Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry*

Serguey Braguinsky, Carnegie Mellon University

Atsushi Ohyama, Hokkaido University

Tetsuji Okazaki, University of Tokyo

Chad Syverson, University of Chicago Booth School of Business and NBER

January 2014

Abstract

We explore how changes in ownership and managerial control affect the productivity and profitability of producers. Using detailed operational, financial, and ownership data from the Japanese cotton spinning industry at the turn of the last century, we find a more nuanced picture than the straightforward "higher productivity buys lower productivity" story commonly appealed to in the literature. Acquired firms' production facilities were *not* on average less physically productive than the plants of the acquiring firms before acquisition, conditional on operating. They were much less *profitable*, however, due to consistently higher inventory levels and lower capacity utilization—differences which reflected problems in managing the uncertainties of demand. When purchased by more profitable firms, these less profitable acquired plants saw drops in inventories and gains in capacity utilization that raised both their productivity and profitability levels, consistent with acquiring owner/managers spreading their better demand management abilities across the acquired capital.

* This research was funded in part by the Berkman Foundation at Carnegie Mellon University, Japan Society for the Promotion of Science (No. 20598825), and Initiative on Global Markets at the University of Chicago Booth School of Business. We thank Xavier Giroud, Jesse Shapiro, Bob Topel, and conference participants at the 2013 ASSA Meetings, the Stanford Conference on Japanese Entrepreneurship, the NBER Summer Institute and the FOM conference for comments.

Email: Braguinsky: sbrag@andrew.cmu.edu; Ohyama: ohyama@econ.hokudai.ac.jp; Okazaki: okazaki@e.u-tokyo.ac.jp; Syverson: chad.syverson@chicagobooth.edu.

1. Introduction

The influence of changes in corporate control of assets on productivity has been a focus of theoretical and empirical research for some time.[1] In principle, mergers and acquisitions can reallocate control of productive assets to entities that are able to apply them more efficiently. Besides increasing the productivity of the individual production units that are merged or acquired, a broader process of such reallocations can also lead to aggregate productivity growth. Such a mechanism therefore has the potential to explain patterns of productivity at both the micro and macro levels. Implicit in the story of this mechanism—though not often treated explicitly in the empirical work on the subject—is the notion that productivity growth occurs when changes in ownership and control put assets in more able managers' hands.[2]

Despite the comfortable intuition of this logic, previous research has not been fully conclusive about the effects of ownership and management turnover, particularly regarding the nature of any measured productivity growth but especially regarding the particular manners in which this growth is obtained. This reflects in part the inherent limitations of the data available in the earlier studies. For instance, this research could not cleanly distinguish between physical (quantity) productivity and revenue productivity. This distinction can be important (Foster et al., 2008). It is not particularly surprising, excepting bounded rationality or agency problems, that acquisition deals could yield expectedly profitable synergies. However, such between-firm synergies need not be tied to improvements in the efficiency with which producers convert inputs to outputs. For example, mergers or acquisitions may increase market power that leads to higher output prices for the merged firm. In the typical revenue-based productivity measures of the literature (separate price and quantity information is rarely available at the producer level), this would be reflected as a productivity gain even absent changes in technical efficiency. These and related measurement issues mean we are still limited in our knowledge of how turnover in asset ownership and management affects the level and growth of producers' efficiency levels.

In this paper, we seek to make progress on this front. A primary advantage of our effort is a data set that allows us to investigate the production and input allocation processes at an unusual

[1] See, for example, Lichtenberg and Siegel (1987), McGuckin and Nguyen (1995), Maksimovic and Phillips (2001), Jovanovic and Rousseau (2002), Bertrand and Schoar (2003), Rhodes-Kropf and Robinson (2008), and David (2012).

[2] The idea that managers or management practices—even independent of any considerations of ownership—shape differences in productivity across plants, firms, and even countries, is itself a focus of a separate, budding literature. Examples include Bloom and Van Reenen (2007 and 2010) and Bloom et. al (2013).

level of detail. We observe the operations, financial reports, management, and ownership of the universe of plants in a growing industry over the course of several decades (the Japanese cotton spinning industry at the open of the 20th Century). These data, which we describe in the next section, contain records in physical units of inputs employed and output produced at each plant in the years it operated as well as plant-specific output prices and wages and firm-level financial data. We have matched these production and financial data with business histories of the industry's firms to let us identify all major ownership and/or management turnover events and the personalities involved. These combined data let us measure directly how ownership and management turnover events were reflected in plants' physical productivity levels, profitabilities, prices, and other operational and financial metrics.

Our findings draw a more nuanced picture of the effects of ownership and management turnover than the straightforward "higher productivity buys lower productivity" story that has motivated much of the previous theoretical and empirical work. In our sample, acquired firms' production facilities were *not* on average any less physically productive than the plants of the acquiring firms before acquisition; both parties were equally adept at transforming physical inputs into physical outputs, at least conditional on operating. Acquired firms were much less *profitable* than acquiring firms, however. This profitability gap did not result from any output price differences between the firms. Instead, as we show, it reflected systematically lower unit capital costs among acquirers, coming from two sources: lower average inventory levels and systematically higher capacity utilization. When these better acquirers bought less profitable establishments, the acquired plants saw drops in inventories and gains in capacity utilization that raised both their productivity and profitability levels. The pre-acquisition equality in physical productivity between the acquired and the acquiring arose because, as we document below, acquired plants were newer and had more productive capital of younger vintages. This canceled out their capital utilization disadvantages in productivity terms.

Therefore ownership/management turnover in the industry is best characterized as "higher profitability buys lower profitability." More profitable companies took over firms with capital that was actually better, but that was being used suboptimally. The new management took control of this superior capital and, by improving the manner in which it was employed, raised the acquired plants' productivity *and* profitability.

As to the specific source of the better owners' and managers' advantage, the explanation most consistent with the data is that better firms have a superior ability to manage the vagaries of demand in the industry. (We describe just what this means in our context in the next section.) This explanation is consistent not just with the productivity and profitability levels and changes we observe, but also with the differences in inventory levels and capacity utilization. This link between demand management, productivity, and profitability is to our knowledge a new mechanism in the literature examining how management can affect business performance. We present below a simple theoretical framework of managerial time allocation that offers one possible mechanism through which this demand management difference might operate.

This ownership and management reallocation process helped drive considerable productivity growth in the industry. Between 1897 and 1915, industry TFP growth averaged an impressive 2.3 percent per year, while over 3/4 of industry capacity changed hands during our sample. And while acquirers were fairly concentrated—the asset reallocation process resulted in the emergence of several very large firms (we look more closely at these "serial acquirers" below)—what set the leading firms apart was not their market power (we show there was little during the sample) but rather the ability to acquire and fully utilize the most productive capital.

While we focus our analysis on a single industry case study to take advantage of the available data and unique setting, we believe our setting offers broader lessons. The mechanisms we discover here could easily operate in other industries, countries, and time periods; they might just be difficult to isolate in standard datasets. Certainly, the structures of ownership control and the scope of managers to influence outcomes in our setting are very much like the structures and scope that exist today. Furthermore, the data span a time of critical economic development and industrialization for Japan, which at the time was less than two decades removed from the completion of a difficult and often violent process of transition to modernity after 250 years of an isolated, traditionalist society. Information as detailed as our data is unusual even for producers in today's advanced countries, to say nothing of developing countries whose situation might be more similar to that of Japan at the time of our analysis. Hence, we believe that broader lessons regarding the development of an advanced industrial economy can be drawn from this study. By digging deep into the micro-evidence, we aim to complement past empirical work and provide fresh insights for further development of economic theory about resource reallocation.

2. Entry and Acquisitions in the Japanese Cotton Spinning Industry: Background Facts

The development of the Japanese cotton spinning industry in the late 19th and early 20th centuries has long fascinated economists because of its unique nature "as the only significant Asian instance of successful assimilation of modern manufacturing techniques" before World War II (Saxonhouse, 1971; 1974).[3] The historical circumstances surrounding this development made the story even more intriguing. Japan unexpectedly opened up to foreign trade in the 1860s after 250 years of autarky. Cotton yarn, in particular, experienced the combination of the largest fall in relative price from autarky to the free trade regime and the highest negative net exports (Bernhofen and Brown, 2004). But starting from the late 1880s, the domestic cotton spinning industry began a remarkable ascendance. Net exports turned positive for the first time in late 1896, and two decades after that Japan was exporting a sizeable fraction of its output while imports became negligible (see Figure 1).

[Figure 1 about here]

Figure 1 reveals that the development went through several stages. During the first stage, Japanese knowledge of the technology was rudimentary, and as a result spinning mills were small and had low productivity. In 1887, the industry included 21 firms, but the average equipment capacity was just 3,292 spindles and the average number of factory floor workers employed per day was 137.

The second stage, involving the explosive growth of the 1890s, was ushered in by two major technological breakthroughs: the switch to imported raw cotton, and the adoption of a newer type of cotton spinning machinery. By 1896, the total number of active firms in the industry had reached 63 (with 17 more in the process of being set up), with the average plant having a capacity of 12,789 spindles and employing 719 workers. Thus the number of firms tripled over the first decade of growth, average plant size almost quadrupled, and average employment per plant rose fivefold. Industry output in physical units increased 17 fold over during the same period (*Nihon Choki Tokei Soran*, Vol. 2, pp. 346).

Industry entrants of earlier cohorts that set up their production facilities before the major innovations of the 1890s found themselves stuck with older vintage machines. However, an important advantage some of them had developed by the time the technological breakthroughs happened was a superior ability to "manage sales." Since this will play an important role in

[3] To save some space, we present here a "bare-bones" sketch of these facts. More details can be found in Saxonhouse (1974). See also Ohyama, Braguinsky, and Murphy (2004) and Braguinsky and Rose (2009).

mergers and acquisitions analysis below, we dwell upon this in some detail here.

Japanese cotton spinners at the time generally faced a very competitive market (see, e.g., Saxonhouse, 1971 and 1977). Most of the yarn was purchased and distributed by trading houses based in the largest commercial centers of Osaka and Tokyo (Takamura, 1971, Vol. 1, pp. 322-328). The market power of even the largest cotton spinning firms was on par or below that of trading houses, so no producer could exercise much influence over the price at which its yarn was being sold (*ibid*., p. 325).[4] This does not mean, however, that the playing ground was equal for all firms. Especially during anticipated business downturns, large established trading houses often applied rationing, in which they would limit their purchases to reputable producers with whom they had a long-term relationship (Takamura, 1971, Vol. 2). Going outside of the network of reputable trading houses entailed risks of its own, as unscrupulous traders could renege on contracts or their promissory notes could bounce, failing to deliver real cash. We will see below that these problems were indeed quite severe, and that the most successful early entrants (who later became major acquirers in the mergers and acquisition market) managed these sales-related issues better than other firms early on.

This superior ability to manage sales may not have been crucial during the rapid expansion phase, but we show in Section 4 that it started playing a major role in firms' fortunes when the industry's development entered its third stage at the start of the 20th century. After driving out imports, the Japanese cotton spinning industry felt the limits of the market size for the first time. Once the Boxer Rebellion effectively shut down the Chinese market in 1900, the first major "overproduction crisis" in the industry was in full swing. Most of the following decade saw industry consolidation with little if any growth on the extensive margin but with a lot of firm-by-firm (and firm-by-outside investor) acquisitions of existing production facilities, the first one of which happened in 1898.

Figure 2 depicts the total capacity of several categories of plants from 1896-1920, our merger and acquisition analysis timeframe. During the first decade of the 20th century especially, almost all capacity growth among existing firms came through acquisitions. While entry and new

---

[4] Cotton yarn was also traded on the Osaka exchange. The gross transaction volume on the exchange was very large—sometimes several times larger than the amount of output—and the prices set there strongly influenced the prices trading houses were willing to pay even in seemingly isolated local markets (*ibid*., p. 327). Cotton spinning firms did take collective action to support prices by enacting output restriction measures during slow years. By their nature, however, these restrictions affected all firms uniformly and they were enforced by on-site inspections conducted by the All-Japan Cotton Spinners Association.

construction eventually resumed, acquisitions continued to play an important role.[5]

[Figure 2 about here]

The fact that acquisitions assumed such a prominent role in firm growth process so early on also seems at first glance to be at odds with the established theoretical view that investment by purchasing new capital should come before acquisitions (e.g., Jovanovic and Rousseau, 2002, who find support for the theory in U.S. data). However, the intuition behind the underlying theory is simply that new capital purchases do not involve fixed costs, while acquisitions do. This was less true in the early Japanese spinning industry. Because the industry had to import almost all its capital equipment from England at considerable financial and time costs, taking over existing plants at the right price was a potentially cheaper way for Japanese firms to expand.

These factors led to the consummation of 73 distinct acquisition deals involving 95 plants (some changed hands more than once) between 1898 and 1920. Fifteen more plants were consolidated under a single ownership in the deal that in 1914 created Toyo Cotton Spinning Company (Toyobo) from an equal merger of Osaka Cotton Spinning Company (Osaka Boseki) and Mie Cotton Spinning Company (Mie Boseki). All in all, 50 of the 78 plants (64 percent of plants and 76 percent of capacity) that were in operation in the industry in 1897, the year before the first acquisition took place, were subsequently acquired by another company at least once.

Several large firms emerged from this process, mostly through serial acquisitions. These were Kanegafuchi Cotton Spinning Company (Kanebo), Mie Boseki, Osaka Boseki (as already mentioned, the latter two competed an equal merger in 1914 to form Toyobo), Settsu Cotton Spinning Company (Settsu Boseki) and Amagasaki Cotton Spinning Company (Amabo; the latter two merged in 1918 to form Dainippon Boseki).[6] These five firms (which shrank to four after the 1914 merger and to three after the 1918 merger) went from owning 10 percent of the plants and 25 percent of industry capacity and output to 40 percent of plants and half of capacity and output over the 25-year period of our analysis (see also Figure A1 in Appendix C). This

[5] A lack of trust outside immediate family members who operate the business can make it difficult for superior firms in today's developing countries to increase their spans of control through acquisitions (Bloom, Sadun and Van Reenen, 2012). The Japanese cotton spinning industry avoided this problem because the large majority of its firms were set up and run as joint stock companies with easily transferable ownership. In Appendix B we present an example where a new CEO turned around a struggling firm by implementing a set of measures whose description reads amazingly similar to the script laid out by outside consultants for Indian firms in Bloom et al. (2013). How a functioning market for assets emerged so early in the process of economic modernization is a subject for a separate study; see Miwa and Ramseyer (2000) for some insights on this issue.

[6] All these firms were founded before the technological breakthroughs of the early 1890s. Mie Boseki was founded in 1880, Osaka Boseki in 1882, Kanebo in 1887, while Settsu Boseki and Amabo were both founded in 1889.

concentration of ownership could in principle be due to multiple factors, but as our empirical analysis below will show, it appears to be mostly due to their superior ability to manage sales and as a consequence improve the productivity and profitability of the plants they acquired.

3. Data

Our main data source is the plant-level data gathered on the annual basis by governments of various Japanese prefectures and available in historical prefectural statistical yearbooks.[7] For this paper, we have collected and processed all the available data between 1899 and 1920. Since the first acquisition of an operating plant in the industry happened in 1898, we added similar data for 1896-1898 using the annualized monthly data published in the "Geppo" bulletin of the All-Japan Cotton Spinners' Association. Our data thus cover 1896 to 1920. Saxonhouse (1971, p. 41) declares that "the accuracy of these published numbers is unquestioned."[8]

Our data contain inputs used and output produced by each plant in a given year in physical units. In particular, the data contain the number of days the plant operated, the average daily numbers of spindles in operation, and of factory floor workers (male and female separately), average daily wages for each gender, data on intermediate inputs such as the consumption of raw cotton, output of the finished product (cotton yarn) in physical units and its average count, and the average price per unit of yarn produced.[9] We observe which firm owns each plant at a given time, so we can see plant-level variables before and after ownership changes.

We match these plant-level data with financial data from semi-annual reports issued by the firms that owned the plants. Those reports, which we were the first to systematically digitize, contain detailed balance sheets and profit-loss statements as well as lists of all shareholders (with the number of shares they held) and executive board members. Some firm-level financial data were also published in the semi-annual publication "*Reference on Cotton Spinning*" ("*Menshi Boseki Jijo Sankosho*") which started in 1903. We use these data to supplement company reports

[7] Here we describe only the most important features of our data; a more detailed description is in Appendix A.

[8] We checked anyway. We found occasional, unsystematic coding errors as well as obvious typos, which we could often correct by comparing them with annualized monthly data from Geppo. In the vast majority of cases, however, the annual data in statistical yearbooks and the annualized monthly data did correspond very closely (any discrepancies were only a few percentage points). We dropped about 5 percent of observations where the annual data contained in government statistical reports could not be corrected.

[9] See Foster et al. (2008) and Syverson (2011) for the discussion of the importance of separating quantity and revenue productivity and the difficulties encountered when using conventional data containing sales and input expenditures but not inputs and output quantities. Atalay (forthcoming) similarly discusses the importance of separating quantities from expenditures when measuring inputs.

for privately held firms and in cases where company reports were missing.[10]

Several unique properties of our research variables need to be explained in some detail. First, cotton yarn is a relatively homogeneous product, but it still comes in varying degree of fineness, called "count."[11] To make different counts comparable for the purpose of productivity analysis, we converted various counts to the standard 20 count using a procedure detailed in Appendix A. Second, we used plant-year-specific female-to-male wage ratios to convert units of female labor to units of male labor.[12] Third, in addition to the number of spindles installed, we also have data on the actual number of spindles in operation for each plant-year. In other words, the data offer us the unusual ability to directly measure the flow of capital services at the plant level rather than to infer it from capital stocks or through the use of other proxies like energy use. This also allows us to measure capacity utilization rates. Finally, we follow Saxonhouse (1971 and 1977) and exclude intermediate inputs (raw cotton) when estimating the production function. As discussed by Saxonhouse, yarn production is essentially Leontief in raw cotton when input and output are measured in units of weight (the raw correlation between the two variables in our data is 0.95). As a practical matter, including raw cotton in a log-linear production function thus renders all other inputs economically and statistically insignificant. One can interpret our production function estimates as relating yarn output to capital and labor flow inputs, conditioning on the use of the physically necessary quantity of raw cotton.

4. Empirical Analysis

Table A1 in Appendix C presents year-by-year counts of acquired plants during our sample. On average, 4.3 percent of the industry's mills were acquired per year, with the afore-mentioned serial acquirers responsible for about 40 percent of all acquisitions.[13] These acquisition episodes form the base of our estimation sample.

---

[10] We checked the correspondence between the data in *Sankosho* and company reports whenever both sources were available and found a 100 percent match.

[11] The yarn count expresses the thickness of the yarn, and its number indicates the length of yarn relative to the weight. The higher the count, the more yards are contained in a pound of yarn. Thus higher-count yarn is thinner (finer) than lower-count yarn.

[12] Using female-to-male wage ratios to aggregate the labor input assumes that wages reflect the marginal productivity of each gender. All our estimates are robust to including the number of male and female workers separately in the production function estimations.

[13] This average acquisition rate is higher than the 3.9 percent acquisition rate for large U.S. manufacturing plants over 1974-1992 reported in Maksimovic and Phillips (2001) or the 2.7 percent in the LED plant sample from 1972-1981 used by Lichtenberg and Siegel (1987, Table 3).

4.1. Differences between Acquirers and Targets

We first use our detailed data to see, *before there were any acquisitions in the industry*, if there were systematic differences among firms that would eventually a) acquire other firms, b) be acquired, and c) exit without either acquiring or being acquired.[14] We compare these firms' plants along several dimensions: physical (quantity-based) productivity, accounting profitability, average output price, the number of days of the year the plant is operational, the average age of the plant's spindles, and the firm's age. To measure plants' physical total factor productivity levels (henceforth TFPQ, for quantity-based TFP), we estimate a Cobb-Douglas production function using the available data on all plants' output in physical units, labor and capital service flows, year dummies, and the plant's change in capacity from the previous year (as a control for possible adverse effects on output of adjustment costs of installing new equipment). The residuals from this production function reflect plants' TFPQ levels relative to the industry-year average.[15] To measure profitability, we calculate shareholders' return on equity; that is, we divide firms' profits by the amount of equity capital paid in by shareholders.[16] Equipment age is calculated as the current year minus the equipment vintage year, where vintage year reflects the composition of the years the plant's machines were purchased. Firm age, on the other hand, is always equal to the calendar year minus the year the firm was founded.[17]

Table 1 shows means and standard deviations of the aforementioned plant characteristics for each group of firms. We further separate plants of future target firms into those that started

[14] There were also a few surviving firms that did not participate in the acquisition market during our sample.

[15] Note, once again, that because we can measure capital service flows separately from capital stocks, a luxury typically unavailable in producer microdata, we can compute productivity either inclusive or exclusive of capacity utilization. (The former uses capital stocks as inputs. The latter uses capital service flows.) As will become clear below, the distinction between these is informative to explaining outcomes, so we compute TFPQ here using capital service flows, effectively measuring the plant's productivity conditional on it operating. We calculate capacity utilization—how often the plant actually operates during the year—separately and explore the two metrics jointly in our analysis. We also show below that our results are robust to alternative production function estimation methods.

[16] We do not have firm balance sheets data for 1896-97, but we do have these for subsequent years, so we will also measure profitability as return on total capital employed. See below.

[17] For example, if the plant's initially installed machines were purchased in year t and then the plant underwent an expansion during which the same quantity of new machines were purchased in year $t+k$, equipment age is calculated as the calendar year minus t until the year new machines are installed, after which it becomes the calendar year minus $[t+(t+k)]/2$, the average vintage age of machines (or the weighted average if the number of spindles installed later were different from the number initially installed). As the plant's capital stock includes also buildings and various elements of infrastructure, equipment (spindles) age adjusted for vintage this way makes the plants look younger than they actually are. Firm age, on the other hand, certainly makes those plants that had added new spindles (or scrapped old ones, which is also captured in our measurement) look older than they are. Equipment age thus provides the lower bound, and the firm age the upper bound, for the true overall plant age.

operating before 1892 (labeled "first cohort") and those that started operating in 1892 or later ("second cohort"), as the former are more likely to have older-vintage capital. The table includes only data from 1896-97—that is, before any acquisitions took place in the industry.

[Table 1 about here]

Looking across the table's top row to compare the average physical productivity levels across the groups of plants, we see that plants in future acquiring firms—at least conditional on the plant operating—are not more physically efficient than those in future acquired firms. Indeed, the most efficient group of plants is the second cohort of the acquired. (On the other hand, the ubiquitous result in the literature that exiting plants are less productive than continuing establishments is borne out in our data.)

This pattern is reversed when we look at profitability. The most profitable establishments (significantly so) are those in firms that will be acquirers. Plants in the first cohort of target firms are the second most profitable, and exiting and second-cohort acquired plants follow up the rear.

The numbers in the table's third row indicate these profitability gaps are not tied to differences in the prices the plants fetch for their output. All firms earn more or less similar price per unit weight of output. (Acquiring plants' average price is slightly higher, though none of the differences in the table are statistically significant at conventional levels. Furthermore, when we adjust for the average count of the plants' yarn, these differences become even smaller.) This result, which we will see repeatedly below, supports what we know about the institutions of the industry's output market: pricing did not reflect large market power differences across industry producers and is unlikely to contribute to firm- or plant-level outcomes examined in this paper.

The days-in-operation and age comparisons at the bottom of the table offer insight into the possible sources of the productivity and profitability patterns. Second-cohort acquired plants are more productive than other plants, yet less profitable. Their productivity advantage is tied to the fact that they have significantly newer capital (whether measured by equipment or firm age), as reflected in the table's final rows.[18] A hint at why their productivity advantage did not yield a profitability advantage can be seen in the comparison of plants' average days in operation. Second-cohort acquired plants only operated about 80 percent of the time that plants in future

[18] In Appendix D, we use additional data on firms' orders of specific pieces of capital equipment to measure how the machines' technical specifications evolved over time. We find clear evidence of pre- and post-early 1890s differences (not sensitive to the choice of a specific cutoff year around this general timeframe) in technological capabilities along multiple dimensions: spindle rotation speed, spindles per frame, the yarn quality for which the machines are calibrated, and the ability to handle multiple yarn counts and cotton types.

acquiring firms did. Plants that were to exit the industry had the worst of both worlds: their capital was old (not only were they the oldest firms, their equipment and firm ages were almost the same, indicating they did almost no upgrading of their equipment), and their factories were often idle. They were unproductive and unprofitable as a result.

4.2 Changes in Productivity and Profitability within Acquired Plants

The analysis in the previous subsection revealed some systematic pre-acquisition differences between acquiring and target firms. In this subsection, we investigate whether and how acquired plants' attributes change when they are taken over by acquiring firms.

Acquisition is, of course, not an exogenous occurrence. As is typical in this literature, we do not have a source of random or even quasi-random assignment to acquisition, so interpreting any of the plant performance changes around acquisition as isolating causal effects should be done with caution. However, our specifications control for the most obvious sources of potential biases by controlling for acquired plant fixed effects (removing any effects of selection into acquisition on persistent plant attributes) and any common movements with various control groups (the acquiring firms' existing plants, for example). We are relying for causal inference in part on the assumption that the causal effect of acquisition creates a discrete change in attributes surrounding the event, whereas any performance trends that might lead to selection into acquisition would be either common to the control plants (and thus partialled out in our control group specifications) or gradual enough to be distinguished from the more discrete direct effect. To that end, we show in Appendix K that there are no obvious pre-trends in acquired plants' relative performance, while at the same time there is a noticeable change in the trajectory of certain performance measures at the time of acquisition.

To measure the changes in acquired plants' attributes, we estimate specifications that regress these attributes on three sets of time dummies defined around each acquisition event: a "late pre-acquisition" dummy that equals 1 for the two years immediately preceding the acquisition and zero otherwise, an "early post-acquisition" dummy that equals 1 for the first three years after the acquisition and zero otherwise, and a "late post-acquisition" dummy that equals 1 for all subsequent post-acquisition years after the first three and zero otherwise. The omitted category therefore includes the period at least three years prior to the acquisition. (We exclude the acquisition year itself from the regression because acquisitions often happen midyear, making it hard to attribute outcomes solely to the acquirer or the acquired.) We include plant fixed effects in the specifications, so the coefficients on the time dummies reflect within-plant changes in attributes. We also include calendar year fixed effects to remove any systematic changes in attributes over the sample and acquisition fixed effects (which are absorbed in plant fixed effects for the majority of plants that were acquired only once, but allow us to control for possible differences in circumstances of acquisition events for plants acquired multiple times).

Thus our estimating equations have the general form:

$$y_{it} = \alpha_0 + \beta_1 lbA_{it} + \beta_2 eaA_{it} + \beta_3 laA_{it} + \eta_i + m_A + \mu_t + \varepsilon_{it}, \quad (1)$$

where y_{it} is the attribute of plant i in year t; lbA_{it} is the "late before acquisition" dummy; eaA_{it} is the "early after acquisition" dummy; laA_{it} is the "late after acquisition" dummy; η_i is a plant fixed effect; m_A is an acquisition episode fixed effect; μ_t is a year fixed effect; and ε_{it} is the error term.

The first numerical column of Table 2 shows the results for TFPQ. Rather than first estimate physical TFP with a production function regression and then use the residual as the left-hand-side variable in (1), we perform the equivalent one-step estimation by using the plant's logged output as the dependent variable and adding the explanatory variables from the production function to the right hand side of (1): the plant's logged number of composite worker-days (the sum of male and female workdays, weighted by the relative plant-level ratio of female to male wages), its spindle-days in operation (flow of capital services), and the change in log plant capacity from the prior year (to control for any equipment installation adjustment costs).

[Table 2 around here]

The results indicate that in the first 3 years after acquisition, acquired plants' TFPQ levels are about 4 percent higher but not statistically different from their levels in the pre-acquisition years. In subsequent years, however, the TFPQ of acquired plants rises more than 13 percent above their pre-acquisition baseline, and we can reject equality of the early and late post-acquisition dummies at the 1 percent level. Thus acquired plants' TFPQ levels do improve considerably following an acquisition, although it takes time for this to manifest itself fully. We have also estimated a regression similar to (1) with the full set of pre- and post-acquisition year dummies and confirmed that there is no discernible pre-acquisition trend in TFPQ, while there is a clear upward trend after acquisitions, becoming particularly pronounced after the first 3 post-acquisition years (see Appendix K).

Table 2's second column looks at acquired plants' profitability around acquisition episodes. Unfortunately, we cannot directly evaluate plant-level changes in profitability that are analogous to the cross sectional comparisons in Table 1, for the obvious reason that there are no separate post-acquisition firm profit accounts. We work around this issue by constructing a measure of plant-level net operating surplus, computed as the difference between the net value of cotton yarn produced by the plant and plant labor and capital costs (see Appendix E for details). We then divide this by the sum of shareholders' capital (equity and retained earnings) and interest-bearing debt, which in case of multiple plant firms is assigned to each plant in proportion to the plant's installed capacity (number of spindles). We call the resulting measure "plant-level return on capital employed"—"plant ROCE" for short—and we use this measure (winsorized at the top 2.5 percent) to compare plant-level profitability before and after acquisition periods.[19]

The ROCE of acquired plants increases in the first 3 years after acquisition by an average of about 4.5 percent, a difference from the excluded pre-acquisition period that is significant at the 10 percent level. ROCE rises further in subsequent years, to a long-run gain of over 7 percent. Hence, both the onset and share of long-run gains in profitability appear earlier than for TFPQ.

Finally, to see if changes in plant-specific prices contributed to profitability changes, we estimate (1) using as the dependent variable the (logged) plant-specific price, divided by the main count of yarn produced by the plant to adjust for quality differences. The results, in Table 2's third column, indicate that post-acquisition prices are statistically and economically indistinguishable from pre-acquisition prices. Thus again the source of improved profitability over and above TFPQ improvement is not related to plants charging higher prices.

We also tested whether these changes within acquired plants are systematically related to the attributes of the acquiring firm. While acquiring firms could be demarcated along a number of dimensions, a natural one is whether they were one of the "serial acquirers" we discussed in Section 2. Thus we repeated the specifications in estimation equation (1), but while limiting the sample to only acquisitions by one of the five serial acquirer firms. The results are in last 3 columns in Table 2. Qualitatively, the patterns are similar, but more pronounced. In particular, acquisitions by serial acquirers correspond to long run improvements in acquired plants' physical

[19] As shown in Appendix E, our constructed plant ROCE is highly correlated with firm-level ROCE data in years preceding acquisition events (that is, when we have independent accounting data on both acquired and acquiring firms). The raw correlation between the two measures is about 0.74, and with the exception of extreme tails, the overall distribution fit is quite good too (see Figure A.2 in Appendix E).

TFPQ of more than 24 percent ($e^{0.217}$= 1.242), while ROCE increases by almost 15 percentage points. The point estimates on the price changes are larger than for all acquisitions, but *t*-tests fail to reject at conventional confidence levels the hypothesis that the coefficients on either of the post-acquisition dummies are equal to the pre-acquisition dummy coefficient.

Overall, the within-plant results in Tables 3 indicate that acquired plants see growth in both their TFPQ and profitability levels after acquisition, though profitability growth occurs sooner. Moreover, both of these changes are larger for plants that are acquired by the most prolific of acquiring firms.

4.3 Changes within Acquisition Episodes

We also look at productivity and profitability changes from before to after acquisition events in a slightly different way, by comparing acquired plants to the incumbent plants of acquiring firms. This in effect uses the incumbent plants as a control group. We lose some data as a result of this because in 37 acquisitions the acquirer came from outside the industry and hence had no incumbent plants. Additionally, timelines of available data on some incumbent plants were too short to be usable. Therefore, the exercise here is limited to 49 acquisitions. The benefit is that this within-acquisition approach allows us to explicitly compare plants' productivity and profitability levels and changes while controlling for specific circumstances surrounding each acquisition by including acquisition-year fixed effects.[20]

The specification is as follows:

$$\bar{y}_{it} = \alpha_0 + \beta_1 AA_{it} + \beta_2 Aquired_{it} + \beta_3 Aquired_i \times AA_{it} + m_{it} + \mu_t + \varepsilon_{it} \ , \qquad (2)$$

where $\bar{y}_{it}$ is the outcome variable of plant *i* at time *t* if it is an acquired plant, while the outcome variables of incumbent plants are collapsed to $\bar{y}_{it} = \frac{1}{\#m_A}\sum_{j \in m_A} \omega_j y_{jt}$, where m_A denotes the particular acquisition case in which plant i was acquired and $\#m_A$ is the number of incumbent plants in acquisition m_A. Thus, $\bar{y}_{it}$ in the case of incumbent plants is the weighted average of outcomes of those plants within the given acquisition. The variable AA_{it} is a dummy equal to 1 if acquisition m_A happened prior to year *t* and zero otherwise, while the variable $Acquired_i$ is equal

[20] To avoid problems stemming from the fact that plants previously acquired by serial acquirers are already "incumbent" plants when another acquisition happens (which can be as early as in the same year), we only label a previously acquired plant as an incumbent after being under the new ownership for five years. The results presented below are not sensitive to other reasonable cutoffs or to leaving only serial acquirers' originally owned plants in the "incumbent" category.

to 1 if plant i is purchased in acquisition case m_A and zero otherwise; m_{it} is the acquisition-year fixed effect, and μ_t is the calendar year fixed effect included when the outcome variable is ROCE to control for inflation (it is not included in the regression with TFPQ as the outcome variable because, by construction, TFPQ is the residual from the production function estimated inclusive of year dummies). In the main text, we assign weights $\omega_j = 1$ to all incumbent plants in a given acquisition m_A, which allows us to interpret coefficients β_1, β_2, and β_3 similarly to the interpretation given to them in standard difference-in-difference estimations. In particular, $\hat{\beta}_3$ reflects the post-acquisition difference-in-difference between acquired and incumbent plants of acquiring firms, accounting for acquisition-case effects. We limit the sample time period to 4 years before and 8 years after the acquisition event, but reasonable alternative timeline cutoffs produce similar results.[21]

The estimation results are presented in Table 3. The first two columns of numbers reflect TFPQ and plant ROCE results (respectively) for all acquisitions, while the latter two columns look only at acquisitions by the five serial acquirers.

In both TFPQ specifications, the estimates of the interaction coefficient β_3 are positive and statistically significant at the 1 percent level. The post-acquisition improvement of TFPQ of acquired plants (this time relative to incumbent plants of the acquirer) averages about 10 percent for all acquisitions and more than 13 percent for acquisitions by serial acquirers. In addition, the acquired plant dummy coefficients are small and statistically insignificant in both samples, confirming the Table 2 result that there is no systematic difference between the physical TFP of acquired and incumbent plants prior to acquisitions (confirmed in year-by-year estimations presented in Appendix K).

[Table 3 about here]

In the profitability regressions, β_3 is also positive and statistically significant. Profit rates of acquired plants rise by over five percentage points relative to acquiring firms' plants in the whole sample and by more than seven percentage points in acquisitions by serial acquirers. Here, acquired plant dummy coefficients are negative, reflecting the profitability deficit of acquired

[21] We also estimated equation (2) employing kernel weights obtained from the Mahalanobis distance measure where acquired and incumbent plants are matched on plant size, age and location, and also using a standard difference-in-difference procedure ignoring acquisition-based matching altogether. The results of these estimations were very similar to those presented in Table 3 (see Tables A5 and A6 in Appendix F).

firms before acquisition.[22]

These results further reinforce what we document above: acquisition was accompanied by growth in the acquired plants' productivity and profitability levels. We see here that this is true relative not only to the acquired plants' own levels before the acquisition, but also relative to changes within incumbent plants owned by their acquiring firms.

While matching by acquisition cases seems to be the most natural approach in our context, we did explore other matching possibilities as well. Specifically, we also matched acquired plants on pre-acquisition characteristics and also on pre-acquisition productivity trend with a control group of plants that were either never acquired or, at least, not acquired within the time window during which we compare them to acquired plants. The results of these estimations were very similar to the ones presented here. See Tables A8 and A9 in Appendix F.

4.4 Profitability Differentials: Decomposition

We have so far found that incumbent plants of acquiring firms have higher profitability (ROCE) than future acquired plants prior to acquisition, but not greater physical productivity (TFPQ). After being purchased, however, acquired plants improve in terms of not just ROCE but also TFPQ. Moreover, neither the pre-acquisition difference in profitability between acquiring and acquired firms nor the post-acquisition improvement in profitability of acquired plants seem to be driven by price differentials that could be attributed to market power. When considered together, these findings present a sort of puzzle: if it is neither prices nor productivity, what makes incumbent plants more profitable than acquired plants before acquisition? And how do acquisitions by more profitable firms lead to improved TFPQ in acquired plants?

To begin exploring this puzzle, we decompose the pre-acquisition profitability differential between acquiring and acquired firms as well as the pre- to post-acquisition profitability changes for acquired plants into their various components. This lets us isolate the most important factors driving profitability differences.

[22] We also estimated a regression similar to (2) with the outcome variable being the count-adjusted plant output price relative to the industry-year average. As in the previous subsection, we did not find any big differences before and after acquisitions in either acquired or incumbent plants, with a possible exception of plants bought by serial acquirers, where β_3 was estimated to be 0.064 with a p-value of 0.10. The coefficient β_2 on the post-acquisition time dummy is economically and statistically indistinguishable from zero, however, indicating that somewhat higher post-acquisition prices of acquired plants are not shared with incumbent plants. It thus appears that any increase in post-acquisition prices of acquired plants might be a reflection of unobserved quality improvement rather than market power, which would presumably be shared by both acquired and incumbent plants.

We first express a plant's ROCE as the net value of cotton yarn produced and the plant's labor and capital costs (all per unit of capital assets):

$$\frac{\pi_i}{C_i} = \frac{(1-\upsilon)Y_i}{C_i} - \frac{w_i L_i}{C_i} - \frac{R_i}{C_i}. \qquad (3)$$

Here, π_i is plant *i*'s operating income. Y_i denotes the value of its output, and υ is the fraction of intermediate input and non-labor operational costs in the value of output (e.g., the costs of raw cotton, energy, etc.). Plant wage costs are $w_i L_i$, R_i is capital cost, and C_i is plant *i*'s share of its owning firm's assets (the sum of shareholders' capital and interest-bearing debt). The details of variable construction are described in Appendix E. In a nutshell, we use plant price and output data to obtain Y and plant-level data on the number of worker-days and daily wages to obtain wL. Capital cost is calculated as the sum of depreciation of fixed capital and interest payments on borrowed capital, with both depreciation and interest rates assumed to be the same for all plants, as is the parameter υ (these values are estimated from the available firm-level and industry-wide data). All nominal values are divided by the consumer price index to account for inflation.

We present the results of decomposition (3) in Table 4. There are three panels, each corresponding to the decomposition of a particular profitability differential. The top panel compares plants of acquired firms ("acquired plants") and those of their future acquirers ("incumbent plants") for up to 4 years prior to acquisition events. The bottom two panels compare acquired plants before and after acquisitions, with the post-acquisition years split as in the regressions above: the middle panel looks at the first 3 years immediately following the acquisition, and the bottom panel looks at the subsequent post-acquisition years.

The top panel of Table 4 shows that incumbent plants' 3.3 percentage point (66 percent) ROCE advantage over acquired plants is mostly explained by the net output value to total assets ratio (the first term on the right hand side of (3)).[23] This ratio is on average 4.1 percentage points higher in acquiring plants. Wage costs per unit of assets are actually higher in incumbent than in acquired plants, reducing the ROCE difference. Capital costs are similar, though statistically smaller for incumbents. The relative similarity of capital costs reflects similarity in the ratios of

[23] The ROCE differential in the top panel of Table 6 is similar in magnitude to the coefficient on acquired plant dummy obtained from the specification in Table 5, where it was 0.031. The same holds when we compare all other differentials below with the corresponding regression coefficients. This similarity between decompositions results using raw data and regression analyses is reassuring. Some discrepancy is to be expected, of course, as the regression analyses are conducted absorbing year and acquisition effects as well as plant fixed effects.

fixed capital to total assets and interest-bearing debt to shareholders capital (not shown).

The bottom two panels of Table 4 show the decomposition of acquired plants' ROCE changes around acquisition episodes. ROCE improves by 3.9 percentage points (64 percent) in the first three post-acquisition years and by 4.6 percentage points (75 percent) in the longer run. As with the cross-sectional differences, the bulk of the profitability changes came from growth in acquired plants' ratios of net output value to total assets.

[Table 4 about here]

The centrality of net output value—essentially, gross margin—in explaining profitability differences leads naturally to a second decomposition. We break the net output to capital ratio into a product of a) price, net of intermediate input and non-labor operation cost per unit output; b) total input of capital and labor services per total assets; and c) TFPQ. Taking logs, we obtain

$$\log\left(\frac{\psi Y_i}{C_i}\right) = \log\left(\psi p_i\right) + \log\left[\frac{\exp\left(\hat{Y}_i\right)}{C_i}\right] + TPFQ_i, \qquad (4)$$

where $\psi \equiv (1-\upsilon)$ denotes the unit price margin (common to all plants and firms), p_i is plant i's output price, $\hat{Y}_i$ is the total combined input of logged capital and labor in the production function (and thus equals logged physical output predicted from the production function regression), and $TFPQ_i$ is the residual from the production function. We use this expression to measure the relative contribution of these three components to the net value of output per unit of shareholders' capital. These decompositions are presented in Table 5.

[Table 5 about here]

Consistent with the regression analyses in the previous section, price and TFPQ differentials contribute relatively little to the stark profitability differences between acquired and incumbent plants before the acquisition (top panel). Most of the difference is instead driven by the ratio of predicted output (or combined total inputs) to total assets, $\exp\left(\hat{Y}_i\right)/C_i$.

This gives us a hint to resolving the puzzle that motivated our decomposition exercise: while prices and physical productivity of plants conditional on operating are similar, the numbers in the top panel imply that for the same amount of capital invested by shareholders, incumbent plants somehow manage to mobilize over 25 percent more of their combined inputs toward production than do acquired plants in pre-acquisition years. (Recall that our TFPQ metric is estimated with respect to capital and labor inputs measured as flows of services; that is, we are measuring the plant's productivity conditional on it operating.)

Acquired plants' changes in gross margin per unit assets in the bottom panels of Table 5 indicate a larger role for TFPQ in profitability growth, at least in the long run. As in Table 3, TFPQ initially improves modestly, again with growth in input use per unit assets explaining most of acquired plants' gains in profitability. In the long run, though, TFPQ growth accounts for 11 points of a 28 log point improvement in net output value per unit assets, while the impact of the input per asset ratio became small. As shown in Appendix G, this is due to a big increase in the retained earnings component of acquired plants' shareholders capital and can thus be interpreted as a consequence of years of accumulated high profits. In contrast to the regression analysis, the contribution of the price margin in the decomposition approach is relatively large and statistically significant.

What is it that allows acquiring firms to systematically mobilize and put to direct use in production process a greater share of their total assets than acquired firms? Table A.12 in Appendix G further decomposes combined input to total assets ratio $\exp\left(\hat{Y}_i\right)/C_i$ into the ratio of combined inputs to available capacity (reflecting capacity utilization rates) and the ratio of capacity to total assets. It turns out that incumbent plants' advantage in input mobilization is almost entirely explained by their higher capacity utilization rates—the rates of capital services extracted from the installed machine capacity and the corresponding rates of labor services applied to this capacity. This ratio jumps 9 percent in acquired plants in the first few years after acquisition and by more than 16 percent in the long run. Thus acquisitions immediately lead to higher rates of employment of available capital resources (and correspondingly higher employment of labor, as can be seen from the increases in total labor cost per total assets for acquired plants in Table 4). Ratios of capacity to total assets, on the other hand, are statistically indistinguishable in pre-acquisition years between incumbent and future acquired plants.

Putting the results of these decompositions together, we find that pre-acquisition profitability (ROCE) differences prior to acquisition arise mostly because acquiring firms' plants utilize their available capital and labor more intensively than do acquired firms' plants. This also contributed the bulk of short-term profitability improvement in acquired plants post-acquisition. In the longer run, on the other hand, TFPQ improvements become more prominent in raising the profitability of acquired plants.

4.5 The Link from Profitability to Productivity: The Role of Demand Management

Why were stronger firms able to utilize their productive capacity so much more than weaker firms? In this section we seek an answer by looking at the ability of stronger companies to manage the industry's inherent demand variations better.

As we discussed in Section 2, a lack of price differentiation does not mean that output-market conditions were equivalent across firms. To quantitatively explore possible differences in firms' demand-facing operations, we investigate patterns in plants' finished goods inventory and accrued revenues on delivered output (that is, the payment for which is in arrears). We choose these metrics because both may indicate that the plant is having difficulty finding buyers in a timely manner or finding buyers who can be relied upon to disburse payments on time. This in turn may explain capital utilization differences, which under this interpretation reflect poor management of matching production to demand or difficulty in finding (reliable) buyers. Indeed, anecdotal evidence from company histories suggests that inventories and payment arrears were intrinsically linked to utilization in the industry, as firms would often halt production as unsold yarn and accrued revenues piled up, and would resume only after the gridlock had cleared.[24]

In Table 6 we present producers' ratios of period-end finished goods inventories, accrued revenues and the sum of these ("unrealized output" for short) to the value of their output over the period, split by different categories of plants in the same way as in the previous sections.[25] Incumbent plants' ratios of unrealized output to their output value were about 60 percent lower than the same ratios of acquired plants. Post-acquisition, acquired plants' inventory to output values ratios fell by 60 percent in the first three years, and by further 10 percent after that. Within-acquisition comparisons of acquired and incumbent plants (not shown) yield very similar patterns.

[Table 6 about here]

There are many possible sources of cotton spinning firms' abilities to manage demand. While many of these are difficult to quantify, one important factor already mentioned in Section

[24] We also looked at stocks of intermediate goods and unfinished products but found no evidence that those systematically influenced outcomes.

[25] Finished goods inventories and accrued revenues are positively correlated in the data, but the correlation is not that high (about 0.22 for both incumbent and acquired plants). There may be a certain substitutability between the two, as having difficulty finding reputable buyers in a timely fashion might lead a firm to reach out to lesser buyers who are more likely to fall into arrears. Therefore, the total "unrealized output" seems to be the best metric to measure demand-facing operations efficiency. Nevertheless, the evidence presented in Table 6 indicates that all the three metrics paint a similar and consistent picture.

2 was that in low-demand times, major trading houses rationed their demand to certain producers. This suggests a mechanism similar to positive assortative matching, where long-term relationships established between reputable industry producers and prominent traders allowed those producers to sustain more consistent operations, resulting in the lower inventories and higher utilization levels observed above.

To explore this possibility quantitatively, we used the 1898 edition of *Nihon Zenkoku Shoukou Jinmeiroku*, a nationwide registry of names of traders and manufacturers, to extract the names of individuals likely to play the most prominent role in cotton spinners' output markets. This yielded a list of 154 individuals.[26] We then matched these individuals to the lists of board members and top 10-12 shareholders of the 67 firms for which we have company reports in 1898 (this is 90 percent of firms operating that year). Of a total of 1,197 board members and top shareholders, 128 were in the list of the 154 most prominent traders described above. Of the 67 firms, 33 had at least one prominent trader among its board members and top shareholders. We create a "trader network" indicator equal to 1 if the firm is one of these 33 or one of two more firms for which firm histories (Kinugawa, 1964) clearly indicated connectedness to major traders at their inception (we refer to these as "in-network" firms) and 0 otherwise (these are "out-of-network" firms).

We then tested whether a producer's relationship to trading houses in this way is reflected in the performance metrics we explored above. Table 7 presents the means of ROCE and ratios of unrealized output to the value of output (both taken directly from firms' accounts), as well as spindle utilization rates and output prices from plant-level data for in-network and out-of-network firms. (Figures A3-A6 in Appendix H plot the corresponding distributions.) Since our in- or out-of-network classification is based primarily on the 1898 shareholders and board composition data, we limit our attention to years 1898-1902 to obtain a reasonable number of observations while not going too far forward, as board and shareholders (as well as traders' importance) of course changed over time.

[Table 7 about here]

The results in Table 7 show that ROCE of "in-network" firms was much higher than that

[26] These individuals fit into groups meeting one of three criteria. One group included 98 cotton yarn and yarn-related traders across Japan who paid more than 50,000 yen in operating tax that year. A second group included 25 individuals listed as board members of the 4 largest incorporated cotton yarn-related trade companies (Naigaimen, Nihon Menka, Nitto Menshi and Mitsui Bussan). Finally, the third group includes the 31 board members and traders registered at the Osaka cotton and cotton yarn exchange.

of "out-of-network" firms (recall that 1898-1902 was when demand constraints first presented themselves very strongly in the young industry's history). The trader network indicator is also associated with a large drop (on the order of 40 percent) in the average of plants' ratios of unrealized to produced output. The distributions of both ROCE and unrealized output ratios of in-network firms are basically rightward shifts of the corresponding distributions of out-of-network firms (see Figures A3 and A4 in Appendix H). In-network firms also have higher capacity utilization and prices, although these differences are smaller than in unrealized output ratios and are not equally pronounced across the whole distributions (Figures A5 and A6 in Appendix H). In particular, the distributions of prices of in- and out-of-network plants are quite similar up to the utmost right tail, where there are no more plants of out-of-network firms but still some plants of in-network firms selling at very high prices.

Overall, these results suggest that close relationships between industry producers and prominent traders may have allowed "matched" producers to manage demand fluctuations more effectively, particularly with regard to being able to operate with lower average inventory levels and often at greater capacity utilization levels as well. Notably, in-network firms were also more likely to acquire other firms in the future (the sample probability of being a future acquiring firm is 0.79 for an in-network firm as opposed to 0.21 for an out-of-network firm). Hence, relationships with traders' networks (along with perhaps other demand management mechanisms) can explain why an initial profitability gap existed, and why it was closed by acquisition. But of course we observed TFPQ gains upon acquisition too (though there was no prior gap). This is consistent with the profitability story above if demand management is correlated with broader managerial ability that raised operational efficiency as well. We explore this connection in Section 5 below.

4.6. Robustness

As already mentioned, we have conducted several robustness checks. We relegate the details to Appendix F for the sake of parsimony, but we briefly describe the exercises here.

Our benchmark results above use TFPQ estimates obtained as residuals from a production function estimated via OLS. However, the classic "transmission bias" problem of a correlation between unobserved productivity shocks and producers' input choices may cause OLS estimates to be biased. Therefore we also ran our specifications using TFPQ values constructed via four alternative methods designed to avoid transmission bias. One estimator included plant fixed effects in the production function, eliminating any bias caused by permanent productivity differences across plants. A second used the Wooldridge (2009) “proxy variable” estimator (which is a generalization of Levinsohn and Petrin, 2003). A third was the Blundell-Bond (1988) “system GMM” estimator, which treats inputs as endogenous variables, allows for autoregressive errors and employs lagged values as GMM-type instruments, together with other instruments that are thought to be orthogonal to fixed effects. The fourth is a Solow-style index number, where TFPQ is constructed as logged physical output minus a weighted sum of inputs. The theoretically correct weights are the elasticities of output with respect to each input; empirically, these are measured as the inputs’ share of total industry costs. In all cases, the results (presented in Appendix F) were qualitatively and quantitatively similar to those above.

We also constructed multiple alternative control groups to the incumbent-owned plants we used above. These alternative approaches used matching techniques to construct a set of control plants that looked like plants that were to be acquired in other respects. In the first matched sample, comparison plants of acquired plant *i* are incumbent plants that had been managed by the same owner who acquired plant *i*. The second matched sample formed matches based on whether a non-acquired plant is similar to acquired plant *i* in terms of pre-acquisition characteristics or trends in outcome variables. The construction of and results from these matched samples are detailed in Appendix F, Section F.3. To summarize, the basic patterns above were qualitatively and quantitatively robust to these alternative control groups.

Finally, we performed a simple placebo test by randomly assigning acquisition status to plants and then estimating the relationships between our outcome variables and this randomly generated acquisition status. The procedure and results are detailed in Appendix F, Section F.4.) We repeated this process 1,000 times and calculated the sample mean of the estimated coefficients relating “acquisition” to outcomes. In most cases, the magnitudes were only fractions of their analogs from the true acquisition samples.

5. A Mechanism

Our empirical results point to some sort of demand management ability (reflected empirically in capital utilization levels, related to unrealized output rates) as driving variation in productivity and profitability across plants, both in the cross section and over time (the latter with regard to acquisition events). In this section, we offer a simple theory that elucidates one channel through which fundamental heterogeneity across owner/managers leads to variations in the ability to manage demand, and through this, TFPQ and profitability. Further, if this heterogeneity is "carried" with the owner/manager in an acquisition into the target plants' operations, it also explains the productivity and profitability changes that surround acquisition events that we estimated above.

The specific mechanism in the model involves a managerial time allocation decision, where owners/managers must trade off spending more time managing demand but at the cost of spending less time managing production. Further, managers and plants are both of heterogeneous quality. We show below how this framework delivers the empirical patterns we document above. That said, it is possible that other possible mechanisms could explain the data, and in any case we cannot test the time allocation model directly because we have no data on owners'/managers' time allocations. Nevertheless, we find it useful to explicitly lay out a set of conditions and economic decisions that can yield the empirical patterns above.

5.1 Plant Production and Demand

For simplicity, we focus on a single plant, though implications from the model remain qualitatively the same if a firm operates several plants. The plant's owner has access to the following production technology:

$$y = g(m)x\omega \tag{5}$$

where ω is the given quality of a plant, and x is the composite input of labor and capital, weighted appropriately. (For example, if the technology is Cobb-Douglas and there are constant returns to scale, the composite would be the plant's inputs raised to their respective input elasticities). The function $g(m)$ is a flow of in-firm services provided by the plant manager to increase outputs from a given level of $x\omega$. The variable m is the manager's time allocated to managing production. This is divided into time spent ensuring that the plant operates at full capacity (therefore affecting input utilization), and time spent improving efficiency of operations themselves. For example, the former use of managerial time may involve making sure that machines are in working condition and that there are always enough workers to operate them.[27] The time spent improving efficiency of operation, on the other hand, would involve monitoring

[27] Saxonhouse (1971) describes the problem of absenteeism in the industry.

the production process, receiving reports from workers and team leaders and improving quality control and organizational structure of the firm.[28] To ease notation, assume that $g(m) = \sqrt{uv}$, where u denotes the time spent improving the frequency of operation (so that utilized input is given by $\tilde{x} = \sqrt{u}x$), while v is the time spent improving plant performance conditional on operating, and thus augments the intrinsic plant productivity, which is thus equal to $\widetilde{\omega} = \sqrt{v}\omega$.[29] The total time spent managing the plant $m = u + v$ is assumed to be bounded between 0 and some $\gamma > 0$, the manager's *effective* time endowment, which we discuss more below.

We assume that the firm first chooses x to minimize the cost of producing a given y and then optimally chooses u, v, and y. Thus the input choice x is

$$x^* = \frac{y}{\sqrt{uv}\omega}, \tag{6}$$

and the plant's cost function is $c(y) = p_x x^* = y/\sqrt{uv}\omega$, where to simplify notation we have normalized the price of x to 1 by an appropriate choice of units.

We assume the plant takes price (determined by the exchanges) as given, but the quantity of demand it faces depends on managerial time allocation. Specifically, it produces and sells the amount of output given by $\gamma - m$, so that revenues are

$$r = p(\gamma - m), \tag{7}$$

where p is the output price. The function $(\gamma - m)$ is the channel through which we introduce the notion of demand management; the plant's demand depends on the time the manager allocates to selling product. Remember that m is the total time the manager devotes toward production; this means that other things equal, a higher value of m means less demand for output. From (6) and (7), profits are

$$\pi = (\gamma - m)\left(p - \frac{1}{\sqrt{uv}\omega}\right).^{30} \tag{8}$$

5.2 Optimal Allocation of Manager's Time

---

[28] Some anecdotal evidence about the importance of this sort of managerial activity can be found in, e.g., Kuwahara (2004). See also an example in Appendix B.

[29] Diminishing returns are not necessary for the results below to hold. In particular, all of the analyses in this section go through if we instead assume input utilization and augmented plant quality are simply proportional to managerial time spent on these activities, so that $\tilde{x} = ux$ and $\widetilde{\omega} = \omega x$, although derivations become more cumbersome.

[30] We assume that p is greater than the plant's marginal cost for at least some $m_0 < \gamma$, so that operation is profitable for all values of m between m_0 and γ. The $(\gamma - m)$ function limits the size of the plant, though it would be easy to introduce upward sloping marginal costs or downward sloping residual demand (say as in a monopolistically competitive structure) if one wanted to further constrain plant size.

The plant's owner allocates his time between managing plant production and managing demand (sales) so as to maximize profit in (8):

$$\max_{u,v}(\gamma - u - v)\left(p - \frac{1}{\sqrt{uv\omega}}\right), \quad (9)$$

where we have made use of the relationship $m = u + v$. The optimal resource allocation problem (9) thus captures the fundamental tradeoff faced by the manager: if he devotes more time to managing sales, frequency and/or efficiency of operation are lost, and vice versa. This tradeoff is mitigated by effective time endowment γ; a higher value of γ reduces the lost revenue from any m. The parameter γ is thus interpreted as "demand management ability," such as a networking relationship with trading houses, reputation for reliability, as well as perhaps the ability to effectively collect debt.

It is easy to see (see Appendix I for the proof) that at the optimum:

$$u = v = m/2. \quad (10)$$

We can thus restate (9) in terms of the optimal choice of the total time allocated to production management, m, as:

$$\max_{m}(\gamma - m)\left(p - \frac{2}{\omega m}\right). \quad (11)$$

The first order condition is sufficient and it yields (after some manipulations):

$$m(\gamma, \omega) = \sqrt{2\gamma/p\omega}. \quad (12)$$

and

$$\pi(\gamma, \omega) = \left(\sqrt{\gamma p \omega} - \sqrt{2}\right)^2/\omega \quad (13)$$

(Note that by assumption, $p > MC$ at the optimum, so the numerator on the right-hand side of (13) is strictly positive.) A simple comparative exercise yields the following results.

Lemma 1:

(i) $\partial m(\gamma, \omega)/\partial \gamma > 0$; Time allocated to managing production at the plant, $m(\gamma, \omega)$, increases with γ. Moreover, both input utilization $\tilde{x}$ and augmented productivity $\widetilde{\omega}$ also increase with γ.

(ii) $\partial \pi(\gamma, \omega)/\partial \gamma > 0$ and $\partial \pi(\gamma, \omega)/\partial \omega > 0$; also, $\partial x^*/\partial \gamma > 0$. That is, profits increase in ability γ and plant quality ω, while total inputs also increase in ability γ.

(iii) $\partial^2 \pi(\gamma, \omega)/\partial \gamma \partial \omega > 0$; ability γ and plant quality ω are complements in the profit function.

Proof: See Appendix I.

Lemma 1 implies increasing returns to demand management ability, which manifest themselves in both an increased span of control in production, x^*, and input utilization, which is an increasing function of m. Augmented plant efficiency also increases in demand management ability, implying that output increases with ability even controlling for total input and its utilization. The first feature is consistent with our decomposition results in the previous section that showed an advantage of more profitable firms (with higher demand management ability) in the combined input employed and capacity utilization rates as compared to less profitable firms (with lower demand management ability). The second feature is consistent with TFPQ measured conditional on capacity utilization increasing once a plant owned by a less profitable firm is acquired by a more profitable firm.

5.3 Mergers and Acquisitions

We employ a setting inspired by the Jovanovic and MacDonald (1994) structure, which was developed with the evolution of the U.S. tire industry in mind but also fits some stark patterns in our data. Specifically, assume that an initial "basic" state of technological knowledge arrives first, offering the possibility of entry by the industry's first cohort of entrants. The "basic" nature of this initial technological knowledge is manifested in the low quality of plants, ω_1, available for the first entry cohort. Later, at some time T, there is an unanticipated jump in the state of technology (aka "refinement" in the Jovanovic-MacDonald model). This is reflected in higher quality, $\omega_2 > \omega_1$, of plants available for new entrants after time T.

Each entrant comes into the industry with some initial level of demand management ability, γ_0. Producers from the early cohort have an opportunity to develop this ability (for instance, by building reputation for consistent delivery) above and beyond the initial level, however. Assume that when the second cohort enters the market at T, the first cohort's ability is already distributed with support $[\gamma_0, \gamma_{max}]$.

Even though all entrants in the second cohort possess only the initial level of sales management ability, the quality of their plants is higher because they incorporate the superior technology. This leads to a new market equilibrium where only plant owners in the first cohort whose ability exceeds a threshold level $\gamma_e \in (\gamma_0, \gamma_{max})$ can remain in the industry; those with ability below this threshold have to exit.[31] Thus, after time T, the industry is comprised of a mixture of incumbents with (differentiated) high ability levels operating low-quality plants and new entrants with only basic ability but operating high-quality plants.

After time T, an opportunity to negotiate a merger or acquisition arrives at a random rate, and plant owners are randomly matched into negotiating pairs. It is clear that under the circumstances described above, whenever an acquisition actually occurs, it involves a higher-ability manager acquiring a plant managed by a lower-ability manager. Let a negotiating pair be formed between a manager with ability γ_H and a manager with ability γ_L, where $\gamma_H > \gamma_L \geq 1$. By Lemma 1, we have $\pi(\omega, \gamma_H) > \pi(\omega, \gamma_L)$. Therefore a manager with ability γ_H has a potential incentive to acquire the plant of a manager with ability γ_L regardless of the plant's quality. The following Proposition, proven and discussed further in Appendix I, summarizes empirical predictions for acquisition patterns:

Proposition 1: In any acquisition, a higher-ability plant owner acquires a plant managed by a lower-ability plant owner. Higher-quality plants are more likely to change ownership than lower-quality plants. Together, these imply that the most common acquisition pattern will involve a high-ability early entrant with a relatively aged plant acquiring a more recent entrant with lower ability but a newer plant.

Proof: See Appendix I.

5.4 Implications for Productivity and Profitability

We now derive implications of the merger and acquisition process outlined above for productivity and profitability of acquired plants. These implications are consistent with the patterns documented in our empirical analyses in Section 4.

To discuss the implications for productivity, note that a plant's TFPQ is given by

$$TFPQ \equiv y/u(\gamma)x = v(\gamma)\omega. \tag{14}$$

Lemma 1(i) implies that, for a given ω, TFPQ will increase with the acquiring firm manager's ability γ. Similarly, Lemma 1(ii) says that profits increase with manager's ability. Proposition 1 says that a higher-ability manager acquires a plant managed by a lower-ability manager. Together, these imply

[31] See Jovanovic and MacDonald (1994). In our data, 10 out of 21 firms that had operated in the industry prior to the late 1880s had remained small and exited by shutting their plants.

Proposition 2: Both the productivity and the profitability of an acquired plant rise after an acquisition.

Lemma 1(iii) implies increasing returns to ability in the plant profit function. Therefore:

Proposition 3: After an acquisition, the acquired plant's profits increase by more than its TFPQ.
Proof: See Appendix I.

The key intuition behind both Propositions 2 and 3 is that the new manager's superior ability to manage demand (sales) allows him to increase the time allocated to managing the production facility without sacrificing actual sales at any given price.

We next derive implications that allow us to compare the pre-acquisition levels of productivity and profitability of acquired plants with those of acquiring plants. The total derivative of the profit can be expressed as

$$d\pi = \frac{1}{\omega}\left[\sqrt{\frac{2p\gamma}{\omega}} - \frac{2}{\omega}\right]d\omega + \left[p - \sqrt{\frac{2p\gamma}{\omega}}\right]d\gamma. \qquad (15)$$

The first term in (15) captures the effect on profit of plant quality differential between acquired and acquiring plants, whereas the second term is the effect of managers' ability differential.

If two incumbents are involved in a merger negotiation, the plant quality is the same, i.e., $d\omega = 0$. In this case, (15) immediately implies that the profit of the acquiring plant is higher in the pre-acquisition period than that of the acquired plant because the acquiring plant has a higher-ability owner; i.e., $d\gamma > 0$. When the acquired plant is owned by a new entrant, on the other hand, the acquiring plant's quality is lower than the acquired plant's quality; i.e., $d\omega < 0$. Therefore relative pre-acquisition profits depend on whether the plant quality effect dominates the manager's ability effect or vice versa. In Appendix I we formally establish the following:

Proposition 4: Under suitable parameter values, pre-acquisition TFPQ of an acquiring plant can be lower than that of an acquired plant even though pre-acquisition profitability of an acquiring plant is higher than that of an acquired plant. Other things equal, this is more likely to happen if the ability of the acquiring manager, γ_H, is high.
Proof: See Appendix I.

Proposition 4 is consistent with the empirical patterns we saw in our data, but it is in contrast to both the assortative matching and the "Q" theories of mergers. A simple numerical example of the model in Appendix J illustrates how the mechanism outlined above can deliver all the patterns observed in our empirical analyses.

We have shown how a managerial time allocation decision, in the presence of heterogeneous quality managers and plants, can yield the empirical patterns documented above. We note again, however, that other possible mechanisms may be able to tie demand management to productivity and profitability levels and changes through acquisition. Further, we do not have data on owner/managers' time allocations, so we cannot test the model directly. Nonetheless, the theoretical framework outlined in this section offers a concrete example against which both the data and other theories can be compared.

6. Discussion and Conclusions

We have used unusually detailed data to investigate how acquisitions and the associated management turnover affect the performance of the firms directly involved in the transaction as well as the broader industry. These effects have been the subject of substantial, if inconclusive, theoretical and empirical research in the prior literature. Because our data allow us to observe outcomes and mechanisms at a typically unavailable level of detail, we were able to make progress toward gaining further insights.

We find in our setting (the Japanese cotton spinning industry during the turn of the 20th century) a more nuanced picture than the straightforward "higher productivity buys lower productivity" story commonly appealed to in the literature. Because they owned systematically newer and better vintages of capital equipment, acquired firms' production facilities were *not* on average any less physically productive than the plants of the acquiring firms before acquisition, at least conditional on operating. However, they were much less *profitable*. This profitability difference appears to reflect acquired firms' problems in managing the inherent demand uncertainties in the industry. These demand management problems resulted in consistently higher inventory levels and lower capacity utilization among acquired producers, raising per-unit capital costs. We show that once purchased by more profitable firms, the acquired plants saw drops in inventories and gains in capacity utilization that raised both their productivity and profitability levels, patterns consistent with acquiring owner/managers spreading their better demand management abilities across the acquired capital. This link between demand management, productivity, and profitability is, to our knowledge, a new mechanism in the literature examining how management can affect business performance.

While our data are historical in nature, we believe the patterns we document in this particular industry and time have broader lessons. They demonstrate that the ties between productivity, profitability, and ownership can be subtle while still providing a clear mechanism to spur an industry's growth. Further, they introduce a new mechanism through which superior managers lead to performance gains. Finally, Japan during the sample was essentially a developing country, less than two decades removed from a difficult transition to modernity. Thus the processes we explore here may offer specific insights into ways in which firms and industries in developing countries might achieve self-sustaining growth.

References

Atalay, Enghin, forthcoming. "Materials Prices and Productivity." *Journal of the European Economic Association*.

Bernhofen, Daniel M., and John C. Brown, 2004. "A Direct Test of the Theory of Comparative Advantage: The Case of Japan," *Journal of Political Economy*, 112 (1), 48-67.

Bertrand, Marianne, and Antoinette Schoar, 2003. "Managing with Style: The Effects of Managers on Corporate Policy," *Quarterly Journal of Economics*, 118 (4), 1169-1208.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie and John Roberts, 2013. "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 128 (1), 1-51.

Bloom, Nicholas, Rafffaella Sadun, and John Van Reenen, 2012. "The Organization of Firms Across Countries," *Quarterly Journal of Economics*, 127 (4), 1663-1705.

Bloom, Nicholas, and John Van Reenen, 2007. "Measuring and Explaining Management Practices Across Firms and Countries," *Quarterly Journal of Economics*, 122 (4), 1351-1408.

Bloom, Nicholas, and John Van Reenen, 2010. "Why Do Management Practices Differ Across Firms and Countries?" *Journal of Economic Perspectives*, 24 (1), 203-224.

Blundell, Richard, and Stephen Bond, 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *Journal of Econometrics*, 87 (1), pp. 115-143.

Braguinsky, Serguey, and David C. Rose, 2009. "Competition, Cooperation, and the Neighboring Farmer Effect," *Journal of Economic Behavior and Organization*, 72 (1), 361-376.

David, Joel, 2012. "The Aggregate Implications of Mergers and Acquisitions," Working Paper, University of South California.

Foster, Lucia, John Haltiwanger, and Chad Syverson, 2008. "Reallocation, Firm Turnover and Efficiency: Selection of Productivity or Profitability?" *American Economic Review*, 98 (1), 394-425.

Fuken Tokei (Prefectural Statistics), 1912-1920, various prefectures (in Japanese).
Geppo (Monthly Bulletin), 1889-1920 (in Japanese), Dainihon Boseki Rengokai (All-Japan Cotton Spinners Association), various issues.
Jovanovic, Boyan, and Peter Rousseau, 2002. “The Q-Theory of Mergers.” *American Economic Review*, 92 (2), 198-204.
Jovanovic, Boyan and Glenn MacDonald, 1994. “The Life Cycle of a Competitive Industry,” *Journal of Political Economy*, 102 (2), 322-47.
Kinugawa, Taiichi, 1964. *Hompo Menshi Boseki Shi* (History of the Japanese Cotton Spinning Industry) in 7 volumes (in Japanese), Tokyo, Hara Shobo, 1991 (reprinted from the original 1942 edition).
Klepper, Steven and Kenneth L. Simons, “Dominance by Birthright: Entry of Prior Radio Producers and Competitive Ramifications in the U.S. Television Receiver Industry,” *Strategic Management Journal*, 21 (10-11), 997-1016.
Kokajo (Company reports), 1896-1920 (in Japanese), Osaka University Library.
Kuwahara, Tetsuya, 2004. “The Development of Factory Management in Japan During the Early Stages of Industrialization,” Chapter 16 in *The Fibre That Changed the World*, Douglas A. Farnie and David J. Jeremy, eds., Oxford University Press, 495-519.
Levinsohn, James and Amil Petrin, 2003. “Estimating Production Functions Using Inputs to Control for Unobservables.” *Review of Economic Studies*, 70(2), 317-341.
Lichtenberg, Frank R., and Donald Siegel, 1987. “Productivity and Changes in Ownership of Manufacturing Firms,” *Brookings Papers of Economic Activity*, 1987 (3), 643-683.
Maksimovic, Vojislav, and Gordon Phillips, 2001. “The Market for Corporate Assets: Who Engages in Mergers and Asset Sales and Are There Efficiency Gains?” *Journal of Finance*, 56 (6), 2019-2065.
McGuckin, Robert H., and Sang V. Nguyen, 1995. “On Productivity and Plant Ownership Change: New Evidence from the Longitudinal Research Database,” *Rand Journal of Economics*, 26 (2), 257-276.
Menshi Boseki Jijo Sankosho (Handbook of the Cotton Spinning Industry), 1903-1920 (in Japanese), Dainihon Boseki Rengokai (All-Japan Cotton Spinners Association.
Miwa, Yoshiro, and Mark J. Ramseyer, 2000. “Corporate Governance in Transitional Economies: Lessons From the Pre-War Japanese Cotton Textile Industry.” *Journal of Legal Studies*, 29 (1), 171–203.
Nihon Choki Tokei Soran (Historical Statistics of Japan), 1987. In 4 volumes (in Japanese), Japan Statistical Association.
Noshokomu Tokei Nempo (Statistics on Agriculture, Commerce and Industry), Noshokomusho (Ministry of Agriculture, Commerce and Industry) (in Japanese), 1896-1911.
Nihon Zenkoku Shoukou Jinmeiroku, (Nationwide registry of names of traders and manufacturers) (in Japanese), edited by Suzuki Kihachi and Seki Itaro, 1898.
Ohyama, Atsushi, Serguey Braguinsky, and Kevin M. Murphy, 2004. “Entrepreneurial Ability and Market Selection in an Infant Industry: Evidence From the Japanese Cotton Spinning

Industry," *Review of Economic Dynamics*, 7 (2), 354-381.
Olley, Steven G., and Ariel Pakes, 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64 (6), pp. 1263-1298.
Rhodes-Kropf, Matthew, and David Robinson, 2008. "The Market for Mergers and the Boundaries of the Firm," *Journal of Finance* 63 (3), 1169–1211.
Saxonhouse, Gary R., 1971. "Productivity Change in the Japanese Cotton Spinning Industry, 1891-1935," Unpublished Ph.D. Dissertation, Yale University.
Saxonhouse, Gary R., 1974. "A Tale of Japanese Technological Diffusion in the Meiji Period," *Journal of Economic History*, 34, 149-165.
Saxonhouse, Gary R., 1977. "Productivity Change and Labor Absorption in Japanese Cotton Spinning, 1891-1935," *Quarterly Journal of Economics*, 91 (2), 195-220.
Syverson, Chad, 2011. "What Determines Productivity at the Micro Level?" *Journal of Economic Literature*, 49(2), 326-365.
Takamura, Naosuke, 1971. *Nihon Bosekigyoshi Josetsu* (Introduction to the History of the Japanese Cotton Spinning, in Japanese). In 2 volumes. Hanawa Shobo, Tokyo.
Wooldridge, Jeffrey M., 2009. "On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables," *Economics Letters*, 104, 112-114.
Wright, Gavin, 2011. Cotton Spinning Machinery Orders, British Textile Machinery Firms, 1878-1933. ICPSR27141-v1. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2011-11-18. doi:10.3886/ICPSR27141.v1.

Figure 1. Domestic output, import and export of cotton yarn (1887-1914)

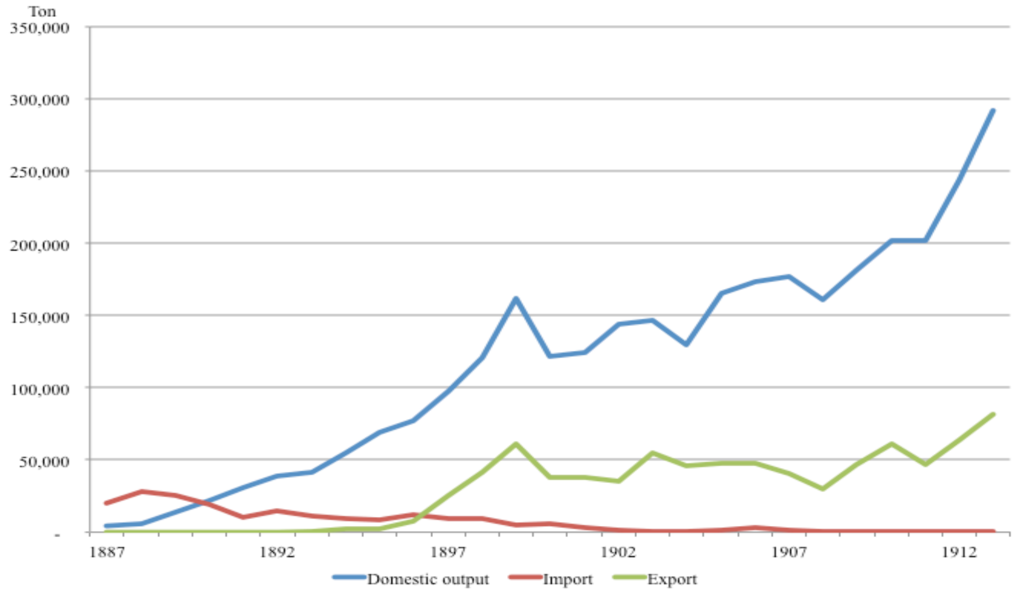

Source: Nihon Choki Tokei Soran, our estimates.

Figure 2. Capacity dynamics of older, acquired, and newer plants

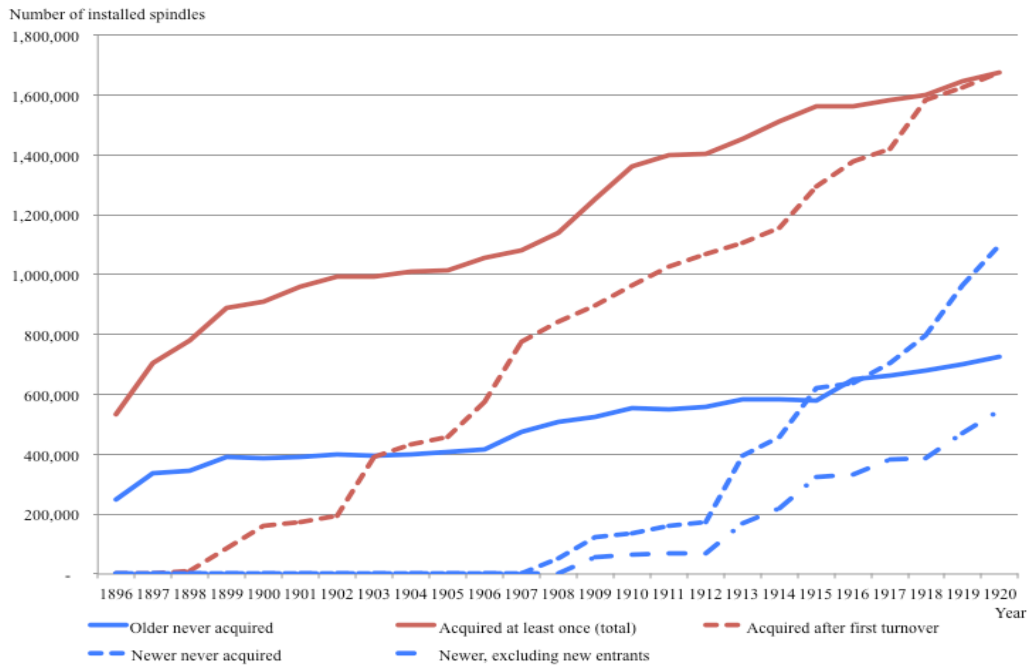

Source: Our estimates using the data described in Section 4 below.[32]

[32] "Older never acquired" are plants that came into operation in 1902 or earlier and were never targets in an acquisition. "Newer never acquired" are plants that started operating in 1908 or later and had not been acquired by 1920. "Acquired plants" is the total capacity of those plants (regardless of whether they had been acquired or not yet), while the dashed line is the capacity of those that had already gone through at least one acquisition.

Table 1. Future acquiring, acquired and exiting plants in 1896-97.

		Acquiring plants	Acquired plants		Exiting plants
			First cohort	Second cohort	
TFPQ	Mean	-0.004	0.003	0.125	-0.162
	(SD)	(0.182)	(0.229)	(0.226)	(0.513)
Earnings per paid-in value of shares	Mean	0.274	0.183	0.148	0.159
	(SD)	(0.205)	(0.076)	(0.136)	(0.101)
Price (yen/400lb)	Mean	94.7	92.8	93.2	91.7
	(SD)	(6.5)	(4.2)	(9.8)	(7.0)
Days in operation	Mean	311	315	253	265
	(SD)	(65)	(29)	(104)	(66)
Equipment age	Mean	5.28	5.87	2.50	11.77
	(SD)	(3.49)	(2.77)	(1.18)	(6.69)
Firm age	Mean	9.13	11.30	2.66	12.54
	(SD)	(5.08)	(3.56)	(1.49)	(7.86)
Observations		32	33	38	24

Note: TFPQ (quantity-based total factor productivity) is estimated as residuals from the Cobb-Douglas production function using all available observations for years 1896-97 as described in the main text. ROCE is return on equity, accounting profits divided by shareholders' paid-in capital. There are only 6 ROCE observations available for exiting plants in these years. Days in operation per year, equipment and firm age are measured in years. First cohort is plants of firms that started operating before 1892, second cohort is plants of firms that started operating in 1892 and after. Acquiring plants refer to plants belonging to future acquiring firms, exiting plants refer to plants belonging to future exiting firms (exiting not through acquisition) that will be scrapped.

Table 2. Within-acquired-plants comparisons of productivity, profitability and prices

	All acquisitions			By serial acquirer		
Dependent variable	Log output	Plant return on capital employed	Log count-adjusted price	Log output	Plant return on capital employed	Log count-adjusted price
Late pre-acquisition dummy	-0.012	0.006	-0.003	-0.014	0.020	0.030
	(0.030)	(0.015)	(0.032)	(0.056)	(0.020)	(0.037)
Early post-acquisition dummy	0.039	0.045*	0.020	0.112*	0.115***	0.092
	(0.039)	(0.026)	(0.042)	(0.065)	(0.024)	(0.069)
Late post-acquisition dummy	0.129**	0.071**	0.024	0.217**	0.146***	0.129
	(0.058)	(0.030)	(0.052)	(0.081)	(0.034)	(0.080)
Log spindles-days in operation	0.736***			0.717***		
	(0.042)			(0.083)		
Log worker-days	0.258***			0.250***		
	(0.038)			(0.058)		
Log capacity change	-0.095*	-0.075***	0.043	-0.149	-0.129**	-0.033
	(0.052)	(0.026)	(0.034)	(0.117)	(0.051)	(0.052)
Constant	-1.320**	0.155***	4.722***	-0.977	0.022	4.771***
	(0.570)	(0.012)	(0.028)	(1.144)	(0.052)	(0.037)
Plant fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Acquisition fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,248	949	1,213	555	468	530
Adjusted R-squared	0.942	0.616	0.835	0.922	0.648	0.860

Note: The omitted category includes period three years or more prior to acquisition. Serial acquirers are Kanegafuchi Boseki, Mie Boseki, Osaka Boseki, Settsu Boseki, and Amagasaki Boseki. The omitted category includes period three years or more prior to acquisition. Robust standard errors clustered at the plant level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 3. Within-acquisition comparisons of productivity and profitability: acquired and incumbent plants

	All acquisitions		By serial acquirers	
	TFPQ	Plant ROCE	TFPQ	Plant ROCE
After acquisition	-0.034	0.001	-0.014	0.003
	(0.027)	(0.016)	(0.026)	(0.019)
Acquired plant	0.025	-0.030*	0.017	-0.035*
	(0.035)	(0.015)	(0.026)	(0.018)
After acquisition x Acquired plant	0.096**	0.050**	0.132***	0.071***
	(0.036)	(0.020)	(0.038)	(0.023)
Constant	0.036	0.111***	0.099***	0.076***
	(0.020)	(0.027)	(0.025)	(0.018)
Year dummies	No	Yes	No	Yes
Acquisition–year dummies	Yes	Yes	Yes	Yes
Observations	1,398	1,261	1,001	905
R-squared	0.185	0.445	0.351	0.460

Note: Serial acquirers are Kanegafuchi Boseki, Mie Boseki, Osaka Boseki, Settsu Boseki, and Amagasaki Boseki. TFPQ is estimated residual from the production function using all available data. Robust standard errors clustered at the acquisition case level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4. Decomposition of plants' returns on capital: incumbent and acquired plants, and acquired plants pre- and post-acquisition

Pre-acquisition means	Acquired plants (A)	Incumbent plants (B)	Difference (B-A)	Percentage difference
ROCE	0.050	0.083	0.033	65.6***
of which				
net output value/total assets	0.152	0.193	0.041	26.7***
minus:				
wage cost/total assets	0.056	0.068	0.012	21.7***
capital cost/total assets	0.046	0.042	-0.004	-9.6***
# of observations	120	213		

Pre- and early post-acquisition means	Pre-acquisition (A)	Early post-acquisition (B)	Difference (B)-(A)	Percentage difference
ROCE	0.062	0.101	0.039	63.7***
of which				
net output value/total assets	0.165	0.206	0.042	25.4***
minus:				
wage cost/total assets	0.057	0.067	0.010	17.8***
capital cost/total assets	0.046	0.038	-0.008	-17.3***
# of observations	137	130		

Pre- and late post-acquisition means	Pre-acquisition (A)	Late post-acquisition (B)	Difference (B)-(A)	Percentage difference
ROCE	0.062	0.108	0.046	74.8***
of which				
net output value/total assets	0.165	0.198	0.033	20.2***
minus:				
wage cost/total assets	0.057	0.059	0.002	3.6
capital cost/total assets	0.046	0.030	-0.015	-33.2***
# of observations	137	231		

Note: The pre-acquisition time period includes observations on up to 4 years prior to acquisition. "Early after acquisition" period includes 3 years immediately following acquisitions. "Late after acquisition" period includes years starting from year 4 after acquisitions. Nominal variables are deflated by the annual consumer price index. Details of variable construction are explained in Appendix E. ***, **, and * indicate that the corresponding difference is statistically significant at the 1 percent level, 5 percent level and 10 percent level, respectively, using a double-sided *t*-test.

Table 5. Decomposition of plants' net output values: incumbent and acquired plants and acquired plants pre- and post-acquisition

Pre-acquisition means	Acquired plants (A)	Incumbent plants (B)	Difference (B)-(A)	Percentage difference
ln(net output value/total assets)	-2.024	-1.704	0.319	37.6***
of which:				
ln(price margin)	-1.354	-1.307	0.047	4.8*
ln(total input/total assets)	-0.654	-0.406	0.248	28.1***
TFPQ	-0.016	0.009	0.024	2.5
# of observations	117	206		

Pre- and early post- acquisition means of logs	Pre-acquisition (A)	Early post-acquisition (B)	Difference (B)-(A)	Percentage difference
ln(net output value/total assets)	-1.971	-1.701	0.270	31.0***
of which:				
ln(price margin)	-1.359	-1.290	0.069	7.2**
ln(total input/total assets)	-0.599	-0.456	0.143	15.4**
TFPQ	-0.012	0.045	0.057	5.9**
# of observations	139	131		

Pre- and late post- acquisition means of logs	Pre-acquisition (A)	Late post-acquisition (B)	Difference (B)-(A)	Percentage difference
ln(net output value/total assets)	-1.971	-1.722	0.248	28.2***
of which:				
ln(price margin)	-1.359	-1.252	0.107	11.3***
ln(total input/total assets)	-0.599	-0.560	0.038	3.9
TFPQ	-0.012	0.091	0.103	10.9***
# of observations	139	231		

Note: The pre-acquisition time period includes observations on up to 4 years prior to acquisition. "Early after acquisition" period includes 3 years immediately following acquisitions. "Late after acquisition" period includes years starting from year 4 after acquisitions. Nominal variables are deflated by the annual consumer price index. Details of variable construction are explained in Appendix E. ***, **, and * indicate that the corresponding difference is statistically significant at the 1 percent level, 5 percent level and 10 percent level, respectively, using a double-sided *t*-test.

Table 6. Inventory and accrued payments to output value ratios: incumbent and acquired plants and acquired plants pre- and post-acquisition

Means	Acquired plants (A)	Incumbent plants (B)	Difference (B-A)	Percentage difference
Inventory/produced output (C)	0.045	0.018	-0.027	-60.9***
Accrued revenues/produced output (D)	0.031	0.015	-0.017	-52.7***
Unrealized/produced output (C)+(D)	0.078	0.033	-0.045	-58.2***
# of observations	111	190		

	Pre-acquisition (A)	Early post-acquisition (B)	Difference (B-A)	Percentage difference
Inventory/produced output (C)	0.046	0.013	-0.033	-71.6***
Accrued revenues/produced output (D)	0.030	0.019	-0.011	-36.0***
Unrealized/produced output (C)+(D)	0.078	0.031	-0.047	-60.3***
# of observations	134	101		

	Pre-acquisition (A)	Late post-acquisition (B)	Difference (B-A)	Percentage difference
Inventory/produced output (C)	0.046	0.009	-0.038	-80.9***
Accrued revenues/produced output (D)	0.030	0.015	-0.015	-49.8***
Unrealized/produced output (C)+(D)	0.078	0.023	-0.054	-70.3***
# of observations	134	121		

Note: The pre-acquisition time period includes observations on up to 4 years prior to acquisition. "Early after acquisition" period includes 3 years immediately following acquisitions. "Late after acquisition" period includes years starting from year 4 after acquisitions. *** indicates that the corresponding difference is statistically significant at the 1 percent level, using a double-sided *t*-test.

Table 7.
Plant and firm perform metrics in 1898-1902 by in-network and out-of network firms

Means	Out-of-network (A)	In-network (B)	Difference (B-A)
ROCE	-0.005	0.044	0.049***
Unrealized output ratios	0.127	0.084	-0.043***
Spindle utilization rates	0.741	0.781	0.040**
Output prices	85.4	77.5	-7.9**
# of observations	100	104	

Appendix—For Online Publication

A. Data Description

Our main data source is plant-level data collected annually by Japan's prefectural governments. The collection of these data started in 1899, and until 1911 they were brought together and published nationally in a single source, the *Statistical Yearbook of the Ministry of Agriculture and Commerce* (Noshokomu Tokei Nempo). Even though the national government discontinued publishing these data after 1911, the subsequent data can still be found in prefectural statistical yearbooks. For this paper we have collected and processed all the available data between 1899 and 1920.

The plant-level annual data record inputs used and output produced by each plant in a given year in physical units. In particular, the data contain the number of spindles in operation, number of days and average number of hours per day the plant operated, output of the finished product (cotton yarn) in physical units, the average count (measure of fineness) of produced yarn, the average monthly price per unit of yarn produced, the number of factory floor workers (subdivided into male and female workers), average daily wages separately for male and female workers, as well as the data on intermediate inputs, such as the consumption of raw cotton, type of engine(s) that powered the cotton spinning mill (steam, water, electrical or gas/kerosene), their total horsepower, etc.

We supplement the plant-level data from prefectural governments' statistics by several other data sources. In particular, we employed the data containing the same variables as above collected at the firm level by the All-Japan Cotton Spinners' Association (hereafter "Boren," using its name's abbreviation in Japanese) and published in its monthly bulletin (Geppo). Even though the data were collected at the firm- and not plant level, there were no acquisitions and mergers to speak of until 1898 and all but 2 firms were single-plant firms, so the data are usable for pre-acquisition plant-level comparisons. We thus converted monthly Geppo data for 1896-1898 to annual data and use these in our estimations alongside government-collected annual plant-level data for 1899 and beyond.

With regard to data reliability, past literature has concluded that "the accuracy of these published numbers is unquestioned." (Saxonhouse, 1971, p. 41). Nevertheless, we scrutinized these numbers ourselves and found occasional, unsystematic coding errors as well as obvious typos. We then used the overlap between the government-collected annual plant-level data and the firm-level monthly data published in Geppo to cross-check the data for single-plant firms. In the vast majority of cases we found that the annual data in statistical yearbooks and the annualized monthly data corresponded very closely (the discrepancy, if any, did not exceed a few percentage points). We were also able to use annualized monthly data to correct above-mentioned coding errors and typos in annual plant-level data in a significant number of cases. In the end, we have not been able clean the annual plant-level data in just about 5 percent of the total number of observations. We elected to drop such observations from our analysis.[34]

Each plant in the records is associated with the firm that owned it in a given year, making it possible to directly compare the plant's physical (quantity) productivity before and after the change in ownership. This feature makes our data particularly attractive for analyzing plant productivity changes following ownership and/or management turnover. We also collected actual stories surrounding each acquisition and ownership turnover case, including but not limited to identities and backgrounds of the most important individuals involved (shareholders, top managers and engineers). Several data sources made this possible. First, almost 90 percent of the

[34] To the best of our knowledge, we were the first to conduct this comprehensive cleaning of published plant-level records for the Japanese cotton spinning industry for 1896-1920. Our cleaned plant-level data tables and the details of the procedure outlined above are available upon request.

Japanese cotton spinning firms (and all significant firms) were public (joint stock) companies, obligated to issue shareholders' reports every half a year. Copies of these reports were also sent to Boren's headquarters in Osaka and those of them that have survived until the present day are currently hosted in the rare books section of Osaka University library. With the permission from the library we have photocopied the total of 1,292 reports on 149 firms, all what was available for the period from the early 1890s until 1920.[35] Each report, in particular, contains a list of all shareholders and board members of the company issuing it, making it possible to see whether shareholders or top management teams had already been substantially overlapping even prior to the formal acquisition event and what were the new positions (if any) of major shareholders and top managers of acquired firms in the new integrated firms. Company reports also contain detailed balance sheets and profit-loss statements as well as qualitative information about shareholders' meetings, deaths, illnesses, resignations and replacements of board members and so on, which we use as appropriate.

We supplement these primary data sources by the information contained in the seven-volume history of the industry written in the 1930s by the Japanese historian Taiichi Kinugawa (Kinugawa, 1964). The book is basically a collection of chapters each of which is dedicated to a particular firm, describing its background, evolution and major personnel involved since the firm entered the industry; in its totality, the chapters cover all but a few firms that entered the industry from its inception in the 1860s and until the beginning of the 20th century. While it appears that Kinugawa had access to the same company reports that we have (in particular, he cites as missing the same reports that we found missing in the Osaka University library), his book nevertheless provides us with a lot of additional insights because he was able to conduct interviews with many important individuals involved in those firms who were still alive at the time he wrote his book. Kinugawa also presents invaluable information about the background of most important shareholders and managers of each firm covered in his book as well as the storyline about how each firm was conceived.

Finally, we also used published company histories of firms that had survived until after World War II (some of them still surviving), although these are of less significance both because the information could be biased and because the level of detail is not nearly as great as in company reports or in Kinugawa's history of the industry. Nevertheless, some qualitative information contained in those company histories proved to be usable and is used in this paper as appropriate.

While physical input and output data give us a unique chance to examine physical plant productivity as opposed to its revenue productivity, estimating the plant's TFPQ still presented several challenges. First, even though cotton yarn is a relatively homogeneous product it still comes in varying degree of fineness, called "count."[36] Output of cotton yarn in our data is measured in units of weight, but there is also information about the average count produced by a given plant in a given year. To make different counts comparable for the purpose of productivity

[35] While some of these company reports had been used in previous research by Japanese historians, we were the first to systematically digitalize them. The Osaka University library plans to launch a web site that will make our digital copies available in the public domain in the near future.

[36] The yarn count expresses the thickness of the yarn and its number indicates the length of yarn relative to the weight. The higher the count, the more yards are contained in the pound of yarn, so higher-count yarn is thinner (finer) than lower-count yarn. Producing higher-count (finer) yarn generally requires more skill and superior technology than producing lower-count (coarser) yarn. High-count yarn is often also improved further by more complex technological processes known as doubling, gassing, and so on, which were quite challenging for the fledgling Japanese cotton spinning mills to master at that time.

analysis, we converted them to the standard 20^{th} count using a procedure in which we first estimated coefficients on different count dummies in the production function regression, with (log) output measured in weight as the dependent variable, including also (logged) spindle and worker input and year dummies. We then used the estimated coefficients on count dummies to convert output of other counts to the 20^{th} count (details are available upon request). We also conducted all our estimations in an alternative way, using output in weight units and including the average count as a separate regressor when estimating the production function and confirmed that the results were similar.

Second, the worker count data include blue-collar workers (by gender—male, "danko" and female, "joko") but do not include white-collar workers ("shyain"). Hence, in our total factor productivity estimates, the residual should be interpreted as reflecting the managerial input in a broad sense, including the input of all white-collar personnel. As the data give us the number of male and female blue-collar workers separately, we used the plant-year-specific ratios of female to male wages to convert one unit of female labor to one unit of male labor.[37] Third, while we have direct measures of capital input in the data in the form of the number of spindles in operation, spinning frames are just one part of capital equipment which accounts for 25-30 percent of the total equipment cost of a mill (Saxonhouse, 1971, p. 55). Correlation between spindles and other equipment (cards, draw frames, slubbing frames, intermediate frames, roving frames, etc.) is, however, extremely high (over 95 percent), so "there is no question that spindles are a good proxy for equipment as a whole" (Saxonhouse, 1971, p. 56). We also have the data on the number of spindles installed in each plant in each year, which allows us to measure capacity utilization rates and follow any plant upgrades as the new equipment is installed.

Finally, when estimating the production function we followed Saxonhouse (1971 and 1977) and excluded intermediate inputs. The reason, already discussed by Saxonouse, is that the coefficient of transformation of raw cotton into cotton yarn is almost fixed, at least when both input and output are measured in weight units (the raw correlation in our data is 0.95), so it renders all other inputs economically and statistically insignificant in the production function. Raw cotton can be added to inputs without running into this problem when output is adjusted for count but such a procedure would still be problematic because finer counts of cotton yarn are typically produced from higher-quality raw cotton (e.g., American or Egyptian cotton instead of Indian cotton) and we do not have plant-level data about the type of raw cotton used. Nevertheless, we did check the robustness of our estimates to including the raw cotton input (and also engine horse power) with output adjusted for count and confirmed that the results pertaining to total factor-productivity presented in this paper still hold, although the estimated magnitude of the coefficients is reduced by about one half (most of them still retain statistical significance, however).

Finally, even though our data also contain records of the average number of hours plants operated per day in a given year, we elected to measure our inputs by worker- and spindle-days in the main specifications in this paper. As is well known, plants in Japan in this period operated in two shifts around or almost around the clock most of the time (e.g., Takamura, 1971), although occasionally the second shift would be suspended and the plant would operate only for half a day.

[37] In the division of labor between sexes in Japanese cotton spinning mills, opening, mixing, carding, repairing and boiler room work were generally (although not exclusively) men's jobs, while tending, drawing, roving and operating ring frames were generally women's work (Clark, Cotton Goods in Japan, pp. 191-194, cited in Saxonhouse, 1971, p. 56). Using female to male wage ratios to aggregate the labor input assumes that wages reflect the marginal productivity of each sex. All our estimates are completely robust to using the number of male and female workers separately in the production function estimations.

Unfortunately, the information about average hours in operation reported in the annual plant-level data turned out to be rather inaccurate (in particular, there are large and apparently random discrepancies with the more accurate monthly firm-level data from firm reports in Geppo). We did repeat all the estimation below using the information on hours in operation and the results remained very much the same, with the impact of acquisitions on TFPQ even more strongly pronounced than reported in Tables 3 and 4 in the main text.

B. An example of management turnover in our data

In August 1898, the shareholders of the decade-old struggling Onagigawa Menpu (Onagigawa Cotton Fabrics) company in Tokyo, Japan appointed a new board member. His name was Heizaemon Hibiya, a cotton trader and also founder and CEO of Tokyo Gasu Boseki (Tokyo Gassed Cotton Spinning) company, one of the more recent and successful high-tech entrants in the Japanese cotton spinning industry at the time. When Hibiya first toured the Onagigawa factory, he was reportedly in shock at what he saw. Workers brought portable charcoal stoves and smoked inside the plant. Women cooked and ate on the factory floor, strewing garbage. Cotton and other materials were everywhere, blocking hallways, while workers in inventory room gambled. Managerial personnel were out at a nearby river fishing (Kinugawa, 1964, Vol. 5).

Hibiya, who was promoted to company president in early 1899, wasted no time in introducing much needed change. All work-unrelated and hazardous activities on factory premises were immediately banned. Plant deputy manager tried to stir workers' unrest and was quickly fired, together with the head of the personnel department and the chief accountant (an off-duty police officer was temporarily stationed inside the plant as a show of new management's determination). But Hibiya did not stop at just introducing disciplinary measures. Even though he had another plant of his own to take care of, he and his right-hand man from Tokyo Gasu Boseki came to the Onagigawa factory and personally inspected equipment and checked output for defects on a daily basis, while also teaching workers how to do it on their own. During these visits, Hibiya reportedly engaged workers in conversations related to technology and production practices, taking questions, writing down those that he couldn't answer immediately and coming back the next day with answers obtained from outside sources. Having determined that one reason for poor quality was that factory resources were spread too thinly, he concentrated production in just a few key areas, shutting down some workshops and switching from in-house production of finer counts of cotton yarn to procuring those from his other newer and more high-tech plant. Other measures included selling older equipment and purchasing more modern machines.

The above account reads remarkably similar to the description of the experiment in modern Indian textile industry conducted by Bloom et al. (2013). The results of Hibiya's restructuring effort were also equally or perhaps even more impressive. Using our data described in detail below, we estimate that the plant's TFPQ relative to the industry average more than doubled in the 3 years after Hibiya took over compared to 3 years before that while labor productivity (measured as output in physical units per worker-hours) increased on average by 70 percent. Over the same period, labor productivity in two other comparable plants in the same Tokyo area increased by just 6 percent. It is also worth noting that Hibiya was not part of an international aid effort; he was hired through an internal decision-making process of the shareholders, dishing out their own money.[38]

[38] Hibiya's story is typical of industrialization pioneers in Japan and shows how much it was a land of opportunity at the time. Born Kichijiro Ohshima, third child of the owner of a hotel in a small provincial town, the future Heizaemon Hibiya was noticed by a cotton trader who stayed at the hotel when the boy

C. Acquisitions over time and the concentration of ownership in 3 largest firms, 1898-1920.

Table A1. Number of acquired plants by year

Year	Number of acquired plants	Fraction of total	Of which: acquired by largest acquirers	Fraction of total number of acquisitions
1896	0	0.000	0	0.000
1897	0	0.000	0	0.000
1898	1	0.012	0	0.000
1899	5	0.060	0	0.000
1900	7	0.085	3	0.429
1901	1	0.012	0	0.000
1902	2	0.025	1	0.500
1903	15	0.188	7	0.467
1904	2	0.025	0	0.000
1905	3	0.038	0	0.000
1906	5	0.062	3	0.600
1907	11	0.136	6	0.545
1908	2	0.025	0	0.000
1909	1	0.011	0	0.000
1910	1	0.012	0	0.000
1911	6	0.069	4	0.667
1912	5	0.057	2	0.400
1913	0	0.000	0	0.000
1914	0	0.000	0	0.000
1915	4	0.038	2	0.500
1916	5	0.048	2	0.400
1917	3	0.028	0	0.000
1918	11	0.100	7	0.636
1919	3	0.026	0	0.000
1920	2	0.017	0	0.000
Total	95	0.043	37	0.389

Note: Largest acquirers are Kanebo, Mie Boseki, Osaka Boseki, Settsu Boseki and Amabo. Excluding 15 plants that were consolidated in 1914 in the equal merger between Mie Boseki and Osaka Boseki.

was 13 and went to Tokyo to become the trader's apprentice. At the age of 20 he was doing trades on his own. He went on to grow one the most successful cotton trading houses in the Tokyo area, while also playing a major role in several prominent cotton spinning and other firms and eventually becoming vice-chairman of the Tokyo Chamber of Commerce.

Figure A1. Ownership concentration in three largest firms

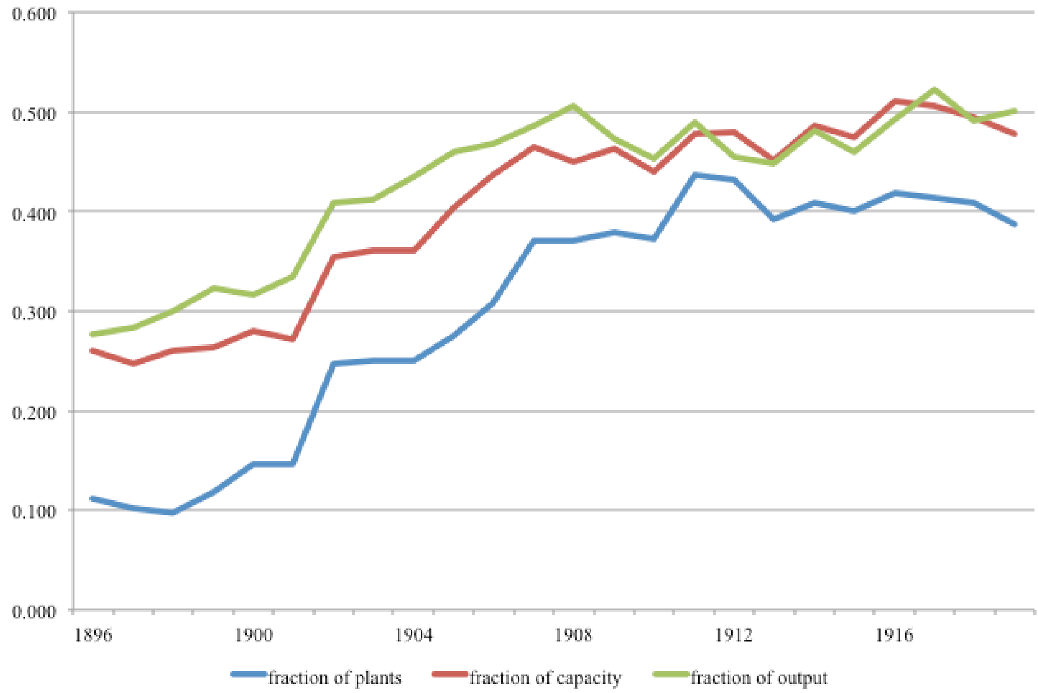

Note: The figure depicts the evolution of the fraction of plants owned by the three largest firms in 1920 (Kanebo, Toyobo, Dainippon Boseki) and these plants' capacity and output as a fraction of the industry total. Toyobo data include that of its predecessor firms (Osaka Boseki and Mie Boseki) prior to their 1914 merger. Dainippon Boseki includes the data of its predecessor firms (Amabo and Settsu Boseki) prior to their 1918 merger.

D. Evidence of capital vintage effects as reflected in machine characteristics

We extracted the data on a number of specific orders made by Japanese cotton spinning firms during our sample for capital equipment from British suppliers from the general file on world-wide orders from British manufacturers in 1879-1933 compiled by Gary Saxonhouse and archived at the ICSPR (Wright, 2011).[39] We used these data to measure the average values of numerous technical characteristics of the machines that were shipped in each year. These characteristics are (1) average spindle speed (sometimes highest and lowest speeds are also available but mostly the data are on average speed); (2) average (and also highest and lowest) count of cotton yarn to produce which the machine was designed for; (3) number of spindles per frame; (4) how many different types of raw cotton the machine was designed to work with (from 1 to 4); and (5) dummies equal to 1 if the machine was designed to work with Indian cotton and 0 otherwise, and the same for American and Egyptian cotton (the omitted category would be machines designed to work only with inferior-quality Japanese or Chinese cotton).

This yielded a file of vintage-specific machine characteristics for each year in our data. We then merged this file with our main data file which contains vintage age of machines in all plants (calculated as the weighted average of spindle capacity installed in a given year—in practice we subtract one year from the year machines were equipped to allow for delivery and installation time). This makes it possible to assign average vintage-year characteristics (1)-(5)

[39] We thank Patrick McGuire for helping us with these data.

above to all individual plants in our data.

Table A2 shows the degree of technological progress in machine characteristics from an early vintage to a later vintage during the first waves of large-scale entry into the Japanese cotton spinning industry. Even though we have the data by each year, there are just a few orders prior to 1887, at which point orders pick up (14 in 1887, 16 in 1888, and 11 in 1889). There are only 8 orders in 1890 and only 2 orders in 1891, but the orders pick up again, and quite dramatircally so staring in 1892; there were 14 orders in that year, 25 in 1893, 35 in 1894, 18 in 1895, 39 in 1896 and 24 in 1897. Despite this large number of observations, machine characteristics are remarkably similar throughout these later years, so we lump them all together into the single 1892-97 vintage. (t-tests on mean differences across different subperiods within this period were all insignificant.)

The differences in average characteristics of the machines belonging to pre-1892 vintage where our first cohort firms (started operating prior to 1892) entered the industry and the later vintage which was ordered by the second cohort of entrants (and also by those of the first-cohort firms that attempted to modernize) are rather large and are all statistically significant at the 1 percent level using double-sided t-test. We also used 1890 or 1891 as the cutoff year and the results were basically the same.

Table A2. Average machine characteristics by two vintages

	Pre-1892 vintage	1892-97 vintage
Spindle rotation speed (RPM x 1000)	7.10	7.71
Cotton yarn count designed for	17.53	19.96
Number of spindles per ring frame	332.25	377.71
Number of cotton types designed for	1.06	2.47
Designed for Indian cotton	0.00	0.56
Designed for US cotton	0.04	0.44

Along all dimensions, the newer machines embody more technological capabilities. First, the increase in spindle rotation speed means that the same number of spindles operating the same number of hours can produce more cotton yarn if employed at full speed. The differences in average speed over the period would allow output per operating spindle to increase by 6.4 percent. However, on top of this there was also an 11.4 percent increase in the count of cotton yarn machines are designed for, resulting in a total potential boost to count adjusted output per spindle of 17.8 percent. The number of spindles per frame also increased from the older to the newer vintage, by 8 percent. Because the frame size remains the same, it is reasonable to assume that the same amount of workers attend to one frame as before (or at least the number of workers attending to a frame does not grow anywhere near proportionately to the increase in the number of spindles and their speed), resulting in a potential of up to 8 percent improvement in labor productivity per machine from spindle density. Finally, the newer machines were more versatile. While older machines were almost exclusively designed to work with just one type of cotton (Japanese or Chinese), new machines could work with an average of 2.47 cotton types. Moreover, about half of the new machines were designed to work with Indian or US cotton as compared to virtually none of the older machines.

As already mentioned, second-cohort entrants all had access to these new and better machines. However, first-cohort entrants—especially those of them who later became our acquiring firms—also ordered new machines and gradually removed old machines from service. Therefore, the gap in machine quality between different firm types is not as dramatic as the difference in vintages may indicate, but it is still considerable as shown in Table A3. Table A3 follows the same format as Table 1 in the main text, but it shows differences in machine characteristics and therefore differences in potential rather than actual productivity across these categories (recall that these figures are computed for 1896-97, when no acquisition had yet taken place).

Table A3. Technical characteristics of machines by types of plants, 1896-97

		Acquiring plants	Acquired plants		Exiting plants
			First cohort	Second cohort	
Spindle rotation speed (RPM x 1000)	Mean	7.46	7.44	7.70	7.01
	(SD)	0.34	0.29	0.14	0.33
Cotton yarn count designed for	Mean	18.57	18.35	20.32	17.80
	(SD)	1.46	1.87	2.24	0.84
Number of spindles per ring frame	Mean	365.91	357.01	379.92	314.69
	(SD)	22.58	33.43	8.60	47.46
Number of cotton types designed for	Mean	1.89	1.57	2.48	1.29
	(SD)	0.69	0.70	0.22	0.61
Designed for Indian cotton	Mean	0.32	0.17	0.59	0.11
	(SD)	0.30	0.25	0.15	0.25
Designed for US cotton	Mean	0.28	0.21	0.43	0.11
	(SD)	0.24	0.25	0.13	0.14
Observations		32	31	38	23

Notes: See Table 1 in our main text.

Comparing newer (second-cohort) future acquired plants to future acquiring plants, we can see that the average spindle rotation speed was about 3.3 percent higher among newer plants, while the count they were designed to produce was about 9.4 percent higher (both differences are statistically highly significant). Together, thus, potential increase in count-adjusted output due to machine superiority alone was 12.7 percent. The increase in the number of spindles per ring frame was 3.8 percent, again statistically highly significant, and there are huge differences in machines' versatility (number of cotton types they can work with and the fraction designed to work with better-quality imported cotton). Again, as we saw in the main text, exiting plants are the worst on all aspects in these technical characteristics (which is also reflected in very old equipment age of those plants in Table 1 in the main text).

Thus we have direct evidence of technological superiority of younger future acquired plants compared to future acquiring plants in those years. In the language of our model, the younger plants' omega was indeed higher (by 13-16 percent overall perhaps) than that of the acquiring plants. The fact that acquired plants didn't exhibit big TFPQ differences compared to acquiring plants before their acquisition (even though they did exhibit this difference in 1896-97, which were very good years for the industry without few worries about demand management) suggests that after the onset of industry-wide demand problems starting around 1898, these plants started squandering their potential productivity advantage. It was only regained after acquisition and the influence of new management.

E. Construction of plant-level profitability measure

We construct a plant-level analogue to ROCE (return on capital employed) according to the following procedure. Output of cotton yarn, output price, and the number of male and female work-days as well as the corresponding daily wages are observed directly at the plant level. Capital cost is defined as the sum of depreciation and interest cost of debt. For deprecation, we use firm-level accounting data and a standard depreciation rate of 0.05 to compute the amount of depreciation of fixed capital, and we assign it to each plant in a multiple-plant firm in proportion to the plant's installed capacity in the firm capacity. The interest cost of borrowed capital is imputed to each plant as its share of the firm-level interest-bearing debt, multiplied by the economy-wide interest rate (proxied by the Bank of Japan discount rate), times 1.31. The latter number is the coefficient on the economy-wide interest rate estimated from a regression using all available firm-level data, with actual interest payments to the amount of interest-bearing debt reported in firms' accounting statements as the independent variable, and economy-wide interest rate and year dummies as the explanatory variables (the estimated coefficient is statistically highly significant).

To complete the construction of plant-level ROCE, we also need a proxy for the margin on the gross value of output (parameter ψ in the first decomposition equation (4) in the main text). To this effect, we in turn need to estimate the cost of intermediate input (raw cotton) and other non-labor operation expenses (packing, shipping, engine fueling, etc.). Since there were also markets for yarn and raw cotton wasted in the production process and subsequently recovered, we also need to add the amount of sales of waste yarn and recovered waste cotton as those are the by-products of the spinning process.

As already mentioned, the production of cotton yarn uses the main input of raw cotton in almost fixed proportion to output (the coefficient of correlation between output of cotton yarn and input of raw cotton, both measured in physical units, is 0.997 for our acquired plants both before and after acquisitions). Data from profit-loss statements suggest that non-labor expenses were also a more less constant fraction of sales. We thus assume a fraction of intermediate inputs and other operational expenses in the value of output to be a common parameter for all plants, and we calculate it from available firm-level profit-loss statements. Physical volume of waste yarn and recovered raw cotton are observed at the plant level, and we estimate the sales of these by-products by multiplying these amounts by the available data on their yearly market prices. The main parameters obtained in this way are presented in Table A13, and they lead to calculated value of ψ=0.15, which is employed in constructing plant-level ROCE measure and our first decomposition analysis.[40]

[40] While we assume these to be the same for all firms, it is possible that less successful future acquired firms may have had higher (non-wage) operating costs than future acquiring firms. Available data from company profit-loss statements do not, however, indicate that this was the case. Future acquired firms may have also faced higher interest rates on their borrowings than more successful future acquiring firms. Based on available data from company reports, we cannot reject this possibility; the ratio of interest payments to the amount of borrowing is indeed considerably (and statistically significantly) higher for target firms in pre-acquisition years than for the firms that eventually acquired them in the same years. The impact of this on our overall profitability differential measure is fairly small, but inasmuch as it is present, our plant-level ROCE measure would actually understate the profitability disadvantage of acquired plants relative to plants of acquiring firms. The decomposed differentials reported in the main text should therefore be considered lower bounds.

Table A4. Parameters in cost calculations

Cotton input to output ratio	1.162
Relative cotton price	0.677
Waste yarn to output ratio	0.012
Relative waste yarn price	0.294
Recovered cotton to input ratio	0.113
Relative recovered cotton price	0.438
Net input cost to total output value ratio	0.746
Non-labor operating expenses rate	0.105
Margin before labor and capital cost	0.150

The plant-level ROCE measure obtained in this way (and winsorized at the top 2 percent) is highly correlated with firm-level ROCE measure available for pre-acquisition years (the coefficient of correlation is 0.75). Figure A2 plots the density of plant-level ROCE distribution and the corresponding firm-level ROCE in the whole sample, and visually confirms that our measure of plant-level profitability is a reasonable proxy for profitability as reported in firm accounts.

Figure A2.
Distributions of constructed plant-level ROCE measure and ROCE from firm accounts (all plants and years)

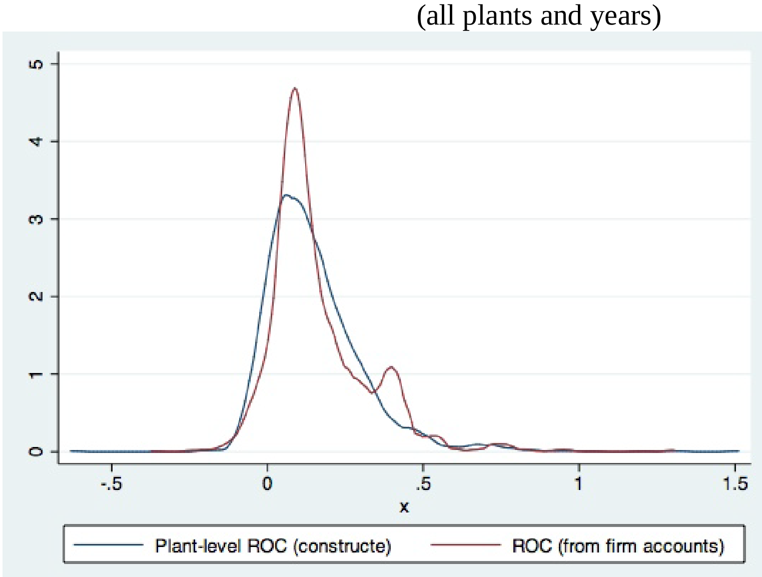

F. Robustness Checks

In this section we describe the details of the design and the results of robustness checks summarized in Section 4.6 of the main text.

We are interested in estimating the following parameters:

$$\beta_1 = \frac{1}{N_M}\sum_{i\in M}\left\{\frac{1}{\#m_i}\sum_{j\in m_i}\omega_j\left(y_{ja}^C - y_{jb}^C\right)\right\} \quad \text{(A1)}$$

$$\beta_2 = \frac{1}{N_M}\sum_{i\in M}\left\{y_{ib}^A - \frac{1}{\#m_i}\omega_j\sum_{j\in m_i}y_{jb}^C\right\} \quad \text{(A2)}$$

$$\beta_3 = \frac{1}{N_M}\sum_{i\in M}\left\{\left(y_{ia}^A - y_{ib}^A\right) - \frac{1}{\#m_i}\sum_{j\in m_i}\omega_j\left(y_{ja}^C - y_{jb}^C\right)\right\} \quad \text{(A3)}$$

where M is a set of matches, and acquired plant i is matched with "comparison" plants to form match m_i. Outcome variables y_{ib}^{A} are the physical TFPQ measures of acquired plant *i* and the (log of) the plant gross operating surplus rate (GOS) of acquired plant *i* before an acquisition event, and outcome variables y_{jb}^{A} are these variables after the acquisition event. Superscript C indicates the corresponding variables for comparison plants. N_M is the total number of matches, $\#m_i$ is the number of comparison plants within match m_i, and ω_j is a weight attached to the outcome variables, y_{ja}^{C} and y_{jb}^{C}.

The parameters β_1, β_2 and β_3 can be estimated by

$$\bar{y}_{it} = \alpha_0 + \beta_1 AA_{it} + \beta_2 Aquired_{it} + \beta_3 Aquired_i \times AA_{it} + \mu_t + \varepsilon_{it} \qquad \text{(A4)}$$

where $\bar{y}_{it}$ is the outcome variable of plant *i* at time t if it belongs to a group of acquired plants. The outcome variables of comparison plants within the match m_i are collapsed to $\bar{y}_{it} = \sum_{j \in m_i} \omega_j y_j$, the weighted average of outcomes of comparison plants within the match m_i. The variable AA_{it} is a dummy equal to 1 if acquisition m_i happened prior to year *t* and zero otherwise, while the variable $Acquired_i$ is equal to 1 if plant *i* is purchased in acquisition case m_i and zero otherwise. μ_t is an acquisition-year fixed effect. The estimate $\hat{\beta}_3$ reflects the post-acquisition difference-in-difference between acquired and incumbent plants of acquiring firms by accounting for acquisition-case effects.[41]

F.1 Alternative TFPQ measures

In the main text, we used TFPQ estimates obtained as residuals from a production function estimated via OLS to compare outcomes of acquired and incumbent plants. As already mentioned, the classic "transmission bias" problem of a correlation between unobserved productivity shocks and producers' input choices may cause OLS estimates to be biased. Therefore we also ran our specifications using TFPQ values constructed via four alternative methods designed to avoid transmission bias.

Our basic specification of the production function is given by

$$q_{it} = \gamma_0 + \gamma_1 l_{it} + \gamma_2 k_{it} + \gamma_3 x_{it} + \mu_t + \varepsilon_{it},$$

where q_{it} is plant's logged output, l_{it} is the plant's logged number of composite worker-days, k_{it} is its number of spindle-days in operation, x_{it} is a vector of control variables that include the change in log plant capacity from the previous year and (logged) age of the plant's machines, and μ_t are year dummies. The first measure of TFPQ is residuals from the OLS regression of the production function as in the main text. The second measure utilizes residuals from the Wooldridge (2009) GMM estimation method where a proxy variable is included to control for unobserved firm-level productivity shocks. This method is a generalization of earlier approaches using investment (Olley and Pakes, 1996) or intermediate inputs (Levinsohn and Petrin, 2003) as proxy variables, but it allows for the use of any proxy variable that is positively associated with unobserved productivity (Woodridge, 2009). In our case, as implied also by our theoretical mechanism, we can employ capacity utilization rate (a variable not normally available to an econometrician) as such a proxy. Our second measure of TFPQ (Wooldridge) is thus the residual from the production function estimation by the Wooldridge (2009) GMM method with the (logged) capacity utilization rate serving as the proxy variable. To construct the third measure of TFPQ we follow the system GMM approach proposed by Blundell and Bond (1998). Specifically, we conduct a

[41] We can also write equations (1) to (3) as $\beta_1 = \frac{1}{N_M}\sum_{i \in M}\left(\bar{y}_{m_i a}^{C} - \bar{y}_{m_i b}^{C}\right)$, $\beta_2 = \frac{1}{N_M}\sum_{i \in M}\left\{y_{ib}^{A} - \bar{y}_{m_i b}^{C}\right\}$, and $\beta_3 = \frac{1}{N_M}\sum_{i \in M}\left\{\left(y_{ia}^{A} - y_{ib}^{A}\right) - \left(\bar{y}_{m_i a}^{C} - \bar{y}_{m_i b}^{C}\right)\right\}$. These expressions give us an interpretation of the parameters similar to the one from the standard dif-in-dif estimations.

two-step implementation of the Blundell and Bond estimator with two-period lags, treating the number of worker- and spindle-days as endogenous variables, alongside with output, and generating GMM-style instruments for them. The estimations again include also year dummies, the change in log plant capacity from the previous year, and (logged) age of the plant's machines as additional instruments. Finally, we construct the fourth measure of TFPQ by using the standard index approach. We calculate the ratio of labor cost to revenue for each firm and use its industry average as the labor elasticity. We compute the capital elasticity by subtracting the labor elasticity from 1.

F.2 Same owner matching

We construct two different matched samples to estimate equation (A4). In the first matched sample, which is the one we use in the main text, a match is made based on whether an incumbent plant of an acquiring firm belongs to the same owner who acquired plant *i*. Thus, comparison plants of acquired plant *i* are incumbent plants that had been managed by the same owner who acquired the plant *i*. We call this matched sample "Same owner matching" sample.

For this matched sample, we use two different weights to estimate (A4). In the main text, we use a simple weight by setting $\omega_j = 1$ for all j so that all incumbent plants of an acquiring firm carry an equal weight. The other weight is the kernel weight. We calculate the distance between an acquired plant and each incumbent plant within the match by using mahalanobis distance. Plant size, plant age, and plant location are used to calculate this distance. Then, we generate a weight for an incumbent plant by using this distance and normal kernel. A large weight is assigned to an incumbent plant when it is similar to the acquired plant in terms of these variables.

Tables A5 and A6 report estimation results using this matched sample with different weighting schemes as above. For comparison, we also include results from the standard difference-in-difference estimation where we just ignore matching. All specifications include acquisition and calendar year fixed effects, as in the main text.

Table A5: Estimation Results from Same owner matching – All acquisitions

	All acquisitions					
	Simple weights		Kernel weights		Standard DID estimation	
	TFPQ (OLS)	Plant ROCE	TFPQ (OLS)	Plant ROCE	TFPQ (OLS)	Plant ROCE
After acquisition	-0.034	0.001	-0.052*	0.0004	-0.043	0.003
	(0.027)	(0.016)	(0.026)	(0.017)	(0.029)	(0.014)
Acquired plant	0.025	-0.030*	-0.014	-0.040**	0.006	-0.028**
	(0.035)	(0.015)	(0.037)	(0.015)	(0.033)	(0.013)
After acquisition x	0.096**	0.050**	0.113***	0.049**	0.104***	0.052***
Acquired plant	(0.036)	(0.020)	(0.037)	(0.022)	(0.039)	(0.019)
Constant	0.036	0.111***	0.061***	0.108***	0.049**	0.107***
	(0.020)	(0.027)	(0.020)	(0.029)	(0.020)	(0.027)
Observations	1,413	1,269	1,127	1,008	1,413	1,269

Note: Robust standard errors clustered at the acquisition-case level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. These symbols apply to all the tables below.

Table A6: Estimation Results from Same owner matching – Serial acquirers

	Serial acquirers					
	Simple weights		Kernel weights		Standard DID estimation	
	TFPQ (OLS)	Plant ROCE	TFPQ (OLS)	Plant ROCE	TFPQ (OLS)	Plant ROCE
After acquisition	-0.014	0.003	-0.064**	-0.002	-0.037	0.005
	(0.026)	(0.019)	(0.026)	(0.021)	(0.034)	(0.016)
Acquired plant	0.017	-0.035*	-0.044	-0.052***	-0.004	-0.026*
	(0.026)	(0.018)	(0.031)	(0.018)	(0.027)	(0.015)
After acquisition x	0.132***	0.071***	0.180***	0.071**	0.152***	0.069***
Acquired plant	(0.038)	(0.023)	(0.040)	(0.028)	(0.043)	(0.023)
Constant	0.099***	0.076***	0.164***	0.093***	0.106***	0.067***
	(0.025)	(0.018)	(0.024)	(0.018)	(0.029)	(0.015)
Observations	1,016	916	787	689	1,016	916

Table A7 presents the estimation results using different measures of TFPQ as described in Section F.1 and simple weights (results using other types of weights are similar). The main results are robust to alternative weights and alternative measures of TFPQ.

Table A7: Estimation Results from Same owner matching –Several TFPQ measures

	All acquisitions and Simple weights			
	Dependent variable: TFPQ			
	OLS	Wooldridge	Blundell-Bond	Index
After acquisition	-0.034	0.064**	-0.045**	0.065**
	(0.027)	(0.027)	(0.018)	(0.026)
Acquired plant	0.025	-0.023	0.018	-0.079**
	(0.035)	(0.033)	(0.024)	(0.031)
After acquisition x	0.096**	0.096***	0.082***	0.111***
Acquired plant	(0.036)	(0.035)	(0.025)	(0.034)
Constant	0.036	2.441***	0.075***	-2.276***
	(0.020)	(0.020)	(0.018)	(0.019)
Observations	1,413	1,413	1,332	1,460

F.3 Pre-acquisition characteristics and trend matching

While matching on the same ultimate owner seems to be the most natural procedure in our case, we also created an alternative matched sample to estimate equation (A4) by forming matches based on whether a non-acquired plant is similar to acquired plant i in terms of pre-acquisition characteristics or pre-acquisition trends of outcome variables. To construct this matched sample, we first specify a group of non-acquired plants that could be potentially matched with each acquired plant. Potential non-acquired plants include all those plants that were owned by acquiring firms and were never acquired themselves, but it also includes plants of firms that just did not participate in the acquisition process at all and also previously or future acquired plants far enough from the time they were actually acquired so that we can consider them to be not affected by acquisition events.[42] Then, we calculate a distance between a particular acquired plant and each non-acquired plant by using mahalanobis distance. Two sets of variables are used for calculating mahalanobis distance. One set is pre-acquisition average value of plant size, plant age, and plant location. The other set is pre-acquisition average value of TFPQ growth and plant ROCE growth. We calculated TFPQ growth rates of each plant by using the four measures of TFPQ and use them to calculate mahalanobis distance. Thus, a small value of this distance indicates that an acquired plant and a non-acquired plant are similar with respect to pre-acquisition TFPQ growth rates. In a similar way, growth rates of the plant-level return on capital employed (plant ROCE) are used to calculate mahalanobis distance. After calculating these distances, a non-acquired plant is included in a particular match only if its distance is below the median distance of the overall sample.[43] The simple weight (i.e., $\omega_j = 1$) is used for this estimation.

Tables A8, A9, and A10 present estimation results using this matched sample. Table A11 shows estimation results for the ratio of plant-level inventory to output, the ratio of fixed capital cost to output, days in operation, and capacity utilization rate. The main results are robust to alternative matching criteria and alternative measures of TFPQ.

Table A8: Estimation Results from pre characteristics and trend matching – All acquisitions

	All acquisitions			
Matching criteria	Plant age, size, location		TFPQ growth rate	Plant ROCE growth rate
	TFPQ (OLS)	Plant ROCE	TFPQ (OLS)	Plant ROCE
After acquisition	-0.020**	-0.009	-0.018**	-0.008
	(0.008)	(0.009)	(0.008)	(0.012)
Acquired plant	-0.011	-0.050***	0.014	-0.033***
	(0.023)	(0.011)	(0.022)	(0.011)
After acquisition x Acquired plant	0.080***	0.050***	0.079***	0.042**
	(0.028)	(0.015)	(0.029)	(0.018)
Constant	0.046***	0.079***	0.032***	0.074***
	(0.009)	(0.011)	(0.010)	(0.011)
Observations	11,521	8,908	11,225	4,155

[42] More specifically, acquired plants in 3 years prior to and 5 years after their own acquisition events are excluded. A plant was also excluded when it does not have any usable observation both before or after the acquisition event.

[43] Other cutoff values such as the mean and lower quartile are used for this estimation, and the results remain unchanged qualitatively.

Table A9: Estimation Results from pre characteristics and trend matching – Serial acquirers

	Serial acquirers			
Matching criteria	Plant age, size, location		TFPQ growth rate	Plant ROCE growth rate
	TFPQ (OLS)	Plant ROCE	TFPQ (OLS)	Plant ROCE
After acquisition	-0.012	-0.008	-0.014	-0.003
	(0.011)	(0.013)	(0.010)	(0.011)
Acquired plant	-0.010	-0.030**	0.012	-0.029*
	(0.026)	(0.015)	(0.029)	(0.015)
After acquisition x	0.137***	0.060***	0.139***	0.054**
Acquired plant	(0.033)	(0.020)	(0.035)	(0.022)
Constant	0.118***	0.052***	0.099***	0.070***
	(0.016)	(0.006)	(0.022)	(0.015)
Observations	5,647	4,444	5,365	2,684

Table A10: Estimation Results from pre characteristics and trend matching – Several TFPQ measures

	All acquisitions			
	Matching criteria: Plant age, size, location			
	Dependent variable: TFPQ			
	OLS	Wooldridge	Blundell-Bond	Index
After acquisition	-0.020**	0.083***	-0.015***	0.084***
	(0.008)	(0.011)	(0.005)	(0.010)
Acquired plant	-0.011	-0.008	-0.010	-0.011
	(0.023)	(0.021)	(0.018)	(0.021)
After acquisition x	0.080***	0.079***	0.060***	0.087***
Acquired plant	(0.028)	(0.026)	(0.019)	(0.025)
Constant	0.046***	2.386***	0.019***	-2.406***
	(0.009)	(0.010)	(0.005)	(0.011)
Observations	11,521	11,521	10,591	11,607

F.4 Placebo test

We also performed a simple placebo test to check that our main findings are not spurious. We randomly assign acquisition status to plants in our sample and estimate how the outcome variables are related to this randomly generated acquisition status. More specifically, we use the same owner matched sample[44] and generate a random variable from the uniform distribution for each plant in the whole matched sample. We assign an acquired plant status to a plant that obtained the maximum value within a particular match. Then, we estimate parameters in equation (4) by using all acquisition cases and simple weights. We repeat this procedure at 1000 times, and calculate a sample mean of estimated coefficients from 1000 time simulations, and their standard errors.

[44] The results were qualitatively unchanged if we conduct the same placebo test using a pre characteristics and trend matched sample.

Table A11: Placebo test

	TFPQ			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	-0.0138	0.0004	-0.0146	-0.0129
Acquired plant	-0.0004	0.0010	-0.0024	0.0017
After acquisition x Acquired plant	0.0037	0.0011	0.0015	0.0058
Constant	0.0624	0.0005	0.0614	0.0634
	Plant ROCE			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	0.0151	0.0003	0.0145	0.0157
Acquired plant	-0.0032	0.0005	-0.0041	-0.0023
After acquisition x Acquired plant	0.0033	0.0006	0.0022	0.0045
Constant	0.1042	0.0002	0.1038	0.1047

Table A11 reports the results from this placebo test. The magnitudes of acquisition and its interaction with the after acquisition dummy effects on outcome variables approach toward zero, and they are economically insignificant.

G. More Decompositions and Input Utilization Details

We can express the logged ratio of total combined capital and input to shareholders' invested capital (the middle term on the right-hand side of equation (4) in the main text) as the sum of two logged ratios:

$$\log\left[\frac{\exp\left(\hat{Y}_i\right)}{C_i}\right]=\log\left[\frac{\exp\left(\hat{Y}_i\right)}{M_i}\right]+\log\left(\frac{M_i}{C_i}\right),$$

where M_i is plant i's total available annual capacity (the number of spindles in installed machines, times 365 days).

The results of this decomposition are presented in Table A12. The top panel indicates that higher combined input to total assets ratios among incumbent plants are mostly explained by higher ratios of combined input to available capacity, the first term on the right-hand side of the decomposition equation above. Given that we measure both capital and labor inputs as flows conditional on operation, this ratio reflects capacity utilization rates—the rates of capital services extracted from the installed machine capacity and the corresponding rates of labor services applied to this capacity. The bottom panels of the table indicate that this capacity utilization metric jumps 9 percent in acquired plants in the first few years after acquisition and by more than 16 percent in the long run. Thus acquisitions immediately lead to higher rates of employment of available capital resources (and correspondingly higher employment of labor, as can be seen from the increases in total labor cost per total assets for acquired plants in Table 6). Ratios of capacity to total assets, on the other hand, are statistically indistinguishable in pre-acquisition years between incumbent and future acquired plants. Moreover, these ratios change little in acquired plants after they are bought, and indeed actually fall in the long run (which, once again, reflects an increase in total assets coming from accumulated retained earnings).

Table A12. Decomposition of plants' total input to total assets ratios:
incumbent and acquired plants and acquired plants pre- and post-acquisition

Pre-acquisition means	Acquired plants (A)	Incumbent plants (B)	Difference (B)-(A)	Percentage difference
ln(total input/total assets)	-0.654	-0.406	0.248	28.1***
of which:				
ln(total input/plant capacity)	-2.564	-2.391	0.174	19.0***
ln(plant capacity/total assets)	1.910	1.984	0.074	7.7
# of observations	115	205		

Pre- and early post- acquisition means of logs	Pre-acquisition (A)	Early post-acquisition (B)	Difference (B)-(A)	Percentage difference
ln(total input/total assets)	-0.599	-0.456	0.143	15.4**
of which:				
ln(total input/plant capacity)	-2.552	-2.463	0.089	9.3***
ln(plant capacity/total assets)	1.953	2.007	0.054	5.6
# of observations	137	130		

Pre- and late post- acquisition means of logs	Pre-acquisition (A)	Late post-acquisition (B)	Difference (B)-(A)	Percentage difference
ln(total input/total assets)	-0.599	-0.560	0.038	3.9
of which:				
ln(total input/plant capacity)	-2.552	-2.401	0.151	16.3***
ln(plant capacity/total assets)	1.953	1.841	-0.112	-10.6***
# of observations	137	231		

Note: The pre-acquisition time period includes observations on up to 4 years prior to acquisition. "Early after acquisition" period includes 3 years immediately following acquisitions. "Late after acquisition" period includes years starting from year 4 after acquisitions. Details of variable construction are explained in the Appendix. . ***, **, and * indicate that the corresponding difference is statistically significant at the 1 percent level, 5 percent level and 10 percent level, respectively, using a double-sided *t*-test.

The differences in the combined input actually deployed in the production process discussed in the main text can be measured even more directly in our data by looking at a plant's spindles utilization rate (defined as the number of total spindle-days the plant operated during a given year divided by the product of the plant's number of spindles installed and 365 days). Table A13 presents this measure separately for acquired and incumbent plants in the four years preceding an acquisition. We also look at the same difference for two labor-capital ratios: one divides the plant's total worker-days by its machine capacity (the number of spindles installed times 365 days), and the other divides it by the actual number of spindle-days the plant operated during a given year.

The top two rows in Table A13 show that both spindle and labor utilization of machine capacity are significantly higher, by about 12-13 percent, among incumbent plants than in acquired plants before acquisition. Our production function estimates suggest modest increasing returns to scale (the sum of the log spindle-days and work-days coefficients is 1.12). Hence the implied output differential per unit of machine capacity is 14-15 percent, which accounts for most of the differential in the input-to-capacity ratios in Table 8. In contrast, to these capacity utilization differences, acquiring and acquired plants operating capital-labor ratios were essentially the same before acquisition. This further confirms the findings that, conditional on operating, the plants operated in similar ways. These results also indicate that the 50 percent higher wage cost rate per unit of shareholders' equity observed among incumbent plants in the first ROCE decomposition in Table 4 in the main text is largely a reflection of higher employment rates, not just higher wages per worker. Daily real wages of both male and female workers (not shown) were also about 14-15 percent higher in incumbent than in acquired plants.

Table A13.
Pre-acquisition incumbent and acquired plants spindle utilization rates and labor to capital ratios

	Acquired plants (A)	Incumbent plants (B)	Percent difference (B-A)
Spindles utilization rate	0.851	0.755	12.6***
Labor to capacity ratio	0.015	0.017	12.3***
Labor to operating spindles ratio	0.020	0.020	-0.5
# of observations	171	227	

Note: The pre-acquisition time period includes observations on up to 4 years prior to acquisition event. Spindle utilization rates are computed as the total number of spindle-days the plant operated during a year, divided by the total spindle capacity (the product of the total number of spindles installed and 365 days). Labor to capacity ratio is the number of worker-days in a given year divided by the total spindle capacity. Labor to operating spindles ratio is the number of worker-days in a given year divided by the total number of spindle-days the plant operated during a year. *** indicates that the corresponding difference is statistically significant at the 1 percent level, using a double-sided *t*-test.

H. In- and out-of-network firms distribution densities of ROCE, unrealized output rates, capacity utilization and prices

Figure A3. Return on capital employed, 1898-1902

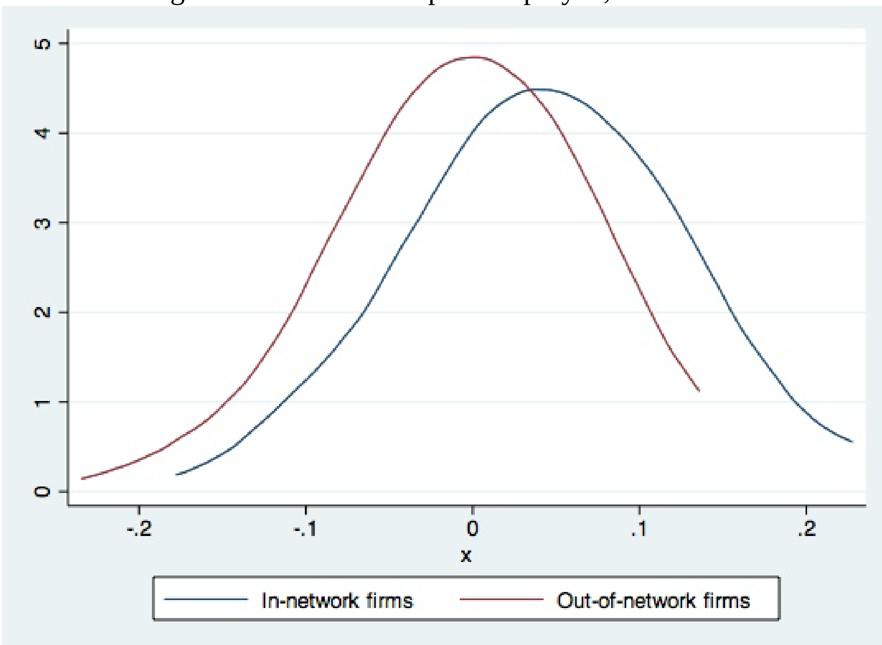

Figure A4. Unrealized output to produced output ratios, 1898-1902

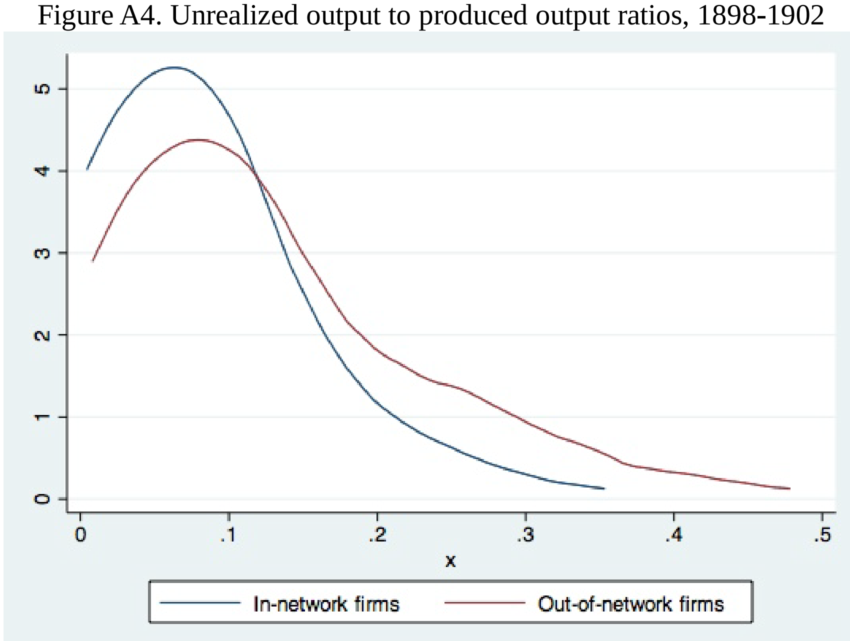

Figure A5. Spindle utilization rates, 1898-1902

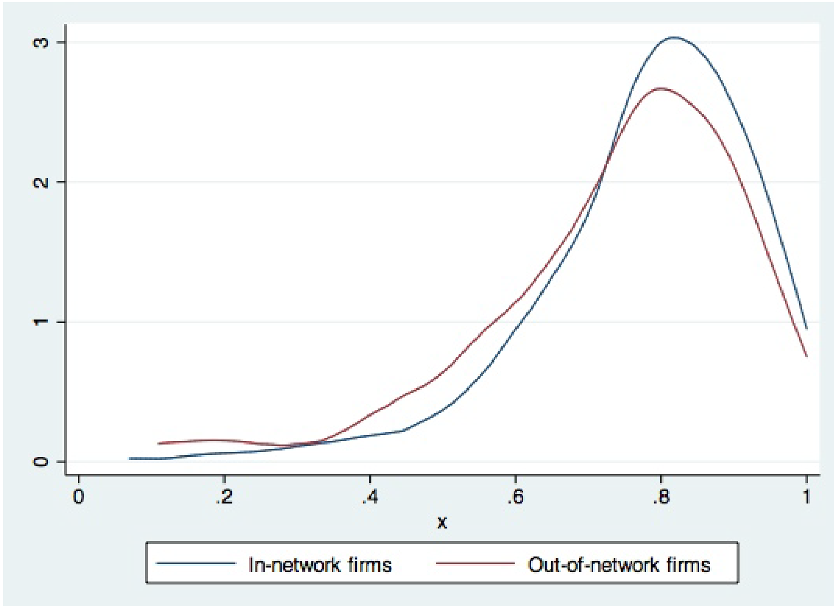

Figure A6. Output prices (yen per unit weight), 1898-1902

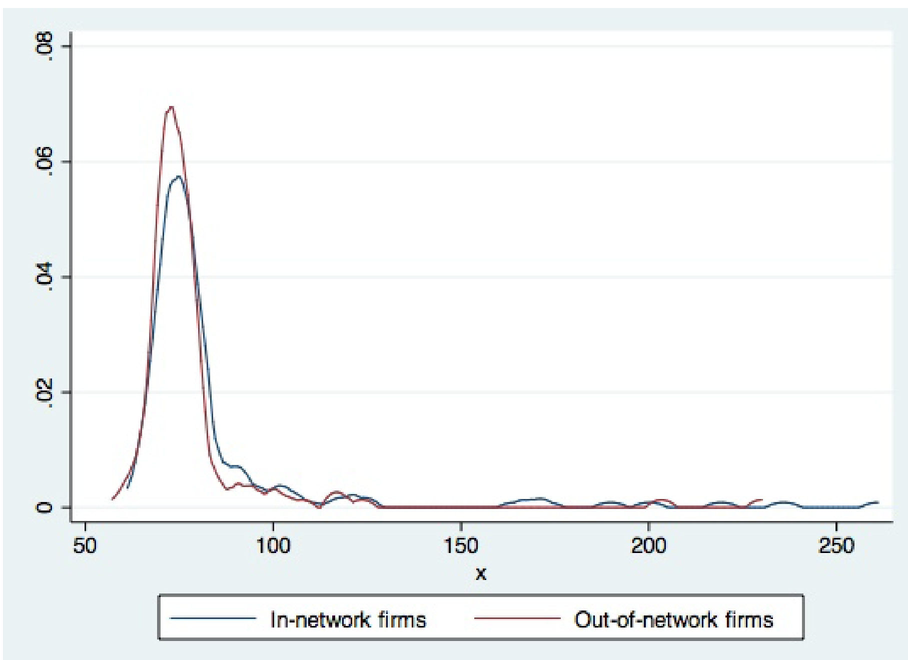

I. Proofs of the results in Section 5

Proof that $u = v$ at the optimum (equation (10) in the main text)

The two first order conditions for the maximization of (9) are given by:

$\frac{\partial \pi}{\partial u} = 0 \Rightarrow \left(p - \frac{1}{\sqrt{uv\omega}}\right) = (\gamma - u - v)\frac{1}{2u\sqrt{uv\omega}}$, and

$\frac{\partial\pi}{\partial v} = 0 \Rightarrow \left(p - \frac{1}{\sqrt{uv\omega}}\right) = (\gamma - u - v)\frac{1}{2v\sqrt{uv\omega}}$.
The claim follows immediately.

Proof of Lemma 1(i): Straightforward from (10) and (12) in the main text.

Proof of Lemma 1(ii): We have

$$\pi(\gamma,\omega) = p\gamma + \frac{2}{\omega} - 2\sqrt{\frac{2p\gamma}{\omega}}.$$

$$\frac{\partial\pi(\gamma,\omega)}{\partial\gamma} = \frac{p}{\gamma}\left(\gamma - \sqrt{\frac{2\gamma}{p\omega}}\right) = \frac{p}{\gamma}(\gamma - m) > 0.$$

$$\frac{\partial\pi(\gamma,\omega)}{\partial\omega} = \frac{m}{\omega}\left(\sqrt{\frac{2p\gamma}{\omega}}\frac{1}{m} - \frac{2}{\omega m}\right) = \frac{m}{\omega}\left(p - \frac{2}{\omega m}\right) > 0.$$

The first two claims follow immediately. Also, $x^* = \frac{2(\gamma-m)}{m\omega} = \sqrt{\frac{2\gamma p}{\omega}} - \frac{1}{\omega}$, which is also clearly increasing in γ.

Proof of Lemma 1(iii): We have

$$\frac{\partial^2\pi(\gamma,\omega)}{\partial\gamma\partial\omega} = \sqrt{\frac{p}{2\gamma\omega^3}} > 0.$$

Proof of Proposition 1: Denote the highest bid price for a potential target plant of quality ω by $\delta\pi(\omega,\gamma_H)$, where $\delta \in [0,1)$ captures any possible transfer costs (rent dissipation) associated with an acquisition. This highest bid price is thus the profit that a manager with ability γ_H can obtain by taking over this plant. The lowest asking price, on the other hand, is given by $\pi(\omega,\gamma_L)$. Assuming that the actual price will be somewhere in-between, an acquisition will be consummated whenever

$$\delta\pi(\omega,\gamma_H) > \pi(\omega,\gamma_L). \qquad \text{(A.4)}$$

Lemma 1(iii) implies that potential gains from an acquisition will be higher if ω is higher, so, other things equal, condition (A.4) is more likely to be met if the potential target is owned by a recent entrant (with quality ω_2) than if it is owned by a first-cohort entrant (with quality ω_1). Condition (A.4) also implies that for any plant quality, an acquisition is more likely to happen when the difference between γ_H and γ_L is large. Once again, this is more likely to happen when an incumbent (first-cohort) firm meets a new entrant (second-cohort) firm than when two incumbents meet. Also, since the ability level of new entrants never exceeds that of incumbents, new entrants never act as acquirers. [45]

Proof of Proposition 3: We show that $\frac{\partial ln[\pi(\omega,\gamma)]}{\partial\gamma} > \frac{\partial ln[TFPQ]}{\partial\gamma}$ for any given ω. We have:

$$ln[\pi(\omega,\gamma)] = ln\left[p\gamma + \frac{2}{\omega} - 2\sqrt{\frac{2p\gamma}{\omega}}\right].$$

Differentiating with respect to γ yields

[45] In reality there were a few cases in our data where new entrants acted as acquirers in the industry. However, in all these cases the acquiring entrants were actually spinoffs from firms of early entry cohort, which apparently inherited the demand management ability of their parent firms (cf. Klepper and Simon, 2000).

$$\frac{\partial ln[\pi(\omega,\gamma)]}{\partial\gamma} = \frac{p - \frac{1}{\gamma}\sqrt{\frac{2p\gamma}{\omega}}}{\pi}.$$

Also, $ln[TFPQ] = \frac{1}{2} ln\left[\frac{\gamma\omega}{2p}\right]$.

Differentiating with respect to γ yields

$$\frac{\partial ln[TFPQ]}{\partial\gamma} = \frac{1}{2\gamma}.$$

Comparing the two,

$$\frac{\partial ln[\pi(\omega,\gamma)]}{\partial\gamma} - \frac{\partial ln[TFPQ]}{\partial\gamma} = \frac{p\gamma + \sqrt{\frac{2p\gamma}{\omega}} - \frac{2}{\omega}}{2\pi\gamma} = \frac{p\gamma\omega - 2 + \sqrt{2p\gamma\omega}}{2\omega\pi\gamma} > 0,$$

because $p\gamma\omega - 2 = \left(\sqrt{p\gamma\omega} + \sqrt{2}\right)\left(\sqrt{p\gamma\omega} - \sqrt{2}\right) > 0$ by (13) in the main text.

Proof of Proposition 4:

Let subscripts A and T denote acquiring and target plants, respectively. The TFPQ difference between the acquiring and the target plants is given by

$$\sqrt{\frac{\gamma_A\omega_A}{2p}} - \sqrt{\frac{\gamma_T\omega_T}{2p}}. \qquad \text{(A.5)}$$

The difference in profits between the acquiring and target plants, on the other hand, is given by

$$\left(p\gamma_A + \frac{2}{\omega_A} - 2\sqrt{\frac{2p\gamma_A}{\omega_A}}\right) - \left(p\gamma_T + \frac{2}{\omega_T} - 2\sqrt{\frac{2p\gamma_T}{\omega_T}}\right)$$

$$= p(\gamma_A - \gamma_T) + \left(\frac{2}{\omega_A} - \frac{2}{\omega_T}\right) - 2\left(\sqrt{\frac{2p\gamma_A}{\omega_A}} - \sqrt{\frac{2p\gamma_T}{\omega_T}}\right)$$

$$= p(\gamma_A - \gamma_T) + 2\left(\frac{1}{\omega_A}\left(1 - \sqrt{2p\gamma_A\omega_A}\right) - \frac{1}{\omega_T}\left(1 - \sqrt{2p\gamma_T\omega_T}\right)\right). \qquad \text{(A.6)}$$

Assume now that the difference (A.5) above is zero. This means that the difference (A.6) boils down to

$$p(\gamma_A - \gamma_T) + 2\left(\frac{1}{\omega_A} - \frac{1}{\omega_T}\right) > 0,$$

which is positive because by Proposition 1, $\gamma_A > \gamma_T$, while $\omega_T > \omega_A$ by the assumption that the target plant has higher quality. We have thus shown that if the TFPQ of the acquiring and target plants are the same, the profit of the acquiring firm will be higher than the profit of the target firm (this also follows directly from Proposition 3, of course). By continuity, the profit of the acquiring firm will still be higher than that of the target firm even for some range of parameters where TFPQ(acquirer) < TFPQ(target). It is also clear from the expression above that this range will be larger when the difference $\gamma_A - \gamma_T$ is larger.

J. A Numerical Example of the Model

Set the value of model's parameters as follows: $p = 3, \omega_{entrant} = 1.5, \omega_{incumbent} = 1, \gamma_0 = 2$. Assume that surviving incumbents' ability, $\gamma_{incumbent}$, is uniformly distributed over the interval [2.45,3.5]. The choice of the lower bound in the distribution of $\gamma_{incumbent}$ makes sure that the lowest-ability incumbent attains the same profits as all entrants, while the upper bound gives the highest-ability incumbent the profit that is twice as large as entrants' profits.

Under these parameters, the optimal choice of *m*, the maximized profit, input utilization and TFPQ are given by the values in Table A14 below:

Table A14. Numerical example: New entrant, low- and high-ability incumbents

	New entrant	Low-ability incumbent	High-ability incumbent
Time managing production	0.94	1.28	1.53
Total input	1.50	1.83	2.58
Input utilization	0.69	0.80	0.87
TFPQ	1.03	0.80	0.87
Profit	1.68	1.68	3.33
Profit/total input	1.12	0.92	1.29

As can be seen from Table A14, high-ability incumbent's profit is double the profit of both new entrant and low-ability incumbent, but his plant's TFPQ is lower than that of a new entrant. Input utilization is the lowest for a new entrant, higher for a low-ability incumbent, and the highest for the high-ability incumbent. These are exactly the patterns we saw in the data.

What happens after a high-ability incumbent acquires a new entrant or a low-ability incumbent in the setup above? Recalculating optimal m using acquirer's ability level, $\gamma = 3.5$ yields the changes presented in Table A-15 below.

Table A15. Numerical example:
New entrant and low-ability incumbent from before to after acquisition by a high-ability incumbent

Acquired:	New entrant		Low-ability incumbent	
	Pre-acquisition	Post-acquisition	Pre-acquisition	Post-acquisition
Time managing production	0.94	1.25	1.28	1.53
Total input	1.50	2.41	1.83	2.58
Input utilization	0.69	0.79	0.80	0.87
TFPQ	1.03	1.18	0.80	0.87
Profit	1.68	4.35	1.68	3.33
Profit/total input	1.12	1.81	0.92	1.29

Under the new, more capable ownership, plants of both new entrants and low-ability incumbents improve input utilization and TFPQ. Profits jump by even more—they double for the low-ability incumbent plant from before to after acquisition, and increase 2.6 times for the plant formerly owned by a new entrant. Even when normalized by total input, the profit rate improves by more than TFPQ, again consistent with the patterns we discovered in our sample.

K. Year-by-year estimates of acquired plants' TFPQ changes before and after acquisition events and within-acquisition comparisons between incumbent and acquired plants

We first present here the results of estimating an equation similar to (1) in the main text but with the full set of pre- and post-acquisition time dummies, so that as to see both pre- and post-acquisition time trends. The estimation equation is:

$$y_{it} = \alpha_0 + \beta_{1t} k_{it} + \beta_{2t} l_{it} + \gamma_- D_{i-} + \sum_{t=T-3}^{t=T-1} \gamma_{it} D_{it} + \sum_{t=T+1}^{t=T+5} \gamma_{it} D_{it} + \gamma_+ D_{i+} + h_i + \mu_t + \varepsilon_{it}, \qquad \text{(A.7)}$$

where T is the year of acquisition, and D_{i-} and D_{i+} are the dummies equal to 1 for plant i in years up to 4 years prior to the acquisition year and zero otherwise, and equal to 1 for plant i in years 6 and beyond after the acquisition and zero otherwise, respectively. The estimates are carried out using observations on all productive establishments that changed ownership between 1898 and 1920.

Figure A7. Within-acquired plants changes in TFPQ from before to after acquisition

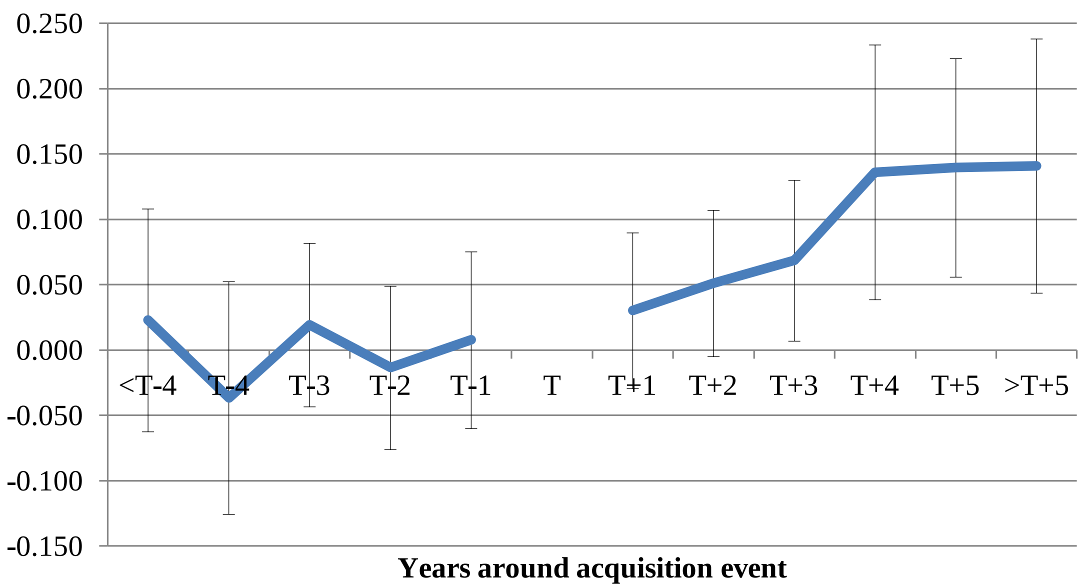

Figure A7 presents the coefficients on the dummies from the equation (A.7) along with the 95 percent confidence intervals (the omitted year is the year of acquisition). Consistent with the results in Table 2 in the main text, the TFPQ of acquired plants are statistically and economically indistinguishable from zero during the pre-acquisition period. After acquisitions, TFPQ starts improving, gradually in the first 3 years and then plateaus after the fourth year at around 0.12-0.14.

We also estimated a regression similar to the "difference-in-difference" regression (2) in the main text, also with a full set of time dummies:

$$\bar{y}_{it} = \alpha_0 + \sum_{t=T-4, t\neq o}^{8} \beta_s \overline{Inc_{is}} + \sum_{t=T-4}^{8} \beta_s Acq_{is} + m_{it} + \varepsilon_{it}\,, \qquad \text{(A.8)}$$

where, as in the main text, $\bar{y}_{it}$ is TFPQ (relative to industry-year average) of plant i at time t if it is an acquired plant, while TFPQs (also relative to industry-year average) of incumbent plants are collapsed to $\bar{y}_{it} = \frac{1}{\#m_A} \sum_{j \in m_A} y_{jt}$, where m_A denotes the particular acquisition case in which plant i was acquired and $\#m_A$ is the number of incumbent plants in acquisition m_A. The timeline is, once again, from 4 years before to 8 years after acquisitions.

Figure A8. Within-acquisition TFPQ of acquired and incumbent plants

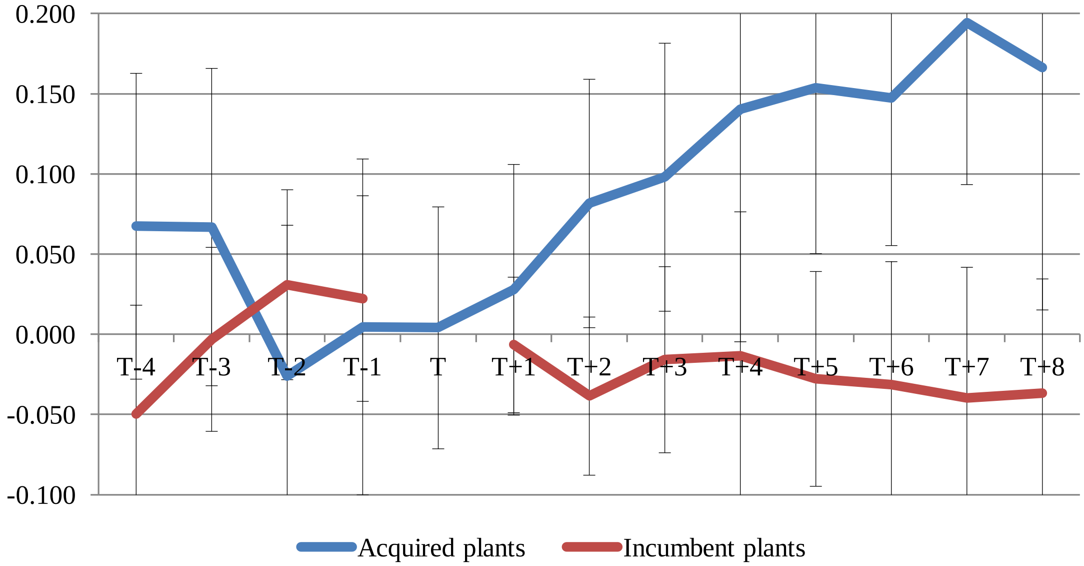

Figure A8 presents the coefficients on the dummies from the equation (A.8), along with the 95 percent confidence intervals. (The omitted category is TFPQ of incumbent plants in the year of acquisition, so all other variables are measured relative to the incumbent plants' average TFPQ in the acquisition year.) Consistent with the results in Table 3 in the main text, TFPQ of acquired plants is somewhat higher than TFPQ of incumbent plants before acquisition, but the difference is not statistically significant. There is no particularly pronounced trend in incumbent plants' TFPQ both before and after acquisitions, but acquired plants clearly diverge upward from incumbent plants (and the rest of the industry—recall that TFPQ are the residuals from production function estimates using all available data for all years, including also year dummies) after acquisitions.