firm size and age, and the results would be virtually unchanged.) Second, due to the staggered nature of the introduction of new airline routes, eventually treated plants are first in the control group and only later—once they have been treated—in the treatment group. This is the same approach as in Bertrand and Mullainathan (2003). Third, again following Bertrand and Mullainathan (2003), I show in robustness checks that my results are unchanged if I restrict the sample to the 70,467 plant-year observations of eventually treated plants. Fourth, as is shown in Table 3 below, I obtain virtually 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.

The 70,467 plant-year observations in column [2] of Table 1 correspond to 10,533 treated plants. In **Table 2**, 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 (e.g., the new direct flight uses an airport that is closer to headquarters or the plant); 2) “Indirect to Indirect”: a new indirect flight using a different route replaces a previously optimal indirect flight (e.g., 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 (e.g., the

new direct flight from BOS to MEM in the above example); 4) “Road to Flight”: a new direct or indirect flight replaces car travel as the previously optimal means of transportation.

For all treated plants (column [1]) and separately also for each of the above four categories (columns [2] to [5]), Table 2 reports the average distance in miles between headquarters and plants, the average travel time before and after the introduction of the new airline route, and the average travel time reduction, both in absolute and relative terms. As column [1] shows, the average travel time reduction across all treated plants is 1 hour and 43 minutes, which amounts to a reduction of 25%. The breakdown in columns [2] to [5] shows that there is heterogeneity across the four categories. The category “Indirect to Indirect” accounts for the largest reduction in travel time (2 hours and 26 minutes), followed by the category “Indirect to Direct” (2 hours and 7 minutes) and the category “Direct to Direct” (1 hour and 12 minutes). Larger reductions in travel time are associated with longer physical distances (1,211, 942, and 726 miles, respectively). Finally, the category “Road to Flight” only applies 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 reduction in travel time is small for this category (47 minutes).

3 Results

3.1 Main Results

Table 3 presents the main results. All regressions include plant and year fixed effects. Standard errors are clustered at the MSA level.²² Column [1] shows the effect of the introduction of new airline routes on investment. Investment is defined as capital expenditures divided by capital stock and is industry-adjusted at the 3-digit SIC level. As is shown, the coefficient on the treatment dummy is 0.008, which implies that investment increases by 0.8 percentage point on average. The coefficient is statistically highly significant. It is also significant in economic terms. Given that the sample mean of investment is 0.10 (treated and non-treated plants alike), an increase of 0.8 percentage points implies that investment increases by 8%.

In columns [2] and [3], I examine the robustness of this result to using alternative speci-

²²As mentioned earlier, 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 in Bertrand, Duflo, and Mullainathan (2004).

fications. 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). I also control for plant age and size. MSA- and firm-year controls are defined in Section 2.2. Age is the logarithm of one plus the number of years since the plant is covered in the LBD. Size is the logarithm of the number of employees. As is shown, the results are not sensitive to the inclusion of control variables. In fact, the coefficient on the treatment dummy is slightly larger (0.009). 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.²³

Overall, I find that the introduction of new airline routes leads to an increase in plant-level investment. This result is consistent with models of irreversible investment predicting a negative association between information uncertainty and investment (e.g., Arrow, 1968; Bernanke, 1983; McDonald and Siegel, 1986; Pindyck, 1988; Dixit and Pindyck, 1994) as well as capital budgeting models predicting that information asymmetry between plants (or divisions) and headquarters leads to underinvestment (e.g., Bernardo, Cai, and Luo, 2001, 2004).

An important question is whether investment is efficient. While neoclassical models of investment generally presume efficiency, a class of models in the organizational and capital budgeting literature argues that investment decisions within firms are distorted by lobbying and rent-seeking activities (e.g., Milgrom, 1988; Meyer, Milgrom, and Roberts, 1992; Rajan, Servaes, and Zingales, 2000; Scharfstein and Stein, 2000). For instance, it could be the case that investment does not increase because information uncertainty (or asymmetry) is reduced, but rather because the introduction of new airline routes makes it easier for plant-level managers to lobby headquarters for bigger capital budgets. In this case, however, the increase in investment should be *inefficient*, a result that has been emphasized by the rent-seeking literature.²⁴

To examine the implications for efficiency, I re-estimate the specifications in columns [1] to [3] with TFP as the dependent variable. The results, which are shown in columns [4] to [6], are inconsistent with theories of inefficient investment. Across all three specifications, the coefficient on the treatment dummy is highly significant. The coefficient lies between 0.013 and 0.014, which

²³I also obtain similar results if I interact year dummies with other plant characteristics from Table 1 (excluding capital expenditures and capital stock, which enter into the definition of the dependent variable).

²⁴In a follow-up paper (Giroud and Mueller, 2010), I explore this issue further by looking at investment spillovers across different plants of the same firm. The follow-up paper also looks at aggregate effects at the firm level, showing that TFP increases not only at the treated plant but at the firm level overall.

implies an increase in plant-level productivity of 1.3% to 1.4%. In robustness checks, I show how this effect on TFP also translates into other measures of performance, such as cash-flow based measures (return on capital and operating margin). Further evidence is also provided in Table 4 below, where I show that TFP increases only with a lag of six months after the increase in investment.

In the remainder of this paper, 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. However, all my results are similar if I exclude these four controls (see, e.g., columns [1] and [4] of Table 1), if I include only a subset of the controls, or if I additionally control for firm age and size.

3.2 Dynamic Effects of New Airline Routes

An important concern is that the introduction of new airline routes might be the outcome of lobbying by individual plants or firms experiencing positive investment shocks (“reverse causality”). As discussed earlier, the inclusion of MSA- and firm-year controls accounts for the possibility of investment shocks at the MSA and firm level, respectively. It remains the possibility that the investment shock is specific to a single plant who lobbies for the introduction of a new airline route to its headquarters. In the following, I test an implication of this hypothesis (see Bertrand and Mullainathan (2003) for a similar test). Precisely, if a plant experiences a positive investment shock prior to the introduction of a new airline route, then I should find an “effect” already before the new route is introduced.

To examine this hypothesis, 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 end month of each plant-year observation is December. Since 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. Thus, 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 example from Section 2.2, where a new direct flight between MEM and BOS is introduced in October 1986. In that case, the 1985 observation of the Memphis plant with headquarters in Boston is recorded nine months before the treatment, the 1986 observation of the same plant is recorded three months after the treatment, the 1987 observation of the same plant is recorded 15 months after the treatment, and so on.

By exploiting the variation in 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 plant-year observation and treatment. Specifically, I use a set of eight dummies. The first dummy, “Treatment (-12m, -6m),” equals one if the plant-year observation is recorded between 12 and six months before the treatment. The other dummies are defined accordingly with respect to the intervals (-6m, 0m), (0m, 6m), (6m, 12m), (12m, 18m), (18m, 24m), (24m, 30m), and 30 months and beyond (“30m +”), respectively.

Table 4 shows the results. The main variables of interest are Treatment (-12m, -6m) and Treatment (-6m, 0m), which measure the “effect” of new airline routes prior to their introduction. Finding such an “effect” would be symptomatic of reverse causation. As column [1] shows, the coefficients on both variables are small and insignificant. What is more, the coefficient on Treatment (0m, 6m), which captures the effect of new airline routes within the first six months after their introduction, is also insignificant. The only way to reconcile a lobbying story with these results is to argue that plants lobby for the introduction of new airline routes in anticipation of investment shocks that occur in the distant future. For instance, suppose it takes, say, six months between the lobbying and the introduction of a new airline route. In this case, reverse causation would imply that plants lobby for new airline routes in anticipation of investment shocks that occur 12 to 18 months later. While not impossible, this is unlikely. Nevertheless, in Section 4.1, I provide further evidence that is inconsistent with a lobbying story. Specifically, I show that my results are robust to considering only new airline routes that are the outcome of a merger between two airlines or the opening of a new hub. Arguably, it is less likely that a plant is so powerful that it can successfully lobby for an airline merger or the opening of a new hub.

The other estimates in column [1] provide further insights into the dynamic effects of the introduction of new airline routes. While the effect becomes significant after six months, it remains initially small in economic terms (0.005). It is only after 12 months that the effect becomes large and highly significant. Specifically, the coefficients on Treatment (12m, 18m), Treatment (18m, 24m) and Treatment (24m, 30m) are between 0.013 and 0.014, which implies that investment increases by 13% to 14%. In the longer run (i.e., 30 months and beyond), the magnitude of the coefficient reverts to a lower level (0.009), suggesting that the effect of the introduction of new airline routes becomes weaker over time.

Column [2] reports the results for TFP as the dependent variable. The dynamic pattern

is similar to above, except that the increase in TFP occurs with a lag of six months after the increase in investment. As above, the effect on TFP becomes weaker in the longer run.

3.3 Small versus Large Reductions in Travel Time

Any new airline route that reduces the travel time between headquarters and plants is coded as treatment, regardless of the magnitude of the travel time reduction. Arguably, the treatment effect may be stronger for larger reductions in travel time. To examine this hypothesis, 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).²⁵ The results are shown in **Table 5**. In column [1], the dependent variable is investment. As is shown, the introduction of new airline routes has a small and insignificant effect on investment if the travel time reduction is less than one hour. Accordingly, if the time saving associated with the introduction of new airline routes is small, managers do not seem to find it worthwhile to change their existing itineraries. Once the reduction in travel time exceeds one hour, the effect becomes significant. Moreover, it is monotonic in the magnitude of the travel time reduction. The strongest effect is observed when the reduction in travel time is more than two hours. In this case, investment increases by 15%, which is almost twice as large as the average treatment effect reported in Table 3.

In column [2], the dependent variable is TFP. The results mirror those in column [1]. The effect is again monotonic in the reduction in travel time, is strongest when the travel time reduction exceeds two hours, and is small and insignificant when the travel time reduction is less than one hour.

4 Robustness

4.1 Hub Openings and Airline Mergers

In Section 3.2, I addressed concerns that the introduction of new airline routes might be the outcome of lobbying by plants or firms experiencing positive investment shocks (“reverse causality”). I now present further evidence suggesting that this possibility is unlikely. Specifically, I

²⁵I obtain similar results if instead of using 30-minute cutoffs, I use quintiles based on the empirical distribution of the reduction in travel time.

consider new airline routes that are the outcome of a merger between two airlines or the opening of a new hub.²⁶ Arguably, it is less likely that a plant or firm is so powerful that it can successfully lobby for an airline merger or the opening of a new hub.

Table 6 provides a list of all 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). Specifically, 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. As a result, the 1980s witnessed a substantial increase in the number of hub openings (Ivy, 1993).

Table 7 provides a list of airline mergers that were completed during the sample period. 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. While many airline mergers were completed during the sample period, I only consider mergers that account for at least one treatment in my sample. Mergers of small commuter airlines servicing few locations often fail to satisfy this criterion.²⁷ As can be seen, 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

²⁶I thank Adair Morse for suggesting the idea to look at hub openings.

²⁷I apply three additional criteria when compiling the list of airline mergers. First, I only consider 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 listed in Table 7 is the year in which the airlines 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 7, the 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.

consolidation phase was largely over. As a result, industry-wide concentration increased sharply, with the nine largest airlines representing a total market share of over 90% of domestic revenue passenger miles (Goetz and Sutton, 1997).

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.²⁸ In total, my sample includes 1,761 hub treatments and 535 merger treatments, which combined account for 22% of all treated plants.²⁹

In **Table 8**, 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. In column [1], the dependent variable is investment. In column [2], the dependent variable is TFP. As is shown, the coefficients on all three dummies are economically large and statistically significant. They are largest for hub treatments, slightly smaller for merger treatments, and smallest for the “other” treatments. These differences among the coefficients are 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 was shown in Table 5, larger travel time reductions are associated with stronger treatment effects. The average travel time reduction for hub treatments is 2 hours, and the average travel time reduction for merger treatments is 1 hour and 51 minutes.³⁰ In contrast, the average travel time reduction for the “other” treatments is 1 hour and 39 minutes. Most importantly, however, that all three coefficients are large and significant is inconsistent with reverse causality arguments based on individual plants or firms lobbying for the introduction of new airline routes.

²⁸If 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 at the same time.

²⁹Due to Census Bureau’s disclosure policy, I cannot report treatment counts for individual hub openings and airline mergers.

³⁰The average reduction in travel time of two hours for hub treatments is likely a lower bound. By construction, layover times are assumed to be constant across airports. However, one of the very benefits of hub airports is to reduce the layover time of indirect flights. Hence, the actual reduction in travel time due to new hub openings is likely larger than two hours.

4.2 Alternative Control Groups

In my baseline regressions, the control group consists of all plant-year observations that have not been treated by time t . In **Table 9**, I examine the robustness of my results to using alternative control groups.

A. Multi-unit Firms

In columns [1] and [2], I exclude single-unit firms, meaning the sample is comprised exclusively of multi-unit firms. Single-unit firms consist of a single establishment (“unit”), which implies that the ZIP code of the plant and the ZIP code of headquarters are the same by definition. Consequently, single-unit firms cannot possibly be affected by the introduction of new airline routes, which implies they are never in the treatment group. Since the control and treatment groups should ideally be similar except for the treatment, it is an open question of whether single-unit firms should be in the control group. As columns [1] and [2] show, my results are virtually unaffected when single-unit firms are excluded.

B. Eventually Treated Plants

A more serious concern is that eventually treated plants and plants that are never treated during the sample period differ in important ways. As is shown in Table 1, eventually treated plants are larger, belong to larger firms, and are located farther away from headquarters. As discussed previously, these differences make sense intuitively. In order to be treated, a plant needs to be sufficiently far away from headquarters, so that air travel is the optimal means of transportation. However, plants that are located far away from headquarters are often part of larger companies that own more, and bigger, plants. While these differences are intuitive, they raise the question of whether the control group is an appropriate one. There are several reasons why this should not be a serious concern. First, I control for plant size and age in all my regressions. (I could additionally control for firm size and age, and the results would be unchanged.) Second, due to the staggered introduction of new airline routes, eventually treated plants are first in the control group and then in the treatment group. Third, the staggered introduction of new airline routes allows me to exclude plants that are never treated during the sample period altogether, so that the control group consists only of eventually treated plants that have not yet been treated by time t (see Bertrand and Mullainathan (2003) for a similar robustness check). Columns [3] and

[4] show the results from re-estimating (1) using only the sample of eventually treated plants. As is shown, the results are qualitatively similar to my baseline results. While the effect on investment is slightly larger than before (0.011, compared to 0.009 in column [2] of Table 3), the effect on TFP is slightly smaller (0.010, compared to 0.013 in column [5] of Table 3).

C. Increase in Travel Time

In my baseline specification, I consider only the introduction of new airline routes, not the termination of existing routes. By construction, plants whose travel time to headquarters increases due to the termination of existing airline routes are therefore in the control group. Under the hypothesis that an increase in travel time leads to a decrease in investment and productivity, this “misclassification” will bias upward the coefficient on the treatment dummy, meaning I will overestimate the effect of the introduction of new airline routes. While this is a valid concern, it is very much minimized here, for two reasons. First, terminations of existing airline routes are only relevant for 3,071 plants in my sample. Accordingly, they only account for a small subset of the control group. Second, my results are unchanged if I assign these plants to a separate treatment group consisting of plants whose travel time to headquarters increases due to the termination of existing airline routes. (Analogous to plants whose travel time is reduced due to the introduction of new airline routes, the plants are first in the control group and then—once they are treated—in the treatment group.)

Specifically, I augment the specification in equation (1) by including a second treatment dummy (“increase in travel time”), which equals one whenever the termination of an existing airline route leads to an increase in travel time between the plant and its headquarters. The results are reported in columns [5] and [6].³¹ Not surprisingly, the coefficient on the “increase in travel time” dummy is of the opposite sign as the coefficient on the main treatment dummy. (That the former coefficient is statistically weaker is likely due to the fact that it is identified with much fewer observations.) Most importantly, however, the coefficient on the main treatment dummy is identical to that in my baseline regressions, both for investment and TFP. Hence,

³¹The number of observations decreases slightly from 1,291,280 to 1,282,228. In the main specification, plant-year observations of treated plants are restricted to five years before and after the treatment. For consistency, I apply the same rule to plants that are treated due to an increase in travel time. The results are unchanged if I use the entire sample of 1,291,280 plant-year observations instead.

whether I account for increases in travel time is immaterial for my results.³²

4.3 Alternative Measures of Productive Efficiency

My main measure of productivity is TFP. In **Table 10**, I consider alternative measures of productive efficiency. In the first two regressions, I use margin-based measures. In column [1], the dependent variable is return on capital (ROC). In column [2], the dependent variable is operating margin (OM). In either case, the results are similar to before. Specifically, the introduction of new airline routes leads to a significant increase in ROC (OM) by 1.3 (0.9) percentage points. Given the high correlation of TFP with both ROC (60%) and OM (50%), this is not surprising.

In columns [3] and [4], I use the structural techniques of Olley and Pakes (OP, 1996) and Levinsohn and Petrin (LP, 2003), respectively, to estimate the coefficients of the Cobb-Douglas production function. As discussed previously, these methods account for the endogeneity of input choices as well as selection effects. In either case, the results are similar to before. Again, this is not surprising, given that the correlation between TFP estimated by OLS and TFP estimated using OP’s (LP’s) technique is 81% (84%).³³

5 Interaction Effects

In this paper, a treatment is the introduction of a new airline route that reduces the travel time between plants and headquarters. A natural interpretation is that this reduction in travel time makes it easier for headquarters to monitor plants or gather “soft” information—i.e., information that cannot be transmitted through non-personal means (see footnote 3 for a definition). However, even if we take as given the result that the introduction of new airline routes leads

³²Similar to the analysis in Section 4.1, I can focus exclusively on terminations of existing airline routes due to exogenous reasons, such as hub closings, airline mergers, and bankruptcies. Analogous to what I find in Table 8, the effect is strongest for hub closings and airline mergers. That said, the effect is small and insignificant for bankruptcies, which is likely due to the limited scope of the bankrupt airlines’ operations. (The vast majority of bankrupt airlines are small regional carriers, meaning the resulting increase in travel time, and thus the treatment effect, is relatively modest.)

³³Since OP’s and LP’s techniques require non-missing lag values of the production factors, the sample size in columns [3] and [4] of Table 10 is smaller than in Table 3. In addition, OP’s method requires non-zero investment values, which further reduces the sample size. OP’s and LP’s techniques are described in Appendix C.

to an increase in plant-level investment, there are alternative stories that have little to do with monitoring or information gathering. For instance, in the example of the Memphis plant with headquarters in Boston, suppose the new airline route between MEM and BOS leads to an increase in commerce between the two cities, with the effect that plants located in Memphis have now access to cheaper input goods from Boston-based suppliers. Or suppose it triggers an increase in tourism with more visitors now coming from Boston, stimulating local growth in the Memphis area that benefits, in some indirect way, local plants. In either case, investment at the Memphis plant might increase for reasons that have nothing to do with monitoring or information gathering by headquarters.

While I cannot completely rule out such alternative stories, I believe they are unlikely to explain my results. First, if investment also increases at other Memphis plants, then this possibility is already accounted for by the inclusion of MSA-year controls. Second, I provide below auxiliary results that are supportive of an information-type story while, collectively, they are difficult to reconcile with alternative stories. In particular, I show that my results are stronger:

1. for plants whose headquarters are more time constrained, consistent with the notion that the reduction in travel time relaxes the time constraints of managers (who must regularly travel to plants to monitor them, gather (soft) information, etc.);
2. for plants operating in soft-information industries;
3. during the earlier years of my sample period, where other, non-personal, means of monitoring or transmitting information (internet, corporate intranet, video conferencing) have been either unavailable or less developed;
4. for plants where the information uncertainty is likely to be greater, such as smaller plants, “peripheral plants” operating in industries that are not the firm’s main industry (and where headquarters is likely less knowledgeable), and plants operating in industries with more volatile sales or wages.

A. Headquarters’ Time Constraints

Most alternative stories, including the two stories sketched above, do not involve headquarters. Accordingly, measures of headquarters’ time constraints should have no effect on the results. In

contrast, if the introduction of new airline routes makes it easier for headquarters to visit plants, then my results should be stronger for plants whose headquarters are more time constrained. To examine this hypothesis, I construct two measures of headquarters’ time constraints. The first measure is the number of employees at headquarters divided by the total number of plants of the company (“Managers/Plants”). The second measure is the number of employees at headquarters divided by the total distance (in miles) between headquarters and all of the company’s plants (“Managers/Total Distance”). In either case, the idea is that managers based at headquarters must regularly visit the company’s plants. The lower is the ratio of managers to plants—or the greater is the average distance the managers must travel—the more time constrained are headquarters.

Though employees at headquarters are white-collar employees, not all of them are “managers.” The number of employees at headquarters also includes, e.g., secretaries and clerical employees working at headquarters. Nevertheless, as long as the number of other employees is roughly proportional to the number of managers, this measurement error is unlikely to affect my results. Bearing this limitation in mind, I sort treated plants into two categories, “high time constraint” and “low time constraint,” indicating whether headquarters’ time constraint lies above or below the median value of all treated plants in the year prior to the treatment. Using pre-treatment values to sort plants mitigates concerns that the classification is affected by the treatment itself.

To examine whether my results are stronger for plants whose headquarters are more time constrained, I interact the treatment dummy in equation (1) with two dummies indicating whether headquarters’ time constraint is low or high. The results are presented in **Table 11**. As is shown, the effect of a reduction in travel time on investment is stronger when headquarters are more time constrained. In column [1], where headquarters’ time constraint is 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 (the p -value of the F -statistic is 0.028). The difference becomes even more pronounced when the measure of time constraints takes into account the geographic dispersion of plants. In column [3], where headquarters’ time constraint is 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 significant at the 1%

level ($p = 0.002$). The results when TFP is the dependent variable (columns [2] and [4]) mirror those for investment. In sum, my results are stronger for plants whose headquarters are more time constrained, consistent with the notion that the reduction in travel time relaxes the time constraints of managers.

B. Soft- versus Hard-information Industries

If the introduction of new airline routes makes it easier for headquarters to gather soft information about plants, then my results should be stronger for plants operating in soft-information industries. In contrast, under most alternative stories, whether a plant operates in a soft- or hard-information industry is immaterial for the results.

To test this hypothesis, I sort treated plants into soft- and hard-information industries using the classification in Landier, Nair, and Wulf (2009), which is based on Petersen and Rajan (2002). In columns [1] and [2] of **Table 12**, I interact the treatment dummy in equation (1) with two dummies indicating whether a plant operates in a soft- or hard-information industry. As is shown, the treatment effect is stronger for plants in soft-information industries. When investment is the dependent variable, the coefficient on the interacted treatment dummy is 0.012 in soft-information industries but only 0.007 in hard-information industries. The difference is significant at the 10% level ($p = 0.079$). The results when TFP is the dependent variable mirror those for investment. Overall, my results are stronger for plants in soft-information industries, consistent with the notion that the reduction in travel time makes it easier for headquarters to gather soft information about plants.

C. Innovations in Information Technology

The period from 1977 to 2005 witnessed major innovations in information technology. These innovations (e.g., internet, corporate intranet, video conferencing) played an important role in facilitating information flows both within and across company units (for a review, see Dewett and Jones, 2001), thus reducing the need for headquarters to “physically” travel to plants. Consequently, if the introduction of new airline routes matters because it makes it easier for headquarters to visit plants, then my results should be stronger during the earlier years of the sample period.

To examine this hypothesis, I interact the treatment dummy in equation (1) with three dummies indicating different time periods: before 1986 (nine years), between 1986 and 1995

(ten years), and after 1995 (ten years). The results are presented in columns [3] and [4] of Table 12. As is shown, the treatment effect is stronger during the earlier years of the sample period. When investment is the dependent variable, the coefficient on the interacted treatment dummy is 0.013 for the pre-1986 period, 0.010 for the period between 1986 and 1995, and only 0.005 for the post-1995 period. The difference between the pre-1986 and post-1995 coefficients is significant at the 5% level ($p = 0.012$). The results when TFP is the dependent variable are similar. In sum, as innovations in information technology reduce the need for headquarters to “physically” travel to plants, the treatment effect becomes gradually weaker.

D. Information Uncertainty at the Plant and Industry Level

If the introduction of new airline routes reduces information uncertainty between plants and headquarters, then my results should be stronger for plants where the information uncertainty is likely to be greater. To test this hypothesis, I sort treated plants along five different attributes that are likely to be correlated with the information uncertainty faced by headquarters.

The first attribute is the relative size of a plant within the company. The idea is that headquarters has more knowledge about larger plants within the company, and less knowledge about relatively smaller plants. Relative plant size is defined as the number of employees of the plant divided by the total number of employees of the company (from the LBD). I sort treated plants into “small” and “large” plants depending on whether the relative plant size is above or below the median value across all of the company’s plants in the year prior to the treatment.

The second attribute is whether a plant operates in a “main” or “peripheral” industry of the company. Main industries are 3-digit SIC industries that account for at least 25% of the company’s employees in the year prior to the treatment. All other industries are classified as peripheral.³⁴ The idea is that headquarters has more knowledge about plants operating in the company’s main industries, for two reasons. First, companies are likely to spend more time and

³⁴The 25% cutoff is the same as in Maksimovic and Phillips (2002), except that I use employees instead of shipments. Shipments is only available for manufacturing plants in the CMF and ASM. Thus, computing industry shares based on shipments measures the share of a given industry across all manufacturing industries, not across all industries within a company. Using employees overcomes this limitation since the LBD reports employees for *all* establishments. Nevertheless, I have verified that my results are similar if I restrict the sample to “pure” manufacturing firms (i.e., firms which exclusively have manufacturing establishments) and determine main and peripheral industries on the basis of shipments.

resources developing their core businesses. Second, companies have typically more experience in their core industries.

The third attribute is whether a plant was acquired by the company during the sample period (but prior to the treatment). The idea is that headquarters faces higher information uncertainty with respect to acquired plants as opposed to “own” plants, which the company built from scratch (or acquired before the sample period).

The fourth attribute is sales volatility within a plant’s industry. Sales volatility is closely related to the notion of information uncertainty in the theoretical literature on uncertainty and investment, which often models uncertainty as “demand uncertainty” (e.g., Abel, 1983; Pindyck, 1988). Sales volatility is computed as the standard deviation of annual shipment growth across all plants in a given 3-digit SIC industry. Treated plants are sorted into industries with high or low sales volatility depending on whether sales volatility in the plant’s industry lies above or below the median value across all industries in the year prior to the treatment.

The fifth attribute is wage volatility in a plant’s industry. While sales volatility represents uncertainty in output markets, wage volatility represents uncertainty in input markets.³⁵ Wage volatility is computed as the standard deviation of the hourly wage of production workers (production worker payroll divided by total production worker hours) across all plants in a given 3-digit SIC industry. Treated plants are sorted into industries with high or low wage volatility depending on whether wage volatility in the plant’s industry lies above or below the median value across all industries in the year prior to the treatment.

For each of the five attributes, I interact the treatment dummy in equation (1) with two dummies indicating the two possible realizations of the attribute (e.g., small and large plants). **Table 13** presents the results for the first three attributes, which all measure information uncertainty at the plant level. In columns [1] and [2], the attribute is the relative size of a plant within the company. As is shown, the treatment effect is stronger for relatively smaller plants. When investment is the dependent variable, the coefficient on the interacted treatment dummy is 0.013 for small plants but only 0.005 for large plants. The difference is significant at the 1% level ($p = 0.001$). The results are similar when TFP is the dependent variable.

In columns [3] and [4], the attribute is whether a plant operates in a “main” or “peripheral” industry of the company. As is shown, the treatment effect is stronger for plants in periph-

³⁵The relationship between wage uncertainty and investment is modeled in, e.g., Hartman (1972).

eral industries. When investment is the dependent variable, the coefficient on the interacted treatment dummy is 0.014 for plants in peripheral industries but only 0.004 for plants in main industries. The difference is significant at the 1% level ($p = 0.000$). The results when TFP is the dependent variable mirror those for investment.

In columns [5] and [6], the attribute is whether a plant was acquired by the company during the sample period. Above I argued that headquarters may face greater information uncertainty with respect to acquired plants, meaning the treatment effect should be stronger. However, this hypothesis is not supported by the data. As is shown, the coefficients on the interacted treatment dummies are almost identical for acquired and own plants, both for investment and TFP. A possible explanation is that plants are not acquired randomly. Rather, any acquisition involves due diligence, so that eventually headquarters may know as much about acquired plants as it does about own plants. In fact, headquarters may *choose* to acquire only plants where the information uncertainty is small (selection effect).

Table 14 presents the results for the last two attributes, which measure information uncertainty at the industry level. In columns [1] and [2], the attribute is sales volatility within a plant’s industry. As is shown, the treatment effect is stronger when sales volatility is high. When investment is the dependent variable, the coefficient on the interacted treatment dummy is 0.013 when sales volatility is high but only 0.006 when it is low. The difference is significant at the 1% level ($p = 0.004$). The results are similar when TFP is the dependent variable.

In columns [3] and [4], the attribute is wage volatility within a plant’s industry. As is shown, the results are qualitatively similar to those for sales volatility. When investment is the dependent variable, the coefficient on the interacted treatment dummy is 0.011 when wage volatility is high but only 0.007 when it is low. However, while qualitatively similar, the results are weaker in the sense that the difference between the two coefficients is only significant at the 10% level ($p = 0.100$). The results when TFP is the dependent variable mirror those for investment, except that the difference between the two coefficients is insignificant.

6 Conclusion

Using the introduction of new airline routes as a source of exogenous variation in headquarters’ information about plants, I examine the effects on plant-level investment and productivity. Consistent with models of irreversible investment—as well as capital budgeting and contracting

models predicting underinvestment—I find that a reduction in travel time between plants and headquarters leads to an increase in plant-level investment and productivity. I also provide auxiliary results suggesting that the reduction in travel time makes it easier for headquarters to monitor and gather soft information about plants, as opposed to alternative stories that have little to do with monitoring or information gathering. Specifically, I show that my results are stronger for plants whose headquarters are more time constrained, plants operating in soft-information industries, and plants where information uncertainty is likely to be greater, such as smaller plants, “peripheral plants” operating in industries that are not the firm’s main industry, and plants operating in industries with more volatile sales or wages.

In this paper, I focus on plants that are *directly* affected by the introduction of new airline routes, i.e., plants whose travel time to headquarters is reduced as a consequence. However, better monitoring or access to soft information at one plant may have *indirect* effects on other plants of the same company. For instance, if the company is financially constrained, it may have to fund investment at the treated plant by withdrawing funds from other plants. Another example are “information spillovers” across plants of the same company. For instance, information gathered at the treated plant could be relevant for investment at other plants, e.g., plants operating in the same industry. More generally, that the introduction of new airline routes might have indirect effects on other plants of the same company raises a whole set of interesting questions: Does investment at the other plants increase or decrease? And *which* other plants are affected (plant and industry attributes)? And does the answer depend on whether the company is financially constrained? And what is the aggregate effect on investment and TFP at the firm level? These, as well as many related, questions are currently being examined in a follow-up paper (Giroud and Mueller, 2010).

Appendix

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, p. 133). Auxiliary establishments include head-

quarters, warehouses, garages, and other facilities primarily engaged in servicing a company’s operating establishments.

To distinguish headquarters from other auxiliary establishments, I use the selection criteria of Aarland et al. (2007). Specifically, in the years 1997 and 2002, headquarters are identified by the 6-digit NAICS industry code 551114.³⁶ In the former years (1977, 1982, 1987, and 1992), prior to the introduction of NAICS codes, headquarters are identified as those establishments 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 corporate headquarters and regional or divisional administrative offices. As a result, these criteria may return more than one “headquarters” per company. In my manufacturing sample, 20% of the multi-unit companies have multiple headquarters in the AES. To identify the main headquarters, I supplement the AES with information from the SSEL. The SSEL contains name and address information for 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 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.

Finally, not all multi-unit companies have headquarters data in the AES. Since by definition auxiliary establishments are physically separated from the 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

³⁶The 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).

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 information from the SSEL is not sufficient to determine the corporate headquarters, I select the establishment with the highest payroll from the LBD (or highest white-collar payroll from the CMF if all establishments are manufacturing plants). Arguably, the latter criterion is subject to misclassification if, e.g., the headquarters is located in the smallest plant of the company. Fortunately, the impact of such misclassification is mitigated. 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 plants. As a consequence, companies without stand-alone headquarters account for only 7% of the treated plants. The results are unaffected if these plants are excluded from the sample.

To assess the accuracy of the headquarters location obtained from the Census micro data, I merge my dataset 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 only available for the latest available year of the database (2009), and may therefore be an incorrect benchmark for companies whose headquarters has moved since the last AES year (2002). 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. While this match may be considered satisfactory, I have also verified that my results are qualitatively similar if I restrict my sample to the publicly traded Compustat firms and use 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 total value of shipments plus changes in the value of inventories for finished goods

and work in process, divided by the 4-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 4-digit SIC material deflator from the NBER-CES Manufacturing Industry Database. *Labor* is measured in “production worker-equivalent hours,” following the procedure 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 non-production 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 2-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 4-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 4-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 2-digit SIC depreciation rate from the BEA, and I is capital expenditures divided by the 4-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.

Appendix C: Structural Estimation of Total Factor Productivity

In this appendix, I describe how the coefficients in equation (2) are estimated using LP and OP’s techniques. I first describe LP’s technique, which relies on variation in m_{it} to proxy for

the productivity shock ω_{it} . How to implement OP's technique is discussed subsequently.

Consider the log-linear Cobb-Douglas production function in equation (2). By decomposing the error term ε_{it} into ω_{it} and η_{it} , (2) can be rewritten as

$$y_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \omega_{it} + \eta_{it}. \quad (3)$$

Suppose that the demand for material inputs m_{it} is strictly monotonic in ω_{it} conditional on k_{it} and l_{it} .³⁷ Accordingly, it can be expressed as

$$m_{it} = f_t(\omega_{it}, k_{it}, l_{it}). \quad (4)$$

The function $f(\cdot)$ is indexed by t to account for time-varying variables that are part of the state space, but do not vary across plants (e.g., input prices). Since (4) is strictly monotonic in ω_{it} , it can be inverted to generate

$$\omega_{it} = f_t^{-1}(m_{it}, k_{it}, l_{it}). \quad (5)$$

Inserting (5) into the production function (3) yields

$$\begin{aligned} y_{it} &= \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + f_t^{-1}(m_{it}, k_{it}, l_{it}) + \eta_{it} \\ &= \Phi_{it}(k_{it}, l_{it}, m_{it}) + \eta_{it}. \end{aligned} \quad (6)$$

Estimating equation (6) is the first stage of the procedure.³⁸ Since the function $f_t(\cdot)$, and hence the inverse function $f_t^{-1}(\cdot)$, depend on the primitives of the model, the functional form of Φ is unknown. To overcome this limitation, (6) can be estimated non-parametrically. Specifically, I estimate (6) by OLS using a third-order polynomial in $\{k_{it}, l_{it}, m_{it}\}$. The predicted values from

³⁷LP state conditions under which the demand for material inputs is strictly increasing in ω_{it} .

³⁸Equation (6) is not identical to the first stage equation in the original LP article. In LP, the demand for material inputs $f_t(\cdot)$ does not depend on l_{it} . As a result, the first stage equation is of the form $y_{it} = \beta_l l_{it} + \phi_{it}(k_{it}, m_{it}) + \eta_{it}$. By treating $\phi_{it}(\cdot)$ non-parametrically, LP are able to identify β_l in the first stage equation. In a recent article, Akerberg, Caves, and Frazer (ACF, 2006) argue that l_{it} is collinear with the non-parametric terms used to approximate $\phi_{it}(\cdot)$, which may impair the identification of β_l . To avoid this collinearity problem, they propose a modification of LP's procedure in which β_l is not estimated in the first stage, but in the second stage using l_{it-1} as instrument. This procedure accommodates the more general case in which labor enters the demand function for material inputs. In this paper, I overcome the collinearity problem by using ACF's modification of LP's technique. ACF's critique also applies to OP's methodology and is addressed similarly.

this regression are estimates of Φ_{it} , denoted by $\hat{\Phi}_{it}$. Combining (3) and (6), the productivity shock can now be expressed as

$$\omega_{it}(\boldsymbol{\beta}) = \hat{\Phi}_{it} - \beta_0 - \beta_k k_{it} - \beta_l l_{it} - \beta_m m_{it},$$

where $\boldsymbol{\beta} = (\beta_0, \beta_k, \beta_l, \beta_m)'$. In the second stage of the procedure, I assume that the productivity shock follows a first-order Markov process. Accordingly, ω_{it} is generated by

$$\omega_{it}(\boldsymbol{\beta}) = g(\omega_{it-1}(\boldsymbol{\beta}), p_{it}) + \xi_{it},$$

where ξ_{it} is the innovation in the productivity shock, and p_{it} is the survival propensity score obtained from a probit regression of the exit choice on a third-order polynomial in $\{k_{it-1}, l_{it-1}, m_{it-1}\}$. To obtain an expression for ξ_{it} , I regress $\omega_{it}(\boldsymbol{\beta})$ on a third-order polynomial in $\{\omega_{it-1}(\boldsymbol{\beta}), p_{it}\}$. The residuals from this regression are the implied $\xi_{it}(\boldsymbol{\beta})$'s. Under the assumption that capital stock at time t is determined at time $t - 1$, $\xi_{it}(\boldsymbol{\beta})$ is orthogonal to k_{it} . Moreover, labor and material inputs at time $t - 1$ are uncorrelated with the innovation at time t . The corresponding moment conditions can be written compactly as

$$\mathbf{E} \left[\xi_{it}(\boldsymbol{\beta}) \begin{pmatrix} 1 \\ k_{it} \\ l_{it-1} \\ m_{it-1} \end{pmatrix} \right] = \mathbf{0}.$$

These four moment conditions are used to estimate the four coefficients in $\boldsymbol{\beta}$ using the GMM criterion.

The estimation technique of OP is similar. Instead of using m_{it} to “proxy” for the productivity shock ω_{it} , OP use the logarithm of investment (total capital expenditures divided by the investment deflator from the NBER-CES Manufacturing Industry Database), denoted by ι_{it} . Equation (4) is replaced by the optimal investment decision $\iota_{it} = f_t(\omega_{it}, k_{it}, l_{it})$. OP state conditions under which ι_{it} is strictly monotonic in ω_{it} . Provided these conditions hold, $f_t(\cdot)$ is invertible and the resulting expression for ω_{it} can be plugged into (3). The remainder of the procedure mirrors LP’s procedure, with ι_{it} and ι_{it-1} entering the third-order polynomials used to estimate Φ_{it} and p_{it} , respectively.

References

- Aarland, Kristin, James C. Davis, J. Vernon Henderson, and Yukako Ono, 2007, Spatial Organization of Firms: The Decision to Split Production and Administration, *RAND Journal of Economics* 38, 480-494.
- Abel, Andrew B., 1983, Optimal Investment under Uncertainty, *American Economic Review* 73, 228-233.
- Abel, Andrew B., Avinash K. Dixit, Janice C. Eberly, and Robert S. Pindyck, 1996, Options, the Value of Capital, and Investment, *Quarterly Journal of Economics* 111, 753-777.
- Ackerberg, Daniel A., Kevin Caves, and Garth Frazer, 2006, Structural Identification of Production Functions, *Mimeo*, UCLA.
- Ackerberg, Daniel A., C. Lanier Benkard, Steven Berry, and Ariel Pakes, 2007, Econometric Tools for Analyzing Market Outcomes, in: James J. Heckman and Edward E. Leamer (eds.), *Handbook of Econometrics*, Amsterdam: North Holland, 4171-4276.
- Aghion, Philippe, and Jean Tirole, 1997, Formal and Real Authority in Organizations, *Journal of Political Economy* 105, 1-29.
- Alesina, Alberto, and Roberto Perotti, 1996, Income Distribution, Political Instability and Investment, *European Economic Review* 40, 1202-1229.
- Arrow, Kenneth J., 1968, Optimal Capital Policy with Irreversible Investment, in: James N. Wolfe (ed.), *Value, Capital and Growth: Essays in Honour of Sir John Hicks*, Edinburgh, UK: Edinburgh University Press, 1-19.
- Bernanke, Ben S., 1983, Irreversibility, Uncertainty, and Cyclical Investment, *Quarterly Journal of Economics* 98, 85-106.
- Bernardo, Antonio E., Hongbin Cai, and Jiang Luo, 2001, Capital Budgeting and Compensation with Asymmetric Information and Moral Hazard, *Journal of Financial Economics* 61, 311-344.
- Bernardo, Antonio E., Hongbin Cai, and Jiang Luo, 2004, Capital Budgeting in Multidivision Firms: Information, Agency, and Incentives, *Review of Financial Studies* 17, 739-767.

- Bertrand, Marianne, and Sendhil Mullainathan, 2003, Enjoying the Quiet Life? Corporate Governance and Managerial Preferences, *Journal of Political Economy* 111, 2003, 1043-1075.
- Bertrand, Marianne, Esther Dufo, and Sendhil Mullainathan, 2004, How Much Should We Trust Differences-in-Differences Estimates?, *Quarterly Journal of Economics* 119, 249-275.
- Bloom, Nicholas, Stephen R. Bond, and John M. Van Reenen, 2007, Uncertainty and Investment Dynamics, *Review of Economic Studies* 74, 391-415.
- Burkart, Mike, Denis Gromb, and Fausto Panunzi, 1997, Large Shareholders, Monitoring, and the Value of the Firm, *Quarterly Journal of Economics* 112, 693-728.
- Caballero, Ricardo J., 1991, On the Sign of the Investment-Uncertainty Relationship, *American Economic Review* 81, 279-288.
- Caballero, Ricardo J., and Robert S. Pindyck, 1996, Uncertainty, Investment and Industry Evolution, *International Economic Review* 37, 641-662.
- Coval, Joshua D., and Tobias J. Moskowitz, 1999, Home Bias at Home: Local Equity Preference in Domestic Portfolios, *Journal of Finance* 54, 2045-2074.
- Coval, Joshua D., and Tobias J. Moskowitz, 2001, The Geography of Investment: Informed Trading and Asset Prices, *Journal of Political Economy* 109, 811-841.
- Crémer, Jacques, 1995, Arm's Length Relationships, *Quarterly Journal of Economics* 110, 275-295.
- Dewett, Todd, and Gareth R. Jones, 2001, The Role of Information Technology in the Organization: A Review, Model, and Assessment, *Journal of Management* 27, 313-346.
- Dixit, Avinash K., and Robert S. Pindyck, 1994, *Investment under Uncertainty*, Princeton, NJ: Princeton University Press.
- Foster, Lucia, John C. Haltiwanger, and Chad Syverson, 2008, Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review* 98, 394-425.

- Garmaise, Mark J., and Tobias J. Moskowitz, 2004, Confronting Information Asymmetries: Evidence from Real Estate Markets, *Review of Financial Studies* 17, 405-437.
- Giroud, Xavier, and Holger M. Mueller, 2010, Do Internal Capital Markets (Re-)Allocate Capital Efficiently? Evidence from Investment Spillovers across Plants, Mimeo, New York University.
- Goetz, Andrew R., and Christopher J. Sutton, 1997, The Geography of Deregulation in the U.S. Airline Industry, *Annals of the Association of American Geographers* 87, 238-263.
- Guiso, Luigi, and Giuseppe Parigi, 1999, Investment and Demand Uncertainty, *Quarterly Journal of Economics* 114, 185-227.
- Harris, Milton, and Artur Raviv, 1996, The Capital Budgeting Process: Incentives and Information, *Journal of Finance* 51, 1139-1174.
- Harris, Milton, and Artur Raviv, 1998, Capital Budgeting and Delegation, *Journal of Financial Economics* 50, 259-289.
- Hartman, Richard, 1972, The Effects of Price and Cost Uncertainty on Investment, *Journal of Economic Theory* 5, 258-266.
- Ivy, Russell L., 1993, Variations in Hub Service in the US Domestic Air Transportation Network, *Journal of Transport Geography* 1, 211-218.
- Jarmin, Ron S., and Javier Miranda, 2003, The Longitudinal Business Database, CES Working Paper.
- Landier, Augustin, Vinay B. Nair, and Julie Wulf, 2009, Trade-offs in Staying Close: Corporate Decision Making and Geographic Dispersion, *Review of Financial Studies* 22, 1119-1148.
- Leahy, John V., and Toni M. Whited, 1996, The Effect of Uncertainty on Investment: Some Stylized Facts, *Journal of Money, Credit and Banking* 28, 64-83.
- Lerner, Josh, 1995, Venture Capitalists and the Oversight of Private Firms, *Journal of Finance* 50, 301-318.
- Levinsohn, James, and Amil Petrin, 2003, Estimating Production Functions Using Inputs to Control for Unobservables, *Review of Economic Studies* 70, 317-341.

- Lichtenberg, Frank R., 1992, *Corporate Takeovers and Productivity*, Cambridge, MA: MIT Press.
- Maksimovic, Vojislav, and Gordon M. Phillips, 2002, Do Conglomerate Firms Allocate Resources Inefficiently across Industries? Theory and Evidence, *Journal of Finance* 57, 721-767.
- McDonald, Robert L., and Daniel R. Siegel, 1986, The Value of Waiting to Invest, *Quarterly Journal of Economics* 101, 707-727.
- Meyer, Margaret, Paul R. Milgrom, and John Roberts, 1992, Organizational Prospects, Influence Costs, and Ownership Changes, *Journal of Economics and Management Strategy* 1, 9-35.
- Michaels, Guy, 2008, The Effect of Trade on the Demand for Skill - Evidence from the Interstate Highway System, *Review of Economics and Statistics* 90, 683-701.
- Milgrom, Paul R., 1988, Employment Contracts, Influence Activities, and Efficient Organization Design, *Journal of Political Economy* 96, 42-60.
- Olley, Steven G., and Ariel Pakes, 1996, The Dynamics of Productivity in the Telecommunications Equipment Industry, *Econometrica* 64, 1263-1297.
- Ozbas, Oguzhan, and David S. Scharfstein, 2010, Evidence on the Dark Side of Internal Capital Markets, *Review of Financial Studies* 23, 581-599.
- Petersen, Mitchell A., and Raghuram G. Rajan, 2002, Does Distance Still Matter? The Information Revolution in Small Business Lending, *Journal of Finance* 57, 2533-2570.
- Petersen, Mitchell A., 2004, *Information: Hard and Soft*, Mimeo, Northwestern University.
- Petersen, Mitchell A., 2009, Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches, *Review of Financial Studies* 22, 435-480.
- Pindyck, Robert S., 1988, Irreversible Investment, Capacity Choice and the Value of the Firm, *American Economic Review* 78, 969-985.
- Pindyck, Robert S., and Andres Solimano, 1993, Economic Instability and Aggregate Investment, *NBER Macroeconomics Annual* 8, 259-303.

- Rajan, Raghuram G., Henri Servaes, and Luigi Zingales, 2000, The Cost of Diversity: The Diversification Discount and Inefficient Investment, *Journal of Finance* 55, 35-80.
- Scharfstein, David S., and Jeremy C. Stein, 2000, The Dark Side of Internal Capital Markets: Divisional Rent-seeking and Inefficient Investment, *Journal of Finance* 55, 2537-2564.
- Schoar, Antoinette, 2002, The Effect of Diversification on Firm Productivity, *Journal of Finance* 62, 2379-2403.
- Stein, Jeremy C., 2002, Information Production and Capital Allocation: Decentralized vs. Hierarchical Firms, *Journal of Finance* 57, 1891-1921.
- Syverson, Chad, 2004, Market Structure and Productivity: A Concrete Example, *Journal of Political Economy* 112, 1181-1222.
- U.S. Census Bureau, 1996, History of the 1992 Economic Census, Washington, DC: U.S. Government Printing Office.
- U.S. Census Bureau, 2000, 1997 Economic Census: Management of Companies and Enterprises, Washington, DC: U.S. Government Printing Office.

Table 1
Summary Statistics: Plants and Parent Companies

“All Plants” refers to all plants in the sample. “Eventually New Airline Route” refers to plants that are treated during the sample period, i.e., 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, capital stock, and capital expenditures are expressed in 1997 dollars using 4-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 shortest route and fastest means of transportation (car or plane) between the plant’s ZIP code and the ZIP code of headquarters (in minutes). Employees (firm-level) is the total number of employees of the parent company to which the plant belongs. Number of plants (firm-level) and number of industries (firm-level) are the number of LBD establishments and 3-digit SIC industries, respectively, of the parent company to which the plant belongs. All figures are sample means across all plant-year observations. Standard deviations are in parentheses. The sample period is from 1977 to 2005.

	All Plants [1]	Eventually New Airline Route [2]	No New Airline Route [3]
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)
Capital Expenditures	1,791 (26,610)	2,770 (12,423)	1,735 (27,202)
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)
Employees (firm-level)	9,203 (32,351)	22,211 (47,696)	8,452 (31,070)
Number of Plants (firm-level)	111 (431)	279 (730)	101 (405)
Number of Industries (firm-level)	8.15 (13.01)	17.41 (15.52)	7.62 (12.65)
Number of Observations	1,291,280	70,467	1,220,813

Table 2
Summary Statistics: Travel Time Reductions

“All” refers to all treated plants, i.e., 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 shortest route and fastest 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, either in minutes or as a percentage of the travel time before. All figures are sample means across all plants. The sample period is from 1977 to 2005.

	All	Indirect to Indirect	Indirect to Direct	Direct to Direct	Road to Flight
	[1]	[2]	[3]	[4]	[5]
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

Table 4
Dynamic Effects of New Airline Routes

Treatment (-12m, -6m) is a dummy variable that equals one if the plant-year observation is recorded between six and 12 months before the introduction of the new airline route. Treatment (-6m, 0m), Treatment (0m, 6m), Treatment (6m, 12m), Treatment (12m, 18m), Treatment (18m, 24m), Treatment (24m, 30m), and Treatment (30m +) are defined analogously. All other variables are defined in Table 3. 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.

Dependent Variable:	Investment	TFP
	[1]	[2]
Treatment (-12m, -6m)	-0.000 (0.003)	-0.001 (0.005)
Treatment (-6m, 0m)	-0.001 (0.002)	-0.001 (0.004)
Treatment (0m, 6m)	0.003 (0.003)	0.001 (0.005)
Treatment (6m, 12m)	0.005** (0.002)	0.006 (0.005)
Treatment (12m, 18m)	0.013*** (0.003)	0.012** (0.005)
Treatment (18m, 24m)	0.014*** (0.002)	0.020*** (0.004)
Treatment (24m, 30m)	0.014*** (0.003)	0.020*** (0.005)
Treatment (30m +)	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

Table 5
Small versus Large Reductions in Travel Time

($\Delta t \leq 30$ 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) are dummy variables indicating the magnitude of the travel time reduction. All other variables are defined in Table 3. 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.

Dependent Variable:	Investment	TFP
	[1]	[2]
Treatment \times ($\Delta t \leq 30$ min)	0.003 (0.004)	0.002 (0.009)
Treatment \times ($\Delta t > 30$ min and $\Delta t \leq 1$ hr)	0.002 (0.003)	0.004 (0.005)
Treatment \times ($\Delta t > 1$ hr and $\Delta t \leq 1$ hr 30 min)	0.006** (0.003)	0.012** (0.006)
Treatment \times ($\Delta t > 1$ hr 30 min and $\Delta t \leq 2$ hr)	0.014*** (0.003)	0.017*** (0.006)
Treatment \times ($\Delta t > 2$ hr)	0.015*** (0.002)	0.019*** (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

Table 6
Summary Statistics: Hub Openings

The table provides a list of airline hubs that are opened during the sample period. “Airline” is the name of the airline carrier. “City” is the city of the airport in which the new hub is opened (Federal Aviation Administration (FAA) 3-letter airport codes are in parentheses). “Year” is the year of the hub opening. The list is compiled from newspaper reports and airlines’ annual reports. The sample period is from 1977 to 2005.

Airline	City	Year
American Airlines	Dallas (DFW)	1981
Piedmont Airlines	Charlotte (CLT)	1981
American Airlines	Chicago (ORD)	1982
Piedmont Airlines	Dayton (DAY)	1982
Trans World Airlines	St. Louis (STL)	1982
Western Airlines	Salt Lake City (SLC)	1982
Piedmont Airlines	Baltimore (BWI)	1983
Republic Airlines	Detroit (DTW)	1984
Republic Airlines	Memphis (MEM)	1985
America West Airlines	Las Vegas (LAS)	1986
American Airlines	Nashville (BNA)	1986
Eastern Airlines	Philadelphia (PHL)	1986
Northwest Airlines	Detroit (DTW)	1986
Northwest Airlines	Memphis (MEM)	1986
Northwest Airlines	Minneapolis (MSP)	1986
Piedmont Airlines	Syracuse (SYR)	1986
United Airlines	Washington (IAD)	1986
American Airlines	Raleigh-Durham (RDU)	1987
Continental Airlines	Cleveland (CLE)	1987
Delta Airlines	Salt Lake City (SLC)	1987
American Airlines	San Jose (SJC)	1988
Braniff	Kansas City (MCI)	1988
American Airlines	Miami (MIA)	1989
Delta Airlines	Orlando (MCO)	1989
US Airways	Baltimore (BWI)	1989
US Airways	Charlotte (CLT)	1989
America West Airlines	Colombus (CMH)	1991
Trans World Airlines	Atlanta (ATL)	1992
United Airlines	Los Angeles (LAX)	1997
Midwest Airlines	Kansas City (MCI)	2000

Table 7
Summary Statistics: Airline Mergers

The table provides a list of airline mergers that are completed during the sample period and that account for at least one treatment during the sample period. The list only includes mergers that result in a merger of the airlines' operations. "Acquirer Airline" is the name of the acquiring airline carrier. "Target Airline" is the name of the acquired airline carrier. "Year" is the year in which the operations of the two airlines are merged. This list is compiled from newspaper reports, airlines' annual reports, and the Securities Data Corporation (SDC) database. The sample period is from 1977 to 2005.

Acquirer Airline	Target Airline	Year
North Central Airlines	Southern Airways	1979
Pan American World Airways	National Airlines	1980
Republic Airlines	Hughes Airwest	1980
Continental Airlines	Texas International Airlines	1982
People Express Airlines	Frontier Airlines	1985
Alaska Airlines	Jet America Airlines	1986
American Airlines	Air California	1986
Northwest Airlines	Republic Airlines	1986
Piedmont Airlines	Empire Airlines	1986
Trans World Airlines	Ozark Airlines	1986
Continental Airlines	New York Air	1987
Continental Airlines	People Express Airlines	1987
Delta Airlines	Western Airlines	1987
Braniff	Florida Express	1988
US Airways	Pacific Southwest Airlines	1988
US Airways	Piedmont Airlines	1989
Air Wisconsin	Aspen Airways	1990
Delta Airlines	Pan American World Airways	1991
Southwest Airlines	Morris Air	1994
AirTran Airways	ValueJet Airlines	1997
American Airlines	Reno Air	1999
American Airlines	Trans World Airlines	2001

Table 8
Hub Openings and Airline Mergers

Treatment (Hub) and Treatment (Merger) are dummy variables that equal one if the treatment dummy equals one and the new airline route is introduced by an airline in the same year as it opens a new hub or merges operations with another airline, respectively. Treatment (Other) is a dummy variable that equals one if the treatment dummy equals one and the treatment is not a hub or merger treatment as defined above. All other variables are defined in Table 3. 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.

Dependent Variable:	Investment	TFP
	[1]	[2]
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

Table 9
Alternative Control Groups

In columns [1] and [2], the sample is restricted to plants that belong to multi-unit firms consisting of more than one establishment. In columns [3] and [4], the sample is restricted to plants that are eventually treated—i.e., 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 one 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 3. 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.

Dependent Variable:	Multi-unit Firms		Eventually Treated Plants		Increase in Travel Time	
	Investment [1]	TFP [2]	Investment [3]	TFP [4]	Investment [5]	TFP [6]
Treatment	0.009*** (0.001)	0.012*** (0.003)	0.011*** (0.002)	0.010*** (0.003)	0.009*** (0.001)	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

Table 10
Alternative Measures of Productive Efficiency

Return on capital is the total value of shipments minus labor and material costs, divided by capital stock. Operating margin is the total value of shipments minus labor and material costs, divided by total value of shipments. All dollar values are expressed in 1997 dollars using industry-level deflators from the NBER-CES Manufacturing Industry Database and the Bureau of Economic Analysis. Return on capital and operating margin are industry-adjusted by subtracting the industry median across all plants in a given 3-digit SIC industry and year. TFP (Olley and Pakes) and TFP (Levinsohn and Petrin) are computed using the methodology of Olley and Pakes (1996) and Levinsohn and Petrin (2003), respectively, as described in Appendix C. All other variables are defined in Table 3. 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.

Dependent Variable:	Return on Capital	Operating Margin	TFP (Olley and Pakes)	TFP (Levinsohn and Petrin)
	[1]	[2]	[3]	[4]
Treatment	0.013*** (0.004)	0.009*** (0.003)	0.010*** (0.004)	0.013*** (0.003)
MSA-year	0.153*** (0.018)	0.181*** (0.069)	0.247*** (0.014)	0.197*** (0.017)
Firm-year	0.219*** (0.008)	0.309*** (0.018)	0.399*** (0.007)	0.355*** (0.008)
Age	0.014*** (0.003)	0.021*** (0.003)	-0.018*** (0.005)	-0.012*** (0.004)
Size	0.070*** (0.002)	-0.007*** (0.002)	0.054*** (0.002)	0.059*** (0.002)
Plant Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
R-squared	0.61	0.63	0.62	0.61
Number of Observations	1,291,280	1,291,280	940,064	1,081,893

Table 11
Headquarters' Time Constraints

In columns [1] and [2], headquarters' time constraint is measured as the number of employees at headquarters divided by the total number of plants of the company ("Managers/Plants"). In columns [3] and [4], headquarters' time constraint is measured as the number of employees at headquarters divided by the total distance (in miles) between headquarters and all of the company's plants ("Managers/Total Distance"). High time constraint is a dummy variable that equals one 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 constraint is defined analogously. All other variables are defined in Table 3. 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.

Dependent Variable:	Managers/Plants		Managers/Total Distance	
	Investment	TFP	Investment	TFP
	[1]	[2]	[3]	[4]
Treatment × High Time Constraint	0.012*** (0.002)	0.015*** (0.004)	0.013*** (0.002)	0.015*** (0.003)
Treatment × Low Time Constraint	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

Table 12
Soft- versus Hard-information Industries and Innovations in Information Technology

Soft- and hard-information industry are dummy variables that equal one if the plant operates in a soft- and hard-information industry, respectively, in the year prior to the treatment. The classification of industries into soft- and hard-information industries is adopted from Landier, Nair, and Wulf (2009). Pre 1986, Between 1986 and 1995, and Post 1995 are dummy variables that equal one if the plant-year observation lies within the specified time interval. All other variables are defined in Table 3. 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.

Dependent Variable:	Investment	TFP	Investment	TFP
	[1]	[2]	[3]	[4]
Treatment × Soft-information Industry	0.012*** (0.002)	0.017*** (0.004)		
Treatment × Hard-information Industry	0.007*** (0.002)	0.009** (0.004)		
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)	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

Table 14
Information Uncertainty at the Industry Level

High sales volatility is a dummy variable that equals one if the plant operates in a 3-digit SIC industry whose sales volatility lies above the median value across all 3-digit SIC industries in the year prior to the treatment. Sales volatility is the standard deviation of annual shipment growth across all plants in a given 3-digit SIC industry. High wage volatility is a dummy variable that equals one if the plant operates in a 3-digit SIC industry whose wage volatility lies above the median value across all 3-digit SIC industries in the year prior to the treatment. Wage volatility is the standard deviation of the hourly wage of production workers (production worker payroll divided by total production worker hours) across all plants in a given 3-digit SIC industry. Low sales volatility and low wage volatility are defined analogously. All other variables are defined in Table 3. 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.

Dependent Variable:	Investment	TFP	Investment	TFP
	[1]	[2]	[3]	[4]
Treatment × High Sales Volatility	0.013*** (0.002)	0.018*** (0.004)		
Treatment × Low Sales Volatility	0.006*** (0.002)	0.008** (0.003)		
Treatment × High Wage Volatility			0.011*** (0.002)	0.015*** (0.004)
Treatment × Low Wage Volatility			0.007*** (0.002)	0.010*** (0.003)
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