CIRJE-F-945

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

Serguey Braguinsky
Carnegie Mellon University
Atsushi Ohyama
Hokkaido University
Tetsuji Okazaki
The University of Tokyo
Chad Syverson
University of Chicago and NBER
October 2014

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

September 2014

Abstract

We explore how changes in ownership and managerial control affect the productivity and profitability of producers. Using detailed operational, financial, management, 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 any 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 inherent uncertainties of demand in the industry. When these less profitable plants were purchased by more profitable establishments, the 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.

---

* We thank conference participants at the 2013 ASSA Meetings, the Stanford Conference on Japanese Entrepreneurship, and the NBER Summer Institute 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.[2] 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.[3]

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 in increased productivity measures 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

---

[2] 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).

[3] 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).

data set that allows us to investigate the production and input allocation processes at an unusual 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 opening of the 20th Century). These data, which we describe in the next section and in the appendix, 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. 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. While illustrative, we note this time allocation mechanism is not the only possible source of differences in demand management, nor is it directly testable in our data, as we have no information on the allocation of owners' or managers' time. Any mechanism that creates disparity in companies' abilities to manage the demand uncertainty inherent in this industry can create the patterns we document.

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.

Our setting offers additional advantages besides detailed data and a novel mechanism. 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).[4] 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. As late as in 1887, domestically produced output was still a fraction of imports, but it exploded over the following decade. 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).[5]

[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 industry was also hampered by low quality of domestically grown cotton (Chinese-grown cotton was also used, but it was not much better) and by the choice of what turned out to be inferior equipment (mules instead of ring spinning frames).

The second stage, that of explosive growth in the 1890s, was ushered in by two major technological breakthroughs: the switch to higher-quality raw cotton imported from India, and the adoption of ring spinning frames (a new type of cotton spinning machinery first invented in the U.S.).[6] The success of early experiments using the new technological paradigm led to wide-spread emulation, with entrants purchasing only ring frames for their newly constructed mills and using almost exclusively Indian cotton (and later, even higher quality raw cotton from the U.S. and

---

[4] See also Saxonhouse (1974), (1977); Miwa and Ramseyer (2000).

[5] Import tariffs on cotton yarn remained negligible throughout this period as Japan only regained full control over its import tariffs in 1911. Export tariffs on cotton yarn were abolished in the early 1890s, and this did help exports somewhat. However, the tariff rates had been just a few percent to begin with, so they were hardly a decisive factor.

[6] See, e.g., Ohyama, Braguinsky and Murphy (2004) and Braguinsky and Rose (2009) for detailed accounts of these developments.

Egypt).[7] 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, the average size of the plant almost quadrupled, and the average number of workers per plant rose fivefold. Industry output in physical units increased 17 times 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.[8] 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 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).[9] This does not mean, however, that the playing ground was equal for all firms. In order to sell their output, firms had to connect to trading houses. This wasn't an easy task, especially during anticipated business downturns when large established trading houses would often limit their purchases to reputable producers with whom they had a long-term relationship (Takamura, 1971, Vol. 2). Of course, firms could (and did) sometimes offer to sell at discounted

---

[7] The Japanese cotton spinning industry eventually led the world in the speed of ring frames adoption, with even the U.S. behind it; see Otsuka, Ranis, and Saxonhouse, 1988.

[8] Jovanovic (1998) and Jovanovic and Yatsenko (2011), among others, offer models where firms rationally stick to older-vintage capital even after newer vintage has become available.

[9] 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.

prices in situations like this, but this practice appears to have been limited by reputational considerations (all price information was public and even hidden discounts, such as increasing the amount of yarn delivered over and above the contracted amount, were often reported—and decried—by the journal published monthly by the All-Japan Cotton Spinners' Association; hereafter "Geppo"). 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 (Section 4) that these problems were indeed quite severe, and that the most successful early entrants (who later became major acquirers of production facilities 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. A temporary reprieve came in the form of burgeoning exports (mostly to China), but 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, which is 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. There was virtually no new entry or plant construction. While entry and new construction eventually resumed by the end of the decade, acquisitions continued to play an important role.

[Figure 2 about here]

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.[10] In the appendix we present an example

[10] How a functioning market for assets emerged so early in the process of economic modernization is a subject for a

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). 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). However, the intuition behind the underlying theory (for which Jovanovic and Rousseau find support in U.S. data) 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 actually a potentially cheaper alternative for Japanese firms looking to expand.[11]

These factors led to the consummation of 73 distinct acquisition deals in the industry between 1898 and 1920, during which 95 plants changed hands (sometimes more than once). 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 out 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).[12] Figure 3 traces the dynamics of the fraction of

---

separate study; see Miwa and Ramseyer (2000) for some insights on this issue.

[11] In support of this, Saxonhouse estimates that the time lag between receipt of spinning mill orders from Japan and shipments of equipment ranged from one to two years for most of the 1890s and the 1900s. It increased to 3-5 years after the start of World War I (Saxonhouse, 1971, p. 51).

[12] 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 Bosseki and Amabo were both founded in 1889. The former three firms also struggled with older-vintage capital equipment for some time.

the industry's plants, spindle capacity, and output that were owned by the above five firms (which shrank to four after the 1914 merger and to three after the 1918 merger) over the 25-year period of our analyses. These five firms went from owning 10 percent of the plants and 25 percent of industry capacity and output at the beginning of the period to 40 percent of plants and half of capacity and output by the end. As no other firm owned more than 10 plants or had more than 10 percent of industry-wide capacity under its ownership at the end of our period, these serial acquirers stood out by their sheer size and importance. The 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. It is also worth noting that the three super large firms all survived long into the post-World War II era.

[Figure 3 about here]

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.[13] For this paper, we have collected and processed all the available data between 1899 and 1920 (prior to 1912 these data are also available in the *Statistical Yearbook of the Ministry of Agriculture and Commerce* ("*Noshokomu Tokei Nempo*")). Since the first acquisition of an operating plant 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 (hereafter "Boren"). Our data thus cover 1896 to 1920. Saxonhouse (1971, p. 41) writes that "the accuracy of these published numbers is unquestioned."[14]

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 spindles in operation, number of days the plant

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

[14] We scrutinized these numbers ourselves and 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 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 dropped about 5 percent of observations where the annual data contained in government statistical reports could not be corrected.

operated and the average number of hours per day, the number of factory floor workers (male and female workers separately), average daily wages separately for male and female workers, 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) and their total horsepower, output of the finished product (cotton yarn) in physical units, the average count (measure of fineness) of produced yarn, and the average price per unit of yarn produced.[15] 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 coming from semi-annual reports issued by the firms that owned the plants. Those reports 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. The reports also contain qualitative information about the deaths, illnesses, resignations, and replacements of board members. For the purpose of this study we have digitalized and processed the data from all the reports between 1896 and 1920 currently preserved by the Osaka University library (over 1,200 total).[16] Some basic (although not as detailed as in company reports) firm-level financial data were also published in the semi-annual Boren's publication "*Reference on Cotton Spinning*" ("*Menshi Boseki Jijo Sankosho*") which started in 1903. These data were used to supplement company reports for privately held firms and also in cases where company reports were missing.[17] We also collected detailed information surrounding each acquisition and ownership turnover event, including but not limited to identities and backgrounds of the most important individuals involved (shareholders, top managers) from company reports, supplementing with the information in the seven-volume history of the industry written in the late 1930s by the Japanese historian Taiichi Kinugawa (Kinugawa, 1964) as well as company histories.

---

[15] See Foster et al. (2008) and Syverson (2011) for the discussion of the importance of separating quantity and revenue productivity and the difficulties encountered by researchers trying to do it using conventional data that contain sales and expenditures on inputs but not direct evidence on the quantity of inputs and outputs. Atalay (forthcoming) similarly discusses the importance of separating quantities from expenditures when measuring inputs.

[16] Some selected reports have been used in previous research but to the best of our knowledge, we were the first to systematically digitalize and process all the available reports, including those by smaller firms.

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

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."[18] While output of cotton yarn in our data is measured in weight, we also observe the average count produced by a given plant in a given year. To make different counts comparable for the purpose of productivity analysis, we converted various counts to the standard 20th count using a procedure detailed in the appendix. Second, we used plant-year-specific female-to-male wage ratios to convert units of female labor to units of male labor.[19] 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.[20] 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 Saxonouse, yarn production is essentially Leontief in raw cotton and other inputs when both input and output are measured in units of weight (the raw correlation between the two variables in our data is 0.95). As

---

[18] 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. Producing higher-count (finer) yarn generally requires more skill and superior technology than lower-count (thicker) yarn. High-count yarn is often 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.

[19] 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. 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 gender. All our estimates are robust to including the number of male and female workers separately in the production function estimations.

[20] While spinning frames (and their spindles) account for only 25-30 percent of the total equipment cost of a mill, the correlations in our data between spindles and other equipment (cards, draw frames, slubbing frames, intermediate frames, roving frames, etc.) are over 95 percent. Thus spindles are a good proxy for equipment as a whole (Saxonhouse, 1971, p. 55-56).

a practical matter, including raw cotton in a log-linear production function therefore renders all other inputs economically and statistically insignificant. One can therefore 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 1 presents year-by-year counts of acquired plants during our sample. The total is broken out separately by plants taken over by the five largest (serial) acquirers. On average, 4.3 percent of the industry's mills were acquired per year, with the largest acquirers responsible for about 40 percent of all acquisitions.[21] These acquisition episodes form the base of our estimation sample.

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. 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 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.[22] To measure profitability, we calculate

---

[21] 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).

[22] Note 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

shareholders' return on equity; that is, we divide firms' profits by the amount of equity capital paid in by shareholders.[23] 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 spindles were purchased.[24] Firm age, on the other hand, is always equal to the calendar year minus the year the firm was founded.[25]

Table 2 shows the means and standard deviations of the aforementioned plant characteristics for each group of firms: future acquiring firms, future target firms, and future exiting (not by acquisition) firms. We further separate plants of target firms into those that started 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.[26] The table includes only data from 1896-97—that is, before any acquisitions took place in the industry.

[Table 2 about here]

---

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 then calculate capacity utilization—how often the plant is actually operating during the year—separately and explore the two metrics jointly in our analysis.

[23] 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 invested. See below.

[24] For example, if the plant's initially installed spindles were purchased in year t and then the plant underwent an expansion during which the same quantity of new spindles were purchased in year $t+k$, plant age is calculated as the calendar year minus t until the year new equipment is installed, after which it becomes the calendar year minus $[t+(t+k)]/2$, the average vintage age of equipment (or the weighted average if the number of spindles installed later were different from the number initially installed).

[25] Since the plant's capital stock includes various pieces of machinery as well as buildings, engines, and various elements of infrastructure, equipment (spindles) age adjusted for vintage as above 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.

[26] In the appendix, we use additional data on firms' orders of specific pieces of capital equipment to measure how the machines' technical specifications evolve over time. We find clear evidence of pre- and post-early 1890s differences in technological capabilities along multiple dimensions: spindle rotation speed, spindles per frame, the quality of yarn for which the machines are calibrated, and the ability to handle multiple yarn counts and cotton types.

Looking across the table's top row to compare the average physical productivity levels across the groups of plants, we see that in contrast to the "higher productivity buys lower productivity" scenario, 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 third row of the table 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: the pricing process did not reflect large differences in market power 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. As we described in Section 2, capital vintage effects (particularly for equipment spanning 1890, the year we use to split cohorts) were important in the industry. 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 as plants in future acquiring firms. 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.

To measure these changes, we estimate specifications that regress acquired plants' 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 do not include the acquisition year itself in the regression because acquisitions often happen mid-year, 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 differences in attributes. We also include calendar year fixed effects to remove any systematic changes in attributes over the sample.

Thus our estimating equations have the general form:

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

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; μ_t is a year fixed effect; and ε_{it} is the error term.

The first numerical column of Table 3 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)[27], its number of spindle-days in operation (flow of capital services), and the change in log plant capacity from the previous year (control for equipment installation adjustment costs).[28]

---

[27] We also conducted all estimates using male and female work-days separately and the results were almost identical.

[28] Even though our data also contain records of the average number of hours plants operated per day in a given year,

The results indicate that in the first 3 years after acquisition, acquired plants' TFPQ levels are 3.2 percent higher and not statistically different from their levels in the pre-acquisition years. In subsequent years, however, the TFPQ of acquired plants rises to a level 11 percent above their pre-acquisition baseline, and the hypothesis that the coefficients on the early and late post-acquisition dummies are equal is rejected at the 1 percent confidence level. Thus acquired plants' TFPQ levels do improve considerably following an acquisition, although it takes time for this to manifest itself fully.

[Table 3 around here]

Table 3's second column looks at acquired plants' profitability around acquisition episodes. Unfortunately, we cannot directly evaluate plant-level changes in profitability from before to after acquisition that are analogous to the cross sectional comparisons in Table 2. This is for the obvious reason that there are no separate post-acquisition firm profit accounts. We work around this issue by using plant-level gross operating surplus, computed as the difference between the value of output produced (output times the plant-specific price) and input costs, where the latter equals the sum of capital costs (capital invested, including equity, various borrowings, and corporate bonds, times the market interest rate) and labor costs (number of male and female work-days, multiplied by the corresponding daily wages). We convert this gross surplus to a rate by dividing it by input costs. Since the data on capital invested are available at the firm level only, for multiple-plant firms we assign firm-level capital cost to each plant in proportion to the fraction of its revenue (plant-specific output, times plant-specific price) in total firm revenue. We estimate (1) with (logged) gross operating surplus as the dependent variable.[29]

---

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. 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 3-5 below.

[29] The raw correlation between the plant-level gross operating surplus rates constructed in this way and the ROC measure (defined in Section 4.2.1 above) for pre-acquisition years is 0.7, so our measure is a reasonably good proxy

In contrast to the TFPQ patterns in the first column, acquired plants' gross operating surplus jumps immediately after acquisition, and by a lot: over 18 percent within the first 3 years post acquisition. It rises yet another 13 percent in subsequent years (the difference between the early and late post-acquisition coefficients is also statistically significant). Hence, profitability improves even more than productivity following acquisition and, in contrast to TFPQ improvement, more than half of the change is observed soon after the acquisition.

Finally, to see if changes in plant-specific price contributed to profitability changes, we estimated a regression similar to (1) with the dependent variable being the (logged) plant-specific price, divided by the main count of yarn produced by the plant to adjust for quality differences. The table's last column shows 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 next see whether these changes within acquired plants are systematically related to the attributes of the acquiring firm. While acquiring firms could be quantitatively 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 repeat the specifications in Table 3, but with the sample limited to only acquisitions conducted by one of the five serial acquirer firms. The results are presented in Table 4.

[Table 4 about here]

The picture offered by the table is qualitatively similar to Table 3, but all the changes are indeed more pronounced. In particular, acquisitions by serial acquirers correspond to long run improvements in acquired plants' physical TFPQ of more than 20 percent, while profitability increases by more than 60 log points. The point estimates on the price changes are larger than in Table 3, but *t*-tests fail to reject at conventional confidence levels the hypothesis that the coefficient on either of the post-acquisition dummies are equal to the pre-acquisition dummy coefficient.

Overall, the within-plant results in Tables 3 and 4 indicate that acquired plants see growth in both their TFPQ and profitability levels after they are acquired, though the latter occurs sooner

---

for plant-level profitability. Using logs does not result in losing observations because while firm-level profits can be negative, plant-level operating surplus gross of expenses on intermediate inputs, as constructed here, is always positive in the data.

and is larger in magnitude. 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, in effect using the incumbent plants as a control group. Including incumbent plants results in the loss of data because in 37 acquisitions the acquirer came from outside the industry and hence had no incumbent plants. Additionally, the timelines of available data on some incumbent plants were missing or too short to be usable. Therefore, the exercise here is limited to only 58 out of 95 total acquisitions in the sample. 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.[30]

The specification is as follows:

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

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 $\bar{y}_{it} = \frac{1}{\#m_i} \sum_{j \in m_i} \omega_j y_j$, where m_i denotes the particular acquisition case in which plant *i* was acquired and $\#m_i$ is the number of incumbent plants in acquisition m_i. 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_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 the acquisition-year fixed

[30] We also conducted several robustness checks using other control groups created using various matching techniques. These yielded similar results and are described below and detailed in the appendix. Here, to avoid problems stemming from the fact that previously acquired plants by serial acquirers are already "incumbent" plants when another acquisition happens (which can be as early as in the same year), we impose a rule that a previously acquired plant only becomes labeled 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, however.

effect. In the main text, we assign weights $\omega_j = 1$ to all incumbent plants in a given acquisition m_i, which allows us to interpret coefficients β_1, β_2, and β_3 similar to the interpretation given to them in the standard difference-in-difference estimations (see Appendix for explicit derivations). In particular, estimate $\hat{\beta}_3$ 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 acquisition event, but reasonable alternative timeline cutoffs produce similar results. [31]

The estimation results are presented in Table 5. The first two columns of numbers reflect TFPQ and gross surplus 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 $\hat{\beta}_3$ are positive and statistically significant at the 1 percent level. The results indicate that the post-acquisition improvement of TFPQ of acquired plants (this time relative to incumbent plants) is about 12.2 percent on average for all acquisitions and 16.9 percent for acquisitions by serial acquirers. (Excluding serial acquirers, the estimated value of $\hat{\beta}_3$ is just about 5 percent and statistically not significant at conventional levels; partly this is due to the fact that we have relatively few observations on acquisitions carried out by non-serial acquirers where the acquiring firm also had usable incumbent plant data.) In addition, the coefficients on the acquired plant dummy 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.

[Table 5 about here]

In the profitability regressions, $\hat{\beta}_3$ is also positive and statistically significant. Profit rates of acquired plants rise by about 22 percent relative to the plants in the acquiring firms in the whole sample and by about 28 percent in the sample of acquisitions by serial acquirers. Here, acquired plant dummy coefficients are large and negative, reflecting the profitability deficit of acquired

---

[31] We also estimated equation (2) in two more different ways, where we employ (1) kernel weights obtained from the mahalanobis distance measure where acquired and incumbent plants are matched on plant size, age and location; and (3) a standard difference-in-difference estimation procedure ingoring acquisition-based matching altogether. The results of all these estimations are very similar to those presented in Table 5 in the main text (see Tables A3 and A4 in the Appendix).

firms prior to acquisition.[32]

These results further reinforce what we document above: acquisition is accompanied by growth in the acquired plants' productivity and profitability levels. We see here that this is true relative not only to the plants' own levels before the acquisition, but also relative to changes within incumbent plants owned by their acquiring firms.[33]

4.4 The Link from Profitability to Productivity: Demand Management

In principle, the pre-acquisition gap in profitability without a comparable gap in TFPQ might suggest price differences due to market power between acquiring and acquired firms. The results above, however, offer no support for this hypothesis. As discussed in Section 2, this is consistent with what we know about how the industry's output market operated at the time.

But as also mentioned in Section 2, a lack of price differentiation does not mean output-market conditions were equivalent across firms. Stronger companies may have been able to manage the industry's inherent demand variations betters.

To quantitatively explore possible differences in firms' demand-facing operations, we investigate plants' finished goods inventory and capital utilization metrics before and after acquisitions. We choose these metrics because holding finished product in inventory may indicate that the plant is having difficulty finding buyers in a timely manner, and because capital utilization differences—which may also reflect poor management of matching production to demand or difficulty in finding buyers—may explain a considerable portion of the profitability differences across plants (as hinted at in Table 2 above). Indeed, anecdotal evidence from company histories suggests that inventories and utilization were intrinsically linked in the industry, as firms would often halt production as unsold yarn and uncollected revenues piled up, and would resume only

[32] We also estimated a regression similar to (2) with the outcome variable being count-adjusted plant output price (relative to the industry-year average). As in the previous subsection, we did not find any significant differences between pre- and post-acquisition trends in either acquired or incumbent plants.

[33] While matching by acquisition cases seems to be the most natural way in our context, we also conducted robustness checks where we 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 A7 and A8 in the Appendix.

after the gridlock had cleared. We again use the within-acquisition difference-in-difference specification (2) so as to characterize whether acquired plants converged toward incumbent plants after they were bought.[34]

Table 6 presents the estimation results. The first column looks at producers' ratios of period-end finished goods inventories to the value of their output over the period. (Plant-level inventories are imputed from firm-level reports in post-acquisition observations by assigning a firm's inventories across its plants in proportion to the plant's share of firm output.) The second through fourth columns of the table look at three measures of plants' capital utilization: the ratio of their fixed capital costs to their output (fixed capital costs are imputed from firm-level reports in post-acquisition observations in proportion to the plant's share in the firm total capacity), the number of days the plant is in operation during the year, and the plant's capacity 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). The latter two measures are observed directly at the plant level and do not involve any imputations (nor does the pre-acquisition difference between acquired and incumbent plants on the former two measures).

[Table 6 about here]

The estimates in the table show that acquired plants were notably inferior to incumbent plants on each of these dimensions before they were purchased. The "acquired plant" indicator coefficients imply that, prior to acquisition, these plants held higher inventories, had higher unit fixed capital costs, operated fewer days out of the year, and had lower capacity utilization. All these deficits are statistically significant and economically large. These are strong indicators of these plants' lower pre-acquisition profitability levels, and are consistent with the demand management mechanism discussed above.

Once the plants were purchased, the gaps closed along every dimension. In fact, the acquired plants completely caught up to incumbent plants in inventories-to-output ratio and days in operation, as witnessed by the fact that the interaction term's coefficients are of similar

[34] We also looked at unit labor costs (measured as the plant-level wage bill over output) and found them to be remarkably similar between acquired and acquiring plants (details are available upon request). This leaves unit capital costs (capital outlays per unit of output produced) with both fixed and inventory costs as primary components as our main "suspects." We note for the sake of completeness that we found no evidence that intermediate goods and unfinished products systematically influenced outcomes, and thus we do not consider them here.

magnitude to the acquired plant indicator. Thus it appears that while acquired plants suffered a relative inability to manage sales, their purchase by firms with higher-ability owners and managers led to the transference of this capability to the acquired operations, leading to the gains in profitability observed in the previous sections.[35]

The ability of cotton spinning firms to manage demand could arise from many possible sources. While many of these are difficult to quantify, our reading of the industry's history suggests that one important factor was how well the firms were connected to the trading houses that purchased their output. As discussed above, while prices varied little across industry producers, in low-demand times demand was rationed by the trading houses delaying or refusing to make purchases from certain producers. It is possible, then, that having close relationships with trading houses could allow some industry producers to sustain more consistent operations, resulting in the lower inventories and higher utilization levels observed above.

To explore this possibility further, 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 be best connected to cotton spinners' output markets. This yielded a list of 154 individuals.[36] 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.[37] Of a total of 1197

---

[35] Once again, we conducted robustness checks using different weighting schemes and different matching methodologies. The results, presented in Tables A6 and A10 in the Appendix are similar to those in Table 6.

[36] 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 worth of operation tax in that year. Not surprisingly, geographical distribution is heavily biased, with 35 of these traders located in the Osaka prefecture, 19 more in Tokyo, 7 in the Aichi prefecture (Nagoya) and 6 more in Kyoto. No other location had more than 3 large traders. 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; the first 3 firms had headquarters in Osaka, while Mitsui Bussan's headquarters were in Tokyo). Finally, the third group includes the 31 board members and traders registered at the Osaka cotton and cotton yarn exchange.

[37] This sample covers almost 90% of all firms that operated in that year. For 2 of those companies the earliest available reports are in 1899 and 1900, respectively. After consulting company histories (Kinugawa, 1964) we determined that board members and top shareholders linked to traders and trading houses had been associated with these companies since they were founded. No single firm drives the results in any event.

board members and top shareholders, we found 128 who were in the list of the 154 top traders described above. Of the 67 firms, 33 had at least one connected trader among its board members and top shareholders. We created a "trader network" indicator equal to 1 if the firm was one of these 33 (we refer to these as "in-network" firms) and 0 otherwise (these are "out-of-network" firms). For in-network firms, we also calculated the percentage of shares in those firms that they held (including family members who are not treated as distinct individuals but whose shares are summed up in the total number of shares controlled by traders). Conditional on being positive, this percentage varies from 0.7 to 51.2 percent.

We test whether these measures of a producer's connectedness to trading houses is reflected in the demand performance metrics we explored above. Specifically, we regress plant-level performance metrics for the years 1898-1902 separately on the trader network indicator and the percentage of shares held by connected traders located in Osaka or Tokyo (by far the two most central cities for cotton trade). We include the (log of) firm age and year dummies as additional controls.

Table 7 presents the estimation results, with panel A using the trader network indicator to measure connectedness, panel B using the asset shares of connected traders. The trader network indicator is associated with a large (on the order of 40 percent) drop in a plant's average inventory-to-output ratio. While the point estimates indicate being in-network is tied to lower unit capital costs and higher utilization, none of these are statistically significant. Measuring connections to trading houses using the percentage of shares held by traders, as in panel B, shows that connectedness is again related to significantly lower inventory levels and, here, higher capacity utilization as well.

[Table 7 about here]

These results suggest that having connections to trading houses and exchanges does allow a producer to manage demand fluctuations more effectively, particularly with regard to being able to operate with lower average inventory levels and perhaps at greater capacity utilization levels as well.[38] That said, it also appears that trade network connects are not the sole source of demand management differences across firms, as considerable variation in capital costs and utilization rates remain unexplained.

---

[38] We also tested whether connections were related to the price at which a plant sold its output. As with all of our findings on prices, we found no systematic relationship.

This and 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 also consistent with the profitability story above if demand management is correlated with broader managerial ability. In this case, acquisition by higher-ability firms also acted to raise the productive efficiency of acquired plants, as reflected in the TFPQ results.

4.5. Do Better Acquiring Firms Improve Acquired Plants More?

In the above results that focus on acquisitions by the five serial acquirers, we have already seen that, at least at a rough cut level, being bought by a higher ability acquirer results in larger productivity and profitability improvements. In this subsection we take one step further and try to measure acquirer ability in a more continuous way and see if it is correlated with productivity improvements.

If, as consistent with the results in the previous subsection, the inventory and fixed capital cost to output ratios are functions of owner/manager ability, then this suggests these values may be used to proxy for such ability. Specifically, if the ratios are monotonically decreasing in ability, we can invert this function to control for ability. That is, $ability = g_1^{-1}\left(\frac{inventory}{output}\right)$ and $ability = g_2^{-1}\left(\frac{fixed\_capital\_cost}{output}\right)$. By estimating regressions similar to (1) while including functions of the two ratios for the acquirer, we can see the extent to which the estimated post-acquisition improvement in acquired plants' TFPQ is explained by acquirer ability.

Table 8 presents the results of this estimation. It shows regression specifications of the form (1), comparing results both with and without including log of the acquirer's inventory or fixed capital cost to output ratios as controls for acquirer ability.

[Table 8 about here]

When controls for acquirer ability are included (the table's second and third columns), the magnitudes of coefficients on all post-acquisition dummies drop noticeably and become economically and statistically indistinguishable from zero. Additionally, the coefficients on the ability control functions are negative and economically and statistically significant, indicating that increases in both ratios reduce productivity, which is consistent with the notion that TFPQ is positively associated with ability.

4.6. Robustness

As already mentioned, we have conducted several checks on the robustness of our results. We relegate the details to the appendix 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 that have been used in the literature to avoid the 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 the Levinsohn and Petrin, 2003, estimator). 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 the Appendix) 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, 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*. 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 the appendix. 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. (We conducted the exercise on both our same-owner and pre-acquisition-characteristics matched samples. The procedure and results are explained in detail in the appendix.) We repeated this process 500 times and calculated the sample mean of the estimated coefficients relating "acquisition" to outcomes. In all cases, the magnitudes were only fractions of their analogs from the true acquisition samples and were economically insignificant (see Appendix).

5. A Mechanism

Our empirical results point to some sort of demand management ability (reflected empirically in inventory and capital utilization levels) 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 (and increasing sales in the process) but at the cost of spending less time managing production (decreasing productivity in the process). 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:

$$y = u\omega x \tag{3}$$

where $u \in [0,1]$ is the fraction of the manager's time allocated to managing the plant, 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). We assume that the firm first chooses x to minimize the cost of producing a given y and then optimally chooses u and y. Since the input x is the only choice variable in the cost minimization problem, the optimal choice of x is given by

$$x^* = \frac{y}{u\omega},$$

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

We assume the plant faces the following isoelastic demand:

$$q = \lambda(u;\gamma)p^{-\sigma}, \qquad (4)$$

where q is quantity demanded, p is the output price, and $\sigma > 1$. Note that demand also depends on u, the manager's time allocation. This is the channel through which we introduce the notion of demand management. From (3) and (4), standard profit maximization solution leads to the optimal price:

$$p = \left(\frac{\sigma}{\sigma-1}\right)\frac{1}{u\omega},$$

and the optimal profit:

$$\pi(u;\gamma) = \mu(\omega;\sigma)\lambda(u;\gamma)u^{\sigma-1}, \qquad (5)$$

where $\mu(\omega;\sigma) = \sigma^{-\sigma}(\sigma-1)^{\sigma-1}\omega^{\sigma-1}$.

5.2 Optimal Allocation of Manager's Time

The plant's owner, who is also its manager, allocates his time between managing the plant and managing demand (sales) so as to maximize profit in (5). This optimal time allocation problem can thus be written as

$$Max_u\ \lambda(u;\gamma)u^{\sigma-1}. \qquad (6)$$

The function $\lambda(u;\gamma)$, is assumed to satisfy the following properties:

(a) $0 \leq \lambda(u;\gamma) \leq 1$,

(b) $\frac{\partial\lambda(u;\gamma)}{\partial u} < 0$,

(c) $\frac{\partial\lambda(u;\gamma)}{\partial\gamma} > 0$, and

(d) $\frac{\partial}{\partial \gamma}\left(\frac{\partial \lambda(u;\gamma)}{\partial u}\right) > 0$.

Properties (a) and (b) clarify the interpretation of λ as a fraction of demand that the firm loses if the manager's time is allocated away from managing sales, 1 – u, to managing the production facility, u. Property (c) says that higher "sales management ability," such as networking relationship with trading houses, quality certification, as well as perhaps the ability to effectively collect debt (see the previous section) leads to smaller loss of demand for any u. Finally, (d) says that the demand loss due to a marginal increment in u decreases with the manager's ability. The optimal resource allocation problem (6) thus captures the fundamental tradeoff faced by the manager: if he devotes more time to managing sales, production is lost, and vice versa. This tradeoff, however, is mitigated by higher ability.

The first order condition for an interior solution is

$$\frac{\partial \lambda(u;\gamma)}{\partial u} u = -(\sigma - 1)\lambda(u;\gamma). \tag{7}$$

Let $u(\gamma)$ denote an optimal solution from (7). A simple comparative exercise yields the following results.

Lemma 1: Assume that the function $\lambda(u;\gamma)$ satisfies properties (a)–(d). Then,

(i) $\frac{du(\gamma)}{d\gamma} > 0$; Time allocated to managing production at the plant, $u(\gamma)$, increases with ability γ.

(ii) $\frac{\partial \pi(u(\gamma);\gamma)}{\partial \gamma} > 0$ and $\frac{\partial \pi(u(\gamma);\gamma)}{\partial \omega} > 0$; Profits increase in ability γ and plant quality ω.

(iii) $\frac{\partial^2 \pi(u(\gamma);\gamma)}{\partial \gamma \partial \omega} > 0$; Ability γ and plant quality ω are complements in the profit function.

Proof: See appendix.

5.3 Mergers and Acquisitions

This subsection describes one way the merger and acquisition process can work in this framework. In doing so, we employ a setting inspired by the structure in Jovanovic and MacDonald (1994), which was developed with the evolution of the U.S. tire industry in mind but also fits some stark patterns in our data.

We assume that an initial "basic" state of technological knowledge arrives which offers 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 technological knowledge (aka "refinement" in the Jovanovic-MacDonald model; this would be knowledge about ring spinning frames and imported raw cotton in our context). 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 sales management ability, normalized to be equal to 1 for all cohorts. Producers from the early cohort, however, have an opportunity to develop this ability above and beyond the initial level. How much any given producer develops his ability before time *T* is a random process, but its outcome is that at time *T*, when the second cohort enters the market, the first cohort's ability is distributed with support $[1, \gamma_{max}]$.

Even though all entrants in the second cohort possess only the initial level of sales management ability, the quality of their plants is higher. We assume that this leads to a new market equilibrium where only plant owners in the first cohort whose ability exceeds a threshold level $\gamma_e \in (1, \gamma_{max})$ can remain in the industry, while those with ability below this threshold have to exit.[39] 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 matched into negotiating pairs also randomly. 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 potential incentive to acquire the plant of the manager with ability γ_L regardless of the plant's quality. Also, the acquisition is more likely to be consummated if, in addition, the plant quality is high than if it is low. To see this more explicitly, note that the highest bid price is given by $\delta\pi(\omega, \gamma_H)$, which is the profit that the manager with ability γ_H can obtain if he takes over this

[39] See the appendix for an explicit derivation of one such industry-wide equilibrium. In our data, 10 out of 21 firms that operated in the industry in 1887 had remained small and exited by shutting their plants.

plant.[40] The lowest asking price, on the other hand, is given by $\pi(\omega, \gamma_L)$. Assuming that the actual price will be somewhere between, an acquisition will be consummated whenever

$$\delta\pi(\omega, \gamma_H) > \pi(\omega, \gamma_L). \quad (8)$$

Other things equal, Lemma 1(iii) implies that the potential gains from an acquisition will be higher when ω is higher, so that plants owned by recent entrants (with quality ω_2) will indeed be more likely to be acquired than plants owned by first-cohort entrants (with quality ω_1).

Condition (8) also implies that for any plant quality, an acquisition is more likely to happen when a 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.[41]

We summarize the empirical predictions implied by this discussion in Proposition 1.

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 be an high-ability early entrant with a relatively aged plant acquiring a more recent entrant with lower ability but a newer plant.

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 can be defined as

$$TFPQ \equiv \frac{y}{x} = \omega u(\gamma). \quad (9)$$

---

[40] The parameter $\delta \in [0,1)$ captures any possible transfer costs (rent dissipation) associated with an acquisition.

[41] In reality there were a few cases where new entrants acted as acquirers. A detailed examination of the data revealed, however, that in all such cases new entrants were actually "spinoffs" from the early entry cohort, that is, they were founded by entrepreneurs with previous experience in incumbent firms. It is well known that such spinoff entrants tend to inherit the ability of their parent firms (see e.g., Klepper and Simons, 2000), so such cases actually render additional support to our theory.

This measure of productivity is predicated on the notion that manager's optimal time allocation is not directly observed by an econometrician.

Lemma 1(i) implies that, for a given ω, TFPQ will increase with the acquiring firm manager's ability γ. Proposition 1 says that a higher-ability manager acquires a plant managed by a lower-ability manager. Thus, we have the following.

Proposition 2: The productivity of an acquired plant rises after an acquisition.

Similarly, Lemma 1(ii) says that profits increase with manager's ability. Combining it with Proposition 1, we have the following.

Proposition 3: The profits of an acquired plant increase after an acquisition.

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

Corollary: Under suitable functional form restrictions on $\lambda(u;\gamma)$, the ratio of effective sales to output in a given plant increases following an acquisition, and its profitability increases by more than its TFPQ.[42]

Proof: See appendix.

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. We first discuss profitability. Write the logarithm of the profit as $\ln[\pi(\omega,\gamma)] = (\sigma-1)\ln\omega + \ln[\lambda(u(\gamma);\gamma)] + (\sigma-1)\ln[u(\gamma)]$, where we drop the constant term containing σ. Combined with the first-order condition (7), the total derivative of this expression reduces to

$$d\ln[\pi(\omega,\gamma)] = (\sigma-1)\frac{1}{\omega}d\omega + \varepsilon_{\lambda\gamma}\frac{1}{\gamma}d\gamma \qquad (10)$$

---

[42] The example above shows that this property holds for some simple and reasonable functional forms and, as shown in the next section, this appears to be the empirically relevant case. We have not yet been able to establish Corollary 1 for a general function form satisfying properties (a)–(c) above, and we suspect that it may not be true in general.

where $\varepsilon_{\lambda\gamma} = \frac{\partial\lambda}{\partial\gamma}\frac{\gamma}{\lambda}$. The first term in (10) captures the effect of plant quality differential between acquired and acquiring plants on profitability, whereas the second term is the effect of managers' ability differential on profitability.

If two incumbents are involved in a merger negotiation, the plant quality is the same, i.e., $d\omega = 0$. In this case, (10) 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., $d\gamma > 0$ and $\varepsilon_{\lambda\gamma} > 0$. 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, the relative pre-acquisition profits depend on whether the plant quality effect dominates the manager's ability effect. From equation (10), a sufficient condition for the pre-acquisition profitability of the acquiring plant to be higher than that of the acquired plant is $\frac{\omega_1}{\omega_2} > \delta^{\frac{1}{\sigma-1}}$, since equation (8) implies that $\varepsilon_{\lambda\gamma}\frac{1}{\gamma}d\gamma \geq -ln\delta$. Furthermore, (10) shows that managerial ability effects are increasing with the ability differential between acquired and acquiring plants, so incumbents who have developed large enough ability to deal with sales will be more profitable than their target plants even though their plant's quality may be considerably lower.

To summarize, the profit of the acquired plant is lower than that of the acquiring plant before acquisition in all acquisitions involving two incumbents. When an incumbent is matched with a new entrant, the profit of the acquired plant is lower than that of the acquiring plant either if $\omega_1/\omega_2 > \delta^{\frac{1}{\sigma-1}}$, or if the acquirer's sales ability is high enough.

Similarly, we can write the total derivative of the logarithm of TFPQ to compare pre-acquisition levels of productivity:

$$dln[TFPQ(\omega,\gamma)] = \frac{1}{\omega}d\omega + \varepsilon_{u\gamma}\frac{1}{\gamma}d\gamma \quad (11)$$

where $\varepsilon_{u\gamma} = \frac{du}{d\gamma}\frac{\gamma}{u}$. Again, the sign of equation (11) is determined by two opposing effects, and the basic relation in (11) is similar to the relation in (10).

To connect pre-acquisition level of profitability with pre-acquisition level of TFPQ, we can use the fact that $\pi = \sigma^{-\sigma}(\sigma-1)^{\sigma-1}(TFPQ)^{\sigma-1}\lambda$ and obtain the following relationship:

$$dln[\pi(\omega,\gamma)] = (\sigma-1)dln[TFPQ(\omega,\gamma)] + dln[\lambda(u(\gamma);\gamma)],$$

which can be written as

$$dln[\pi(\omega,\gamma)] = (\sigma-1)dln[TFPQ(\omega,\gamma)] + \frac{1}{\gamma}\left[\varepsilon_{\lambda u}\varepsilon_{u\gamma} + \varepsilon_{\lambda\gamma}\right]d\gamma, \qquad (12)$$

where $\varepsilon_{\lambda u} = \frac{\partial\lambda}{\partial u}\frac{u}{\lambda}$. Note that $\varepsilon_{\lambda u} < 0$, $\varepsilon_{u\gamma} > 0$, and $\varepsilon_{\lambda\gamma} > 0$. Thus equation (12) indicates that it is possible that $dln[\pi(\omega,\gamma)]$ is strictly positive but $dln[TFPQ(\omega,\gamma)]$ can be zero or negative when ability effects are sufficiently large, i.e., the term $\varepsilon_{\lambda\gamma}$ is sufficiently large. We then have

Proposition 4: Under suitable parametric restrictions, pre-acquisition TFPQ of an acquiring plant can be lower than that of an acquired plant even if pre-acquisition profitability of an acquiring plant is higher than that of an acquired plant.

This property is consistent with the empirical patterns we saw in our data and it is in sharp contrast to the assortative matching theory of mergers and/or Q-theory of merger.

To illustrate that the empirical patters in our empirical analysis can be consistent with the mechanism outlined above, we work through a simple numerical example of the model in the appendix.

We have shown how a managerial time allocation decision, in the presence of heterogeneous quality managers and plants, can yield the empirical patterns document 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 of the Japanese Cotton Spinners' Association), 1889-1920 (in Japanese), Dainihon Boseki Rengokai, various issues.

Jovanovic, Boyan, 1998. "Vintage Capital and Inequality," *Review of Economic Dynamics*, 1 (2), 497-530.

Jovanovic, Boyan and Yuri Yatsenko, 2012. "Investment in Vintage Capital," *Journal of Economic Theory*, 147 (2), 551-569.

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.

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.

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.

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.

Otsuka, Keijiro, Ranis, Gustav, and Gary Saxonhouse. *Comparative technology choice in development: The Indian and Japanese cotton textile industries*. New York: St. Martin’s Press, 1988.

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., and Galvin Wright, 1984. *Technique, Spirit and Form in the Making of the Modern Economies: Essays in the Honor of Willam N. Parker*. London: JAI Press.

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.

Figures, Tables and Appendix

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

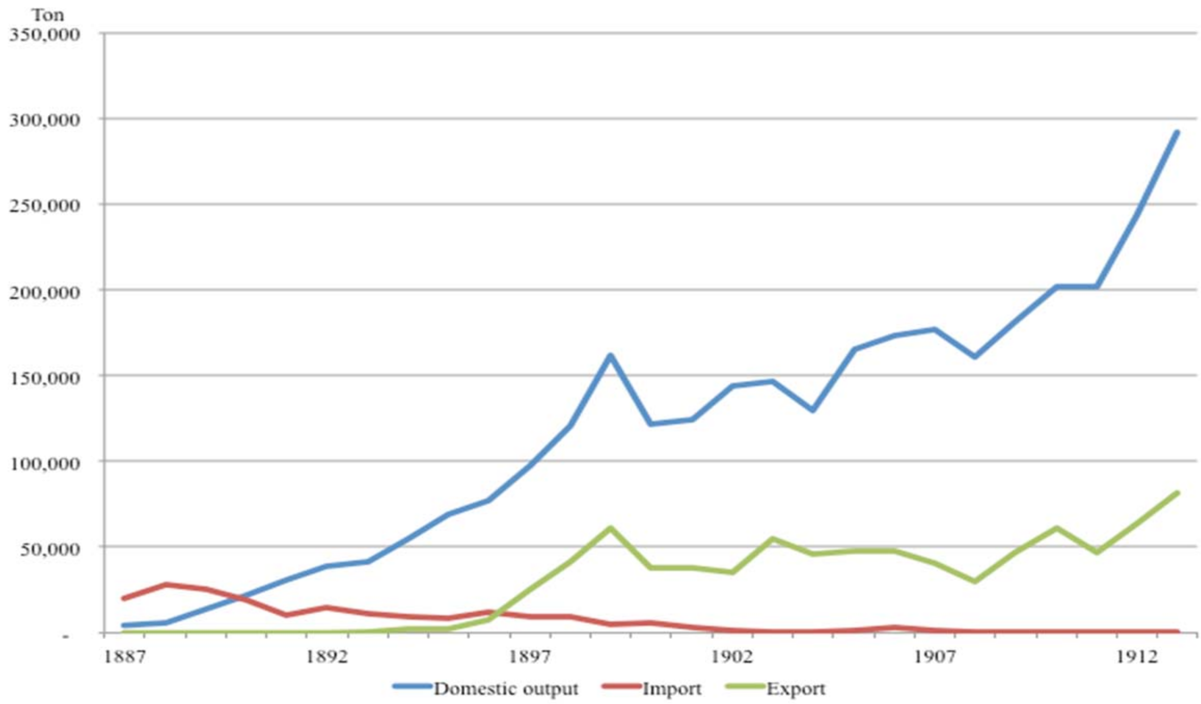

Source: Nihon Choki Tokei Soran, our estimates.

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

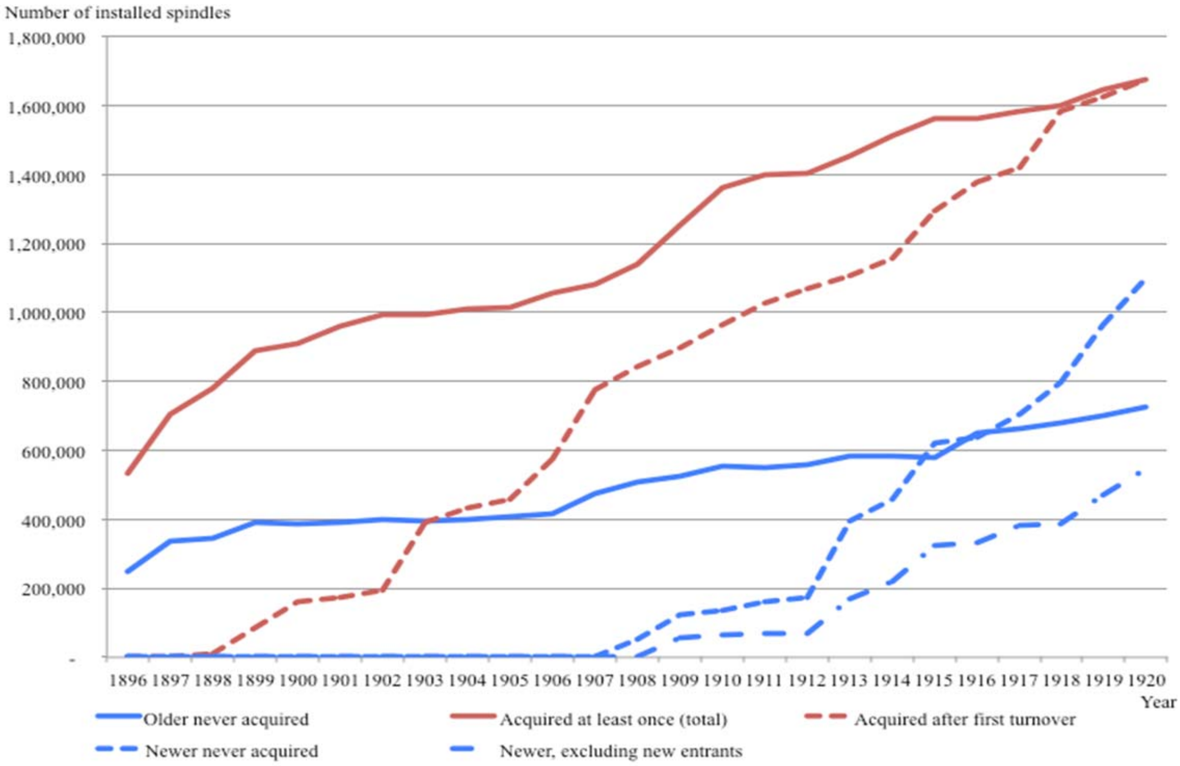

Source: Our estimates using the data described in Section 4 below.[43]

[43] "Older never acquired" represents plants that came into operation in 1902 or earlier and were never targets in an acquisition. "Newer never acquired" represents 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.

Figure 3. Ownership concentration in three largest firms

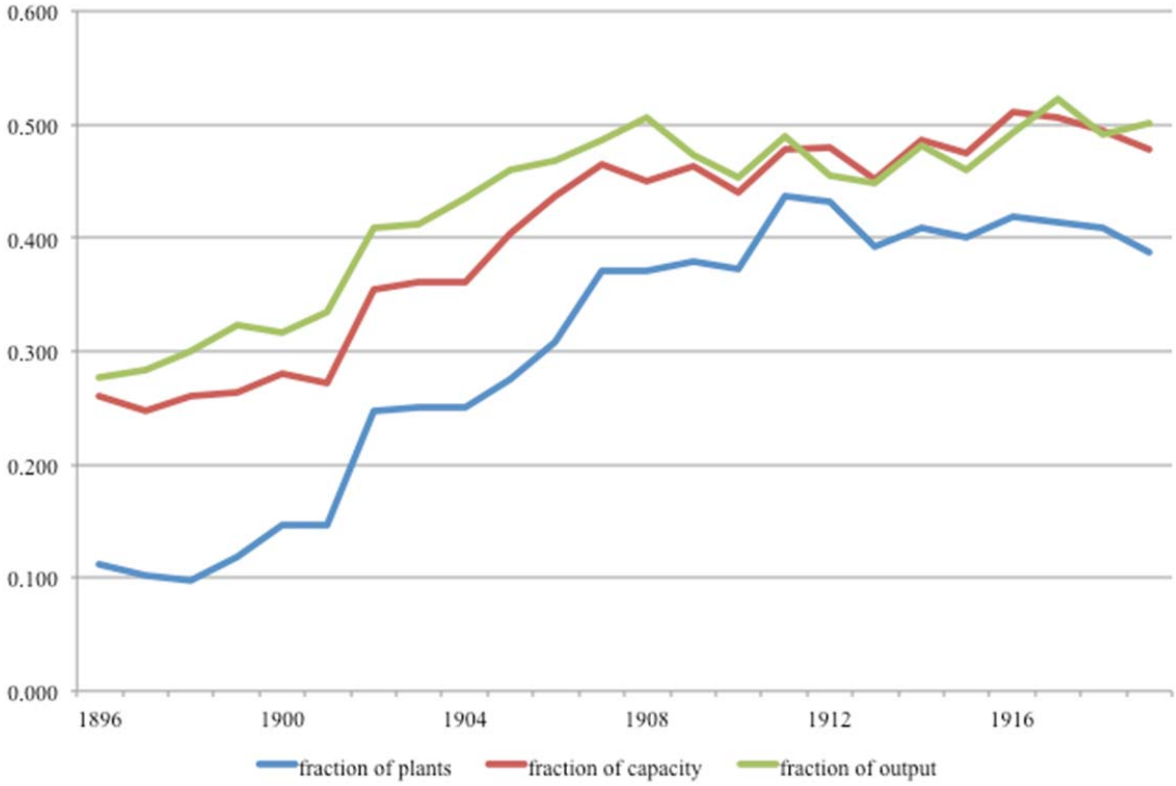

Note: The figure depicts the evolution of the fraction of plants owned by the three largest firms in 1920 (Kanebo, Toyobo, Dainippon Boseki) and these plants' capacity and output as a fraction of the industry total. Toyobo data include that of its predecessor firms (Osaka Boseki and Mie Boseki) prior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

Table 1. 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.

Table 2. 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)
ROE	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. ROE is return on equity, accounting profits divided by shareholders' paid-in capital. There are only 6 ROE 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 3. Within-acquired-plants comparisons of productivity, profitability and prices

	Dependent variable		
	Log output	Log gross operating surplus rate	Log count-adjusted price
Late pre-acquisition dummy	-0.017	-0.008	-0.006
	(0.028)	(0.063)	(0.030)
Early post-acquisition dummy	0.032	0.181**	0.012
	(0.033)	(0.078)	(0.035)
Late post-acquisition dummy	0.111**	0.317***	0.007
	(0.048)	(0.094)	(0.038)
Log spindles-days in operation	0.715***		
	(0.039)		
Log worker-days	0.248***		
	(0.037)		
Log capacity change	-0.090*	-0.205*	0.026
	(0.050)	(0.108)	(0.031)
Constant	-0.903*	2.310***	1.722***
	(0.535)	(0.043)	(0.027)
Plant fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Observations	1,248	962	1,213
R-squared	0.944	0.649	0.844

Note: 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 4. Within-acquired-plants comparisons of productivity, profitability and prices: acquired by serial acquirers

	Dependent variable		
	Log output	Log gross operating surplus rate	Log count-adjusted price
Late pre-acquisition dummy	-0.009	0.043	0.002
	(0.049)	(0.076)	(0.032)
Early post-acquisition dummy	0.110*	0.450***	0.054
	(0.060)	(0.085)	(0.070)
Late post-acquisition dummy	0.204**	0.609***	0.094
	(0.077)	(0.102)	(0.077)
Log spindles-days in operation	0.695***		
	(0.073)		
Log worker-days	0.245***		
	(0.060)		
Log capacity change	-0.140	-0.553***	-0.045
	(0.114)	(0.194)	(0.051)
Constant	-0.605	1.049	1.754***
	(0.999)	(0.871)	(0.042)
Plant fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Observations	555	475	530
R-squared	0.930	0.648	0.866

Note: 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 5. Within-acquisition comparisons of productivity and profitability: acquired and incumbent plants.

	All acquisitions		By serial acquirers	
	TFPQ	Log normalized GOS rate	TFPQ	Log normalized GOS rate
After acquisition	-0.047*	0.036	-0.043	0.022
	(0.025)	(0.038)	(0.026)	(0.030)
Acquired plant	0.007	-0.187***	-0.005	-0.166**
	(0.030)	(0.056)	(0.022)	(0.069)
After acquisition x Acquired plant	0.116***	0.202***	0.157***	0.249***
	(0.032)	(0.063)	(0.034)	(0.069)
Constant	0.054***	0.322***	0.121***	0.039
	(0.018)	(0.037)	(0.024)	(0.034)
Acquisition–year dummies	Yes	Yes	Yes	Yes
Observations	1,522	1,329	1,105	948
R-squared	0.150	0.209	0.330	0.226

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. GOS rate is gross operating surplus rate of the plant as defined in the main text, divided by industry-year average for each year. Robust standard errors clustered at the acquisition case level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 6. Within-acquisition inventory/output value ratio, fixed capital cost/output ratio, days in operation per year and capacity utilization rate comparisons.

	Logged inventory/output	Logged fixed capital cost/output	Days in operation per year	Capacity utilization rate
After acquisition	-0.309***	-0.051	-0.192	-0.002
	(0.083)	(0.063)	(1.080)	(0.008)
Acquired plant	0.660***	0.251***	-11.917***	-0.057***
	(0.124)	(0.069)	(3.950)	(0.013)
After acquisition x Acquired plant	-0.690***	-0.278**	12.264***	0.041**
	(0.126)	(0.081)	(3.936)	(0.016)
Constant	-2.491***	-2.758***	329.488***	0.854***
	(0.084)	(0.059)	(1.124)	(0.006)
Acquisition-year dummies	Yes	Yes	Yes	Yes
Observations	810	1,503	1,574	1,571
R-squared	0.383	0.273	0.095	0.069

Inventory/Output value ratio is the ratio of finished yarn, plus accounts accruable on the company balance, divided by the value of output produced (output, times plant-specific price). Unit fixed cost is the book value of fixed assets, multiplied by the interest rate, divided by count-adjusted output produced. The time period is up to 4 years prior to acquisition and up to 8 years after acquisition. Robust standard errors clustered at the acquisition case level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 7. Producer's connections to trade houses and performance metrics, 1898-1902

A. Using trader network indicator, excluding already acquired plants

	Logged inventory/output	Logged fixed capital cost/output	Days in operation per year	Capacity utilization rate
Trader network dummy	-0.463**	-0.601	4.451	0.027
	(0.181)	(0.577)	(6.715)	(0.024)
Logged firm age	0.350**	-0.113	14.030**	-0.012
	(0.173)	(0.123)	(6.242)	(0.028)
Constant	-2.609***	-1.894***	269.151***	0.735***
	(0.368)	(0.254)	(15.084)	(0.073)
Year dummies	Yes	Yes	Yes	Yes
Observations	222	245	306	306
R-squared	0.181	0.108	0.048	0.025

B. Using percentage of shares owned by connected traders from Osaka and Tokyo areas, excluding already acquired plants

	Logged inventory/output	Logged fixed capital cost/output	Days in operation per year	Capacity utilization rate
Pct of shares held by traders from Osaka and Tokyo	-2.700***	-0.631	6.804	0.157*
	(0.705)	(0.687)	(19.328)	(0.089)
Logged firm age	0.353*	-0.121	14.912**	-0.013
	(0.184)	(0.132)	(6.317)	(0.029)
Constant	-2.744***	-1.890***	269.380***	0.743***
	(0.382)	(0.273)	(15.01)	(0.074)
Year dummies	Yes	Yes	Yes	Yes
Observations	222	229	306	306
R-squared	0.204	0.127	0.046	0.028

Inventory/Output value ratio is the ratio of finished yarn, plus accounts accruable on the company balance, divided by the value of output produced (output, times plant-specific price). Unit fixed cost is the book value of fixed assets, multiplied by the interest rate, divided by count-adjusted output produced. The time period is 1898-1902. Robust standard errors clustered at the firm level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 8. Within-acquired plants comparisons of productivity, not including and including ability control functions

	Dependent variable: log output (adjusted for count).		
Late pre-acquisition dummy	-0.017	0.003	-0.053
	(0.028)	(0.028)	(0.033)
Early post-acquisition dummy	0.032	-0.021	-0.025
	(0.033)	(0.036)	(0.045)
Late post-acquisition dummy	0.111**	0.028	-0.002
	(0.048)	(0.048)	(0.054)
Log spindles-days in operation	0.715***	0.751***	0.702***
	(0.039)	(0.050)	(0.086)
Log worker-days	0.248***	0.115***	0.228***
	(0.037)	(0.032)	(0.053)
Log capacity change	-0.090*	-0.003	-0.173*
	(0.050)	(0.050)	(0.092)
Log fixed cost to output ratio		-0.276***	
		(0.040)	
Log inventory to output ratio			-0.075***
			(0.020)
Constant	-1.075*	-0.500	-0.742
	(0.544)	(0.722)	(1.194)
Plant fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Observations	1,248	1,027	572
R-squared	0.944	0.946	0.943

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.

Appendix

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.[44]

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 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.[45] 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

---

[44] 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.

[45] 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.

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."[46] 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 analysis, we converted them to the standard 20th 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 20th 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.[47] Third,

---

[46] 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.

[47] In the division of labor between sexes in Japanese cotton spinning mills, opening, mixing, carding, repairing and

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).

B. An example of ownership 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

---

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.

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. (2012). 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.[48]

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

---

[48] 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 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.

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).[49] 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) above to all individual plants in our data.

Table A1 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 dramatically 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.

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

Table A1. 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 A2. Table A2 follows the same format as Table 2 in our 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 A2. 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 our 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.

D. Robustness Checks

We investigate the robustness of our results in the within-acquisition-episode specifications in the main text to alternative controls groups. We construct two different matched samples to do so.

We are interested in estimating the following parameters:

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

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

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

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.[50]

D.1 TFPQ measures

As an additional robustness check, we measure TFPQ in the tests above using five alternative methods. The first four of these use residuals from the estimations of production function. Our basic specification of the production function is given by

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

[50] 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.

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. 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 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 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. The fourth measure of TFPQ is a residual from the fixed-model (FE) estimation. Finally, we also construct the fifth 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.

D.2 Same owner matching

As mentioned above, we construct two different matched samples to estimate equation (A4). In the first matched sample, 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). We first 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 a 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 A3 and A4 report estimation results using this matched sample with different weighting schemes as above. For comparison, we also include results from the standard dif-in-dif estimation where we just ignore matching. Table A5 presents the estimation results using different measures of TFPQ as described in Section D.1 and simple weights (results using other types of weights are similar). Table A6 shows estimation results of equation (A4) when outcome variables are the ratio of plant-level inventory to output, the ratio of fixed capital cost to output, days in operation, and capacity utilization rate. As we can see, the main results are robust to alternative weights and alternative measures of TFPQ.

Table A3: Estimation Results from Same owner matching – All acquisitions

	All acquisitions					
	Simple weights		Kernel weights		Standard dif-in-dif estimation	
	TFPQ (OLS)	Log normalized GOS rate	TFPQ (OLS)	Log normalized GOS rate	TFPQ (OLS)	Log normalized GOS rate
After acquisition	-0.047*	0.036	-0.055*	0.103**	-0.056**	0.018
	(0.025)	(0.038)	(0.029)	(0.046)	(0.027)	(0.029)
Acquired plant	0.007	-0.187***	-0.005	-0.119*	-0.010	-0.213***
	(0.030)	(0.056)	(0.036)	(0.063)	(0.028)	(0.052)
After acquisition x Acquired plant	0.116***	0.202***	0.114***	0.107	0.125***	0.214***
	(0.032)	(0.063)	(0.038)	(0.069)	(0.034)	(0.058)
Constant	0.054***	0.322***	0.065***	0.269***	0.065***	0.347***
	(0.018)	(0.037)	(0.021)	(0.043)	(0.018)	(0.029)
Observations	1,522	1,329	966	832	1,522	1,329

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

Table A4: Estimation Results from Same owner matching – Serial acquirers

	Serial acquirers					
	Simple weights		Kernel weights		Standard dif-in-dif estimation	
	TFPQ (OLS)	Log normalized GOS rate	TFPQ (OLS)	Log normalized GOS rate	TFPQ (OLS)	Log normalized GOS rate
After acquisition	-0.043	0.022	-0.070**	0.057	-0.059*	0.022
	(0.026)	(0.030)	(0.032)	(0.042)	(0.032)	(0.031)
Acquired plant	-0.005	-0.166**	-0.040	-0.115	-0.017	-0.165**
	(0.022)	(0.069)	(0.032)	(0.068)	(0.023)	(0.070)
After acquisition x	0.157***	0.249***	0.182***	0.161**	0.171***	0.241***
Acquired plant	(0.034)	(0.069)	(0.037)	(0.071)	(0.038)	(0.072)
Constant	0.121***	0.039	0.161***	0.032	0.124***	0.034
	(0.024)	(0.034)	(0.029)	(0.040)	(0.027)	(0.032)
Observations	1,105	948	708	601	1,105	948

Table A5: Estimation Results from Same owner matching –Several TFPQ measures

	All acquisitions and Simple weights				
	Dependent variable: TFPQ				
	OLS	Wooldridge	Blundell-Bond	FE	Index
After acquisition	-0.047*	0.026	-0.059***	-0.047*	0.029
	(0.025)	(0.028)	(0.018)	(0.024)	(0.027)
Acquired plant	0.007	-0.034	0.023	-0.065**	-0.079***
	(0.030)	(0.028)	(0.022)	(0.025)	(0.027)
After acquisition x Acquired plant	0.116***	0.113***	0.099***	0.107***	0.122***
	(0.032)	(0.032)	(0.023)	(0.031)	(0.031)
Constant	0.054***	2.246***	-0.009	0.021	-2.260***
	(0.018)	(0.019)	(0.018)	(0.015)	(0.018)
Observations	1,522	1,524	1,453	1,522	1,571

Table A6: Estimation Results from Same owner matching – Inventory, Fixed capital, Days in operation, Capacity utilization rates

	All Acquisitions			
	Simple weights			
	Logged inventory/output	Logged fixed capital/output	Days in operation per year	Capacity utilization rate
After acquisition	-0.309***	-0.051	-0.192	-0.002
	(0.083)	(0.063)	(1.080)	(0.008)
Acquired plant	0.660***	0.251***	-11.917***	-0.057***
	(0.124)	(0.069)	(3.950)	(0.013)
After acquisition x Acquired plant	-0.690***	-0.278***	12.264***	0.041**
	(0.126)	(0.081)	(3.936)	(0.016)
Constant	-2.491***	-2.758***	329.488***	0.854***
	(0.084)	(0.059)	(1.124)	(0.006)
Observations	810	1,474	1,574	1,571
	Kernel weights			
	Logged inventory/output	Logged fixed capital/output	Days in operation per year	Capacity utilization rate
After acquisition	-0.317**	-0.066	0.784	0.010
	(0.089)	(0.078)	(1.503)	(0.013)
Acquired plant	0.656***	0.240***	-15.072***	-0.050***
	(0.137)	(0.082)	(4.942)	(0.017)
After acquisition x Acquired plant	-0.670***	-0.244**	14.723***	0.026
	(0.139)	(0.095)	(4.921)	(0.020)
Constant	-2.491***	-2.754***	329.964***	0.848***
	(0.089)	(0.072)	(1.536)	(0.009)
Observations	631	940	1,011	1,008
	Standard dif-in-dif estimation			
	Logged inventory/output	Logged fixed capital/output	Days in operation per year	Capacity utilization rate
After acquisition	-0.369***	0.050	-0.421	-0.008
	(0.074)	(0.060)	(0.903)	(0.008)
Acquired plant	0.642***	0.307***	-10.722***	-0.061***
	(0.117)	(0.068)	(3.970)	(0.013)
After acquisition x Acquired plant	-0.623***	-0.380***	12.622***	0.047***
	(0.126)	(0.085)	(4.042)	(0.015)
Constant	-2.455***	-2.835***	328.916***	0.858***
	(0.073)	(0.056)	(1.161)	(0.006)
Observations	810	1,474	1,574	1,571

D.3 Pre-acquisition characteristics and trend matching

We created the second 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.[51] 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 GOS 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 gross operation surplus (GOS) 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.[52] The simple weight (i.e., $\omega_j = 1$) is used for this estimation.

Tables A7, A8, and A9 present estimation results using this matched sample. Table A10 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 A7: Estimation Results from pre characteristics and trend matching – All acquisitions

	All acquisitions			
Matching criteria	Plant age, size, location		TFPQ growth rate	GOS growth rate
	TFPQ (OLS)	Log normalized GOS rate	TFPQ (OLS)	Log normalized GOS rate
After acquisition	-0.015*	0.043**	-0.012	0.012
	(0.008)	(0.019)	(0.007)	(0.035)
Acquired plant	-0.009	-0.196***	0.003	-0.155**
	(0.021)	(0.062)	(0.022)	(0.067)
After acquisition x Acquired plant	0.079***	0.194***	0.075***	0.227***
	(0.026)	(0.058)	(0.026)	(0.060)
Constant	0.042***	0.077**	0.042***	-0.066
	(0.008)	(0.030)	(0.008)	(0.040)
Observations	12,431	9,959	12,113	7,784

[51] 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.

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

Table A8: Estimation Results from pre characteristics and trend matching – Serial acquirers

	Serial acquirers			
Matching criteria	Plant age, size, location		TFPQ growth rate	GOS growth rate
	TFPQ (OLS)	Log normalized GOS rate	TFPQ (OLS)	Log normalized GOS rate
After acquisition	-0.005	0.090***	-0.014	0.037
	(0.010)	(0.022)	(0.009)	(0.048)
Acquired plant	0.006	-0.038	-0.005	-0.027
	(0.024)	(0.064)	(0.028)	(0.066)
After acquisition x	0.123***	0.187***	0.129***	0.253***
Acquired plant	(0.030)	(0.062)	(0.031)	(0.054)
Constant	0.104***	-0.132***	0.120***	-0.132**
	(0.015)	(0.025)	(0.037)	(0.050)
Observations	6,353	5,151	6,030	4,133

Table A9: 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	FE	Index
After acquisition	-0.015*	0.067***	-0.013**	-0.017***	0.072***
	(0.008)	(0.012)	(0.006)	(0.005)	(0.010)
Acquired plant	-0.009	-0.001	-0.002	-0.057***	0.0004
	(0.021)	(0.020)	(0.017)	(0.016)	(0.020)
After acquisition	0.079***	0.075***	0.061***	0.081***	0.079***
x Acquired plant	(0.026)	(0.025)	(0.019)	(0.023)	(0.023)
Constant	0.042***	2.176***	-0.036***	0.020***	-2.404***
	(0.008)	(0.010)	(0.008)	(0.005)	(0.010)
Observations	12,431	12,435	11,630	12,431	12,521
	Matching criteria: TFPQ growth rate				
	Dependent variable: TFPQ				
	OLS	Wooldridge	Blundell-Bond	FE	Index
After acquisition	-0.025**	0.053***	-0.014**	-0.032***	0.061***
	(0.010)	(0.014)	(0.007)	(0.009)	(0.014)
Acquired plant	-0.014	-0.012	-0.001	-0.064***	-0.012
	(0.023)	(0.023)	(0.018)	(0.017)	(0.022)
After acquisition	0.090***	0.091***	0.065***	0.096***	0.093***
x Acquired plant	(0.028)	(0.027)	(0.021)	(0.024)	(0.026)
Constant	0.032*	2.157***	-0.115***	0.005	-2.404***
	(0.019)	(0.020)	(0.014)	(0.012)	(0.019)
Observations	12,134	12,134	11,516	12,134	12,184

Table A10: Estimation Results from pre characteristics and trend matching – Inventory, Fixed capital, Days in operation, Capacity utilization rates

	All acquisitions			
Matching criteria	Plant age, size, location			
	Logged inventory/output	Logged fixed capital/output	Days in operation per year	Capacity utilization rate
After acquisition	-0.430***	-0.227***	2.095**	0.013***
	(0.032)	(0.034)	(0.823)	(0.005)
Acquired plant	0.170*	0.143*	-13.140***	-0.054***
	(0.093)	(0.084)	(4.079)	(0.013)
After acquisition x Acquired plant	-0.672***	-0.151*	13.974***	0.043***
	(0.114)	(0.083)	(3.986)	(0.014)
Constant	-2.160***	-2.204***	322.487***	0.811***
	(0.058)	(0.039)	(1.101)	(0.005)
Observations	6,639	10,324	12,698	12,694

D.4 Placebo test

We perform a simple placebo test to check that our main findings are not artificial. In doing so, 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[53] 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 500 times, and calculate a sample mean of estimated coefficients from 500 time simulations, and their standard errors

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.

[53] We also conduct the same placebo test by using pre characteristics and trend matched sample. The results remain unchanged qualitatively.

Table A11: Placebo test

	TFPQ			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	-0.009	0.001	-0.011	-0.008
Acquired plant	-0.002	0.001	-0.005	0.001
After acquisition x Acquired plant	0.006	0.001	0.003	0.009
Constant	0.066	0.001	0.065	0.067
	Log normalized GOS rate			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	0.082	0.001	0.080	0.084
Acquired plant	-0.011	0.002	-0.015	-0.007
After acquisition x Acquired plant	0.016	0.002	0.011	0.021
Constant	0.282	0.001	0.280	0.283
	Logged inventory/output			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	-0.824	0.004	-0.831	-0.817
Acquired plant	-0.017	0.008	-0.032	-0.002
After acquisition x Acquired plant	0.020	0.008	0.004	0.035
Constant	-2.104	0.004	-2.111	-2.097
	Logged fixed capital/output			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	-0.121	0.001	-0.124	-0.118
Acquired plant	0.003	0.003	-0.003	0.008
After acquisition x Acquired plant	-0.008	0.003	-0.014	-0.002
Constant	-2.753	0.001	-2.756	-2.751
	Days in operation per year			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	5.166	0.071	5.027	5.306
Acquired plant	0.114	0.155	-0.190	0.418
After acquisition x Acquired plant	0.017	0.156	-0.289	0.323
Constant	323.856	0.074	323.709	324.002
	Capital utilization rate			
	Mean	Std. Err	95% Conf. Interval	
After acquisition	0.016	0.0003	0.015	0.016
Acquired plant	0.0004	0.001	-0.001	0.002
After acquisition x Acquired plant	-0.0002	0.001	-0.002	0.001
Constant	0.826	0.0003	0.826	0.827

F. Proofs of Lemma 1 and Corollary

Proof of Lemma 1(i): It follows from the first order condition that $\frac{\partial \lambda(u(\gamma);\gamma)}{\partial u}u(\gamma) = -(\sigma-1)\lambda(u(\gamma);\gamma)$.

Differentiating this equation with respect to γ and rearranging yields

$$\frac{du(\gamma)}{d\gamma} = \frac{-(\sigma-1)\frac{\partial \lambda}{\partial \gamma} - \frac{\partial^2 \lambda}{\partial \gamma \partial u}u}{\frac{\partial^2 \lambda}{\partial u^2}u + \frac{\partial \lambda}{\partial u} + (\sigma-1)\frac{\partial \lambda}{\partial u}u} > 0.$$

The denominator of the right hand side of the equation above is negative because of the second order condition. The numerator is negative because of the assumptions that $\frac{\partial \lambda}{\partial \gamma} > 0$ and $\frac{\partial^2 \lambda}{\partial \gamma \partial u} \geq 0$.

Proof of Lemma 1(ii): By the envelope theorem, we have

$$\frac{\partial \pi(u;\gamma)}{\partial \gamma} = \mu(\omega;\sigma)\frac{\partial \lambda(u;\gamma)}{\partial \gamma}u^{\sigma-1} > 0.$$

We also have

$$\frac{\partial \pi(u;\gamma)}{\partial \omega} = \sigma^{-\sigma}(\sigma-1)^{\sigma}\omega^{\sigma-2}\lambda(u;\gamma)u^{\sigma-1} > 0.$$

Proof of Lemma 1(iii): The results in Lemma 1(ii) immediately imply

$$\frac{\partial^2 \pi(u;\gamma)}{\partial \gamma^2} = \sigma^{-\sigma}(\sigma-1)^{\sigma}\omega^{\sigma-2}\frac{\partial \lambda(u;\gamma)}{\partial \gamma}u^{\sigma-1} > 0.$$

Proof of the Corollary: The total derivative of the logarithm of the profit function is given by

$$dln[\pi(\omega,\gamma)] = (\sigma-1)\frac{1}{\omega}d\omega + \varepsilon_{\lambda\gamma}\frac{1}{\gamma}d\gamma$$

where $\varepsilon_{\lambda\gamma} = \frac{\partial \lambda}{\partial \gamma}\frac{\gamma}{\lambda}$. Similarly, the total derivative of the logarithm of TFPQ is written as

$$dln[TFPQ(\omega,\gamma)] = \frac{1}{\omega}d\omega + \varepsilon_{u\gamma}\frac{1}{\gamma}d\gamma$$

where $\varepsilon_{u\gamma} = \frac{du}{d\gamma}\frac{\gamma}{u}$. When an acquisition takes place, plant quality is held constant, i.e., $d\omega = 0$. If we impose the condition that $\varepsilon_{\lambda\gamma} > \varepsilon_{u\gamma}$ on the functional form of λ, then we have $dln[\pi(\omega,\gamma)] > dln[TFPQ(\omega,\gamma)]$.

E. A Numerical Example of the Model

Assume that the function λ takes a particularly simple form $\lambda(u;\gamma) = \gamma - u$ and set the value of model's parameters as follows: $\omega_{entrant} = 1.5\omega_{incumbent}$, $\delta = 0.5$, and $\sigma = 2$. Thus, we assume that the quality of an entrant's plant (potential target) is 50 percent higher than the quality of an incumbent's plant (potential acquirer), and that the potential profit loss due to the acquisition is 50 percent. The entrants' ability level is set to 1, $\gamma_{entrant} = 1$, and the surviving incumbent ability, $\gamma_{incumbent}$, is uniformly distributed over the interval [1.2,2].

Under the assumptions above, profits are written as $\pi = \frac{1}{4}\omega(\gamma - u)u$. The optimal choice of u, the maximized profit, and TFPQ are given, respectively, by

$$u^* = \frac{1}{2}\gamma,$$

$$\pi(\omega,\gamma) = \frac{1}{16}\omega\gamma^2,$$

$$TFPQ(\omega,\gamma) = \frac{1}{2}\omega\gamma.$$

These equations show that Lemma 1 and Propositions 1 to 3 hold.

To see whether or not the Corollary holds, note that $ln\pi(\omega,\gamma) = ln\frac{1}{16} + ln\omega + 2ln\gamma$ and $lnTFPQ(\omega,\gamma) = ln\frac{1}{2} + ln\omega + ln\gamma$. Differentiating both equations with respect to γ, we get

$$\frac{\partial ln\pi}{\partial\gamma} = 2\frac{1}{\gamma},$$

$$\frac{\partial lnTFPQ}{\partial\gamma} = \frac{1}{\gamma}.$$

In other words, a percentage change of profits is higher than a percentage change of TFPQ when a plant owner with higher ability acquires a plant managed by a lower ability owner.

To compare pre-acquisition levels of profitability and productivity, write $dln\pi = \frac{1}{\omega}d\omega + 2\frac{1}{\gamma}d\gamma$, and $dlnTFPQ = \frac{1}{\omega}d\omega + \frac{1}{\gamma}d\gamma$. In a discrete variable case, these equations can be written as

$$ln\pi_{acquiring} - ln\pi_{target} = \left(1 - \frac{\omega_{target}}{\omega_{acquiring}}\right) + 2\left(1 - \frac{\gamma_{target}}{\gamma_{acquiring}}\right),$$

$$lnTFPQ_{acquiring} - lnTFPQ_{target} = \left(1 - \frac{\omega_{target}}{\omega_{acquiring}}\right) + \left(1 - \frac{\gamma_{target}}{\gamma_{acquiring}}\right)$$

Recall that the acquisition takes place if

$$\frac{\pi(\omega;\gamma_{acquiring})}{\pi(\omega;\gamma_{target})} = \left(\frac{\gamma_{acquiring}}{\gamma_{target}}\right)^2 \geq \frac{1}{\delta}$$

Under the assumed parameter values, this inequality can be written as

$$0.5 \leq \frac{\gamma_{target}}{\gamma_{acquiring}} \leq \sqrt{\delta} \approx 0.71$$

Since the incumbents are a potential acquirer and the entrants are a potential target, a range of ability distribution of actual acquirers is determined by

$$\frac{1}{\sqrt{\delta}} \approx 1.41 \leq \gamma_{acquiring} \leq 2$$

Under the assumed parameter values, we have

$$ln\pi_{acquiring} - ln\pi_{target} = -0.5 + 2\left(1 - \frac{\gamma_{target}}{\gamma_{acqurring}}\right) \geq 0,$$

$$lnTFPQ_{acquiring} - lnTFPQ_{target} = -0.5 + \left(1 - \frac{\gamma_{target}}{\gamma_{acqurring}}\right) \leq 0$$

Hence, pre-acquisition TFPQ of an acquiring plant can be lower than that of an acquired plant even if pre-acquisition profit of an acquiring plant is higher than that of an acquired plant.