

Very Preliminary—Please do not cite or circulate without permission

Ownership Turnover, Management Change and Productivity: Evidence from the Japanese Cotton Spinning Industry

Serguey Braguinsky, Carnegie Mellon University

Atsushi Ohyama, Hokkaido University

Tetsuji Okazaki, University of Tokyo

Chad Syverson, University of Chicago Booth School of Business

December 2012

1. Introduction

The influence of changes in corporate control of assets on productivity has been a focus of theoretical and empirical research for some time.¹ In principle, mergers and acquisitions can reallocate control of productive assets to entities that are able to apply them more efficiently. Besides increasing the productivity of the individual production units that are merged or acquired, a broader process of such reallocations can also lead to aggregate productivity growth. Such a mechanism therefore has the potential to explain patterns of productivity at both the micro and macro levels.

Implicit in the story of this mechanism—though not often treated explicitly in the empirical work on the subject—is the notion that productivity growth occurs when changes in ownership and control put assets in the hands of more able managers. 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.²

Despite the comfortable intuition of this logic, previous research has not been fully conclusive about the effects of ownership and management turnover, particularly regarding the nature of any measured productivity growth and the particular manners in which this

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

² Examples include Bloom and Van Reenen (2007 and 2010) and Bloom et. al (2012).

growth is obtained. This reflects in part the inherent limitations of the data available in the earlier studies. For instance, this work often had to treat all mergers and acquisitions as basically the same, ruling out the possibility that changes in ownership are prompted for various reasons that might be associated with heterogeneous productivity effects. Further, this work could not cleanly distinguish between physical (quantity) productivity and revenue productivity. This distinction can be important (Foster et al., 2008). It is not particularly surprising, excepting bounded rationality or agency problems, that acquisition deals could yield expectedly profitable synergies. However, such between-firm synergies need not be tied to improvements in the efficiency with which producers convert inputs to outputs. For example, mergers or acquisitions may increase market power that leads to higher output prices for the merged firm. In the typical revenue-based productivity measures of the literature (separate price and quantity information is rarely available at the producer level), this would be reflected in increased productivity measures despite there being no change in technical efficiency. These and related measurement issues mean we are still limited in our knowledge of how turnover in asset ownership and management affects the level and growth of producers' efficiency levels.

In this paper, we seek to make progress on this front. A primary advantage of our effort is a data set that allows us to investigate the production process at an incredible level of detail. We observe the operations, 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 turn of the 20th Century). These data, which we describe in detail in the next section, contain records in physical units of inputs employed and output produced at each plant in the years it operated. We have matched these production data with business histories of the firms that operated in the industry to allow us to identify all major ownership and/or management turnover events (including but not limited to mergers and acquisitions) as well as the personalities involved. This combined data let us measure directly how ownership and management turnover was reflected in plants' physical productivity levels, and the extent to which these relationships differed across turnover events that appear to be triggered by lack of performance, those that followed unforeseen events (such as sudden death of the owner or CEO), and those driven by other motives.

While our data is historical in nature, it holds additional advantages besides its detail. Specifically, it spans a time of critical economic development and industrialization for Japan, which at the time was less than two decades removed from the completion of a difficult and often violent process of transition to modernity after 250 years of an isolated, traditionalist society. (Indeed, two-thirds of the firms in the industry were less than 5 years old at the time our panel begins.) 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 there are many important lessons that can be learned from our exercise in this study. By digging deep into micro-evidence, we aim at complementing past empirical work and providing fresh

insights for further development of economic theory about resource reallocation.

Our main findings can be summarized as follows. (To be completed.)

2. The data and research variables

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

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

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

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

observations from our analysis.³

Each plant in the records is associated with the firm that owned it in a given year, making it possible to directly compare the plant's physical (quantity) productivity before and after the change in ownership. This feature makes our data particularly attractive for analyzing plant productivity changes following ownership and/or management turnover.⁴ We also collected actual stories surrounding each acquisition and ownership turnover case, including but not limited to identities and backgrounds of the most important individuals involved (shareholders, top managers and engineers). Several data sources made this possible. First, 90 percent or more of the Japanese cotton spinning firms were public (joint stock) companies, obligated by law to issue shareholders' reports every half a year. Copies of these reports were also sent to Boren's headquarters in Osaka and those of them that have survived until the present day are currently hosted in the rare books section of Osaka University library. With the permission from the library we have photocopied the total of 1,292 reports on 149 firms, all what was available for the period from the early 1890s until 1920.⁵ Each report, in particular, contains a list of all shareholders and board members of the

³ To the best of our knowledge, we were the first to conduct this comprehensive cleaning of published plant-level records for the Japanese cotton spinning industry for 1896-1920. Our cleaned plant-level data tables and the details of the procedure outlined above are available upon request.

⁴ See Foster et al. (2008) and Syverson (2011) for the discussion of the importance of separating quantity and revenue productivity and the difficulties encountered by researchers trying to do it using conventional data that contain sales, values of inputs and prices but not direct evidence on the quantity of inputs and outputs.

⁵ While some of these company reports had been used in previous research by Japanese historians, we were the

company issuing it, making it possible to see whether shareholders or top management teams had already been substantially overlapping even prior to the formal acquisition event and what were the new positions (if any) of major shareholders and top managers of acquired firms in the new integrated firms. Company reports also contain detailed balance sheets and profit-loss statements as well as qualitative information about shareholders' meetings, deaths, illnesses, resignations and replacements of board members and so on, which we use as appropriate.

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

first to systematically digitalize them. The Osaka University library plans to launch a web site that will contain these digital copies and introduce them into the public domain in the near future.

Kinukawa also presents invaluable information about the background of most important shareholders and managers of each firm covered in his book as well as the storyline about how each firm was conceived.

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

While physical input and output data give us a unique chance to examine physical plant productivity as opposed to its revenue productivity, estimating the plant's physical productivity still presented several challenges. First, even though cotton yarn is a relatively homogeneous product it still comes in varying degree of fineness, called "count."⁶ Output of cotton yarn in our data is measured in units of weight, but there is also information about the average count produced by a given plant in a given year. To make different counts comparable for the purpose of productivity analysis, we converted various counts to the

⁶ The yarn count expresses the thickness of the yarn and its number indicates the length of yarn relative to the weight. The higher the count, the more yards are contained in the pound of yarn, so higher-count yarn is thinner (finer) than lower-count yarn. Producing higher-count (finer) yarn generally requires more skill and superior technology than producing lower-count (thicker) yarn. High-count yarn is often also improved further by more complex technological processes known as doubling, gassing, and so on, which were quite challenging for the fledgling Japanese cotton spinning mills to master at that time.

standard 20th count using a procedure described in detail in the Appendix. We also conducted all our estimations in an alternative way, using output in weight units and including the average count as a separate regressor when estimating the production function and confirmed that the results were similar.

Second, the worker count data include blue-collar workers (by gender—male, “danko” and female, “joko”) but do not include white-collar workers (“shyain”). Hence, in our total factor productivity estimates, the residual should be interpreted as reflecting the managerial input in a broad sense, including the input of all white-collar personnel. As the data give us the number of male and female blue-collar workers separately, we used the plant-year-specific ratios of female to male wages to convert one unit of female labor to one unit of male labor.⁷ Third, while we have direct measures of capital input in the data in the form of the number of spindles in operation, spinning frames are just one part of capital equipment which accounts for 25-30 percent of the total equipment cost of a mill (Saxonhouse, 1971, p. 55). Correlation between spindles and other equipment (cards, draw frames, slubbing frames, intermediate frames, roving frames, etc.) is, however, extremely high (over 95 percent), so “there is no question that spindles are a good proxy for equipment

⁷ In the division of labor between sexes in Japanese cotton spinning mills, opening, mixing, carding, repairing and boiler room work were generally (although not exclusively) men’s jobs, while tending, drawing, roving and operating ring frames were generally women’s work (Clark, *Cotton Goods in Japan*, pp. 191-194, cited in Saxonhouse, 1971, p. 56). Using female to male wage ratios to aggregate the labor input assumes that wages reflect the marginal productivity of each sex. All our estimates are completely robust to using the number of male and female workers separately in the production function estimations.

as a whole” (Saxonhouse, 1971, p. 56). We also have the data on the number of spindles installed in each plant in each year, which allows us to measure capacity utilization rates and follow any plant upgrades as the new equipment is installed.

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

2.1. An example of ownership turnover in our data

In August 1898, the shareholders of the decade-old struggling Onagigawa Menpu

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

Hibiya, who was promoted to company president in early 1899, wasted no time in introducing much needed change. All work-unrelated and hazardous activities on factory premises were immediately banned. Plant deputy manager tried to stir workers' unrest and was quickly fired, together with the head of the personnel department and the chief accountant (an off-duty police officer was temporarily stationed inside the plant as a show of new management's determination). But Hibiya did not stop at just introducing disciplinary measures. Even though he had another plant of his own to take care of, he and his right-hand man from Tokyo Gasu Boseki came to the Onagigawa factory and personally inspected equipment and checked output for defects on a daily basis, while also teaching workers how to do it on their own. During these visits, Hibiya reportedly engaged workers in conversations related to technology and production practices, taking questions, writing down

those that he couldn't answer immediately and coming back the next day with answers obtained from outside sources. Having determined that one reason for poor quality was that factory resources were spread too thinly, he concentrated production in just a few key areas, shutting down some workshops and switching from in-house production of finer counts of cotton yarn to procuring those from his other newer and more high-tech plant. Other measures included selling older equipment and purchasing more modern machines.

The above account reads remarkably similar to the description of the experiment in modern Indian textile industry conducted by Bloom et al. (2012). The results of Hibiya's restructuring effort were also equally or perhaps even more impressive. Using our data described in detail below, we estimate that the plant's TFP relative to the industry average more than doubled in the 3 years after Hibiya took over compared to 3 years before that while labor productivity (measured as output in physical units per worker-hours) increased on average by 70 percent. Over the same period, labor productivity in two other comparable plants in the same Tokyo area increased by just 6 percent. It is also worth noting that Hibiya was not part of an international aid effort; he was hired through an internal decision-making process of the shareholders, dishing out their own money.⁸

⁸ Hibiya's story is typical of industrialization pioneers in Japan and shows how much it was a land of opportunity at the time. Born Kichijiro Ohshima, third child of the owner of a hotel in a small provincial town, the future Heizaemon Hibiya was noticed by a cotton trader who stayed at the hotel when the boy was 13 and went to Tokyo to become the trader's apprentice. At the age of 20 he was doing trades on his own. He went on to grow one the most successful cotton trading houses in the Tokyo area, while also playing a major role in several

3. The development of the Japanese cotton spinning industry and acquisitions and ownership turnover: some basic facts

3.1. Industry growth and acquisition process

The development of the Japanese cotton spinning industry in the late 19th - early 20th century has long fascinated economists because of its unique nature “as the only significant Asian instance of successful assimilation of modern manufacturing techniques” before the Second World War (Saxonhouse, 1971). (For further discussion of the industry, see Saxonhouse, 1974 and 1977; Miwa and Ramseyer, 2000; Ohyama, Braguinsky, and Murphy, 2004; and Braguinsky and Rose, 2009.) 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 complete autarky. In line with the neoclassical comparative advantage theory, this resulted in most pre-industrial domestic production of manufactured goods being wiped out by foreign competition. Cotton yarn in particular experienced the combination of the largest fall in relative price from the autarky to the free trade regime and the highest negative net exports (Bernhofen and Brown, 2004). Starting in the late 1880s, however, the cotton spinning industry began a remarkable ascendance, which, within a historically very short period, led it to become the first globally competitive industry

prominent cotton spinning and other firms and eventually becoming vice-chairman of the Tokyo Chamber of Commerce.


to emerge from Asia.

As late as in 1887, domestic output was still a fraction of imports but it exploded in the following decade. Net export turned positive for the first time in late 1896, the year we start our analyses, and over the next two decades the country was exporting a sizeable fraction of its output while imports became negligible (see Figures 1 and 2). In 1887 there were 21 firms operating in the industry, with the average plant capacity (installed spindles) of 3,292 and the average number of factory floor workers of 137. Ten years later, in 1896 there were 63 firms with the average plant capacity of 12,789 spindles employing on average 719 workers. Thus, while the number of firms (plants) tripled, the average size of the plant almost quadrupled and the average number of workers per plant increased by more than 5 times. The overall industry output in physical units increased by more than 17 times over the same period (Nihon Choki Tokei Soran, Vol. 2, pp. 346).

Past research has identified a crucial role played in this process by superior entrepreneurial talent and rapid diffusion of new technologies (Ohyama, Braguinsky, and Murphy, 2004; Saxonhouse, 1974; Braguinsky and Rose, 2009). The success of initial few leading firms led to wide-spread emulation. By the second half of the 1890s, virtually all Japanese cotton spinning mills used the same capital equipment supplied by a single English manufacturer, Platt Bros. of Oldham (Saxonhouse, 1971). Best technological practices were also diffused through Geppo, the monthly journal published by the already-mentioned industry association, and directly through firm-to-firm contacts, where leading firms shared

technical knowledge and even proactively helped lagging firms and new entrants with adopting the best practices and training their workers.


Figure 1
Domestic output, import and export of cotton yarn (1888-1897)



Source: Nihon Choki Tokei Soran, our data.

Figure 2

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




Source: Nihon Choki Tokei Soran, our estimates.

Despite this, not all entrants were equally successful. In 1896, the average capacity of cotton spinning mills in the top quartile of the plant size distribution was 32,500 spindles, each employing on average 1,600 workers, quite respectable size by the international standards of the time (and even by today's standards in most developing countries). The average capacity of mills in the bottom quartile of the size distribution, on the other hand, was just 1,908 spindles and 126 workers. The gap between the average labor productivity and capital productivity (measured as output in physical units per factory worker-days and per spindle-days in operation, respectively) between these two categories of cotton spinning mills was 50 percent and 40 percent, respectively.

As the industry's extensive growth (temporarily) peaked out at the turn of the 20th

century, leading firms looking to increase their span of control increasingly turned attention to plants whose quality of management appeared to be subpar. In late 1901, Sanji Muto, plant manager and future chief executive of one the largest companies, Kanegafuchi Boseki (Kanegafuchi Cotton Spinning Co., aka Kanebo) published a book entitled “A Treatise on the Cotton Spinning Industry Grand Merger” (Boseki Dai Godo Ron, 1901) where he advocated merging independent cotton spinning firms into one grand coalition. While this idea has never been realized, the process of mergers and acquisitions became the most important source of growth for most firms in the industry for over a decade (Figure 3).

Figure 3
 Dynamics of capacity of older plants, acquired plants and newer plants



Source: our estimates using the data described in Section 2.

The category of plants labeled “Older never acquired” in Figure 3 represents plants that came into operation in 1902 or earlier. The category “Newer never acquired” represents plants that started operating in 1908 or later. No plants started operating between 1903 and 1907, and in fact only one small plant started operating between 1900 and 1907, so for eight years there was basically no new plant construction in the industry. The category “Acquired plants” is shown in Figure 3 in two versions: the solid line is the total capacity of those plants (regardless of whether they had been acquired or not yet), while the dashed line is the capacity of those that had already gone through at least one acquisition. The data show that capacity expansion of continuing firms went through an expansion phase exclusively through acquisitions and then re-commenced expanding their older plants and started constructing new plants. New entry also picked up once again after 1907, and new entrants were responsible for about the same amount of new capacity added to the industry as continuing firms.

Bloom et al. (2012) point out how difficult it is even in today’s India for superior firms to increase their span of control through acquisitions. The biggest difficulty appears to be the lack of trust outside immediate family members. In contrast, most Japanese cotton spinning mills were joint stock companies with easily transferable ownership, so it seems that investor protection and lack of business trust were never really an issue.⁹

⁹ How did a functioning market for assets emerge so early in the process of economic modernization is a subject for a separate study (see Miwa and Ramseyer, 2000, for some insights on this issue).

The fact that acquisitions assumed such a prominent role in firm growth process so early on also seems at first glance to be at odds with the established theoretical view, according to which investment by purchasing new capital should come before acquisitions (e.g., Jovanovic and Rousseau, 2002). The intuition behind the underlying theory (which finds support in the evidence gathered in the U.S.), however, is simply that new capital can be purchased without incurring a fixed cost, while acquisitions entail a fixed cost of consummating the deal. Since Japan in those years had to import capital equipment from England, taking over existing plants (which had the same equipment already shipped and installed) was potentially a cheaper (or at the very least a more easily accessible¹⁰) alternative to new investment for Japanese firms looking to expand.

The combination of factors mentioned in the previous two paragraphs led to a large number of mergers and acquisitions in the industry in the first two decades of the 20th century. Some 74 distinct acquisition deals were consummated between 1898 and 1920, during which 93 plants changed hands (often more than once). Fifteen more plants were consolidated under a single ownership in the megadeal that in 1914 created Toyo Boseki Spinning Company on the basis of equal merger of two of the largest firms at the time, Osaka Boseki Spinning Company and Mie Boseki Spinning Company. All in all, 50 out of 78 plants (64 percent), which were in operation in the industry in 1897, the year before the first acquisition

¹⁰ Using Platt Bros. records, Saxonhouse estimates that the time lag between receipt of spinning mill orders from Japan and shipments of equipment ranged from one to two years for most of the 1890s and the 1900s, increasing to 3-5 years after the start of World War I (Saxonhouse, 1971, p. 51).

took place, were subsequently acquired by a different firm at least once. In terms of plant capacity, the fraction acquired at least once was even larger, 76.2 percent. There were also 35 plants where ownership and/or management changed completely without going through a formal acquisition by a different firm.

Table 1
Number of acquired and turnover plants by year

Year	Number of acquired plants	Fraction of total	Number of turnover plants	Fraction of total
1896	0	0.000	0	0.000
1897	0	0.000	0	0.000
1898	1	0.012	2	0.024
1899	5	0.060	6	0.071
1900	5	0.061	5	0.061
1901	1	0.012	4	0.049
1902	2	0.025	3	0.037
1903	15	0.188	16	0.200
1904	2	0.025	2	0.025
1905	3	0.038	4	0.050
1906	5	0.062	5	0.062
1907	11	0.136	11	0.136
1908	2	0.025	4	0.049
1909	2	0.023	2	0.023
1910	1	0.012	1	0.012
1911	6	0.069	7	0.080
1912	5	0.057	9	0.102
1913	0	0.000	0	0.000
1914	0	0.000	15	0.153
1915	3	0.029	4	0.038
1916	5	0.048	5	0.048
1917	3	0.028	3	0.028
1918	12	0.109	13	0.118
1919	2	0.017	3	0.026
1920	2	0.017	2	0.017
Total	93	0.042	126	0.056

Note: Acquired plants are those acquired by another firm in the industry or an outside investor forming a new firm. Turnover plants include also plants that were part of changes in ownership without forming a new firm and one large equal merger (in 1914).

Table 1 presents year-by-year numbers of acquired plants as well as the number of plants that were not necessarily acquired by a different firm but were either part of an equal merger or otherwise underwent complete change of ownership or top management team (“turnover plants”). Each year over the 25-year period from 1896-1920, an average 4.2 percent of cotton spinning mills were acquired, while an average 5.6 percent experienced either an acquisition or another form of ownership turnover. The average acquisition rate is higher than the 3.89 percent acquisition rate for large U.S. manufacturing plants over 1974-1992 reported in Maksimovic and Phillips (2001) or 2.69 percent in the LED plants from 1972-1981 sample employed by Lichtenberg and Siegel (1987, Table 3). Indeed, almost all plants that were operating in the industry in 1896 and had not been shut down by 1920 were either acquired at least once or otherwise belonged to firms that participated as acquirers in the market for assets at least once over these 25 years.

3.2. Acquisitions and productivity – an overall view

Table 2 compares acquired plants to non-acquired continuing plants and exiting plants in the first and last three years of our data. Not acquired continuing plants are those that had never experienced ownership turnover and were still in operation as of 1920. Not acquired exiting plants are plants that were shut down (and their equipment either scrapped or relocated) between 1898 and 1920. These are mostly plants belong to firms that exited the industry, but also include a few plants that belonged to multiple-plant continuing firms.

Table 2
Comparing acquired, not acquired continuing and exiting plants

1896-1898	Acquired at least once	Not acquired continuing	Not acquired exiting
Capacity	14,447	24,867	3,310
Output per worker	1.376	1.359	0.991
Output per spindle	0.095	0.097	0.076
Capital/labor ratio	21.870	19.954	23.214
TFP relative to industry-year avg.	0.035	-0.015	-0.118
Average return on equity	0.118	0.190	0.128
Average dividend rate	0.080	0.114	0.081
Female daily wage	13.7	14.5	13.0
Main count produced	18.6	19.3	15.9
Average price	89.8	89.4	90.7
Number of observations	140	32	51
<hr/>			
1918-1920			
Capacity	23,559	58,342	
Output per worker	2.500	2.591	
Output per spindle	0.119	0.113	
Capital/labor ratio	27.553	29.896	
TFP relative to industry-year avg.	0.057	-0.051	
Female daily wage	86.2	96.1	
Main count produced	24.3	23.7	
Average price	434.3	388.9	
Number of observations	183	36	

Note: “Exiting plants” are those that had been shut down before 1920. Return on equity is annualized semi-annual net profit, divided by shareholders’ invested capital; dividend rate is annualized semi-annual dividends paid, divided by shareholders’ invested capital.

Comparing the columns in the upper panel of Table 2, we can see that in terms of capacity, acquired plants were on average much smaller than not acquired continuing plants but several times larger than exiting plants (both differences are statistically significant at 1 percent level using the double-sided *t*-test). Acquired plants were also much more

productive as measured by output per worker, output per spindle and relative TFP¹¹ than exiting plants (all the differences are statistically highly significant). But when comparing acquired plants with not acquired continuing plants, we can see that was not too much difference between these two categories in terms of productivity measures – while capital productivity and the average count of cotton yarn produced by the “to be acquired” plants were somewhat lower compared to continuing plants that were not to be acquired, their labor productivity was higher and so was their relative TFP (none of the differences is statistically significant at conventional levels, however). Thus, apart from their smaller size, plants that were subsequently acquired were not inferior to the plants of acquiring firms in terms of productivity. This is consistent with McGuckin and Nguyen (1995) and Rhodes-Kropf and Robinson (2008) and seems to be not consistent with the Q-theory where acquisitions are driven by productivity gaps.

Comparing returns on equity and dividends rates across the firms that owned these three categories of plants reveals, however, that continuing not acquired firms are on average already 61 percent more profitable than firms that would be acquisition targets and their dividend rate is on average 42 percent higher (both differences are statistically significant at the 1 percent level). There is no difference, on the other hand, between the profitability and dividends rate of future acquired and future exiting firms.

¹¹ Relative TFP in Table 2 is measured as residuals from the regression $\ln y = \alpha + \beta \ln k + \gamma \ln l$ estimated at the plant level, separately for each year (see Lichtenberg and Siegel, 1987; McGuckin and Nguen, 1995), averaged over 1896-98 and 1918-20.

It is also worth noting that both acquired and non-acquired continuing plants produced yarn of higher counts and paid higher wages to their female factory operators than did exiting plants. The capital-labor ratio (the ratio of the number of spindles to the number of workers) and the average price per unit of output, on the other hand, are higher in exiting plants than in any other category. Combined with exiting plants' small size and poor productivity, this probably is an indicator of bad management decisions. In particular, higher prices per unit of output in this case probably reflect not the higher quality of output (which is captured separately in our data by the average count of yarn produced) but high operational costs, something that made these plants not viable in the long run.

The data in the bottom panel of Table 2 show that both acquired and continuing non-acquired plants grew much larger by the end of the period. Not acquired continuing plants especially more than doubled their capacity, but the increase in acquired plants' capacity was also considerable. Both categories of plants also became much more productive in absolute terms (output per worker, output per spindle, main count produced), while acquired plants increased their advantage over non-acquired plants in relative TFP (the difference in relative TFP between these two categories is statistically significant at the 10 percent level in 1918-1920). Overall, however, the picture suggests that acquired and non-acquired surviving plants improved almost in parallel over these years.

We now take a more formal first look at how outcomes of acquired (turnover) plants changed with acquisitions. To this effect, we estimate the production function using output

and inputs in physical units across all years by pooled OLS, including three sets of dummies for acquired plants, one set equal to 1 for years before a given plant was acquired, the second set equal to 1 for years in which the plant was acquired and the third equal to 1 for years after the plant was acquired. For plants that were acquired more than once the period in-between acquisitions was split equally into “after” and “before” years. In addition to (logs of) spindles in operation (capital input) and composite labor input (the sum of male and female workdays, weighted by the relative plant-level ratio of female to male wage) we also include the (log) of plant age (to control for equipment depreciation)¹² and the difference in logs of plant capacity from the previous year (to capture possible adjustment costs of installing new equipment). The regression is

$$(1) \ln y_{it} = \alpha_0 + \alpha_1 \ln k_{it} + \alpha_2 \ln l_{it} + \alpha_3 \ln age_{it} + \alpha_4 d \ln cap_{it} + \beta_1 BA_{it} + \beta_2 A_{it} + \beta_3 AA_{it} + \mu_t + \varepsilon_{it},$$

where y_{it} is the output of plant i in year t (adjusted for count), k_{it} is the number of spindle-days in operation, l_{it} is the number of worker-days, age_{it} is the plant (vintage-adjusted) age, $d \ln cap_{it}$ is change in log plant capacity from the previous year, BA_{it} (“before acquisition dummy”) is equal to 1 if the plant was acquired in some year after t and zero otherwise, A_{it} is the dummy equal to 1 if the plant was acquired in year t and zero otherwise, AA_{it} (“after acquisition dummy”) is equal to 1 if the plant was acquired in some year before t and zero

¹² We have data on all equipment installations and/or removals in each plant, so plant age is adjusted for the vintage of capital. For example, if a plant was first built in 1890 but added an equal amount of new equipment in 1898, its vintage age is 1890 until 1897 while in 1898 and beyond it is calculated as the average of 1890 and 1898 (that is, it will be 1894). A similar procedure is applied to removals of older equipment.

otherwise, μ_t is the set of year dummies and ε_{it} is the error term. We also estimated regression (1) including plant fixed-effects (and omitting plant age) as well as including other instances of ownership or top management team turnover, not necessarily by acquisition. The estimation results are presented in Table 3.

Table 3.

Productivity changes after acquisitions and turnovers

Dep. variable	Log output		Dep. variable	Log output	
Log spindle days	0.772*** (0.043)	0.743*** (0.035)	Log spindle days	0.766*** (0.043)	0.744*** (0.034)
Log worker days	0.340*** (0.053)	0.241*** (0.031)	Log worker days	0.346*** (0.053)	0.241*** (0.031)
Log plant age	-0.037 (0.022)		Log plant age	-0.040* (0.023)	
Dlog plant capacity	-0.152*** (0.043)	-0.091** (0.040)	Dlog plant capacity	-0.147*** (0.042)	-0.089** (0.040)
Before acquired	0.065** (0.031)	-0.000 (0.026)	Before turnover	0.037 (0.032)	-0.024 (0.025)
After acquired	0.138*** (0.032)	0.071*** (0.025)	After turnover	0.140*** (0.033)	0.057** (0.024)
Constant	-3.084*** (0.258)	-1.439*** (0.421)	Constant	-3.053*** (0.253)	-1.451*** (0.419)
Plant fixed-effects	No	Yes	Plant fixed-effects	No	Yes
Year fixed-effects	Yes	Yes	Year fixed-effects	Yes	Yes
Observations	2063	2,064	Observations	2063	2,064
R-squared	0.941	0.965	R-squared	0.941	0.965

Note: robust standard errors clustered at plant level in parentheses. ***, ** and * indicate that the coefficients are statistically significant at the 1 percent, 5 percent, and 10 percent levels, respectively.

The sum of the coefficients on log capital and log labor input exceeds one, suggesting the presence of some degree of increasing returns in the production function. As expected, plant vintage age negatively affects productivity in the pooled OLS regression,

although the effect is statistically not significant. Plant expansions, on the other hand, are strongly and robustly negatively associated with productivity in the corresponding years. Our coefficients of interest are the coefficients on before and after acquisition (turnover) dummies. In the pooled OLS specification, acquired plants are significantly relatively more productive even before they are acquired (the omitted category is all non-acquired plants, whether continuing, exiting, or new), and they become relatively even more productive after being acquired, with the difference between the two coefficients indicating an improvement in relative productivity of about 7 percent, statistically significant at the 5 percent level using double-sided *t*-test with unequal variance (at the 1 percent level with one-sided *t*-test). The specification including plant fixed-effects estimates the improvement in productivity after acquisition of almost the same magnitude.

The last two specifications (which compare pre- and post-turnover productivity for all 126 turnover events in Table 1) indicate that the productivity of pre-turnover plants is similar to industry-year average but it improves by 14 percent compared to industry year-average after turnover. With plant fixed-effects included, the magnitude of improvement is very similar to that for acquisitions only.

We next turn to a more in-depth examination of different types of acquisitions and ownership/management turnover.

4. Acquisitions, management turnover and productivity: detailed analyses

4.1. Acquisitions and ownership concentration

When Kanegafuchi Boseki (Kanebo) was founded in 1887 in Tokyo it strived to become the largest cotton spinning firm in Japan right away. The strategy backfired at the time as the firm almost went bankrupt. It was rescued by Mitsui bank, the main bank of the surging Mitsui zaibatsu (financial, mining and trading conglomerate) and became one of the first ones to expand aggressively through acquisitions at the turn of the 20th century. Over the period from 1899-1911, Kanebo acquired eight different firms to add the total of 13 plants all over the country to its own original two plants.

At about the same time, another early entrant, Mie Boseki (founded in 1880 but completely reorganized in 1887), operating from the cities of Yokkaichi and Nagoya in central Japan was busy acquiring one firm after another in the areas where its main plants were based. By the early 1910s it had also acquired 8 different firms and found itself owning 11 plants (2 of its own, plus 9 acquired). In 1914, Mie Boseki and Osaka Boseki, another large firm with 4 plants accounting for 6.5 percent of industry-wide capacity at the time, completed an equal merger. The new integrated company, named Toyo Boseki instantaneously became the largest cotton spinning firm in Japan.

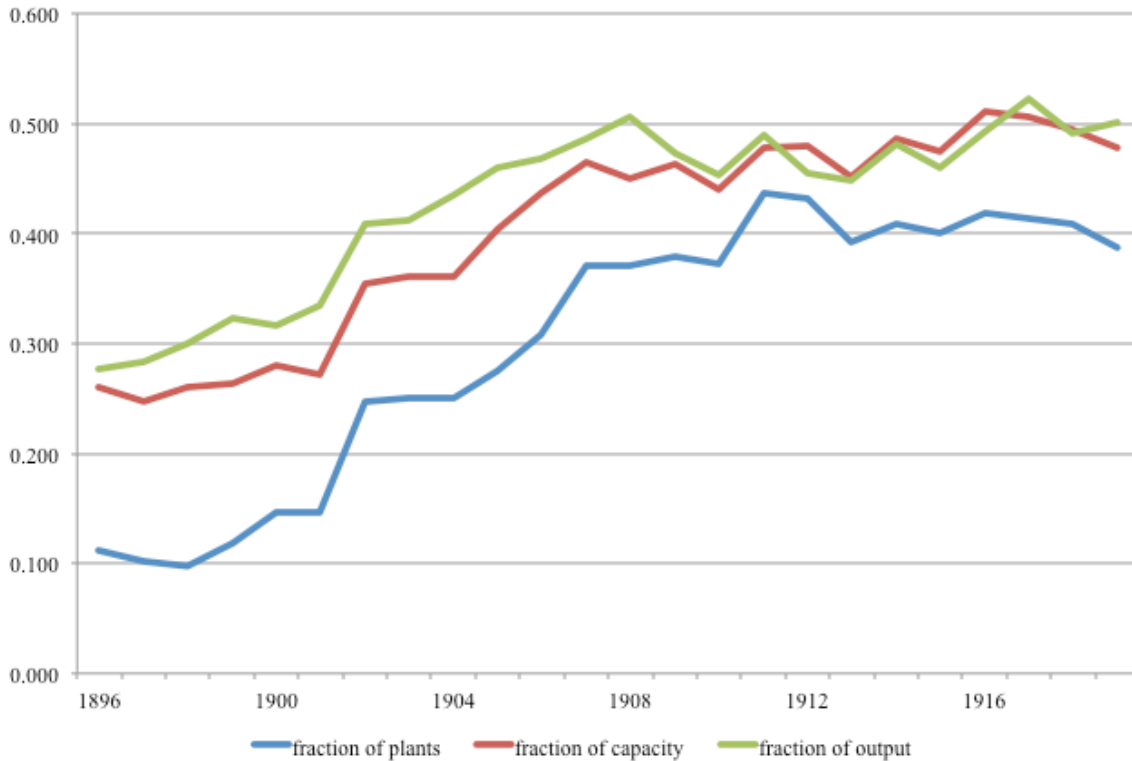
In contrast to Kanebo and Mie, the Osaka area-based Amagasaki Boseki joined the mergers and acquisition market relatively late. The company, founded in 1889, pioneered the production of high-quality 42nd count cotton yarn in Japan and for a long time was the only

competitive producer of this type of cotton yarn in the domestic market and for export. As its rivals were growing through acquisitions, Amagasaki Boseki finally decided it needed to expand too. In the end of 1914 it acquired Tokyo Boseki, an established firm in the capital city area that had fallen into disarray following the sudden death of its founder and CEO a few years earlier. Amagasaki Boseki then added another firm that had also lost its founding owner in 1916 while its CEO, Kyozo Kikuchi in the same year became the CEO of Settsu Boseki, another large acquirer with 7 plants and 7.3 percent of industry capacity. In 1918, Amagasaki Boseki formally acquired Settsu Boseki to establish Dainihon Boseki that had about the same size as Toyo Boseki and Kanebo.

Figure 5 traces the dynamics of the fraction of plants, their capacity and industry-wide output owned by the above 3 firms (Kanebo, Toyo Boseki and Dainihon Boseki, including their predecessor firms) over the 25-year period of our analyses. We can see that at the beginning of the period, the fraction of plants jointly owned by these 3 firms was just about 10 percent of all plants but the same fraction exceeded 40 percent towards the end of the period. Their combined share of industry-wide capacity and output was already close to 30 percent at the beginning of the period but it increased to over 50 percent towards the end. No other firm at the end of our period owned more than 10 plants or had more than 10 percent of industry-wide capacity under its ownership, so these 3 firms stood out by their sheer size and importance.

Figure 5

Ownership concentration in three largest firms



Note: the figure depicts the evolution of the fraction of plants owned by three largest firms in 1920 (Kanegafuchi Boseki, Toyo Boseki, Dainihon Boseki) and the fraction of these plants' capacity and output in the industry total. Toyo Boseki data include the data of its predecessor firms (Osaka Boseki and Mie Boseki) prior to their 1914 merger. Dainihon Boseki includes the data of its predecessor firms (Amagasaki Boseki and Settsu Boseki) prior to their 1918 merger.

The concentration of ownership in the hands of these 3 largest firms could in principle be due to multiple factors. The most interesting question for the purpose of this study is how much it was related to physical productivity of their plants. To gain some preliminary insights into this question we first estimated the production function regression

(not shown) similar to (1) but including the dummy equal to 1 if the plant was owned by one of the 3 largest firms in a given year and 0 otherwise. The estimated coefficient on the dummy reflecting ownership by one of the 3 largest firms was 0.092, statistically significant at the 1 percent level. Moreover, when we split the plants owned by these firms into their original plants and the plants they acquired later, the coefficient on the dummy equal to 1 if the plant was an original plant and zero otherwise is indistinguishable from zero, while the coefficient on the dummy equal to 1 if it was an acquired plant and zero otherwise is estimated to be 0.11. Thus, it appears all the relative productivity advantage of the 3 largest firms during this period came from their acquired plants.

These preliminary findings raise an interesting question – did the superior productivity of the plants acquired by the 3 largest firms (and hence the superior productivity of those firms overall) come from actual improvement in the productivity of the plants they acquired or were the plants acquired by these largest firms more productive than the industry average to begin with? In the former case we can conclude that ownership concentration contributed to industry-wide improvement in productivity. But if the latter is actually the case, this could mean that those firms were simply good at selecting and successfully targeting plants that were better regardless of their ownership.

To begin answering this question, in Table 4 we presents some summary statistics comparing plants acquired by the 3 largest firms above and all other firms at the beginning and end of our period of analysis in the format analogous to Table 2. At the beginning of the

period (in 1896-98) the only statistically significant difference between plants that would be acquired by the 3 largest firms and all other firms is the plants' capacity – the plants that were later targeted by the 3 largest firms have 63 percent larger number of spindles installed. The firms that will be acquired by the largest firms are also somewhat more profitable than firms to be acquired by other firms, although the difference is only marginally statistically significant. There is no significant difference in any other metric – for example plants to be acquired by the 3 largest firm have slightly higher labor productivity than other subsequently acquired plants and they produce higher counts of cotton yarn, but their capital productivity and relative TFP are actually lower.

The picture changes dramatically after acquisitions. While plants acquired by other firms somewhat caught up in size with plants acquired by the 3 largest firms, the plants whose ownership was taken over by the 3 largest firms are now much more productive than other acquired plants on all dimensions, with all the differences economically and statistically highly significant. The relative TFP of plants that had been acquired by these 3 largest firms was 13 percent higher than the industry average in 1918-1920, while it was statistically and economically indistinguishable from the industry average in 1896-98. This comparison shows that acquisitions by the 3 largest firms did indeed actually improve the productivity of the plants they acquired as compared to other acquired plants as well as to industry average. The fact that those plants were also larger than industry average (and other acquired plants) underlines the importance of this improvement even more from the

industry-wide productivity perspective.

Table 4.

Comparing plants acquired by large and other firms before and after acquisitions.

	Acquired by large firms	Acquired by other firms	Difference
1896-1898			
Capacity	18,339	11,263	62.8%***
Output per worker	1.384	1.370	1.0%
Output per spindle	0.094	0.096	-1.7%
Capital/labor ratio	22.297	21.519	3.6%
TFP rel. to industry-year average	0.016	0.051	-3.5%
Average profit rate	0.150	0.088	6.2%*
Average dividend rate	0.085	0.075	1.0%
Female daily wage	13.9	13.6	2.8%
Main count produced	19.6	17.9	9.8%
Average price	90.7	89.2	1.7%
Number of observations	63	77	
1918-1920			
Capacity	27,570	20,373	35.3%***
Output per worker	2.880	2.186	31.7%***
Output per spindle	0.127	0.113	12.4%***
Capital/labor ratio	29.914	25.515	17.2%***
TFP rel. to industry-year average	0.128	-0.003	13.1%***
Average profit rate	0.955	0.893	15.6%* (1-sided only)
Average dividend rate	0.596	0.550	10.6%**
Female daily wage	84.1	88.0	-4.4%
Main count produced	26.0	22.9	13.9%* (1-sided only)
Average price	474.1	401.4	18.1%**
Number of observations	81	102	

Note: relative TFP is measured as residuals from the regression

$\ln y = \alpha + \beta \ln k + \gamma \ln l$ estimated at the plant level, separately for each year (see Lichtenberg and Siegel, 1992; McGuckin and Nguen, 1995). The residuals are summed over all to-be-acquired plants (to-be-acquired by large firms plants) in “Before acquired” and “Before large acquired” categories, respectively, for three years 1898-1900. The residuals are summed over all acquired plants (acquired by large firms plants) in “After acquired” and “After large acquired” categories, respectively, for

three years 1898-1900.

***, ** and * indicate that the corresponding differences are statistically significant at the 1 percent, 5 percent, and 10 percent levels, respectively using double-sided *t*-test with unequal variance.

In Table 5 we present the results of estimating the same regression (1) as in Table 3 but with the before and after acquisition dummies equal to 1 only if the corresponding plants were acquired by the 3 largest firms (or its predecessor firms), zero for all other plants. The regressions exclude plants acquired by other firms to keep the omitted category the same as in Table 3.

Table 5.

Productivity changes after acquisitions by the 3 largest firms

Dep. variable	Log output	
		0.727*** (0.054)
Log spindle days	0.400*** (0.069)	0.257*** (0.040)
Log worker days	-0.028 (0.025)	0.003 (0.023)
Log plant age	-0.168*** (0.063)	-0.110* (0.064)
Dlog plant capacity	0.076* (0.041)	-0.004 (0.037)
Before acquired	0.176*** (0.045)	0.077** (0.030)
After acquired	-3.211*** (0.319)	-1.948*** (0.577)
Constant		
Plant fixed-effects	No	Yes
Year fixed-effect	Yes	Yes
Observations	1256	1257
R-squared	0.943	0.969

Note: The 3 largest firms are Toyo Boseki, Dainihon Boseki and Kanegafuchi Boseki and their predecessor firms as explained in the main text. Robust standard errors clustered at plant level in parentheses. ***, ** and * indicate that the coefficients are statistically significant at the 1 percent, 5 percent, and 10 percent levels, respectively.

Comparing the estimation results in Table 5 with those in Table 3 confirms that the after-acquisition improvement in productivity was much bigger for plants acquired by the largest firms than by all firms taken together, at least in the pooled OLS specification (including plant fixed-effects makes the estimated coefficient on the after acquisition dummy rather similar to the one estimated in Table 3).

We also looked at within-plant changes in productivity around acquisition events by the largest 3 firms. The specification presented in Table 6 has the following form:

$$(2) \ln y_{it} = \alpha_0 + \alpha_1 \ln k_{it} + \alpha_2 \ln l_{it} + \alpha_3 \ln age_{it} + \alpha_4 \ln cap_{it} + \sum_{\substack{k=-3 \\ k \neq 0}}^{k=5} \gamma_k AcqYear_{ik} + \mu_t + \pi_i + \varepsilon_{it}.$$

$AcqYear_{ik}$ is a dummy variable indicating the relative year around the acquisition event (year zero is the year of acquisition and the omitted category). The regression is estimated using OLS and includes plant fixed effects (π_i) and year fixed effects (μ_t), with standard errors are clustered at the plant level. Because plant and year fixed effects are included, the coefficients on year dummies capture within-plant productivity in a given year with acquisition year as the omitted category, while controlling for differences in average productivity across different plants as well as for average productivity in a given calendar year.

Table 6.

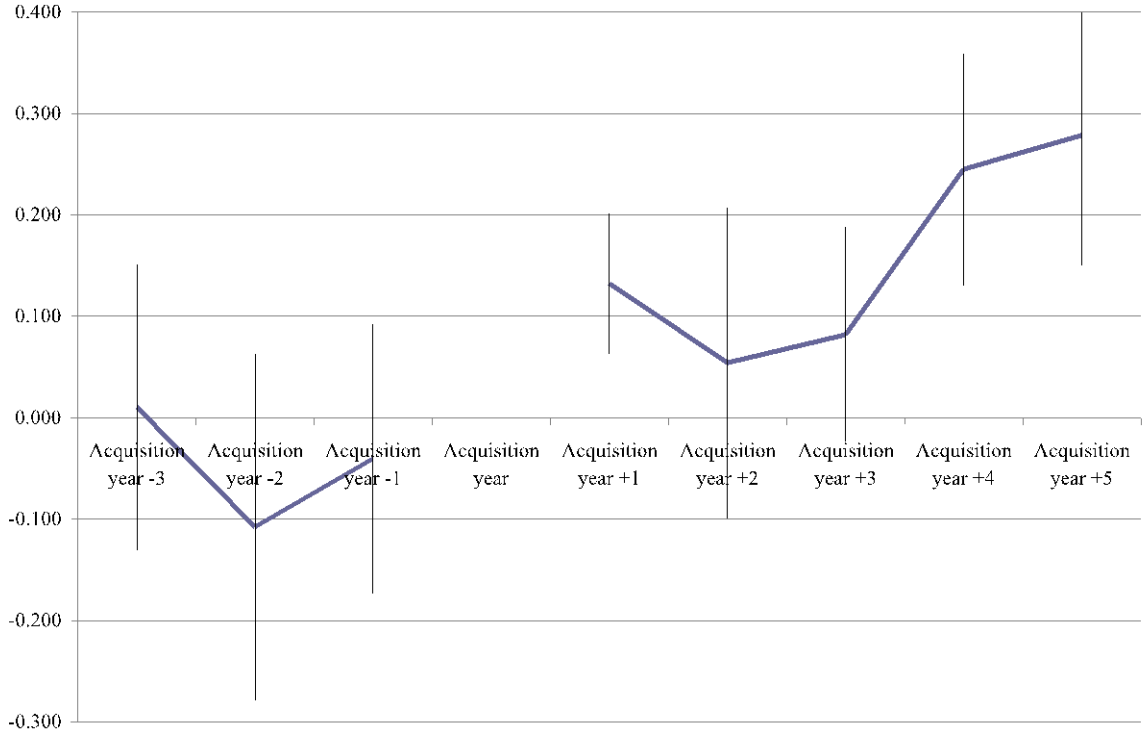
Within-plant relationship between acquisitions by the 3 largest firms and productivity

	Log output
Acquisition year -3	0.011 (0.069)
Acquisition year -2	-0.108 (0.084)
Acquisition year -1	-0.040 (0.065)
Acquisition year	[Omitted]
Acquisition year +1	0.132*** (0.034)
Acquisition year +2	0.054 (0.075)
Acquisition year +3	0.082 (0.052)
Acquisition year +4	0.245*** (0.056)
Acquisition year +5	0.279*** (0.063)
Log spindle days	0.753*** (0.084)
Log worker days	0.112** (0.051)
Dlog plant capacity	0.002 (0.130)
Plant fixed effects	Yes
Year fixed effects	Yes
Observations	245 (on 30 plants)
R-squared	0.946

Note: The table presents within-plant changes in productivity around the acquisition by one of the largest 3 firms (or one of their predecessor firms as explained in the main text). “Acquisition year” are dummy variables indicating the relative year around the acquisition event (the omitted category is the year of acquisition). Estimations are done by OLS with standard errors clustered at the plant level. ***, ** and * indicate that the coefficients are statistically significant at the 1 percent, 5 percent, and 10 percent levels, respectively.

We also estimated regression (2) with different time windows and corresponding sets of plants and the results were very similar as long as at least 4 years after acquisition events were included. Figure 6 presents a graph of the coefficients on year dummies, together with 95 percent confidence intervals.

Figure 6.
TFP around acquisitions by the 3 largest firms.



The figure presents the changes in within-plant TFP in the years around acquisitions by the 3 largest firms (or one of their predecessor firms as explained in the main text) and the 95 percent confidence intervals. The chart estimates and confidence intervals are taken from the acquisition year dummy variables in the first column of Table 6.

The estimation results show that in the years preceding acquisition by one of the largest 3 firms, plant productivity was overall lower than in the omitted year (the acquisition year itself), although due to high standard errors the estimated coefficients are statistically not significant at conventional levels. After the acquisition, within-plant productivity jumps about 10 percent in the first three post-acquisition years and then increases further to 25-30 percent in subsequent years. This pattern is robust to adding more post-acquisition years to the analysis but it appears that the effect of acquisition levels off after the fifth year at somewhat less than 30 percent improvement in productivity.

Plants that were acquired by the 3 largest firms were responsible for 20.1 percent of industry-wide output at the end of our period, 1918-1920. Our estimates of 30 percent improvement in productivity in these plants as a result of being acquired by those firms imply that without such acquisitions the industry-wide output in 1918-1920 would have been reduced by about 6.42 percent.

4.2. Investment underperformance, disruptions, and effects of acquisitions

Hakata Kinuwata, founded in 1896 by a local capitalist on Japan's southernmost island of Kyushu was among relatively new entrants, in pursuit of its share of what appeared to be the ever-growing pie offered by the booming cotton spinning industry at the time. But its business did not go well from the outset and it soon found itself losing money and having to borrow heavily from banks. In February 1900 a revolt by shareholders led to the

resignation of the founder and company president Seizo Ohta but things did not turn around. In December 1902 the company sold itself to Kanebo, which was actively expanding its presence in the Kyushu area, with Hakata Kinuwata shareholders taking a loss of about 50 percent on their invested capital (Kanegafuchi Boseki Kokajo, No. 32, 1902; Okamoto, 1993, pp. 307-308). Our plant-level data estimates show that following this acquisition, the labor productivity of the plant increased by 35 percent, while the TFP which had been about the same as industry average prior to acquisition exceeded it by 25-30 percent starting from the second post-acquisition year.

Table 7.
Labor productivity and relative TFP dynamics at Nihon Boseki plant
around the death of the founder-CEO and acquisition

Year	Output/worker	Relative TFP
1914	5.834	0.069
1915	5.435	0.008
1916	3.240	-0.212
1917	2.853	-0.261
1918	3.499	-0.018
1919	4.004	0.039

Source: our estimates.

The acquisition of Nihon Boseki by Amagasaki Boseki in late 1915 presents a very different picture. The long-time CEO of Nihon Boseki was taken ill and had to resign in May 1915. His successor was also frail and started searching for an acquirer. In November 1915 the deal was struck with Amagasaki Boseki whose CEO Kyozo Kikuchi had been a board member of Nihon Boseki and long-term advisor to the company. Despite this, and despite the fact that Amagasaki Boseki was one of the most prominent firms in the industry (see above),

the transition apparently did not go very smoothly. Table 8 presents the dynamics of labor productivity and relative TFP of the main Nihon Boseki plant for two years immediately preceding the acquisition (consummated in 1916) and 3 years after.

The impact of the disruption caused by unanticipated shock to the top management can be clearly seen in the big drop in productivity in 1916 and 1917. It took the new ownership two years to straighten the ship.

The above two examples are typical. They illustrate the fact that acquisitions happening for different reasons can have very different effects on plants' productivity even where the acquiring firms are technologically and managerially sophisticated top-notch firms such as Kanebo or Amagasaki. In this section we focus on the two most interesting acquisition types from this perspective, those triggered by investment underperformance (as in the Hakata Kinuwata case) and those triggered by unforeseen disruptions (as in the Nihon Boseki case).

Using primarily the data from financial statements contained in shareholders' reports we classified 35 acquisition cases as triggered by investment underperformance.¹³

¹³ In order to be classified as underperforming, the acquired firm had to be posting return on equity below industry average in at least two-thirds of half-years in three years prior to acquisition for which profitability data were available and it had to be posting return on equity at or below the 25th percentile of the corresponding industry-year distribution in at least 1/3 of those half-years. We experimented with other reasonable criteria and confirmed that the results were similar. We also classified as underperforming financially a small number of cases where shareholders' reports were not available because acquired firms were privately held and therefore not obligated to issue those reports. In those cases we relied on the accounts of circumstances surrounding

We classified, using the information from the same shareholders' reports, 17 acquisitions as triggered by a disruption if they followed unforeseen events, such as illness, death of the firm's owner and/or CEO or his abrupt resignation due to personal reasons unrelated to the firm's performance (hereafter "disruption acquisitions"). Altogether, 52 out of 93 acquisitions fall into either "underperforming" or "disruption" category. We will examine the remaining acquisitions separately in the next subsection.

Table 8
Probability of acquisitions after disruptions and loss years

	Probability of acquisition in year t	Probability of turnover in year t
Disruption event in years $t-3$ to $t-1$	0.103*** (0.030)	0.151*** (0.031)
At least 2 loss years in years $t-3$ to $t-1$	0.111*** (0.056)	0.124*** (0.056)
Year dummies	Yes	Yes
Number of observations	2241	2241
R-squared	0.071	0.082

Probit regression showing marginal effects, with standard errors clustered at the plant level. OLS and logit regressions produce very similar results. *** indicates that the coefficients are statistically significant at the 1 percent level.

The importance of focusing on investment underperformance and/or disruptions can be seen in that these events are strong predictors of acquisitions and ownership turnover. In Table 8 we present the results of a simple regression where the probability of acquisition (turnover) happening in year t is regressed on at least one "disruption" event happening in

acquisitions in Kinukawa (1964). Details are available upon request.

years from $t-3$ to $t-1$, at least 2 loss years in years from $t-3$ to $t-1$, and on year dummies. Given that the unconditional average probability of acquisition in our time period is 0.042 and the average probability of turnover is 0.056, we can see that going through a disruption event in one of the 3 previous years or experiencing at least 2 financial loss years more than doubles these probabilities in any given year.¹⁴

Table 9.

Productivity changes after “underperforming” and “disruption” acquisitions

	Underperforming acquisition		Disruption acquisition	
	Log output		Log output	
Log spindle days	0.766*** (0.043)	0.744*** (0.034)	0.761*** (0.044)	0.745*** (0.035)
Log worker days	0.353*** (0.054)	0.242*** (0.030)	0.355*** (0.056)	0.244*** (0.032)
Disruption event			-0.022 (0.039)	-0.067* (0.035)
Before acquisition	-0.001 (0.033)	-0.032 (0.033)	0.046 (0.047)	0.013 (0.049)
After acquisition	0.145*** (0.028)	0.124*** (0.035)	0.047 (0.079)	-0.007 (0.056)
Constant	-3.125*** (0.260)	-1.467*** (0.405)	-3.094*** (0.261)	-1.509*** (0.419)
Plant fixed effects	No	Yes	No	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	2063	2064	2063	2064
R-squared	0.941	0.966	0.939	0.965

OLS with standard errors clustered at the plant level. ***, ** and * indicate that the coefficients are statistically significant at the 1 percent, 5 percent, and 10 percent levels, respectively. Regressions include also log plant age (only in specification without plant

¹⁴ We also estimated the acquisition probability using the dummy equal to 1 if the firm was below the 25th percentile of the profit rate (return on equity) distribution in at least 2 of the 3 previous years instead of posting losses and the results were similar.

fixed-effects) and log capacity in year t – log capacity in year $t - 1$.

Table 9 presents the results of estimating regression (1) with dummies set equal to 1 for years before and after “underperforming” acquisitions (columns 1 and 2), and with dummies set equal to 1 for years before, at and after “disruption” acquisitions (columns 3 and 4). As can be seen from the table, relative TFP of plants acquired after they underperformed financially compared to the industry as a whole is not different from industry average prior to acquisition but is 12-15 percent higher after they are acquired. For plants acquired due to disruptions, on the other hand, there no difference between their productivity before and after acquisitions.

We next conduct the same exercise as in the previous section for within-plants changes in productivity among “underperforming acquired” plants and compare these with all other acquisitions (including but not limited to “disruption acquisitions”). Table 10 presents the estimation results (we have also tried longer time windows as well as two-year averages of pre- and post-acquisition year dummies and the coefficients were essentially unchanged). It appears that controlling for plant-level and year fixed effects, acquisition triggered by investment underperformance increase the productivity of acquired plants by 10-15 percent on average, starting with the second after-acquisition year (we confirmed that this effect persists beyond year 5 after acquisition). The results are also very similar to those presented in Table 9 where we include all plants, not just acquired ones.

In contrast, there is nothing going on around other acquisitions, so whatever the

reasons why firms were acquired in those instances, they did not result in improved plant-level physical productivity or perhaps were not intended for this purpose to begin with.

Table 10.

Within-plant relationship between acquisitions and productivity

	Underperforming acquisitions	Other acquisitions
Acquisition year -3	0.088* (0.050)	-0.018 (0.048)
Acquisition year -2	0.004 (0.059)	-0.010 (0.047)
Acquisition year -1	0.031 (0.056)	0.001 (0.050)
Acquisition year +1	0.051 (0.046)	-0.045 (0.041)
Acquisition year +2	0.106* (0.057)	-0.032 (0.042)
Acquisition year +3	0.102 (0.062)	-0.032 (0.057)
Acquisition year +4	0.162*** (0.057)	0.038 (0.095)
Acquisition year +5	0.142** (0.061)	0.033 (0.081)
Log spindle days	0.708*** (0.050)	0.695*** (0.089)
Log worker days	0.225*** (0.052)	0.155* (0.083)
Plant fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Observations	295	309
R-squared	0.949	0.959

Note: The table presents within-plant changes in productivity around the “underperforming” and all other acquisition, as explained in the main text. “Acquisition year” are dummy variables indicating the relative year around the acquisition event (the omitted category is the

year of acquisition). Estimations are done by OLS with standard errors clustered at the plant level. ***, ** and * indicate that the coefficients are statistically significant at the 1 percent, 5 percent, and 10 percent levels, respectively.

It appears from examining Table 10 that most of the effect of acquisitions triggered by underperformance is the level, with just a slight upward trend after acquisition event. To test this more formally, we have estimated a “structural break” regression where we used only acquired plants and regressed (the log of) output on (the log of) capital input, (the log of) labor, difference between (the log of) capacity between the previous and current years, as well as on pre- and post-acquisition dummies and pre- and post-acquisition dummies interacted with time trend, while omitting the constant term. The difference between the coefficients on pre- and post-acquisition dummies (essentially, the two constants in before- and after- regressions) measures the level change from before- to after-acquisition, while the difference between the coefficients on the interaction terms between pre- and post-acquisition dummies with time trend will measure the change in trend.

The results, separately for “underperforming” acquisitions and all other acquisitions are presented below in Tables 11a and 11b, each in three specifications. In the first specification (top panel in each table) we use all available observations, while in the second panel we limit the time window to years where we have at least 5 observations on different plants (from year Acquisition-10 to year Acquisition+18 for “underperforming” acquisitions and from year Acquisition-10 to year Acquisition+17 for other acquisitions). Finally, in the

bottom panel we present the results from the same regression as in the second panel but without interacting capital and labor with pre- and post-acquisition dummies (essentially assuming the same underlying production function).

Table 11a.

	Underperforming acquisitions		
All observations	Before	After	Difference
Trend	0.005	0.009	0.003
	0.007	0.003	
Constant	-3.043	-2.839	0.204
	0.470	0.430	
Number of Observations	196	385	
<hr/>			
Between year-10 and +18 (at least 5 observations per year)			
Trend	-0.004	0.010	0.014
	0.009	0.003	
Constant	-2.981	-2.823	0.158
	0.489	0.467	
Number of Observations	185	375	
<hr/>			
Between year-10 and +18 (at least 5 observations per year)			
Without interacting other controls			
Trend	-0.007	0.010	0.017*
	0.008	0.003	
Constant	-3.038	-2.852	0.187***
	0.383	0.379	
Number of Observations	185	375	

Standard errors adjusted for 33 clusters in plant id

The estimation results in Table 11a suggest that the effect of underperforming

acquisitions is generally positive on both productivity level and trend (at least when we limit the time window to years where we have at least 5 observations per year). The standard errors are high compared to the magnitude of the differences in estimated coefficients, so none of the differences in the top two panels is statistically significant at conventional levels. The estimated magnitudes in the second and third panels, however, are very similar and both post- acquisition level and trend are statistically different from their pre-acquisition counterparts, although the difference in trend is barely significant.

Table 11b
Other acquisitions

All observations	Before	After	Difference
Trend	0.013	0.003	-0.009*
	0.004	0.003	
Constant	-2.629	-0.929	1.700***
	0.266	0.408	
Number of Observations	284	283	
<hr/>			
Between year-10 and +17 (at least 5 observations per year)			
Trend	0.017	0.002	-0.015*
	0.006	0.006	
Constant	-2.606	-0.845	1.761*
	0.689	0.716	
Number of Observations	265	278	
<hr/>			
Between year-10 and +17 (at least 5 observations per year)			
Without interacting other controls			
Trend	0.018	0.001	-0.017*
	0.006	0.006	
Constant	-2.024	-2.048	-0.024
	0.518	0.528	
Number of Observations	265	278	

Standard errors adjusted for 36 clusters in plant id


The estimations in Table 11b consistently show that the productivity trend, which is positive and statistically significant prior to acquisition, totally disappears after acquisition. The difference between small positive upward trend before acquisition and the (lack of) post-acquisition trend is also statistically significant in all three specifications (at the 10 percent level). Thus, those other acquisitions may indeed be driven by selection on the part of acquiring firms. The big difference between the pre-acquisition and post-acquisition productivity levels is harder to interpret, especially that this difference seems to be sensitive to the specification.

Finally, in Figure 7 we have plotted the residuals from demeaned and detrended plant-specific production function regressions before and after acquisitions triggered by investment underperformance, disruption events and all other events in three distinct categories. In order to do so, we first estimate a regression similar to (1), including also plant fixed-effects, time trend and all interaction terms between plant dummies and time trend, that is,

$$(3) \quad \ln y_{it} = \alpha_0 + \mathbf{aX} + \beta_1 trend_t + \beta_2 \pi_i + \beta_3 \pi_i \times trend_t + \eta_{it},$$

where \mathbf{X} is the vector of our usual controls and π_i s are plant dummies. The residual term in this regression, η_{it} , thus measures by how much the total factor productivity of a given plant in a given year deviated from its demeaned overall trend.

Figure 7.
Residuals from regression (3)
“underperforming acquired”, “disruption acquired,” and other acquired plants



As expected, disruption events (which by construction occur during years -3 to -1 in Figure 7) pull down plant productivity compared to the plant-specific time trend. We can see that following the acquisition event, the decline is halted and reversed beyond year 3 (compare to the Nihon Boseki case in Table 7 above) but at no point do such acquisitions improve the plant productivity compared to its long-term trend. This is, of course, what

should be expected in cases like this – since the acquisition was triggered by an event that disrupted an otherwise performing plant, the new ownership on average is no better match for the orphaned plant than was its deceased or otherwise retired previous top manager. Thus, it should not be surprising that acquisitions triggered by unforeseen natural causes do not have any significant impact on the productivity of plants involved.

Other, non-disrupted and not underperforming plants can be seen to be somewhat below their long-term trend in years 3 and 2 before being acquired but they are already at their long-term trend at the acquisition event and remain there until year 4 after acquisition. Their productivity then finally starts increasing again (and keeps the somewhat elevated level seen in year Acquisition +5 for several more years).

The picture is very different for plants acquired from financially underperforming firms. As can be seen from Figure 7, such firms tend to underperform not just financially but also in terms of physical plant productivity compared to the long-term trend of the plants they owned by about 5 percent on average in years Acquisition -3 to year Acquisition +1. In other words, relatively low profitability is indicative of the plant's ownership (top management team) inability to operate it up to its full potential. Once taken over, such plants tend to outperform their long-term trend by about 5-10 percent during the first 3-5 years after acquisition and continue to outperform their long-term trend for a few more years beyond that. The total improvement in productivity attributable to acquisition is therefore estimated to be about 10-15 percent.

There were 73 plants acquired in the “underperforming” category and their total output in 1918-1920 amounted to 17.3 percent of total industry output. In the counterfactual world where they had not been acquired, we estimate that their productivity would have stayed 15 percent lower, implying a loss of industry output of about 3 percent. This loss could, of course have been much larger if those plants had to be shut down altogether, so the 3 percent number is the lower bound (the upper bound is the 17.3 percent of industry-wide output they accounted for at the end of our time frame). Exactly half of these plants had been acquired by the 3 largest firms which we examined in the previous section and the total output produced by plants acquired in the “underperforming” category by the largest firms and other firms was also split about equally. Thus, the contribution of “underperforming”-driven acquisitions by firms other than the 3 largest ones to industry-wide productivity growth was about half of the total.

5. Discussion

To be completed.

References

- Bernhofen, Daniel M., and John C. Brown, 2004. "A Direct Test of the Theory of Comparative Advantage: The Case of Japan," *Journal of Political Economy*, 112 (1), 48-67.
- Bertrand, Marianne, and Antoinette Schoar, 2003. "Managing with Style: The Effects of Managers on Corporate Policy," *Quarterly Journal of Economics*, 118 (4), 1169-1208.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie and John Roberts, 2012. "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, forthcoming.
- Bloom, Nicholas, and John Van Reenen, 2010. "Why Do Management Practices Differ Across Firms and Countries?" *Journal of Economic Perspectives*, 24 (1), 203-224.
- Braguinsky, Serguey, and David C. Rose, 2009. "Competition, Cooperation, and the Neighboring Farmer Effect," *Journal of Economic Behavior and Organization*, 72 (1), 361-376.
- David, Joel, 2012. "The Aggregate Implications of Mergers and Acquisitions," Working Paper, University of South California.
- Foster, Lucia, John Haltiwanger, and Chad Syverson, 2008. "Reallocation, Firm Turnover and Efficiency: Selection of Productivity or Profitability?" *American Economic Review*, 98 (1), 394-425.

- Fuken Tokei (Prefectural Statistics), 1912-1920, various prefectures (in Japanese).
- Geppo (Monthly Bulletin of the Japanese Cotton Spinners' Association), 1889-1920 (in Japanese), Dainihon Boseki Rengokai, various issues.
- Hsieh, Chiang-Tai and Pete Klenow, 2009. "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124 (4), 1403-1448.
- Jovanovic, Boyan, and Peter Rousseau, 2002.
- Kinukawa, Taiichi, 1964. *Hompo Menshi Boseki Shi* (History of the Japanese Cotton Spinning Industry) in 7 volumes (in Japanese), Tokyo, Hara Shobo, 1991 (reprinted from the original 1942 edition).
- Kokajo (Company reports), 1896-1920 (in Japanese), Osaka University Library.
- Lichtenberg, Frank R., and Donald Siegel, 1987. "Productivity and Changes in Ownership of Manufacturing Firms," *Brookings Papers of Economic Activity*, 1987 (3), 643-683.
- Maksimovic, Vojislav, and Gordon Phillips, 2001. "The Market for Corporate Assets: Who Engages in Mergers and Asset Sales and Are There Efficiency Gains?" *Journal of Finance*, 56 (6), 2019-2065.
- McGuckin, Robert H., and Sang V. Nguen, 1995. "On Productivity and Plant Ownership Change: New Evidence from the Longitudinal Research Database," *Rand Journal of Economics*, 26 (2), 257-276.
- Miwa, Yoshiro, and Mark J. Ramseyer, 2000. "Corporate Governance in Transitional Economies: Lessons From the Pre-War Japanese Cotton Textile Industry." *Journal of*

Legal Studies, 29 (1), 171–203.

Nihon Choki Tokei Soran (Historical Statistics of Japan), 1987. In 4 volumes (in Japanese), Japan Statistical Association.

Nosho Komu Tokei (Statistics on Agriculture, Commerce and Industry), Noshokomusho (Ministry of Agriculture, Commerce and Industry) (in Japanese), 1896-1911.

Ohyama, Atsushi, Serguey Braguinsky, and Kevin M. Murphy, 2004. “Entrepreneurial Ability and Market Selection in an Infant Industry: Evidence From the Japanese Cotton Spinning Industry,” *Review of Economic Dynamics*, 7 (2), 354-381.

Okamoto, Yukio, 1993. *Chiho Boseki Kigyo-no Seiritsu-to Tenkai* (Setup and Development of Local Cotton Spinning Firms, in Japanese). Kyushu University 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.

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

Saxonhouse, Gary R., 1977. “Productivity Change and Labor Absorption in Japanese Cotton Spinning, 1891-1935,” *Quarterly Journal of Economics*, 91 (2), 195-220.

Saxonhouse, Gary R., 1974. “A Tale of Japanese Technological Diffusion in the Meiji Period,” *Journal of Economic History*, 34, 149-165.

Saxonhouse, Gary R., 1971. “Productivity Change in the Japanese Cotton Spinning Industry, 1891-1935,” Unpublished Ph.D. Dissertation, Yale University.

Schoar, Antoinette, 2002. "Effect of Corporate Diversification on Productivity," *Journal of Finance*, 57 (6), 2379-2403.

Siegel, Donald S., and Kenneth L. Simons, 2010. "Assessing the Effects of Mergers and Acquisitions on Firm Performance, Plant Productivity, and Workers: New Evidence From Matched Employer-Employee Data," *Strategic Management Journal*, 31 (8), 903-916.

Syverson, Chad, 2011. "What Determines Productivity at the Micro Level?" *Journal of Economic Literature*,