

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

Serguey Braguinsky, Atsushi Ohyama, Tetsuji Okazaki, and Chad Syverson^{*}

Abstract

We explore how changes in ownership affect the productivity and profitability of producers. Using detailed data from the Japanese cotton spinning industry at the turn of the last century, we find that acquired firms' production facilities were *not* on average less physically productive than the plants of the acquiring firms before acquisition. They were much less *profitable*, however, due to higher inventory levels and lower capacity utilization—differences that reflected problems in managing the uncertainties of demand. After acquisitions, less profitable acquired plants saw drops in inventories and gains in capacity utilization that raised both their productivity and profitability levels.

^{*} Braguinsky: Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA, e-mail: sbrag@andrew.cmu.edu; Ohyama: Graduate School of Economics and Business Administration, Hokkaido University, Kita 9, Nishi 7, Kitaku, Sapporo, Hokkaido 060-0809, Japan, e-mail: ohyama@econ.hokudai.ac.jp; Okazaki: Faculty of Economics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033, Japan, e-mail: okazaki@e.u-tokyo.ac.jp; Syverson: University of Chicago Booth School of Business and NBER, 5807 S. Woodlawn, Chicago, IL 60637, USA, e-mail: chad.syverson@chicagobooth.edu. This research was funded in part by the Berkman Foundation at Carnegie Mellon University, Japan Society for the Promotion of Science (grants No. 25780155 and No. 22223003), and Initiative on Global Markets at the University of Chicago Booth School of Business. We thank three anonymous referees, Nick Bloom, Xavier Giroud, David Greenstreet, Jesse Shapiro, Bob Topel, Tatsuo Ushijima, and conference participants at the ASSA Meetings, the Stanford Conference on Japanese Entrepreneurship, NBER Summer Institute, FOM conference, and the NBER Japan Project for comments. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

The influence of changes in corporate control of assets on productivity has been a focus of theoretical and empirical research for some time. 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.¹

Despite the comfortable intuition of this logic, previous research has not been fully conclusive about the effects of ownership and management turnover. One clear cleft in the literature (spanning both theory and empirics as well as multiple fields) is whether ownership changes are indeed a mechanism to raise the productivity of inputs (Lichtenberg and Siegel, 1987, Maksimovic and Phillips, 2001, Jovanovic and Rousseau, 2002, Schoar, 2002, and Nguyen and Ollinger, 2006, are more recent examples of work supporting this view) or instead driven by non-efficiency considerations like managerial hubris, market power, or investor irrationality (examples backing such viewpoints include Roll, 1986, and Shleifer and Vishny, 2003).²

While there could well be multiple motives for and consequences of ownership changes, part of the literature’s ambiguity no doubt also reflects the inherent limitations of the data available to earlier studies. For example, most datasets do not allow researchers to cleanly distinguish between physical (quantity) productivity and revenue productivity, which can lead to mismeasurement and incorrect interpretations (e.g., Foster, Haltiwanger, and Syverson, 2008, Katayama, Lu, and Tybout, 2009, Syverson, 2011; Atalay, 2014, discusses the importance of separating quantities from expenditures when measuring inputs). In particular, mergers or acquisitions that increase market power will tend to lead to higher output prices for the merged firm. In the typical revenue-based productivity measures of the literature, this would be reflected

¹ 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).

² The literature’s size precludes comprehensive citation. Surveys include Jensen and Ruback (1983) and Andrade, Mitchell, and Stafford (2001). See also the collected works in Kaplan (2000).

as a measured productivity gain even absent changes in technical efficiency. These and related measurement issues mean we are still limited in our knowledge of how turnover in asset ownership and management affects producers' efficiency levels.

In this paper, we seek to make progress on this front. A primary advantage of our effort is a data set that allows us to investigate the production and input allocation processes at an unusual level of detail. We observe the operations, financial reports, management, and ownership of the universe of plants in a growing industry over the course of several decades (the Japanese cotton spinning industry at the open of the 20th century). These data, which we describe in the next section, contain records in physical units of inputs employed and output produced at each plant in the years it operated as well as plant-specific output prices and wages and firm-level financial data. A unique feature of these data is that we observe capacity utilization and can thus measure plant productivity conditional on operation (as well as, of course, without conditioning on operation). We also collected information on all major ownership and/or management turnover events. These combined data let us measure directly how such events were reflected in plants' physical productivity levels, profitabilities, prices, and other operational and financial metrics.

Our first set of findings draws 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 on efficiency-enhancing mergers. Using our best measure of productivity described below (with physical output and input quantities, the latter measured as service flows) we find that 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. We also find, however, that acquired firms were much less *profitable* than acquiring firms prior to being acquired. These findings echo an important strand in previous research that emphasized the role played by assortative matching and profit-enhancing (but not necessarily efficiency-enhancing) synergies (e.g., McGuckin and Nguyen, 1995, Rajan, Volpin, and Zingales, 2000, Rhodes-Kropf and Robinson, 2008, David, 2014).

Therefore ownership/management turnover in the industry is best characterized as "higher profitability buys lower profitability." We use the uniquely detailed nature of our data to dig deeper into the sources of pre-acquisition profitability differentials and to open the "black

box” of post-acquisition profitability improvement by disentangling its various components. We find that pre-acquisition profitability gap did not result from large output price differences between the firms. Nor do we see much evidence of increased market power contributing to higher post-acquisition profits. Instead, as we show, the profitability gap reflected systematically lower unit capital costs among acquirers, coming from two sources: lower average unrealized output levels (inventories and sales for which payment had not been received) and systematically higher capacity utilization. When these better acquirers bought less profitable establishments, the acquired plants saw drops in unrealized output, gains in capacity utilization, and increases in both their productivity and profitability. The pre-acquisition equality in physical productivity between the acquired and the acquiring arose because, as we document below, acquired plants had more productive capital of younger vintages. This canceled out their other disadvantages.

We thus show that despite similar initial productivity levels, efficiency gains along several dimensions contributed to profitability growth for acquired establishments. Essentially, more profitable companies took over firms that had better capital but were using it suboptimally. By taking control of this superior capital and improving the manner in which it was employed, the new management 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. We present a simple model that offers one possible mechanism through which this demand management difference might operate.

The ownership and management reallocation process helped drive considerable productivity growth in the industry. Between 1897 and 1914, industry TFP growth averaged an impressive 2.5 percent per year, while about 70 percent 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—what set the leading firms apart was not their market power (we show there was little) but rather the ability to acquire and fully utilize the most productive capital.

While we focus our analysis on a single industry case study to take advantage of the

available data and unique setting, we believe that we offer broader lessons that shed light on the current literature. It is worth noting that economic environment in Japan during our sample was largely that of more or less unfettered capitalism, with much less government intervention than became common later, and with corporations predominantly relying on equity to raise capital (see, e.g., Miwa and Ramseyer, 2000). In particular, most Japanese firms in our sample (and all important acquiring firms) were joint stock companies with diffused ownership, so that the structures of ownership control and the scope of managers to influence outcomes were very much like the structures and scope that exist today. Thus, the mechanisms we discover here could easily operate in other industries, countries, and time periods; they might just be difficult to isolate in standard datasets.

Furthermore, our data span a time of critical economic development and industrialization for Japan, which was undergoing transition to modernity after 250 years of an isolated, traditionalist society in what can be aptly described as a “self-discovery” process of development (see Hausmann and Rodrik, 2003). 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.

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

The development of the Japanese cotton spinning industry in the late 19th and early 20th centuries has long fascinated economists because of its unique nature “as the only significant Asian instance of successful assimilation of modern manufacturing techniques” before World War II (Saxonhouse, 1971; 1974).³ The historical circumstances surrounding this development made the story even more intriguing. Japan unexpectedly opened up to foreign trade in the 1860s after 250 years of autarky. Cotton yarn, in particular, experienced the combination of the largest fall in relative price from autarky to the free trade regime and the highest net imports (Bernhofen

³ To save some space, we present here a “bare-bones” sketch of these facts. More details can be found in Saxonhouse (1974) and Braguinsky and Hounshell (2014), building upon and expanding Saxonhouse’s study. See also Ohyama, Braguinsky, and Murphy (2004) and Braguinsky and Rose (2009).

and Brown, 2004). But starting from the late 1880s, the domestic cotton spinning industry began a remarkable ascendance. Net exports turned positive for the first time in late 1896, and soon after Japan was exporting a sizeable fraction of its output while imports became negligible.

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 unproductive. In 1887 there were 21 one-mill firms in the industry, with the average mill containing only 4,022 spindles and employing 137 workers on the factory floor. By way of comparison, average mill sizes were much larger in the United States (15,691 spindles), India (25,022), and Britain (38,619)—see Rose (2000, p. 192) and Murayama (1961, p. 340).

[Figure 1 about here]

The second stage involved the explosive growth of the 1890s and was ushered in by two major innovations: the switch to longer-stapled raw cotton imported from India and the U.S., and the adoption of a newer type of cotton spinning machinery. These two innovations were actually closely linked. When Japanese producers were confined to short-stapled cotton grown domestically or imported from China, they had to use specially adapted machines with below state-of-the-art rotation speeds and other characteristics. (Thread spun from short-stapled cotton is prone to breakage, and breakage rates rise with the spinning machinery's speed and power levels.) The switch to Indian and U.S. cotton allowed Japanese mills to import state-of-the-art machines for the first time, making it an episode of technological "refinement" extensively studied in the general growth literature (see the discussion below in Section IV and in Appendix D). By 1896, the average plant already had a capacity of 12,767 spindles and employed 719 workers. Over this decade of growth the number of firms and average plant capacity both tripled while average plant employment rose fivefold. Combined with productivity growth, this caused industry output in physical units to increase 17 fold during the same period.

Early industry entrants that had set up their production facilities before the major innovations of the 1890s faced a disadvantage of being stuck with older vintage machines. However, an important advantage some of them had developed by the time the innovations 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). 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 (Takamura, 1971, I: 325).⁴ This does not mean, however, that the playing ground was level across firms. Especially during slow demand, established trading houses often limited their purchases to reputable producers with whom they had long-term relationships (Takamura, 1971, II: 60-62). Selling outside of the network of large 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 show below that these problems were indeed severe, and the most successful early entrants (who later became major acquirers in the mergers and acquisition market) managed these sales-related issues better than other firms early on.

This superior ability to manage sales may not have been crucial during the rapid expansion phase, but we show in Section III 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 most imports, the Japanese cotton spinning industry felt the limits of the market size for the first time. Once the Boxer Rebellion effectively shut down the Chinese market in 1900, the industry's first major "overproduction crisis" 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 acquisitions of existing production facilities, the first of which occurring in 1898 (Figure A1 in Appendix C). Acquisitions were preferred over purchases of new machinery in part because the average delivery lag for imported machine orders was 21.7 months during our sample, with a lot of variance from year to year (Saxonhouse, 1971, p. 51).

These factors led to the consummation of 73 distinct acquisition deals involving 95 plants (some changed hands more than once) between 1898 and 1920. All in all, 49 of the 78 plants—63 percent of plants and 68 percent of capacity—that were in operation in the industry in 1897, the year before the first acquisition took place, were subsequently acquired at least once.

Several large firms emerged from this process, mostly through serial acquisitions. These were Kanegafuchi Boseki, Mie Boseki, Osaka Boseki (the latter two completed an equal merger in 1914 to form Toyo Boseki), Settsu Boseki, and Amagasaki Boseki (the latter two merged in

⁴ Cotton yarn was also traded on the Osaka exchange, with gross transaction volumes being several times larger than output. Exchange prices strongly influenced what trading houses were willing to pay even in seemingly isolated local markets (Takamura, 1971, I: 327). Cotton spinning firms did take collective action to support prices by enacting output restrictions in slow years. By their nature, however, these restrictions affected all firms uniformly.

1918 to form Dainippon Boseki). These five firms, which shrank to four after the 1914 merger and to three after the 1918 merger, went from owning 10 percent of the plants and 25 percent of industry capacity and output to 40 percent of plants and half of capacity and output over the 25-year period of our analysis (Figure A2 in Appendix C). This concentration of ownership could in principle be due to multiple factors, but as our empirical analysis below will show, it appears to be sourced mostly in their superior ability to manage sales and as a consequence improve the productivity and profitability of the plants they acquired.

II. Data

Our main data source is plant-level data gathered annually by various Japanese prefecture governments and available in historical statistical yearbooks.⁵ For this paper, we have collected and processed all the available data between 1899 and 1920. Because the first acquisition of an operating plant in the industry happened in 1898, we added similar data for 1896-1898 using annualized monthly data published in the *Geppo* bulletin of the All-Japan Cotton Spinners' Association. Our data thus cover 1896 to 1920. Saxonhouse (1971, p. 41) declares that “the accuracy of these published numbers is unquestioned.”⁶

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

We match these plant-level data with financial data from semi-annual reports issued by the firms that owned the plants. Those reports, which we were the first to systematically digitize,

⁵ We describe only the most important features of our data here. A more detailed description is in Appendix A.

⁶ We checked anyway. We found occasional, unsystematic coding errors as well as obvious typos that we could often correct by comparing them with annualized monthly data from *Geppo*. In the vast majority of cases, however, the annual data in statistical yearbooks and the annualized monthly data did correspond very closely (any discrepancies were only a few percentage points). We dropped about 5 percent of observations where the annual data contained in government statistical reports could not be corrected.

contain detailed balance sheets and profit-and-loss statements as well as lists of all shareholders (with the number of shares they held) and executive board members. Select financial data from company reports were also published in the semi-annual publication *Reference on Cotton Spinning (Menshi Boseki Jijo Sankosho)* which started in 1903. We use these data to supplement company reports where they were missing.

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.”⁷ To make different counts comparable for the purpose of productivity analysis, we converted various counts to the standard 20 count using a procedure detailed in Appendix A. Second, we used plant-year-specific female-to-male wage ratios to convert units of female labor to units of male labor.⁸ Third, in addition to the number of installed spindles and total employment, we also have data on the actual number of days of the year the plant was operating. In other words, the data offer us the unusual ability to directly measure the flow of capital and labor services at the plant level rather than to infer them from capital and employment stocks or through other proxies like energy use. This also allows us to measure input utilization rates.

III. Empirical Analysis

On average, 4.3 percent of the industry’s mills were acquired per year during our sample, with the aforementioned serial acquirers responsible for about 40 percent of all acquisitions.⁹ These acquisition episodes form the base of our estimation sample.

A. Differences between Acquirers and Targets before Acquisition

⁷ The yarn count expresses how many yards are contained in a pound of yarn, so it reflects the yarn thickness. Higher-count yarn is thinner (finer) and sells at a higher price per pound than lower-count yarn.

⁸ 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.

⁹ Table A2 in Appendix C presents year-by-year counts of acquired plants during our sample. 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) and the 2.7 percent rate in the LED plant sample from 1972-1981 used by Lichtenberg and Siegel (1987).

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.¹⁰ We compare these firms' plants along several dimensions: physical (quantity-based) productivity, accounting profitability, average output price, main count of yarn produced, the number of days of the year the plant is operational, the average age of the plant's spindles, and the firm's age.¹¹

We compute plants' physical total factor productivity levels (henceforth TFPQ, for quantity-based TFP) using capital and labor input flows, effectively measuring the plant's productivity conditional on it operating. Being able to measure input service flows separately from stocks is a luxury typically unavailable in producer microdata (especially for capital inputs), and as will become clear below, the distinction between this TFP measure and a more typical one that uses input stocks instead is informative about the nature of our results. We compute TFPQ by estimating a production function using the method proposed by De Loecker (2013), with the residuals reflecting plants' TFPQ levels.¹² To measure profitability, we use firms' reported net earnings, divided by the amount of paid-in shareholders' capital.¹³ Equipment age is calculated as the current year minus the equipment vintage year, where vintage year reflects the composition of the years the plant's machines were purchased. Firm age, on the other hand, is always equal to the calendar year minus the year the firm was founded (defined as the year the firm came into existence, which mostly coincides with the year it was incorporated).¹⁴

Table 1 shows means and standard deviations of the aforementioned plant characteristics

¹⁰ There were also a few surviving firms that did not participate in the acquisition market during our sample.

¹¹ In all but a couple of cases (and in all cases that are part of our analysis below), acquisitions involved all plants of acquired firms.

¹² The adaptation of this method to our setting is described in detail in the following section. We also show in Appendix F that our results are robust to alternative production function estimation methods.

¹³ We do not have firm balance sheets data for 1896-97, but we do have these for subsequent years, so we will also measure profitability as return on total capital employed. See Sections III.B and III.C.

¹⁴ As the plant's capital stock includes also buildings and various elements of infrastructure, equipment (spindles) age adjusted for vintage this way makes the plants look younger than they actually are. Firm age, on the other hand, certainly makes those plants that had added new spindles (or scrapped old ones, which is also captured in our measurement) look older than they are. Equipment age thus provides the lower bound, and the firm age the upper bound, for the true overall plant age.

for each group of firms. We separate plants of future 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. The table includes only data from 1896-97, before any acquisitions took place in the industry, and it excludes observations on a few second-cohort plants whose first, partial year of operation was in 1896 or 1897.

[Table 1 around here]

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

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

The numbers in the table’s third through fifth variable rows indicate these profitability gaps are not tied to differences in the prices the plants fetch for their output. As seen in the third row, all firms earn more or less similar prices per unit weight of output. Furthermore, future acquirers produce higher (finer) counts of yarn. When we adjust for this fact by regressing the logged unit-weight prices on indicators for the plant’s main count produced (counts were aggregated into deciles and year dummies were included), we see from the fourth row that acquirers’ count-adjusted prices (the residual from this regression) are even somewhat lower than those of other firms. None of the groups’ average price residuals are significantly different from zero, however. Thus profitability is not about plants earning supernormal prices relative to other similar producers. This result, which we will see in other guises below, supports what we know about the industry’s output market institutions: pricing did not reflect large market power differences across industry producers and is unlikely to contribute to firm- or plant-level outcomes examined in this paper.

The days-in-operation and age comparisons at the bottom of the table offer insight into the possible sources of the productivity and profitability patterns. We saw that 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.¹⁵ 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 operated almost a full working month less than plants in future acquiring firms did. They were efficient while operating, but they were operating considerably less often. 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.

B. Empirical Specifications

The analysis in the previous subsection revealed some systematic pre-acquisition differences between acquiring and target firms. In particular, we saw that although acquiring firms were more profitable, their plants were not necessarily physically more productive, conditional on operating. Now we begin investigating whether and how acquired plants' performance metrics change when they are taken over by acquiring firms.

To measure plants' productivity, we first estimate a production function. As shown in Appendix F (Table A6), even a naïve calculation of TFPQ using residuals from an OLS production function regression shows a substantial post-acquisition TFPQ increase (Table A6 in Appendix F). Capacity utilization also rises (Appendix G). The fact that input use appears to systematically adjust when ownership changes means that standard approaches to measuring productivity effects of acquisitions, which assume productivity evolves exogenously, could bias the estimates by attributing too much of any output gains to input use rather than changes in productivity.¹⁶ Hence as already mentioned we employ the productivity estimation method proposed in De Loecker (2013). This approach accommodates endogenous productivity processes and corrects for any simultaneous shifts in input use and productivity around

¹⁵ In Appendix D, we use additional data on firms' orders of specific pieces of capital equipment to measure how the machines' technical specifications evolved over time. We find clear evidence of pre- and post-early 1890s differences (not sensitive to the choice of a specific cutoff year around this general timeframe) along multiple dimensions: spindle rotation speed, spindles per frame, ability to handle multiple yarn counts and cotton types, etc.

¹⁶ We thank an anonymous referee for pointing this out.

acquisitions, analogous to plants entering into exporting status in De Loecker’s investigation of “learning by exporting”. Comparisons of the estimates below and those obtained using alternative methods in Appendix F suggest that such a phenomenon may indeed be operating in our setting in the period soon after acquisition, although estimated long-term acquisition effects are similar across all methods.

Following De Loecker (2013), we assume that the production function for plant i at time t is given by

$$(1) \quad y_{it} = \beta_k k_{it} + \beta_l l_{it} + \beta_i i_{it} + \beta_a a_{it} + \omega_{it} + \varepsilon_{it},$$

where y is logged output, k and l are respectively logged capital and labor *flows* (i.e., spindle-days and worker-days), i is the change in logged plant capacity—total number of installed spindles—from the previous to the current year (a control for any adjustment costs reflected in production), and a is the logged age of plant capital. The term ω_{it} captures productivity and subsumes the constant, and ε_{it} is a standard i.i.d. error. Productivity evolution is governed by

$$(2) \quad \omega_{it+1} = g(\omega_{it}, \mathbf{acq}_{it}) + \xi_{it+1},$$

where acq_{it} is a vector relating to a plant’s acquisition experience, and ξ_{it+1} is an exogenous productivity shock. In the baseline specification we assume that

$$(3) \quad g(\omega_{it}, \mathbf{acq}_{it}) = \sum_{j=1}^3 \gamma_j \omega_{it}^j + \theta_1 lb_acq_{it} + \theta_2 ea_acq_{it} + \theta_3 la_acq_{it},$$

where we employ three sets of time dummies defined around each acquisition event: a “late pre-acquisition” dummy (lb_acq_{it}) that equals 1 in the two years immediately preceding the acquisition and zero otherwise, an “early post-acquisition” dummy (ea_acq_{it}) equal to 1 for the first three years after the acquisition and zero otherwise, and a “late post-acquisition” dummy (la_acq_{it}) that equals 1 for all subsequent post-acquisition years after the first three and zero otherwise.¹⁷ The predicted output in the first stage of De Loecker’s method is obtained by a polynomial approximation using all inputs in (1) along with the proxy variables including three acquisition timing dummies as above and (logged) cotton consumed in the production process. Capital and labor input coefficients are identified from the following moment conditions:¹⁸

¹⁷ We also estimated the production function with the cubic in specification (3) replaced by a linear approximation as well as by a cubic interacted with acquisition dummies (as in De Loecker, 2013, equation (10)). The results were very similar in all cases; see Appendix F.

¹⁸ This corresponds to equation (26) in Akerberg, Caves, and Fraser (2006) and the timing assumptions discussed therein. Since both capital and labor inputs are measured as service flows in our baseline specification, it is natural

$$(4) \quad E \left[\xi_{it} \middle| \begin{matrix} k_{it} \\ l_{it} \end{matrix} \right] = 0.$$

The coefficients on labor and capital services flow inputs obtained from this specification are estimated to be 0.323 and 0.738, respectively.¹⁹

In our main specification below we use the residual from the production function estimated using this approach as our measure of plant-level TFPQ.²⁰ We use these productivity estimates along with other plant performance measures to investigate how acquisitions are related to changes in plant operations and performance. Because of not enough number of post-acquisition observations on plants that were acquired very late in the sample, acquisitions that happened in 1918 or later are excluded from estimations below. We first look at changes *within* acquired plants. The estimating equations have the general form:

$$(5) \quad y_{it} = \alpha + \theta_1 lb_acq_{it} + \theta_2 ea_acq_{it} + \theta_3 la_acq_{it} + m_A + \mu_t + \varepsilon_{it},$$

to assume that these inputs are chosen simultaneously at the start of production. The quantity of cotton consumed in production, on the other hand, is inseparable from actual output produced, so it reflects all subsequent unobserved productivity shocks like stoppages due to breaking yarn, adjustments made to spindles rotation speeds, and so on.

¹⁹ As a check on the plausibility of these production function estimates, we compared this estimated labor elasticity to labor's share of value added as computed from firms' financial accounts. Assuming input adjustment costs aren't too large, cost minimization implies these two values should be of similar magnitude. They were. While there is some ambiguity as to which line items in our cost data should be excluded from value added, the most inclusive assumptions imply an average wage share in our plants of 0.232, while the most exclusive imply a share of 0.485. Our estimated labor coefficient falls roughly halfway between these bounds. In addition to this check, we estimated the production function using several other approaches and found similar input elasticities. See Appendix F.

²⁰ De Loecker's method allows for estimation of acquisition-driven plant-specific productivity changes directly from equation (2). However, there are two reasons why we use for our benchmark analysis the two-step approach of first computing residual TFPQ measures with De Loecker's method and then using regression specifications to compute the average effect of acquisition (though this limits our analysis to average productivity effects across plants). One, we want to look at productivity changes from before to after acquisition events not just for acquired plants only, but also in comparison to a control group (such as incumbent plants of acquiring firms—see below) in a framework similar to difference-in-difference estimation. Such comparisons require the two-step approach. Two, when conducting TFP measure decompositions below (see Section III.D) we want to use a consistent set of production function parameters to be able to meaningfully compare our estimates of TFPQ conditional on operation with more conventional measures of TFPR and TFPQ not conditioning on capacity utilization. We did, however, confirm that estimates of plant-specific TFPQ (conditional on operation) changes using equation (2) directly (including all cubic interaction terms). The average estimated effects are qualitatively and quantitatively very similar to the benchmark results. These estimates are presented in Appendix F.

where y_{it} is a performance measure of plant i in year t . The key right-hand-side variables are the indicators for the three time periods discussed above: late pre-acquisition, early post-acquisition, and late post-acquisition (the excluded early pre-acquisition period is also the same). We exclude the acquisition year itself from the regression because acquisitions often happen mid-year, making it hard to attribute outcomes solely to the acquirer or the acquired.²¹ The coefficients on these period indicators will reflect how acquired plants' performance measures change around acquisitions. Because we are interested in looking at changes within plants, we include acquisition fixed effects m_A in the specification. These are identical to plant fixed effects for plants that were acquired only once, the majority of our sample, but they allow us to control for possible differences across acquisition events for plants acquired multiple times. We also include year fixed effects μ_t to capture any industry-wide performance shifts over the sample.

In a second specification, we look at productivity changes from before to after acquisition events in a slightly different way. Namely, we compare acquired plants to the incumbent plants of acquiring firms. This in effect uses the incumbent plants as a control group. We lose some data as a result of this (namely, the cases where the acquirer came from outside the industry and hence had no incumbent plants), limiting the exercise to 49 acquired plants. The benefit is that this within-acquisition approach lets us explicitly compare plants' productivity and profitability changes while controlling for any specific circumstances of an acquisition.²²

The estimating equations in this case have the following form:

$$(6) \quad \bar{y}_{it} = \alpha_0 + \beta_1 AA_{it} + \beta_2 Acquired_{it} + \beta_3 Acquired_i \times AA_{it} + m_{it} + \mu_t + \varepsilon_{it} ,$$

where \bar{y}_{it} is the outcome variable of plant i at time t if it is an acquired plant, while the outcome variables of incumbent plants are collapsed to $\bar{y}_{it} = \frac{1}{\#m_A} \sum_{j \in m_A} \omega_j y_{jt}$, where m_A denotes the particular acquisition case in which plant i was acquired and $\#m_A$ is the number of incumbent

²¹ We included all observations when estimating the production function using De Loecker (2013) method because it employs lagged values of various variables, making a time gap undesirable. In those estimations, the acquisition year is treated as part of the late pre-acquisition period. All our estimation results are robust to including the acquisition year into the (early) post-acquisition period or to dropping acquisition-year observations altogether.

²² To avoid problems stemming from the fact that plants previously acquired by serial acquirers are already "incumbent" plants when another acquisition happens, we only label a previously acquired plant as an incumbent after being under the new ownership for five years. The results presented below are not sensitive to other reasonable cutoffs or to using only serial acquirers' originally owned plants in the "incumbent" category.

plants in acquisition m_A . Thus, \bar{y}_{it} for incumbent plants is the weighted average of outcomes of the plants within the given acquisition. The variable AA_{it} is a dummy equal to 1 if acquisition m_A happened prior to year t and zero otherwise, and $Acquired_i$ equals 1 if plant i is purchased in acquisition case m_A and zero otherwise. The acquisition-year fixed effect is m_{it} , and μ_t is the calendar year fixed effect. In the main text, we assign weights $\omega_j = 1$ to all incumbent plants in a given acquisition m_A , which allows us to interpret coefficients β_1 , β_2 , and β_3 similarly to that in standard difference-in-difference estimations. In particular, $\hat{\beta}_3$ reflects the change in acquired plants' performance around their acquisitions relative to the performance changes experienced by the existing plants of their acquirers. We limit the sample time period to 4 years before and 8 years after the acquisition event, but reasonable alternative cutoffs produce similar results.²³

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

C. Changes in Productivity and Profitability

Table 2 shows the results from estimating the within-acquired-plant specification (5) for

²³ We also estimated equation (6) employing kernel weights obtained from the Mahalanobis distance measure where acquired and incumbent plants are matched on plant size, age and location, and also using a standard difference-in-difference procedure ignoring acquisition-based matching altogether. The results of these estimations were very similar to those presented in Table 3 (see Tables A7 and A8 in Appendix F).

three outcome variables: TFPQ, plant profitability, and the count-adjusted price residuals described in Section III.A. It does so for the entire sample of acquisitions (the first three numerical columns of the table) as well as the subsample of acquisitions done by the “serial acquirers” discussed previously (the last three columns).

[Table 2 about here]

The results for TFPQ in the first numerical column indicate that in the first three years after acquisition, acquired plants’ quantity-based TFP levels (conditional on operating) rose about 4.5 percent above their pre-acquisition levels, a marginally statistically significant difference. Subsequent years saw much more productivity growth, with the average TFPQ of acquired plants in the late post-acquisition period (i.e., more than three years after acquisition) being more than 13 percent ($e^{0.126} = 1.134$) higher than their pre-acquisition baseline and significantly higher than in the early post-acquisition period. Thus acquired plants’ TFPQ levels improve considerably following acquisition, though it takes time for this to manifest itself fully.

The next column looks at acquired plants’ profitability around acquisition episodes. We cannot directly evaluate plant-level profitability levels analogously to the cross sectional comparisons in Table 1, for the obvious reason that there are no separate post-acquisition firm profit accounts. We work around this issue by constructing a measure of plant-level net operating surplus equal to the difference between the net value of cotton yarn produced by the plant and plant labor and capital costs (see Appendix E for details). We then divide this by the sum of shareholders’ capital (equity and retained earnings) and interest-bearing debt, which in case of multiple plant firms is assigned to each plant in proportion to the plant’s installed spindle capacity. We call the resulting measure “plant-level return on capital employed”—“plant ROCE” for short—and we use this measure, Winsorized at the top two percent, to compare plant-level profitability before and after acquisition periods.²⁴

The results in the table indicate that the ROCE of acquired plants increases by an average of about six percentage points in the first three years after acquisition. ROCE rises further in subsequent years to a long-run gain of almost nine percentage points. Thus as a share of total

²⁴ As shown in Appendix E, our constructed plant ROCE is highly correlated with firm-level ROCE data in years preceding acquisition events, when we have independent accounting data on both acquired and acquiring firms. The raw correlation between the two measures is about 0.7, and with the exception of extreme tails, the overall distribution fit is quite good too (Figure A4 in Appendix E).

long-run gains, profitability growth occurs faster than the relatively back-loaded growth in productivity. These are big changes in profit rates; the mean pre-acquisition ROCE of acquired plants is about seven percent.

Finally, to see if changes in plant-specific prices contributed to profitability changes, we estimate (5) using as the dependent variable the residuals from the regression of (logged) plant-specific price on the deciles of yarn counts produced and year dummies. As already mentioned, this reflects by how much the price of a given plant was above or below the average plant making yarn of that count in a given year. The results, in Table 2's third column, indicate that post-acquisition prices are statistically indistinguishable from and economically similar to pre-acquisition prices. Prices again do not explain profitability differences.

We also test whether these productivity and profitability changes within acquired plants are systematically related to the attributes of the acquiring firm. While acquiring firms could be demarcated a number of ways, a natural one is whether they were one of the "serial acquirers" we discussed in Section I. We therefore run specification (5) while limiting the sample to acquisitions by one of the five serial acquirer firms. The results are in the three rightmost columns of Table 2. The patterns are qualitatively similar while being slightly more pronounced in magnitude. Acquisitions by serial acquirers correspond to long run improvements in acquired plants' physical TFPQ of about 17 percent ($e^{0.159} = 1.172$) and ROCE increases of 14 percentage points. The point estimates for price changes are larger than in the entire sample, but t -tests fail to reject at conventional confidence levels equality of the coefficient on the pre-acquisition indicator with either of the post-acquisition coefficients.

Overall, the within-plant results in Table 2 indicate that acquired plants see growth in both their TFPQ and profitability levels after acquisition, though a greater share of long-run growth occurs early on for profitability. These productivity and profitability changes are larger for plants that are acquired by the most prolific of acquiring firms.

Table 3 presents similar comparisons using the within-acquisition difference-in-difference framework of equation (6). Now the key variable of interest is β_3 , the coefficient on the interaction of the indicators for an acquired plant and for the post-acquisition period. This coefficient shows how productivity, profitability, and prices change for acquired plants *relative to* their average levels among the incumbent plants of the firm that acquires them. We again estimate the specification for all acquisitions as well as the subsample done by serial acquirers.

[Table 3 about here]

In both TFPQ specifications, the estimates of the interaction coefficient β_3 are positive and statistically significant at the 1 percent level. The post-acquisition improvement of TFPQ of acquired plants (this time relative to incumbent plants of the acquirer) averages about nine percent for all acquisitions and about 12 percent for acquisitions by serial acquirers. In addition, the acquired plant dummy coefficients are small in both samples, suggesting once again that there is little systematic difference between the physical TFP of acquired and incumbent plants prior to acquisitions (this is also observed in year-by-year estimations presented in Appendix K).

In the profitability regressions, $\hat{\beta}_3$ is also positive and statistically significant. Profit rates of acquired plants rise by four percentage points relative to acquiring firms' plants in the whole sample and by about six percentage points in acquisitions by serial acquirers. Here, the acquired plant main effect is both statistically and economically negative, reflecting acquired firms' profitability deficits before acquisition.

Once again, there are no differences to speak of in prices charged by acquired and incumbent plants, both before and after acquisition events, although point estimates suggest a small post-acquisition increase.

These results further reinforce what we saw in Table 2: acquisition was accompanied by growth in the acquired plants' productivity and profitability levels. We see here that this is true relative not only to the acquired plants' own levels before the acquisition, but also relative to changes within incumbent plants owned by their acquiring firms.

D. Decomposing Profitability Differentials

When considered together, the findings above present a sort of puzzle. If it is neither prices nor productivity, what makes incumbent plants more profitable than acquired plants before acquisition? How do acquisitions by more profitable firms improve TFPQ in acquired plants?

Accounting Decompositions. We begin digging into this puzzle by decomposing plants' profitability differences using our detailed financial data. Specifically, we decompose the pre-acquisition profitability differential between acquiring and acquired firms as well as the pre- to post-acquisition profitability changes for acquired plants into their various components. This lets us isolate the most important factors driving profitability differences.

We first express a plant's ROCE as the net value of cotton yarn produced and the plant's

labor and capital costs (all per unit of capital assets):

$$(7) \quad \frac{\pi_i}{C_i} = \frac{(1-\nu)Y_i}{C_i} - \frac{w_iL_i}{C_i} - \frac{R_i}{C_i}.$$

Here, π_i is plant i 's operating income (here, we drop time subscripts for the sake of parsimony). Y_i denotes the value of its output, and ν is the fraction of intermediate input and non-labor operational costs in the value of output (e.g., the costs of raw cotton, energy, etc.). Plant wage costs are w_iL_i , R_i is capital cost, and C_i is plant i 's share of its owning firm's capital employed (the sum of shareholders' capital and interest-bearing debt), where the share equals the plant's share of its firm's installed capacity (number of spindles). The details of variable construction are described in Appendix E. In a nutshell, we use plant price and output data to obtain Y and plant-level data on worker-days and average daily wages to obtain wL . Capital cost is the sum of depreciation of fixed capital and interest payments on borrowed capital, with both depreciation and interest rates assumed to be the same for all plants, as is the parameter ν (these values are estimated from the available firm-level and industry-wide data). All nominal values including capital employed are divided by the consumer price index to account for inflation. Note that we did not have to do this in our regression analysis because our specifications include year fixed effects.

We present the results of decomposition (7) in Table 4. The three panels each correspond to the decomposition of a particular profitability differential. The top panel compares plants of acquired firms ("acquired plants") and those of their future acquirers ("incumbent plants") for up to 4 years prior to acquisition events. The bottom two panels compare acquired plants before and after acquisitions, with the post-acquisition years split as in the regressions above: the middle panel looks at the first 3 years immediately following the acquisition, and the bottom panel looks at the subsequent post-acquisition years up to the 10th year.

[Table 4 about here]

The top panel of Table 4 shows that incumbent plants' 5.1 percentage point ROCE advantage over acquired plants is mostly explained by a net output value to total assets ratio (the first term on the right hand side of (7)) that is on average 6.5 percentage points higher.²⁵ Wage

²⁵ The ROCE differential in the top panel of Table 4 is somewhat larger in magnitude than the acquired plant dummy coefficient in Table 3, where it was -0.03. The ROCE differentials in the middle and bottom panels of Table 4, however, correspond very closely to regression coefficients in Table 2, where they were 0.06 and 0.09, respectively.

costs per unit of assets are actually higher in incumbent than in acquired plants, reducing the ROCE difference. Capital costs are similar in size though statistically smaller for incumbents.

The bottom two panels of Table 4 show the decomposition of acquired plants' ROCE changes around acquisition episodes. ROCE improves by 6.3 percentage points, and grows 10 percentage points in the longer run. As with the cross-sectional differences, most of the changes came from growth in acquired plants' ratios of net output value to total assets.

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

$$(8) \quad \log\left(\frac{\psi Y_i}{C_i}\right) = \log(\psi p_i) + \log\left[\frac{\exp(\hat{Y}_i)}{C_i}\right] + TFPQ_i,$$

where $\psi \equiv 1 - \nu$ is the unit price margin (common to all producers), p_i is the plant's output price, \hat{Y}_i is the predicted output from the production function, and $TFPQ_i$ is the production function residual.²⁶ This expression lets us measure the contribution of these three components to the net value of output per unit of shareholders' capital. These decompositions are presented in Table 5.

[Table 5 about here]

As in the regression analyses, price and TFPQ differentials contribute relatively little to the stark profitability differences between acquired and incumbent plants before the acquisition (top panel). Most of the difference is instead driven by the ratio of predicted output (or combined total inputs) to total assets, $\exp(\hat{Y}_i)/C_i$. The numbers in the top panel imply that for the same amount of capital employed, incumbent plants manage to mobilize almost 30 percent more of their combined inputs toward production than do acquired plants in pre-acquisition years.

The decompositions of changes in acquired plants' gross margins in the table's bottom

Reassuringly, the same holds when we compare most other computed differentials with the corresponding regression coefficients. Some discrepancy is to be expected, of course, as the regressions include acquisition fixed effects.

²⁶ As we calculate TFPQ using output adjusted to a standard 20-count yarn as explained in Appendix A, we similarly adjust plants' prices (which again are expressed per unit weight in the data). Specifically, we use the inverses of the conversion coefficients we use to adjust output. Adjusted output is obtained as $\hat{y} = ky$, where y is output measured in weight and k is the conversion coefficient applied, and the adjusted price for the same count is $\hat{p} = (1/k)p$. This procedure ensures adjusted plant revenues remain the same as in the original data.

two panels indicate input use intensity dominates early post-acquisition profitability growth, with TFPQ growth mattering relatively more in the long run. This is similar to what we observed in Table 2. In contrast to the regressions, price margins have a relatively large and statistically significant long-run contribution, and TFPQ's contribution is substantially larger than implied by Table 2.²⁷ The impact of the inputs-to-assets ratio, on the other hand, falls compared to the early post-acquisition period, although it still contributes about a third of total increase in net output value per unit assets.²⁸

TFP Measure Decompositions. As a complement to the accounting decompositions, we compare the TFPQ patterns we document above to what one would find if one had more conventional producer microdata. Recall that our TFPQ metric has two distinguishing characteristics: it measures output in physical units and it measures inputs as service flows rather than stocks. Typical producer microdata contains only revenues as an output measure and capital and labor stocks for inputs. As such, standard TFP measures tend to confound price and output differences and embody variations in input utilization rather than conditioning on the plant actually operating. Because the accounting decompositions above suggest a prominent role for input utilization in explaining profitability differences across mills, this latter distinction between our TFPQ and standard TFP metrics may be salient in our results.

We compute two alternative measures of TFP to explore this issue. One measures TFPQ

²⁷ The reason for this difference is that the year fixed effects in regressions estimations effectively remove a time trend in productivity, while the TFPQ measure presented in Table 5 is best interpreted as inclusive of industry-wide productivity growth over time (which is itself partly a consequence of the acquisition process). Thus the regression coefficients give us a lower bound for TFPQ's contribution to profitability growth (as they are stripped of any effect acquisitions may have on industry-wide productivity improvement over time), while the differentials in Table 5 represent the upper bound ("loading" all industry-wide productivity improvement into acquisition effects). We recomputed Table 5 using residuals from the production function estimations demeaned by industry-year averages and confirmed that TFPQ differentials in that case are closely aligned in magnitude with the regression coefficients.

²⁸ We show in Appendix G that this is not driven by a decline in capacity utilization rates. These in fact increase further in the long run, though at a more modest rate (we see this in another setting immediately below). The fall in the input-per-asset ratio observed in the bottom panel of Table 5 is instead an accounting phenomenon explained by a drop in the ratio of plant capacity to total firm assets. This drop is in turn driven by a big increase in acquired plants' retained earnings (and therefore their shareholder capital). More detailed analysis of balance sheets (see Appendix G) indicates that retained earnings growth is related to firms' increasing use of accumulated profits to finance new construction toward the end of the sample, where many of our late post-acquisition observations fall.

without conditioning on the plant actually operating. Specifically, when computing the residual of the production function (3) to obtain TFPQ, instead of the input flows (spindle-days and worker-days) used in our benchmark TFPQ metric, we use capital and labor stocks (spindles and workers). This measure, which we call TFPQU (“U” for “unconditional” on operating), is shifted by disparities in input utilization. Higher (lower) input utilization shows up as higher (lower) TFPQU for a plant.

Our second alternative TFP measure further modifies TFPQU by adding to it the plant’s logged output price. This mimics the revenue-based output measure typically used in the literature. By construction, any difference between patterns in this productivity measure (which we refer to as TFPR, using the standard nomenclature for revenue-based productivity) and TFPQU comes from price differences across producers.

Using TFPR in specifications (5) and (6) reveals how our productivity results would look if we had only standard producer-level microdata. Any contrast between such results and those obtained above using our benchmark TFPQ metric reveals the combined influence of plant-level heterogeneity in prices and input utilization. We can further use TFPQU to decompose this contrast into the separate influences of price and input utilization differences.

The estimates of (5) and (6) with our three TFP measures are in Table 6. The left half of the table shows the within-plant specification (5), the right half the within-acquisition difference-in-difference specification (6). The results for our benchmark productivity measure TFPQ are the same as those in Tables 2 and 3. We report them again here for convenience.

[Table 6 about here]

The TFPR results indicate a roughly 18 percent rise in this productivity measure relative to the pre-acquisition baseline and an almost 34 percent increase in the longer term. These changes are 2-3 times the size of the TFPQ gains estimated above. The difference-in-difference results for TFPR tell a similar story. The interaction terms indicate acquired plants’ TFPR levels rose post-acquisition about 16 percent more than among their acquiring firms’ incumbent plants. The same gap in TFPQ terms was only about nine percent. Also unlike the TFPQ regressions, both main effects are significant. Before acquisition, purchased plants had on average about eight percent lower TFPR than their acquirers’ plants.

The specifications using TFPQU offer insights as to the source of the differences in the TFPR and TFPQ results. In both the within-plant and difference-in-difference specifications, the

estimated TFPQU changes are quantitatively closer to their TFPR analogs than their TFPQ counterparts. In fact, we cannot reject the hypothesis that the TFPR and TFPQU coefficients are equal. Because TFPQU is shifted by variation in input utilization but is not affected by price differences, the close tracking of TFPR by TFPQU implies that input utilization heterogeneity explains most of the difference between our benchmark TFPQ results and those obtained using the TFPR metrics typical of the literature. Price heterogeneity across plants, on the other hand, explains little. Both of these results are consistent with both the regression and accounting decomposition exercises above, which found few price differentials but substantial variation in capacity utilization.

Putting these results together offers an explanation for the patterns documented in Sections III.A and III.B. Profitability and productivity conditional on operating both rise at acquired plants after acquisition. In the short run, almost all profitability increases are the result of increased input utilization rates rather than greater productivity conditional on operating. In the longer run, conditional productivity TFPQ plays a larger role in raising profitability, though the contribution of increased utilization is of similar size. This connection can be seen even more clearly in Figure A12 in Appendix K where we present estimated effects of acquisitions on TFPQ and TFPQU using a full set of *annual* pre- and post-acquisition year dummies.

E. The Link from Profitability to Productivity: The Role of Demand Management

Why were stronger firms able to utilize their inputs so much more than weaker firms? In this section we tie these utilization differences to companies' abilities to manage the industry's inherent demand variations.

As we discussed in Section I, a lack of price differentiation does not mean that output-market conditions were equivalent across firms. To quantitatively explore possible differences in firms' demand-facing operations, we investigate patterns in plants' finished goods inventory and accrued revenues on delivered output (that is, the payment for which is in arrears). We choose these metrics because they may indicate when a plant is having difficulty finding buyers in a timely manner or finding buyers who can be relied upon to disburse payments on time. These conditions in turn may explain capital utilization differences.

Table 7 shows producers' ratios of period-end finished goods inventories, accrued revenues, and the sum of these ("unrealized output" for short) to their output over the period. We

split the sample by the same plant categories as in the previous decompositions.²⁹

[Table 7 about here]

The top panel shows that incumbent plants' ratios of unrealized output to their total produced output value were about 60 percent lower than that of acquired plants before acquisition. The bottom two panels indicate that after acquisition, acquired plants' unrealized output ratios fell 60 percent within the first three years and another 10 percent after that. Within-acquisition comparisons of acquired and incumbent plants (not shown) yield similar patterns. Thus whatever management abilities allowed acquirers to sustain lower unrealized output was transferred to their acquired mills after purchase.

As to the specific sources of cotton spinning firms' abilities to manage demand, there are several potential explanations. While many of these are difficult to quantify, one important factor already mentioned in Section I was that in low-demand times, major trading houses appeared to limit their purchases by "sticking" with certain producers rather than cutting prices. At the time, big trading houses were still much stronger financially than most spinning firms, and they often had to extend credit to the latter (either directly or through forward purchases) during business downturns (Takamura, 1971, I: 323-325; II: 60-62). High risks associated with this led the traders to favor reputable and well run industry producers with whom they had established long-term relationships. In turn, this allowed those producers to sustain more consistent operations, resulting in the lower inventories and higher utilization levels observed above.

To explore this possibility quantitatively, we used the 1898 edition of *Nihon Zenkoku Shoukou Jinmeiroku*, a nationwide registry of names of traders and manufacturers, to extract the names of individuals likely to play the most prominent role in cotton spinners' output markets. This yielded a list of 154 individuals.³⁰ We then matched these individuals to the lists of board

²⁹ Finished goods inventories and accrued revenues are positively correlated in the data, but the correlation is modest, about 0.22 for both incumbent and acquired plants. There may be some direct connection between the two, as having difficulty finding reputable buyers in a timely fashion might lead a firm to reach out to lesser buyers who are more likely to fall into arrears. Therefore, total unrealized output seems to be the best metric to measure demand-facing operations efficiency. Nevertheless, all three metrics paint a consistent picture in Table 7.

³⁰ These individuals fit into groups meeting one of three criteria. One group included 98 cotton yarn and yarn-related traders across Japan who paid more than 50,000 yen in operating tax that year. A second group included 25 individuals listed as board members of the 4 largest incorporated cotton yarn-related trade companies (Naigaimen, Nihon Menka, Nitto Menshi and Mitsui Bussan). Finally, the third group includes the 31 board members and traders

members and top 10-12 shareholders of the 67 firms for which we have company reports in 1898 (this is 90 percent of firms operating that year). Of a total of 1,197 board members and top shareholders, 128 were on the list of the 154 most prominent traders described above. 33 of the 67 firms had at least one prominent trader among its board members and top shareholders. We create an indicator equal to 1 if the firm is one of these 33 or one of two more firms for which firm histories (Kinugawa, 1964) clearly indicated connectedness to major traders at their inception (we refer to these as “in-network” firms) and 0 otherwise (“out-of-network” firms).

We then tested whether a producer’s relationship to trading houses is reflected in the performance metrics we explored above. Table 8 compares the means for in-network and out-of-network firms of TFPQ, TFPQU, ROCE, ratios of unrealized output to the value of output, spindle utilization rates, and count-adjusted prices residuals. (Figures A6-A11 in Appendix H plot the corresponding distributions.) Since our in- or out-of-network classification is based primarily on the 1898 shareholders and board composition data, we limit our attention to years 1898-1902 to obtain a reasonable number of observations while not going too far forward, as board and shareholders as well as traders’ importance of course changed over time.

[Table 8 about here]

The results in Table 8 show that both average TFPQ levels and especially average TFPQU levels—which register variations in capacity utilization as productivity differences—of in-network firms’ plants are significantly higher than those of out-of-network firms. We observe large ROCE differences across the two sets of plants as well. Furthermore, being in-network is associated with a roughly 40 percent drop in plants’ unrealized output ratios. These mean effects are reflected broadly across the distribution of plants: both the ROCE and unrealized output ratio distributions of in-network firms are basically shifts of the corresponding out-of-network distributions (see Figures A8-A9 in Appendix H). In-network firms also have higher capacity utilization and prices, although these differences are relatively small and are not equally pronounced across the distributions. The distributions of price residuals of in- and out-of-network plants in particular are quite similar except for their far left and right tails, where some plants of in-network firms sell at very high prices (Figures A10 and A11 in Appendix H).

Overall, these results suggest that close relationships between industry producers and prominent traders allowed connected producers to manage demand fluctuations more effectively,

registered at the Osaka cotton and cotton yarn exchange.

particularly with regard to being able to operate with lower average inventories and greater capacity utilization levels. Notably, in-network firms were also more likely to acquire other firms in the future; the sample probability of being a future acquiring firm is 0.79 for in-network firms as opposed to 0.21 for out-of-network firms. Hence, relationships with traders' networks can help explain why initial profitability gaps existed, and why they were closed by acquisition. The accompanying TFPQ gains—improvements in efficiency even conditioning on operating—are consistent with this mechanism if demand management is correlated with broader managerial abilities that raised operational efficiency. We explore this connection in Section IV below.

Another related factor that contributes to better plant and firm performance is having chief engineers with formal technical education. Such engineers were scarce in Japan at the time. Indeed, in 1898 we counted only 14 educated engineers supervising operations at 18 of the 76 firms for which we have operational data in that year.³¹ (Two engineers provided their services to multiple firms located near one another with overlapping shareholders' interests.) We created an indicator variable for whether the firm had a formally educated engineer in charge in 1898 and repeated the comparisons conducted in the main text with regard to in- and out-of-network producers. The results are presented in Table 9.

[Table 9 about here]

The table shows that having formally educated engineers in charge has effects similar to being in-network, but even more strongly pronounced in TFPQ. Estimating regressions (not shown) including both in-network and educated-engineer indicators also shows that the performance differences associated with being in-network and having an educated engineer in charge are largely independent. Still, it is worth noting that 12 of the 18 firms with educated engineers in charge were also in-network firms, including 8 of the 14 acquiring firms and all five “serial acquirers.” Examining the interaction between demand management and technical competence is a fascinating task for future research as more complete data presents itself.

F. Robustness

As already mentioned, we have conducted several robustness checks. We relegate the

³¹ Saxonhouse (1977) was the first to analyze the role of educated engineers in this industry but the main data source he used starts in the 1910s. We have matched the data he used with the firms' histories in Kinugawa (1964) to obtain the list of educated engineers at the firm level around 1898. See Braguinsky and Hounshell (2014) for more details.

details and presentation of the results to Appendix F for the sake of parsimony, but we briefly describe the exercises here.

Our benchmark results above use TFPQ estimates obtained from a production function estimated via one of the three specifications discussed by De Loecker (2013). While this presents a way to deal with the classic transmission bias arising from a correlation between unobserved productivity changes and producers' input choices, we also estimated our specifications with TFPQ constructed via alternative methods, including simple OLS, the Blundell and Bond (1998) "system GMM" estimator, and two other specifications suggested by De Loecker (2013). In all cases, the results were qualitatively and quantitatively similar to those above.

While matching by acquisition cases seems to be the most natural approach in our context, we did explore other matching strategies. We matched acquired plants on pre-acquisition characteristics and on pre-acquisition productivity trends 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, presented in Tables A10 and A11 of Appendix F, are very similar to the ones presented here.

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 repeated this process 1,000 times and calculated the sample mean of the estimated coefficients relating "acquisition" to outcomes. In most cases, the magnitudes were only fractions of their analogs from the true acquisition sample.

IV. A Mechanism

Our empirical results point to some sort of demand management ability, reflected empirically in capital utilization levels and unrealized output rates, as being related to productivity and profitability variation in both the cross section and over time within acquisition events. Here we offer a simple theory that elucidates one channel through which fundamental heterogeneity across owners/managers leads to variations in such ability, and through this, TFPQ and profitability. If this heterogeneity is "carried" in acquisitions by owners/managers into target plants' operations, it explains the productivity and profitability changes surrounding acquisition events estimated above. That said, it is possible that other possible mechanisms could explain the data, and we cannot test the model's time allocation implications directly because we do not

observe 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.

A. Plant Production and Demand

For simplicity, we focus on a case where each firm initially operates a single plant before an acquisition opportunity arrives. A firm has access to the following production technology:

$$(9) \quad y = g(m)x\omega,$$

where ω is the given quality of a plant, and x is the composite input of appropriately weighted labor and capital. For example, if the technology is Cobb-Douglas, the composite would be the plant's inputs raised to their respective input elasticities. The function $g(m)$ is a flow of in-firm services provided by the plant manager to increase outputs from a level of $x\omega$. The variable m is the manager's time allocated to managing production. This is divided into time spent ensuring that the plant operates at full capacity (therefore affecting input utilization), and time spent improving efficiency of operations themselves. For example, the former may involve making sure that machines are in working condition and that there are always enough workers to operate them.³² The time spent improving operational efficiency, on the other hand, would involve monitoring the production process, receiving and acting upon reports from workers and improving quality control.³³ To ease notation, assume $g(m) = \sqrt{uv}$, where u denotes the time spent improving the frequency of operation (so that utilized input is given by $\tilde{x} = \sqrt{u}x$), and v is the time spent improving plant performance conditional on operating, thus augmenting the intrinsic plant productivity, equal to $\tilde{\omega} = \sqrt{v}\omega$.³⁴ We assume the total time spent managing the plant $m = u + v$ is bounded between 0 and some $\gamma > 0$, the manager's *effective* time endowment. We discuss this more below.

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

$$(10) \quad x^* = \frac{y}{\sqrt{uv}\omega},$$

³² Saxonhouse (1971) describes the problem of absenteeism in the industry.

³³ Anecdotes about the importance of this sort of managerial activity are in, e.g., Kuwahara (2004) and Appendix B.

³⁴ Diminishing returns are not necessary for the results below to hold. In particular, all of the analyses in this section go through if we instead assume input utilization and augmented plant quality are simply proportional to managerial time spent on these activities, so that $\tilde{x} = ux$ and $\tilde{\omega} = \omega x$, although derivations become more cumbersome.

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

The plant takes output price p (determined by the exchanges) as given, but its quantity sold depends on managerial time allocation. Namely, it sells $\gamma - m$ units. Revenues are then

$$(11) \quad r = p(\gamma - m).$$

The quantity sold $\gamma - m$ is the channel through which we introduce the notion of demand management; the plant's demand depends on the time the manager allocates to selling product. Because m is the total time the manager devotes toward production, other things equal, a higher m means less demand for output.

B. Optimal Allocation of Manager's Time

From (10) and (11), the plant owner's time allocation problem is

$$(12) \quad \max_{u,v} (\gamma - u - v) \left(p - \frac{1}{\sqrt{uv}\omega} \right),$$

where we have made use of $m \equiv u + v$. That is, the plant's owner allocates his time between managing plant production and managing demand (sales) so as to maximize profits.³⁵ The optimal resource allocation problem (12) captures the fundamental tradeoff faced by the manager: devoting more time to managing sales results in lower operational frequency and/or efficiency, and vice versa. The constraint is set by the effective time endowment γ ; a higher γ reduces the lost revenue from any m . The parameter γ is thus interpreted as "demand management ability"; this can include skill at building networking relationships with trading houses, a reputation for reliable delivery, and perhaps the ability to effectively collect debt. It might also be enhanced by having an educated engineer in charge of the plant, which presumably allows the owner to spend more time managing sales and less on technical productivity issues.

It is easy to see (see Appendix I for the proof) that at the optimum, $u = v = m/2$. We can thus restate (12) in terms of the optimal time allocated to production management, m :

$$(13) \quad \max_m (\gamma - m) \left(p - \frac{2}{\omega m} \right).$$

³⁵ We assume that p is greater than the plant's marginal cost for at least some $m_0 < \gamma$, so that operation is profitable for all $m_0 < m < \gamma$. The $(\gamma - m)$ function limits the size of the plant, though it would be easy to introduce increasing marginal costs or downward sloping residual demand (say as in a monopolistically competitive structure) if one wanted to further constrain plant size.

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

$$(14) \quad m(\gamma, \omega) = \sqrt{2\gamma/p\omega} \quad \text{and} \quad \pi(\gamma, \omega) = (\sqrt{\gamma p \omega} - \sqrt{2})^2 / \omega.$$

A simple exercise yields the following results.

Lemma 1:

- (i) $\partial \tilde{x} / \partial \gamma > 0$ and $\partial \tilde{\omega} / \partial \gamma > 0$. Input utilization \tilde{x} and augmented productivity $\tilde{\omega}$ increase in γ .
- (ii) $\partial \pi(\gamma, \omega) / \partial \gamma > 0$ and $\partial \pi(\gamma, \omega) / \partial \omega > 0$; also, $\partial x^* / \partial \gamma > 0$. Profits increase in ability γ and plant quality ω , while total inputs also increase in γ .
- (iii) $\partial^2 \pi(\gamma, \omega) / \partial \gamma \partial \omega > 0$. Ability γ and plant quality ω are complements in the profit function.

Proof: See Appendix I.

Lemma 1 implies increasing returns to demand management ability that are manifested in both an increased span of control in production, x^* , and input utilization, \tilde{x} . Augmented plant efficiency $\tilde{\omega}$ also increases in demand management ability, implying that output increases with ability even conditioning on inputs and their utilization. The first feature is consistent with our decomposition results that showed more profitable firms (with higher demand management ability) had higher input utilization rates. The second feature is consistent with TFPQ, measured conditional on operating, increasing once a plant owned by a less profitable firm is acquired by a more profitable firm. We explore this point more below. It is also in line with the findings in Tables 8 that capacity utilization and TFPQ are higher for “in-network” and 9 that capacity utilization and TFPQ are higher for “in-network” firms and in Table A17 in Appendix L showing the same for firms with educated engineers in charge.

C. Mergers and Acquisitions

We employ a model of asset reallocation through acquisitions similar to Jovanovic and Braguinsky (2004) and Jovanovic and MacDonald (1994). Since our focus is on plant-level profitability and productivity changes, we limit the exposition in the main text to the basics. See Appendix I for the full setup and formalization of industry equilibria described intuitively below.

The industry evolves in three stages. In the first two, each firm can manage at most one plant. In the first stage, an initial “basic” state of technological knowledge arrives, offering entry by the industry’s first cohort of firms. The basic nature of this initial technological knowledge is

manifested in the low quality of plants, ω_1 , available for this first entry cohort. Each entrant comes into the industry with some initial demand management ability level, γ_0 . First-cohort producers have an opportunity to develop this ability above the initial level (for instance, they make connections with traders or are able to hire an educated engineer). In equilibrium at the end of the first stage, the first cohort's ability is distributed with support $[\gamma^*, \gamma_{max}]$, where γ^* is a threshold ability level and $\gamma^* \geq \gamma_0$.

The second stage begins with an unanticipated change in the state of technology (a "refinement," in Jovanovic and MacDonald, 1994). As mentioned, such a refinement occurred in Japanese cotton spinning when the industry developed new sources for raw cotton (imported from India and the U.S.). This made it possible to import state-of-the-art machines from England for the first time; see Appendix D. In the model, this is captured by a higher plant quality, $\omega_2 > \omega_1$, available to the second cohort of entrants.³⁶ In the new industry equilibrium at the end of this stage, the industry contains 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 (recall that each firm can only manage one plant at this stage). The threshold ability of a marginal surviving firm in the second-stage equilibrium, γ^{**} , is greater than the first-stage threshold γ^* . Hence some first-cohort firms exit at this stage.³⁷

The third stage is characterized by an unanticipated opening of the market for acquisitions. In this stage, each firm can potentially manage more than one plant and can replicate its plant manager quality in a newly acquired plant.³⁸ There is no new entry during this

³⁶ Jovanovic and MacDonald (1994) present a detailed account of one such refinement, the invention of the Banbury mixer, and how it affected the entry and exit of firms in the U.S. tire industry. Rajan, Volpin, and Zingales (2000) describe how another refinement, the advent of the radial, in the same industry more than half a century later led to its eclipse in the U.S. through acquisitions by foreign producers. More generally, a refinement can be interpreted as any investment-specific technological change embodied in new vintage capital or a new type of input (or both, as in our case). The issues related to such changes have been extensively studied in the macro growth literature (see, for example, Cooley, Greenwood, and Yorukoglu, 1997, and Jovanovic and Yatsenko, 2012), and may account for a major part of economic growth (Greenwood, Hercowitz, and Krusell, 1997).

³⁷ See Jovanovic and MacDonald (1994). In our data, 10 out of 21 firms that had operated in the industry prior to the refinement of the early 1890s remained small and eventually exited by shutting down their plants.

³⁸ This is consistent with a situation where management quality is tied primarily to a set of practices (e.g., Bloom and Van Reenen, 2007) rather than person-specific human capital.

stage. Profitability and productivity growth is attained through the reallocation of production facilities from firms with low demand management ability to those with high ability.

In Appendix I we construct and formally solve for the asset reallocation equilibrium in this stage. The key characteristics of this equilibrium are intuitive and can be summarized as follows: (i) all second-cohort owners of high-quality plants sell their assets and their firms exit; (ii) among first-cohort owners of low-quality plants, those with higher ability buy plants from those with lower ability; (iii) because profits $\pi(\omega, \gamma)$ are increasing in γ , the gains from acquisitions are the highest when first-cohort entrants with especially high ability acquire high-quality ω_2 plants formerly managed by the low-ability second cohort entrants.

D. Implications for Productivity and Profitability

We now derive implications of the merger and acquisition process outlined above for productivity and profitability of acquired plants. As we will show, the implications are consistent with the patterns we document in Section III.

To discuss the implications for productivity, note that a plant's TFPQ in the model is $\text{TFPQ} \equiv y/u(\gamma)x = v(\gamma)\omega$. Lemma 1(i) implies that for a given ω , TFPQ will increase with the acquiring firm's managerial ability γ . Similarly, Lemma 1(ii) says that profits increase with this ability. Because all acquisitions involve firms with higher ability acquiring a plant managed by a firm with lower ability, these imply

Proposition 1: Both the productivity and the profitability of acquired plants rise after acquisition.

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

Proposition 2: After an acquisition, the acquired plant profits increase by more than TFPQ.

Proof: See Appendix I.

The key intuition behind these propositions is that the new manager's superior ability to manage demand/sales allows more time to be allocated to managing the production process without sacrificing sales at any given price.

We next derive implications that allow us to compare the pre-acquisition levels of

productivity and profitability of acquired plants with those of acquiring plants. We can express the total derivative of the profit function as

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

The first term in (15) reflects how plant quality differentials between acquired and acquiring plants affect profits, while the second term reflects the effect of demand management ability differentials. An acquiring plant has a higher-ability owner—i.e., $d\gamma > 0$ —while its quality is equal to or lower than an acquired plant’s quality—i.e., $d\omega \leq 0$. The nature of the equilibrium implies, however, that the profit of an acquiring plant is always higher in the pre-acquisition period than that of an acquired plant. To see this, suppose that $d\omega < 0$, so a first-cohort firm acquires a second-cohort plant. Because low-ability first-cohort firms (that achieved the same profit as the second cohort firms) in the pre-acquisition period also exit the industry, a first-cohort acquirer must have an ability level greater than that which generates profits just equal to that of a second-cohort acquired firm. In Appendix I we formally establish the following:

Proposition 3: The pre-acquisition TFPQ of an acquiring plant could be less than that of an acquired plant even though the pre-acquisition profitability of an acquiring plant is always higher than that of an acquired plant.

Proof: See Appendix I.

A simple numerical example of the model in Appendix J illustrates how the mechanism outlined above can deliver all the patterns observed in our empirical analyses.

V. 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 around the start 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 vagaries of demand in the industry. These demand management problems resulted in consistently higher inventory and unrealized output levels along with lower capacity utilization among acquired producers, reducing returns on capital. We show that once purchased by more profitable firms, the acquired plants saw drops in inventories and unrealized output, gains in capacity utilization, and growth in both productivity and profitability. These patterns are 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 that may plausibly operate in many markets. Finally, 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

- Akerberg, Daniel, Kevin Caves, and Garth Frazer, 2006. “Structural Identification of Production Functions.” MPRA Paper 38349, University Library of Munich, Germany, available online at <http://mpra.ub.uni-muenchen.de/38349/1/ACF20withtables.pdf>
- Andrade, Gregor, Mark Mitchell, and Erik Stafford, 2001. “New Evidence and Perspective on Mergers.” *Journal of Economic Perspectives*, 15(2), 103-120.
- Atalay, Enghin, 2014. “Materials Prices and Productivity.” *Journal of the European Economic Association*, 12(3), 575-611.
- Bernhofen, Daniel M., and John C. Brown, 2004. “A Direct Test of the Theory of Comparative

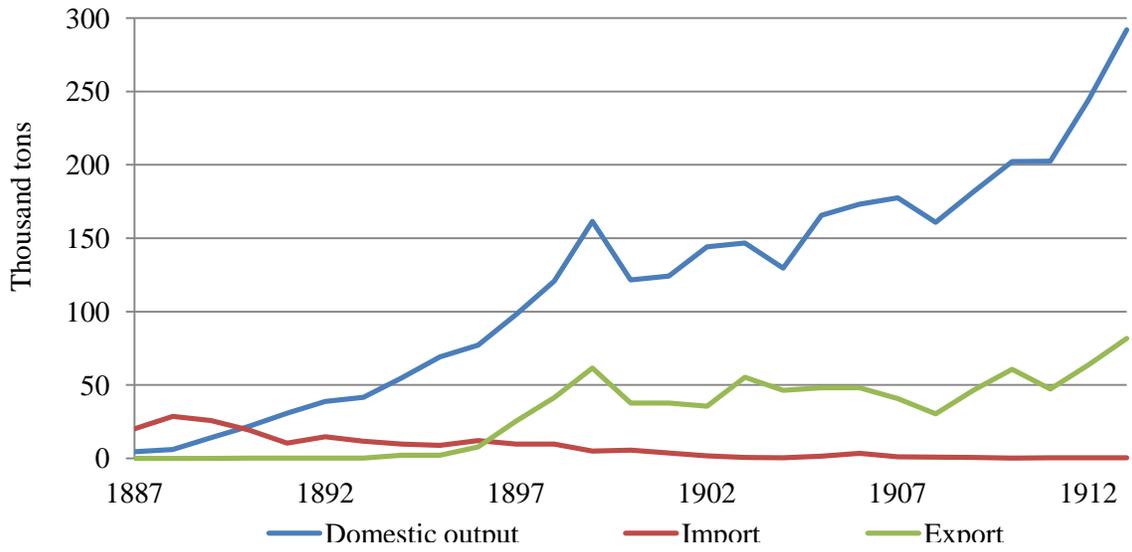
- Advantage: The Case of Japan.” *Journal of Political Economy*, 112(1), 48-67.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie and John Roberts, 2013. “Does Management Matter? Evidence from India.” *Quarterly Journal of Economics*, 128(1), 1-51.
- Bloom, Nicholas, and John Van Reenen, 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 A. Hounshell, 2014. “Spinning Tales about Japanese Cotton Spinning: Saxonhouse (1974) and Lessons from New Data.” http://www.andrew.cmu.edu/user/sbrag/spinning_tale.pdf
- Braguinsky, Serguey, and David C. Rose, 2009. “Competition, Cooperation, and the Neighboring Farmer Effect.” *Journal of Economic Behavior and Organization*, 72(1), 361-376.
- Cooley, Thomas F., Jeremy Greenwood, and Mehmet Yorukoglu, 1997. “The Replacement Problem.” *Journal of Monetary Economics*, 40, 457-499.
- David, Joel, 2014. “The Aggregate Implications of Mergers and Acquisitions.” <https://sites.google.com/site/joelmichaeldavid/research>
- De Loecker, Jan, 2013. “Detecting Learning by Exporting.” *American Economic Journal: Microeconomics*, 5(3), 1-21.
- Foster, Lucia, John Haltiwanger, and Chad Syverson, 2008. “Reallocation, Firm Turnover and Efficiency: Selection of Productivity or Profitability?” *American Economic Review*, 98(1), 394-425.
- Fuken Tokei (Prefectural Statistics), 1912-1920, various prefectures (in Japanese).
- Geppo* (Monthly Bulletin), 1889-1920 (in Japanese), Dainihon Boseki Rengokai (All-Japan Cotton Spinners Association), various issues.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, 1997. “Long-Run Implications of Investment-Specific Technological Change.” *American Economic Review*, 87(3), 342-362.
- Hausmann, Ricardo, and Dani Rodrik, 2003. “Economic Development as Self-discovery.” *Journal of Development Economics*, 72(2), 603-633.

- Jensen, Michael C., and Richard S. Ruback, 1983. "The Market for Corporate Control. The Scientific Evidence." *Journal of Financial Economics*, 11, 5-50.
- Jovanovic, Boyan, and Serguey Braguinsky, 2004. "Bidder Discounts and Target Premia in Takeovers." *American Economic Review*, 94(1), 46-56.
- Jovanovic, Boyan and Glenn MacDonald, 1994. "The Life Cycle of a Competitive Industry." *Journal of Political Economy*, 102 (2), 322-347.
- Jovanovic, Boyan, and Peter Rousseau, 2002. "The Q-Theory of Mergers." *American Economic Review*, 92 (2), 198-204.
- Jovanovic, Boyan, and Yuri Yatsenko, 2012. "Investment in Vintage Capital." *Journal of Economic Theory*, 147, 551-569.
- Kaplan, Steven (ed.), 2000. *Mergers and Productivity*. Chicago: University of Chicago Press.
- Katayama, Hajime, Shihua Lu, and James R. Tybout, 2009. "Firm-level Productivity Studies: Illusions and a Solution." *International Journal of Industrial Organization*, 27(3), 403-413.
- Kinugawa, Taiichi, 1964. *Hompo Menshi Boseki Shi* (History of the Japanese Cotton Spinning Industry) in 7 volumes (in Japanese), Tokyo: Hara Shobo (reprinted from the original 1942 edition).
- Kokajo (Company reports), 1896-1920 (in Japanese), Osaka: Osaka University Library.
- Kuwahara, Tetsuya, 2004. "The Development of Factory Management in Japan During the Early Stages of Industrialization." Chapter 16 in *The Fibre That Changed the World*, Douglas A. Farnie and David J. Jeremy, eds., Oxford: Oxford University Press, 495-519.
- Lichtenberg, Frank R., and Donald Siegel, 1987. "Productivity and Changes in Ownership of Manufacturing Firms." *Brookings Papers of Economic Activity*, 1987(3), 643-673.
- 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 J. Mark Ramseyer, 2000. "Corporate Governance in Transitional Economies: Lessons from the Prewar Japanese Cotton Textile Industry." *Journal of Legal Studies*, 29(1), 171-203.

- Menshi Boseki Jijo Sankosho (Handbook of the Cotton Spinning Industry)*, 1903-1920 (in Japanese), Dainihon Boseki Rengokai (All-Japan Cotton Spinners Association).
- Murayama, Takashi, 1961. *Sekai Mengyo Hattatsu Shi (History of Cotton Industry Worldwide)*, in Japanese). Osaka: Japan's Cotton Spinners Association.
- Nihon Choki Tokei Soran (Historical Statistics of Japan), 1987. In 4 volumes (in Japanese), Tokyo: Japan Statistical Association.
- Noshokomu Tokei Nempo (Statistics on Agriculture, Commerce and Industry), Noshokomusho (Ministry of Agriculture, Commerce and Industry) (in Japanese), 1896-1911.
- Nihon Zenkoku Shoukou Jinmeiroku*, (Nationwide registry of names of traders and manufacturers) (in Japanese), edited by Suzuki Kihachi and Seki Itaro, 1898.
- Nguyen, Sang V. and Michael Ollinger, 2006. "Mergers and Acquisitions and Productivity in the U.S. Meat Products Industries: Evidence from the Micro Data." *American Journal of Agricultural Economics*, 88(3), 606-616.
- 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.
- Rajan, Raghuram, Paolo Volpin, and Luigi Zingales, 2000. "The Eclipse of the U.S. Tire Industry." in Steven Kaplan, ed., *Mergers and Productivity*, Chicago: University of Chicago Press.
- Rhodes-Kropf, Matthew, and David Robinson, 2008. "The Market for Mergers and the Boundaries of the Firm." *Journal of Finance* 63(3), 1169-1211.
- Roll, Richard, 1986. "The Hubris Hypothesis of Corporate Takeovers." *Journal of Business* 59(2): 197-216.
- Rose, Mary B., 2000. *Firms, Networks and Business Values: The British and American Cotton Industries since 1750*, Cambridge: Cambridge University 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.

- Schoar, Antoinette, 2002. "Effects of Corporate Diversification on Productivity." *Journal of Finance*, 57(6). 2379-2403.
- Shleifer, Andrei, and Robert Vishny, 2003. "Stock Market Driven Acquisitions." *Journal of Financial Economics*, 70, 295-311.
- 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. Tokyo: Hanawa Shobo.

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



Source: Nihon Choki Tokei Soran, our estimates.

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

	Acquiring plants	Acquired plants		Exiting plants
		First cohort	Second cohort	
TFPQ	0.066 (0.156)	0.034 (0.225)	0.156 (0.229)	-0.211 (0.552)
Profit per paid-in share	0.274 (0.205)	0.185 (0.074)	0.159 (0.149)	0.159 (0.101)
Price (yen/400lb)	94.7 (6.5)	92.4 (3.8)	92.8 (7.4)	91.7 (7.0)
Logged price residual	-0.041 (0.151)	0.014 (0.040)	0.012 (0.041)	0.021 (0.062)
Main count of yarn produced	21.5 (11.5)	17.5 (2.6)	17.2 (4.7)	13.9 (5.6)
Days in operation	323.7 (29.8)	315.9 (29.5)	300.6 (55.6)	278.6 (56.8)
Equipment age	5.28 (3.49)	5.88 (2.76)	2.79 (1.00)	11.77 (6.69)
Firm age	9.13 (5.08)	11.06 (3.81)	3.31 (2.05)	12.54 (7.86)
Observations	32	33	32	24

Source: Authors' calculations. This table reports plant-level means and standard deviations (in parentheses) of data for various subsamples of plants. "Acquiring plants" are those owned by future acquiring firms, "acquired plants" are owned by firms that will be acquired in the future, and "exiting plants" are owned by firms that will exit in the future (not through acquisition) and be scrapped. "First cohort" are plants of firms that started operating before 1892, "second cohort" is plants of firms that started operating in or after 1892. 1896 and 1897 observations for second-cohort plants that began operations in those years are excluded. TFPQ (quantity-based total factor productivity) is estimated using De Loecker's (2013) method described in the main text. Profit per paid-in value of shares is net revenue from company reports divided by shareholders' paid-in capital. There are only 6 observations on net revenue available for exiting plants in these years. The log price residuals are from a regression of log plant-level price on count dummies and year dummies, as described in the main text. Equipment and firm age are measured in years.

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

Dependent variable	All acquisitions			By serial acquirer		
	TFPQ	Plant ROCE	Log price res.	TFPQ	Plant ROCE	Log price res.
Late pre-acquisition dummy	-0.003 (0.019)	0.020 (0.013)	0.011 (0.013)	-0.016 (0.030)	0.025 (0.016)	0.018 (0.030)
Early post-acquisition dummy	0.045* (0.026)	0.060*** (0.022)	0.036 (0.027)	0.053 (0.046)	0.106*** (0.023)	0.065 (0.063)
Late post-acquisition dummy	0.126*** (0.033)	0.089*** (0.025)	0.044 (0.034)	0.159** (0.062)	0.140*** (0.032)	0.089 (0.068)
Constant	0.603*** (0.032)	0.102*** (0.013)	0.029*** (0.010)	0.356*** (0.025)	0.079** -0.031	0.041*** (0.008)
Acquisition fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,078	891	1,118	512	472	521
Adjusted R-squared	0.734	0.639	0.097	0.695	0.625	0.082

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

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

Dependent variable	All acquisitions			By serial acquirer		
	TFPQ	Plant ROCE	Log price res.	TFPQ	Plant ROCE	Log price res.
After acquisition	-0.055*** (0.013)	-0.004 (0.012)	-0.031** (0.013)	-0.048*** (0.008)	-0.012 (0.016)	-0.026* (0.015)
Acquired plant	-0.025 (0.021)	-0.030*** (0.011)	-0.019 (0.014)	-0.032* (0.017)	-0.032** (0.013)	-0.015 (0.015)
After acquisition x Acquired plant	0.091*** (0.023)	0.040*** (0.014)	0.024 (0.017)	0.113*** (0.028)	0.058*** (0.017)	0.033 (0.025)
Constant	0.480*** (0.034)	0.145*** (0.018)	0.038*** (0.008)	0.410*** (0.008)	0.069*** (0.013)	-0.007 (0.008)
Acquisition fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,487	1,392	1,528	1,067	994	1,091
Adjusted R-squared	0.347	0.433	0.108	0.489	0.455	0.164

Source: Authors' calculations. See note for Table 2.

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

Pre-acquisition means	Acquired plants (A)	Incumbent plants (B)	Difference (B-A)	Percentage difference
ROCE	0.053	0.104	0.051	95.3***
<i>of which:</i>				
net output value/capital employed	0.193	0.257	0.065	33.5***
<i>minus:</i>				
wage cost/capital employed	0.077	0.094	0.018	22.9***
capital cost/capital employed	0.062	0.059	-0.004	-6.2***
Observations	133	269		

Pre- and early post-acquisition means	Pre-acquisition (A)	Early post-acquisition (B)	Difference (B)-(A)	Percentage difference
ROCE	0.062	0.126	0.063	101.7***
<i>of which:</i>				
net output value/capital employed	0.202	0.286	0.084	41.6***
<i>minus:</i>				
wage cost/capital employed	0.078	0.103	0.025	32.2***
capital cost/capital employed	0.062	0.058	-0.004	-7.0**
Observations	163	159		

Pre- and late post-acquisition means	Pre-acquisition (A)	Late post-acquisition (B)	Difference (B)-(A)	Percentage difference
ROCE	0.062	0.163	0.100	161.1***
<i>of which:</i>				
net output value/capital employed	0.202	0.317	0.114	56.6***
<i>minus:</i>				
wage cost/capital employed	0.078	0.103	0.025	31.7***
capital cost/capital employed	0.062	0.051	-0.011	-17.3***
Observations	163	280		

Source: Authors' calculations. The pre-acquisition time period includes observations on up to 4 years prior to acquisition. "Early post-acquisition" period includes 3 years immediately following acquisitions. "Late post-acquisition" period includes years starting from year 4 after acquisitions. Nominal variables are deflated by the annual consumer price index. Details of variable construction are explained in Appendix E. ***, and ** indicate that the corresponding difference is statistically significant at the 1 percent level, and 5 percent level, respectively, using a double-sided *t*-test.

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

Pre-acquisition means of logs	Acquired plants (A)	Incumbent plants (B)	Difference (B)-(A)	Percentage difference
ln(net output value/capital employed)	-1.791	-1.436	0.355	42.6***
<i>of which:</i>				
ln(price margin)	-1.407	-1.377	0.030	3.1
TFPQ	0.500	0.568	0.069	7.1***
ln(total input/capital employed)	-0.883	-0.627	0.256	29.2***
Observations	129	262		

Pre- and early post- acquisition means of logs	Pre-acquisition (A)	Early post-acquisition (B)	Difference (B)-(A)	Percentage difference
ln(net output value/capital employed)	-1.735	-1.392	0.343	40.9***
<i>of which:</i>				
ln(price margin)	-1.438	-1.367	0.071	7.4**
TFPQ	0.499	0.568	0.069	7.2***
ln(total input/capital employed)	-0.795	-0.593	0.202	22.4**
Observations	157	157		

Pre- and late post- acquisition means of logs	Pre-acquisition (A)	Late post-acquisition (B)	Difference (B)-(A)	Percentage difference
ln(net output value/capital employed)	-1.735	-1.275	0.460	58.4***
<i>of which:</i>				
ln(price margin)	-1.438	-1.316	0.122	13.0***
TFPQ	0.499	0.685	0.187	20.5***
ln(total input/capital employed)	-0.795	-0.644	0.151	16.3***
Observations	157	278		

Source: Authors' calculations. The pre-acquisition time period includes observations on up to 4 years prior to acquisition. "Early post- acquisition" period includes 3 years immediately following acquisitions. "Late post-acquisition" period includes years starting from year 4 after acquisitions. Nominal variables are deflated by the annual consumer price index. Details of variable construction are explained in Appendix E. *** and ** indicate that the corresponding difference is statistically significant at the 1 percent level and 5 percent level, respectively, using a double-sided *t*-test.

Table 6. Total factor productivity changes around acquisition events.

Within-acquired plants estimations				“Difference-in-difference” estimations			
Dependent variable	TFPR	TFPQU	TFPQ	Dependent variable	TFPR	TFPQU	TFPQ
Late pre-acquisition dummy	0.020 (0.051)	-0.027 (0.044)	-0.003 (0.019)	After acquisition	-0.083** (0.035)	-0.042* (0.024)	-0.055*** (0.013)
Early post-acquisition dummy	0.168*** (0.058)	0.104*** (0.038)	0.045* (0.026)	Acquired plant	-0.075** (0.035)	-0.093*** (0.029)	-0.025 (0.021)
Late post-acquisition dummy	0.290*** (0.080)	0.211*** (0.050)	0.126*** (0.033)	(After acquisition) x (Acquired plant)	0.148*** (0.042)	0.139*** (0.033)	0.091*** (0.023)
Constant	0.750*** (0.065)	0.304*** (0.053)	0.603*** (0.032)	Constant	1.197*** (0.083)	0.393*** (0.042)	0.480*** (0.034)
Acquisition fixed effects	Yes	Yes	Yes	Acquisition fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Year fixed effects	Yes	Yes	Yes
Observations	1,047	1,077	1,078	Observations	1,430	1,486	1,487
Adjusted R-squared	0.824	0.478	0.734	R-squared	0.636	0.318	0.347

Source: Authors’ calculations. TFPQ is our benchmark TFP measure that uses capital and labor services flows as inputs. TFPQU is “unconditional TFPQ,” using instead plants’ total capacity and labor input. TFPR equals TFPQU plus logged plant-specific output price. The omitted category includes period three years or more prior to acquisition. The omitted category includes period three years or more prior to acquisition. Robust standard errors clustered at the acquisition event level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 7. Inventory and accrued payments to output value ratios: incumbent and acquired plants and acquired plants pre- and post-acquisition

Means	Acquired plants (A)	Incumbent plants (B)	Difference (B-A)	Percentage difference
Inventory/produced output (C)	0.046	0.018	-0.028	-61.0***
Accrued revenues/produced output (D)	0.031	0.015	-0.016	-50.6***
Unrealized/produced output (C)+(D)	0.078	0.033	-0.045	-57.4***
Observations	113	195		

	Pre-acquisition (A)	Early post-acquisition (B)	Difference (B-A)	Percentage difference
Inventory/produced output (C)	0.048	0.013	-0.034	-72.0***
Accrued revenues/produced output (D)	0.029	0.020	-0.010	-32.4**
Unrealized/produced output (C)+(D)	0.078	0.032	-0.046	-59.4***
Observations	139	100		

	Pre-acquisition (A)	Late post-acquisition (B)	Difference (B-A)	Percentage difference
Inventory/produced output (C)	0.048	0.009	-0.039	-81.5***
Accrued revenues/produced output (D)	0.029	0.015	-0.014	-48.1***
Unrealized/produced output (C)+(D)	0.078	0.023	-0.055	-70.6***
Observations	139	124		

Source: Authors' calculations. The pre-acquisition time period includes observations on up to 4 years prior to acquisition. "Early after acquisition" period includes 3 years immediately following acquisitions. "Late after acquisition" period includes years starting from year 4 after acquisitions. *** and ** indicate that the corresponding difference is statistically significant at the 1 percent level and 5 percent level, respectively, using a double-sided *t*-test.

Table 8. Plant and firm performance metrics (1898-1902), in-network and out-of network firms

Outcome	Out-of-network (A)	In-network (B)	Difference (B-A)
TFPQ	0.433	0.488	0.055***
TFPQU	0.117	0.241	0.123***
ROCE	0.023	0.059	0.037***
Unrealized output ratios	0.127	0.084	-0.043***
Spindle utilization rates	0.739	0.781	0.043**
Logged price residuals	-0.025	0.018	0.044***
Observations	127	170	

Source: Authors' calculations. *** and ** indicate that the corresponding difference is statistically significant at the 1 percent level and 5 percent level, respectively, using a double-sided *t*-test.

Table 9. Plant and firm performance metrics in 1898-1902
by firms with and without educated engineers

Outcome	No formally educated engineer (A)	Formally educated engineer (B)	Difference (B-A)
TFPQ	0.435	0.517	0.082***
TFPQU	0.131	0.286	0.156***
ROCE	0.024	0.072	0.047***
Unrealized output ratios	0.119	0.078	-0.042***
Spindle utilization rates	0.746	0.792	0.046***
Logged price residuals	-0.014	0.021	0.035**
Observations	188	109	

Source: Authors' calculations. *** and ** indicate that the corresponding difference is statistically significant at the 1 percent level and 5 percent level, respectively, using a double-sided *t*-test.